Controlling the World's

As the National Ignition Facility takes shape, so does the computer system to control the giant laser.

THE extraordinary size, complexity, and goals for the National Ignition Facility, now under construction at Lawrence Livermore, make the project one of the most ambitious in high-energy-density physics history. No less ambitious is the computer control system being developed in parallel with the giant $1.2-billion facility. This system, distributed over an area roughly the size of two football fields, will choreograph the workings of thousands of parts to help ensure the reliability of America's nuclear stockpile and keep the United States the world leader in laser and inertial fusion research.

The Livermore engineering team assigned to design the NIF control system faces an immense challenge: every shot will require computerized monitoring and control of some 60,000 points, including mirrors, lenses, motors, sensors, cameras, amplifiers, capacitors, and diagnostic instruments. By working in computerized synchrony, the components will ensure that NIF's 192 beams propagate a 20-nanosecond-long burst of light along a 1-kilometer optical path, arriving within 30 picoseconds of each other at the center of a target chamber 10 meters in diameter. Here, they will strike—within 50 micrometers of their assigned spot—a target measuring less than one centimeter long.

In helping to ensure the success of each shot, the control system will supervise shot setup and countdown; oversee machine interlocks to protect hardware, data, and personnel; generate reports on system performance; provide operators with graphical interfaces for control and system status displays; perform alignment, diagnostics, and control of power conditioning and electro-optic subsystems; and monitor

Most Powerful Laser

the health of all subsystems and components, advising operators of any abnormal conditions. (The box at the right describes how the control system performs beam alignment.)

This system will operate around the clock and be able to control the firing of target shots every 8 hours or less (see the box on p. 9), with an allotted downtime of 7.5 days per year for unscheduled maintenance. Finally, as a major capital investment representing about 10 percent of the total facility cost, the entire system must be easy to maintain, extend, and upgrade.

Flexible Architecture

Paul VanArsdall, project leader for the NIF integrated control system, says the control system architecture is "a truly flexible design because we know, based on over a decade of experience with Nova, that the facility will evolve over the 30 years of its projected lifetime." Lead software architect John Woodruff says the team was well aware of the risks associated with controlling highly complex systems with inflexible system architectures.

The NIF control system is an event-driven control system, as opposed to a continuous system like those in many manufacturing plants or an airport baggage-handling system. Team members believe that the NIF control system will prove so flexible that control managers who are developing event-driven computer control systems for other facilities will be able to adapt the NIF system to their needs.

Project designers also understand that to control costs, they must purchase as many off-the-shelf components as possible. Although most of NIF's hardware is by necessity one-of-a-kind designs, the computer control system incorporates proven microprocessors, workstations, networks, interfaces, and other hardware. Finally, to assure NIF's long operational life, the team has elected to adopt concepts in software development and deployment that meet rigorous engineering standards.

The architecture for the NIF Integrated Computer Control System features two main layers: a lower layer of front-end processors (FEPs)

Deploying 10,000 Motors to Align NIF's Laser Beams

Developing a way to align NIF's 192 laser beams automatically so that they precisely converge on a minuscule target is a formidable task for NIF engineers. The computerized alignment system, now under advanced development, must automatically point the beam through pinholes, center the beam on apertures of several optical sections, adjust the square beam's orientation, focus the beam onto the target, and adjust the angle of the KDP (potassium dihydrogen phosphate) frequency conversion crystals to match beam pointing, all in less than one hour.

Control system project leader Paul VanArsdall notes that the design of the alignment system reflects years of experience working on laser systems. The alignment control system is one of NIF's largest systems. It consists of 600 video cameras distributed at 20 points along the length of each of the 192 beamlines, 10,000 stepping motors, 3,000 actuators, 110 racks, 150 miles of cable, a high-speed network for transmitting digitized video images, and software to integrate all devices.

