01ICES208

Advanced Control Techniques For Efficient And Robust Operation Of Advanced Life Support Systems

C.W. Pawlowski, A.M. Bell, S. Crawford
Orbital Sciences Corporation

W.A. Sethares
University of Wisconsin, Madison

C. Finn
NASA Ames Research Center

ABSTRACT

This paper examines the structure and performance of three control strategies for a regenerative life support system constrained by mass balance equations. A novel agent-based control strategy derived from economic models of markets is compared to two standard control strategies, proportional feedback and optimal control. The control systems require different amounts of knowledge about the underlying system dynamics, utilize different amounts of information about the current state of the system, and differ in their ability to achieve system-wide performance goals. Simulations illustrate the dynamic behavior of the life support system after it is perturbed away from its equilibrium state or nominal operating point under the three different control strategies. The performance of these strategies is discussed in the context of system-wide performance goals such as efficiency and robustness.

INTRODUCTION

Part of the systems modeling research at NASA Ames Research Center has centered on the use of advanced control techniques to actively manage resources in advanced life support systems (ALSs). One important class of resources is the class of scarce, common-use resources. Power is an example of a scarce common-use resource—most ALS subsystems will use power to process and/or cycle mass, and for heating and cooling. Moreover, power in ALS systems is limited by the power generating capacity of the power plant.

Management of ALS systems requires the integration of diverse yet tightly coupled system elements. Because system elements are coupled, scarce, common-use resources pose an additional challenge for ALS management systems. The added difficulty stems from the fact that system performance depends on meeting both subsystem and system-wide performance goals. For example, reference [3] focuses on the problem of eliminating system-wide power surges while meeting individual subsystem life support requirements under power constraints. The system-wide goal of eliminating power surges may reduce the required size of the power supply by reducing the need for excess capacity.

It is with a view towards meeting performance goals, at both the subsystem and system levels, that we examine three control strategies aimed at intelligent resource allocation. The control strategies differ in their information structure—the amount of information necessary to calculate their controls.

At one extreme is a proportional feedback control about individual system states. This is a completely decentralized information control strategy in that each system state is concerned with its own performance without regard to that of others. When the control task is such that independent subsystem operation can be tolerated, decentralized control is easy to implement and can integrate a diversity of system elements. However, one of the drawbacks of a highly decentralized control system is that it can be difficult to manage common-use resources like power and to achieve system-wide performance goals. These types of considerations require a certain amount of coordination that is not present in a decentralized information structure.

At the other extreme is optimal control, which uses a global cost function to calculate feedback controls. This approach is centralized in the sense that information about all states is used in the calculation of each individual control. While in theory this approach perfectly integrates system elements to achieve a global system performance goal, in practice it is brittle (i.e. sensitive to uncertainty in or changes to system components) since it requires extensive knowledge of the control object. Furthermore, it can be computationally intensive, and it is not necessarily easy to specify a global cost function that captures the system performance criteria.

Market based methods fall in between these two informational extremes. In the market based approach examined here a decentralized control strategy, in this case a proportional feedback around individual states, is supplemented with one or more signals of global scope. These signals are termed prices, and they communicate scarcity of common-use resources. The local controls are designed to respond to these price signals by adjusting their demand for the resources. In this way information about the global state of the system impacts individual control operation. The combination of centralized and decentralized elements make this approach more 'plastic' (tolerant to structural changes) while at the same time capable of addressing global performance concerns. However, this approach often requires balancing the costs of additional communication and computation against the global performance that can be achieved under a centralized information structure.

The purpose of this paper is to examine example control systems with these three information structures and to describe their performance characteristics. The approach is simulation based.

SYSTEM MODEL

The systems modeling group at the NASA Ames Research Center has developed a suite of detailed simulation models of the BIO-Plex Advanced Life Support Test Bed [4] Here we consider a simplified mass balance model of the Air Revitalization System (ARS) (see Figure 1) as our simulation testbed for analyzing the performance of the different control systems.

The key elements of the model are the crew chamber, which shares its atmosphere with the solids processing system; the biomass production chamber, which contains wheat grown with a 24 hour photoperiod and with a constant profile of crop ages; a solid polymer electrolysis (SPE) unit, which produces Oxygen from water; and two buffer tanks which hold Oxygen and Carbon Dioxide. The model tracks Oxygen and Carbon Dioxide only and does not account for system pressure. All flows are on a per hour basis.

STATE EQUATIONS – The underlying dynamics of the system can be written

$$\dot{\mathbf{x}} = f(\mathbf{x},\mathbf{u})$$

where $\mathbf{x}$ represents the vector valued state and $\mathbf{u}$ represents the (vector of) inputs. Given an initial starting

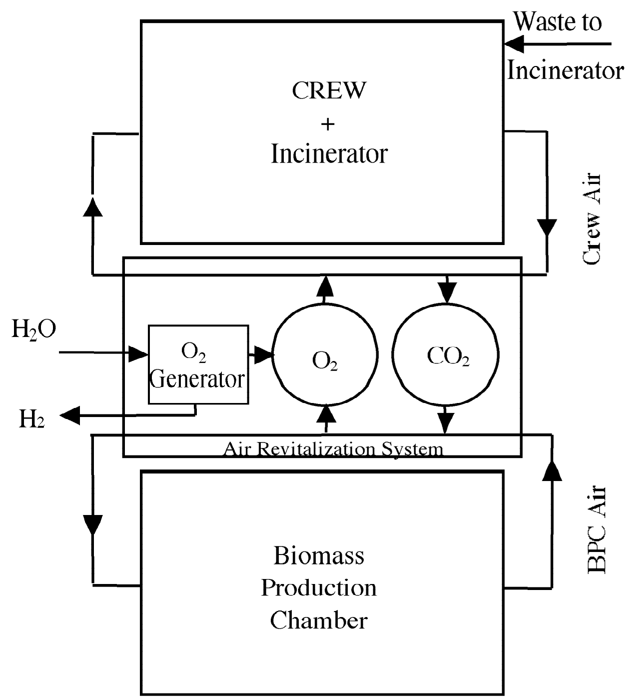

Figure 1 ARS Representation

condition, the equations of state completely determine the evolution of the system and the mass in the various compartments of the model system at any point in time. There are seven state variables and seven controls.

