Time-changed Linear Quadratic Regulators

Andrew Lamperski and Noah J. Cowan

***Abstract*— Many control methods implicitly depend on the assumption that time is accurately known. For example, the finite-horizon linear quadratic regulator is a linear policy with time-varying gains. Such policies may be infeasible for controllers without accurate clocks, such as the motor systems in humans and other animals, since gains would be applied at incorrect times. Little appears to be known, however, about control with imperfect timing. This paper gives a solution to the linear quadratic regulator problem in which the state is perfectly known, but the controller's measure of time is a stochastic process derived from a strictly increasing Lévy process. The optimal controller is linear and can be computed from a generalization of the classical Riccati differential equation.**

I. Introduction

This paper studies the finite-horizon linear quadratic regulator (LQR) with temporal uncertainty. Solutions to finite-horizon optimal control problems are typically time-varying feedback policies. Implicitly, it is assumed that the controller has perfect knowledge of time, even in cases of imperfect state information. This paper examines the consequences of assuming that controller's measure of time is stochastic.

A stochastic process can be *time-changed* by replacing its time index, *objective time*, by a monotonically increasing stochastic process, *subjective time* [1]. Time-changed stochastic processes arise in finance, since changing the time index to a measure of economically relevant events, such as trades, can improve modeling [2]–[4]. This new time index is, however, stochastic with respect to calendar time.

Similar notions of objective time and subjective time arise in the study of time estimation in the nervous system. Human timing is subject to neural noise and environmental perturbation [5]. Furthermore, humans rationally exploit the statistics of their temporal noise during simple timed movements, such as button pushing [6] and pointing [7]. To analyze more complex movements, a theory of feedback controllers that compensate for temporal noise is desirable.

The study of control systems with temporal uncertainty is, however, relatively unexplored. Attention has been devoted to the problem of clock synchronization in distributed systems [8], [9]. Less effort has been devoted studying the impact of asynchronous clock behavior on common control issues, such as stability [10] and optimal performance [11]. Only a limited amount of work focuses on general consequences of temporal uncertainty for control [12], [13].

This paper focuses on the linear quadratic regulator problem with perfect state information, but a stochastically time-changed control process. The optimal cost-to-go function is derived by dynamic programming, based on a generalization of the classical LQR Riccati differential equation. The results apply to a wide class of stochastic time changes given by strictly increasing Lévy processes.

Section II reviews the classical linear quadratic regulator, and demonstrates how temporal uncertainty can lead to poor behavior. Next, in Section III, basic ideas from Lévy processes are reviewed, and the class of temporal noise models is defined. The problem of interest and its solution are stated in Section IV, while the solution is derived in Section V. Finally, a conclusion is given in Section VI.

II. Classical LQR and Time

The classical linear quadratic regulator chooses inputs $u_t = g_t(x_t)$ that minimize the quadratic integral

$$\int_0^{t_f} \left(x_t^T Q x_t + u_t^T R u_t\right) dt + x_{t_f}^T \Phi x_{t_f},$$

subject to linear dynamics

$$\dot{x}_t = A x_t + B u_t$$

and initial condition $x_0 = x$. Here Q and Φ are positive semidefinite, while R is positive definite.

The optimal solution is given by

$$u_t = L_t x_t, \tag{1}$$

where

$$L_t = -R^{-1} B^T P_t, \tag{2}$$

and P_t satisfies the backward differential equation

$$-\dot{P}_t = Q + A^T P_t + P_t A - P_t B R^{-1} B^T P_t, \tag{3}$$

with final condition $P_{t_f} = \Phi$.

The linear quadratic regulator requires perfect knowledge of the state and, perhaps more subtly, perfect knowledge of time. While linear quadratic control with imperfect state information is commonly studied, little seems to be known about optimal control with imperfect time information. The following example shows that using the standard LQR policy with noisy time measurement can lead to poor behavior.

Example 1: Consider the system defined by the state matrices

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \qquad B = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

with cost matrices given by $\Phi = I$, $Q = 0$, and $R = 1$.