Says engineer Allyn Saroyan, "We realized that with 192 beams, we needed to automate as much as possible and have operator assistance only when we got into trouble." To determine an alignment error, a high-speed computer performs computations that analyze the digitized video images. Saroyan explains that the system compares the location of the observed beam with where it ought to be, requests a front-end processor (FEP) to move certain motors that precisely shift the beam, asks for another image, makes a second round of calculations, and if necessary, repeats the process. While normally automated, the process can also be performed by a NIF operator in the control room with joystick in hand and eyes trained on a video image.

In all, more than 4,000 devices are manipulated during the alignment process, including mirrors, pinhole wheels, waveplates, lenses, polarizers, shutters, light sources, reticles, and attenuators. The process continually readjusts itself over the course of the allotted time, during which some 48,000 messages are expected to be relayed between FEPs and the supervisory system.

Guiding the development of the system is the staff of the alignment concepts laboratory. This facility simulates a portion of the main laser segment of the NIF beamline, with prototype FEPs used to demonstrate automatic alignment.

interacting directly with laser and target equipment and a higher, supervisory layer to control and integrate the FEPs (Figure 1). The FEPs are distributed throughout the facility, while supervisory computers are located in the NIF central control room. Most FEPs will feature Power PC microprocessors running on the VxWorks operating system, while the supervisory system will be hosted on Sun workstations running a variant of UNIX.

The more than 300 FEPs constituting the bottom layer have been tailored for 19 applications, such as power conditioning and target diagnostics. FEPs will be installed in water-cooled, 19-inch racks and linked to thousands of components, such as laser energy sensors and motorized actuators, as well as to more complex precision instruments (Figure 2).

FEPs will report to seven supervisory control applications, each corresponding to a primary NIF subsystem such as alignment or power conditioning. The subsystems will incorporate several databases for both experimental data and data used during operations and maintenance. The subsystems are integrated by an eighth application, the shot director. In the control room, the system's nerve center, about a dozen operators will be assigned to consoles corresponding to each subsystem (Figure 3).

A distinct segment of the system contains industrial controls. A network of programmable logic controllers will be connected to the industrial control FEPs attached to devices governing, for example, vacuum systems for the target chamber and spatial filters. This segment also includes the Safety Interlock System, which monitors doors, hatches, shutters, and other sensors in order to display hazard levels and ensure personnel safety.

In addition to the 300 FEPs, the NIF computer system will consist of 16 dual-monitor workstations in the control room, several hundred embedded controllers, file servers, and two major data networks featuring both Ethernet and Asynchronous Transfer Mode (ATM) technologies (Figure 4). ATM, running at 155 megabits per second, will connect those systems requiring high throughput for tasks such as transmitting digitized video, timing signals, target diagnostics data, and optics inspection data. Ethernet, running at both 10- and 100-megabit-per-second speeds (depending on traffic requirements), will connect all other systems.

