FEATURE

Award-Winning Control Applications

By James Antaki, Brad E. Paden, Michael J. Piovoso, and Siva S. Banda

Control systems technology is incredibly diverse. Applications of the technology are also diverse, ranging from very large, complex systems to relatively simple, microlevel devices. System costs may range from millions for large industrial plants to just a few dollars for mass-produced products. Performance is also diverse, ranging from ultra precision to only modest levels that are nevertheless completely satisfactory for their unique requirements. Furthermore, many diverse methodologies are available to control designers for use in increasing performance, reducing cost, improving robustness, and/or achieving a variety of other benefits.

Virtually all dynamical systems—whether mechanical, electrical, chemical, economic, or social—can be improved with control technology. Most practical implementations, however, still require challenging innovations by designers to fully realize the potential of our diverse methodologies. This article reviews three recent examples of such innovations.

Streamliner Artificial Heart by James Antaki and Brad Paden

The quest for a mechanical artificial heart is motivated by a tremendous human need. In the United States alone, 700,000 deaths per year are attributed to heart disease [1], and it is estimated that 35,000 to 70,000 of these lives can be saved with some form of mechanical circulatory assist. Worldwide, artificial hearts could save an overwhelming

IEEE Control Systems Technology Award

The applications presented here have won the IEEE Control Systems Technology Award in the past three years. The CST Award recognizes "outstanding contributions in either design and implementation or project management, pertaining to control systems technology." This award especially focuses on achievements that yield measurable benefits, so reviewing award winners helps illustrate the real-world impact of advanced control.

Recent winners of this award are:

2001	James Antaki and Brad Paden	1994	Rajan Suri and Gregory Diehl
2000	Siva S. Banda	1993	Alan Laub
1999	Michael Piovoso	1992	George Meyer
1998	Guy Dumont	1991	Gunner Bengsston
1997	Joe Chow	1990	Ken Lorell and the Lockheed Team
1996	Tom Banks		
1995	Suresh Joshi	1989	Edgar Bristol

In this article, each summary is written by the engineers who were recognized for these award-winning applications. This article was coordinated by Michael K. Masten.

This article was edited by Michael K. Masten (m.masten@ieee.org).

0272-1708/02/$17.00©2002IEEE

200,000 lives per year. Many of these people could also be saved through heart transplantation, but the supply of donor hearts is limited—a mere 2,400 in the United States. In addition, transplant patients suffer the side effects of immunosuppressive drugs needed to prevent transplant rejection. Hence, mechanical circulatory assist is viewed as a very promising treatment for end-stage heart disease.

Artificial hearts have a long history dating back to the early animal studies of Kolff and Akutsu in 1957 [2] and the first Congressional funding in 1964 [3]. The phrase "artificial heart" is applied to a number of treatment technologies incorporating mechanical circulatory assist devices. One technology is the total artificial heart (TAH), which involves the removal and replacement of the patient's own heart with a mechanical pump. The second—and leading—approach uses "helper hearts" called ventricular assist devices, with the *left* ventricular assist devices (LVADs) being the most common. In the heart, the right ventricle pumps blood to the lungs at an average pressure of roughly 10 mmHg, and the left ventricle does ten times that amount of work by pumping blood to the body at an average pressure of roughly 100 mmHg. Since the ventricular pressure is seen in the major arteries throughout the body, we are more aware of our left ventricular pressure (typically 80 to 100 mmHg for a resting adult). More importantly, disease of the left ventricle has the greatest health implications by a large margin, so LVAD development has been the focus of our artificial heart research.

LVAD contributions previous to our own include the Novacor LVAS (Baxter Healthcare) solenoid-activated LVAD, an FDA-approved implantable pulsatile pump with a flexible blood chamber and mechanical valves. More recently, researchers have shown that pump pulsatility is not essential, so compact turbodynamic pumps are being developed (imagine having no pulse!). The Nimbus HeartMate-II (now owned by Thoratec, Inc.) and the Jarvik 2000 are examples of rotary pumps that have reached clinical trials. A critical issue with all of these pumps is the risk of clotting and the release of clot material into the bloodstream in the form of emboli, which can cause stroke or other organ damage. The tendency of these pumps to form clots is reduced with anticoagulant drugs, but these drugs often cause hemorrhaging. Current rotary pumps use blood as a lubricant within hydrodynamic bearings. Consequently, the red blood cells in these pumps may be damaged and, in turn, send biochemical signals that instigate clotting. In simple terms, the patient is between a rock and a hard place. Despite these challenges, a recent clinical study shows a 48% reduction in the risk of death when LVADs are used relative to more conventional medical treatments [4]. The technical issues of clotting, blood damage, mechanical reliability, and the promise of LVAD for treating heart disease is the backdrop against which we developed the Streamliner magnetically levitated LVAD. We used an actively controlled magnetically levitated pump impeller to radically improve reliability and the quality of the blood flow through this axial flow pump.

In the United States, 700,000 deaths per year are attributed to heart disease; 35,000 to 70,000 of these lives can be saved with some form of mechanical circulatory assist.

Design of the Streamliner

The overall Streamliner circulatory assist system is depicted in Fig. 1. The pump draws blood from the apex of the left ventricle and pumps blood, in parallel with the left ventricle, into the aorta. The pump has been the primary focus of our work, since it contacts the blood and is the primary power consumer in the system. Other subsystems include an implanted battery to enable patients to take showers without the external battery pack, a transcutaneous electrical transmission system (TETS) for transmitting power and signals induc-

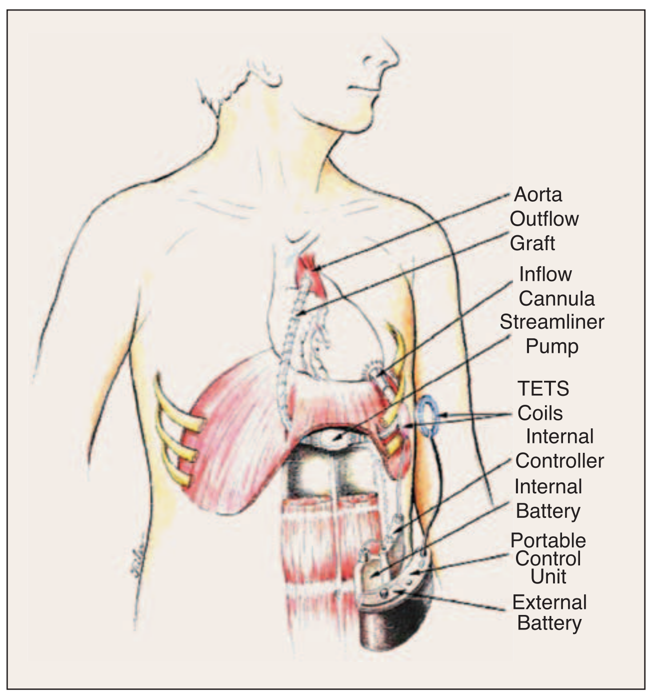

Figure 1. *Configuration of the Streamliner artificial heart system as planned for implantation in humans.*

tively through the skin, an internal controller for levitation and pump speed control, and an external controller and battery for power and system monitoring. Other blood-contacting components are the inflow and outflow cannulae, which serve as artificial blood vessels. At the time of animal trials, the Streamliner system consisted of an internally implanted magnetically levitated pump, cannulae, and an external controller and power source. This was the first fully magnetically levitated system to be tested in animals (Fig. 2).

