

An Information Space View of "Time": From Clocks to Open-Loop Control *

Steven M. LaValle¹ and Magnus Egerstedt²

¹ Department of Computer Science, University of Illinois, Urbana, IL 61801 USA
lavalle@uiuc.edu

² School of Electrical and Computer Engineering, Georgia Institute of Technology,
Atlanta, GA 30332, USA.
magnus@ece.gatech.edu

Abstract. This paper addresses the peculiar treatment that time receives when studying control systems. For example, why is the ability to perfectly observe time assumed implicitly in virtually all control formulations? What happens if this implicit assumption is violated? It turns out that some basic control results fall apart when time cannot be perfectly measured. To make this explicit, we introduce information space concepts that permit imperfect time information to be considered in the same way as imperfect state information. We then argue that classical open-loop control should be reconsidered as perfect *time-feedback* control. Following this, we introduce a notion of *strongly open-loop* control, which does not require perfect time observations. We provide some examples of these concepts and argue that many fascinating directions for future controls research emerge.

1 Introduction

“Until the mid 1750s, navigation at sea was an unsolved problem due to the difficulty in calculating longitudinal position. Navigators could determine their latitude by measuring the sun’s angle at noon. To find their longitude, however, they needed a portable time standard that would work aboard a ship. However, to find their longitude, they needed a portable time standard that would work aboard a ship. The purpose of a chronometer is to keep the time of a known fixed location, which can then serve as a reference point for determining the ship’s position. Conceptually, by comparing local high noon to the chronometer’s time, a navigator could use the time difference to determine the ship’s present longitude” [7].

As known to anyone who has ever been exposed to a course on control theory, a sharp distinction is made between closed-loop and open-loop control. This distinction is drawn between control laws that involve references to the state

* The authors thank Roger Brockett and Dan Koditschek for helpful comments and discussions. LaValle is supported in part by the DARPA STOMP program (DSO HR0011-07-1-0002). Egerstedt is supported in part by the National Science Foundation through NSF-CAREER award (grant # 0237971).

of the system and control laws that are specified only in terms of time. The idea is that time is somehow readily available to the controller, independent of the state of the system. Of course, time imperfections are often addressed in discrete-time or delayed-measurement models; however, a rich variety of other time uncertainty models can be imagined, rather than just quantizations or delay.

This paper came about because the special treatment of time seemed to the authors to be somewhat arbitrary in the context of robotics, particularly when considering that all information comes from sensors. As a simple example, consider the two systems: $\dot{x} = f(x, u)$, $u = g(t)$ and $\dot{x} = f(x, u)$, $\dot{z} = 1$, $u = g(z)$. Here one typically asserts that the first case is an open-loop controller and the second is closed-loop. However, they are mathematically the same control law (modulo initial conditions) and both seem to involve some kind of feedback.

The fundamental issues concerning time are not limited to simple mathematical technicalities. For example, in robotics, a number of sensors, including compasses, gyroscopes, accelerometers, wheel-encoders, GPS antennas, IR, sonar and laser-range finders, and cameras (just to name a few) are employed to estimate the internal state of the robot as well as characterize the environment in which it is deployed. However, clocks are also routinely employed, not just for synchronizing the executions, but also, for example, for smoothing the pose estimates or for establishing distances through ultrasonic range sensors. As such, it appears that clocks should be treated as any other sensor because they provide measurements of an important quantity for specifying control laws (whether implicitly or explicitly).

Rather than establish a collection of new results, this paper emphasizes the reconsideration of time and its role in control theory. This may lead to exciting new avenues for research, ultimately providing a better understanding of how time information affects control. Furthermore, systems can be designed that utilize minimal amounts of time information, thereby achieving greater robustness and affordability.

Of course, the study of time within a controls context is certainly not new. For example, in [1] it was pointed out that a temporally driven sampling strategy (so-called Riemann sampling) could be advantageously replaced (in some contexts) by a state-driven strategy (Lebesgue sampling). Similarly, by allowing for time to be controlled, dynamic-time warping has become a standard controls tool, for example for speed regulation in robotics [6, 9, 14]. Moreover, the view that open-loop control is potentially problematic is also not new, as illustrated by the fact that jitter in the clock is known to cause instabilities. As another example, going from open-loop control signals to corresponding closed-loop control signals is a research topic of continued interest in optimal control [4, 5]. Another relevant topic is asynchronous protocols for distributed systems (e.g., [8]). All of these research areas, although disjoint in explicit focus, share the feature of treating time in a non-standard manner.

