A CANON COMMUNICATIONS LLC PUBLICATION

SEPTEMBER 2001

MICRO

Defect Reduction and Yield Enhancement Strategies
for Semiconductors and Advanced Microelectronics

www.micromagazine.com

Brave New Materials

Real-time system classifies defects in copper CMP module

Tool Automation

Integrated process controller manages wafer-handling systems

Surface Cleaning

Statistical methods help develop recipes for dry residue removal

Product Handling and Automation

Using an integrated controller to manage wafer-handling systems

Mario Lento, *Berkeley Process Control*

The transition from 200- to 300-mm wafers will bring a 30–40% reduction in die costs, but the high cost of building the new fabs and the greater unit value per wafer combine to make yields, defect incidence, tool throughput, and downtime critical factors in achieving an acceptable return on investment. To achieve their financial goals, chip manufacturers must demand the highest levels of sophistication in the tools used throughout the manufacturing process, including during wafer handling, transport, and storage.

Unlike in 200-mm-wafer production, the handling of heavier and more expensive 300-mm wafers must be automated to ensure the safety of the wafers and of the fab personnel who handle heavy lots of 25 wafers, which can weigh more than 40 lb. This process change requires the use of overhead or floor-based closed pod–type carriers. Furthermore, with ever-decreasing linewidths, the manufacture of 300-mm wafers necessitates significantly tighter process tolerances than the manufacture of 200-mm wafers, requiring process tool improvements. Equipment also must run longer without failure, experience less downtime, and be more efficient with consumable materials than the systems used in today's 200-mm fabs.

The availability of a single system to coordinate all machine control tasks will improve tool performance, throughput, and reliability.

This article describes the BX*i* controller, a single integrated process control system from Berkeley Process Control (Richmond, CA), which has been developed to meet the automation challenges of 300-mm wafer fabrication. After comparing traditional process control technology with integrated single-controller technology, the article presents two case studies of the new technology's incorporation into process tools.

Limitations of Multiple-Controller Architecture

In the 300-mm fab, automation and process control have become critical determinants of return on investment, requiring new control architectures. In

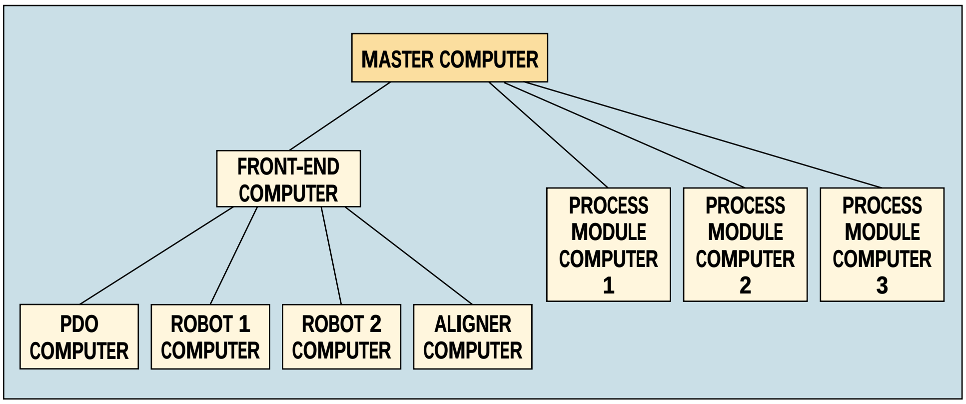

Figure 1: Schematic diagram of a multiple-controller system that relies on a master computer to manage the front-end computer and a number of process chamber computers.

300-mm wafer processing, the traditional multiple-controller architecture used in 200-mm fabs is not a viable option. Traditionally, the control system for a semiconductor tool consists of a computer or PC that serves as the master computer of the wafer-handling system (front-end) and process control (process module) systems. The master computer manages the activities of all the devices that the front-end and process module computers control.

