EUROPEAN FUSION DEVELOPMENT AGREEMENT
JET

EFDA–JET–CP(02)05/22

J.A. How, J.W. Farthing and V. Schmidt
and JET–EFDA contributors

Trends in Computing Systems for Large Fusion Experiments

Trends in Computing Systems for Large Fusion Experiments

J.A. How[1], J.W. Farthing[2] and V. Schmidt[3]
and JET–EFDA contributors*

[1]*Association Euratom-CEA,* CEA/DSM/DRFC-Cadarache,F-13108 *Saint Paul Lez Durance, France*
[2]Association Euratom-UKAEA, Culham Science Centre, Abingdon, OX14 3EA, UK
[3]Associazione Euratom-ENEA, Consorzio RFX, Corso Stati Uniti 4, I-35127 Padova, Italy
* *See annex of J. Pamela et al, 'Overview of Recent JET Results and Future Perspectives', Fusion Energy 2000 (Proc. 18th Int. Conf. Sorrento, 2000), IAEA, Vienna (2001).*

Preprint of Paper to be submitted for publication in
Proceedings of SOFT 2002 (Helsinki, Finland, 9–13 September 2002)

INTRODUCTION

This paper gathers together some of the lessons learned within the computer control and management systems, from the European fusion programme, and in particular from the JET project. A substantial design document for the ITER Supervisory Control System was released in 1998 giving the basic specification and the planning timetable for the "CODAC" Supervisory Control System [1] (although very little work has taken place in this domain since that time). In consequence we concentrate in this present paper on the many strategic and practical experiences that could complement this design.

The CODAC Specification document [1] gives the following timetable, relative to the start of construction, for the entire CODAC system:

- After 36 months: Start final software design.
- After 48 months: Start final hardware design.
- After 63 months: Complete final software design & fabrication.
- After 73 months: Complete final hardware design and fabrication.

We largely agree with such a timetable (adapted of course to the new reduced ITER design) as a compromise between using state-of-the-art techniques and avoidance of delays to the project. However there are many strategic and philosophical principles that should be adopted before the practical implementation is launched (for example hardware and software aspects related to the licensing and safety case). We have highlighted a number of these in this paper.

1. TECHNOLOGY TRENDS

At the opening of the 2001 IAEA computer TCM in Padua [2], F. Bombi recalled the "early days" of JET. Its Control and Data Acquisition System (CODAS) was conceived in the mid-70s [3,4]. At that time the ARPANET [5] was in its infancy and there were no TCP/IP protocols and no PCs. In some ways this made the design easier rather than more difficult. The only standard cited in the original design documents was the use of CAMAC as a computer-independent standard interface to the plant. Today the world is quite different, driven by the PC and home-user commodity market and by the growth of the Internet. The equivalent of software that was developed in-house in the early 1980s is now available off-the-shelf; for example, office computing, databases, data analysis tools and web technology. We assume that technology will continue to develop, at least, at the same rate. Just as in 1978, it is difficult to predict today what technologies will still be around in 25 years time. The requirement is to select the most enduring technologies and standards and only plan to have developed what is fusion-specific.

1.1 HARDWARE

The enduring electronics technology at JET has been CAMAC and EUROCARD. However maintenance of these old systems does become an issue. Standardising on a small number of designs and then deploying them repeatedly across the project - a "step-and-repeat" policy - can greatly reduce these problems. Care must be taken when selecting off-the-shelf hardware solutions - will they still be available in a few years time? For the inevitable custom-made products, experience has shown that ownership of the designs, allowing further production runs in the future, is a good policy. Although initially more expensive, this pays in the long run. Care must also be taken with the selection of components: if possible choose components that are available from more than one source.

The general trend is to move away from bespoke hardware to more flexible software-based systems. The original highly complex analogue plasma control systems are now technically straightforward in software. In the early 90s the long-term availability of some standard CAMAC modules was already uncertain. A move to FPGA [6] technology allowed the development of multi-purpose programmable modules to replace many of the older fixed designs.

Cabling is expensive and is unlikely to be replaced, but the interfaces to it may change. The same JET digital I/O EUROCARD modules, which were originally driven from CAMAC, can now be accessed via VME and could be accessed via PCI, if required. The CAMAC Serial Highway Driver was originally accessed from Norsk Data computers, was ported to SUN SBUS technology and is about to be ported to SUN PCI technology.