To formulate time uncertainty in a general way, we extend standard machinery that was developed mainly for state uncertainty: *information spaces*. The earliest ideas are due to Kuhn in the context of sequential games [10], and were

found to be a convenient, unified way to express state uncertainty in dynamic game theory [2], stochastic control [3, 11], and planning algorithms [12]. Our approach is to treat time as “just another state variable” and consequently consider observations of time, parameterized by an internal, continuous index set.

2 Illustrative Examples

2.1 Brittle Time Sensitivity in Linear Systems

As a first example of the reliance on perfect time measurements, consider the standard, linear time-invariant system

$$\dot{x}(t) = Ax(t) + Bu(t),$$

in which the uncontrolled system is unstable, and (A, B) is a completely controllable pair. Using a static feedback law $u = -Kx$ for stabilizing this system yields $\dot{x}(t) = (A - BK)x(t)$, and $x(t) = e^{(A-BK)t}x_0$, in which $x(0) = x_0$. Now, in the absence of uncertainty we can of course implement this exact control law with an “open-loop” controller

$$u(t) = -Ke^{(A-BK)t}x_0.$$

We now assume that time has to be measured, and instead of t we observe $h(t)$, in which h is an output mapping. We obtain the measurement error $|h(t) - t|$, which is in fact due to a bias in the clock (in that it goes too quickly or too slowly) while it never causes “time” to go backwards; i.e., we assume that $dh(t)/dt > 0$.

Hence, the open-loop controller takes on the form $u(t) = -Ke^{(A-BK)h(t)}x_0$, and subsequently

$$\dot{x}(t) = Ax(t) - BK e^{(A-BK)h(t)}x_0.$$

Now, consider the system

$$\dot{\tilde{x}}(t) = dh(t)/dt(A - BK)\tilde{x}(t), \quad \tilde{x}(t) = x_0.$$

This system is globally, asymptotically stable as long as $dh(t)/dt > 0$, which was assumed. The solution is

$$\tilde{x}(t) = e^{\int_0^t (dh(s)/ds) ds (A-BK)} x_0 = e^{(A-BK)(h(t)-h(0))} x_0.$$

Under the additional assumption that the initial time measurement is correct, i.e. that $h(0) = 0$, we thus have that $u(t) = -K\tilde{x}(t)$, which in turn implies that

$$\begin{aligned} \dot{x}(t) &= Ax(t) - BK\tilde{x}(t) \\ \dot{\tilde{x}}(t) &= dh(t)/dt(A - BK)\tilde{x}(t). \end{aligned}$$

The first of these equations is an unstable linear system, driven by an input that will decay to zero because \tilde{x} will tend to zero, and, as a result, the x -system is unstable. The only situation in which this will not happen is when $x(t) = \tilde{x}(t)$

for all t , which directly implies that $dh/dt = 1$, i.e., $h(t)$ is equal to $t + c$ for an arbitrary constant. However, since $\tilde{x}(0) = x(0) = x_0$ this means that $c = 0$. Hence, the only way in which the “open-loop” system will remain stable is with perfect time measurements. As an example, consider the situation depicted in Figure 1, where

$$A = \begin{bmatrix} 2 & 4 & 3 \\ 1 & 0 & 0 \\ 0 & -1 & -1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 2 \\ -0.5 \end{bmatrix},$$

$$K = [3.5 \ 3.2 \ 3.8].$$

The system starts at $x(0) = (0.1, 0.1, 0.1)^T$. For a measurement model of the form $h(t) = t + f(t)$ for some function f , the upper plot shows the case in which $f(t) = 0$, and the lower shows a small disturbance of $f(t) = 0.01 \sin(t)$.




Fig. 1. Depicted is the evolution of a third order system. The upper plot corresponds perfect time measurements and the lower plot shows the evolution for the time disturbance $h(t) = t + 0.01 \sin(t)$.

2.2 A Nonlinear System with Time Robustness

In contrast to the previous time-sensitive system, consider the nonlinear system defined over $X = \mathbb{R}$,

$$\dot{x} = \prod_{i=1}^5 (a_i - x),$$

with real constants $a_1 < a_2 < \dots < a_5$. This system has three stable equilibrium points $\{a_1, a_3, a_5\}$ and two unstable $\{a_2, a_4\}$.