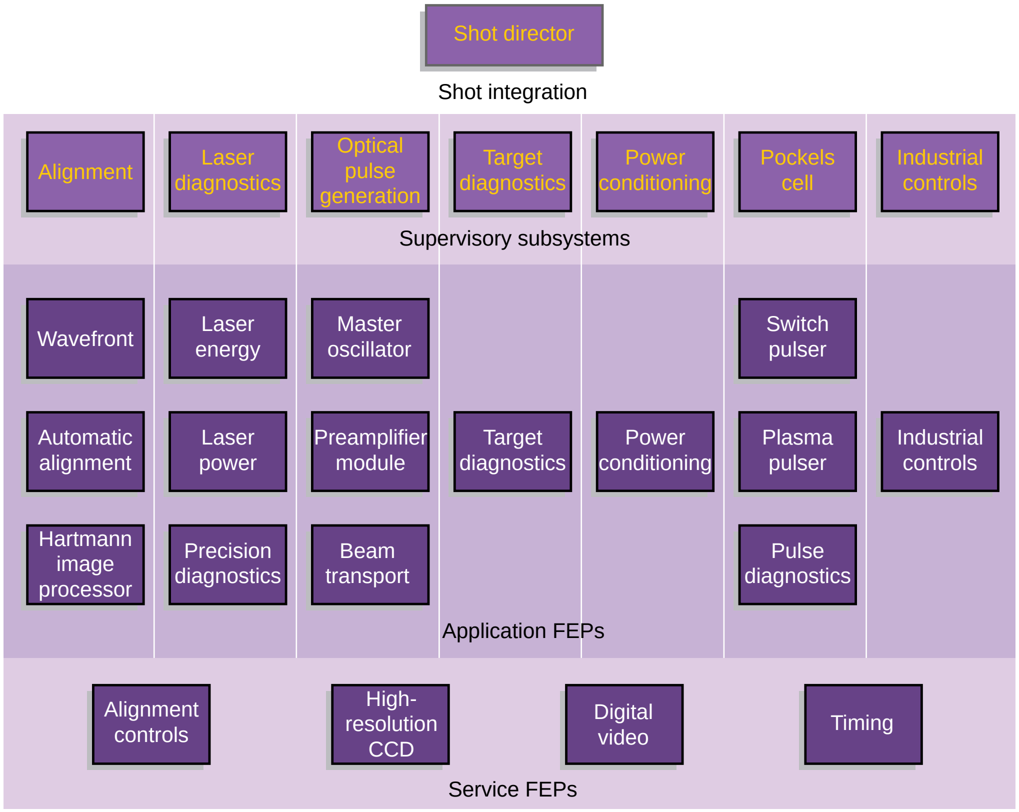

Figure 1. The architecture of the National Ignition Facility's integrated computer control system has four layers. The two main layers are the application front-end processor (FEP) layer (300 FEPs governing 11 applications) and the supervisory subsystems layer, which governs 7 applications. The FEP layer, distributed throughout the NIF laser, reports to the supervisory layer, located in the NIF control room, where all subsystems are integrated by the shot director application.

Owning the Software's "Soul"

The NIF control system, for all its complexity, is designed to "look like one amorphous computer," according to VanArsdall. He says that the development team considered contracting out the control system software, but industry consultants warned that Livermore needed to "own the soul" of the system. Hiring contractors would mean a steep learning curve for people not knowledgeable about lasers and the goals of the complex project. "The key to longevity of software is to own it and be able to replace parts of it as needed," VanArsdall says.

In that respect, he says, "We're building a system that has to last 30 years using a modest team of about 40 people." However, compared with nearby Silicon Valley firms, Livermore is unique because personnel turnover is low. The team also benefits from a wealth of experience with previous Livermore lasers. For example, Fred Holloway, manager of the FEP development effort, has been involved with every major Livermore laser, including the pioneering lasers of the 1970s such as Shiva.

VanArsdall also worked on the early machines, as well as on the automatic alignment system of Livermore's Nova,

