AUTOMATIC ALIGNMENT SYSTEM FOR THE NATIONAL IGNITION FACILITY *

K. Wilhelmsen#, A. Awwal, W. Ferguson, B Horowitz, V. Miller Kamm, C. Reynolds
LLNL, Livermore, CA 94551-0808, U.S.A.

Abstract

The automatic alignment system for the National Ignition Facility (NIF) is a large-scale parallel system that directs all 192 laser beams along the 300-m optical path to a 50-micron focus at target chamber in less than 30 minutes. The system commands 9,000 stepping motors to adjust mirrors and other optics. Twenty-two control loops per beamline request image processing services running on a LINUX cluster to analyze high-resolution images of the beam and references. Process-leveling assures the computational load is evenly spread on the cluster. Algorithms also estimate measurement accuracy and reject off-normal images. One challenge to achieving rapid alignment of beams in parallel is the efficient coordination of shared laser devices, such as sensors that are configurable to monitor multiple beams. Contention for shared resources is managed by the Component Mediation System, which precludes deadlocks and optimizes device motions using a hierarchical component structure. A reservation service provided by the software framework prevents interference from competing instances of automated controls or from the actions of system operators. The design, architecture and performance of the system will be discussed.

INTRODUCTION

The National Ignition Facility (NIF) is a 192-beam pulsed laser system currently being constructed and tested at the Lawrence Livermore National Laboratory (LLNL). NIF will become an international center for the study of inertial confinement fusion and the physics of extreme energy densities and pressures. NIF experiments will allow the study of physical processes at temperatures approaching 100 million K and pressures 100 billion times atmospheric. These conditions exist naturally in the interior of stars and in nuclear weapons explosions [1]. When completed in 2009, NIF will provide energetic laser beams to compress deuterium-tritium fusion targets to conditions where they will ignite and burn. In September 2004, the first four NIF beams (a "quad") were commissioned to the center of the target chamber, which demonstrated end-to-end functionality for all major subsystems. By September 2007, thirteen of the twenty-four laser bundles (96 beams) have been commissioned, making NIF the most energetic laser in the world. Ignition experiments beginning in 2010 are planned to liberate more energy than is required to initiate the fusion reactions. The NIF building consists of 2 laser bays, 4 capacitor areas, 2 laser switchyards, the target area, and a core area that contains the control room and master laser oscillator. The laser is configured in 4 clusters of 48 beams per cluster. Each laser bay contains two clusters. Each cluster has 6 sets of 8 beams called bundles (NIF has 24 bundles), which is the fundamental beam grouping in the laser bay. At the switchyard, each bundle is split into 2 quads, with 1 quad from each bundle directed toward the top of the chamber and the other directed toward the bottom.

NIF is controlled by a large-scale integrated computer control system (ICCS) [2]. ICCS is a layered architecture with the lowest layer being comprised of 750 front-end processors (FEP). At the upper layer, ICCS is coordinated by supervisory subsystems including automatic beam control, laser and target diagnostics, pulse power, and shot control. ICCS software is based on an object-oriented framework using CORBA that incorporates services for archiving, machine configuration, graphical user interfaces, monitoring, event logging, scripting, alert management, and access control [3]. Coding in a mixed-language environment of Java and Ada is 85% complete with over 1.5 million source lines deployed. ICCS operates the laser and target area equipment to automatically set up and fire shots every 4-hours [4].

The automatic alignment system is responsible for aligning all 192 beamlines along the 300-meter optical path to focus precisely at spots on the 10-mm-sized target with a tolerance of 10 microns. In an analogy to baseball, this is like hitting the strike zone with pitch thrown from 350 miles away. Automatic alignment runs within the apportioned 30 minutes of the 4-hour shot cycle, which requires the system to function reliably and quickly without operator intervention. To date, automatic alignment has successfully participated in 500 system and 1500 preamplifier shots with no reported misalignments.

ALIGNMENT REQUIREMENTS

Twenty-two separate optical adjustments are required on each of the 192 beams prior to the shot (Fig. 1). Each optical adjustment is managed by a control loop. A control loop coordinates device movements and image processing tasks while mediating resources shared between loops. The automatic alignment system is comprised of 25 separate control systems for the 24 independent bundles and the target area. Each beam is further organized into three parallel segments that can be independently aligned. In total, there are 3,800 closed loop adjustments using 12,000 devices.