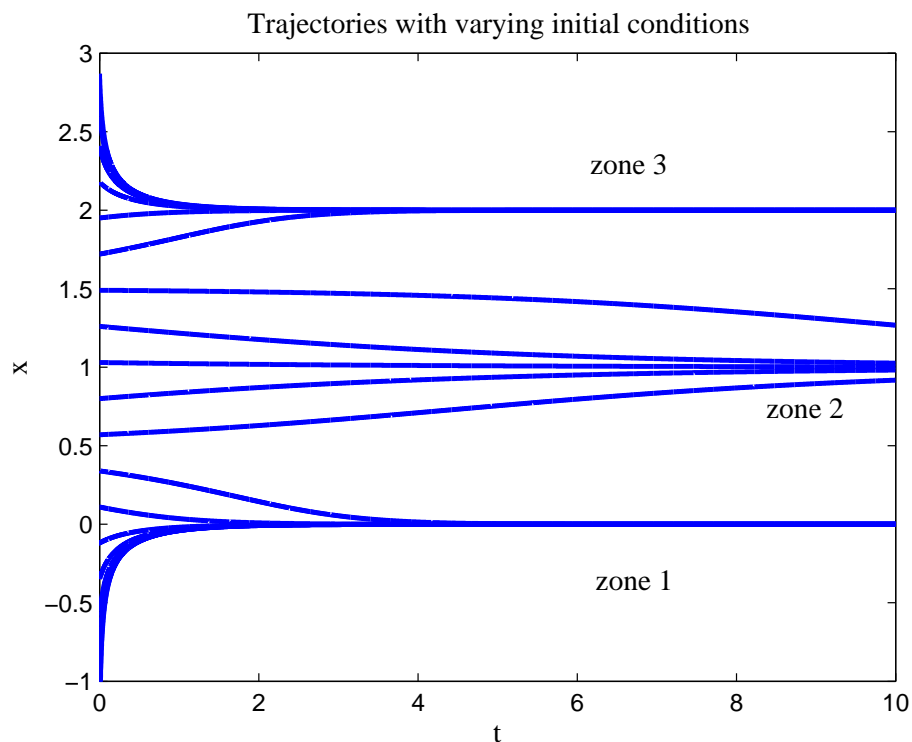


Fig. 2.

We now assume that the controller consists of an impulse of height d . The control task is to apply this impulse at a distinct time τ ; i.e.,

$$x(\tau^+) = x(\tau^-) + d.$$

Under the assumption that $x(0)$ is significantly smaller than a_1 , this controller will move the system from the domain of attraction for the first stable equilibrium point (Zone 1 in Figure 2) to Zone 2 given certain conditions on τ and d . However,

if the controller acts too quickly, this will not happen. On the other hand, if it acts too slowly, it may move the system to Zone 3 instead of Zone 2. As such, this system is robust to incorrectly measured time up to some level, after which the qualitative behavior of the system changes significantly.

This is shown in Figure 3, with

$$a_1 = 0, a_2 = 0.5, a_3 = 1, a_4 = 1.5, a_5 = 2, \\ x(0) = -2, d = 1.7.$$

In the figure, three different scenarios are depicted, in which $\tau = 0.005$ results in a system that asymptotically approaches a_1 , $\tau = 0.2$ results in $x \rightarrow a_3$, and $\tau = 1.1$ gives $x \rightarrow a_5$. Also depicted (with dotted lines as boundaries) are the regions where qualitatively different behaviors are obtained.




Fig. 3.

The conclusion to be drawn from this example is that, as opposed to the previous example, the controller is robust to small errors associated with the time sensor. However, the controller is not robust to arbitrarily large time measurement errors. This leads to the next example.

2.3 Strongly Open-Loop Stabilization

Consider the system

$$\dot{x} = (1 - u)A_1x + uA_2x,$$

in which $x \in \mathbb{R}^n$, $u \in \mathbb{R}$. Furthermore, assume that $\dot{x} = A_1x$ is unstable and $\dot{x} = A_2x$ is asymptotically stable. For any fixed τ , the control law

$$u(t) = \begin{cases} 0 & \text{if } t < \tau \\ 1 & \text{if } t \geq \tau \end{cases}$$

renders the controlled system asymptotically stable. This is indeed true because the controller will switch from the unstable to the stable subsystem at time τ . Moreover, it does not matter when this switch occurs, as long as it does occur, which means that time measurements are really not needed to implement the controller. This is an example of stabilization using a *strongly open-loop* controller, which is defined precisely in Section 5.