A. Lamperski is with the Department of Engineering, University of Cambridge, UK a.lamperski@eng.cam.ac.uk

N. J. Cowan is with the Department of Mechanical Engineering, Johns Hopkins University, USA ncowan@jhu.edu

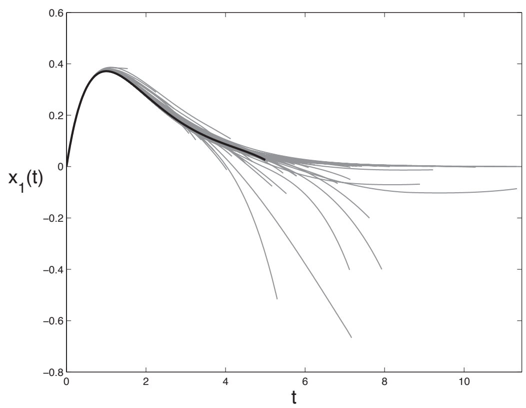

Fig. 1. Plots of the first coordinate, $x_1(t)$, with initial condition $x(0) = [0, 1]^T$, and $t_f = 5$. The black line depicts the trajectory for the LQR policy, (1), while the gray lines are trajectories under the policy (5) for 100 realizations of ζ_t. The final times under (5) varied from 2 to 11, and some trajectories deviate significantly from the nominal trajectory because gains are applied at the incorrect times.

Now say that instead of knowing time t, the controller only has access to a noisy measure of time is given by

$$\zeta_t = \sup\{s + B_s : 0 \le s \le t\}, \tag{4}$$

where B_t is the standard unit Brownian motion. The process, ζ_t, can be interpreted as a noisy clock. A realization of ζ_t is plotted in Figure 2(b), cf. Example 4.

Figure 1 shows the result of using the policy

$$u(\zeta_t, x_t) = L_{\zeta_t} x_t \tag{5}$$

over the time horizon $\{t : \zeta_t \in [0, t_f]\}$. Intuitively, the controller uses ζ_t as though it were the correct clock reading, but this can result in a suboptimal gain when $\zeta_t \ne t$. Note that the time horizon varies with the realization of ζ_t.

For some choices of A, B, Q, R, and Φ, the policy, (5), can perform arbitrarily worse than (1).

This paper generalizes (1), (2), and (3) in order to optimally compensate for uncertainty generated from a relatively large class of temporal noise models.

III. Lévy Processes and Time Changes

This section presents the class of temporal noise models used in this paper. Two notions of time are used: *subjective time* and *objective time*. Subjective time is modeled as a stochastic process parametrized by objective time. The noisy clock process from (4) is an example of a subjective time model.

A. Background on Lévy Processes

Basic notions from Lévy processes required to define the general class of subjective time models are now reviewed. The definitions and results can be found in [14].

A real-valued stochastic process X_t is called a *Lévy process* if

- $X_0 = 0$ almost surely (a.s.).
- X_t has independent, stationary increments: If $0 \le s \le t$, then X_s and $X_t - X_s$ are independent and $X_t - X_s$ has the same distribution as X_{t-s}.
- X_t is stochastically continuous: For all $a > 0$ and all $s \ge 0$, $\lim_{t\to s} \mathbb{P}(|X_t - X_s| > a) = 0$.

It will be assumed that Lévy processes in this paper are right-continuous with left-sided limits, i.e. they are *càdlàg*. No generality is lost since, for every Lévy process, X_t, there is a càdlàg Lévy process, $\tilde{X}_t$, such that $X_t = \tilde{X}_t$ for almost all t.

A monotonically increasing Lévy process is called a *subordinator*. The following theorem, adapted from [14], implies that every subordinator can be uniquely characterized by a function on $\mathbb{C}$.

Theorem 1: *Let X_t be a subordinator. There exists a unique complex-valued function ψ, called the* Laplace exponent *of X_t, such that for all $z \in \mathbb{C}$ with $\operatorname{Re} z > 0$, ψ is analytic at z and*

$$\mathbb{E}\left[e^{-zX_t}\right] = e^{-t\psi(z)}. \tag{6}$$

The function, ψ, is called the Laplace exponent because (6) is the Laplace transform of the distribution of X_t. A converse to Theorem 1 holds, but constraints on the form of ψ are required. Equation (6) is derived from the famous Lévy-Khintchine formula for Lévy processes.

So far, the discussion in this section has been completely standard. This paper will employ (6) for the less common task of computing expectations of matrix exponentials, $\mathbb{E}\left[e^{AX_t}\right]$, where X_t is a subordinator. To reason about unstable A, the definition of the Laplace exponent must be extended, when possible, to values of z with $\operatorname{Re} z \le 0$.

