Systems-within-Systems: A Unifying Paradigm

Wayne J. Davis
Professor Emeritus
Industrial and Enterprise Systems Engineering
University of Illinois at Urbana-Champaign

Abstract

This paper introduces a new paradigm for the design and operation of complex systems. Traditional approaches, including the system-of-systems perspective, attempt to reduce the overall into its fundamental components. Such approaches inherently seek to differentiate, and often isolate, the components. The proposed paradigm seeks a shared mission for all components to exploit recursive design practices. An underlying feature of the proposed recursive designs is containment, where identified components are characterized as a system within another system.

While developing the proposed paradigm, the core system technologies—mechanics, controls and planning—were also unified. This unification was accompanied with an unanticipated unification of time—the past, present and future. In fact, a second temporal axis has been introduced to facilitate the on-line concurrent implementation of planning and control responsibilities. The paradigm discusses the inherent deficiencies of planning in general, and the specific limitations of applying optimization in a real-world planning situations. It establishes an insurmountable need for another agent to implement an entity's plan while refuting the subordinate stature that traditional hierarchical architectures would assign to an implementing agent. Rather, this interdependency establishes the need for expanded interaction among the planner and implementing agent, including the necessity of collaborative planning among the interacting systems.

Introduction

This paper explores a new paradigm for design and operations of systems, one that I suspect many will find disturbing. This proposed paradigm rejects many fundamental principles underlying traditional system engineering. It seeks to unify technologies. It seeks to unify time while placing severe constraints upon one's ability to observe the present. It views planning and control as collaborative activities while severely diminishing the potential contributions that established approaches as optimization or feedback control might provide.

Given the expansive nature of the proposed paradigm and its associated consequences, it cannot be adequately portrayed within a single paper. Thus, this paper itself is only an introduction. Much of what is presented has been substantiated by over a hundred pages of comprehensive mathematical derivations and extensions to existing theories and principles. Only the most basic mathematical arguments are presented here. As an introduction, the paper focuses upon listing concerns and issues.

Two Introductory Examples

> *Example 1*: You are about to enter a limited access highway. Before accelerating into traffic, however, you carefully study the situation before you, including the flow of traffic and the highway's geometry. After formulating your plan, you close your eyes and begin implementing the plan as you merge into traffic. With closed eyes, you continue implementing the plan until it is no longer possible to do so, (i.e. until you encounter a barrier, run off the highway or collide with another car.) At that point, if you are capable of doing so, you open your eyes and initiate the next planning cycle. Obviously, the above scenario is not quite realistic because planners seldom execute their plan.
>
> Let us further assume that you are a blindfolded passenger who must rely upon the driver's description in your assessment of the situation. After planning your response, you relate it to the driver. The driver then interprets and implements your plan until it fails. Meanwhile, you wait for the next planning cycle.

The above example depicts planning as a task. The planner first formulates the problem, often without immediate knowledge of the considered situation. After solving the problem, the planner typically relies upon another agent to manage the implementation. The discovery of the optimal solution terminates the current planning task. The planner then initiates the next planning task, which often revisits a prior problem, when its solution can no longer be implemented.

The implications associated with seeking an optimal solution are seldom discussed. First, the planner needs to consider a perfect statement of the situation to be faced. Should an alternative situation be encountered during the solution's implementation, the current solution will probably not be optimum, and may cease to be feasible. On the other hand, seeking an optimal solution becomes an irrelevant pursuit if one considers the inherent uncertainties. One might minimize the expected cost of addressing the situation, but any solution that is generated will still be optimal to a particular scenario, which has little probability of occurring. When uncertainties are considered, there are several criteria for seeking an optimal solution. A pessimist might seek to minimize the maximum cost while the optimist might maximize the maximum payoff.

There are also additional facets of planning that are seldom addressed. In order to specify its current problem, the planner must quantify any future interactions with another entity during the adopted planning horizon. These quantifications isolate the planner from the other agents with which he routinely interacts. The same situation also arises with respect to the other planners. These planners are also trying to anticipate their interactions with the other agents. In short, every planner is trying to optimize their portion of the overall problem from their own perspective. Unfortunately, one cannot demonstrate that the ensemble of optimal solutions to the individual problems can be assembled into an optimal solution to the overall problem.

In competitive situations, the notion of a global optimal solution becomes confused. In non-competitive cases, however, the potential benefits of collaborative planning are obvious. Moreover, it is probable that each planner interacts with a unique ensemble of other planners. In such situations, a network of concurrent collaborations evolves among the planners.

> *Example 2.* A person in Champaign discovers that she must drive to a particular address in Chicago by 3 p.m. the next day. Several questions emerge immediately. When should she leave? What route should she take? In this case, the obvious solution is to take I-57 from Champaign to Chicago. After this decision has been made, a trio of tasks emerges. Task 1 is to drive from the starting point in Champaign to a selected entrance onto I-57. Task 2 involves driving along I-57 to the selected exit. Task 3 is to drive from the selected exit to the specified downtown destination.
>
> Task 1 starts with her departure. A route to the desired entrance onto I-57 must have been selected. This first task then evolves into a sequence of subtasks of driving along specified streets between appropriate intersections. Now consider the driving along a given street between two intersections. She enters the street by turning into a particular lane and later exits the street from a particular lane. During the intervening drive, she can transfer between lanes within reason to expedite her movement in traffic. While driving in any lane, she must observe appropriate constraints upon her speed as follow other cars. Because these tasks are usually executed with little or no conscious thought, her thoughts might be anticipating what lies ahead. The lead car in her lane of traffic suddenly brakes. Without making a conscious choice, she instinctively brakes while her conscious effort turns to the car which is braking. Sometime later, she encounters a detour and quickly reconfigures the remaining route to her selected entrance. She might even decide to travel to another entrance. When she finally enters the interstate highway, the second task begins. While driving along I-57, she manages her flow with the traffic. She also continues to compare her prior success against her original driving plan. With this comparison, she can update the anticipated arrival time.

In this example, the rationale for optimal planning is less obvious. There are numerous tradeoffs to be considered. The inherent uncertainties cannot be ignored. The continued utility and feasibility of any plan must be reassessed. Implementing any plan involves several concurrent activities, each addressing different elements of the response at different resolutions.

Clearly, the latter example involves the most sophisticated planning. It also involves an intelligent planning capability which current planning methodologies cannot replicate. As humans, we plan automatically and incessantly throughout our existence. We continually reassess our situations, choose among our available options and respond. Our responses usually can be partitioned into tasks. While implementing any particular task, a sequence of subtasks often emerges. Our dynamic response crosses numerous time scales, ranging from milliseconds to hours, hours, days, and beyond. Almost everyone has formulated five- and ten-year goals at one point or another.