State variables:
x_1, x_2 are the crew chamber molar fractions of O_2 and CO_2, respectively
x_3, x_4 are the plant chamber molar fractions of O_2 and CO_2, respectively
x_5 is the amount (in mols) of O_2 in the O_2 tank
x_6 is the amount (in mols) of CO_2 in the CO_2 tank
x_7 is a moving average of light received by the crop in PPF

Control variables:
u_1 is the molar flow of O_2 from the O_2 tank to the crew chamber
u_2 is the molar flow of crew air to the CO_2 scrubber (we assume that the scrubber does not saturate.)
u_3 is the molar flow of plant air to the O_2 scrubber (we assume that the scrubber does not saturate.)
u_4 is the molar flow of CO_2 from the CO_2 tank to the plant chamber
u_5 is the molar flow of feces to the SPS (incinerator)
u_6 is the molar flow of water to the SPE
u_7 is the light level (PPF)

Parameters:
υ_h is the aggregate constant rate of human O_2 uptake in mols/hr (= 6)

V_h is the volume (in mols) of the crew chamber (34366)
V_p is the volume (in mols) of the plant chamber (16768)

State equations:

$$V_h\dot{x}_1 = u_1 - 52.75u_5 - \upsilon_h$$
$$V_h\dot{x}_2 = 0.8894\upsilon_h + 42u_5 - x_2u_2$$
$$V_p\dot{x}_3 = 1.1\frac{0.00294x_4}{(0.0002^2 + x_4^2)^{1/2}}u_7 - x_3u_3$$
$$V_p\dot{x}_4 = u_4 - \frac{0.00294x_4}{(0.0002^2 + x_4^2)^{1/2}}u_7 \qquad \text{Eq. 1}$$
$$\dot{x}_5 = x_3u_3 - u_1 + 0.5u_6$$
$$\dot{x}_6 = x_2u_2 - u_4$$
$$\dot{x}_7 = -0.5x_7 + 0.5u_7$$

These equations of state (with the exception of state 7) represent a mass balance on the system of Figure 1. The relationship between plant chamber molar fraction of CO_2, light level and CO_2 uptake, given by:

$$\frac{0.00294x_4}{(0.0002^2 + x_4^2)^{1/2}}u_7$$

is derived from an energy cascade model of plant photosynthesis [5]. (add this reference)

SYSTEM EQUILIBRIA – The equilibria of the system $\mathbf{x}^e, \mathbf{u}^e$ occur when $\dot{\mathbf{x}} = \mathbf{0}$. Although there are many equilibria, the requirements of human and plant physiology, as well as considerations of system buffering and operational margins, narrow the relevant number. For example, $(\mathbf{x}^e, \mathbf{u}^e = \mathbf{0})$ is an equilibrium, albeit not a particularly interesting one. Consequently, choice of a nominal operating point is a key element in system performance. We consider the following equilibrium as the baseline or benchmark for the mode:

$$\mathbf{x}^e = [0.23\ 0.0003\ 0.23\ 0.001\ 500\ 500\ 2080]^T$$
$$\mathbf{u}^e = [6.8299\ 19988\ 28.679\ 5.9964\ 0.015715\ 0.46579\ 2080]^T \qquad \text{Eq. 2}$$

Equilibrium states one through four reflect crew and crop physiology, while the equilibrium buffer states are half full. The equilibrium light level is the daily target PPF level required by the crop. The equilibrium controls then follow from the system equations and the constraint $\dot{\mathbf{x}} = \mathbf{0}$.

LINEARIZED EQUATIONS OF STATE – The optimal control problem relies on linearized versions of the above equations of state for analytical tractability. The system dynamics are linearized about the equilibria given by Eq. 2. The linearized system has the form:

$$\dot{\mathbf{z}} = \mathbf{A}\mathbf{z} + \mathbf{B}\mathbf{v}$$

where $\mathbf{z}$ and $\mathbf{v}$ are the states and controls with the equilibria shifted to the origin ($\mathbf{z} = \mathbf{x} - \mathbf{x}^e, \mathbf{v} = \mathbf{u} - \mathbf{u}^e$) and

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -0.58163 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -0.0017103 & 0.015130 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.013754 & 0 & 0 & 0 \\ 0 & 0 & 28.679 & 0 & 0 & 0 & 0 \\ 0 & 19988 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.5 \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} 2.9099e{-5} & 0 & 0 & 0 & -1.5349e{-3} & 0 & 0 \\ 0 & -8.7296e{-9} & 0 & 0 & 1.2221e{-3} & 0 & 0 \\ 0 & 0 & -1.3713e{-5} & 0 & 0 & 0 & 3.9337e{-4} \\ 0 & 0 & 0 & 5.9637e{-5} & 0 & 0 & -3.5761e{-4} \\ -1 & 0 & 0.23 & 0 & 0 & 0.5 & 0 \\ 0 & 3e{-4} & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.5 \end{bmatrix}$$

PERFORMANCE CRITERIA

System performance can be judged by many criteria. Here we discuss two broad classes of system performance goals: Efficiency and Robustness. We use the term efficiency in a general context to signify performance measures that capture resource costs. These generally are of secondary importance. Robustness refers to performance measures that describe the effectiveness of strategies in maintaining system integrity. These generally are of primary importance. For example, power usage is of secondary importance to crew safety.

EFFICIENCY – There are several ways in which the problem of efficient use of ALS resources can be formulated. For example, one way to examine the use of common-use resources such as power is to track their total usage. A given system is deemed more efficient than another if the size of the common pool required during system operation is smaller. In the case of power, this reduces the launch cost significantly.