There is always a pressure to collect more data (the JET raw data collected per pulse has doubled in size every two years since 1983), but network and storage system capacities always manage to keep up.

The most obvious trend in computer and network hardware is that prices have continued to drop. It always pays to delay purchasing. It should also be noted that in recent years the PC commodity market has greatly influenced the adoption of Linux as the platform of choice for running physics codes.

1.2 SOFTWARE

The successful and long-lasting programs are those that were data-driven from the beginning and so could adapt to change. Hard-coded programs generally have a much shorter lifetime. Software maintenance is as big an issue as hardware maintenance. At JET there is a vast amount of legacy FORTRAN code from the 80s. Apart from the complexity of maintaining this code, it is also difficult to keep the interest of high-calibre staff in this role. The analogy of the step-and-repeat hardware policy is to apply reusable code via plug-in modules. To make this possible, software engineering standards and quality assurance must be implemented and enforced all across the project from the beginning. It should also be noted that the experience of employing outside companies to implement turnkey systems is mixed. Is the same company going to be employed to maintain the software over the lifetime of the project? Will the company exist for the lifetime of the project? If the outside company does not maintain the software then who does? A clear succession policy should be introduced for the ownership of software. Who is responsible for it and who will take it over if that person leaves?

Software licensing issues should also be considered from the beginning. With a likely dispersed development team at a mixture of universities, national laboratories and private industry, central software purchasing should be considered, preferably with international agreements for beneficial "research" pricing status. A decision also needs to be made about the use of Open Source Software [7]. At JET this is very much the trend and has proved very cost-effective.

2. CENTRALISED DATABASE

A major project requires a centralised database to hold all project data and documentation. Data is valuable and its value increases when related to other data. It should be captured only once and used throughout the lifetime of the project. Essential is a well-designed data model and hence a set of naming and numbering schemes from the outset. Database technologies may change with time, but the information that is stored will remain the same. Many products are in use in the European fusion labs, none of which fulfil all requirements.

The two approaches have been to either buy commercial products, or to develop in-house systems. The major problem encountered with the commercial approach is that these products have never been broad enough to integrate all the needs of Fusion (an experience common also in other research institutions). For example, good systems are available for the drawing-office management (including materials databases, version control, etc) but they are closed systems that do not readily interface with other areas (e.g. site administration tools), resulting either in incomplete data, or in duplication of databases. Despite manufacturer claims, it is not easy to integrate diverse software packages. Similarly there are good documentation systems such as the ESA Official Documents Management Service (EDMS) [8], but they are limited in scope. The CERN "Engineering & Equipment Data Management System" (also "EDMS") is perhaps the most comprehensive of the systems in use today in the research domain [9].

Within Fusion there is notably the Electra database and associated access tools at JET [10], and the EDITEX documentation and configuration management system at the DRFC Cadarache [11]. Both of these have been very successful but have not extended wide enough, and have not received sufficient levels of mandate or support; neither is a truly project-wide database. These aspects were also considered in the NET project and a substantial proposal was made, but not followed up further [12].

We consider this subject so crucial to the success of a major project that it must be given very high priority and support as early in the project as possible.

3. SUBSYSTEM INTERFACE SPECIFICATION

At JET, like other tokamaks, a three level (supervisory (1), subsystem (2) and component (3)) hierarchical and modular structure for control was defined in order to simplify the design, and to allow the use of common solutions throughout the system [3,4]. It should be noted that this was the basis of the highly successful step-and-repeat policy. The subsystems collected related functions together. They were designed to be autonomous to allow independent operation, commissioning and maintenance. They were made up of a number of components each performing a specific function. The supervision of the whole experiment was to be performed at a higher level, which co-ordinated the operation of all subsystems. The basic components of the supervisory level - central timing, interlocks, data storage etc - were implemented early in the project, but the supervisory experimental control system (level-1) [13] was a later and gradual addition. This system reached a level of sophistication, which far outstripped the individual subsystems' capabilities, and more and more there was a tendency to run local tests from the top level. In hindsight it would have been desirable to equip the local subsystems with far more of the sophisticated control options and possibilities.