Wafer-handling systems are typically designed with a front-end computer to oversee a number of subcomputers associated with mechanisms that deliver wafers to and from various process chambers. To accomplish this oversight, the front-end computer issues macrocommands to a number of robots and other wafer-handling subsystems. All of the wafer-handling subsystems, which include robots, wafer aligners, and pod door openers (PDOs), are sold with specially designed subcomputers, which are connected to the front-end computer via RS-232 or other serial interfaces. (In some applications, the front-end computer must control some motion and input/output (I/O) functions to run specific devices and mechanisms used in the transfer of wafers to and from the process modules.) The use of subcomputers arose because of the unfeasibility of creating a custom multiaxis motion control system out of various I/O boards, which are programmed in their own languages and plug into one or more of the PCs.

Typically, each process module employs one dedicated computer per process chamber to receive, process, and return wafers to the handling system. Toolmakers usually establish different departments to design front-end systems and process modules, even though the process control department is responsible for creating the master control system to manage the front-end and process module computers.

Figure 1 is a schematic diagram of a system that relies on a master computer to manage the front-end computer and a number of process chamber computers. The master computer also is responsible for a series of other functions, such as connecting the tool to the factory network to provide detailed tracking data on each wafer, exchanging front-opening unified pods (FOUPs) with other tools to ensure that unprocessed wafers enter the tool and processed wafers leave it, and providing tool-system diagnostics.

Figure 2 illustrates a multiple-controller system in which a master computer commands the front-end system to transfer aligned wafers into and out of the process chambers that the master computer oversees. In a hypothetical example using this type of system, the master computer decides to use robot 1 to transfer a wafer from PDO 1 to a notch aligner. It then signals the aligner to align the wafer and subsequently signals robot 2 to pick up the aligned wafer and transfer the wafer into process chamber 1.

This simple process requires several steps. After the master computer has performed a number of checks and has determined that it is safe to begin the transfer of a wafer from PDO 1 into process chamber 1, the front-end computer must query the PDO 1 computer to ensure that it is safe for robot 1 to reach past its open door and get a wafer. Then the front-end computer queries the robot 1 computer to assess whether the latter can be instructed to pick up the wafer (i.e., by determining that there is no wafer on its end-effector or that it is in an idle state). Once the front-end computer is satisfied that it can safely obtain a wafer from one of the slots in PDO 1, a "get" command is issued to robot 1 to pick up a wafer from the PDO and wait in a safe position. Before the front-end computer can issue a command to transfer the wafer onto the aligner, however, it must query the state of robot 2 to ensure that it is not within the collision zone when robot 1 places a wafer onto the aligner.

In 300-mm fabs, automation and process control have become critical determinants of return on investment, requiring new control architectures.

Next, the front-end computer queries the aligner computer to ensure that it is safe to place a wafer on it. After the aligner reports its state, the front-end computer has enough state information to determine whether it is safe for the transfer to take place. Now the front-end computer can issue the "put" command to robot 1 to place the wafer onto the aligner. After robot 1 has put the wafer onto the aligner, it reports its state to the front-end computer. The front-end computer then instructs robot 1 to report its state, and wait.

Once the front-end computer assesses that it is safe to proceed, the align command is sent to the aligner. After the