Now let's address a new paradigm for coordinating several interacting systems. The new paradigm arose from my professional goal to unify three principal technologies: mechanics, planning and control. This goal emerged during my graduate studies and was later reinforced while engineering several complex real-world studies that ranged from designing Illinois' emergency response plans to assessing the efficacy of a first-generation, computer-integrated-manufacturing hierarchy for an integrated steel mill to assisting the National Institute Standards and Technology in developing new on-line planning capabilities for their experimental manufacturing execution systems.

Before introducing the new paradigm, however, I will first list several assertions upon which the new paradigm is based. Given their scope, it will be difficult to rigorously demonstrate their validity.

No Time like the Present

The Heisenberg Uncertainty Principle places insurmountable constraints upon one's ability to simultaneously measure a physical entity's position and momentum. It also implicitly limits one's ability to observe the present. In fact, the new paradigm assumes that it is impossible to observe the present. As a mental exercise, let us ponder what we might see if we could instantaneously freeze the moment, i.e. stop time. In our universe, physical objects are comprised of vibrating molecules. The molecules, in turn, are comprised of vibrating atoms. The atoms are comprised of vibrating atomic particles. The atomic particles are comprised of vibrating subatomic particles. This recursion continues until one eventually encounters what some believe to be vibrating strings of energy. Given this conjectured universe, an instantaneous image of the universe would likely be empty (void of any physical entity). Obviously, this is purely conjecture because it is impossible to take an instantaneous snapshot of any physical entity. Our observations of any physical entity arise from the intrinsic blurring that arises from observing the trajectories of constituent elements over a finite time interval. Moreover, the perceived image is comprised of photons emitted across a spectrum of times. To better appreciate the consequences that arise from this, consider one of the many pictures captured by Hubble telescope as it peers into distant space. Often, such images are accompanied with a caption that asserts one is viewing the universe so many billions of years ago. In reality, the recorded image never occurred. Because the recorded objects are different distances from the earth, we are seeing each object at different time.

Pondering such issues is beyond my knowledge of cosmology. However, one would do well to accept the reality that the present cannot be observed. Observations are simply approximations or virtual images of what has occurred. Moreover, any physical response we seek to manage must reside in the object's virtual future. Our image of the present might be best described as the difference between two open sets: our current conceptualizations of the past and the future. Should you question this assertion, you might recall a time when you were totally engrossed in a particular activity, say reading a book. Whenever your activity is temporarily interrupted, you will likely not be able to estimate the current time without first consulting a clock. On the other hand, had your attention been less focused, you likely could estimate the current time. The human instinct continuously assesses its recent observations as it anticipates the immediate future. This instinct caused the example's driver to react subconsciously to the car's braking in front of her even before her conscious attention returned to the emerging situation.

What's Next?

The notion of a present pervades the traditional system technologies. In the field of controls, one often assumes that the current state is known. When planning, the temporal horizon typically originates from the present. Both assertions are physically impossible. With regard to controls, one observes the recent past while seeking to manage the immediate future. The size of the deadband about the present depends upon several physical factors. Consider the factors that arose while managing the Mar's explorer, Rover. The delay between recording an image on Mars and receiving that image on earth prohibits the employment of earth-based feedback control measures to control Rover's real-time activity on Mars. The earth-based control is restricted to assigning future tasks or goals which the Rover's on-board controllers subsequently executed using its on-board feedback control systems.

It is also impossible to plan for the present because time continues to pass while the plan is generated. Planners may only consider the future. After the earth-based planners specified future assignments for Rover, they relied upon its on-board controllers to implement their assigned tasks. Rover consequently served as an agent for implementing the planner's assignments. This situation is not unique for every planner must rely upon a local agent to execute the response that it had previously planned for the current moment.

Planners must plan between two future times. Unfortunately, the planner does not have total control upon the state from which it initiates its plan or the goal state it should seek. Because the planner must rely upon a local agent to act upon its behalf, the outcome of the local agent's execution determines the future state from which the planner initiates its planning. Moreover, the planner must anticipate that outcome even while the local agent's executes upon its behalf. If the planner postpones its planning until its agent completes its current assignment, then the agent must wait while the planners decide which goal should be pursued next.

The planner is also an agent that attempts to achieve its assigned goals. Consider Rover's handler. The manager must decide which tasks should be pursued next. Several alternatives often exist. The quality of each alternative must be evaluated. Usually, such evaluations assess the given alternative's contribution toward achieving a more comprehensive goal or assignment. Realistically, the necessary goals must be assigned by another entity. One should not conclude, however, that the agent's role is subordinate to the planner. This is not a hierarchical relationship. Rather, it is a symbiotic relationship. The planner accomplishes nothing without the implementing agent. The agent lacks direction or mission without a goal. Interestingly, however, a given entity cannot be characterized exclusively as a planner or an agent. Each entity serves both roles. As an agent, it assigned goals direct its planning. As a planner, each entity relies upon an agent to execute its plans. The entities differ not with their roles, but rather in their mission and in the time horizons they address.

One Size Fits All

Reductionism currently provides the primary modus operandi for system analysis. The recent characterization of large systems as systems-of-systems further substantiates the fundamental contribution that reductionism provides as it attempts to identify the fundamental components comprising the overall system. Once the components are known, each can be singularly analyzed. One needs only to design a central interface with which each component interacts after the solution for each component has been engineered.

Reductionism has at least three critical flaws: it emphasizes the difference among the components rather than their similarities; it isolates an individual component from direct interactions with the other components; and it does not distinguish between the core systems that physically exist and the virtual systems that manage their existence.

On the other hand, by emphasizing the similarities among the system's components, one can employ one of the most useful concepts in algorithmic design—recursion. Because most entities function both as planner and executing agent, a common design for their concept of operation and their interaction with other entities emerges. More importantly, the entities interact directly with each other rather than through a central communication device.

The three principal system technologies—mechanics, controls and planning—can also be unified under the singular phenomenon of equilibration. This technological unification was accompanied with a temporal unification of the past, present and future. Although the temporal unification was unexpected, both unifications were essential. Mechanics provides the principles for modeling past performance. Planning chooses a desirable future response from the potential responses. Control manages the instantiation of the immediate future into the immediate by bridging the temporal threshold, which we designate as the present.

From the combined technological and temporal unifications, a second temporal axis naturally emerged, which the paradigm characterizes as virtual time. The origin of the virtual time axis is fixed to the current real time. Future plans continuously evolve for a designated future time interval along the positive virtual time axis while system identification concurrently models the system's observed response along the negative virtual axis. Control straddles the origin of virtual time axis where it manages the instantiation of the future into the past.