Hence we suggest devolving more of the ideas of the JET level-1 system down to the subsystem level, to allow more complete off-line testing (principle of subsidiarity!). The ITER "Test and Commissioning State" (TCS [1]) makes provision for each subsystem to be autonomous for such testing. Such off-line testing should be as comprehensive as possible, including pulse simulation, pulse recall and as complete countdown checks as possible when off-line. It should be noted that as long as both the supervisory and subsystem control system databases are both derived from the same central database then all the attributes defined for one will automatically be available for the other.

It is apparent that, over time, the pulse-related level-1 database has grown to include a high proportion of the 80,000 JET hardware signals. These two databases could be merged, and the "pulse snapshot" principle could be applied to producing regular snapshots of the entire plant. A consequence of this would be that all the plant synoptic screens (mimics) could be interrogated to compare the real-time status with historic data; a major improvement to control room facilities. We note that the JET level-1 system [10] is a good example of where in-house software development is required. It is unlikely that such a system will ever be available off-the-shelf.

4. STEADY-STATE REQUIREMENTS

Most of the fusion research devices constructed around the world have been pulsed machines, with data collection after the pulse. However the distinction between "continuous data" and""pulse data" will be much less apparent in ITER than in present machines. Three of the worlds long-pulse fusion machines, Tore Supra, LHD, and W7-X, are working towards marrying the pulse and continuous systems [14, 15, 16]. This work has been further enhanced by recent real-time physics calculations reducing post-pulse data treatment to a minimum.

It is important for the next generation of tokamaks to fully accept the message that the traditional distinction between continuous machine subsystems and pulse-recorded diagnostic subsystems is no longer valid.

Some of the areas most affected by the change to steady-state are:

a) Continuous archive and availability of *all* signals (real-time data collection streaming, and viewing).
b) Module restart and parameter change *during* a pulse.
c) As much on-line real-time processing as possible.
d) Data compression before storage where possible (e.g. recording only signal changes)
e) Data reduction on retrieval of large signals – people tend initially to want an overview of whole signals, but long pulse plotting can exceed local memory limits leading to excessive swap-times
f) Absolute time stamps applied to all data and alarms (for example, satellite time references in all cubicles). Data analysis and retrieval can be referenced to absolute, pulse or event times.

g) Common Data Platform: many fusion machines worldwide have adopted the common MDSplus data platform. However this will require substantial modification to adapt it to long-pulse or continuous data [17].

5. INTEGRATED CONTROL AND FEEDBACK

Throughout the world, the fusion research machines are reaching a new level of integrated control and feedback, where pre-programmed scenarios are used only to reach the initial conditions of a feedback-controlled experiment. Most machines have feedback circuits designed around specific experiments. JET took a far more flexible approach, wiring some hundreds of signals into a real-time signal server (RTSS) linked to a powerful and flexible feedback central controller (RTCC) that serves control signals to local units (fuelling, position, current, and heating systems) [18]. This system is coupled to the level-1 supervisor, allowing replay of all pulses, off-line set-up and simulation, and rapid, safe parameter loading. Nevertheless the JET feedback system, like all others, was an add-on, not contemplated in the original design, which traditionally separated diagnostic systems from control systems. Next-generation machines should have integrated systems. Subsystems must be designed with the view that all signals might be needed for feedback, and therefore will need built-in confidence factors, validation flags, watch-dogs, abort scenarios, and facilities for off-line testing. An array of central signal servers will access the diagnostic information, producing both "physics signals" and control signals from a variety of inputs, and pass these to control systems and real-time displays.

6. SAFETY LICENSING

Following the example of JET, ITER intends that the majority of the "Supervisory Control System" will not be part of the regulatory scrutiny to ensure safety [ref 1, section 1.3]. Rather, the protective actions will be taken by the supervisory and subsystem level "Interlock Systems" to maintain the machine in a safe state, and the "Access Control System" will handle the safety aspects of personnel control. We include here some specific points from the JET experience.

6.1 The Plasma Operation will require envelope backstops within the safety interlock system for many of the control loops. For example, if the interlock system has a hard-wired beam-power envelope according to well-defined "Operating Instructions", then the physicist-controlled feedback can be allowed substantial operator flexibility without safety implications.

