ITER Aims and Technical Challenges - II

Jo Lister

CRPP-EPFL

Centre de Recherches en Physique des Plasmas

Ecole Polytechnique Fédérale de Lausanne

Switzerland

Operating and exploiting ITER

- If we get all that right, we have constructed a very big, very complex and very expensive piece of equipment, ready and capable of fulfilling its physics goals
- What then…..
- We need to measure what is happening
- We need to control the full plant
- We need to control the plasma operation
- We need to handle all the data generated

Measuring plasma parameters - TCV examples

- Diagnostics and analysis approaches for resolving the spatial distribution of plasma properties across the variable TCV cross-section were developed over several years

140 channel
AXUV Bolometer

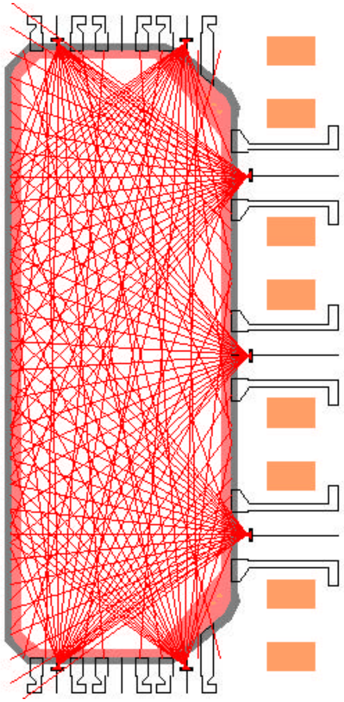

200 channel
Soft-X

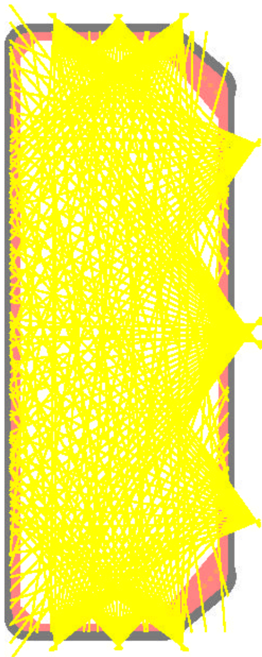

54 channel
Bolometry

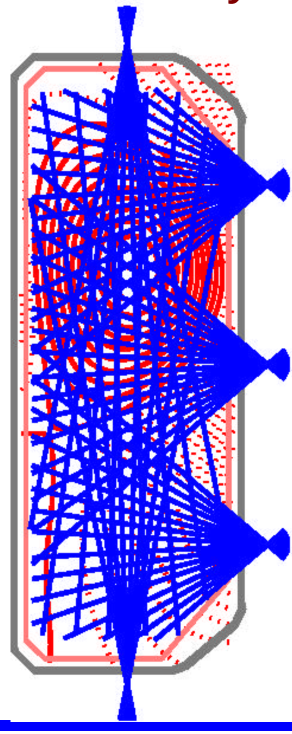

14 channel FIR
25 channel Th.Scat.

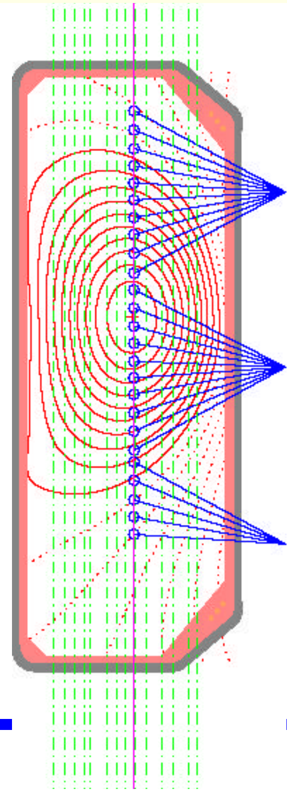

Measuring plasma parameters - ITER challenges

- Where are the ITER challenges ? - just a few examples
  - Neutron damage to semiconductors
  - Neutron damage to windows
  - Radiation induced and thermoelectric emf in mineral insulated cabling
  - Thermo-mechanical distortion
  - Deposition and erosion of mirror surfaces
  - Blackening and annealing of fibre-optics

Measure the plasma parameters - OK, but with difficulties

Optimally using all the diagnostic information, still a challenge

Controlling the full plant

- "Slow controls" will generate significant data, but ITER will have slower time-scales than present tokamaks
- Standard SCADA systems are well-developed
- We can imagine 500'000 to 1'000'000 channels of information
- This is a data handling issue rather than technical issue (return to this)
- Generating the working control software will require inventiveness, like for LHC

