

Multivariate Control in Fusion Energy: Need for a Virtual Fusion Controls Development Laboratory

T.L. Weaver, D.E. Turner, and J.A. Ramsey
*The Boeing Company, P. O. Box 516, St. Louis, Missouri 63166-516,
thomas.l.weaver@boeing.com*

The Need

Two lines of evidence are converging to indicate that control software will be one of the key areas in which further progress must be made before practical magnetically confined fusion energy will be a reality. The first of these is lines of evidence is the observation that progress in the field of fusion energy has been so great that the temperature and density requirements needed for a burning plasma are close to being achieved¹; but, that stable plasma confinement time is a long way from being sufficient to make fusion energy practical.

The second is that present understanding of confinement methods and control in tokamak configurations indicates that in a power reactor, impacts of parts of the plasma on the first wall, such as from ELMs, could occur several times a minute, and that given the properties of all presently known suitable first wall materials, and the energy present in a power reactor, as few as ten plasma impacts near the same spot will remove sufficient material to compromise the proper function of the wall².

That means that for magnetically confined fusion to mature from the regime the 1000s of second pulses planned for ITER to the six months of steady operation that will be required of an electricity-producing power reactor, either new first wall materials must be developed, or the control system must be capable of preventing the vast majority of impacts of the plasma against the first wall.

The Controls organization within Boeing cannot comment on the likelihood of suitable new materials being developed in the next decade or two that could serve as first wall materials in a power-producing fusion reactor. However, we can comment on advancements in control methods in the aerospace industry, which if adapted to fusion energy applications, have the potential to provide the degree of high bandwidth control needed to prevent most wall impingements by tokamak plasmas.

Controls Divergence Background

The issues present in the control of magnetically confined plasmas and the control of efficient, but marginally-stable aircraft, are remarkably similar. Both problems consist of multiple, interlinked, marginally-stable control loops, and individually stable control loops which can form an unstable higher order system when linked. However, progress in the two fields has developed at widely divergent rates, mostly due to the severe funding limitations in the field of fusion energy controls relative to the funding available for aerospace control system development. As a result, there is now expertise available

in the aerospace industry that could be beneficially tapped to accelerate the development of fusion energy controls if the proper development environment were to be provided.

System Stability

System stability is a major criterion for aircraft as well as a plethora of other dynamical systems, including magnetically confined plasmas. Because stability is a relative term, it is critical to establish a reference point or trajectory about which the notion of stability can be defined and verified in a practical manner. For example, aircraft-related phenomena such as aero-elastic flutter and uncontrolled pitching possess distinct characteristics as they relate to stability, or lack thereof. In the case of flutter, instability is represented by “large” uncontrolled vibration of the wing arising from resonant coupling between unsteady aerodynamic loading and one or more natural frequencies of the wing(s) (or bending moments interacting with wing torsion). The deviation of wing tip displacement, speed, and/or acceleration from nominal operation provides a physical measure of stability of the aircraft as it pertains to breaking or exceeding effective physical operating limits. Measuring the aircraft pitch attitude and pitch rate provides another means of determining the relative stability of the aircraft; but, for this case the issue involves stability in the longitudinal axis. Similar stability concerns exist for the roll and yaw axes as well for other aircraft dynamics such as actuator behavior, engine stall characteristics, etc.

Classical aircraft control, as discussed above, may give the impression that each phenomenon can be addressed independently of others. That is sometimes true; however, controls based on such treatments require large built-in stability margins and tend to yield vehicles with mediocre, but stable, performance. In reality, device control characteristics are all coupled to some degree. That coupling must be understood if performance is to be improved, whether that improvement is higher force maneuvering ability of an aircraft or higher density in a plasma. In aircraft, the amalgamation of data is essential to the realization of a set of upper and lower thresholds of multiple parameters used to specify the stability boundaries at any given instant. Aircraft stability thresholds are generally a function of several dynamic parameters; such as, angle of attack, velocity, dynamic pressure, altitude, and weight, just to name a few. Current trends suggest that the next generation of controllers will include ever-increasing capability, thereby further improving the real-time, automated maintenance of aircraft stability despite unforeseen, destabilizing perturbations created by noise, modeling uncertainty, internal and external environmental factors, system failures, and/or high agility maneuver-based aircraft loads.