6.2 The diagnostic and machine signals necessary for safety licensing must be identified, and "hardened" early in the design. An example of such a signal at JET is the Bremsstrahlung detector [19], used as a primary safety interlock to prevent beams firing into an empty or under-dense vessel. This is part of the Fast Beam Interlock System [20], formally designated as a JET Integrated Operation Protection System (IOPS) [21]. As such, the system design was subject to review and approval through a formally constituted JET Machine Protection Working Group. Such IOPS systems form part of the "JET Safety Case" and are subject to periodic re-commissioning according to approved commissioning procedures resulting in a formal declaration of readiness for operation. All such cases for ITER require early identification and rugged, safe design.

6.3 It is of course essential that all hardware interlock protection levels are made known to the subsystem and supervisory control systems. One of the tasks of the supervisory control system is to make sure that the control circuits do not approach the hardware trip limits during a pulse, and so such hardware limits must have a read-back into the control software. For example, as with JET, the ITER design makes provision for two hard-wired protection levels on each coil circuit [22]. JET went through a painful period of retrofitting read-back units on the hardware limit potentiometers.

7. REMOTE PARTICIPATION

The ITER documentation [1] states: "ITER will provide remote experiment capability to maximise its experiment efficiency, enable participation of other off-site control rooms, and reduce the number of personnel located at the ITER site."

7.1 SECURITY WITH TRANSPARENCY

JET has demonstrated that it can operate a major nuclear experiment securely, and yet provide full visual transparency both from local offices and off-site. EFDA-JET created a strategic "Remote Participation Project" which has since been widened to encompass all of European Fusion [23]. We strongly recommend that the present JET philosophy should be applied to future collaborative fusion projects.

The JET control systems are fully and absolutely protected as required by licensing authorities, but nevertheless it has still been possible to build in transparency of all information, and full remote "participation". The main factor to achieve this was the strictly-controlled connection between the control systems and the office computer network. The extension to the outside world off-site is then a question of politics and protocol rather than a safety issue. Staff in on-site offices and off-site have access to all real-time data, and the pulse-definition "level-1" tools. For example, from off-site, one can capture the template for the last pulse, modify any parameters, and store the new pulse specification in a location available for verification and loading by the on-site Engineer-in-Charge.

On top of this internal local security structure, three levels of external (off-site) access are available.

a) The public web site, freely available to all.

b) A password-protected "Users" web site and data access, containing a wide spectrum of information, including Task Force information, visitor information, selected operations information and some data access.

c) Highly secure office and analysis computer access whereby a remote user has the same tools and access rights as a local office-based user.

7.2 REMOTE OPERATION

There is a strong move worldwide towards off-site contracts for subsystem maintenance and off-site diagnostic support and control. Networks are now sufficiently good to be able to ignore the distance aspects, and have support teams located anywhere. Manual intervention, when necessary, can be by non-specialist staff.

In European Fusion, one of the diagnostics on the TJ-II Heliac in Madrid is now fully operated from Barcelona [24]. Also recently the Accelerator community has produced a recommendations paper for remote operation of future large accelerators [25].

While there is nothing intrinsically wrong with such an approach, and indeed it has some major advantages for an international machine such as ITER, very careful analyses would need to be carried out on the safety, licensing and down-time implications, as well as on network security and redundancy requirements. Communications failure during remote maintenance could lead to substantial machine downtime. It should be noted that JET does not share this general view of remote control, and follows very successfully a strong policy of "Remote Participation" rather than""Remote Control".

7.3 REMOTE PARTICIPATION ARCHITECTURE

ITER intends to have extensive remote participation [1]. We note that there are two distinct architectures that could be used:

a) Remote Participation to individuals, controlled by the on-site staff. This is the situation presently at JET. Approximately 450 people, mostly within Europe have highly-secure "remote-office" access to JET (100 of which are JET staff who use the service for home or travel access), and a further 350 have access to the lower-security "Users" web (a significant proportion of the professional fusion staff in Europe!). Such an infrastructure requires substantial technical and administrative support.

b) An alternative to consider is to create one "mirror" site in each partner organisation, whereby *all information* is mirrored from the central machine to the mirror sites, and the responsibility for individual remote participation is devolved to the local partners. Such mirror sites could also house remote experiment participation rooms[1] and cater for 24-hour ITER operation. Although this would be expensive to the partners (buildings, equipment and support) it could be both politically desirable (reduction of central staff, close partner participation, duplication of archives, etc) and efficient in terms of utilisation of central machine resources.