Efficiency can also be judged relative to a cost function integrated over time. For example, one standard formulation of an optimal control problem uses the cost function:

$$\mathrm{J}(\mathbf{z},\mathbf{v}) = \int_{t_0}^{t_1} (\mathbf{v}^T\mathbf{R}\mathbf{v} + \mathbf{z}^T\mathbf{Q}\mathbf{z})dt$$

which penalizes deviation of the state ($\mathbf{z} = \mathbf{x} - \mathbf{x}^e$) and deviation of the control ($\mathbf{v} = \mathbf{u} - \mathbf{u}^e$) from their nominal values. The matrices $\mathbf{Q}$ and $\mathbf{R}$ are used to weight the contributions of the state and control. Thus the efficiency problem can be cast as the problem of returning to the nominal operating point from a perturbed state while minimizing $\mathrm{J}(\mathbf{z},\mathbf{v})$. When the system equations are linear, this is known as the Linear Quadratic Regulator (LQR) problem.

ROBUSTNESS - Robustness can also be defined and measured according to several criteria. One form of robustness is structural, and pertains to the informational structure of the designed controllers. For instance, the Linear Quadratic Regulator which explicitly solves the problem of minimizing $\mathrm{J}(\mathbf{z},\mathbf{v})$, requires feedback from all states. Such information intensive controls tend to be rigid, i.e. not easily changed or modified, and thus are not robust to changes in the control object or to inadvertent changes in the operating point.

Another criterion for robustness (as applied to ALS) involves the system's response to perturbations, such as its ability to recover from or tolerate unexpected events. These may take many forms. A sensor failure may cause the state to be misread, and a controller to assume incorrect values. How does such an error propagate through the system? For information intensive controllers the error will appear in all the controls, while it may appear in only a single control when using a decentralized strategy. Another form of unexpected event may involve persistent deviations or mis-measurements, i.e., noisy measurements of the state. Again, a controller will be more robust the less such deviations impact system performance. Here we do not examine our controllers by this robustness measure. Rather, we use a much simpler form of robustness criterion.

A third form of perturbation can be viewed as single, short duration perturbations, and the job of the controller is to return the system to its nominal operating point. Such events can be modeled as initial conditions that deviate from equilibrium, and can be studied by examining the trajectories of the system as the controls return the state to equilibrium. The robustness requirement is then to return the system to equilibrium following the disturbance. This is a 'degenerate' form of robustness known as stability.

CONTROL STRATEGIES.

PROPORTIONAL CONTROL – One common and well understood control strategy sets the control proportional to the error between the current and desired state. Proportional control is a special case of Proportional—Integral—Derivative or PID control. The decentralized proportional controllers used here are of the form:

$$\mathbf{u} = \mathbf{K}_p(\mathbf{x} - \mathbf{x}^e) + \mathbf{u}^e \qquad \text{Eq. 3}$$

$$\mathbf{K}_p =$$

$$\begin{bmatrix} -1500 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 50000 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2500 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -5000 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -0.00075 & 0 \\ 0 & 0 & 0 & 0 & -0.075 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.2 \end{bmatrix}$$

These coefficients were determined to achieve system stability. This control is applied to the full nonlinear model (Eq. 1). Note that each control attempts to move towards the equilibrium relying on information about the local state only. P control does not require detailed knowledge of the system dynamics nor does it attempt to meet specified system-wide goals besides stability.

MARKET BASED AGENTS - The market-based approach that we propose and develop starts with the same basic structure and gain matrix of the P controllers discussed above. It is decentralized in that control decisions are made autonomously by individual software agents that rely primarily on local information about their own state. In addition, a market and price are introduced to manage a common resource (power in this case) [2].

Building on previous work on allocation of power in LSSs [3], the market-based control method proposed here seeks to smooth out surges in power demand and to meet power consumption constraints. Not surprisingly, the market only involves those controls which use power. The first four controls regulate the flow of gasses to and from chambers and tanks and do not consume appreciable amounts of power. Consequently, the "agents" for these controls are identical to the proportional controllers—they are entirely autonomous and do not participate in any markets or have any access to system-wide information. The last three controls run processes that use large amounts power. Their choice of controls is given by a function that depends both on the distance of their state from its desired target and on the price of power, which reflects the state of the system and the requirements of other power-using agents. The simplest version of individual agent control choices, those most readily compared to proportional control, have the following form:

$$u_5 = \frac{(u_5{}^e + k_{5,6}(x_6 - x_6{}^e))(M-p)}{\sqrt{p}}; \qquad u_5 \geq 0$$

$$u_6 = \frac{(u_6{}^e + k_{6,5}(x_5 - x_5{}^e))(M-p)}{\sqrt{p}}; \qquad u_6 \geq 0 \quad \text{Eq. 4}$$

$$u_7 = \frac{(u_7{}^e + k_{7,7}(x_7 - x_7{}^e))(M-p)}{\sqrt{p}}; \qquad 2288 \geq u_7 \geq 0$$

where p is the price of power; k_{ij} is the same gain factor used in the proportional controllers, and $M = 2$ is the price at which the process turns itself off entirely.

These functions give the agents' choice of control. When $p = 1$ the market-based agents are identical to the proportional controllers. Other things being equal, when the price is greater than 1 the market-based controls are set to a lower level than their proportional counterparts, and vice versa. To get agents' demand for power, the value of these controls are multiplied by the amount of power used by each process. Summing over agents gives an expression for the total power demanded, which is a function of the price and the states of the agents.