These simple examples illustrate interesting phenomena that should be carefully studied in increasing degrees of generality. We hope that this would lead to a greater understanding of control systems in general and the particular role that time should play. To begin this quest to improve our understanding, we introduce definitions in the remainder of the paper that permit great modeling flexibility in allowing time measurement uncertainty.

3 Basic Definitions

The definitions in this section build on standard control terminology and information space concepts that have previously been included into control formulations (e.g., [2]). Some concepts and terminology are also borrowed from robotics, particularly in the treatment of planning under sensing uncertainty ([12], Chapter 11). Readers familiar with various treatments of uncertainties in control systems will observe similarities with input-output representations, behavioral systems theory [15], quantized control systems [13], and imperfect state information in stochastic control [3, 11]. Throughout the coming presentation, the control systems are particularly inspired by robotics applications.

3.1 Inputs, Observations, and History

Let U and Y be the *input space* and *observation space*, respectively; i.e., the spaces in which the inputs and outputs take on values. It is now conceivable that the controllers and sensors are evolving asynchronously, or in other ways are driven by device-specific counters, which may be continuous. As such, we define index sets P and S , which parameterize the inputs and outputs, respectively. If we let $P = S = \mathbb{N}$, then we would obtain a classical discrete-time control model. However, we prefer to let P and S both be closed intervals of \mathbb{R} . For convenience, suppose $P = S = [0, \infty)$. This intuitively corresponds to having a continuous

index set that starts at zero and increases monotonically during execution. Let \tilde{u}_p denote a function, $\tilde{u}_p : [0, p] \rightarrow U$, called the *input history*, and let \tilde{y}_s denote a function, $\tilde{y}_s : [0, s] \rightarrow Y$, called the *observation history*. These histories are considered as partial functions on P and S , respectively, because $[0, p] \subset P$ and $[0, s] \subset S$.

The histories naturally “grow” during execution. For example, suppose during execution, \tilde{y}_s is obtained, followed by $\tilde{y}_{s'}$ at some later stage. We then have that $s' > s$, together with the requirement is that $\tilde{y}_{s'}(s'') = \tilde{y}_s(s'')$ for all $s'' \in S$.

Now, given input and output histories associated with a particular evolution of a control system, we let $\eta = (\tilde{u}_p, \tilde{y}_s)$ denote the *history I-state* (“I-state” is short for “information state”). Moreover, we let the *history I-space* be the set of all possible I-states, η , which includes any permissible \tilde{u}_p and \tilde{y}_s for any $p, s \in [0, \infty)$.

3.2 Interacting with a State Space

Now introduce a state space X . Let \tilde{x} denote a *state trajectory*, which is a time-parameterized function into X . We now need to define a transition function, and for this we let $\tilde{w}_t : [0, t] \rightarrow U$ be a time-parameterized *control function*. Let Φ denote the *state transition function*:

$$\tilde{x}(t' + t) = \Phi(\tilde{x}(t'), \tilde{w}_t).$$

For example, Φ could be defined in the usual way if X is a differentiable manifold with an associated *control system* $\dot{x} = f(x, u)$, in which $x \in X$ and $u \in U$. In this case, the state transition function Φ becomes:

$$\tilde{x}(t' + t) = \Phi(\tilde{x}(t'), \tilde{w}_t) = \tilde{x}(t') + \int_{t'}^{t'+t} f(\tilde{x}(\tau), \tilde{w}_t(\tau)) d\tau,$$

assuming appropriate integrability conditions.

3.3 The State-Time Space

Even though the way control signals interact with state variables above is not only well-known and supposedly unproblematic, we need to incorporate time more explicitly to understand the role of time in a more direct manner. For this, consider incorporating time into the state space definition to form the *state-time space*. Let $T \subseteq \mathbb{R}$ be the maximal *time interval* for which the system is defined. Let $Z = X \times T$ denote the *state-time space*. Elements of Z will be denoted as z or (x, t) .

We are comfortable with time-parameterized paths through X . However, what is not entirely clear is: What is the best way to parameterize paths through Z ? Consider a path $\tau : R \rightarrow Z$, in which R is a closed interval. For example, imagine rigid bodies moving in \mathbb{R}^3 according to the laws of mechanics. To show the motions of the bodies, we could parameterize τ in many ways, much in

the same way as varying parameters in a computer-generated animation. We could vary the speed, play parts of it backward, and so on. Since there are many possibilities, we would like to choose one that is most convenient for the coming formulations. Returning to the analogy of an animation, we assume that animations are played forward in “real time”. Rather than allow any arbitrary path τ , this means that we require that $d\tau/dr = 1$ (for the case in which Z is a differentiable manifold), in which $r \in R$ is the path parameter.