* This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.
# Wilhelmsen1@llnl.gov

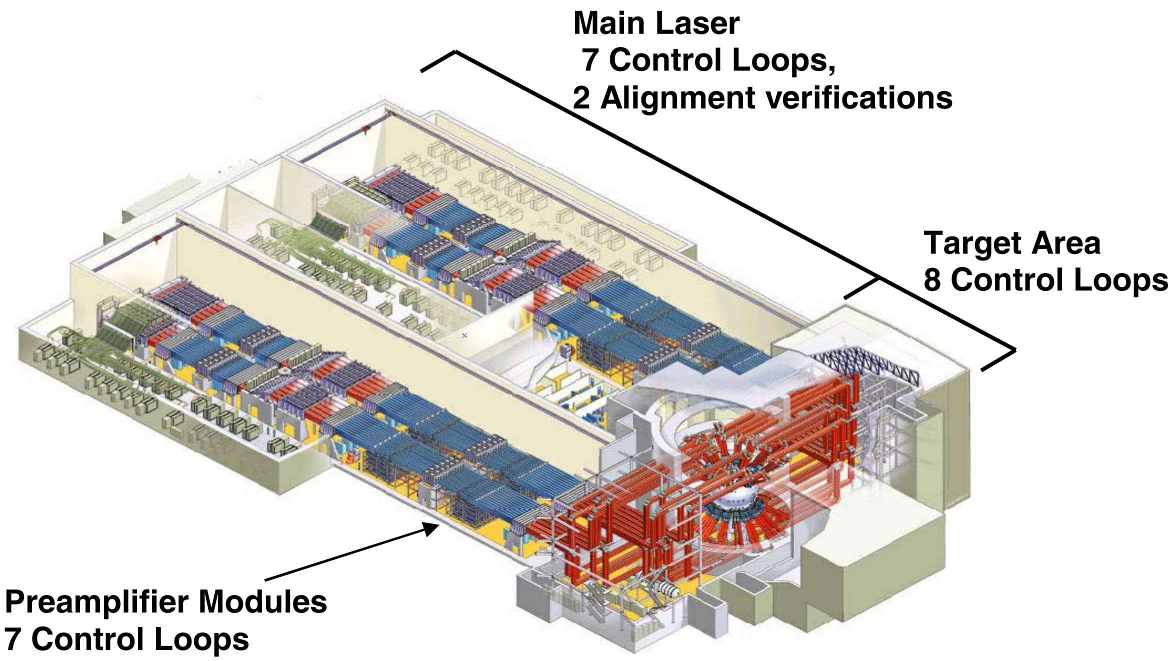

Figure 1: The National Ignition Facility incorporates 3,800 closed-loop alignment operations.

Image processing algorithms are required to locate and characterize alignment features. Many image processing calculations must deliver sub-pixel accuracy. The algorithms are required to be robust to many commonly encountered laser aberrations including wavefront distortion, diffraction effects, and changing light levels.

During routine shot operations, only one operator and control room console manages NIF alignment using the automatic system (Fig. 2). Off-normal image detection is required to qualify images before feature location is attempted [5]. In off-normal situations, the operator is alerted to bring manual controls online for the affected beam to correct the problem. The operator then resumes the automatic process.

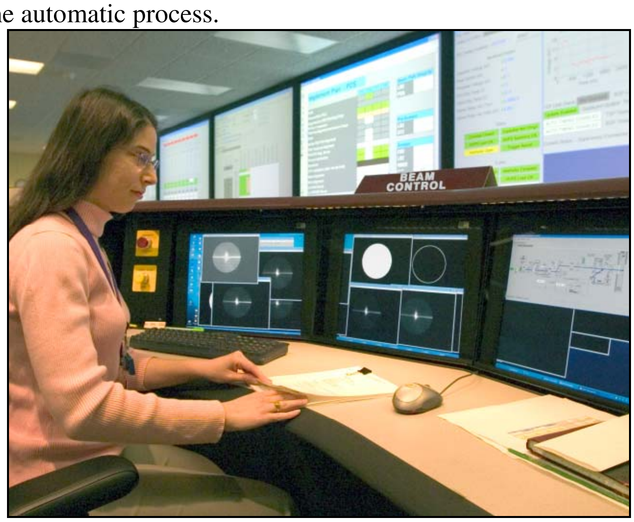

Figure 2: Beam control operator in the NIF control room.