In a market, prices are jointly determined by the demand for and the supply of a good. The specification of the supply side of the power market plays a key role in determining the behavior of the system and consequently must be tailored to the desired operational goals. For example, suppose the supply of power is fixed and constant at all times. Setting this constant supply equal to the demand results in:

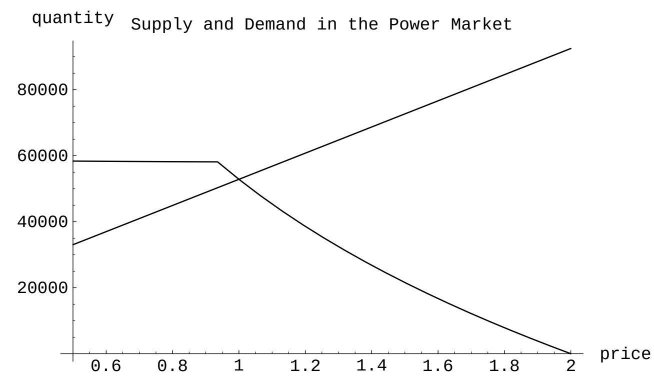

Figure 2 Supply and demand

$$L_5 u_5 + L_6 u_6 + L_7 u_7 = \overline{S} \qquad \text{Eq. 5}$$

where L_5, L_6, L_7 are the power usage of the respective processes and $\overline{S}$ is the fixed supply of power. The constraint on total power use is exogenously determined, and the market clearing price simply allocates the fixed amount of power among agents. The goal of surge management or preventing power usage above a fixed amount can be achieved with this supply strategy.

Another possibility is to have the supply of power also depend on the price. When the states associated with the power using controls are below their targets, demand for power will be higher at every price. Consequently, having the power supply depend positively on the price tends to increase the use of power consuming processes when the states are below their targets and to decrease it when states approach or exceed the targets. As the simulation results discussed below confirm, this smoothes out power and control usage relative to the proportional controls but is more responsive to deviations of the state away from the equilibrium than a fixed supply of power. In order for the market-based control system to have the same equilibrium as the proportional and optimal controls, the supply of power at a price of 1 must equal the quantity of power required to exactly support the equilibrium control use. For example, a simple linear function can be used:

$$S(p) = m\,p + b;\ m = 39634.1;\ b\ =\ 13211.4. \qquad \text{Eq. 6}$$

Putting the two sides of the market together highlights the connection with economic models (See Figure 2). The intersection of the supply and demand curves shows the power usage at the price of 1 when all states are at their equilibrium values. Figure 3 shows a shift in demand that occurs when some of the states are below their target values, increasing overall demand (and price).

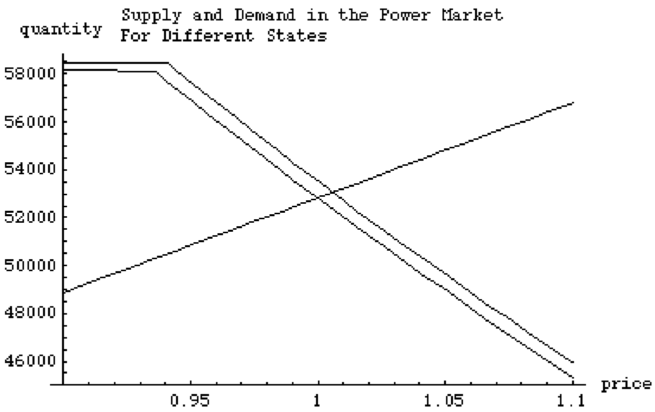

Figure 3 Supply and demand shifted

The preceding discussion and graphs assume a market-clearing price that sets supply exactly equal to demand at each point in time. Alternatively, the price can be expressed as a time varying process that is part of the system dynamics, increasing when demand is greater than supply and decreasing when demand is less than supply. This allows for easy integration of the market with the continuous time system dynamics, and can be less computationally intensive. For example, price dynamics could depend linearly on the difference between supply and demand.

$$\dot{p} = .001(L_5 u_5 + L_6 u_6 + L_7 u_7 - S(p)). \qquad \text{Eq. 7}$$

<u>Alternative Performance Goals for Market Based Controllers</u>—One of the advantages of the market-based approach is its flexibility in incorporating different operating goals involving the common resource. One straightforward way to do this is by altering the supply side of the market for power. The desired power usage is then achieved indirectly through the price. For example, we can trade off between power usage and maintaining the desired states by contracting the supply of power so that 2.5% less power is available at a price of 1 (Eq. 8).

$$S(p) = m\,p + b;\, m = \ 40955.2;\, b \ = \ 10569.1. \quad \text{Eq. 8}$$

The system will converge to an operating point that uses less power but note that it does not return to the target state values. Instead, the system will settle down to an alternative equilibrium balancing a new demand and a new supply at a new equilibrium price.

Similarly, suppose the supply of available power varies over time due to the requirements of other users or other parts of the system. Consider the power supply described by Eq. 9.

$$S(t) = L^e + \ .05\, L^e\, Sin(\frac{\pi}{6} t); \quad L^e = 52845.4 \quad \text{Eq. 9}$$

This supply varies 10% around the equilibrium level on a 12 hour cycle. When power is relatively scarce the market clearing price would be higher than 1 and all of the agents would reduce their demands accordingly. When power is relatively abundant a lower price signals agents to increase their use of controls at that time.

OPTIMAL CONTROL - Optimal control seeks to minimize (or maximize) an explicit cost (or reward) function of the system states and controls. Because ARS requires a nominal power level, attempting to minimize use of power does not a well-posed problem make. That is, power can indeed be reduced, but at the expense of not achieving equilibrium. The speed at which the system returns to equilibrium determines its transient power consumption. Rather than state the problem in terms of power minimization, we examine a secondary cost as captured in the linear quadratic regulator (LQR). The canonical LQR optimal control problem chooses the controls **u** so as to minimize

$$\mathrm{J}(\mathbf{z},\mathbf{v}) = \min_{\mathbf{v}} \int_{t_0}^{\infty} (\mathbf{v}^T \mathbf{R} \mathbf{v} + \mathbf{z}^T \mathbf{Q} \mathbf{z}) dt \qquad \text{Eq. 10}$$

subject to the linear dynamics

$$\dot{\mathbf{z}} = \mathbf{A}\mathbf{z} + \mathbf{B}\mathbf{v}$$

In other words, linear quadratic optimal control minimizes the weighted deviation of the states and of the controls needed to bring the system back to its nominal (equilibrium) state. The weighting matrices $\mathbf{Q}$ and $\mathbf{R}$ are chosen to be diagonal matrices with

$$q_{i,i} = \left(\frac{1}{x_i^e}\right)^2; \quad r_{i,i} = \left(\frac{1}{u_i^e}\right)^2;$$

which assigns equal costs to percent deviations squared in both states and controls.