Since time progresses forward monotonically, any path τ could be reparameterized using t , even though t is a coordinate of Z . This would be convenient, but note that it is somewhat abusive because t is serving both as a coordinate of Z and the parameter of a path, τ . Nevertheless, we assume that all paths through Z are parameterized by time t and are referred to as *(state-time) trajectories*. A trajectory is denoted by \tilde{z} and refers to a mapping $\tilde{z} : [t_1, t_2] \rightarrow Z$ in which t_1 and t_2 are the starting and ending times, respectively.

If the dimension of X is n , then let z_{n+1} denote the last component of z , i.e., the time component. Under the observation that $\dot{z}_{n+1} = 1$, it is straightforward to convert a system $\dot{x} = f(x, u)$ into $\dot{z} = f'(z, u)$ by simple extension. In general, Φ can be extended in the straightforward way to obtain Φ' .

A trajectory \tilde{z} is called *valid* if there exists some control function \tilde{w} such that $\dot{z} = f'(z, u)$ is satisfied for all times along \tilde{z} . Let \tilde{Z} denote the space of all valid trajectories.

3.4 Initial Conditions on the State-Time Space

Once the transition has been made to viewing time as something that must be measured and is part of the state-time variable, time-dependent initial conditions become important. In fact, we would like to express various initial conditions on Z ; therefore, the history I-state will be expanded to include the initial conditions.

Let $Z_0 \subseteq Z$ denote a given *initial condition*, which indicates the set of possible initial state-times. The expanded history I-state is defined as $\eta = (Z_0, \tilde{u}_p, \tilde{y}_s)$.

Three important special cases of initial conditions are:

1. **Known initial state-time:** Z_0 is a singleton, $\{(x_0, t_0)\}$, which means that the initial time is t_0 and the initial state is x_0 . Thus, (6) considers only trajectories for which $\tilde{z}(t_0) = (x_0, t_0)$.
2. **Known initial time:** Suppose it is known that t_0 is the starting time. In this case,

$$Z_0 = \{(x, t) \in Z \mid t = t_0\}. \quad (1)$$

3. **Known initial state:** Suppose it is known that x_0 is the starting state. In this case,

$$Z_0 = \{(x, t) \in Z \mid x = x_0\}. \quad (2)$$

Whenever the initial time is given, we will typically have $t_0 = 0$.

3.5 Sensor Mappings

A sensor is defined in terms of a mapping from Z (or trajectories on Z) into Y , the observation space. A sensor may be either: 1) *instantaneous*, which means that the observation depends only on the current state-time, 2) *history-based*, which means that the observation may depend on any portion of the entire state-time trajectory.

An instantaneous *sensor mapping* is defined as $h : Z \rightarrow Y$, for which we write $y = h(z)$ or $y = h(x, t)$. For a history-based sensor mapping, we have $h : \tilde{Z} \rightarrow Y$, and an observation is given as $y = h(\tilde{z})$.

Several simple and important sensors are now defined.

1. **(Perfect Information)** A *perfect state-time sensor* is defined as any injective mapping $h : Z \rightarrow Y$ because (x, t) can be recovered from any observed y . The simplest case is $Y = Z$ and $y = h(z) = z$.
2. **(State-Only)** A *perfect state sensor* can be defined, for example, as $y = h(x, t) = x$. In this case, we may know the current state but remain uncertain about the particular time.
3. **(Clock)** A *perfect time sensor* or (perfect) *clock* is defined as $y = h(x, t) = t$.
4. **(Chronometer)** An important history-based sensor is the *perfect time odometer* or (perfect) *chronometer*, which yields the total time elapsed for a state-time trajectory. Let $\tilde{z}_{t,t'}$ denote a valid space-time trajectory with endpoints t and t' , and $t < t'$. The chronometer is defined as $h(\tilde{z}_{t,t'}) = t' - t$.

3.6 Disturbances

In this section, we briefly illustrate how to incorporate disturbances into Φ and h ; however, to simplify the presentation, we will not include such disturbances in the subsequent sections.