Aerospace Controls Development

While direct measurements from sensors on an aircraft can provide information on the stability of a dynamical system, Boeing never enters into tests of control systems without first performing extensive analysis, modeling, and simulation. The first step is analysis. Reliability, fault-tree, failure rate, and similar analyses have been critical required Boeing skills since the founding of the oldest branch of the company in 1916. The consequence for an aircraft company of not competently performing reliability and safety analyses is bankruptcy.

The second step is modeling. Having good mathematical models of the system is important in predicting and expanding stability boundaries in a simulation environment prior to actual flight. Presently, in the aircraft industry, it is not uncommon to perform and rely on mathematical-model-based linear and nonlinear analyses as a theoretical justification of the existence of stable, robust systems and subsystems. Classical and modern methods such as pole placement, Nyquist diagrams, Bode plots, optimal linear quadratic schemes, and recently, (over the last decade and a half) the resurgence of nonlinear adaptive techniques consistently find applications in the stabilization and tracking of aerospace systems. For linear time-invariant systems, stability is determined solely by the eigenvalues of the system matrix. However, time-varying linear and nonlinear systems often require more advanced techniques to achieve and verify stability. A typical aircraft control system utilizes some form of feedback and/or feedforward augmentation that is scheduled based on flight conditions in order to guarantee stability under nominal conditions as well as over some acceptable range of perturbation.

The ability to preserve stability in the presence of uncertainty is referred to as robustness and is, in the linear systems case, quantifiable by a set of metrics known as stability margins. These stability margins provide a quantitative representation of how close/far a given system is from neutral stability (the stability boundary for linear time-invariant systems) as a function of the system gains and parameters, external disturbances, and system loop time delays from the various inputs to outputs. The analogous stability metric designation in the nonlinear regime, called the domain of attraction, is more difficult to quantify. Work is on-going to find nonlinear analogs of the linear gain and phase margin. In the meantime, the quasi-linear approach combining chirp signal analysis and Monte Carlo analysis remains the industry standard for nonlinear systems analysis, with generalized energy-based methods such as Lyapunov analysis and describing function methods (where applicable) gaining acceptance in practice.

In addition to modeling, Boeing often performs various types of “Hardware-In-(the)-Loop” or “HIL” tests. Shown below, in Figure 1, is a photograph of the F-15 aircraft “Iron Bird”. An Iron Bird is a functional laboratory mock-up of the flight control system of an aircraft. In the case of the F-15, the Iron Bird consisted of hydraulic actuators and plumbing laid out in the approximate configuration of the aircraft, aircraft flight control system sensors (hydraulic system pressure, temperature, etc.), prime power hydraulic pumps, and actively variable loads on the flight actuators. Using an Iron Bird, it is possible to verify the flight control system performance and stability margins calculated in the modeling process, by taking the control laws developed in the models through a simulated aircraft flight using real hardware and real loads.

As control system technologies have evolved, HIL testing systems have evolved as well. Boeing aircraft with electric actuators now have “Copper Birds” for control law testing, and several experimental programs have had “Glass” or “Fiber Birds” for HIL testing of their photonic control systems. In addition, the time resolution of the event sensors has increased from millisecond resolution to sub-microsecond resolution where necessary.