Control the plant - OK

Develop a suitably powerful finite state machine representation - needs work

Making an ITER pulse scenario

Develop the tools to generate ITER scenarios - OK

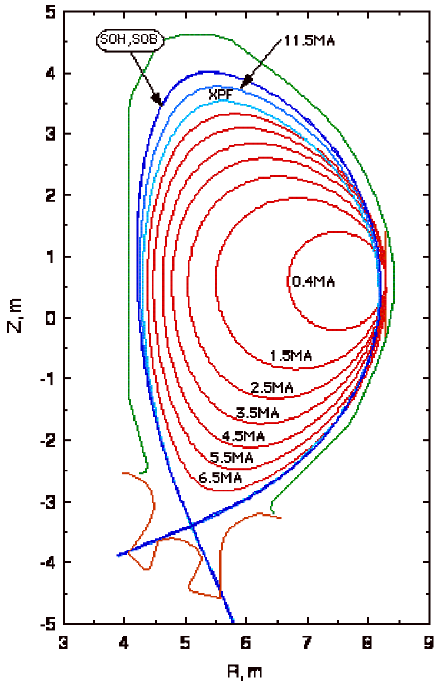

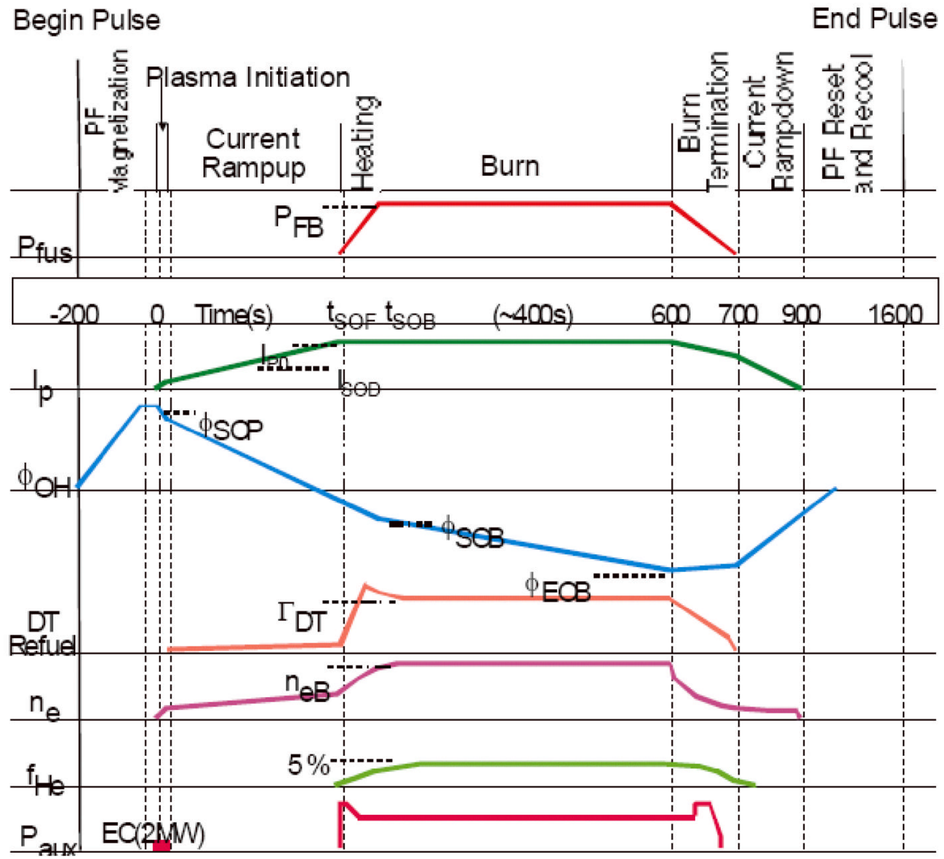

- ITER is pulsed
- A scenario is programmed in advance
- The scenario is then regulated by feedback loops
- This provides a basic mode of operation