In the case of smooth manifolds, f can be extended to obtain $\dot{x} = f(x, u, \theta)$ in which θ is selected from Θ , a set of possible disturbances (or nature inputs).³ In general, the disturbance $\theta \in \Theta$ can be incorporated into Φ to obtain $\tilde{x}(t'+t) = \Phi(\tilde{x}(t'), \tilde{w}_t, \theta)$.

Another disturbance parameter can be defined, to interfere with sensors. For example, $y = h(z, \psi)$, in which $\psi \in \Psi$ and Ψ is an *observation disturbance space*. Imperfect versions of previously defined sensors can be made. For example, an *imperfect clock* is defined as $y = h(x, t, \psi) = t + \psi$ and $\Psi = [-\epsilon, \epsilon]$, in which $\epsilon > 0$ represents the maximum error in the time measurement.

4 Nondeterministic I-spaces

In this section, we introduce an *information mapping* and *derived* I-space, which means that history I-states are mapped into a new I-space that provides some

³ Disturbances in the time direction seem absurd and will therefore not be considered.

interpretation or aggregation of the histories (see [12], Chapter 11, for examples). One possibility is to map η to a posterior pdf $p(z|\eta)$, which would lead to probabilistic I-states (these appear in stochastic control, but usually over X rather than Z). Although there are many possibilities, we exclusively consider a nondeterministic interpretation of the histories: The smallest subset of Z that is consistent with a history I-state η . A set-valued information mapping is thus defined, and is denoted as $Z(\eta) \subseteq Z$. The target of this mapping is a *nondeterministic I-space*, $\mathcal{I}_{ndz} = \text{pow}(Z)$.⁴

4.1 Relating internal parameters to time

The first step is to define the class of mappings that relate the index sets P and S to time. It is assumed that the particular mapping is unknown; otherwise, the precise time could be reconstructed if the mapping is injective. Instead we define Ω as the set of possible mappings from T to P . Similarly, let Λ denote a set of possible mappings from T to S . Although the particular mapping is not given, we assume that the sets Ω and Λ are specified in advance.

Many reasonable definitions are possible for Ω and Λ . Consider defining Ω (the same possibilities exist for Λ). One of the weakest sensible definitions is that Ω contains any mapping $\omega : T \rightarrow P$ for which $\omega(t)$ monotonically increases. This at least ensures that a higher index implies later time.

Another possibility is to restrict Ω to differentiable functions and require bounded derivatives and bounded initial error. For example, each $\omega \in \Omega$ must satisfy $|\omega(0)| \leq c_0$ and $|d\omega/dt| \leq c_1$ for some positive constants c_0 and c_1 . This restricts the possible times to an interval that widens as time increases. Without assuming differentiability, a similar function space could be obtained using Lipschitz constants instead of bounded derivatives.

4.2 The Sensorless Case

We now define the mapping from history I-states into the nondeterministic I-space, \mathcal{I}_{ndz} . Assume that Ω and Λ contain only invertible functions (otherwise, the inverses below can be replaced by preimages to obtain slightly more complicated definitions). First consider the sensorless case, in which $\eta = (Z_0, \tilde{u}_p)$. The nondeterministic I-state (smallest consistent subset of Z) given η is

$$\begin{aligned} Z(\eta) = \{ (x, t) \in Z \mid \exists \tilde{z} \in \tilde{Z} \text{ and } \exists \omega \in \Omega \text{ such that} \\ \tilde{z}(t_1) \in Z_0 \text{ and } \forall t \in [t_1, t_1 + \omega^{-1}(p)), \\ \tilde{z}(t_1 + t) = \Phi'(\tilde{z}(t_1), \tilde{u}_p(\omega(t))) \}, \end{aligned} \quad (3)$$

in which t_1 refers to the starting time of \tilde{z} .

⁴ In most contexts, most elements of $\text{pow}(Z)$ are unnecessary; however, it is simpler to define $\text{pow}(Z)$ than to worry about the precise reachable subset of $\text{pow}(Z)$ (which is an interesting research problem in itself!).