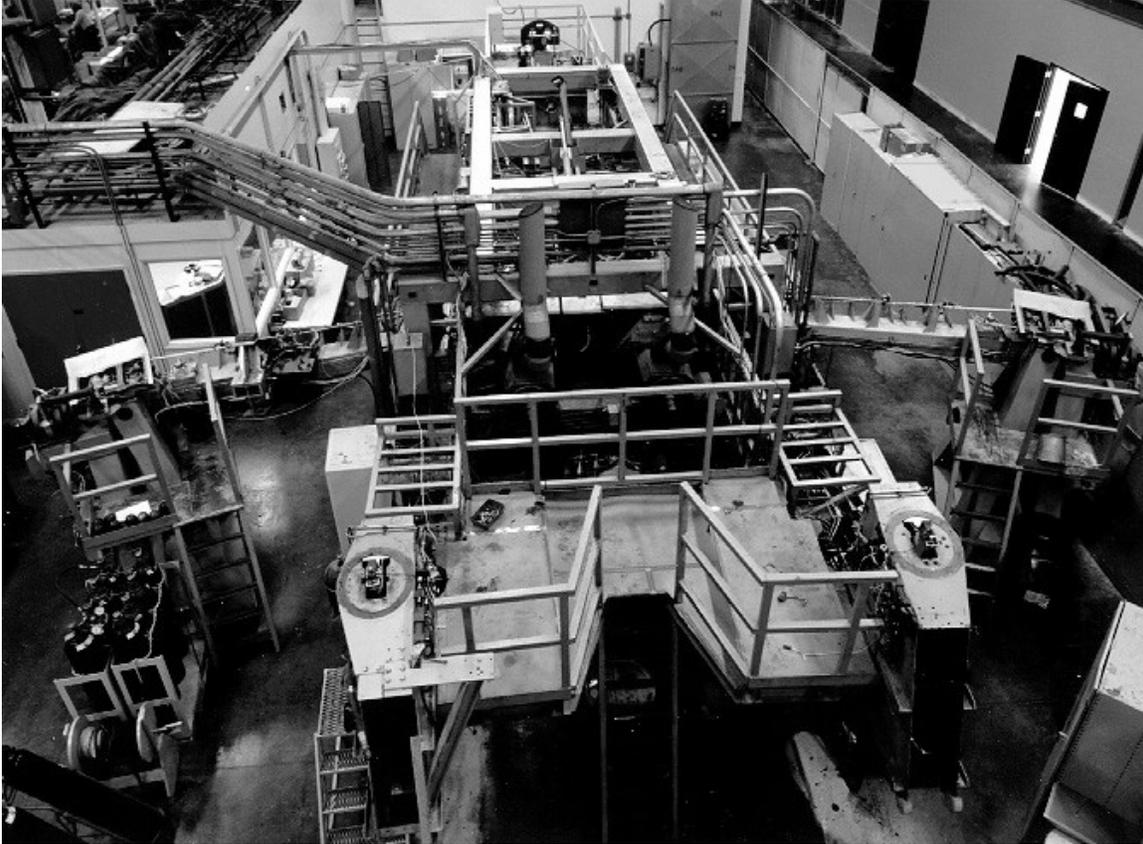


Figure 1 The F-15 Aircraft Iron Bird

In addition to HIL tests that produce only data for analysis, Boeing often studies the stability of control systems by getting pilot input from flight simulators. Shown in Figure 2, is the control room end of a HIL test system that utilizes a pilot. For one of our new design actuators, the vehicle behavior and ten of the eleven flight control actuators were simulated in a computer in the pictured system. The eleventh actuator was real and was working against a load, such as shown in Figure 3, which shows a new-design actuator working against a laboratory actuator which provides the active load. The real actuator was used to verify the ten software model actuators, then the entire system was tested with a pilot at the controls of the simulator pictured in Figure 4.