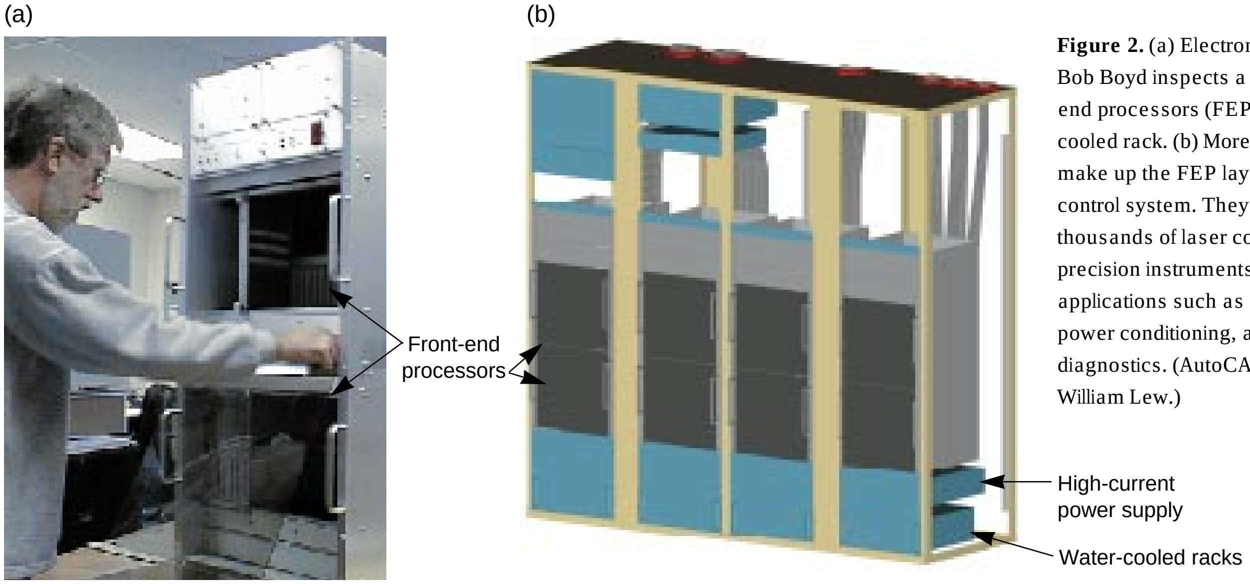

Figure 2. (a) Electronics engineer Bob Boyd inspects a pair of front-end processors (FEPs) in a water-cooled rack. (b) More than 300 FEPs make up the FEP layer of the NIF control system. They are linked to thousands of laser components and precision instruments governing applications such as alignment, power conditioning, and target diagnostics. (AutoCAD drawing by William Lew.)

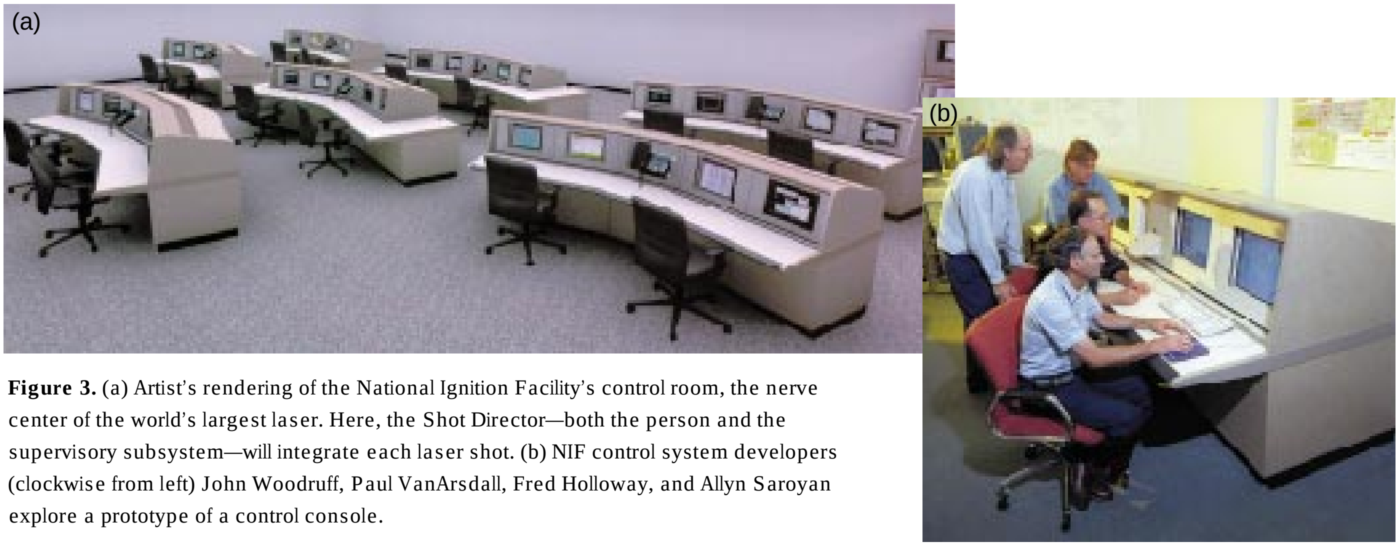

Figure 3. (a) Artist's rendering of the National Ignition Facility's control room, the nerve center of the world's largest laser. Here, the Shot Director—both the person and the supervisory subsystem—will integrate each laser shot. (b) NIF control system developers (clockwise from left) John Woodruff, Paul VanArsdall, Fred Holloway, and Allyn Saroyan explore a prototype of a control console.