Control loop operations can be generalized into 2 types: centering and pointing (Fig. 3). Centering operations are required when the beam is positioned on the optical clear aperture of mirrors and lenses and to travel down the beam tube. Centering tolerances range from 0.01 to 3.0 mm. Pointing operations correct the angle of the beam traveling down the beam tube. Pointing tolerances range from 0.135 to 10 micro radians.

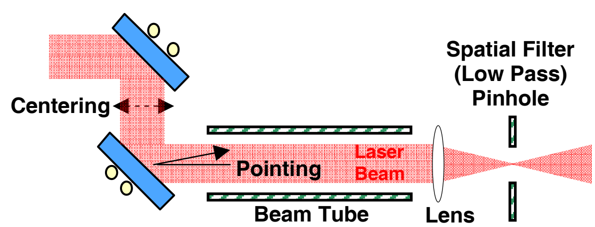

Figure 3: General definition of pointing and centering.

AUTOMATIC ALIGNMENT ARCHITECTURE

The automatic alignment software has three parts: the Segment Manager, Components Manager and an Image Processing Cluster (Fig. 4). Each automatic alignment system (e.g., one bundle) is configured by fetching alignment plans and other definitions from the database. The distributed components communicate to other systems with CORBA. Algorithm tasks from all bundles are distributed to a load leveling LINUX cluster to perform the image processing. The scalability of automatic alignment is assured by demonstrating adequate performance for a single bundle.

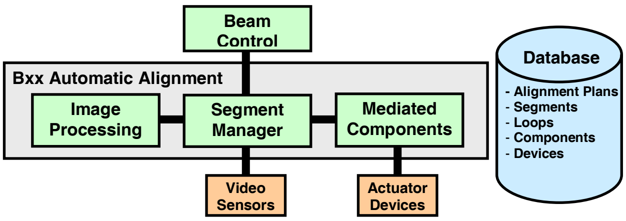

Figure 4: The automatic alignment architecture is replicated to control each of NIF's 24 laser bundles.

To adjust an optic, a loop needs feedback information (e.g., a light source and camera) and control of optics actuator devices in sensor packages and beam path. The three laser segments run in parallel except when using shared components, which requires sequential operations. Sequential operations are initiated when the Component Manager finds resource competition and blocks competing tasks until the resource is free.

Mediated Components

Shared devices are uniquely managed by a special aggregated device object called a Mediated Component (MC) that provides access-controlled commands. These objects coordinate multiple client requests to shared devices by implementing a reservation system and managing a queue of active requests. Accesses to MC objects that are already busy are blocked until free. Note that MC objects can be nested to implement more complex laser configurations. Laser configuration deadlocks are structurally eliminated by enforcing a prescribed order for using MC objects to perform control.

Waiting for device movements to complete dominates the overall alignment time. To optimize performance, the client queue is organized such that clients waiting for a common shared resource are allowed to hop to the front of the queue. This scheduling heuristic yields fewer device movements, thus increasing performance.

Segment Managers

The Segment Manager defines and executes sequences of control loops organized in alignment plans. Alignment plans contain integrated functions for segment initialization, pointing, centering, optics inspection, and alignment verification.

The Segment Manager receives commands from the beam control supervisor or maintenance graphical user interfaces (GUIs). Execution of control loop logic requests device controls directly or, for shared devices, from MC objects. The laser hardware is configured, images taken, and algorithms executed to determine corrective actions. The nominal sequence for an optical adjustment is shown in Fig. 5.

The Segment Manager GUI monitors the status of the entire automatic alignment system. A data concentrator aggregates the status from the 18 segment processes that comprise each bundle. Every 10 seconds, the data concentrator collects all status updates and packages them into a single message for publication to subscribers to minimize communication overhead.