Make the whole plasma operation feedback controlled

Develop the tools for full feedback control - OK

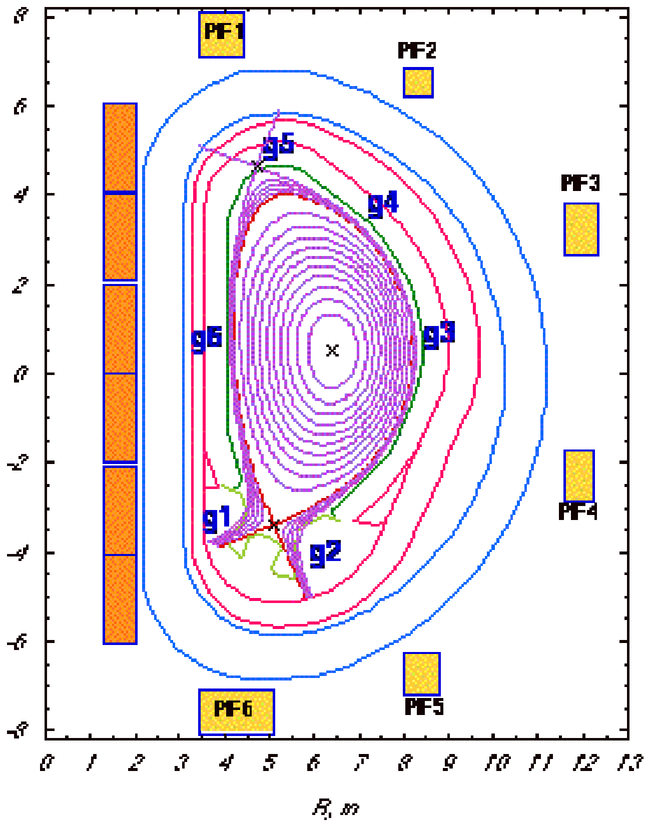

First we need to get the plasma shape correct

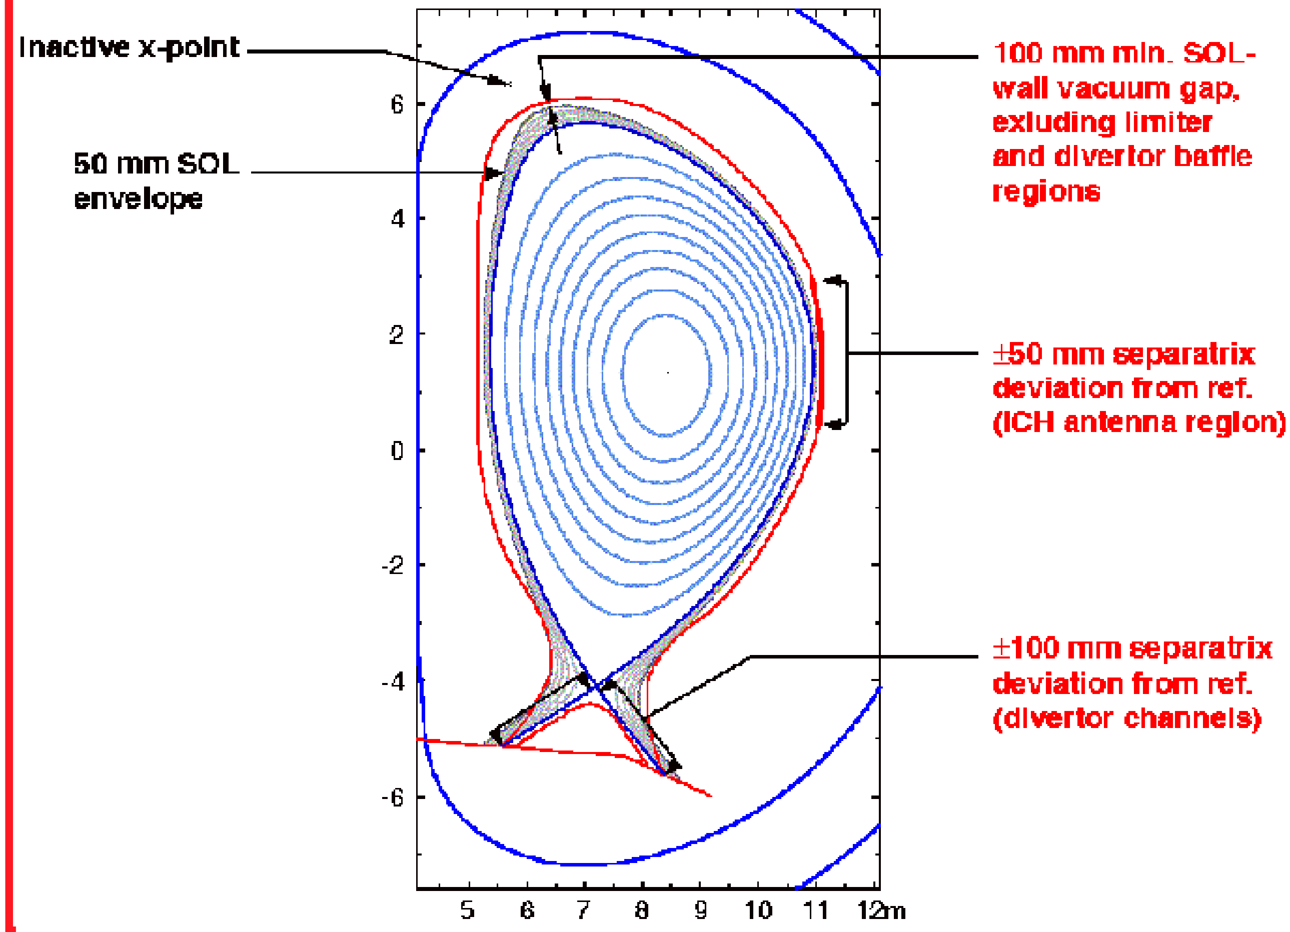

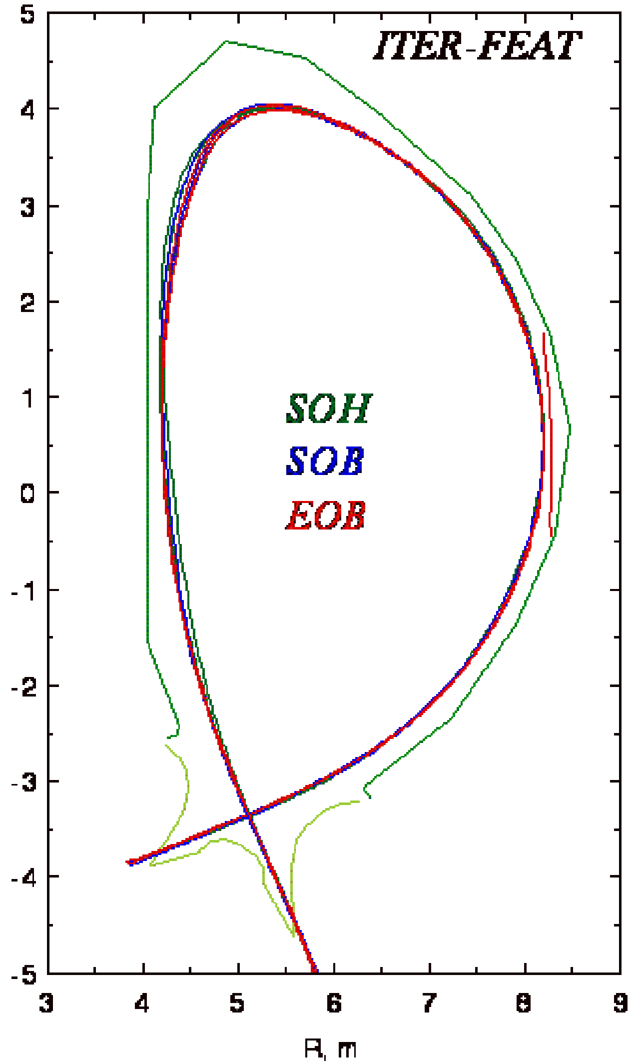

Specify the requirements and develop the algorithms - OK

Feedback control techniques

- Feedback control proceeds in 4 phases
  - Creation of a model of the system to be controlled - actuators and measurements
  - Choice of the parameters to be controlled