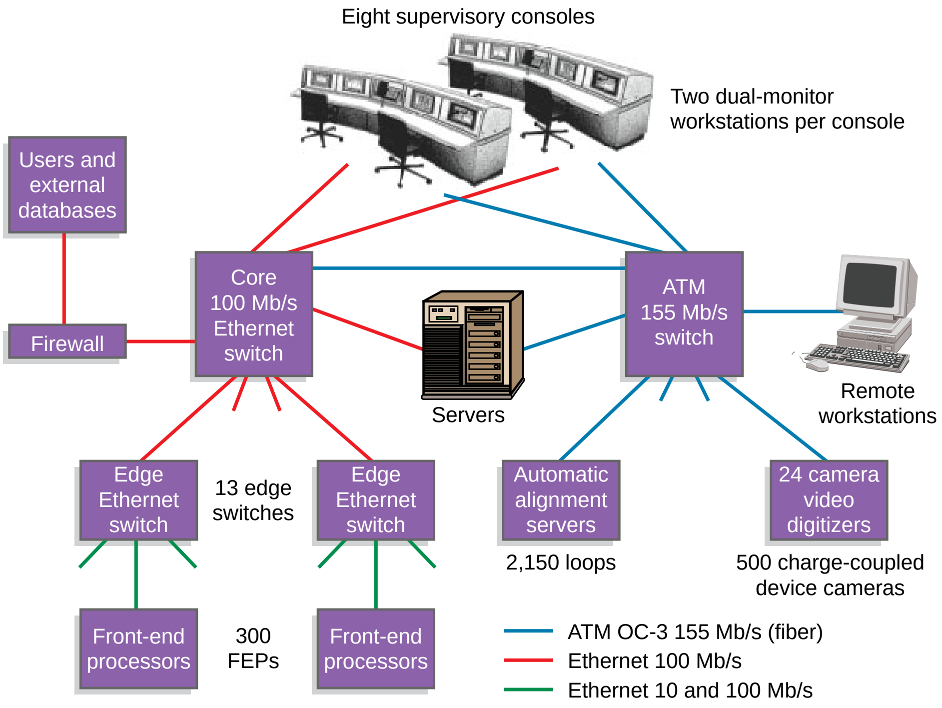

Figure 4. The control system's network employs both Ethernet and Asynchronous Transfer Mode (ATM) technologies to provide flexibility and high-speed data transfer.

for many years the most powerful laser in the world. He notes that Nova was designed in the early 1980s when there were no graphical-user-interface development tools, no fast networks, far less capable computers, and a poorly developed concept of software architecture.

Sustainability of the system over the NIF lifetime is key. Although Nova's control system was upgraded in the early 1990s, simply scaling up its system would not be effective because NIF will be vastly more complex, precise, and powerful. For example, NIF's power conditioning control subsystem (designed by Sandia National Laboratories, New Mexico) will monitor and control 384 high-voltage power supplies, 576 dump switches, 384 gas panels, 384 trigger systems, and 4,800 electrical current monitors distributed throughout four electrical capacitor bays.

Replacing Art with Engineering

Given the overwhelming number of components, the cost constraints, and the unforgiving schedule, the emphasis had to be on increased automation as well as a much greater degree of software development discipline. Woodruff explains, "We are replacing software 'art' with proven engineering techniques that will ensure that the software does what it is designed to do."

For example, the development team has incorporated many of the latest advances in software technology. These advances include CORBA, ADA 95, and so-called object-oriented techniques that together enhance the openness of the architecture and portability of the software from one application to another.

CORBA (Common Object Request Brokerage Architecture)—an international standard developed by the Object Management Group, a consortium of some 500 companies—makes it easy for components and operating systems from different vendors to work with one another. "We have over 400 computers, and they all have to communicate with each other. CORBA allows us to do that. Without it, we'd have to write a lot more software," says Woodruff.