The linear optimal control problem stated above is solved via the Hamilton-Jacobi equations as in [1], and the solution is applied to the full nonlinear system model (Eq. 1). The optimal control is a full state feedback given by:

$$\mathbf{u} = -\mathbf{K}_{lqr}\mathbf{x}$$

Note that in general each control requires knowledge of all the states. Moreover, detailed knowledge of system dynamics is necessary.

INFORMATION STRUCTURE FOR CONTROLLERS

One of the major differences between the optimal, decentralized, and market-based approaches is the amount of information needed to calculate the controls. At one extreme is the optimal approach, for which each element of the input u is a function of all the states and all the inputs. Thus if

$$\dot{\mathbf{z}} = f(\mathbf{z},\mathbf{u})$$

represents the system dynamics, and $\mathbf{z}^e$, $\mathbf{u}^e$ represent the equilibrium solution, the optimal controls can be written as a function

$$\mathbf{u} = \mathbf{g}(\mathbf{z},\mathbf{u},\mathbf{z}^e,\mathbf{u}^e).$$

This follows from differentiation of the Hamiltonian, which shows that it is possible to write the input directly as a function of the adjoint (the Lagrange multiplier). Since the state and adjoint are concatenated into a two point boundary value problem (usually solved numerically, though it can be solved in closed form when $\mathbf{f}(\mathbf{z},\mathbf{u})$ is linear and when the cost function is quadratic), the adjoint is itself a function of all the states and all the inputs. Accordingly, $\mathbf{g}(\mathbf{z},\mathbf{u},\mathbf{z}^e,\mathbf{u}^e)$ is a composition of the two processes.

At the other informational extreme is the decentralized controller. Rewriting the system dynamics element by element gives

$$\dot{z}_i = f_i(z_i,u_i), \quad i = 1,2,...,n$$

and a decentralized approach chooses

$$u_i = g_i(z_i, z^e)$$

to correct deviations from the equilibrium. The simplest of such controllers feeds back the error between the state and its equilibrium in a linear fashion. The only information required for operation is the corresponding state.

The market based approach lies between these two extremes. A new scalar variable called p (price) is introduced, which is (in general) a function of all the states and inputs, and possibly of the equilibrium values. Thus the dynamics are

$$\dot{\mathbf{z}} = f(\mathbf{z},\mathbf{u})$$

and the price is

$$p = h1(\mathbf{z},\mathbf{u},\mathbf{z}^e,\mathbf{u}^e)$$

for a static price determination, or

$$\dot{p} = h2(\mathbf{z},\mathbf{u},\mathbf{z}^e,\mathbf{u}^e)$$

for a first order dynamic price mechanism. The controllers are then chosen as a function of the corresponding state and the price, that is,

$$u_i = g_i(z_i, p).$$

Thus each controller acts independently of the rest, except through the price variable. The decentralized nature of the scheme has certain advantages, namely modularity, reconfigurability, and low computational load. The centralized aspect allows for coordination across otherwise independent agents.

SIMULATION RESULTS

This section presents a series of simulations that investigate the performance of the three control schemes.

We begin by perturbing the system away from the equilibrium given in Eq. 2 and observing the different dynamic behavior under the different control regimes. Figure 6, which appears at the end of the paper, shows the evolution of all the states and controls for each of the above control strategies starting from an initial condition where the first state is approximately 4% below equilibrium, $x_1(0) = .22$, and all other states are at their equilibrium. Note that the disturbance to the first state quickly propagates to the other states as the system adjusts.

In this simulation and in general, the PID controllers return the state most quickly to its equilibrium. This is directly attributable to the choice of the feedback matrix $\mathbf{K}$. The LQR returns the slowest, because it is minimizing its cost over an infinite time, and because its cost includes penalties on large deviations of all of the control signals. Although it is not obvious from the figures, the LQR values do return to the equilibrium after several thousand time steps. Recall that the MBAs have the same gain matrix as the PIDs, though the actual control values are also a function of the price. The market dynamics tend to smooth out control use over time, which causes the states and controls to remain further away from their equilibrium values for longer than the PID, although they still return faster than the LQR trajectories.

The system dynamics with a fixed, a reduced, and a variable supply of power are shown in Figure 7, also at the end of the paper. The reduced power MBAs have the same basic dynamics as the standard MBA's of Figure 6, but they do not return to the same equilibrium. Rather, they converge to a new state that consumes less power. In contrast, the fixed power MBAs return to the same equilibrium while maintaining a specified (constant) power usage. Such a control strategy may find use when power consumption must equal baseload generation. Price dynamics for the standard, reduced and fixed MBAs are shown in Figure 5.

The variable power supply MBA's fluctuate predictably around their non-varying counterparts. One of the major advantages of the market-based control structure is that it easily incorporates changes in system specification. In contrast, the optimal control approach is very sensitive to these changes. The inclusion of inequality constraints on controls, such as appear in Eqs. 4 and 5, or time varying system elements, as in Eq. 9, not only involve re-specifying the entire problem but also often require an entirely new solution technique. Many reasonable problem specifications simply cannot be solved.

Figure 4 shows power usage relative to the equilibrium power usage. The PID uses the most power and has the largest surge. This occurs because the P controller takes no account of the power usage. In contrast, the standard MBAs smooth out power usage over time because the price mechanism forces the agents to conserve energy usage when it is scarce (expensive). The power usage for the LQR is close to equilibrium (and returns slowly to the equilibrium value) since large deviations from the state and control are punished strongly while small deviations are not (a consequence of the quadratic form of the cost function).

The power reduced MBAs clearly use less power. In a steady state the power reduced MBA uses 0.89% less power by reducing light levels for the plants by 0.86%, reducing use of the SPE to generate oxygen by 3.5%, and reducing use of the incinerator by 7.9%. The states differ from $\mathbf{x}^e$ by the following percentages: {-0.001, 0.16, 0.008, -0.12, 0.032, 0.30, -0.86}.