  - Choice of the performance requirements of the feedback control
    - Speed of recovery
    - Actuators limits
    - Power limits
    - Time constants of the feedback (AC losses)
  - Choice of the algorithm for generating the feedback and feedforward control

Develop the algorithms and demonstrate solutions by simulation - OK

Going from simple control to optimising control

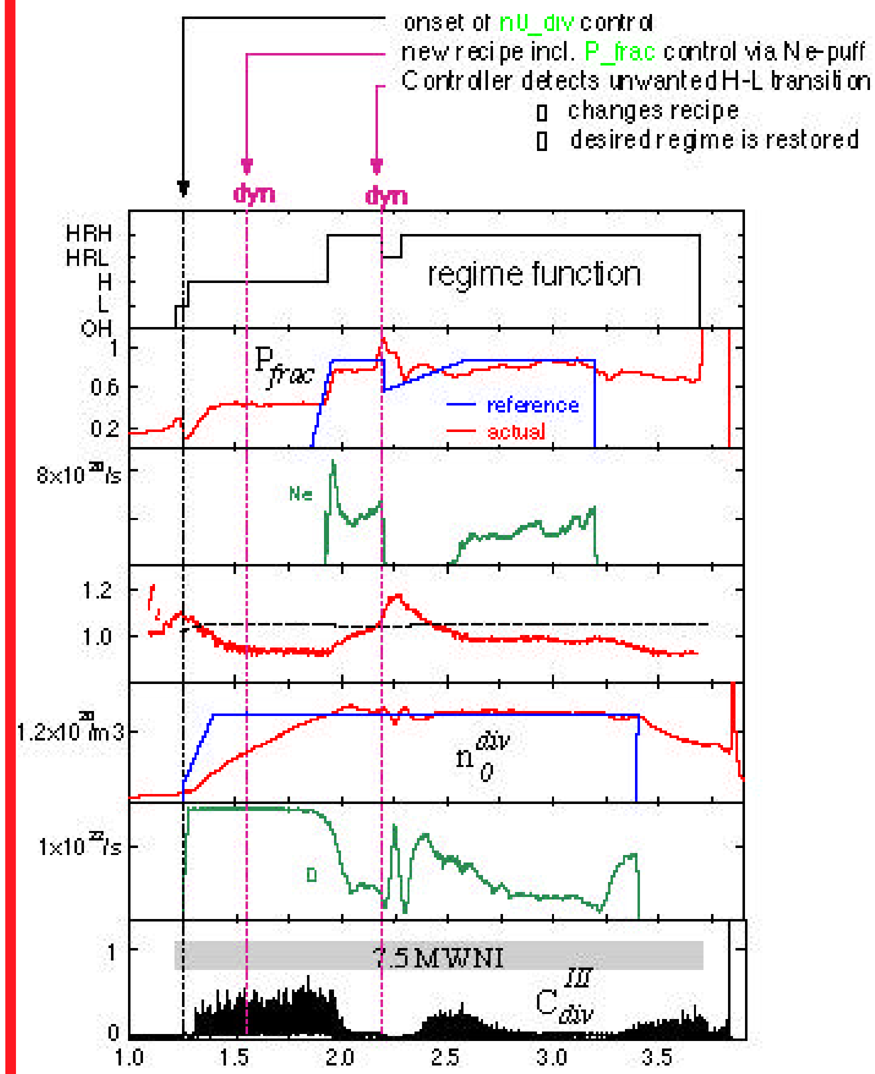

- We need to control the details of the plasma
- Research in the last few years has demonstrated how important this is
- Control becomes optimisation
- 2 old examples