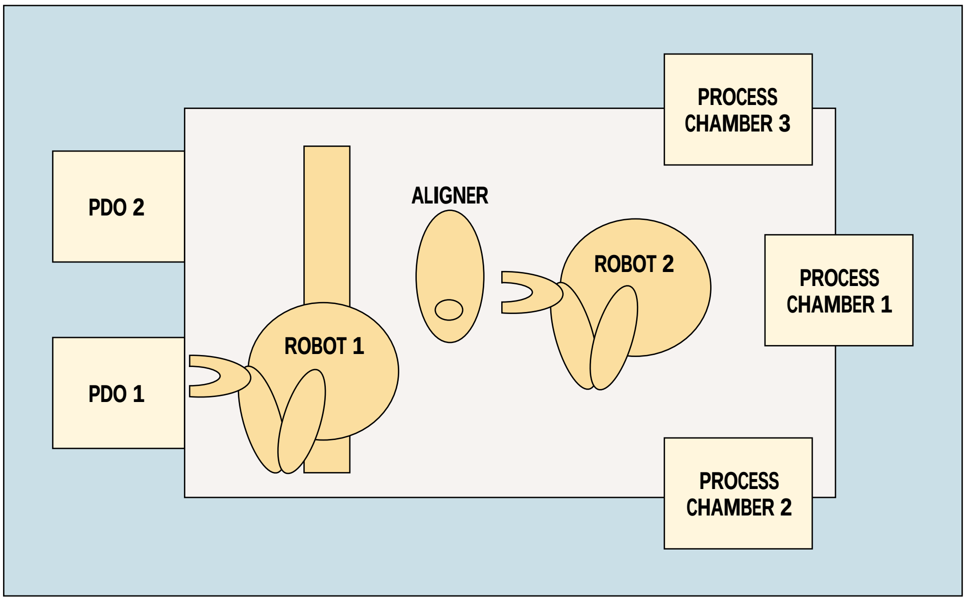

Figure 2: Diagram of a system based on a multiple-controller architecture.

aligner has completed its function, it reports back to the front-end computer. Then the front-end computer sends a pick-and-place command to robot 2 to pick up the wafer from the aligner, transfer it to process chamber 1, place it into the process chamber, move to a safe position, report its state to the front-end computer, and wait.

After robot 2 has sent the message to the front-end computer that it has completed its task, the front-end computer sends a command to the master computer that a wafer has been delivered to the process chamber. Then the master computer sends a command to the process chamber computer to close the process door. After the master computer has been signaled that the door has been closed, it issues a start-process command. At that point, the front-end computer has received enough information to safely issue a command to robot 1 (which has been waiting) to pick up another wafer and place it onto the aligner.

This example illustrates that several devices connected together cannot easily share information about their states. Consequently, control system engineers are forced to write software that attempts to gather all the state information needed so that processes and motions can occur. Unfortunately, the master computer is never aware of the state of any of the devices it is connected to until the device stops and data exchange takes place. In this way, when any device connected to the master computer needs to know the state of any other device, it must wait for each device to relay this information to the master computer. Meanwhile, each device must wait, since any motion or process would alter its state.

Integrated Control Technology

A control system suitable for 300-mm fabrication must be able to manage all machine control tasks, including motion, I/O, scheduling, logic, and networking traffic. In addition, all nodes on the machine network must automatically share state information. The single-controller architecture anticipates that devices in an asynchronous tool must know about one another's states so that they can safely overlap tasks.

The integrated control technology under discussion here can be used to directly manage all of a tool's wafer-handling system I/O devices, servomotors, multiaxis robotic mechanisms, slides, notch aligners, and PDOs. The eight-axis controller is built around a 64-bit RISC processor and an open 100-Mb Ethernet-based I/O architecture, and includes six internally powered drives. It also includes light-curtain and e-stop circuitry with force-guided contactors to the drives. Two or more controllers can be networked and function as one. Once the Ethernet cables are connected and the controllers are powered up, all variable, I/O, and motion states are automatically known to each controller. The same state information is also available via the dynamic link library running on the master PC, which facilitates seamless integration with the fab's information management systems.

Parallel Operations. Unlike traditional multicontroller systems, the real-time state-aware approach to machine control enables subsystems to run in parallel rather than serially. "Dumb" controllers that are not designed to share their real-time states are eliminated, along with the associated serial ports with their inherent communications latencies. For example, as soon as a robot has placed a wafer in a notch aligner and passed through the collision zone on its way back to its starting position, the aligner can begin to perform its task. Similarly, a second robot can remove the wafer from the aligner as soon as alignment is achieved. Should one of the robots fault and stall in the collision zone, the system can respond within microseconds and prevent a crash.