This matter requires early consideration to create the necessary infrastructure as soon as the site team is established, in order to provide off-site support already during the construction phase.

7.4 NETWORK ASPECTS

The network infrastructure of a project requires a well-defined strategy, and can have long lead times and significant financial impact. Extrapolating from existing machines, ITER will need of the order of 10 Gbits/s network connection and 1 Terabyte/day transfer. These figures should be

[1] *The ITER documentation uses the term "Remote Control Room" for this.*

easily attainable by the time ITER is operational, and are modest compared with other communities' needs (in particular next-generation accelerators, and astronomy). Nevertheless the external network architecture will be both a budget and a logistics problem.

Present lack of network reliability, beyond our control, leads us to recommend redundancy with independently routed network connections. Bandwidth management techniques should be used both to guarantee the throughput of priority traffic (e.g. real-time work) and to assure automatic rerouting to provide high availability. However, it is likely in the future that guaranteed bandwidth contracts will be major budget items.

It is probable that a large part of the work during the ITER construction phase is likely to require remote facilities. The network connectivity defined above, and many of the remote participation tools, will already be required early in the construction phase.

8. NON-FUSION TRENDS

We will address only a few key issues here:

a) There is a worldwide trend in physics research laboratories to move away from in-house developed software and hardware and to buy industrially available products, in particular Supervisory Control and Data Acquisition (SCADA) systems, allowing effort to be concentrated towards the overall architecture and control schemes. However, the experience at many laboratories, including JET and CERN, is that such "turn-key" solutions seldom integrate fully with the overall systems, with correct transparency, security and alarm interface, and that they can create a legacy of applications that are difficult and expensive to maintain and support. Extreme care must therefore be exercised in the selection and support contracts for such systems.

b) One notable case is the EPICS Supervisory and Control system [26] that is widely used in the high-energy physics community and beyond. It is a cooperative effort similar to Fusion's MDSplus effort'[17]. It is being used, together with MDSplus, on the Princeton NSTX project [27, 28]. Both these products might be adapted for use by ITER.

c) Remote Participation is receiving substantial development everywhere. The worldwide "Grid" projects [29] allow for coupling of distant databases and computer resources, while the complementary "Access Grid""[30] programme tries to give Remote Participants the look-and-feel of being actually at the local site. ITER could become a partner in these projects. Some organisations, both within and outside Fusion, now use distributed file structures, whereby users have transparent access to files and computers anywhere in the world. For example the Andrews File System as used by IPP and ENEA [31].

d) Many enterprises have maintenance contracts for their office computer networks. However over the last ten years there has been a convergence between office and machine operation computers. The result is that laboratories often find that their on-site machine support staff are considerably more expert than the contractual maintenance team, and in-house computer assembly and maintenance becomes more attractive both on the grounds of cost and exploitation of existing expertise.

CONCLUSIONS.

This paper has tried to highlight some of the many lessons learned, related to the computer control and management systems, in European Fusion, primarily from the JET project. We have had to be both selective and concise with the subjects and we are also aware of some areas that have either been missed or not given sufficient emphasis. Fortunately there are now many ex-JET staff worldwide who will help to carry this experience over to the next machine.

Finally we note that computer support in Fusion has been as much a development area as any of the other subsystem subjects (heating, cryogenics, materials, robotics etc) and there is every reason to think that this will continue into the next generation of machines. ITER, like all previous fusion machines, will be a prototype. Its systems are all unique and purpose-designed, and the level of interaction with other subsystems is much higher than in most other research areas. It will require a strong technical central team to support on-site computer architecture and to verify the conformity of all new contracts. One should note that the duties and responsibilities of the JET CODAS team are well defined and can be readily extrapolated to ITER.

ACKNOWLEDGEMENTS

This work is largely about our experience gained over many years in the JET programme, and we would like to thank the whole JET-EFDA team, both on and off site. However we want also to acknowledge the input from the rest of the fusion community, both from our direct interactions, particularly at Cadarache and Padua, but also from the indirect input from conferences and collaborations. In particular it should be noted that we held a forum discussion at Padua [2] where worldwide representatives from fusion computer teams discussed many of the concerns and issues raised in this paper. Finally we would also like to thank S. Maruyama and B. Green of the ITER team for their helpful comments and suggestions.