Figure 2. *The Streamliner artificial heart as used in a 34-day animal trial.*

A cross section of the Streamliner pump, shown in Fig. 3, depicts the compact design that combines the brushless dc motor, mixed-flow impeller, and magnetic levitation components. The small interior volume of the Streamliner reduces the risk of clot formation in several ways. The path through which blood flows has been designed, and optimized by computational fluid dynamics, to be highly "streamlined." The flow is slow enough to prevent damage to the blood cells, yet fast enough to prevent the buildup of clots. By using magnetic levitation, we are able to design levitation gaps that are wide enough to prevent blood damage and simultaneously eliminate stagnation zones; these are unique features of our design approach. In addition, magnetic levitation eliminates mechanical bearing wear. Indeed, the Streamliner technology is expected to revolutionize mechanical cardiac assist technology.

Although deceptively small, the Streamliner is an immense control engineering challenge. The design objective is to magnetically levitate and rotate a pump impeller in the bloodstream while minimizing pump size, blood damage, battery size, and system weight. The selection of a control system concept is the most critical decision, and all aspects of the design, including the actuator and sensor designs, must be optimized. Further, since cables and connectors are weak links in terms of reliability, the number of control channels must be minimized. In response to these requirements, we developed the pump topology shown in Fig. 3 (from our U.S. patent 6,244,835 [7]). The key design elements are a cylindrical magnetically levitated rotating impeller (12), which is supported on permanent magnet radial bearings (9 and 10). The inner races of these bearings are fixed and supported by the outflow hub (18) and the inlet stator blades (20). The axial position of the impeller is actuated by the voice coils (38 and 40) interacting with the outer race magnets of bearing (9). Sensing of the axial position is accomplished with eddy-current sensor probes (26 and 28). The outputs of these sensors are summed to render pitching motions of the impeller unobservable and decoupled from the axial feedback loop. Obviously, the small package requirements for anatomic fit have forced us to use a very integrated design.

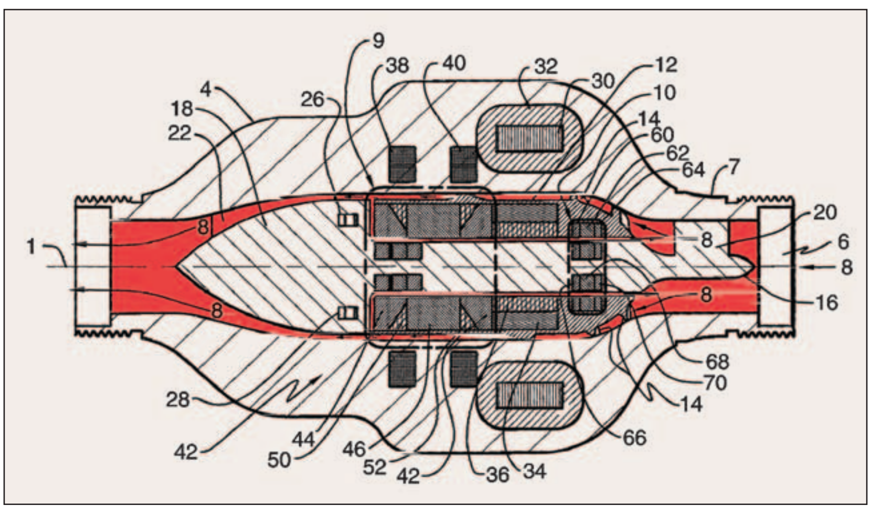

Figure 3. *Section view of the Streamliner artificial heart showing component details: (1) pump axis of symmetry, (4) aft housing, (6) inlet, (7) fore housing, (8) primary blood flow path, (9) aft magnet assemblies (two parts). The inner bearing stack forms the inner race of the permanent magnet bearing. The outer magnets and pole iron stack serve two purposes: a) as the outer race of the aft PM bearing, and b) as the magnet structure for the axial actuator—interacting with coils 38 and 40 to provide thrust force. (10) Fore PM bearing stacks, (12) magnetically levitated impeller (hollow), (14) impeller blades, (16) inlet hub, (18) outflow hub, (20) inlet stator blade and support, (22) outflow stator blade (converts rotational kinetic energy in the blood into pressure), (30) slotless iron laminations of motor stator, (32) toroidal winding on slotless two-pole dc brushless motor stator, (34) two-pole dc brushless motor magnet, (36) motor rotor iron, (38, 40) thrust actuator coils, (26, 28) axial position sensor probes (used for feedback control of the rotor axial position), (42, 44, 46) voice-coil actuator magnets (the aft two magnets also serve as the outer race of the aft PM bearing), (50, 52) iron focusing poles direct field toward coils 38 and 40, (60, 62, 64) outer bearing race magnet, (66, 68, 70) inner race magnet rings.*

Although the rotor is magnetically controlled in six degrees of freedom, only two degrees of freedom are actively controlled: axial and rotational motion. By Earnshaw's theorem [5], the sum of the linear stiffnesses sum to zero, so there is at least one axis along which the position is unstable. By symmetry about the axial axis, we see that the rotational stiffness must be zero. Thus, two active axes are required as a minimum, and in fact all other stiffnesses are positive in the topology shown. Earnshaw's theorem and the symmetry argument just given dictate limits on achievable performance for magnetic levitation in the presence of symmetry, and we have achieved the lower bound on the number of controlled axes.

The Streamliner system has been an exciting challenge from a control engineering standpoint. We have found that sensor and actuator design defines the performance limitations of the overall system. "Zero-placement" has turned out to be an invaluable technique in the active magnetic levitation design, mathematical modeling and optimization are essential, and rendering certain states unobservable to the active controller is required. In addition to actively controlled feedback loops, passive feedback design with permanent magnets has been used. The set of control laws achievable with passive feedback is much less than the field of rational transfer functions, however! In simple terms, the Streamliner is an actuator for the human cardiovascular system that contains internal feedback systems. In the future, external feedback systems responding to the physiologic needs of a patient may further improve biocompatibility.

For the purposes of enumerating alternative topologies for the design process and for patent protection, we developed a design grammar [6]-[8] for our LVAD designs. The phrase "design grammar" is broadly used in the design theory community to describe linguistic approaches to combinatorial design. In our design, we used "sentences" of the following form:

$$\begin{array}{lccccccccc} (\text{FH,AO}) & \text{Sp} & - & \underline{\text{PRB}} & - \text{DCBM} - & \text{ATB} & - & \underline{\text{PRB}} & - & \text{Sp} \\ & \| & & & & & & & & \| \\ & \text{sb} & & - & & \text{ib} & & - & & \text{sb} \end{array}$$

This is an abstract linguistic description of the design of Fig. 3 in which the upper and lower sentences describe the electromechanical and fluid components of the fluid actuation system, respectively. Reading the upper sentence from left to right, we see that the Streamliner has a fixed hub and an axial blood outlet (FH, AO), followed by a support (Sp) and a permanent magnet radial bearing on the interior of the impeller (PRB), where the underline denotes "interior." Next along the flow path is the dc brushless motor toward the exterior of the impeller (DCBM), the active thrust bearing (ATB), the aft permanent magnet radial bearing (PRB), and the aft support (Sp). The lower sentence describes the fluid components of the actuator: stator blades (sb), followed by impeller blades (ib) and then stator blades (sb). The vertical parallel lines connecting "Sp" and "sb" denote hybridization of function: the stator blades also serve as mechanical supports to the inner fixed hub. Variations on the Streamliner are easily communicated with this approach. Interchanging the order of the motor and the active thrust bearings fits in the notation, as does a centrifugal pump with radial outlet (RO). A design grammar is a wonderful notation with which to abstract the topological combinatorics of this control system.