4.3 The Inputless Case

Now consider the case in which $\eta = (Z_0, \tilde{y}_s)$; \tilde{z} is thus completely predictable from any initial $z \in Z$. We have

$$\begin{aligned} Z(\eta) = \{ & (x, t) \in Z \mid \exists \tilde{z} \in \tilde{Z} \text{ and } \exists \lambda \in \Lambda \text{ such that} \\ & \tilde{z}(t_1) \in Z_0 \text{ and } \forall t \in [t_1, t_1 + \lambda^{-1}(s)], \\ & h(\tilde{z}(t)) = \tilde{y}_s(\lambda(t))\}. \end{aligned} \quad (4)$$

4.4 Combining Sensors and Inputs

Now consider the case in which an observation history \tilde{y}_s is also given, yielding $Z(\eta) = Z(Z_0, \tilde{u}_p, \tilde{y}_s)$. The nondeterministic I-state in this case combines the constraints from both (3) and (4), to obtain:

$$\begin{aligned} Z(\eta) = \{ & (x, t) \in Z \mid \exists \tilde{z} \in \tilde{Z}, \exists \omega \in \Omega, \text{ and } \exists \lambda \in \Lambda \\ & \text{such that } \tilde{z}(t_1) \in Z_0 \text{ and} \\ & \forall t \in [t_1, t_1 + \omega^{-1}(p)], \\ & \tilde{z}(t_1 + t) = \Phi^t(\tilde{z}(t_1), \tilde{u}(\omega(t))) \text{ and} \\ & \forall t \in [t_1, t_1 + \lambda^{-1}(s)], h(\tilde{z}(t)) = y(\lambda(t))\}. \end{aligned} \quad (5)$$

5 Defining Control Laws

A wide variety of control laws may be defined in terms of *information feedback*. We could define a control law as a mapping from the history I-space into the input space U , but this would be difficult to manage. Therefore, we consider control laws that map from a derived I-space into U . In particular, we consider in this section control laws of the form $\gamma : \mathcal{I} \rightarrow U$, in which \mathcal{I} is a particular derived I-space \mathcal{I} that is obtained as a mapping from the history I-space. Many possibilities exist; for example, \mathcal{I} may be Z , X , T , P , or \mathcal{I}_{ndz} .

5.1 Strongly open-loop control

Consider sensorless information states of the form $\eta = \tilde{u}_p$ (the initial state-time could be any $z \in Z$). Let the derived I-space $\mathcal{I} = P$ be defined by the information mapping $\tilde{u}_p \mapsto p$. A *strongly open-loop* control law is defined as $\gamma : P \rightarrow U$.

The trajectory obtained by applying \tilde{u}_p could be any \tilde{z} for which there exists an $\omega \in \Omega$ such that $\tilde{z} : [t_1, t_1 + \omega^{-1}(p)] \rightarrow Z$ and

$$\tilde{z}(t_1 + t) = \tilde{z}(t_1) + \int_{t_1}^{t_1+t} f'(\tilde{x}(t'), \tilde{u}(\omega(t'))) dt' \quad (6)$$

for all $t \in [0, \omega^{-1}(p)]$.

Now suppose there is an initial condition $Z_0 \subseteq Z$ so that $\eta = (Z_0, \tilde{u}_p)$. In this case (6) is constrained to consider only those $\tilde{z} \in Z$ for which $\tilde{z}(t_1) \in Z_0$. The last example of Section 2 represents strongly open-loop stabilization, even when $Z_0 = Z$ and Ω contains all time-monotonic functions.

5.2 Perfect time-feedback control

To obtain *perfect time-feedback control* (otherwise classically known as “open loop” control), consider the derived I-space, $\mathcal{I} = T$, in which the exact time t can be derived from η . In this case, the information mapping is $(Z_0, \tilde{u}_p, \tilde{y}_s) \mapsto t$. The most common case occurs with a perfect time sensor, as defined in Section 3.5. The control law is specified as $\gamma : T \rightarrow U$. The most common special case is when Z_0 yields a known initial space-time $(x_0, 0)$.

5.3 Imperfect time-feedback control

Now consider a derived I-space $\mathcal{I} = \text{pow}(T)$, in which a derived I-state is interpreted as the smallest set of times that are consistent with η . A control law is expressed as $\gamma : \text{pow}(T) \rightarrow U$ (we could restrict the domain to consider only reasonably behaved subsets of T , such as closed intervals). In terms of information requirements, this control law lies somewhere in between strongly open-loop control and perfect time-feedback control.

5.4 Other control laws

Many other laws may be considered. If it is possible to reconstruct z from η , then *perfect state-time feedback* is possible, defined as $\gamma : Z \rightarrow U$. In this case Z is considered as a derived I-space, $\mathcal{I} = Z$. Note that *perfect state feedback*, $\gamma : X \rightarrow U$ may also be considered, in which case the precise times are unnecessary (a known benefit of state-feedback control). Finally, using (5), we can define feedback on the nondeterministic I-space, yielding $\gamma : \mathcal{I}_{ndz} \rightarrow U$. This means that subsets of Z are mapped into U .