REFERENCES

[1]. ITER Supervisory Control System, IDoMS:S 45 DDD 18 98-06-10 R 0.1.

[2]. 3rd IAEA TCM on Control, Data Acquisition, and Remote Participation for Fusion Research, Padua, Italy (2001). http://tcm2001.igi.pd.cnr.it

[3]. EUR 5791e, 1977, The JET Project, Scientific and Technical Developments, 1976

[4]. H. van der Beken et al., 1987, "CODAS: The JET Control and Data Acquisition System", Fusion Technology 11(1), pp. 127.

[5]. M. Hauben, History of ARPANET (Advanced Research Projects Agency Network), http://www.dei.isep.ipp.pt/docs/arpa.html.

[6]. Field Programmable Gate Array technology, - see for example M. Barr, A Reconfigurable Computing Primer, Multimedia Systems Design, September 1998 pp. 44-47 reprinted in: http://www.netrino.com/Articles/RCPrimer/

[7]. http://www.opensource.org/

[8]. ESA Official Documents Management Service http://edms.esrin.esa.it/
[9]. CERN, Engineering & Equipment Data Management System, http://edms.cern.ch
[10]. JET Electra DataBase, see for example Codas Internal Report JDN/H(99)18, J.Farthing(ed), 1999.
[11]. B. de Gentile, Internal Report, EDT/NTT-2001.003, DRFC, CEA Cadarache
[12]. Final report for the NET Contract NET/89-214
[13]. P.A. McCullen, J.W. Farthing, The JET level-1 Software, Proc. 20th SOFT Marseille 1998, 575-578
[14]. B. Guillerminet et al., Tore Supra continuous acquisition system, Fusion Eng. Des. Vol 60 (2002) Issue 3, 427-434.
[15]. H. Nakanishi et al., Design for real-time data acquisition based on streaming technology, Proc. 21st SOFT, Madrid 2002, 1011-1016
[16]. H. Lacqua et al., Control System of the Wendelstein 7X Experiment, this conference.
[17]. T.W. Fredian and J. A. Stillerman, MDSplus. Current development and future directions, Fusion. Eng. Des. 60 (2002) 229-233.
[18]. T. Budd et al., 1998, Design and Implementation of a Real-time Power Control System at JET, Proc. 20th SOFT, Marseilles, 1998, 579-582
[19]. S.J. Cox et al, An active phase compatible protection system to prevent excessive neutral beam shine through during JET plasmas, Proc.17th SOFT, San Diego CA (1997)
[20]. D.J. Cooper et al., The Fast Beam Interlock System for JET Neutral Injection, Proc. 12th SOFE, Monterey CA, Vol 2, 1096-1100 (1987)
[21]. D. Stork et al., Systems for the safe operation of the JET tokamak with tritium, FED 47 (1999) 2-3 131-172
[22]. Mondino et al, Integrated Protection for the ITER Coil and Power Supplies, SOFT Marseille 1998, 811-814
[23]. V. Schmidt, J. How, Remote Participation Technical Infrastructure for the JET facilities under EFDA, Proc. 21st SOFT,1039-1044, Madrid 2002
[24]. F. Castejón et al, Results of the Remote Participation on TJ-II NeutralParticle Analyser from UPC, Madison, 14th APS Topical Conference on High Temperature Plasma Diagnostics, July 8-11, 2002 Madison, to be published.
[25]. P. Czarapata et al, The Inter-Laboratory Task Force on Remote Operation, SLAC-R-583, Nov 2001, Stanford Linear Accelerator Center, http://arxiv.org/pdf/physics/0201014
[26]. EPICS Experimental Physics and Industrial Control System http://www.aps.anl.gov/epics/, Argonne national Laboratory
[27]. P. Sichta, et al , “The NSTX Central Instrumentation & Control System ,” 18th SOFE, Albuquerque, NM (1999).
[28]. W. Davis, et al, ‘The use of MDSplus on NSTX at PPPL,’ 3rd IAEA TCM on Control, Data Acquisition, and Remote Participation for Fusion Research, Padua, Italy (2001).

[29]. Grid Computing. See for example : http://www.gridcomputing.com, http://www.egrid.org, http://www.gridforum.org
[30]. AccessGrid http://www.accessgrid.com
[31]. Andrews File System, IBM Open Source Project, http://www.openafs.org