Axial positioning of the impeller is likely the most interesting control problem, and a block diagram is shown in Fig. 4. Working around the loop, the top center block is the rotor axial dynamics determined by the rotor mass and unstable or negative axial stiffness of the radial bearings. By Earnshaw's theorem, the axial stiffness is minus two times the combined radial stiffness of the permanent magnet bearings. The axial motion of the impeller is measured with the two eddy-current sensors, which we designed to have very low noise. The sensor coils are driven with a low-noise crystal oscillator current source, and impedance changes are detected with a coherent demodulator. The input-referred current noise is surprisingly small and roughly 1 nm/root Hz over the controller bandwidth. The sensor output drives a modified PID controller designed by "zero placement." The controller has a zero at zero so that there is no dc gain; if steady state is reached, the voice coil theoretically dissipates no power at all! (In practice, the coils dissipate less than half a watt of power, negligible compared to the 10 W of power to drive the motor.) The controller drives the voice-coil actuator through a current amplifier that in turn applies a force to the impeller. The objective of the controller is to levitate and not to precisely center the rotor axially; hence, there is freedom to let the rotor move axially to save power. As the controller has a zero at zero, it is easy to see that the transfer function from disturbance to rotor displacement at dc is

$$H_{\text{dist-disp}}(0) = -1 / K_b.$$

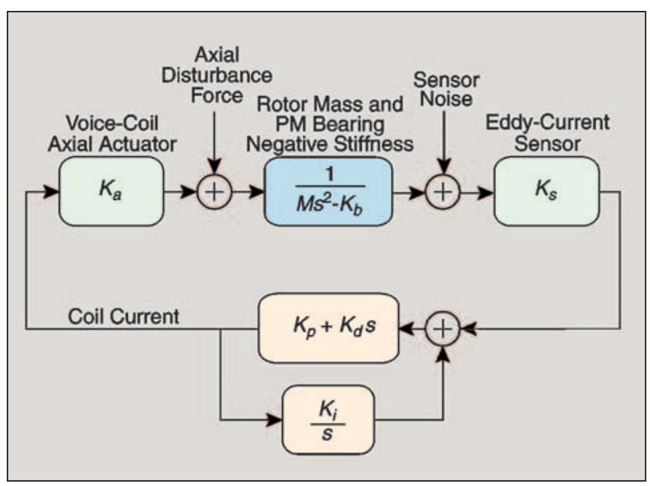

Figure 4. *Axial controller block diagram for Streamliner artificial heart. The transfer function from disturbance to control current has a zero at dc. This is known as a virtual zero power (VZP) controller in the magnetic bearing community.*

When a disturbance force is applied to the rotor, the rotor moves *toward* the disturbance. The more unstable the bearings are in the axial direction, the less motion is required to balance the disturbance force. Instability can be exploited with control! When standing in a windstorm, we use our body (an inverted pendulum) in the same way. All of this supposes that controller gains can be chosen to stabilize the system. In fact, this can be done with a simple pole placement approach, with final tuning accomplished experimentally.

The zero in the controller minimizes power consumption and voice-coil heating and is commonly referred to as virtual zero power (VZP) control. Small voice-coil actuators are quite inefficient, and initial levitation power for this system is on the order of 100 W. Thankfully, levitation only takes a few milliseconds. After levitation, the coil power typically drops to about 0.5 W during pumping. A standard figure of merit for voice-coil actuators is the force per root watt, which in our design is required to be high. We designed the actuator to meet these needs by choosing a geometry that was compatible with the pump geometry and then parametrizing the design for optimization in a magnetic finite element analysis code (ANSYS). Very large improvements in power consumption were accomplished in this way relative to what can be gained with an ad hoc actuator design and optimized controller design alone. The optimization process dictated the location and shape of the iron poles in the outer race of the aft radial bearing. Unusual constraints were placed on the voice-coil design with regard to radial magnetic force. Since any iron in the voice-coil support produces negative stiffness, iron was eliminated from the voice-coil actuator stator design. Similarly, a slotless motor stator was used in the motor to move the motor stator iron as far away from the magnetic impeller as possible.

The motor was modeled directly from Maxwell's equations in 2-D. Closed-form expressions relating the motor geometry and coil current density to power and efficiency were derived and used in the system optimization. This is quite unlike the more common shopping activity used to select a motor in motion control problems. The resulting design has a computed efficiency of 85% at 7,500 rpm, even with the limits on stator iron. Since wires and connectors tend to reduce reliability numbers, we chose to use sensorless motor control. The motor is also a generator, so that the zero crossings from the coil back-EMF voltages can be used to measure rotor angle and, hence, control commutation. A sterling achievement in motor control technology is embodied in the single-chip sensorless motor controllers developed for the disk-drive industry. The tuning of such controllers is time consuming, but we were successful in incorporating such a device into our system with great savings in complexity.

The design of the passive radial bearings was accomplished with new closed-form models for magnetic stiffness. In this case, we are stabilizing the rotor radial position with one hand tied behind our backs in the sense that we only have direct control over stiffness and not damping. Moreover, passive magnetic bearings have essentially no damping so that the blood surrounding the rotor supplies all damping. Significant radial forces are produced by the rotor due to residual imbalance after dynamic balancing and difficult-to-model periodic fluid instabilities within the flow. The rotor of the Streamliner prototype demonstrated two natural frequencies associated with rotor bouncing and pitching motions at 4,000 and 8,000 rpm, respectively. The pump was designed to operate between these two speeds to avoid touchdown of the impeller blades. This design was accomplished with a simple rigid-body model of the impeller and our closed-form expressions for bearing stiffness.

Future of the Streamliner Artificial Heart

We are now working with MedQuest® Products Inc. in Salt Lake City to commercialize a magnetically levitated LVAD. MedQuest has licensed the Streamliner patents from the University of Pittsburgh, and jointly with MedQuest we have developed new technology for a centrifugal pump that promises to be even more efficient than the current axial flow pump. Coming from academic backgrounds, it is rewarding to see control technology making its way to such an important market. Many exciting control issues remain unsolved, and there is much more to contribute in the way of control design methodology. Work is needed in controlling the pump speed in response to physiologic need and in the design of physiologic sensors, and further power savings should be possible.

Industrial Process Monitoring by Michael Piovoso

Advances in automation and distributed control make possible the collection of large quantities of operational data. However, without adequate tools, such data may not be useful to improve understanding or operation of a process. The management of every modern industrial site believes that its data bank would be a gold mine of information, if only the *important* and *relevant* information could be extracted painlessly and effortlessly. Timely interpretation of such data would likely improve quality and safety, reduce waste, and improve business profits. Ideally, this should be done in real time. In addition, analysis of historical databases might provide new insights into and understanding of complex chemical processes that could then be translated into improved operation.

Traditional methods typically use first-principles models to capture the essence of a process. These models may then be used to monitor a process and to detect potential problems such as sensor or controller faults and unusual disturbances. The problem with this approach is that developing such models is both time consuming and costly. Even when completed, many parameters must still be estimated from process data. Such models are useful for process optimization, but they do not necessarily make good models for *monitoring* because they often do not compute information such as the final product quality. Furthermore, many process variables that are important to the operation are not included in the process model. Examples of such variables include power consumption and speeds of motors or agitators that may be critical to the operation but are not necessarily needed for the process model.

A data-driven approach based on multivariate statistical analysis offers the most promise for developing methods for process monitoring.

Because of these limitations, a data-driven approach based on multivariate statistical analysis offers the most promise for developing methods for process monitoring. In this approach, historical data that correspond to *normal* process operations are used to form an initial statistical model. Subsequent operations are then compared to this model to determine if the process is within normal operating conditions. If not, operators are provided with information as to which process variables are behaving differently than expected.