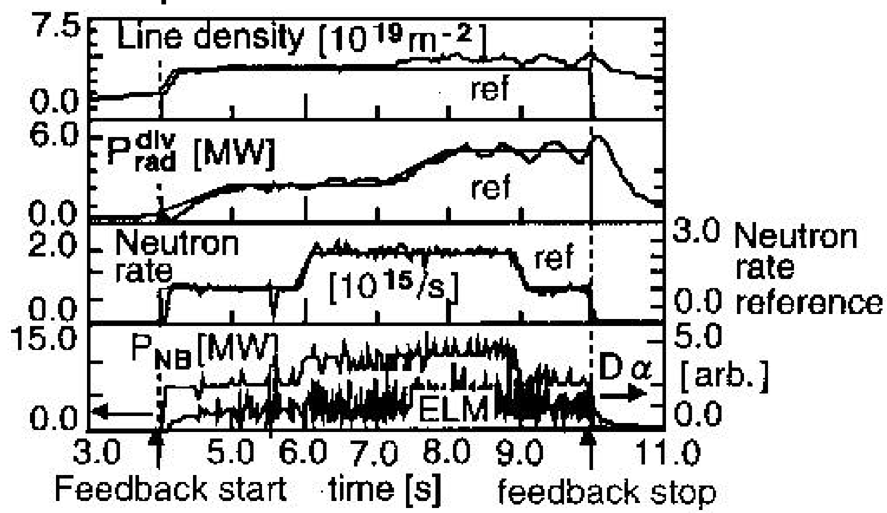

Specify the requirements and develop the algorithms - OK

Where are ITER's technical data handling challenges ?

- Data Source devices - diagnostic measurements ✓
- Data Acquisition Front-ends - ADC's, counters ... ✓
- Data to generate a pulse scenario ✓
- Data for Real-time Feedback Control methods ✓
- Data for Slow Controls ✓
- Data Transfer from the instrument sources ➡
- Data Filtering ➡
- Data Archiving ➡
- Data Validation ➡
- Data Analysis ➡
- Data purpose - existentialism ➡

Specify the requirements - progressing

ITER's likely data rates and volumes ?

- How could we estimate these for ITER ?
  - Extrapolate from the evolution of JET - doubling per year - depends on construction time !
  - Estimate the likely diagnostic requirements - always understated in the past !
  - See what could be done and scale it up with time - explodes !
  - Take existing tokamaks and develop a physics timescale scaling - very small !

	TCV	JET	ITER
Source rate	0.15 GB/s	0.15 GB/s	1-10 GB/s
Volume / pulse	0.25 GB	1.5-10 GB	160-800 GB
Volume / year	100 GB	6 TB	600-2000 TB

Specify our likely requirements - in progress

Are ITER data rates and volumes a challenge ?

- Take the potential source rate 100 GB/s peak

  c.f. LHC-CMS – 40TB / s are planned

- Take 1PB per year

  c.f. LHC – 10PB / year in 2006/7

- Conclusion
  - Rate is probably not a problem with existing technology
  - Volume is not a problem using planned CERN technology - but effort/cost

- Question
  - Should we be solving the rate and volume problem ?