A New Paradigm

The existence of an integrated planning and executing agent creates the potential for a recursive system architecture. The new paradigm assumes that the considered system manages the real-time responses of an ensemble of physical processes or entities. Their physical behavior can be modeled using the principles of mechanics because these entities are physical. We will also assume that the physical entities are time variant, implying that on-line system identification will be essential. We are guaranteed that we can delineate each physical process because no two objects can occupy the same region of space and time. On the other hand, we will assume that individual entities interact with each other for different purposes and upon different time scales.

Processes physically exist in the real-world, and their responses explicitly evolve with real time. That is, each exists in the present. Newton's First Law asserts that every physical object exhibits an inertial response in real time unless an external force accelerates its motion. As an entity traverses its inertial trajectory, no external force performs work upon the system and its energy remains constant.

Newton's Second Law states that an external force can accelerate the entity. Moreover, the entity's acceleration is directly proportional to the applied force. The ratio of force to acceleration is defined to be the entity's inertial mass. Finally, Newton's Third Law roughly states that any external force will be counteracted by an internal force of the same magnitude but opposite direction.

Let us consider an object whose state trajectory between an initial and final time will be denoted as

$$S_{t_i \le t < t_f} = \left\{ \mathbf{s}(t) \,\middle|\, t_i \le t < t_f \right\}$$

Let us further assume that an energy metric exists for this state trajectory such that energy at state $\mathbf{s}(t)$ will be evaluated as $E(\mathbf{s}(t))$. We then assert that the state trajectory $S_{t_i \le t < t_f}$ is an inertial trajectory, if the energy at every state $\mathbf{s}(t)$ along the trajectory initiating from $\mathbf{s}(t_i)$ is constant or

$$E(\mathbf{s}(t))\Big|_{\mathbf{s}(t) \in S_{t_i \le t < t_f}} = E(\mathbf{s}(t_i))$$

In Figure 1, the lower state trajectory prior to the time t_i corresponds to an inertial response, denoted by $S_{t<t_i}$. The energy at any point along the curve prior to t_i (the bold black segment on the upper figure) can be evaluated and is equal to $E(\mathbf{s}(t_i))$ because the energy (upper curve) is constant over this inertial trajectory. On the subsequent time interval between t_i and t, a dynamic force function $F_{t_i \le t}$ is applied causing the considered entity to accelerate from its inertial trajectory.

Since an external force is being applied, work is being performed upon the system and its energy changes as denoted by the red solid trajectory. Letting $\mathbf{F}(t) \in F_{t_i \le t}$ represent the force at time $t \ge t_i$ then the energy along the state trajectory $S_{t_i \le t}$ can be computed as

$$E(\mathbf{s}(t))\Big|_{\mathbf{s}(t) \in S_{t_i \le t}} = E(\mathbf{s}(t_i)) + \int_{t_i}^{t} \mathbf{F}(\tau) \bullet \frac{\partial \mathbf{s}(\tau)}{\partial \tau}\, d\tau$$

Recalling that the derivate of the state vector with respect to time yields the velocity vector $\mathbf{v}(t)$ or

$$\frac{\partial \mathbf{s}(t)}{\partial t} = \mathbf{v}(t)$$

The dot product of the instantaneous force vector with the instantaneous velocity vector or $\mathbf{F}(t) \bullet \frac{\partial \mathbf{s}(t)}{\partial t}$ represents the instantaneous power being delivered to the system or the rate at which work is being performed upon the system at time *t*.

Should one stop performing work upon the system at time $t \ge t_1$, then a new inertial response would continue beyond t (indicated by the red dashed line in the lower figure). Its associated inertial energy curve (indicated by the red dashed line in the upper figure) has a constant energy

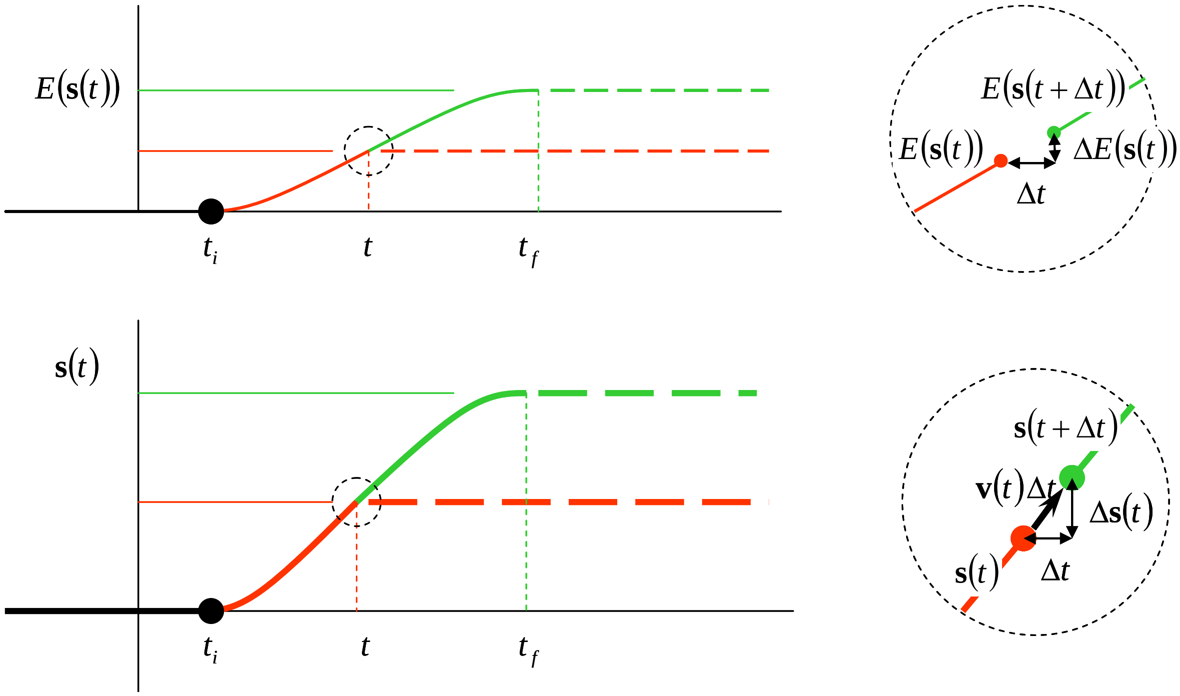

Figure 1—Inertial versus Perturbed System Responses

$$E(\mathbf{s}(t)) = E(\mathbf{s}(t_i)) + \int_{t_i}^{t} \mathbf{F}(\tau) \bullet \frac{\partial \mathbf{s}(\tau)}{\partial \tau}\, d\tau$$