Multivariate Statistics for Process Monitoring

Historical databases used to gather initial modeling information are typically very large, with hundreds of process variables being sampled as often as every few seconds to as infrequently as once a minute, or even longer. The data are highly correlated, have widely varying signal-to-noise ratios, exhibit gross errors due to dropouts or human error, and occasionally are completely missing. Much work is needed to clean the data before the modeling operation itself. Subspace modeling techniques such as principal component analysis (PCA) and projection to latent structures, also known as partial least squares (PLS), have provided a means of handling such data problems.

PCA is an effective technique for modeling data that are highly correlated [11]. Such data occur naturally in most systems since there are typically many more sensors than phenomena, so that the sensor measurements made on various aspects of the phenomena will be correlated. PCA [also known as singular-value decomposition (SVD) and Karhunen-Loeve transformation] can decompose such data into a smaller set of numbers that captures the vast majority of the variability. Conceptually, since the data are not uncorrelated, if one were to plot them in a high-dimensional space, more data could be seen along certain directions than along others. PCA therefore defines the direction along which most of the data points are closest. These form the first eigenvector, or the first *loading* of the data. Each data point can then be represented by how far it is along this eigenvector. These distances are referred to as the *score* of the data along the first principal component. From the loading and score, an approximation of the original data is produced and subtracted from the original data. This forms a residual set from which another principal component is extracted. The analysis continues until a sufficient number of such components are extracted. Several techniques exist for terminating the algorithm [12]-[15].

There are several reasons for treating data in a multivariate way. The PCA loadings define linear combinations by taking a weighted average of the existing data. This linear combination of the data can have less noise than the original data due to the averaging effect. Furthermore, the linear combination can be sensitive to abnormalities that may not be apparent in a univariate analysis.

Once a model has been obtained with the correct number of principal components, it can then be used to *monitor* operation. If there are abnormal data, either the projection in the model space or the residuals of the model should contain that information. Conceptually, the system would check new data for known faults, check the projections onto the model space, and check the residuals. Fig. 5 illustrates these concepts in which three-dimensional data are being described by two directions or principal components. The *normal* data, as illustrated by green points, are consistent and fit the model description well, whereas the red points illustrate two types of *abnormal* situations. The red point that is "in the model plane" is well fit by the plane but is outside the normal region; this abnormality can be discovered by the T^2 statistic. The other red point "projects into the model plane," but the residual or the distance from the plane is abnormally large; the Q statistic is sensitive to unusually large residuals.

Using the Q and T^2 statistics, a multivariable statistical process control scheme can be implemented. This is inherently superior to univariate methods in that the interaction among variables is explicitly described. These two statistical measures are monitored to detect out-of-limit conditions. This can all be done in real time and will provide crucial information to those needing to make mission-critical decisions. Not only can the system detect the presence of an abnormality, but also information as to the source and cause is discerned using contribution plots, which identify which variables contribute to the Q and T^2 statistics. Contribution plots for the Q statistic are easily compiled from the difference between the reconstructed data and the actual data. Plotting the values of the individual components that form the sum to generate Q thereby forms a contribution plot. Although the calculation for the T^2 statistic is a little less direct, the outcome is the same: to identify the respective contributions of each component and to thereby identify the cause of the abnormality.

We have successfully applied these techniques to two types of problems: continuous and batch processes [9], [10]. The primary difference between the two is that batch processes have specific beginning and ending times; whereas continuous data are two dimensional (variable over time), batch data are three dimensional (variables over time over batches). An extension of principal component analysis known as the multiway principal component analysis (MPCA) has proven useful for handing the three-dimensional data. Our application [9] is, to the best of my knowledge, the first commercial application of multivariate statistics for online process monitoring.

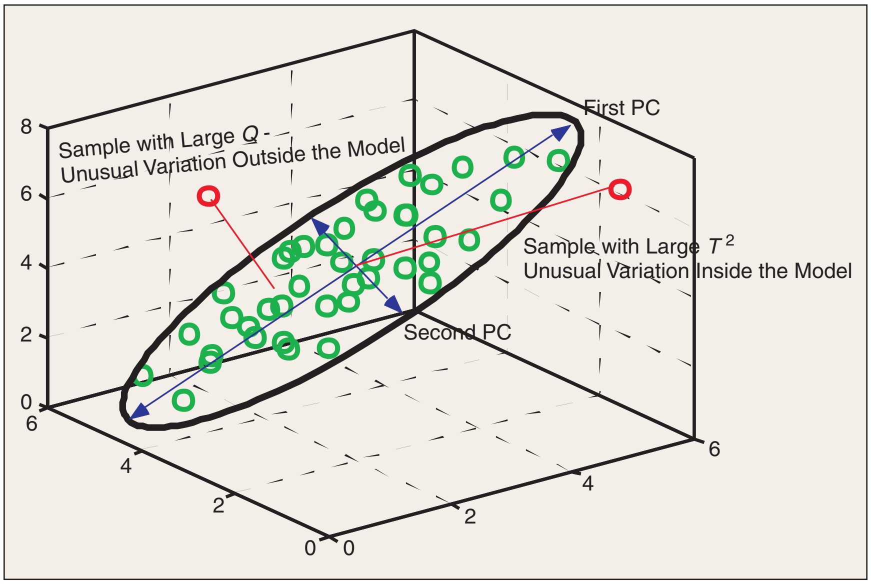

Figure 5. *An example of three-dimensional process data analysis with principal component analysis (PCA).*

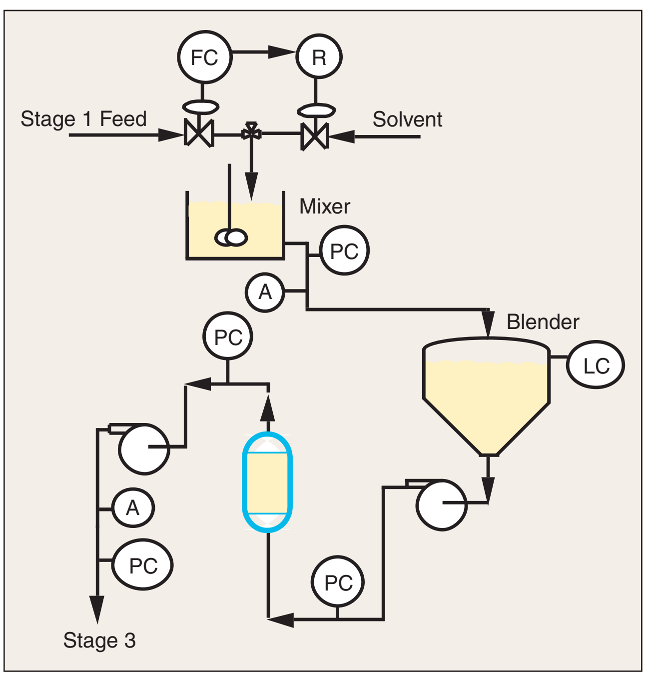

Figure 6. *A second stage for a multistage continuous chemical process.*

Continuous Process

Consider a multistage continuous chemical process. The first two stages where chemical reactions occur have the greatest impact on operations of the downstream stage. Critical properties such as viscosity and density, if altered slightly, will affect the final product, resulting in a loss of revenue or operational difficulties. When yield losses occur, it is often difficult to determine which stage is responsible. A lack of online sensors, a common situation in the chemical industry, makes it impossible to relate any specific changes to a particular stage. As a result, properties are measured in an offline laboratory, and these measurements are sometimes taken on as small a sampling frequency as once every eight hours. Such infrequent measurements make control of the product quality very difficult. At best, the operators may learn a set of heuristics that, if adhered to, usually produce an acceptable product. However, unforeseen disturbances and undetected machine degradations may occur that result in yield loss. There are periods of operation, furthermore, when the final process step produces a degraded product despite near-perfect upstream operations. What also contributes to the control problem is a lack of just-in-time knowledge about the current state of the process.