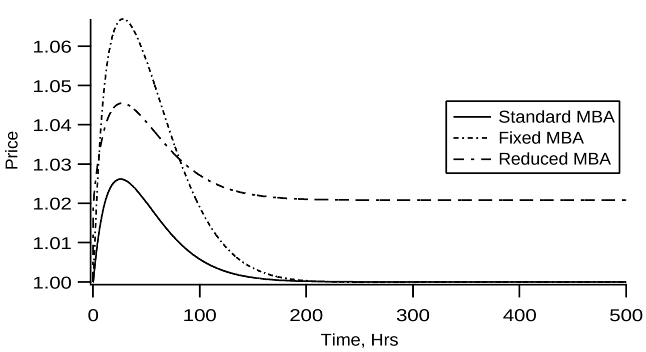

Figure 5 Market Price Dynamics

Figure 5 shows the price dynamics for the three MBA controllers. The price is higher than 1 for the entire simulation. The constant power supply requires a higher price to keep demand for use of controls below the fixed

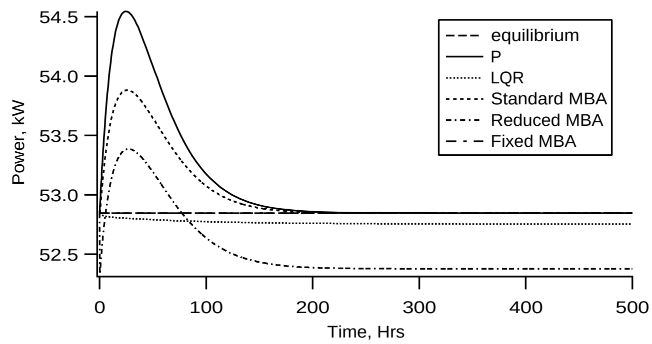

Figure 4 Power Usage

capacity. The higher final price in the power reduced case is a key component in determining the equilibrium of both the states and the controls.

CONCLUSION

This paper examined the performance of three different control structures for an ALS, focusing on a new market-based approach which combines both decentralized and centralized elements. The market based method provides a price signal to independently acting controllers so that constraints on power availability can be met. In addition, the modularity and flexibility of the market-based controllers make it easier to accommodate changes in both agent specification and system goals.

The distributed P control cannot readily accommodate system-wise efficiency or performance goals short of a re-design at the level of individual control gains. This can be time consuming and costly.

The centralized optimal control approach has similar limitations. It can be extremely difficult to incorporate additional constraints or alternative problem specifications such as a system-wide cap on power usage. Furthermore, reducing long term average power is not possible since the state on average remains at its equilibrium (i.e. the optimization has no choice to make). Reductions in power require a new system equilibrium or operating mode and recalculation of the controls.

This does not mean however, that a centralized control approach does not have utility in ALS. Optimal control offers a very powerful method for the realization of precise system goals. In applications where goals are clear and well posed, the computational and communication load light, optimal performance of critical necessity, and the problem or system specifications not likely to change, a centralized scheme likely is preferable.

All three information structures, centralized, decentralized and hybrid, offer utility under certain circumstances. Centralized schemes are favored when the performance goal is of a global nature, and costs of communication and computation are low. A decentralized scheme is appropriate for situations in which system components are independent and modular, and goals tend to be local (e.g. there are no constraints on common resources). A hybrid scheme, such as the market-based approach combines the modularity of decentralized approaches with limited communication to achieve global resource allocation goals. However, more research into the realization of global system goals under a distributed architecture is necessary to be able to fully exploit the market-based approach.

REFERENCES

1. Anderson, B.D.O., Moore, J.B., Optimal control: linear quadratic methods, Prentice Hall: Englewood Cliffs, N.J., 1990.
2. Clearwater, S.H., Costanza, R., Dixon, M., Schroeder, B., (1996) "Saving Energy Using Market-Based Control," in Clearwater, Scott H. (editor). Market-Based Control: A Paradigm for Distributed Resource Allocation. World Scientific, River Edge, NJ, 1996.
3. Crawford, S.S., Pawlowski, C.W., Finn, C., "Power Management in Regenerative Life Support Systems Using Market-Based Control," SAE paper number 2000-01-2259, 30th International Conference on Environmental Systems, Toulouse, France, 2000.
4. Finn, C. 1999. Documentation of the baseline BIO-Plex simulation model. Draft report, NASA Ames Research Center.
5. Volk, T., Bugbee, B., and Wheeler, R. M. 1995 "An Approach to Crop Modeling with the Energy Cascade," *Life Support and Biosphere Science*, 1:119-127.

DEFINITIONS, ACRONYMS, ABBREVIATIONS

ALS: Advanced Life Support

LLS: Life Support System

LQR: Linear Quadratic Regulator

MBA: Market-Based Agents

PID: Proportional—Integral—Derivative

PPF: Photosynthetic Photon Flux (μmols/m^2)