Autocalibration Technology. The calibration of such wafer-handling subsystems as robots and wafer aligners within very tight tolerances is among the challenges faced by 300-mm fab operators. Traditionally, whenever a wafer-handling device is suspected of performing out of specification or actually harms a wafer, or when a component that interacts with the wafer-handling device is worked on or replaced, the device must be recalibrated (often described as "retaught"). A highly trained specialist performs this task by using a handheld teach pendant to jog the robot to the required locations, which are then recorded and stored. Once the locations have been determined (generally within a fraction of a mm), the robot's R (radius), Z (height), and Theta (angular) positions are captured. This entire exercise often requires that the specialist crawl around the tool to achieve visual access (sometimes using mirrors) to all the locations. Since any contact

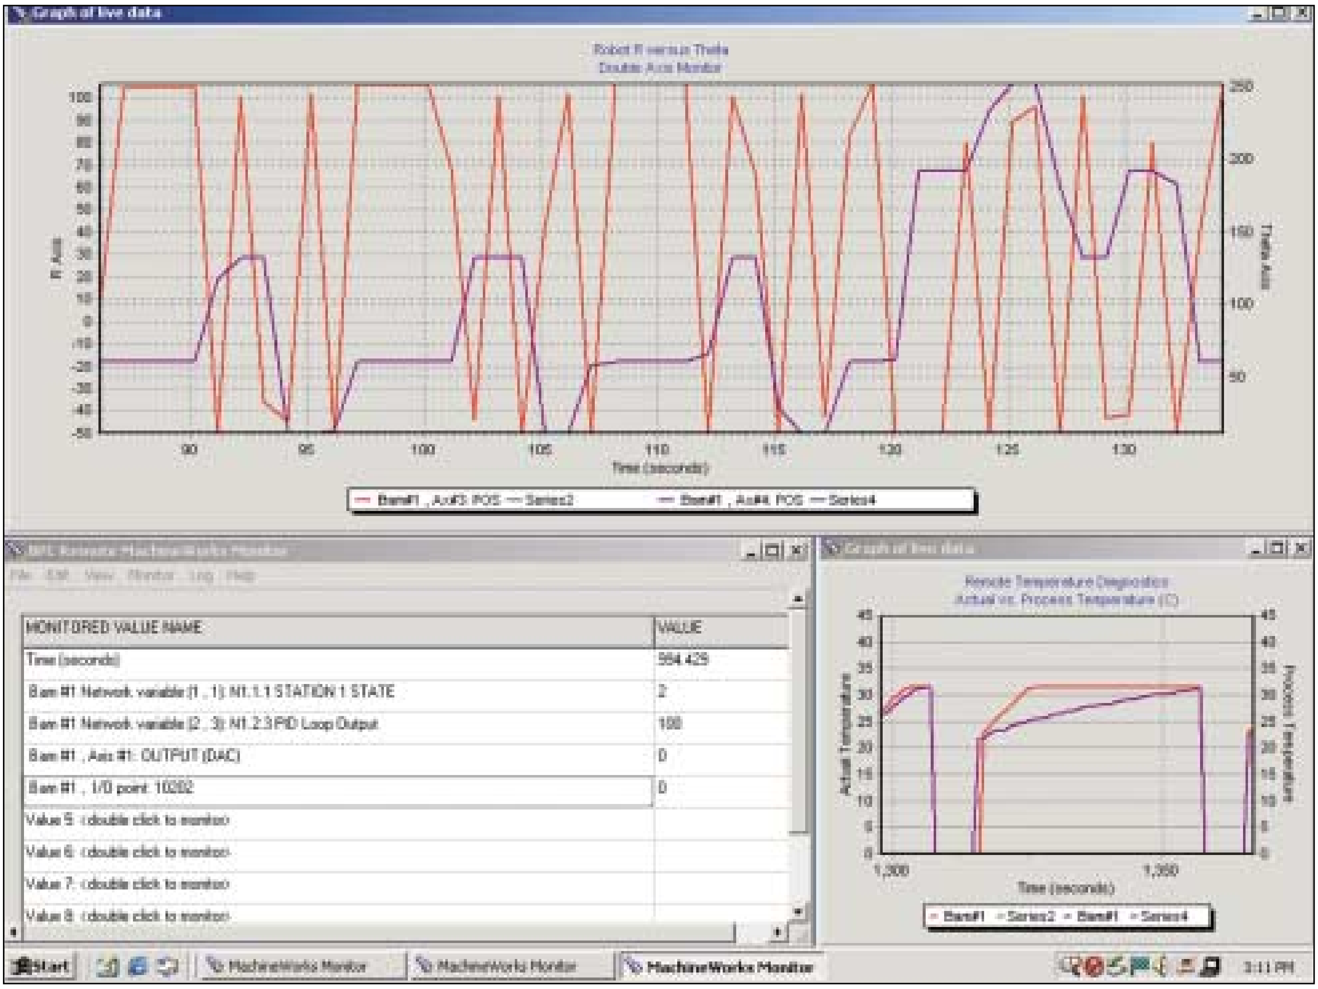

MONITORED VALUE NAME	VALUE
Time (seconds)	554.429
Bam #1 Network variable (1 , 1): N1.1.1 STATION 1 STATE	2
Bam #1 Network variable (2 , 3): N1.2.3 PID Loop Output	100
Bam #1 , Axis #1: OUTPUT (DAC)	0
Bam #1 , I/O point: 10202	0
Value 5: <double click to monitor>	
Value 6: <double click to monitor>	
Value 7: <double click to monitor>	
Value 8: <double click to monitor>	

Figure 3: Screen capture of real-time motion and process data that can be seen by browsing a tool via a laptop PC.

with the tool can generate particles, the technician must use extreme care to minimize contaminating the tool. In the case of a wafer-handling robot, such reteaching typically causes a tool to be out of commission from 2 to 6 hours.