Ask what we should be recording - in progress

Idealised sketch of data reduction / separation

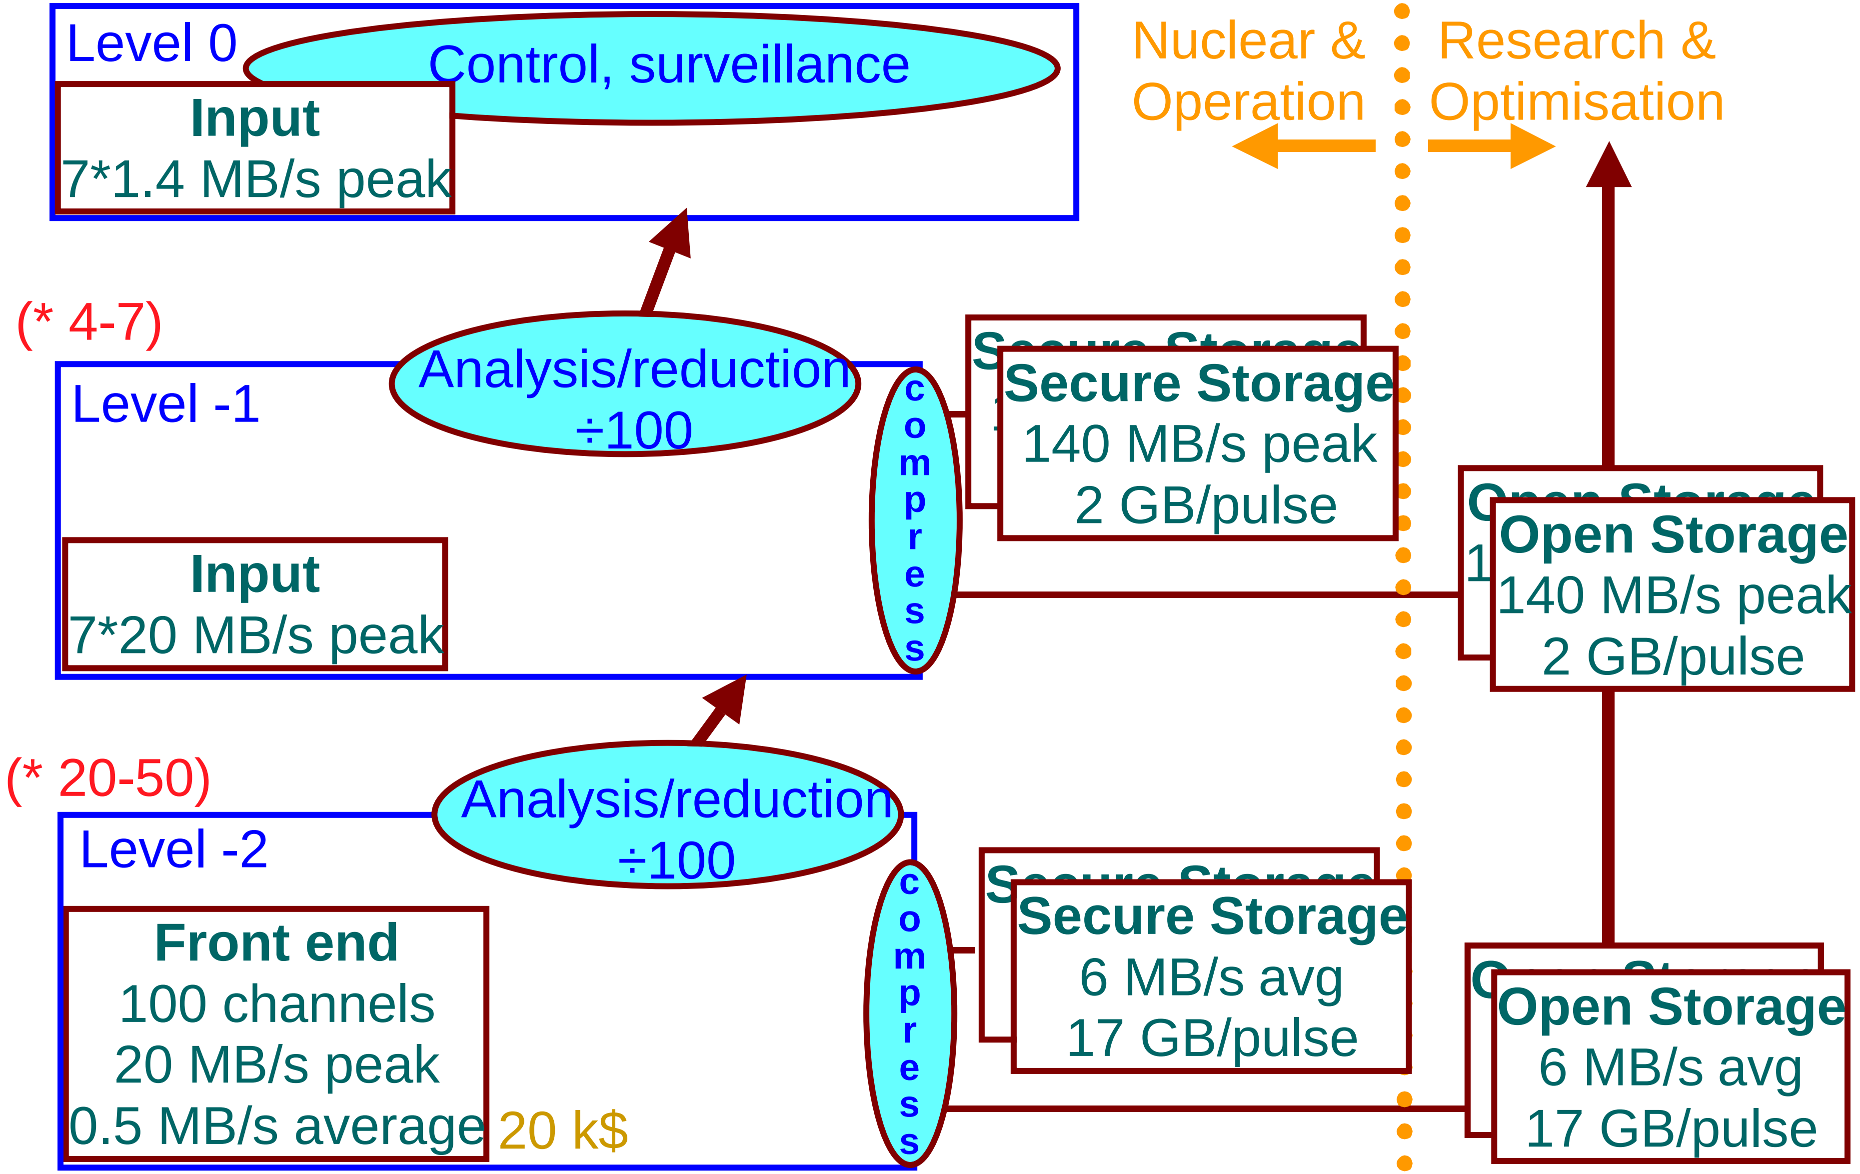

Assumptions made in this sketch

- Data flow reduction is possible in such a "physics" model if
  - Level -2 data can be concentrated with local knowledge + Level 0 data
  - Level -1 data can be concentrated with local knowledge + Level 0 data
  - 10-20 MB/sec Level 0 data is adequate for operating ITER
- No "research" access to the nuclear / operation side
- Loss-free compression - factor 3
- Communication between nodes peak rate ~20MB/s - conventional
- Broadcast information peak ~20 MB/s - conventional
- Streaming rates - conventional
- Storage rates - conventional
- Could do it, but...... storage volume is very large !!