SPE: Solid Polymer Electrolysis

SPS: Solids Processing System

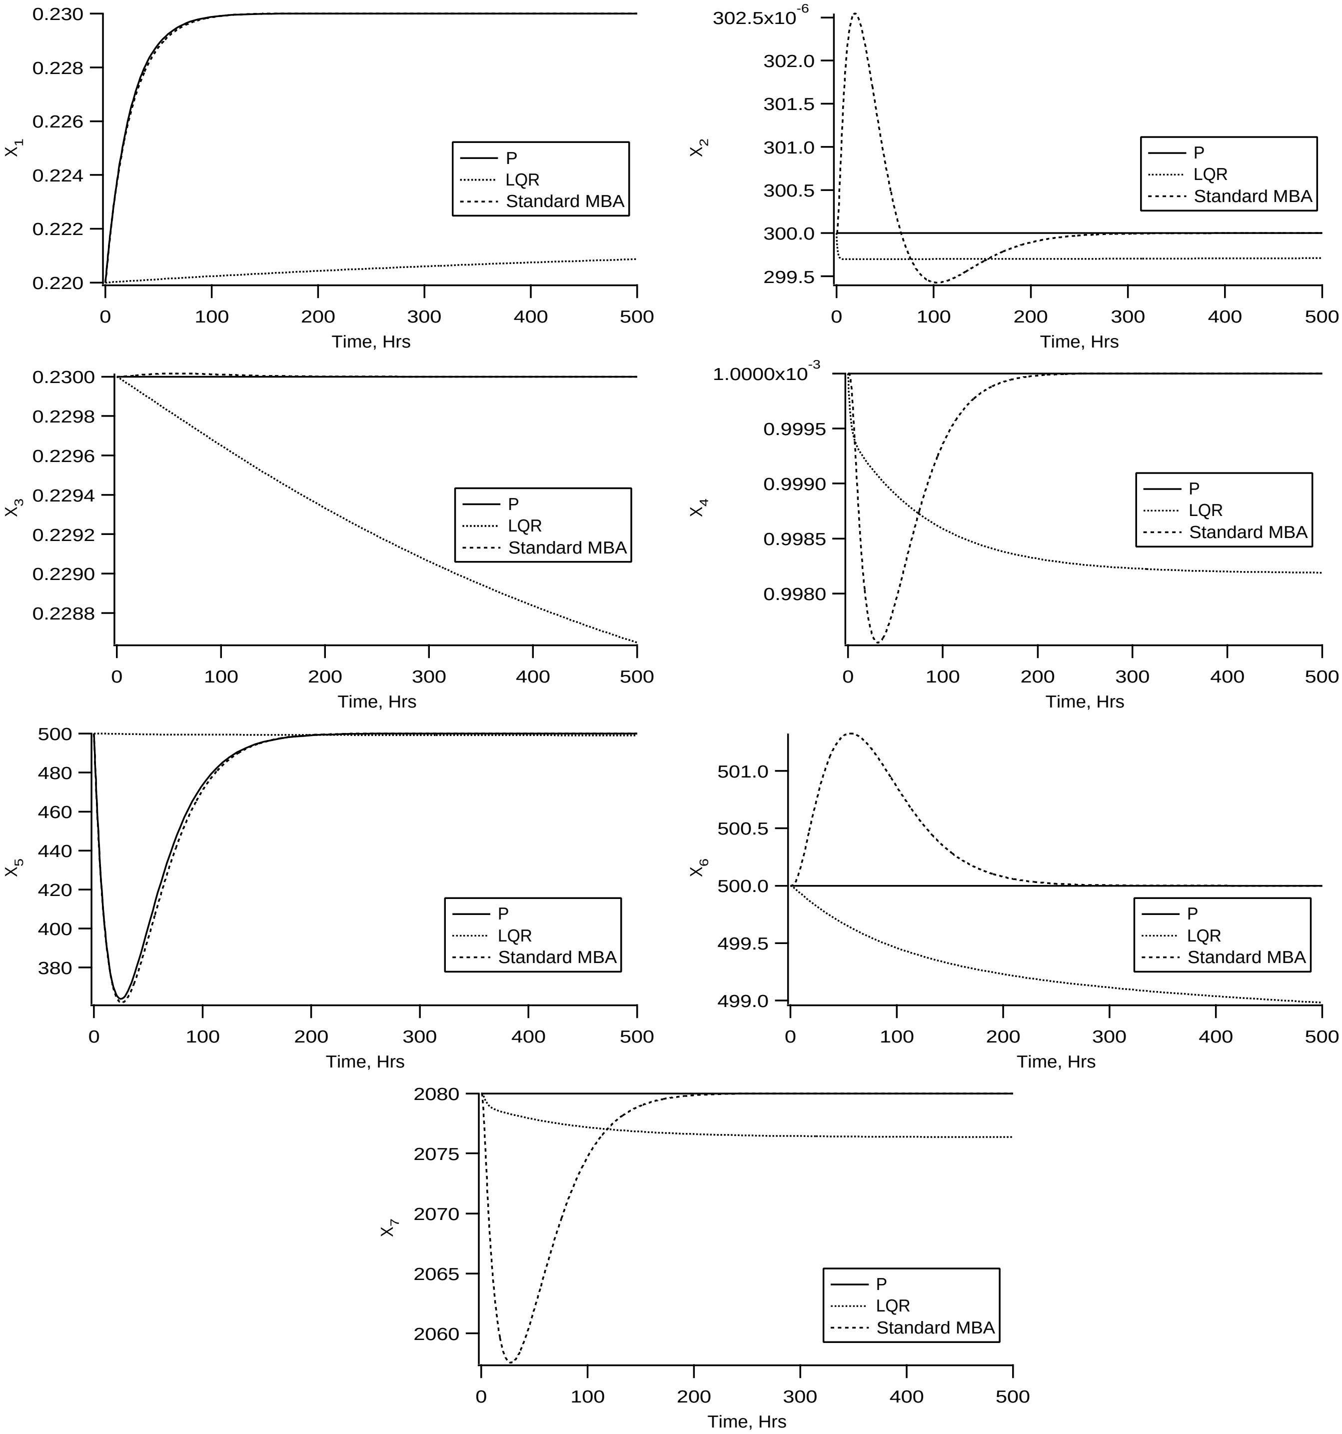

Figure 6 State