In contrast, the integrated control technology enables a robot to be retaught its positions in 2–4 minutes with the push of a button. The desired teach points can be referenced by latching onto the positions of the robot's axes when a through-beam or reflective sensor is triggered by a known feature of the robot. Another reteaching method includes capturing axis information precisely after a change in torque is reflected by the robot's end-effector when it lightly touches a known datum located within a nominal location. Many robots are equipped with a laser that is used for mapping wafers in a cassette. This laser also can be used as a sensing device.

Diagnostics. Another feature of integrated control technology is that the single controller can diagnose the condition of the controlled subsystems using a variety of techniques. For example, the controller can precisely monitor the amount of torque it takes to actuate each of a robot's axes and compare that information with the values of a signature, or benchmark, mechanism. If the robot requires three times the benchmark amount of torque to command a profiled move, there is a good chance that something is wrong with its bearings. The controller can also evaluate the condition of a belt by carefully driving the robot to its internal hard stop, applying a prescribed amount of torque at the precise point of contact, and comparing the amount of travel (stretch) or slippage that occurred to that of a signature device. In both these cases, if the controller determines that a component is no longer performing in specification, the operator can be alerted to take preventive action when the tool undergoes its next scheduled maintenance. This monitoring capability reduces the likelihood that robots will fail catastrophically and cause unscheduled machine downtime or even destroy costly 300-mm wafers.

A feature of integrated control technology is that the single controller can diagnose the condition of the controlled subsystems.