So where are the challenges ?

- ITER is NOT at the cutting-edge of data FLOW requirements
- Data volume - not conventional today, but less than CERN plans
- ITER is NOT at the cutting-edge of data VOLUME requirements

So where are the challenges in this ?

- Data “enquiry” of this data mass is not conventional today
- ITER IS at the cutting-edge of data VOLUME USE
- ITER and LHC have very different data analysis - many small events, easy to filter, vs few, large, dynamical, events
- We want to REDUCE the data recorded because we MUST
  - Remember - 100GB = 5’00’000 screens with 10 traces 1 pixel/point
  - 1 screen/second needs 58 days to flip through a 5 minute pulse
- Otherwise our 1PB/year could be Write Only Memory - cheaper!

How do we concentrate our attention on meaningful data - needs work

CRPP

Why do an experiment - what should it leave ?

To learn something new !

is our model of the ITER correct ?

what are the differences between model and data ?

can we explain them ?

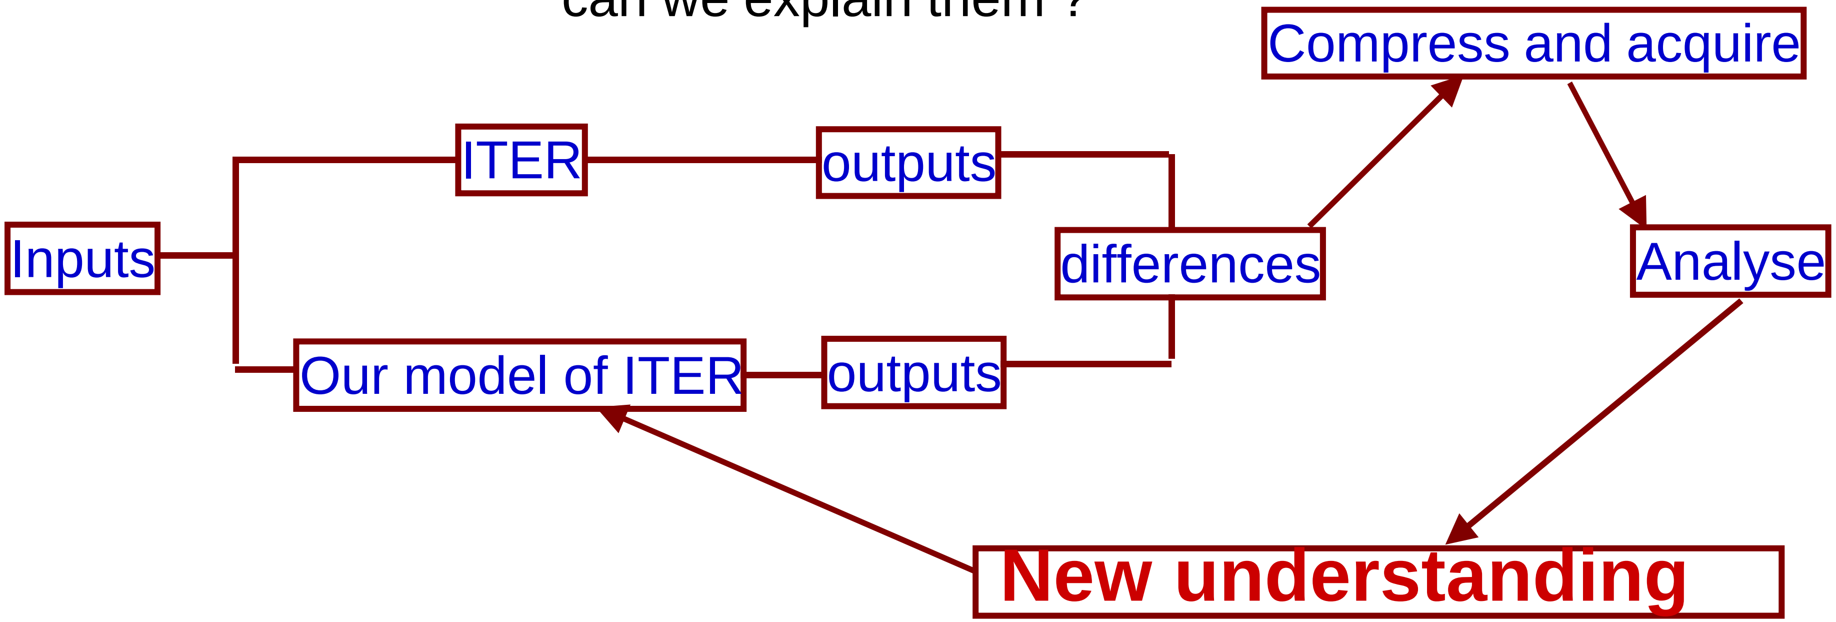

This suggests a “knowledge model filter” for reducing data

Definition of the knowledge model filter

- The knowledge model filter must reject all information which cannot lead to a refinement of the model, therefore of our understanding
- This sounds extremely difficult to generate
- ... But it is exactly what we do with the data today, by eye
- It is naturally adaptive ....