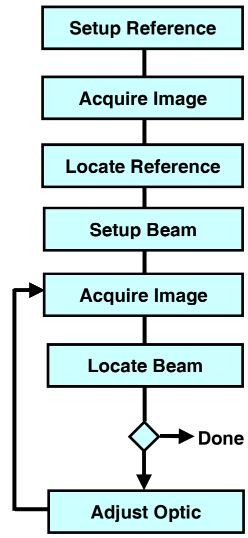

Figure 5: Generic control loop flow diagram.

Loop Execution

Control loops manage all actions required to adjust an optic. Automatic control executes the steps shown in the flow diagram of Fig. 5. First, the loop requests the appropriate MC object to set up the laser and sensor configuration for the reference image. After the MC object manages the shared resources to complete the configuration, the image is acquired and analyzed. This is repeated for the beam image. The error between the reference and beam locations is corrected by adjusting the optics. The loop is repeated until the alignment error is within the specified tolerance, or the maximum retry limit is reached and the loop fails. The most common off–normal condition causing loop failure is poor image quality. In this case, the loop fails and the operator adjusts the loop under manual control. The GUI allows the operator to perform automatic functions step-wise in a manual mode.

All NIF alignment images undergo a set of common algorithmic steps as depicted in Fig. 6. The first step, known as off-normal processing, is used to safeguard the system against accidental misalignment due to processing a false image. Typical off-normal image cases are all-black, all-white, dim, saturated, clipped, and extraneous reflections. The second step processes the image to find the feature location. In the final step, the uncertainty of the determined location is estimated and used to reject the results if the image quality is inadequate to deliver the required accuracy.

IMAGE PROCESSING

Algorithms must balance robustness with reliability, but the results must always be reliable. Algorithms are challenged by images degraded from system effects such as:

- Gradient illumination
- Noise
- Diffraction effects
- Focus variations
- Magnification variations

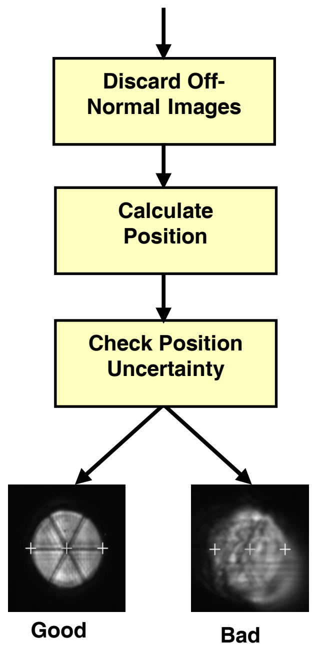

Figure 6: General algorithm block diagram.

Degraded images can cause significant errors in feature location. To be robust, algorithms are specifically designed to mitigate the effects of degraded image quality. Image clipping, saturated images or extraneous blobs may distort features enough that accuracy requirements cannot be assured. To be reliable in these cases, the results must be discarded.

The diverse types of sensor images in the NIF optical system resulted in a suite of twenty-two image processing algorithms (Fig. 7). Algorithms including centroids, Hough transform, templates, and matched filters are responsible for locating features and assessing image quality. The laser optical system may introduce aberrations affecting quality that must be handled. For example, excessive residual error in the wavefront correction system degrades quality by distorting all or portions of the image [6]. Algorithms often successfully process degraded images, and the alignment system is more robust in these cases. Feature location is critical to reliability and accuracy. Therefore, algorithm uncertainty is estimated as a final quality control measure and used to determine reliability.

Algorithm uncertainty is estimated by either of two methods. In the first, a set of centroid calculations is performed by varying the background threshold from 10% to 90% of the signal dynamic range. The result is an ensemble of position estimates for the input image. The variability of the ensemble is called the uncertainty [7]. The threshold method is used for weighted centroid algorithms. In the second, a noise versus uncertainty model is constructed for the image using Monte Carlo techniques and prescribed amounts of noise. The uncertainty is estimated from the model using noise variance of the input image. The noise model method is used for matched filter, Hough transform, and template matching algorithms (Fig. 8). High algorithm uncertainty forces the automatic mode to be abandoned in favor of a manual mode. The intervention of a trained operator to remedy the problem allows normal automatic execution to continue.