VanArsdall says that the best way to think of CORBA is as a universal "software bus" that allows users to access data transparently, that is, without knowing on which software or hardware platform the data reside or where the platform is located on a network. Used in banking and telephone industries, CORBA hadn't been used in scientific control systems before. CORBA thus represented a certain risk. However, it has worked well in advanced NIF control simulations, and it is now considered the critical core of the architecture.

The team also chose the internationally standardized ADA 95 as the central software language. Ada is used in mission-critical applications such as air-traffic and military command and control systems. The major advantages of Ada are that it supports disciplined software engineering and has proven easier to maintain in the long term than other languages are.

Object-oriented design will assure significant reductions in system maintenance, especially in the face of

Counting Down to a NIF Shot

Preparation for a NIF shot will begin many weeks in advance of its occurrence, when a review committee selects experiments submitted by physicists that will be performed during the shot. An experiment may be scheduled for any number of shots, and one shot may support any number of experiments. One or more beams in a single shot may be set for different characteristics and paths than other beams. For example, even when the experiment's primary objective is target physics, it is also possible to fire 1 to 16 beams into calorimeters that measure laser energy, rather than fire them at the target.

All of the experiment's particulars, such as desired beam characteristics, target instrumentation, energy of the beam, and number of beams to be fired, will be entered into the control system software. The system will automatically issue a warning if certain operating parameters could damage key optical components. "NIF will be operating so close to its design limits that the control system has a number of operations to prevent damage," says control system project leader Paul VanArsdall. (Optical damage over time is inevitable, however, and the control system will catalog the damage and automatically schedule corrective maintenance.)

Several hours before the shot, the control system and operational staff will begin preparations of major NIF components—aligning the beamlines, testing diagnostic equipment, and verifying other parameters. Following these preliminary preparations, a formal countdown of 5 to 7 minutes will feature intricate choreography for sequencing all of the operations that must occur before the shot can take place (see figure below). To determine the countdown duration, the system queries each subsystem how long it needs to complete its final preparations. Because some operations may not be performed on every shot (late insertion of cryogenic targets, for example) and the time to complete these operations may vary, the control software will adjust the countdown optimally for each shot.

Many Fail-Safe Points

During the countdown, the supervisory subsystems controlling the plasma electrode Pockels cell, wavefront control, target, and target diagnostics will place a hold on charging of the main amplifiers and those for the preamplifier modules. The holds ensure that the laser will not be charged before all of its subsystems are ready and represent one form of many fail-safe points distributed throughout the countdown. Other fail-safe points, provided by industrial controls, are designed to monitor and control environmental and safety parameters. The industrial controls system will prevent the countdown from continuing if it detects an adverse condition such as high humidity in the KDP (potassium dihydrogen phosphate) crystals or low vacuum levels in the spatial filters.

"We only get three shots a day, so we've provided plenty of fallback positions," says VanArsdall. "It's better to delay a shot than get lower quality science."

Once the holds are lifted, capacitor charging will begin. When the main amplifier capacitor banks reach full charge, the countdown will enter a critical period in which the laser must either be fired within a few seconds or the shot aborted and the capacitors' charge "dumped" to minimize capacitor aging.

For the final second before laser firing, some 360 "t–1" (one second before firing) devices will move into correct position. Many will have been involved in wavefront control and must be removed from the path of the laser. Others will be inserted at the last second to protect laser diagnostic instruments. Industrial controls will be monitoring the positions of every t–1 device; when all are in position, permission to fire will be given. If permission is withheld, the shot will be aborted.

Between t–1 and t-zero (firing of the laser), an avalanche of events will take place, all under the control of the integrated timing system. At t–zero, some 1,200 accurate firing triggers will be distributed, all synchronized to a single master clock source. After all of the months of planning and a painstakingly designed countdown, a NIF laser shot culminates in bombarding the target with a burst of light energy lasting only 20 billionths of a second.