Recall that ψ is always analytic at z with $\operatorname{Re} z > 0$. It could, however, be analytic on a larger region. Define $r_{\min} \in [-\infty, 0]$ to be the minimal value r such that ψ is analytic at all z with $\operatorname{Re} z > r$ and define the *domain of* ψ, $\operatorname{dom}(\psi)$, as

$$\operatorname{dom}(\psi) = \{z \in \mathbb{C} : \operatorname{Re} z > r_{\min}\}. \tag{7}$$

It can be shown that (6) holds on all of $\operatorname{dom}(\psi)$.

Example 2: The simplest non-trivial subordinator is the Poisson process N_t, which is characterized by

$$\mathbb{P}(N_t = k) = e^{-\gamma t}\frac{(\gamma t)^k}{k!},$$

where $\gamma > 0$ is called the rate constant. Its Laplace exponent is given by $\psi(z) = \gamma - \gamma e^{-z}$, which is found by computing the expected value directly. It follows that $r_{\min} = -\infty$ and $\operatorname{dom}(\psi) = \mathbb{C}$.

B. Subjective Time Models

Let τ_s be a strictly increasing subordinator. In other words, if $s < s'$ then $\tau_s < \tau_{s'}$ a.s. (Note that any subordinator can be made to be strictly increasing by adding a drift term bs with $b > 0$.) The process τ_s will have the interpretation of being the amount of objective time that has passed over s units of subjective time. The process ζ_t will be an inverse process that describes how much subjective time has passed over t units of objective time. Formally, ζ_t is defined by

$$\zeta_t = \inf\{\sigma : \tau_\sigma \geq t\}.$$

Note that $\zeta_{\tau_s} = s$ a.s. Indeed, $\zeta_{\tau_s} = \inf\{\sigma : \tau_\sigma = \tau_s\}$, by definition. Since τ_s is right continuous and strictly increasing, a.s., it follows that $\zeta_{\tau_s} = s$, a.s.

Example 3: The case of no temporal uncertainty corresponds to $\tau_s = s$ and $\zeta_t = t$. The Laplace exponent of τ_s is computed directly as $\psi(z) = z$ and $\mathrm{dom}(\psi) = \mathbb{C}$.

Example 4: A more interesting temporal noise model is the inverse Gaussian subordinator. Fix $\gamma > 0$ and $\delta > 0$. Let $C_t = \gamma t + B_t$, where B_t is a standard unit Brownian motion. The inverse Gaussian subordinator is given by

$$\tau_s = \inf\{t : C_t = \delta s\},$$

with Laplace exponent $\psi(z) = \delta(\sqrt{\gamma^2 + 2z} - \gamma)$. Here $\mathrm{dom}(\psi)$ corresponds to $\mathrm{Re}\, z > -\gamma^2/2$. Derivation of ψ, in this case, relies on techniques outside the scope of this paper. See [14].

It can be shown that the inverse process is given by

$$\zeta_t = \sup\left\{\delta^{-1} C_\sigma : 0 \leq \sigma \leq t\right\}.$$

See Figure 2. Note that when $\gamma = \delta = 1$, the subjective time ζ_t becomes the noisy clock process, (4).

In the preceding example, the process τ_s has jumps, but the inverse, ζ_t, is continuous. The next proposition generalizes this observation for any strictly increasing subordinator, τ_s.

Proposition 1: *The process ζ_t is continuous almost surely.*

Proof: Fix $\epsilon > 0$ and $t \geq 0$. Set $s = \zeta_t$. Strict monotonicity of τ_s implies that $[\tau_{\max\{s-\epsilon,0\}}, \tau_{s+\epsilon}]$ is a nonempty interval, a.s. The inverse property of ζ_t implies (almost surely) that $t \in [\tau_{\max\{s-\epsilon,0\}}, \tau_{s+\epsilon}]$ and $\zeta_{t'} \in [\max\{s-\epsilon, 0\}, s+\epsilon]$ for all $t' \in [\tau_{\max\{s-\epsilon,0\}}, \tau_{s+\epsilon}]$. ■

IV. Problem Statement and Main Result

Consider a linear system of the form

$$\dot{y}_t = A y_t + B u(\zeta_t). \tag{8}$$

This system describes a scenario in which dynamics of a physical plant are described by motion in objective time, but the agent choosing u cannot directly measure time, but instead must rely on its subjective measure ζ_t.