  ....at the start, the models are poor, all data are acquired

  .... as our model is refined, less data are acquired, pulses lengthen

We should be able to limit the data volume

Tools to allow us to use the recorded data - needs work

Can we do this ?

- Wrong question - we have to do this !

- Few people are skilled enough to look at 24 hours of Level -2 raw data

  ➔ they will never look at all the Level -2 raw data

- Not developing novelty filters and archiving the data just in case..

  ➘

  guarantees that data which are potentially interesting will never be seen

- The data volume defeats the original purpose of archiving the data !!

How would we structure such a filter?

- Each system develops its methods to reduce data, based on a predictive model
- Data which pass the “as predicted” filter are rejected, replaced by the model
  - Keep a small volume of regularly sampled data to characterise “normality”
- Signal agrees with model, plus noise of predicted properties
  - Nothing new to be seen, decimate or replace by model
- Low-pass filtered signal is unchanged - very compressible
  - The frequency spectrum / xyz analysis is just like it was last time
- Image has the same structure as previous frame, plus predicted noise
  - Nothing new to report
- Residuals after analysing the data are distributed as predicted
  - Nothing new to report

Why not just make an event trigger ?

- An event trigger has to be perfected to recognise all unpredicted events
  - A contradiction !
- The approach suggested is an “anti-non-event” trigger to retain strange data
- A practical analogy of a continuous overwhelming data source with a massive rejection rate is counter-surveillance....
- Anyone any experience with this sort of approach ??
- But… throwing away data meets tremendous opposition, so the lowest level may be archived anyway - a very expensive “just in case”
- But undersampling == throwing away == OK (!)

How do we put the data handling together ?

- Particular challenges
  - Political identity of the partners
  - Multiple international suppliers, not direct vendors
  - Equipment delivered in kind by a political entity - minefield
- Model A
  - Receive the equipment, interfaces, software delivered from suppliers
  - Get enough bright young people to make it work well enough
  - But… time / culture mix / cost / distance from suppliers
- Model B
  - Impose strict, detailed and testable requirements on the suppliers
  - Get enough bright young people to define the requirements
  - But … lever on supplier if wrong ? Political/financial
  - …… ends up like a late model A

Develop a suitable procurement model - needs work

How do we put the data handling together ?

- Use existing technologies
- Use well tried techniques
- Use standard management techniques
- But… Standish group finds that using standard industrial and well tried techniques we obtain
  - ~30% of systems work as defined, on time, on budget
  - ~40% of systems don't meet time/budget or specs
  - ~30% of systems never work and are scrapped
- Listening to experts could guarantee failure - so choose well !!

Develop an appropriate technology - needs work

XML etc as the ideal technologies

- XML technologies, up to Web Services ?

"… permit language agnostic and OS neutral exchange of all data within an enterprise, using well-tested and internationally agreed standards, independently of any single vendor or cartel of vendors. Web technologies have the weighty backing of the business and commercial communities. Web technologies are being successfully exploited in a wide variety of scientific environments"

Go for it ????

But….

- The keyboard ?

"… a world-wide tool used to develop error-free, robust and high-performance code to run on any known computer platform, independently of the computer language and the operating system, including state of the art massively parallel architectures. Also allows simple and direct person-to-person communication, enterprise-wide, via a wide range of conventional tools."

Go for it ????

Sort out the truth from the reputation, benchmark, test - needs work

A possible mix, personal view May 2005

- Use XML technologies, perhaps Web Services, for all exchanges which are compatible with the performance provided. Carefully select components which are useful and suitable. Carefully define and defend all exceptions.
  - Definition of all system properties, initialisation, characterisation, documentation
  - Definition of finite state machine properties
  - Definition of acquisition and control properties and performance
  - Definition of all data within the project
  - Definition of all the data concerning internally and externally produced systems
  - ……. i.e. all exchanges which are data-management-oriented

- Use project-specific products for all performance intensive exchanges
  - High speed data acquisition and storing
  - Slow controls updates and storing
  - Regular screen updates for operations
  - Tight performance in feedback loops

Select the correct tool for the correct problem - needs work

A possible data-driven scenario

- Country is selected to produce a complex product (tritium plant, neutral beam)
- Country receives the database schemas for all the information which will be required for acceptance of their delivery
- Country develops equipment and satisfies all requirements
- Equipment operated remotely from ITER before shipping, according to the specific data supplied on-line
- Any trouble-shooting done before shipping
- Repeated during on-site acceptance
- Integration into operation/control/acquisition also based on the data provided by the system itself

- But… what if it doesn't work - whose fault is it - are the suppliers competent

Select a suitable procurement and integration method - needs work

Summary of ITER challenges

- ITER presents a tremendously exciting variety of scientific and technological challenges

- We are ready to meet all of these