Fig. 6 illustrates the major details of the second stage of a continuous multistage process. It begins with the feed from the first stage, combined with a solvent in a mixer at a carefully controlled speed and temperature. A sample of the mixture is sent to a blender under level and speed control. The mixture is then pushed through a series of pumps and filters before reaching the third stage in the process. The pumps and filters need to be replaced frequently; thus, they are installed in pairs. The load on one can be temporarily increased while the other is being serviced. Likewise, the filters must be changed routinely to avoid plugging and downtime.

Frequent maintenance of the equipment is not the primary source of control problems. Rather, the majority of the abrupt control changes occur due to the demands of the third stage. Whenever there is a decrease in demand, the second stage must reduce throughput because the third stage has very limited storage capacity. In the current control scheme, a change in the third-stage demand is indirectly coupled to the feed flow-rate control valves. When the demand increases, the second stage must ramp up to meet the demand and do so quickly. This causes the process to move around significantly, and it never reaches equilibrium. Clearly, demand is the dominant effect on variability in the sensor values and process performance.

Batch Process

The same concepts used for continuous processes can also be applied to a batch process. Batch and semibatch processes play an important role in the chemical industry, mainly because of their ability to flexibly produce low-volume, high-value products. Examples include reactors, crystallization, distillation, injection-molding processes, and the manufacture of polymers. Batch processing typically involves charging a vessel, processing under controlled conditions, and finally discharging the product. Successful operation means tracking a prescribed recipe and the process variable trajectories with a high degree of reproducibility from batch to batch. Temperature and pressure profiles are implemented with servo-controllers, and tools such as programmable logic controllers ensure precise sequencing operations.

Unfortunately, the main characteristics of batch processes (flexibility, finite duration, and nonlinear behavior) make process control difficult. Control problems are further complicated by a lack of sufficient online instrumentation. Typically, statistical quality control is used to adjust set points for feedback controllers, the length of operations, the duration of heating or cooling, and the duration of the specific stages of the batch.

We have demonstrated that MPCA can be used to improve process understanding and thereby improve the control scheme for the process. We have shown that it can be used effectively to identify the major sources of variability in data taken from an industrial batch process and that these variations are directly related to product quality. We made recommendations that reduced the variability, thereby achieving production of a uniform high-quality product. The economic stake for achieving this is large, with further ramifications in plant operations such as: 1) lower energy costs, 2) lower raw material costs, 3) reduced time to transition between different products, 4) reduced offline product testing, and 5) reduced downtime.

We controlled a reactor in which an aqueous solution is first boiled in an evaporator until the water content is reduced to approximately 20% by weight. The evaporator's contents are then discharged into a reactor where 10-20 lb of polymer residue may be present from processing of previous batches. This batch reactor is operated according to a combination of prespecified reactor and heat source profiles and timed stages. Fig. 7 illustrates the reactor and the various sensor signals that are available for analysis.

The system, first installed in 1990, was composed of three components: the user interface, the statistical engine, and an expert system rule base to help operators decipher the plots. A special computer manufactured by Advanced Computer Applications, Inc. (ACA), of Newtown, PA, handled the user interface, the expert system rule base, and the task-to-task and device-to-device communications. A separate Digital Equipment Corporation VAX computer and an IBM PC computed the statistical analysis.

Every hour, process data taken from the historian are displayed on the ACA computer. After some validation, the data are sent to the PC for statistical analysis. The analysis is based on an offline model used to characterize the process data taken during periods of acceptable operation. The resulting model characterizes the *sweet spot* of the process. The concept is to test the operation at each point in time and determine if the data are consistent with that taken during acceptable operation. For each new sample, the Q and T^2 statistics are computed. The results are visually displayed to the operators on the ACA platform. If the data are outside of the 95% confidence limits in one or both of the statistics, the operator can ask for a contribution plot, which defines how each of the process variables contributes to the values of Q or T^2.

Fig. 8 is a typical plot of scores associated with the first and second principal components. The circles represent data used to calibrate the model. The "+"s display operational data samples, with the arrows indicating increasing time. The point marked "A" is in fact a precursor to pump failure. The pump failed shortly thereafter, at point "B." After it was repaired, the process moved back to the region of acceptable operation.

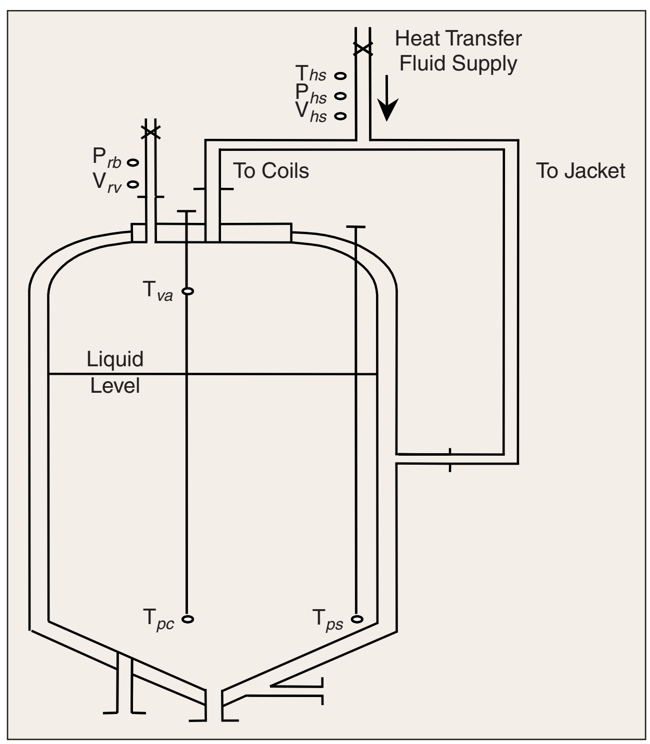

Figure 7. *Example of a batch chemical process reactor.*

In the batch process, we were able to show that multivariate statistical methods, besides being useful for process monitoring, can help validate and improve one's understanding of the process. In this case, we suggested that the control scheme for the process being monitored should be modified to use less pressure and more temperature control within the autoclave. DuPont of Canada implemented these ideas and installed process monitoring software that uses both first-principles and multivariate statistical methods. Improvements of 40% in yield loss and product quality have been estimated.

Future Industrial Process Monitoring Developments

Univariate statistical process control has had a major impact on the discrete parts manufacturing industry. Unfortunately, chemical processes are more complex, and monitoring individual process variables is not fruitful. Multivariate methods hold the promise of monitoring the complex behavior in a fashion that will improve quality, productivity, and safety. In the coming years, this technology will be found in more and more processes. The requirements for production, quality, and safety will drive more and more companies to adapt these concepts. In the coming decades, multivariate statistical quality and process control will become as commonplace as the univariate statistical process control used in many discrete parts manufacturing plants.

This technology is still in its infancy. To achieve wider market acceptance, more robustness must be built into the system. Chemical processes are not static: sensors are moved, modifications are made to existing processes, and new products are frequently introduced. When this happens new models are needed, and model maintenance requirements are significant. To gain wider acceptance, the technology needs to have about the same level of maintenance as does a control system.

Despite the need for further development, some small companies already have products on the market. Anex6 Ltd. (Belfast, U.K.) offers a product that performs multivariate statistical process control for manufacturing industries. This product has been installed in DuPont's manufacturing facility at Maydown, Northern Ireland. Umetrics (based in Sweden, with offices in the United Kingdom and the United States) has a product called SIMCA-P 9.0; although this package is mainly geared toward multivariate data analysis, it has also been used for online process monitoring.