Let $x_s = y_{\tau_s}$ be the time-changed process. Note that (8) can be solved as

$$y_t = e^{At} y_0 + e^{At} \int_0^t e^{-A\sigma} B u(\zeta_\sigma) d\sigma.$$

It follows that x_s can be computed as

$$x_s = e^{A\tau_s} x_0 + e^{A\tau_s} \int_0^{\tau_s} e^{-A\sigma} B u(\zeta_\sigma) d\sigma. \tag{9}$$

Note that x_s depends on the values of $u(r)$ for $r \in [0, s)$. For compactness, $u(r)$ will often be written as u_r. Let

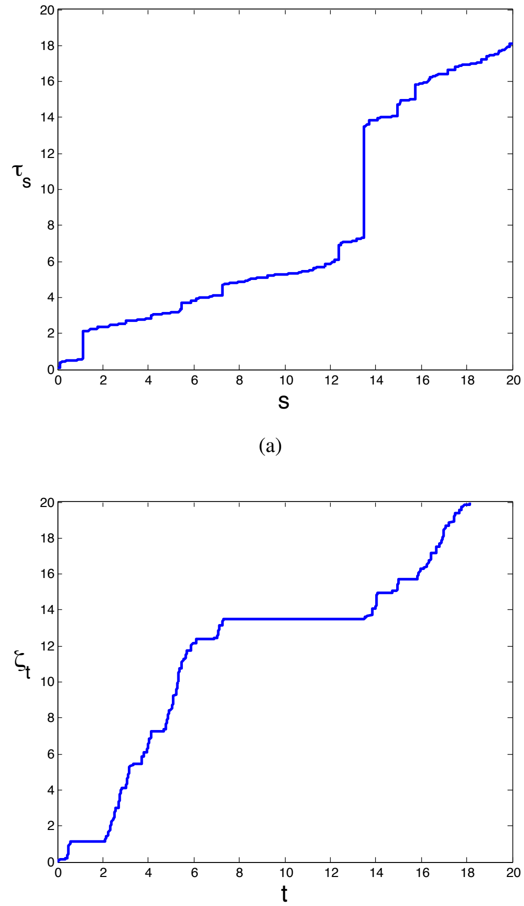

(a)

(b)

Fig. 2. 2(a). The inverse Gaussian subordinator, τ_s, with $\gamma = \delta = 1$. The process was simulated by generating independent inverse Gaussians using the method from [15]. 2(b). The inverse process, ζ_t. Note that the graph of ζ_t can be found from the graph of τ_s by simply switching the axes.

$\mathcal{F} = \{\mathcal{F}_s\}_{s \geq 0}$ be the natural filtration associated with x_s. A control policy is said to be *admissible* if u_s is measurable in $\mathcal{F}_s$ and right-continuous with left limits.

Let $[0, S]$ be a finite interval of subjective time. The time-changed linear quadratic regulator is the admissible control policy that minimizes the expected quadratic cost

$$\mathbb{E}\left[\int_0^S \left(x_s^T Q x_s + u_s^T R u_s\right) ds + x_S^T \Phi x_S\right].$$

Here it is assumed that Q and Φ are positive semidefinite, while R is positive definite.

Given a policy u_s, the *expected cost-to-go* is defined as

$$J_s(x) = \mathbb{E}\left[\int_s^S \left(x_r^T Q x_r + u_r^T R u_r\right) dr + x_S^T \Phi x_S \,\middle|\, x_s = x\right]. \tag{10}$$

It follows that the regulator problem involves choosing the policy that minimizes $J_0(x)$.

The solution to the time-changed regulator problem uses a Riccati-like differential equation whose terms depend on the linear mappings described in the following lemma. The lemma is proved in Section V-B.