In Figure 1, however, the external force $F_{t_1 \le t < t_f}$ continues to act upon the system, allowing the energy to be evaluated

$$E(\mathbf{s}(t))\Big|_{\mathbf{s}(t) \in S_{t_i \le t < t_f}} = E(\mathbf{s}(t_i)) + \int_{t_i}^{t} \mathbf{F}(\tau) \bullet \frac{\partial \mathbf{s}(\tau)}{\partial \tau}\, d\tau \text{ for } t_i \le t < t_f$$

Observe that the projected energy curve beyond the original t is indicated by a solid green line. Because we are assuming that no external force will be applied beyond t_f, the system returns to an inertial response with the constant energy curve (green dashed line) of value

$$E(\mathbf{s}(t_f)) = E(\mathbf{s}(t_i)) + \int_{t_i}^{t_f} \mathbf{F}(\tau) \bullet \frac{\partial \mathbf{s}(\tau)}{\partial \tau}\, d\tau$$

The corresponding state trajectories are on the lower curve with the solid green segment representing the perturbed or forced state trajectory beyond t and the dashed green segment representing the subsequent inertial state trajectory beyond t_f where external forces are no longer being applied.

In Figure 1, the junction between the red and green accelerated segments in both upper and lower curves has been circled and magnified to the right of each respective curve. Looking at the lower right circle, the red state trajectory terminates with a large red dot representing the state $\mathbf{s}(t)$. The subsequent green state trajectory initiates slightly later from the green dot representing the state $\mathbf{s}(t+\Delta t)$. The change in state between over the indicated time interval Δt can be computed as

$$\Delta\mathbf{s}(t)=\mathbf{s}(t+\Delta t)-\mathbf{s}(t).$$

If one takes the limit of the indicated slope with the time interval Δt approaching zero, one obtains the velocity vector or

$$\mathbf{v}(t)=\lim_{\Delta t\to 0}\frac{\Delta\mathbf{s}(t)}{\Delta t}=\frac{\partial\mathbf{s}(t)}{\partial t}$$

Because the velocity vector is tangent to the state trajectory at each time *t*, the velocity vector $\mathbf{v}(t)$ multiplied by the change in time Δt approximates the vector from state $\mathbf{s}(t)$ to state $\mathbf{s}(t+\Delta t)$ or

$$\mathbf{s}(t+\Delta t)=\mathbf{s}(t)+\Delta t\,\mathbf{v}(t)$$

Now consider the circled inset of the upper energy curve. Here the smaller red dot corresponds to the energy of $\mathbf{s}(t)$ or $E(\mathbf{s}(t))$ while the green dot corresponds to the energy of $\mathbf{s}(t+\Delta t)$ or $E(\mathbf{s}(t+\Delta t))$. By now the reader probably has observed that the energy $E(\mathbf{s}(t))$ is being indicated as an explicit function of the state $\mathbf{s}(t)$at a given time rather that a particular time. In the figure, we indicate the corresponding change of energy as

$$\Delta E(\mathbf{s}(t))=E(\mathbf{s}(t+\Delta t))-E(\mathbf{s}(t))$$

which we can compute as

$$\Delta E(\mathbf{s}(t))=\int_{t}^{t+\Delta t}\mathbf{F}(\tau)\bullet\frac{\partial\mathbf{s}(\tau)}{\partial\tau}\,d\tau$$

Taking limit of the above change in energy as Δt approaches zero, yields

$$\begin{aligned}\frac{\partial}{\partial t}E(\mathbf{s}(t))&=\lim_{\Delta t\to 0}\frac{\Delta E(\mathbf{s}(t))}{\Delta t}=\mathbf{F}(t)\bullet\frac{\partial\mathbf{s}(t)}{\partial t}\\&=\mathbf{F}(t)\bullet\mathbf{v}(t)\\&=\|\mathbf{F}(t)\|\|\mathbf{v}(t)\|\cos\theta\end{aligned}$$

Observe that only the component of the force vector along the velocity contributes to a change in the system's energy. Let us review. Beginning at time t_i, we applied the force function $\mathbf{F}(t)$ to the considered system until the final time t_f. During this time interval, the system was accelerated from its inertial trajectory to the state trajectory $S_{t_i\le t<t_f}$ shown as the lower curve in Figure 1. We computed its corresponding value of energy $E(\mathbf{s}(t))$, depicted as the upper curve, for each state $\mathbf{s}(t)$ along this trajectory. Looking in detail at the upper curve, we observed that energy is being accrued by the system

$$\frac{\partial}{\partial t}E(\mathbf{s}(t))=\|\mathbf{F}(t)\|\|\mathbf{v}(t)\|\cos\theta$$

This implies that the system's trajectory actually records the instantaneous component of the force along its instantaneous velocity vector as the force is being applied to the system. Thus, given the energy curve $E(\mathbf{s}(t))$ from t_i to t_f, one can recover the recorded component of the instantaneous force along the velocity vector $\mathbf{v}(t)$ as

$$\mathbf{F}(\mathbf{s}(t)) = \frac{\frac{\partial}{\partial t} E(\mathbf{s}(t))}{\|\mathbf{v}(t)\|^2} \mathbf{v}(t).$$

Observe that we have distinguished the applied force $\mathbf{F}(t)$, a function of t, from the remembered force $\mathbf{F}(\mathbf{s}(t))$, a function of the state $\mathbf{s}(t)$.

Previously we stated that Newton's Third Law asserts for every action there is an equal and opposite reaction. The reaction forces will oppose the applied force when the considered entity is at equilibrium. If the object is at rest, then any applied force must be countered by an equal force from the physical constraints that prevent the system from accelerating. However, when the system is moving along an inertial trajectory at a constant energy, the component of any external force along the instantaneous velocity vector $\mathbf{v}(t)$ must be zero. An object sliding along a horizontal frictionless surface provides an illustrative example. In this case, the upward normal force provided by the table is perpendicular to the object's horizontal velocity vector.

The above observation also establishes that one maximizes the work performed by an external force when one applies any accelerating force along the direction of the object's instantaneous velocity. For example, when one transfers a satellite between orbits, one typically performs the rocket burns at the apogee or perigee of the orbit where the instantaneous velocity vector is perpendicular to instantaneous radius from the satellite to the center of gravity about which the object is orbiting.

The existence of the inertial force $\mathbf{F}(\mathbf{s}(t))$ represents another form of system memory that is often overlooked. One typically focuses upon determining the trajectory for the primal state vector $\mathbf{s}(t)$. However, every system also possesses a dual or adjoint state vector that corresponds to the system's inertial force to further perturbation at any state along its primal state trajectory. Our last derivation demonstrated that the inertial force or dual variable $\mathbf{F}(\mathbf{s}(t))$ could be evaluated using the time derivative of the system's energy. One can also derive an alternative state transition function for computing dual variable $\mathbf{F}(\mathbf{s}(t))$ as a function of time. This formulation, however, is conditioned upon knowing the current primal state trajectory.