Figure 2 Boeing Vehicle Integration Test Stand (VITS)

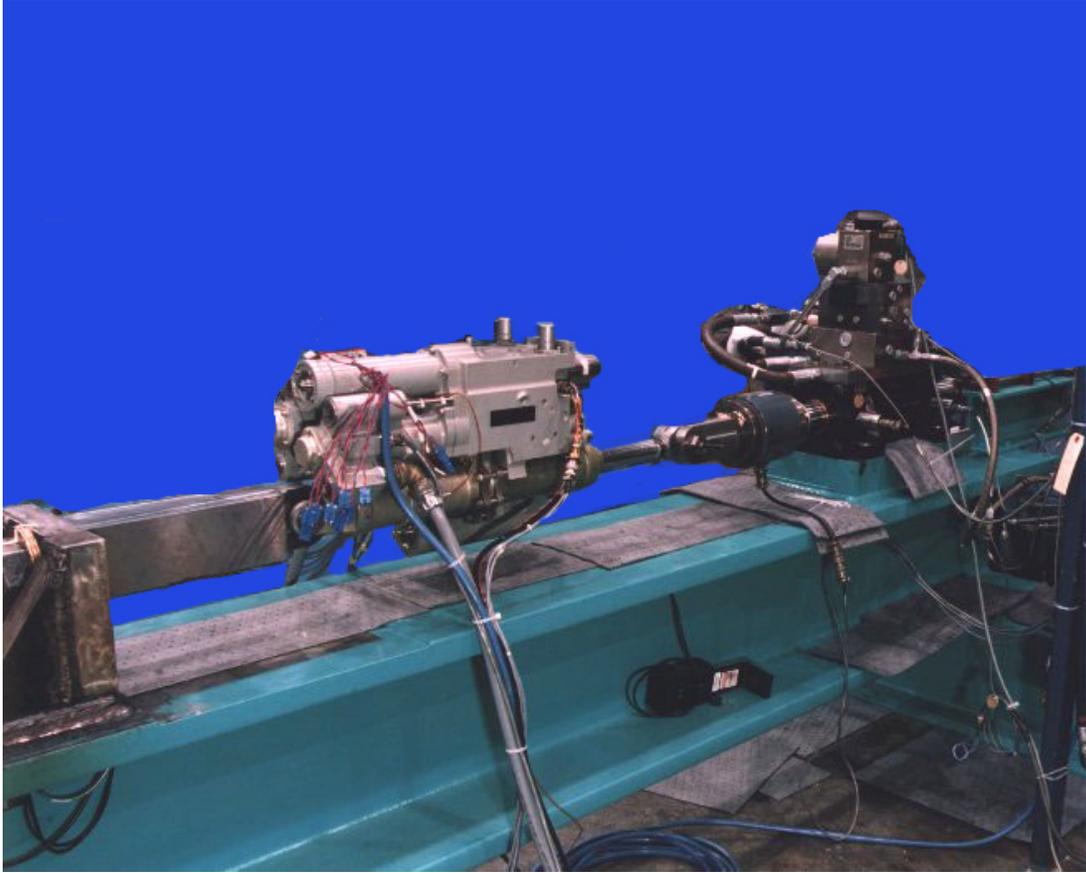


Figure 3 Actuator Under Test working against Active Load Actuator

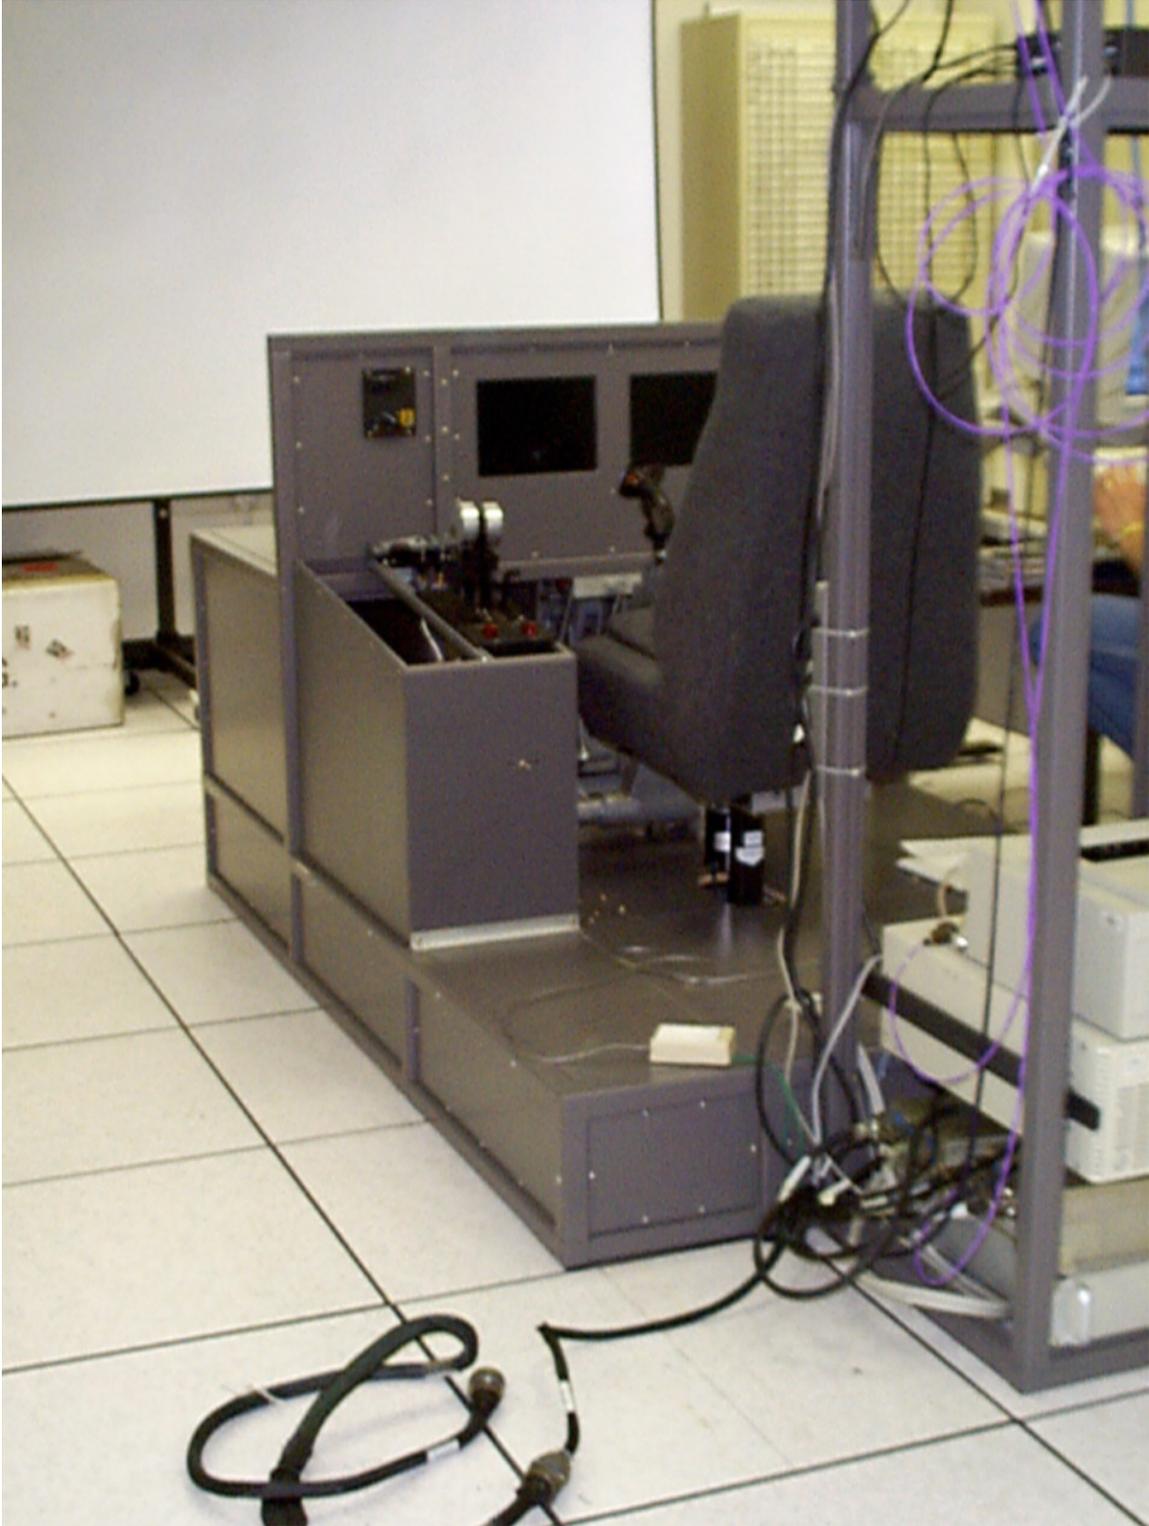


Figure 4 Pilot-Simulator Interface of the VITS

Recommendation

As is apparent from the discussion and images above, much control system work at Boeing is very aircraft-oriented; however, the control issues for fusion experiments are similar to control issues for aircraft, so that by combining the fusion experiment knowledge of universities and National Laboratories with the control system development experience of an aerospace prime contractor, the methods used for developing aircraft control will transfer easily to future fusion energy experiments.

It is our contention that the US magnetic fusion energy community needs to expand the treatment of plasma control from being an adjunct to operation of a plasma confinement experimental facility to being an independent research topic. We believe that without exploration of the interactions among the multiple control loops of a plasma confinement system, the fusion community will never develop confinement systems for burning plasmas capable of producing sufficient plant reliability for it to be the basis of commercial power producing fusion reactors.

The fusion community is already moving in the direction of advanced controls. Considerable work has already been done on making the automatic, closed-loop controls that will be needed for long pulse experiments. However, much of the controls expertise (as opposed to the plasma physics knowledge) needed for further