Finally, the results of the shot will be gathered by the control system and turned over to the experimenter. The laser will return to the idle state and, following a cooldown period, will be automatically inspected by the control system for optical damage in preparation for the next shot.

Of course, the most challenging NIF countdown ends October 1, 2003, when the first full-system NIF experiment is scheduled to occur.

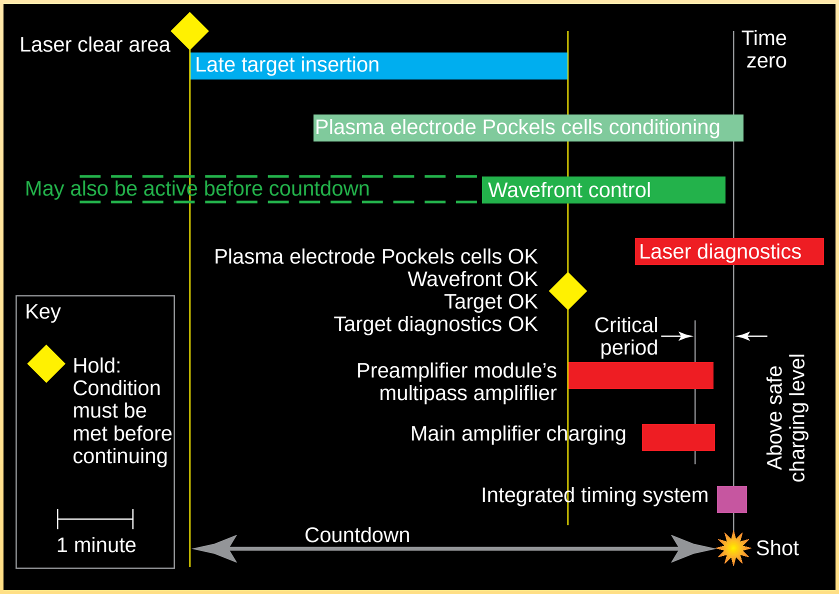

During the final 5- to 7-minute countdown before a NIF laser shot, all operations that must occur before the shot can happen are reviewed by the control system and precisely sequenced.

anticipated periodic replacement of computers and laser hardware. With object-oriented software, for example, supervisory software will not be affected by a change in the type or manufacturer of a motor that it is controlling. That level of detail is hidden in the FEP. With Nova, replacing a motor from a different company required rewriting software at all levels.

The control system software features the use of frameworks, or modular chunks of code. These chunks reduce the amount of coding necessary by providing prebuilt—and tested—components that can be shared for different applications as well as extended to accommodate additional requirements. Says VanArsdall, "With frameworks, we solve each problem once without needing to write huge amounts of detailed code. In effect, we build a software abstraction that we specialize 8 times up in the supervisory system and 19 times down in the FEP layer."

Holloway notes that in developing the FEPs, engineers first came up with a generic FEP with all the frameworks (12) needed for correct operation. For example, the status monitor, a framework reused in many different applications, observes devices (at a periodic rate ranging from 0.1 to 10 seconds) and notifies other elements of the system only when their status changes by a significant amount. This framework's flexibility reduces network traffic significantly.

Figure 5. Rita Bettenhausen and Larry Lagin write and test software in one of many development and simulation laboratories that are interconnected in the NIF control system testbed.

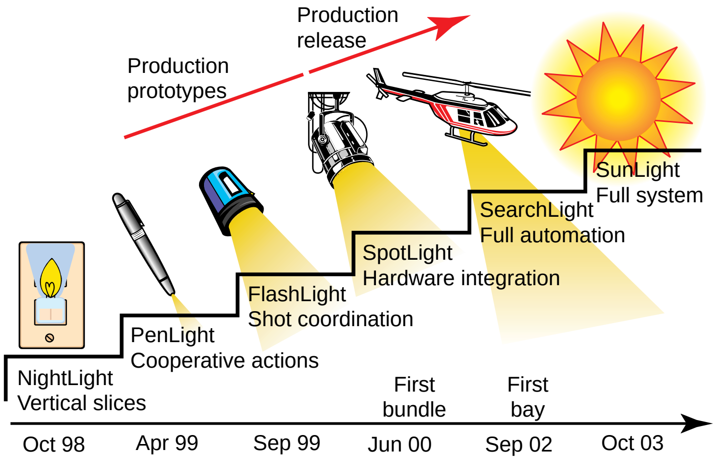

Figure 6. Incremental building of software reduces the risks over the life of the control system's development.