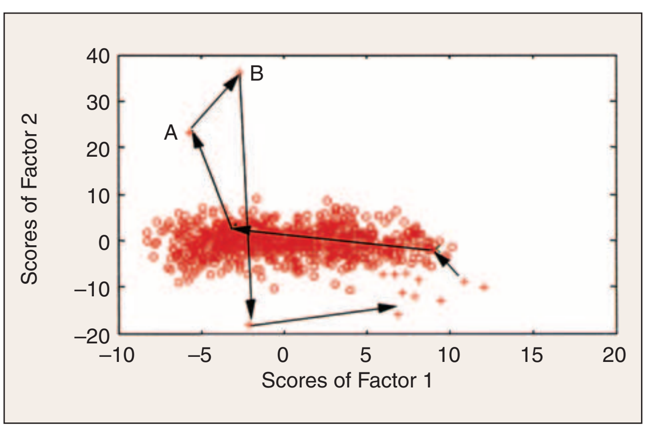

Figure 8. *Performance monitoring for a batch process.*

Multivariable Control for Air Vehicles by Siva S. Banda

Modern military aircraft use fly-by-wire flight control systems to enhance the pilot's ability to control the vehicle. Fly-by-wire systems were first developed in the 1960s and were initially used to alleviate the pilot's workload. Today, practically all modern military aircraft use fly-by-wire-enabled feedback control systems to modify the vehicle dynamics and to provide some level of autonomy.

The process of synthesizing flight control systems using classical single-loop methods requires that aircraft dynamics be linearized at numerous points in the flight envelope. As a result, the feedback gains had to be *scheduled* according to flight condition. Gain scheduling requires that large numbers of flight control gains be precomputed and stored in the aircraft flight computer. Multivariable flight control design methods that have been evolving since the late 1960s promised to alleviate some of the tuning inherent in the classical design process, however. Some of the more popular design techniques, such as those based on the linear quadratic regulator and H-infinity, were optimization-based and were capable of handing multiaxis coupling. Unfortunately, these methods still required a linearized vehicle model as a starting point for the design, and since they were optimization based, deciding exactly *what* should be optimized required considerable engineering judgment. Generally, the ultimate objective was to optimize the handling qualities of the closed-loop aircraft. But what constitutes good handling qualities is largely subjective in nature and not easily captured in a mathematical performance index. Furthermore, flight control designs based on these methods still required that numerous gain sets be computed for the various flight conditions an aircraft is expected to encounter. Nonetheless, one of the major early contributions of multivariable control theory was that it gave designers the ability to analyze the robustness of multiloop control systems to parameter variations in the linearized models.

Flight Control with Feedback Linearization

Over the past three decades, a new control design methodology evolved that came to be known as *feedback linearization* [16], [17]. A special case of the general feedback linearization methodology, called dynamic inversion, is particularly well suited to flight control applications [18]. Dynamic inversion is a control design technique that eliminates the requirement of a linearized model of the vehicle dynamics prior to beginning the control synthesis process. Therefore, control systems designed using this technique can be written directly in terms of the nonlinear equations of motion for the vehicle.

Stability and control information is also required to estimate aerodynamic forces and moments at flight conditions that will be encountered by the vehicle. For aircraft, it is convenient to use this control methodology in conjunction with a model-following scheme because aircraft with desirable handling qualities can be written in terms of simplified dynamic models. When dynamic inversion is used in a model-following framework, the open-loop aircraft dynamics are essentially replaced with the dynamics of a model that is known to have good flying qualities. The technique is used to provide the pilot with the ability to easily command a set of controlled variables. The body-axis angular rates are often chosen as the controlled variables, although they are sometimes blended with aerodynamic angles such as *angle-of-attack* and *sideslip* to avoid problems with unstable zero dynamics [19].

A major advantage of this type of control design is that it can use information from adaptive and reconfigurable control modules to provide tolerance to control effector failures or vehicle damage. The block diagram shown in Fig. 9 is helpful for understanding the interactions between the elements of a dynamic-inversion-based adaptive/reconfigurable control system.

When dynamic inversion is applied to flight control synthesis, it is typically used to decouple—and indeed to cancel out—the rotational aircraft dynamics. These modified dynamics become a bank of decoupled integrators or a rate command system from the perspective of the vector of controlled variables **C**. The input to this modified controlled element is the command variable rates $\dot{\mathbf{C}}_{\text{Desired}}$ (e.g., angular acceleration commands). The model-following approach used to produce $\dot{\mathbf{C}}_{\text{Desired}}$ usually consists of a prefilter that, when combined with the modified controlled element, produces a closed-loop dynamic system that is known to have good flying qualities. Dynamic inversion control laws normally generate a small number of control variable rate commands that must be generated by the available control effector suite $\dot{\mathbf{C}}_{\delta_{\text{Desired}}}$. These pseudo- commands are often expressed in terms of the total desired aerodynamic moment or angular acceleration. Normally, there are more control effectors than there are controlled variables or axes to control. This condition of control effector redundancy gives rise to the control allocation problem.

Multivariable control theory allows flight control designers to analyze the robustness of multiloop control systems to parameter variations in the linearized models.

Control allocation or control mixing can be used to generate any number of control effector commands δ from a small number of pseudocommands $\dot{\mathbf{C}}_{\delta_{\text{Desired}}}$. In the simplest approaches, a control allocator does little more than gang control surfaces together to eliminate the condition of control redundancy that leads to an overdetermined problem. Much more sophisticated control allocation algorithms have evolved over the past decade [20]. Some of the most effective control allocators are based on constrained optimization methods that deliver the desired pseudocommands while optimizing some subobjective such as drag minimization or wing load alleviation. These algorithms generate commands that respect actuator rate and position limits. In other words, these control allocators generate actuator commands that deliver the desired control-induced angular acceleration or moment, so long as it does not violate actuator rate or position limits. When it is not physically possible