We will next focus upon the interaction between the two (red and green) primal state trajectories at time t in Figure 1. Certain constituency relations necessarily exist at their junction. Specifically, the initiating state $\mathbf{s}(t+\Delta t)$ for the green trajectory can be interpreted as acting upon the final state $\mathbf{s}(t)$ for the red trajectory. The rate at which the state $\mathbf{s}(t+\Delta t)$ is performing work upon the state $\mathbf{s}(t)$ can be computed as $\mathbf{F}(\mathbf{s}(t+\Delta t))\mathbf{v}(t+\Delta t)$, while the final state of the red trajectory $\mathbf{s}(t)$ absorbs energy from the state $\mathbf{s}(t+\Delta t)$ at the rate $\mathbf{s}(t)\frac{\partial}{\partial t}\mathbf{F}(\mathbf{s}(t))$. Let the primal state vector be represented by $\mathbf{p}(t)$ (in this case, $\mathbf{s}(t)$) and the dual state vector be represented by $\mathbf{d}(t)$ (in this case, $\mathbf{F}(\mathbf{s}(t))$). If we then allow Δt approach zero, the state trajectories will be continuous at time t if the following constituency relation holds

$$\mathbf{p}(t)\frac{\partial}{\partial t}\mathbf{d}(t) = \mathbf{d}(t)\frac{\partial}{\partial t}\mathbf{p}(t)$$

The above relationship is none other than the complimentary slackness conditions or Kuhn-Tucker-Karush conditions to the optimal control problem for transferring the primal system from initial primal state $\mathbf{p}(t_i)$ to the final primal state $\mathbf{p}(t_f)$.

We have observed that there are two state variables for each system. The primal state variable is traditionally considered because it evaluates the critical attributes that describe the system as a function of time. However, the dual state variable is also important because it provides a temporal description of the inertial forces resulting from the system being accelerated from its inertial trajectory. It can be observed that the residual inertial forces or dual state variables are identically zero for a system traversing its inertial state trajectory.

It was previously claimed that mechanics, controls and optimization could be unified. That demonstration has now been realized. In discussing the state transition mechanisms, we employed the system's energy for characterizing the system's response. Classical mechanics also employs energy-based formulations, particularly the Hamiltonian or the Lagrangian. We further characterized the trajectory planning problem as an optimal control problem and observed that the symmetry relationship between the product of one state variable with the temporal derivative of its complimentary state variable could also be derived by applying the Kuhn-Tucker-Karush necessary conditions for an optimal solution to two-point, boundary-value, optimal control problems. This unification will allow us to conceptualize a system architecture based upon a recursive entity that performs integrated planning and control of the entities that are contained within its domain.

Crossing Boundaries

The proposed architecture begins by specifying the physical processes that are contained within the overall system. In Figure 2a, the controllers to three processes are denoted by the red, green and blue looped arrows bridging the $\tau = 0$ line. On the other hand, the included τ-axis represents a second temporal dimension, virtual time, upon which each process controller implements the essential on-line planning, control and identification functions for managing its assigned process. As discussed earlier, the origin of the virtual-time τ axis is coupled to the current real time, implying that virtual time corresponds to a relative time with respect to real time. Thus a positive value of virtual time relates to a future real time while a negative virtual time denotes a past real time.

The looped arrow in Figure 2a illustrates that each process controller manages the instantiation of its process's imminent future into its immediate past. The looped arrow also recognizes that the process controller cannot assess the process's present value. Rather, the process controller records the process's prior trajectory to the left of its loop, as it concurrently projects the process's near-term state trajectory to its right.

Figure 2b shows each process controller projecting its state trajectory to a common future time, and the projected trajectories have been characterized by their associated energy. These energy curves resemble the potential energy that process will have consumed in traversing to the projected state at the given virtual time because the physical processes perform work as they modify elements of the real world.