Building a Stairway to Success

Another key strategy to minimize risks is an iterative approach to software construction that has proven effective for projects whose final requirements are not fully known until late in the project. Because waiting to begin work on the control system until all NIF hardware designs are complete would add years to the project and large increases in the total cost, writing the code is being done in stages. Using this iterative approach, developers resolve increasingly smaller risks over the life of the project and add function as major NIF hardware designs are finalized.

As each iteration is built, it is tested and evaluated before the next cycle begins. In that respect, says VanArsdall, "Living with the customer (NIF shot operators) is a wonderful advantage as opposed to hiring a software vendor who goes off and then gives us a pile of code." For example, engineers receive immediate feedback on optimizing the displays on workstations and can begin to train NIF operators at the earliest possible date.

Aiding the software development effort is a web of simulation laboratories that include several FEPs, a small tabletop alignment system, and Ethernet and ATM networks (Figures 3b and 5). VanArsdall says that the laboratory has been particularly useful in modeling full-scale communications to answer questions such as "What happens when I have 5,000 devices chattering all at once to one another? Where are the bottlenecks in the system?"

Incremental building of software provides what NIF managers term a "stairway to success" (Figure 6). Within the first step are "vertical slices" of software that contain the supervisory system, appropriate FEPs, and frameworks for a single task using emulated or prototype hardware. For example, the vertical slice for transmitting video images demonstrates capturing an image of the laser beam and delivering that image to a display on an operator console.

In October 1998, the team delivered NightLight, the first step. "NightLight will probably prove to be the hardest," says Holloway, "because it presented the most risks. We had to show we could work as a team and build all the different applications from generic FEPs." In PenLight, scheduled for April 1999, engineers will apply what they learned in NightLight as they begin to build an integrated control system.

FlashLight is scheduled for September 1999 and will implement all parts of a shot life cycle, including shot coordination. For SpotLight (June 2000), the software will control the first bundle of eight laser beams. SearchLight will test full computer automation of 96 beams in September 2002. Finally, SunLight in October 2003 will operate the completed 192-beam machine.

SunLight is a particularly appropriate term for the final step for NIF's computer control system. After all, NIF will be igniting small fusion targets and creating, however briefly, miniature suns.

—Arnie Heller

Key Words: ADA 95, Asynchronous Transfer Mode (ATM) networks, CORBA (Common Object Request Brokerage Architecture), distributed control system, front-end processor (FEP), National Ignition Facility (NIF), NIF integrated computer control system, Nova, object-oriented programming, software frameworks, Unix, VxWorks.

For further information contact Paul VanArsdall (925) 422-4489 (vanarsdall1@llnl.gov).

About the Engineer

PAUL J. VANARSDALL has been involved in the development of the National Ignition Facility since its inception. Currently, he is deputy system engineer for developing control systems for NIF and leads the unified effort to integrate computer controls from all levels in NIF's special equipment. He is also computer engineering group leader in the Electronics Engineering Department's Laser Engineering Division, which supplies many of the control engineers to the NIF project. Prior to his current assignment, he worked on internal confinement fusion (ICF) laser control systems, including Shiva and Nova, specializing in automatic alignment and supervisory controls. After gaining experience upgrading Nova's control system in 1991, VanArsdall led a development effort introducing object-oriented software technology to prepare engineers for the NIF task. He joined the Laboratory in 1976 after receiving his B.S. in computer engineering (1975) and his M.S. in electronics engineering (1976) from the University of Illinois, Urbana.