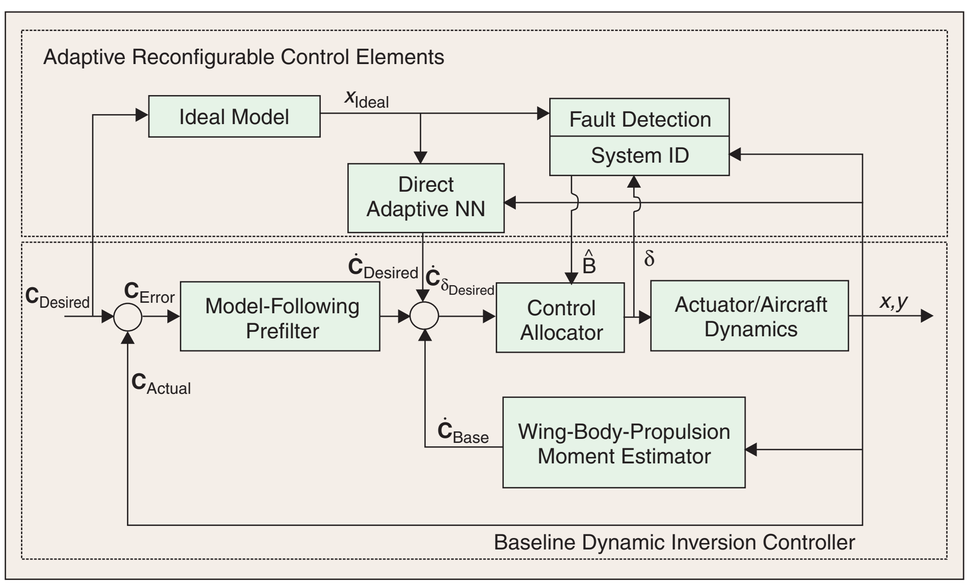

Figure 9. *Dynamic-inversion-based adaptive/reconfigurable control system.*

to deliver the desired moments or accelerations, the optimization-based approaches can minimize the difference between the actual and desired moments or accelerations. The control allocation problem arises quite naturally from the dynamic inversion formulation, but the algorithms may also be used to simplify the synthesis of classical or multivariable control systems that do not generate physically unrealizable actuator commands. Another enormous advantage of online control allocation is the ability to use direct or indirect control effector failure identification to update the information that the control allocator uses to generate actuator commands. Identification of failures or damage can be used to reconfigure the remaining control effectors in order to adapt to the failure. For example, if an elevon fails hard-over on an aircraft, the remaining elevon, ailerons, and rudders may be used to balance the moments on the vehicle and allow the pilot to maintain control of the vehicle. The control allocator can even operate in a mode that takes advantage of a vehicle's control redundancy to enable indirect detection of failures or damage. To detect failures or damage using online system identification, the control surface deflections must be decorrelated and excited. Under nominal conditions, control allocators can make use of control redundancy to achieve decorrelation and excitation without degrading the desired response of the vehicle [21].

Figure 10. *The VISTA F-16 Variable Stability Aircraft demonstrated the ability to safely land with a malfunctioning horizontal tail. (Photo courtesy of U.S. Air Force.)*

Figure 11. *The Boeing X-36 Tailless Fighter Aircraft can recover nominal flying even with control effector or vehicle damage. (Photo courtesy of NASA.)*

The upper portion of Fig. 9 contains several adaptive/reconfigurable control elements that can be used to augment the baseline dynamic inversion control law. The adaptive reconfigurable control elements can be incorporated as a module that enables the vehicle to recover nominal performance to the greatest extent possible if it is damaged or experiences control effector failures. An explicit model of the ideal closed-loop vehicle is normally included so that deviations in the response of the actual vehicle can trigger fault detection logic that determines whether or not failures or damage have occurred. If a fault occurs, the online control allocator can change subobjectives to decorrelate and excite the control effectors. This excitation must provide sufficient signal content to enable the online system identification algorithms to estimate the stability and control characteristics of the vehicle using a blend of prior knowledge, sensor measurements, and control surface deflections. The explicit model can also be used in conjuction with a direct adaptive controller that modifies the pseudocommand when the ideal model and actual vehicle responses differ. The direct adaptive controller and the indirect adaptive controller that relies on online system identification can work together to recover the nominal vehicle performance as closely as possible, given the physical limitations of the vehicle.

The Air Force Research Laboratory's Control Science Center of Excellence (COE) and its predecessors have been instrumental in the development of fly-by-wire flight control technology since the 1960s. Some of the more recent accomplishments made by the COE and its industry partners are the flight demonstration of adaptive/reconfigurable control technology on the VISTA F-16 and the X-36 aircraft shown in Figs. 10 and 11, respectively. A self-designing control (SDC) system was demonstrated on the variable-stability VISTA F-16 in 1995 [22], [23]. The SDC uses online system identification to feed a receding horizon optimal control design method that modifies flight control system gains in response to simulated effector failures. The SDC program resulted in the first-ever landing of a fighter with an emulated *missing* horizontal tail under full reconfigurable control.

The X-36 tailless fighter aircraft used a dynamic inversion control system with a control allocator as a baseline flight control system [24]. Neural-network-based direct adaptive control was successfully used to augment this baseline flight control system. The objective of this direct adaptive method is to recover nominal flying qualities to the greatest extent possible in the presence of control effector failure or vehicle damage. The system was successfully flight demonstrated and underwent extensive hardware-in-the-loop ground testing. The adaptive flight control systems were found to dramatically improve the flying qualities of the vehicle in the presence of failures or damage.

Figure 12. *An integrated adaptive guidance and control system is being evaluated for the X-40A Space Maneuvering Vehicle. (Photo courtesy of NASA.)*

Future Developments in Multivariable Control for Air Vehicles

The success of the VISTA and X-36 flight tests obviously generated increased confidence in adaptive/reconfigurable control technology. Future aircraft and weapon systems will likely benefit from the experience gained in these research and development programs. Legacy munitions whose stability and control characteristics are highly uncertain may also benefit from adaptive control systems.

Our Control Science COE is now working to transition the control technologies developed for fighter aircraft to tomorrow's reusable space-access vehicles. Recent work focuses on developing full-envelope, trajectory-independent flight control systems for reusable launch vehicles. The COE also participated in the NASA Marshall Space Flight Center's advanced guidance and control project for the X-33, where it developed a highly promising adaptive/reconfigurable flight control system for the ascent flight phase [25]. Current research is focusing on the development of integrated adaptive guidance and control for autonomous space-access vehicles. It has been found that even with adaptive/reconfigurable control, the closed-inner-loop performance of autonomous aircraft can degrade in the presence of vehicle failures or damage. Since there is no pilot onboard to adjust the reference trajectory and guidance commands, the autonomous system must perform these tasks online.

The COE and its industry partners are also preparing to flight demonstrate a promising integrated adaptive guidance and control (IAG&C) algorithm using the VISTA F-16 to simulate the X-40A Space Maneuvering Vehicle shown in Fig. 12. The X-40A is an 80% scale version of the X-37 that has been used to flight demonstrate an autonomous landing system for reentry vehicles. The successful demonstration of IAG&C technology on the VISTA F-16 may lead to actual X-40A drop-test demonstrations of the technology in the next few years.

Summary/Conclusion

This review of the three most recent winning applications for the IEEE Control Systems Technology Award has clearly illustrated the diverse nature of control engineering. Our technology is used within environments ranging from the inside of the human body to the inside of large, complex industrial plants, and even to outer space. The future is bright, as we expect devices like the Streamliner artificial heart to benefit thousands of people with heart disease, multivariable statistical process control to improve industrial operations, and modern control laws to enable vastly improved performance for complex air vehicles.

References

Artificial Hearts

[1] J.R. Hogness and M. van Antwerp, Eds. *The Artificial Heart: Prototypes, Policies, and Patients,* Institute of Health. Washington, DC: National Academy, 1991.
[2] W.J. Kolff, "The invention of the artificial heart," *Int. J. Artific. Organs,* no. 7, pp. 386-403, 1990.
[3] E. Marshall, "Artificial heart: The beat goes on," *Science,* vol. 253, pp. 500-502, 1991.
[4] E.A. Rose, et al., "Long-term use of a left ventricular assist device for end-stage heart failure," *New Eng. J. Med.,* vol. 345, no. 20, 2001.
[5] S. Earnshaw, "On the nature of molecular forces which regulate the constitution of the luminiferous ether," *Cambridge Philosoph. Trans.,* vol. 7, pp. 97-112, 1842.
[6] J.F. Antaki, B. Paden, G. Burgreen, and N. Groom, "Magnetically suspended miniature fluid pump and method of designing the same," U.S. Patent 6 015 272, Jan. 18, 2000.
[7] J.F. Antaki, B. Paden, G. Burgreen, and N. Groom, "Blood pump having a magnetically suspended rotor," U.S. Patent 6 244 835, June 12, 2001.
[8] B. Paden, J.F. Antaki, and N. Groom, "Animal trials of a magnetically levitated left-ventricular assist device," in *Proc. 5*th *Int. Symp. Magn. Suspension Technol.,* Santa Barbara, CA, Dec. 1999, pp. 251-258.