6 Open Questions and Issues

Using information space concepts, we have proposed new ways to formulate uncertainty in time measurements. We called classical open-loop control “perfect time-feedback control” and introduced the notion of *strongly* open-loop control, which is robust with respect to massive time distortions. We provided several examples that illustrate the associated issues.

We are fascinated by the numerous exciting questions and issues raised by the examples and formulations developed in this paper, which is intended to open doors to new problems, rather than close them with particular results. Here are some points worth considering:

1. What classes of systems support strongly open-loop stabilization? Stability might be a property too strong to demand of an open-loop controller with imperfect time measurements. Instead, other interesting properties such as state containment or invariance should also be studied.

2. Can local (small time perturbations), in contrast to global strongly open-loop control, be characterized in a meaningful manner? Is there a notion of robustness to time perturbations that can be used for designing control laws for robotic systems that are locally immune to time perturbations?
3. Is there a meaningful notion of time observers? In other words, can we estimate t from the history I-state $\eta = (Z_0, \tilde{u}_p, \tilde{y}_s)$? Also, how would an open-loop controller interact with the plant based on the time estimate?
4. What are the relationships between history-based sensors and histories of observations obtained from instantaneous sensors? For example, by measuring angular velocity and having a perfect chronometer, we can simulate an angular odometer. If the initial angle is given, then we can furthermore simulate a compass. Furthermore, a variety of imperfect versions can be made by replacing the perfect chronometer with a weaker sensor.
5. What other ways can we relate S , P , and T ? So far, we related S and P to T via two mappings. We may instead want, for example, to relate S to P , and then P to T . We might even want to consider mappings that are not invertible.
6. Can useful control laws be defined over \mathcal{I}_{ndz} ? These would choose actions based on particular subsets of $X \times T$.
7. Can control laws that are robust with respect to severe time uncertainty lead to improved approaches to distributed, asynchronous control?

References

1. K. J. Åström and B.M. Bernhardsson. Comparison of riemann and lebesgue sampling for first order stochastic systems. In *Proceedings IEEE Conference Decision and Control*, pages 2011–2016, 2002.
2. T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory, 2nd Ed.* Academic, London, 1995.
3. D. P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. I, 2nd Ed.* Athena Scientific, Belmont, MA, 2001.
4. M. Boccadoro, Y. Wardi, M. Egerstedt, and E. Verriest. Optimal control of switching surfaces in hybrid dynamical systems. *Journal of Discrete Event Dynamic Systems*, 15(4):433–448, December 2005.
5. A. E. Bryson and Y.-C. Ho. *Applied Optimal Control.* Hemisphere Publishing Corp., New York, 1975.
6. M. Egerstedt, X. Hu, and A. Stotsky. Control of mobile platforms using a virtual vehicle approach. *IEEE Transactions on Automatic Control*, 46(11):1777–1782, November 2001.
7. Wikipedia: The Free Encyclopedia. Chronometer. <http://en.wikipedia.org/wiki/Chronometer>, February 2007.
8. A. Ganguli, J. Cortés, and F. Bullo. Distributed deployment of asynchronous guards in art galleries. In *Proceedings American Control Conference*, pages 1416–1421, 2006.
9. J.K. Holm, D. Lee, and M.W. Spong. Time scaling for speed regulation in bipedal locomotion. In *Proc. IEEE International Conf. on Robotics and Automation*, 2007.

10. H. W. Kuhn. Extensive games and the problem of information. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, pages 196–216. Princeton University Press, Princeton, NJ, 1953.
11. P. R. Kumar and P. Varaiya. *Stochastic Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1986.
12. S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006. Also available at <http://planning.cs.uiuc.edu/>.
13. D. Liberzon and J. P. Hespanha. Stabilization of nonlinear systems with limited information feedback. *IEEE Transactions on Automatic Control*, 50(6):910–915, June 2005.
14. G. J. Pappas. Avoiding saturation by trajectory reparameterization. In *Proceedings IEEE Conference Decision and Control*, 1996.
15. J. W. Polderman and J. C. Willems. *Introduction to Mathematical Systems Theory: A Behavioral Approach*. Springer-Verlag, New York, 1998.