Lemma 1: *If the all eigenvalues of* $-2A$ *lie in* $\mathrm{dom}(\psi)$, (7), *then there are linear mappings,* F, G, *and* H *such that*

$$\mathbb{E}\left[e^{A^T\tau_s}Ye^{A\tau_s}\right] = Y + sF(Y) + O(s^2)$$
$$\mathbb{E}\left[e^{A^T\tau_s}Y\int_0^{\tau_s}e^{A\sigma}d\sigma\right] = sG(Y) + O(s^2)$$
$$\mathbb{E}\left[\int_0^{\tau_s}e^{A^T\sigma}d\sigma Y\int_0^{\tau_s}e^{A\rho}d\rho\right] = sH(Y) + O(s^2).$$

The eigenvalue assumption is made so that the equation

$$\mathbb{E}\left[e^{(\overline{\lambda}_1+\lambda_2)\tau_s}\right] = e^{-s\psi(-\overline{\lambda}_1-\lambda_2)}$$

holds whenever λ_1 and λ_2 are eigenvalues of A. Note that if A is Hurwitz, then the eigenvalue requirement always holds.

Example 5: With no temporal noise, the mappings reduce to

$$F(Y) = A^TY + YA, \quad G(Y) = Y, \quad H(Y) = 0. \tag{11}$$

Furthermore, since ψ is analytic everywhere, these formulas are true for any state matrix, A.

Example 6: Consider an arbitrary strictly increasing subordinator with Laplace exponent ψ. Let $A = \lambda$ where λ is a real, non-zero scalar with $-2\lambda \in \mathrm{dom}(\psi)$. Combining (6) with the formula $\int_0^t e^{\lambda\sigma}d\sigma = \lambda^{-1}(e^{\lambda t}-1)$ shows that

$$\mathbb{E}\left[e^{\lambda\tau_s}e^{\lambda\tau_s}\right] = e^{-s\psi(-2\lambda)}$$
$$\mathbb{E}\left[e^{\lambda\tau_s}\int_0^{\tau_s}e^{\lambda\sigma}d\sigma\right] = \lambda^{-1}(e^{-s\psi(-2\lambda)} - e^{-s\psi(-\lambda)})$$
$$\mathbb{E}\left[\int_0^{\tau_s}e^{\lambda\sigma}d\sigma\int_0^{\tau_s}e^{\lambda\rho}d\rho\right] = \lambda^{-2}(e^{-s\psi(-2\lambda)} - 2e^{-s\psi(-\lambda)} + 1).$$

It follows that

$$F(Y) = -\psi(-2\lambda)Y$$
$$G(Y) = \lambda^{-1}(\psi(-\lambda) - \psi(-2\lambda))Y$$
$$H(Y) = \lambda^{-2}(2\psi(-\lambda) - \psi(-2\lambda))Y.$$

Applying similar ideas to state matrices in Jordan form, Lemma 1 can be viewed as a generalization of this example.

With all the required definitions now given, the main theorem of the paper can be stated. The theorem is proved in Section V.

Theorem 2: *Let* τ_s *be a strictly increasing subordinator with Laplace exponent* ψ. *Assume that the eigenvalues of* $-2A$ *are contained in* $\mathrm{dom}(\psi)$ *and* $0 \in \mathrm{dom}(\psi)$. *The optimal expected cost-to-go is of the form* $J_s(x) = x^TY_sx$, *where* Y_s *solves the backward differential equation*

$$-\frac{d}{ds}Y_s = Q + F(Y_s) - G(Y_s)B(R + B^TH(Y_s)B)^{-1}B^TG(Y_s)^T, \tag{12}$$

with final condition $Y_S = \Phi$. *The optimal policy is given by*

$$u_s = K_sx_s,$$

where the gain K_s *is defined by*

$$K_s = -(R + B^TH(Y_s)B)^{-1}B^TG(Y_s)^T. \tag{13}$$

The optimal controller compensates for the temporal uncertainty, since the mappings F, G, and H, from Lemma 1, depend on the distribution of τ_s. Thus, different subjective time models will lead to different optimal controllers.

Example 7: With no temporal noise, by plugging in the matrices from (11), the gain and cost-to-go reduce to (2) and (3), respectively. Thus, the time-changed linear quadratic regulator generalizes the classical linear quadratic regulator.

The following theorem is useful for comparing the expected costs for different linear policies. It will not be proved, since the main ideas of its derivation are contained in the proof of Theorem 2.