Industrial Process Monitoring

[9] M.J. Piovoso, K.A. Kosanovich, and J.P. Yuk, "Process data chemometrics," *IEEE Trans. Instrum. Meas.*, vol. 44, pp. 262-269, 1992.
[10] K.A. Kosanovich, K.S. Dahl, and M.J. Piovoso, "Improved process understanding using multiway principal component analysis," *Ind. Eng. Chem. Res.*, vol. 35, p. 138, 1996.
[11] J.E. Jackson, *A User's Guide to Principal Components*. New York: Wiley-Interscience, 1991.
[12] R.B. Catell, "The Scree test for the number of factors," *Multivariate Behav. Res.*, vol. 1, pp. 245-276, Apr. 1966.
[13] S. Wold, "Cross-validatory estimation of the number of components in factor and principal component analysis," *Technometrics*, vol. 20, no. 4, pp. 397-406, Nov. 1978.
[14] W.R. Zwick and W.F. Velicer, "Comparison of five rules for determining the number of components to retain," *Psych. Bul.*, vol. 99, no. 3, pp. 432-442, 1986.
[15] S.J. Qin and R. Dunia, "Determining the number of principal components for best reconstruction," in *Proc. IFAC Dycops'98*, Greece, June 1998.

Multivariable Control for Air Vehicles

[16] A.J. Krener, "On the equivalence of control systems and linearizations of nonlinear systems," *SIAM J. Control Optim.,* vol. 11, pp. 670-676, 1973.
[17] A. Isidori, "Nonlinear control systems: An introduction," in *Lecture Notes in Control Information Sciences,* vol. 72. New York: Springer-Verlag, 1985.
[18] D. Enns, D. Bugajski, R. Hendrick, and G. Stein, "Dynamic inversion: An evolving methodology for flight control design," *Int. J. Control,* vol. 59, no. 1, 1994.
[19] "Application of multivariable control theory to aircraft control laws," prepared by Honeywell Technology Center and Lockheed Martin Skunk Works, WL-TR-96-3099, Wright-Patterson AFB, OH, 1996.
[20] A.B. Page and M.L. Steinberg, "A closed-loop comparison of control allocation methods," in *Proc. 2000 AIAA Guidance, Navigation, Control Conf.*, AIAA 2000-4538, Aug. 2000.

[21] J.M. Buffington, "Modular control law design for the innovative control effectors (ICE) tailless fighter aircraft configuration 101-3," AFRL-VA-WP-TR-1999-3057, Wright Patterson AFB, OH, 1999.
[22] S.S. Banda, Ed., "A special issue on reconfigurable flight control," *Int. J. Robust Nonlinear Control,* vol. 9, no. 14, Dec. 1999.
[23] P.R. Chandler, M. Pachter, and M. Mears, "System identification for adaptive and reconfigurable control," *AIAA J. Guidance, Control, and Dynam.,* vol. 18, no. 3, pp. 516-524, 1995.
[24] J.S. Brinker and K.A. Wise, "Nonlinear simulation analysis of a tailless advanced fighter aircraft with reconfigurable flight control law," *AIAA J. Guidance, Control, and Dynamics,* vol. 24, no. 5, pp. 275-284, 2001.
[25] D.B. Doman and A.D. Ngo, "Dynamic inversion-based adaptive/reconfigurable control of the X-33 on ascent," *AIAA J. Guidance, Control, and Dynam.,* vol. 25, no. 2, pp. 275-284, 2002.

James Antaki is an associate professor in the department of surgery, University of Pittsburgh. He also has adjunct appointments in the departments of mechanical engineering and bioengineering, as well as an appointment at Carnegie Mellon University. He graduated from Rensselaer Polytechnic Institute in 1985 with a degree in mechanical engineering and minor in electrical engineering. His Ph.D. research at the University of Pittsburgh concerned a study of biomechanics of the human heart and led to the invention of a unique technique for tracking the motion of the heart with MRI imaging. In 1991, he was appointed Director of Research for the Artificial Heart Program, where, in addition to conducting basic and applied research, he led the development of the Pitt artificial heart, known as the Streamliner. He holds seven patents related to artificial organs, three related to the harmonica, and has an additional five patents pending. He has authored or coauthored more than 145 abstracts, refereed papers, and other scientific publications. The central theme of his research in bioengineering is the introduction of rigorous engineering principles into the practice of medical device design. He is one of the first investigators to introduce the concept of numerical optimization to this field. He has dedicated his career to bridging the gap between basic sciences, engineering, and medical practice, with the hope of developing improved therapies, instruments, and devices for the treatment of disease.

Brad E. Paden received his Ph.D. in electrical engineering from U.C.-Berkeley in 1985 and is currently a professor at the University of California-Santa Barbara Department of Mechanical and Environmental Engineering. His research interests focus on nonlinear control theory and its application to electromechanical systems. In 1988, he was a visiting fellow at the Department of Mathematics, University of Western Australia. He was the recipient of the Best Paper Award from the *ASME Journal of Dynamic Systems, Measurement, and Control* in 1993 and the IEEE Control Systems Society Technology Award in 2001. He has consulted for industry on numerous projects, including artificial heart design, maglev vehicle design, and energy storage flywheels, and has served as an associate editor for the *Journal of Robotic Systems.* In 1992, he co-founded Magnetic Moments, LLC, a company specializing in magnetic bearing and electromechanical system design. He directs RoboChallenge, a robotics-based K-12 outreach program involving more than 100 students in mobile robot design.

Michael J. Piovoso is an associate professor of electrical engineering at Penn State University School of Graduate Professional Studies. He received his Ph.D. in electrical engineering in 1969 from the University of Delaware before joining the E.I. DuPont de Nemours & Co. Central Research and Development. Later in 1969, he served as an applied mathematician in the U.S. Army at the Advanced Research Laboratory, returning to the DuPont Company in 1972. While at DuPont, he worked on inspection systems, designing specialized instrumentation, applications of artificial intelligence, and process control. He holds four patents and has two patents pending. After more than 32 years with the DuPont Company, he retired and joined Penn State University. His current research interests are in multivariate methods for system and controller monitoring, application of neural networks to process control, and control of distributed-parameter systems. Dr. Piovoso has published more than 50 peer-reviewed papers. He is a senior member of the IEEE and member of the AIChE. Among his other duties, Dr. Piovoso is an associate editor of *IEEE Control Systems Magazine* and serves as an ABET evaluator. He is a recipient of the 1999 IEEE Control Systems Technology Award and has won an award for teaching at Penn State.

Siva S. Banda leads the Control Science Center of Excellence located at the Air Vehicles Directorate, Air Force Research Laboratory (AFRL), Wright-Patterson Air Force Base, OH, U.S.A. He received his Ph.D. in aerospace engineering in 1980 from the University of Dayton and has been working for the Air Force at AFRL since then. His current research efforts focus on autonomous and cooperative control of multiple unmanned air vehicles, guidance and control of transatmospheric vehicles, aerodynamic flow control, and formation control of microsatellite clusters. Among his many duties, he also serves as a technical advisor to AFOSR, DARPA, ONR, NRC, and NASA on control research for military needs. He served as an associate editor of the *IEEE Transactions on Control Systems Technology,* editor of the *International Journal of Robust and Nonlinear Control,* and associate editor of the *Journal of Guidance, Control, and Dynamics.* He is the recipient of the IEEE Control Systems Technology Award for the year 2000 and the Royal Aeronautical Society's Silver Medal for the year 2001. He is a member of the IEEE Control Systems Society Board of Governors for the year 2002. He is a Fellow of the AFRL, AIAA, and IEEE.