development of advanced fusion plasma controls has already been developed in the aerospace industry, often at government expense. There is no reason the US fusion energy community should not make use of that expertise by teaming plasma physicists with aerospace control engineers.

Such a teaming could be done in several ways. For example, control engineers could be physically moved or electronically connected to fusion experimental facilities. A drawback of that approach is that though such an arrangement would allow the control engineers to accumulate knowledge about plasma physics, they would be limited in their ability to bring their expertise to bear on plasma control problems by the specialization of each facility's control system, and by the inadvisability of performing induced-fault experiments on a real, and valuable, fusion research facility.

A better approach would be to construct a magnetic fusion plasma controls development environment. Interestingly enough, such a facility would not be expensive by the standards of fusion energy experiments. In reality, such a development environment would consist mostly of computers and software, such as the Vehicle Integration Test Stand (VITS) shown in Figure 2, laboratory space for HILs tests, and links to fusion research facilities. A key feature of such a facility is that most of it would be virtual. A control development "campaign" would likely progress as follows. Initially, the control engineers would be linked in non-real-time with physicists at fusion research facilities. They would jointly develop high level models of plasma behavior and control laws for manipulating those models. Testing of the controls would first be done entirely in software in the development environment. Next, if deemed necessary, analogues of real hardware, or elements of real hardware, would be connected to the development system. For example, in place of a mechanical actuator working against a load, an RF source and

tunable antenna would be arranged to work against a tunable RF load under the control of the development system computers and engineers, and at the same time that the development system is running software simulations of magnets and plasmas. This would allow study of the interactions of magnet and RF heating control. A possible penultimate step would be to install real-time data links between the fusion research facility and the controls development system. This would allow controllers at the fusion research facility to feed real-time data from an experiment to newly developed control software for verifications of the operation of the control software. This could also be a way in which planned, but not yet built devices, such as the new antenna from the example above, could be tested without putting the fusion research facility at risk. The last step would be to take the new control software load from the controls development environment and install it in the target experimental facility. It is a significant capability of such a fusion controls development facility that it would provide an environment in which artificial “failures” could be safely introduced so that the effect of failures could be studied, and the lessons learned applied to the software development process.