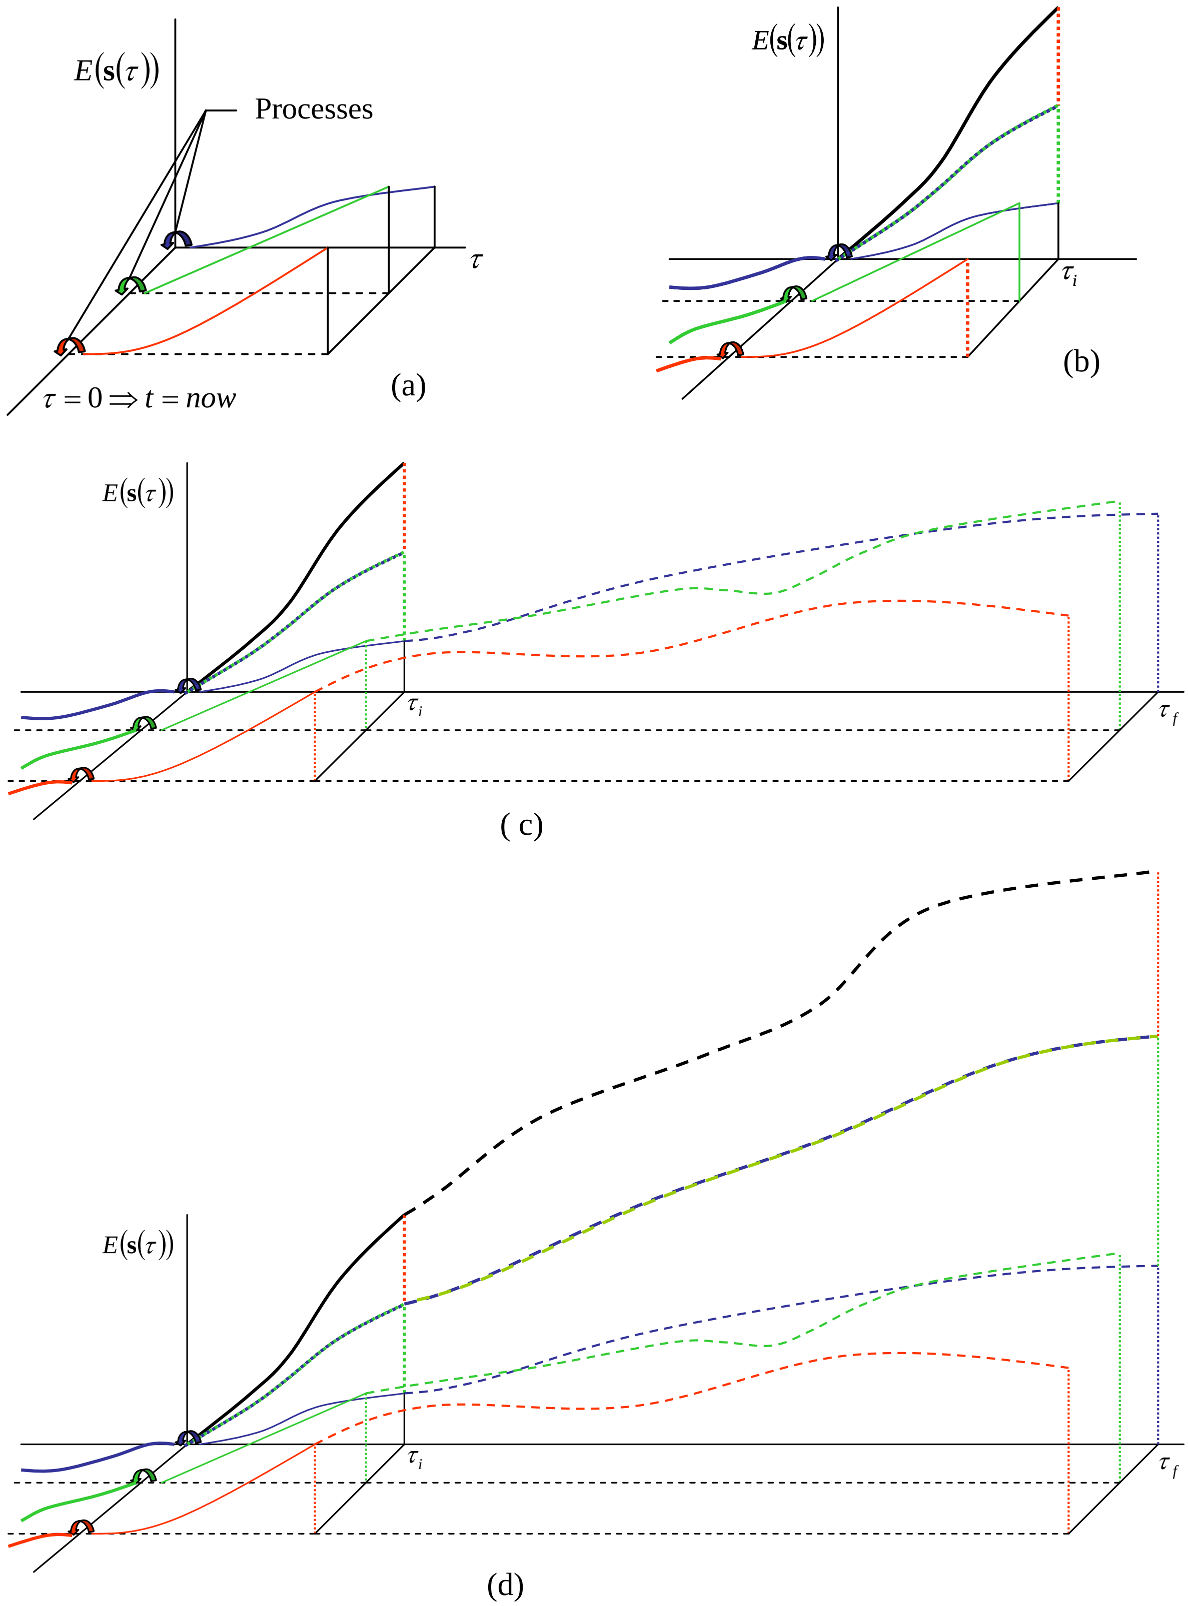

Figure 2—Coupling System Responses

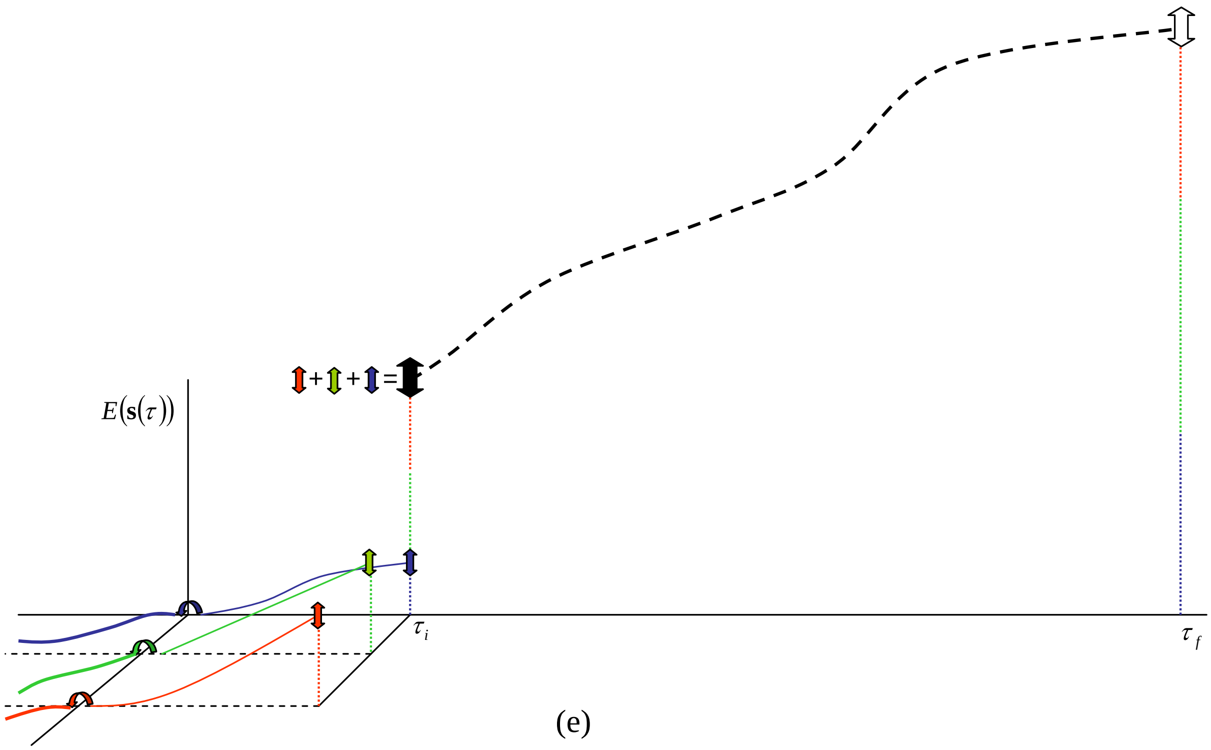

(e)