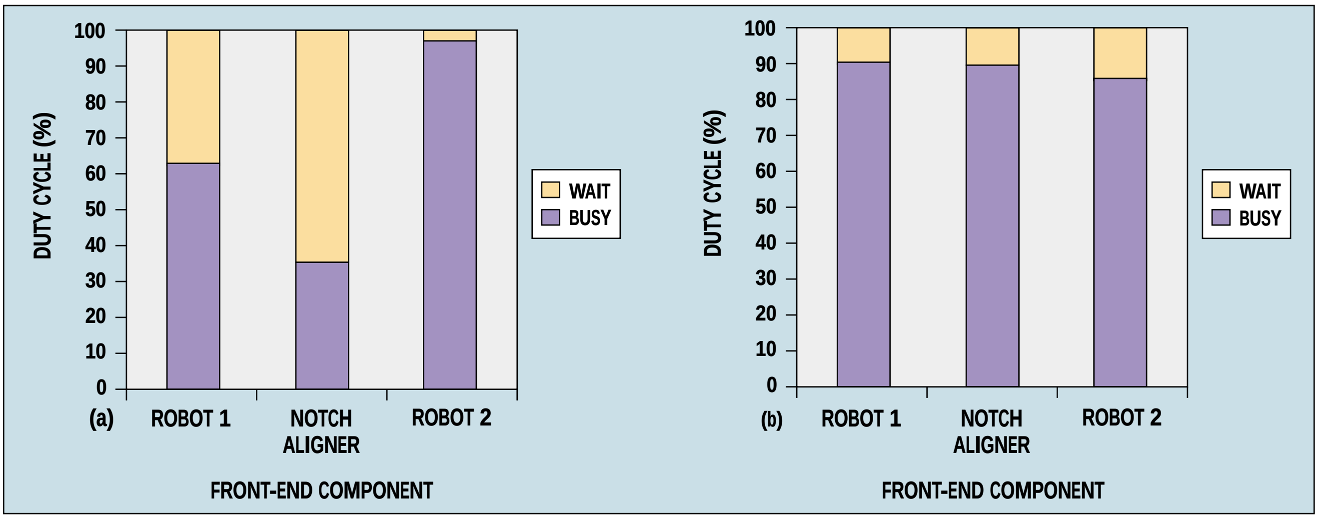

Figure 4: Results of a throughput study comparing (a) multiple-controller architecture and (b) single-controller architecture. The data demonstrate that the new system reduces serial latencies and makes it possible to overlap wafer movements without the risk of collision.

Although robots typically do not fail suddenly, failures can be enormously costly. For example, if a robot crashes in a 300-mm tool that has a throughput of 60 wafers per hour and adds $1000 worth of value to each wafer processed, the crash would contaminate the entire tool with particles, causing at least a day's worth of downtime. It could also destroy a cassette of 25 wafers, which are worth up to $100,000 each. With $1.25 million worth of lost wafers (at an average cost of $50,000 each) and 24 hours of throughput loss valued at $1.44 million, the fab would suffer a loss comparable to the capital outlay required for a typical process tool.

The integrated control technology also satisfies 300-mm fabs' need for Internet-based remote diagnostics (e-diagnostics) down to the device level, presented in SEMI standard E36-0699. The inherently networked nature of the controller enables technicians to browse plant computer systems or individual devices in order to call up error or alarm files, calibrate equipment, install software upgrades, or perform routine maintenance activities without the need to be on-site. Figure 3 is an example of real-time motion and process data that can be seen by browsing a tool via a laptop PC.

Case Study One

Wet bench toolmaker SCP Global Technologies (Boise, ID) has adopted the single-controller architecture for its automated front-end 300-mm wafer-handling system, which incorporates two off-the-shelf R-Z-Theta robots. In the initial product development phase, the company decided to follow industry practice and ran these robots with their accompanying controllers from Equipe (now a division of PRI Automation, Billerica, MA). To load the tool, one robot on a track system took wafers from one of two FOUPs and delivered them to an aligner. A second robot with a stepper-controlled wrist took aligned wafers from the aligner, flipped them to the appropriate vertical orientation, and delivered them to the process area. Unloading was a reverse process in which wafers moved from the process area to the aligner (skipping the align process) and finally to the FOUP.

The original design of the front-end system used a Windows NT–based PC to issue serial commands to the robots, PDOs, and wafer aligners. The multicontroller architecture caused the NT-based PC to lose control of the various devices from the time it sent a command to a specific device until that device completed its task and reported that the command had been executed. Because this approach made it impossible to safely overlap wafer moves, the 300-mm wet bench tool achieved only 25% of the targeted throughput. Tool delays related to the NT-based PC's attempt to manage a number of serial ports while running software and making other decisions also contributed to throughput limitations. After attempting to modify the existing architecture to address these problems, SCP approached Berkeley to revamp the tool's control architecture. The disparate robot computers were replaced with a single integrated controller to run the R-Z-Theta robots. To gain real-time control of the robot's wrist axis, the stepper-controlled wrist was replaced with a servo-controlled counterpart that uses the controller's built-in cam profiler to slave the wrist to the Theta axis.

Figure 4 presents the results of a throughput study of these design modifications. Because the new control system is aware of the states of all wafer-handling subsystems in real time, its adoption eliminated virtually all serial latencies and made it possible to overlap wafer movements without the risk of collision. With the new architecture, if a robot fails to execute a move command, this failure will be immediately known by the controller within microseconds, the mechanisms that would have collided will be stopped, and the situation will be resolved safely.