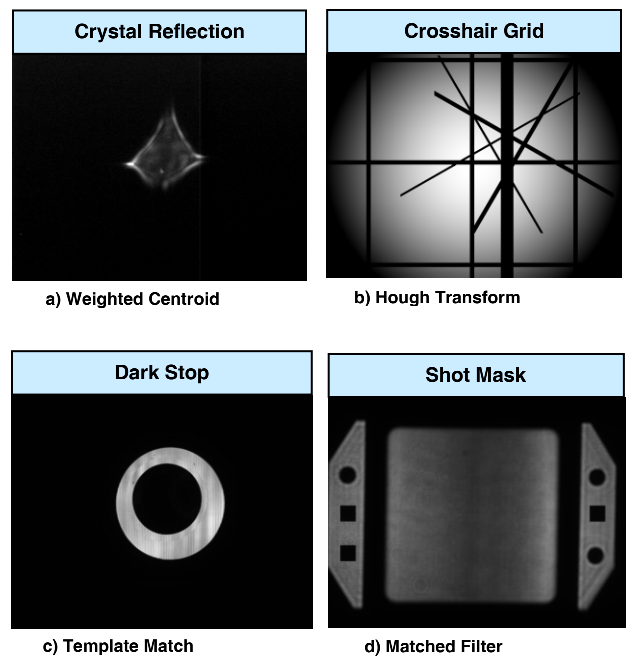

Figure 7: Example sensor images and algorithms.

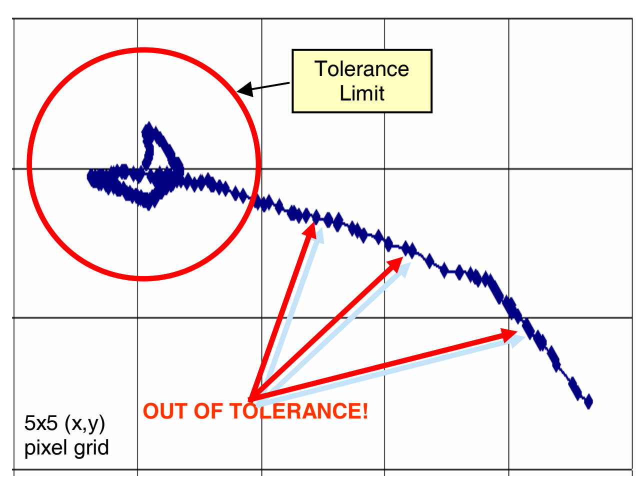

Figure 8: Example demonstrates required sub-pixel accuracy was not met because of high uncertainty.

CONCLUSION

Automatic alignment has successfully demonstrated requirements for speed, accuracy and reliability. Advanced image processing algorithms have delivered robustness to laser environmental conditions, while assuring reliability in the event of off-normal conditions. The alignment system is exercised and proven four times per day on average on the 14 bundles (112 beams) currently commissioned and operational. The independent bundle architecture guarantees the system will continue to successfully scale to the full NIF configuration of 24 bundles and to support precision automatic alignment needs for ignition experiments beginning in 2010.

REFERENCES

[1] E. Moses, "The National Ignition Facility: Path to Ignition in the Laboratory," Fourth International Conference on Fusion Sciences and Applications," Biarritz, France, September 2005.

[2] P. VanArsdall et al., "Status of the National Ignition Facility and Control System," ICALEPCS 2005, Geneva, Switzerland, October 2005.

[3] R. Carey, et al., "Status of the Use of Large-Scale CORBA-Distributed Software Framework for NIF Controls," ICALEPCS 2005, Geneva, Switzerland, October 2005.

[4] L. Lagin et al., "Shot Automation for the National Ignition Facility," ICALEPCS 2005, Geneva, Switzerland, October 2005.

[5] J. Candy et al., "Detection of off-normal images for NIF automatic alignment," in Photonic Devices and Algorithms for Computing VII, Proc. of SPIE Vol. 5907, pp. 59070B-1 - 59070B-12, 2005.

[6] R. Zacharias et al., "Alignment and wavefront control systems of the National Ignition Facility," Optical Engineering, Vol. 43, pp. 2873-2884, 2004.

[7] A. Awwal et al., "Uncertainty Detection for NIF Normal Pointing Images," in Optics and Photonics for Information Processing, Proc. SPIE Vol. 6695, 66950R (Sep. 20, 2007)