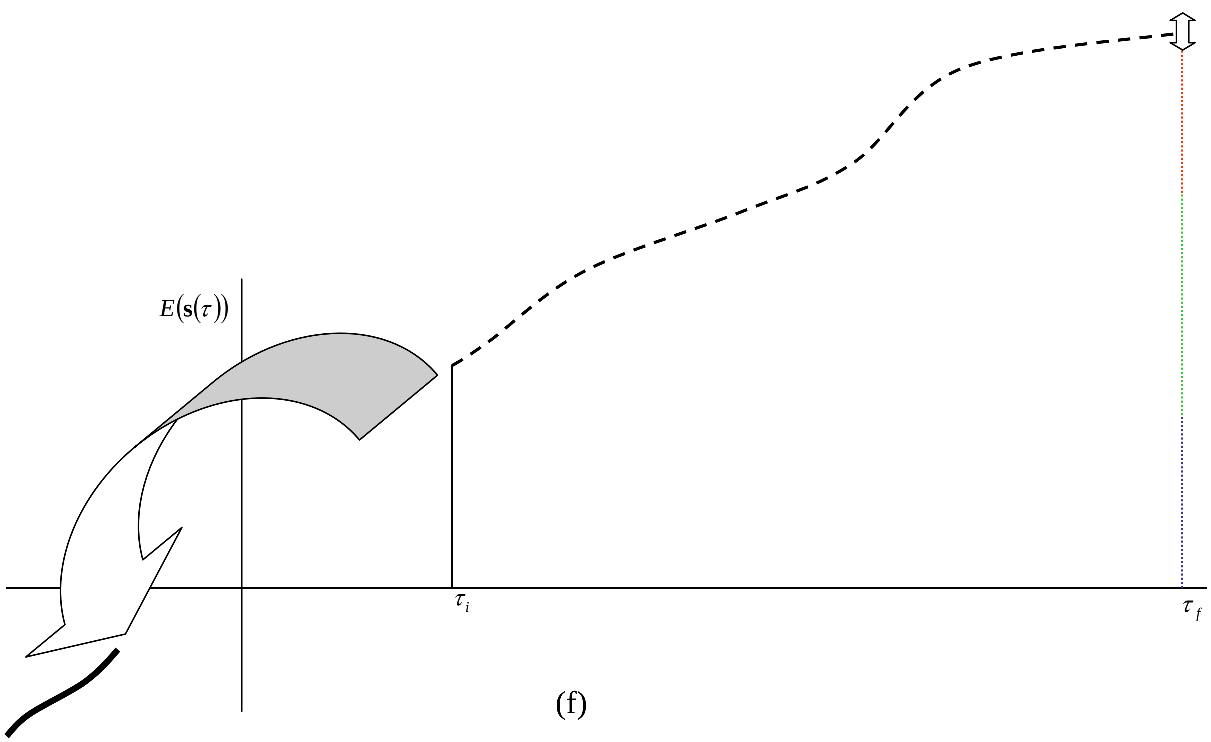

(f)

Figure 2 (cont.)—Coupling System Responses

The three projected energy profiles are summed in Figure 2b. The blue-green curve represents the sum of blue and green energy curves, while the black curve represents the sum of all three energy curves. The black curve represents the total amount of projected work that will be performed by the three processes starting from now into the future because the black total energy curve initiates from the origin of the energy versus virtual time coordinate system,.

An entity other than the process controllers has projected responses for each of the three processes between the virtual times τ_i and τ_f in Figure 2c. Assume that the new entity's response contains or builds upon the responses of the three processes and their dedicated controllers. In Figure 2d, the three projected responses for the individual processes are summed to provide the projected total energy curve for the contained processes, depicted as the dashed black energy curve.

Several of the prior curve segments have been removed to illustrate the proposed recursive algorithm In Figure 2e. First, the composite energy curves have been removed from on the interval between $\tau = 0$ and τ_i because the system will rely upon the projected responses from the dedicated process controllers. During the interval between τ_i and τ_f, the three individual processes will be aggregated into a single composite response. Consequently, the individual projected energy curves during this time interval have been deleted. The projected composite response is initialized to the sum of the projected energies for the three contained processes at τ_i. The composite response is then planned in order to achieve a desired goal state at the virtual time τ_f.

A Systematic Interface

Now we focus on the interface between the process controllers and the aggregate controller at virtual time τ_i. Assume that the tasks, which are currently being executed by the three processes under the supervision of their dedicated controller, were previously assigned by the composite controller. Thus, each process controller's projection for its process's state at time τ_i represents its success in functioning as an implementing agent for the composite controller. Observe that each process controller has initialized its projection to its most recent observation for its process's state. Their projected states at time τ_i will likely changes as real time advances because the future behavior of the processes can be uncertain. The initial state for the composite process trajectory at virtual time τ_i must also be dynamic because the projected states for the processes are dynamic.

Traditionally, one employs a constant definition for the system's state, which behavior is then evaluated as a function of time. Refer to this traditional description as the system's primal configuration. This primal configuration has both a primal and dual formulation, as previously discussed. Now we introduce a dual configuration for a system, which *fixes time* and *allows the definition of state to vary* at that time. The proposed dual configuration might characterize the interface between the individual process controllers and the composite controller at virtual time τ_i.

The detailed state vectors for the component processes at this interface have been aggregated into a composite state vector. That is, at the indicated virtual time, the state definition has been modified. We will designate the dual configuration to be the system's *coupler*, given this potential role for the dual configuration. Figure 2e illustrates the proposed coupling among the state variables of the interacting systems.

The proposed composite controller plans its trajectory from a composite initial state at virtual time τ_i toward an assigned goal state at virtual time τ_f. When the planned composite trajectory is traversed, the difference between its associated final and initial projected energies will be consumed. The composite controller attempts to minimize the energy consumed in achieving its assigned goal state at τ_f. It is important to observe, however, that the composite process trajectory will never be implemented as planned; i.e., it will do no work. Rather, the planned composite trajectory determines when each future task will be executed and which process will execute it—the composite process controller is planning what its component processes will be doing after completing their current individual assignments.

Now we can return to the coupling interface at virtual time τ_i. The composite process controller must have previously assigned the current tasks to its contained processes. Given their current assignments and observed state, each process controller updates it projected trajectory while finishing its current assignment. The composite controller initializes its planned response to projected states of its contained processes at the selected virtual time τ_i. It is also considers additional tasks it composite tasks it must execute in order to achieve its assigned goal state at τ_f.

Given its current plan for scheduling and assigning these tasks, the composite controller projects its composite response from the initial virtual time τ_i to the final virtual time τ_f. The project final state at τ_f will vary with the advancement of real time because the projected process states at virtual time τ_i are dynamic. Moreover, the projected final state will also probably deviate for its assigned goal state.