Case Study Two

A drive for high throughput also prompted Mattson Technology (Fremont, CA) to employ single-controller

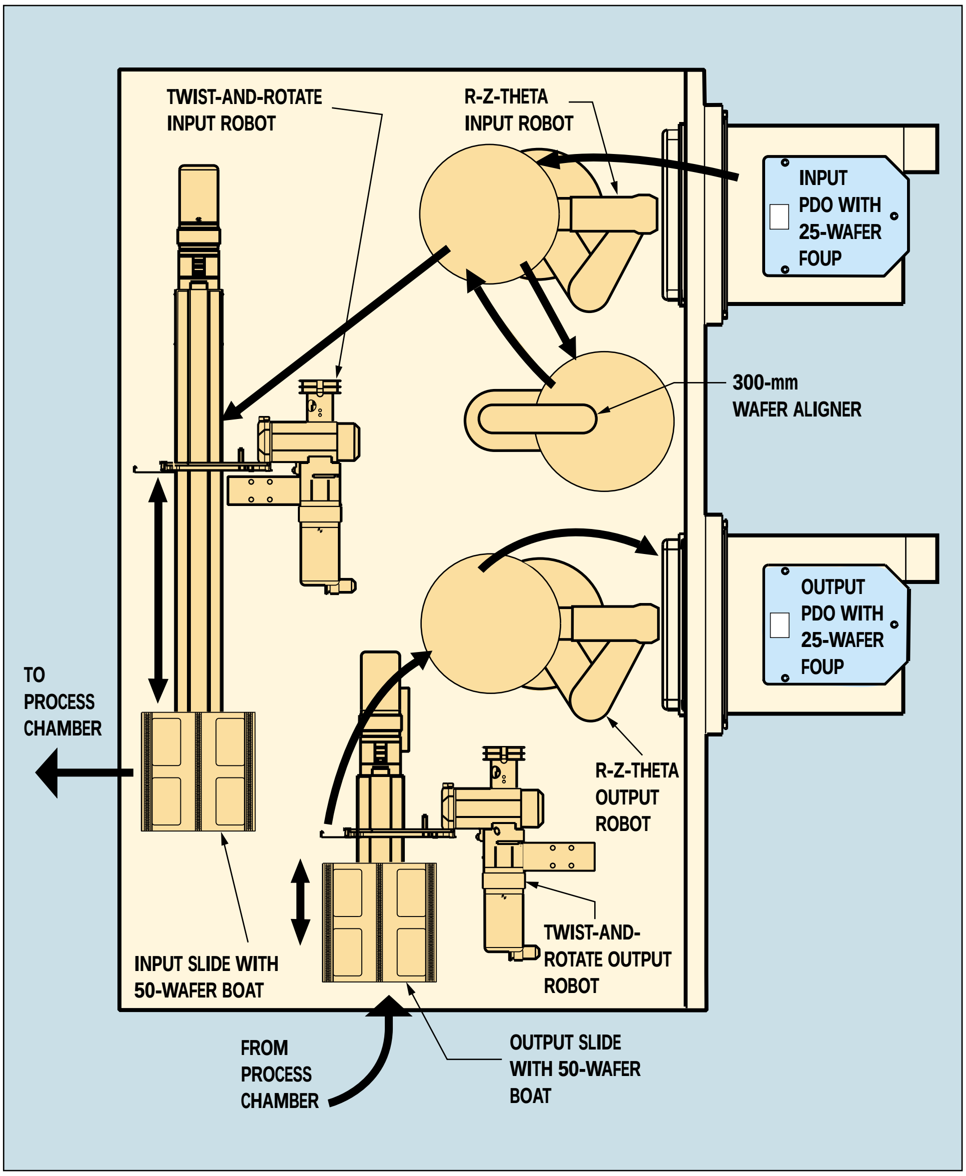

Figure 5: Schematic drawing of a 300-mm wet bench using single-controller architecture.