response under standard equilibrium

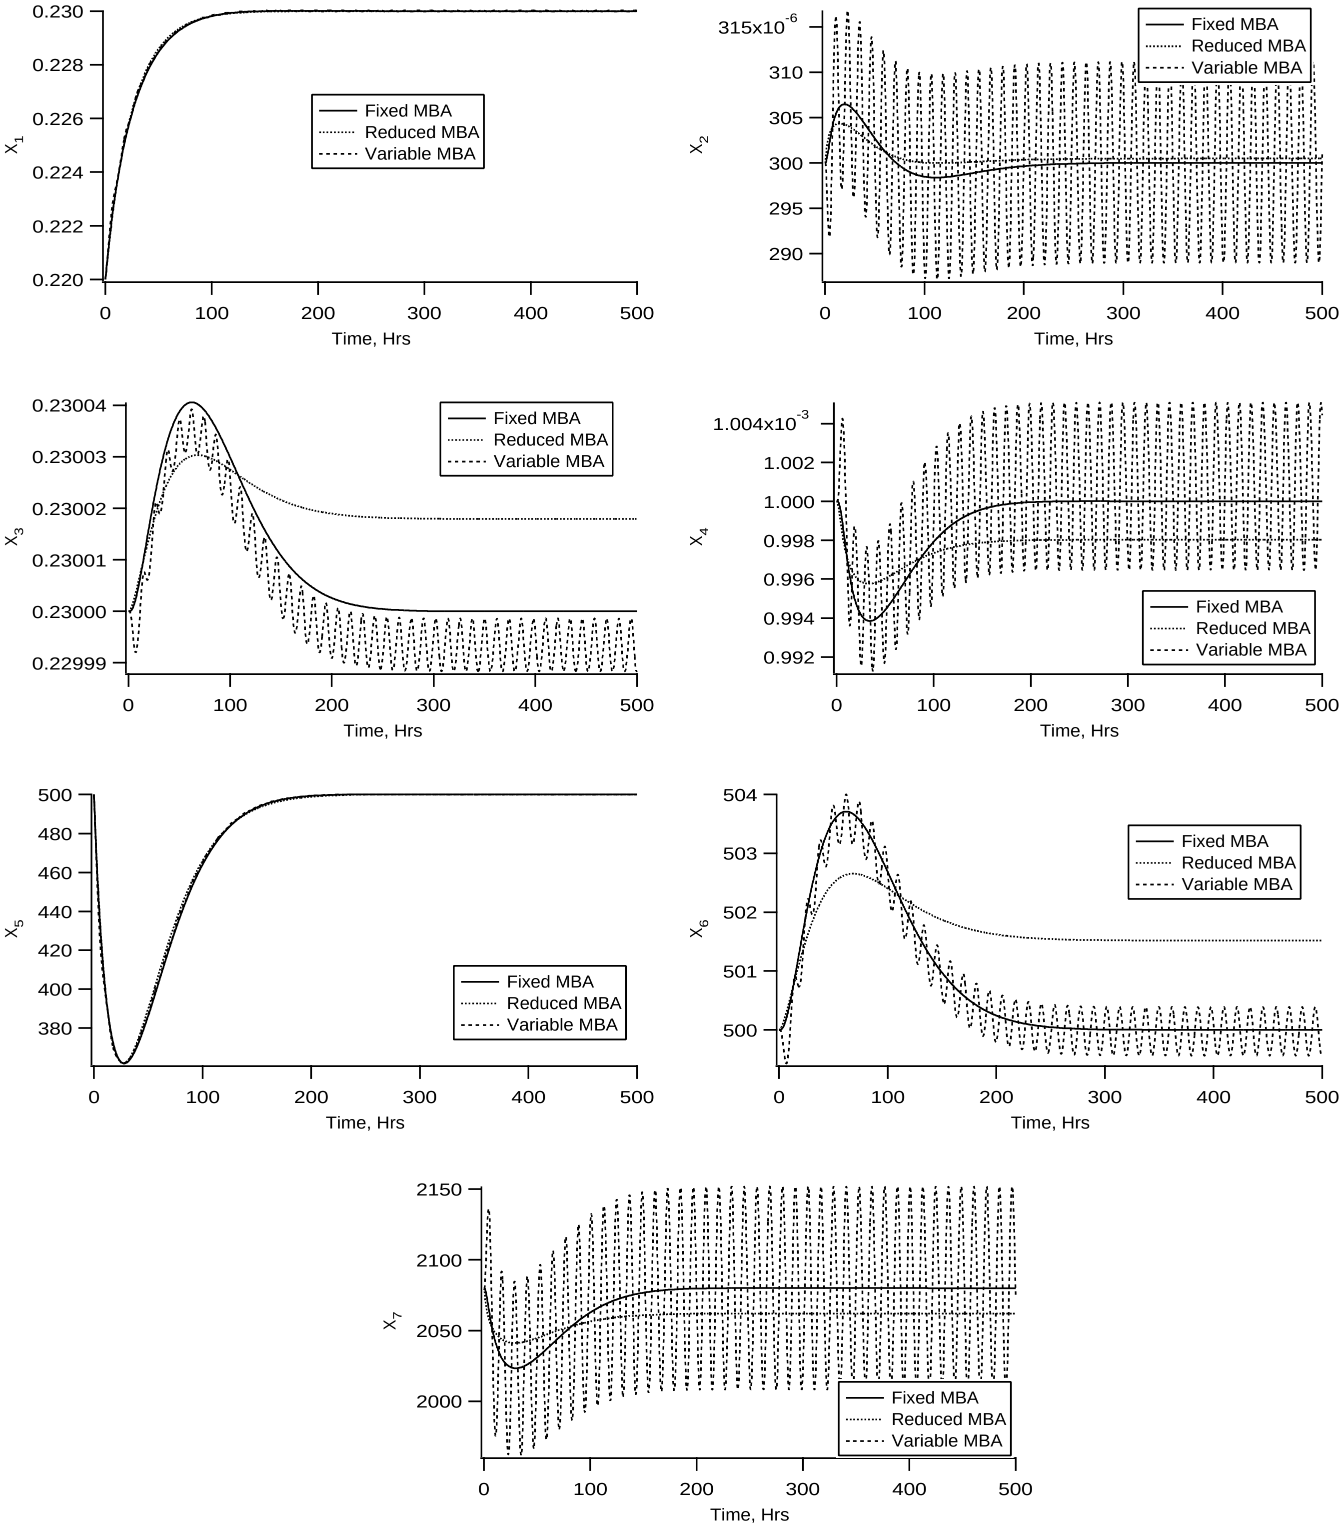

Figure 7 State response, various power constraints