Consider the composite controller's options since it might change the current assignments for one or more of the contained processes. It might improve its planned response between its projected initial state and its assigned goal state. It might also negotiate a new goal state. It is important to observe that any assigned goal states, irrespective of which controller is making the assignment, must be achievable. Whenever any controller either assigns or accepts a goal that cannot be achieved, then the system becomes uncontrollable. The inescapable reality is that the planner must always rely upon another agent to implement its prior plans for the present while it continues to plan for the future. Assigning meaningless goals to any implementing agent is an exercise in futility.

Systems-within-Systems

Observe that the composite controller and its implementing process controllers could be also viewed as a singular unified process controller. Moreover, the unified controller could have been represented as one of the red, blue or green process controllers. Still other composite controllers, not yet specified, might be unified with their constituent process controllers and function as yet another unified process controller. In fact, the red, blue and green process controllers could all be replaced unified controllers. We would only need to change the units of time along the virtual axis, such that one second now corresponded to perhaps one minute. At this new level of resolution, we might again introduce a composite controller for the collection of unified process controllers. The promised recursion has emerged.

When engineering large systems, the first step begins by listing the physical processes that are contained within the overall system. The responses of the physical processes necessarily evolve in real time. For each process, one next identifies a process controller, which observes its process's prior response while managing its future response toward achieving assigned goals. Part of this management function includes projecting its process's future response.

After grouping a subset of the controlled processes into a composite process, we next introduce a composite process controller. The composite process controller plans its response from the coupled responses of it component processes at a given time. Specifying the requisite dynamic coupling between the composite process's state and the states of its contained processes at any virtual time provides the dual configuration for the composite controller.

The recursive process can be repeated by simply treating each composite process controller as a process controller and then forming a composite controller from the set of the former composite controllers. Each time the recursive aggregation is employed, its associated coupling occurs at a later virtual time implying that every composite system must rely upon feedforward (i.e. projected responses) of its coupled processes. Only the most basic process controllers, those that are managing physical processes, can derive any utility from feedback information.

Equilibration versus Optimization

It merits repeating that it is impossible to plan for the present. Every planner must rely on an implementing agent to execute its prior plans regarding the present. Even process control cannot access the present, since it attempts to influence the next state toward achieving its current assignment. It must accept the reality, however, that it can influence but not dictate that next state.

The same situation occurs for any composite controller. It can influence the initial state for its future planning but it cannot impose that state. Although many of my former students were unfamiliar with a Slinky, I often used one to demonstrate the notion of equilibration. I usually asked a student to move both ends of the slinky slowly with respect to each other. Then I urged the other students to watch the behavior of the spanning coil. As its ends moved, the slinky continuously sought a new equilibrium. Whenever the ends stopped moving, its dynamics would decay and a new equilibrium would emerge. The same is true for our composite controller. So long as the considered initial and final states are moving, the controller continues to update its planned trajectory for traversing the two states.

The controller might attempt to optimize the trajectory for fixed boundary states, although that optimization is simply a special case of equilibration. Remember that the necessary condition for optimality, the complimentary slackness condition, yields the constituent relation between the primal and adjoint system trajectories whenever equilibrium is restored. Optimizations seek complementary slackness; dynamic systems seek equilibrium. From that perspective, the proposed system paradigm offers nothing new—it simply mimics nature.

Perhaps the only novel element of the new paradigm is recognition of the dual or coupling configuration between two or more interfaced systems. This coupling permits one system to perform work upon the other while the other distributes the resulting energy to further perturb its state trajectory from its original inertial/equilibrium configuration. It should be further observed that the proposed interactions are not hierarchical. Eventually all changes to the real world must be executed by the physical processes, and the composite controllers are essentially seeking which processors will execute the necessary tasks and when.

A Need for Further Unification

The intent in developing the proposed paradigm was to unify rather than invent. A first unification sought to exploit the similarity among the component subsystems within a system rather than emphasizing their distinguishing characteristics; a second unification sought to support direct interaction among the components rather than providing a central interface through which the interactions must occur; and a third sought to integrate planning, control and identification. The fourth and unanticipated unification sought to integrate the past, present and future.

These unifications established a need for collaborative planning. An individual entity can plan between specified initial and final states. We have been addressing that problem for decades by assuming that the initial and final states are known. The unifying paradigm asserts that these initial final states represent the shared state variables for the coupling dynamics among interacting systems at a given time. As such, any initial state must be established in collaboration, not isolation.

Each composite controller addresses at least three forms of planning concurrently: collaboratively specifying its initial state using the feedforward projections of its implementing agents for alternative goal assignments, collaboratively specifying its goal state while serving as an implementing agent for one or more other composite controllers, and individually determining its current plan for transitioning between its current specification of an initial and final planning states.

There are actually other optimizations to be addressed. Recall that a composite controller's state transition function represents an aggregation of its component's state transition function. We previously assumed that the processes are time variant, implying that their state transition functions change in time. For time-variant systems, another optimization, termed system identification, becomes essential. The composite controller must collaborate with its component processes to update its state transition model because the composite controller's dynamics are derived from the dynamics of its time-variant processes. It must also participate in their collaborative system identification because the composite controller's dynamics are included in the dynamics of any other controller for which it serves as an implementing agent.

System identification is also inherently dependent upon observing prior responses. This necessarily implies that the identification process is constrained by prior usage of the real-world process. The true capabilities of the process can only be ascertained through experimentation; that is, until something different is tried, one cannot ascertain whether the explored behavior is feasible or not. This experimentation obviously represents a form of learning, and one must choose what performance frontier should be explored next. Again, this determination corresponds to another optimization.

As stated above, prior planning/optimization has focused upon traversing the state transition function between specified initial and final states, and that planning is inherently dependent upon the quality of the employed state transition function. The quality of the state transition function is dependent upon identifying the best model for explaining the observed behavior. The information that identification might glean from the observed responses is dependent upon prior efforts to explore the process's operational envelope. Each form of optimization is dependent upon the other. All three optimizations must be addressed concurrently and also collaboratively in the real-world.

Much remains to be accomplished. For that reason alone, the propensity to invent should be restrained, and the need for unification again emerges. We need to determine how advancements in the contributing technologies can be shared. Although, this paper cannot cover all aspects of the proposed unifying perspective, considerable work has already been accomplished. Over a hundred pages of detailed mathematical derivations have been generated to validate this paradigm for linear, nonlinear and discrete-event systems. Yes, much more work is needed but remember much work has been done already. We need only to organize what has been done and consider it from a new perspective, one that is not biased by preconceived notions whose validity has never been challenged.