architecture on a new wet bench tool for 300-mm wafers. In 1999, Steag Electronic Systems (now part of Mattson) wanted to be able to transfer wafers into its tool's wafer-handling system at the same time that the system was being unloaded. Four sets of robotic devices and an aligner were needed to accomplish this task, along with a FOUP wafer buffer, which handles and houses 12 25-wafer FOUPs to ensure a constant supply of wafers to the tool. Because the integrated controller was not yet commercialized, prototypes of the wafer-handling system were built around Berkeley's fourth-generation MWtx 64 controllers. However, these prototypes did not achieve the desired throughput levels of 400 wafers per hour, because the 64-bit RISC controllers did not have enough processing power to perform the sequencing and calculations needed to direct the complex wafer flow, which requires 14 axes of servo motion and hundreds of I/O interfaces.

The subsequent release and incorporation of the fifth-generation controller, which has 10 times the computing capability of the previous model, solved this throughput problem. Since the controller is compatible with prior generations, virtually no software changes were needed. The resulting front-end wafer-transfer system is driven by three Ethernet-networked integrated controllers, which enable wafer-handling sequences to overlap and eliminate process delays.

The controllers' autocalibration capabilities also make it possible to calibrate all mechanical axes in approximately 15 minutes, a significant improvement over manual calibration methods. Real-time machine-state information shared on the 100-Mb machine network and a real-time schedule on the controller make it possible to overlap the scheduling of wafer sequences. An additional controller controls the five servo axes of the FOUP wafer buffer with communication between the buffer and the wafer-transfer system via an SECSII/GEM interface.

As illustrated in the schematic drawing in Figure 5, the input PDO is supplied with wafers from a 25-wafer FOUP. The input robot picks up each of these wafers individually and places it into the aligner. Once the wafer has been centered and aligned, it is again picked up by the input robot, which orients it for transfer to the twist-and-rotate robot. This second robot then places the wafer into a 50-wafer boat. When the boat is filled, an input process handler (not shown) picks up and processes the wafers; meanwhile, an output process handler (not shown) places processed wafers into an output wafer boat. Finally, the output slide indexes this boat so that the output twist-and-rotate robot can transfer each wafer from the boat to the output robot, which places it into the FOUP in the output pod. A photograph of this system appears in Figure 6.

Conclusion

Numerous financial risks are involved in the transition from 200-mm wafers to their larger and more expensive counterparts. For 300-mm fabs to realize sufficient returns

Figure 6: The 300-mm wet bench (illustrated in Figure 5) with the single-controller architecture.

PHOTO COURTESY OF BERKELEY PROCESS CONTROL

on their investments, individual tools will require increased automation and must operate at improved performance and reliability levels. The use of single controllers to manage machine control can help achieve those goals by increasing throughput and minimizing tool downtime, thereby enabling high yields. Furthermore, device-state information from a tool's master computer is accessible to the fab's manufacturing execution system, facilitating total systems integration.

Efforts are also under way to optimize other facets of automation technology. Berkeley is working with General Electric Fanuc Automation North America (Charlottsville, VA) to develop a software-configurable I/O system that uses the integrated control technology. Such a product breakthrough would reduce the need for wiring harnesses and large I/O module inventories, and allow OEMs to assemble tools quickly while minimizing the use of manufacturing floor space. Traditional control architectures require large cabinets for motion, process, and I/O controllers and amplifiers, whereas an integrated controller can be hung directly on a machine's frame.

Mario Lento is a product specialist at Berkeley Process Control (Richmond, CA), where he has spent the last six years working with machine builders in the semiconductor, optical-fiber, and other industries to solve complex machine control problems. Before joining the company, he spent five years in the foundry industry. He received a BS in electrical engineering and industrial management from the University of Massachusetts at Lowell. (Lento can be reached at 510/307-3346 or mariol@berkeleyprocess.com.) M

Berkeley Process Control, Inc.
Your Assurance of Automation Success™

Berkeley Process Control, Inc.
1003 Canal Blvd.
Richmond, CA 94804
PHONE: 510.236.3333
FAX: 510.236.1186
www.berkeleyprocess.com