An important point is that much of the hardware for such a virtual facility already exists. A program to actually create a fusion controls development facility would really consist mostly of the hardware and people needed to modify existing facilities and electronically link them together. For example, an aerospace company, universities, and National Laboratories could work jointly to build the fusion energy controls equivalent of the Vehicle Integration Test Stand, and use it to test interactions among the control software and models of various experiment subsystems and to operate the software and the models with analogues of experiment hardware. That could be easily done, because the VITS was designed to be a general purpose environment for the engineering integration of control software elements with each other and with hardware. A VITS does not care if it pushes an actuator against a load or a magnet current against an inductance. Thus, one of the two VITS, or a VITS modified for fusion, could be used to integrate plant control software (both early versions running in general purpose processors and later versions running in actual experiment control electronics) to assure that the software being developed for each subsystem is compatible with the software for other subsystems. Ideally, modifications to a VITS would be minor enough that run time on the VITS could be split between aerospace and fusion energy work. A VITS could be used to develop control software for fusion energy three days of the week, and used for aerospace on the other two, as represented in Figure 5. That would mean that the fusion energy community would not need to bear the entire cost of the facility, or be under pressure to keep the facility busy to “get their money’s worth” from it.

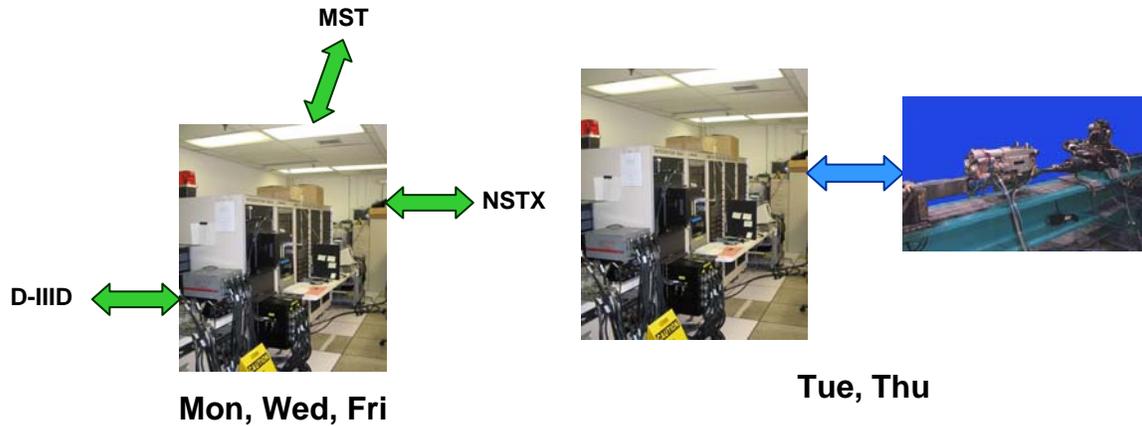


Figure 5 Shared Facility Use

Another useful point about equipment such as that in the VITS is that, if necessary, it can be moved. For example, Boeing built one VITS in St. Louis several years ago. Five years ago, Boeing Commercial Aircraft learned of it and asked the St. Louis group to build one for them. St. Louis built a second VITS, and shipped it to a Boeing facility near Seattle, where it was used to develop the control system for the 787 aircraft. When the 787 program was finished with the VITS two years ago, they shipped it back to St. Louis for use on military programs.

References

1. D. M. Meade; "Some Highlights During 50 Years of Fusion Research"; Presentation; 22nd IAEA Fusion Energy Conference; Geneva, Switzerland; 13-18 October, 2008.
2. A. R. Raffray; "Thermal Effect of Off-Normal Energy Deposition on Bare Ferritic Steel First Wall"; Presentation; Advanced Reactor Innovation Evaluation Study Quarterly Meeting; Madison, Wisconsin; 28-29 May, 2008.