Theorem 3: *If the control input is a linear policy,* $u_s = L_sx_s$, *then the expected cost-to-go is given by* $J_s(x) = x_s^TP_sx_s$, *where* P_s *satisfies the backward differential equation*

$$-\frac{d}{ds}P_s = Q + F(P_s) + L_s^TB^TG(P_s)^T + G(P_s)BL_s + L_s^T(R + B^TH(P_s)B)L_s, \tag{14}$$

with final condition $P_S = \Phi$.

Note that the optimal gain, (13), is recovered by taking the derivative of (14) with respect to L_s and setting to 0. Theorem 2 is not simply a special case of Theorem 3, however, since it gives the optimal input over all $\mathcal{F}_s$-measurable càdlàg policies. Theorem 3, on the other hand, assumes a-priori that the policy is linear.

Example 8: Consider the unstable scalar system described by $A = 1/6$ and $B = 1$ and cost matrices $Q = R = \Phi = 1$. Say that τ_s is the inverse Gaussian subordinator with $\gamma = \delta = 1$. Figure 3 compares the controller from Theorem 2 with the classical LQR controller.

V. Proof of Main Result

The proof of Theorem 2 is a continuous-time stochastic dynamic programming argument. As in the well-known case of differential equations driven by Brownian motion [16], the Hamilton-Jacobi-Bellman equation must be modified from the deterministic case, based on the the behavior of the stochastic process. In this paper, the required corrections are captured by Lemma 1. The dynamic programming argument is given in Subsection V-A and Lemma 1 is proved in V-B.

A. Regulator Derivation

The optimal expected cost-to-go is the minimum of (10) over all admissible policies. As in standard dynamic programming arguments [17], the optimal cost-to-go function is

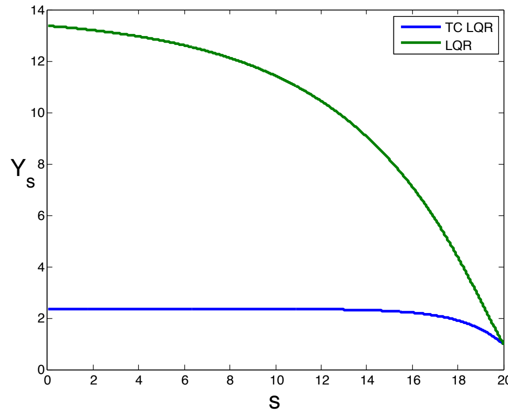

(a)

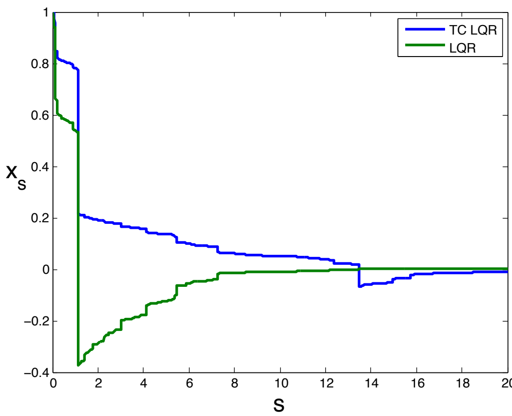

(b)

Fig. 3. 3(a). The blue line shows Y_s, where $x^T Y_s x$ is the optimal expected cost-to-go function when time is perturbed by the inverse Gaussian subordinator. The green line shows P_s, where $x^T P_s x$ is the expected cost-to-go that arises when the classical LQR gains, (2), are substituted into (14). As expected, the $Y_s \le P_s$ for all s. 3(b). Simulations. The blue line is a trajectory x_s generated by using the gain from (13). The green line is a trajectory generated by using the standard LQR gain. Note that the LQR controller induces a large jump near $s = 1$. Often, classical LQR gains are more aggressive than (13) because no H term appears in the inverse.

derived from the identity

$$J_s(x) = \min_u \mathbb{E}\left[\int_s^{s+\delta} \left(x_r^T Q x_r + u_r^T R u_r\right) dr \right. \\ \left. + J_{s+\delta}(x_{s+\delta})\,|\,x_s = x\right] \tag{15}$$

for small $\delta > 0$. Similar to (9), $x_{s+\delta}$ can be written as

$$x_{s+\delta} = e^{A(\tau_{s+\delta}-\tau_s)}x_s + \int_{\tau_s}^{\tau_{s+\delta}} e^{A(\tau_{s+\delta}-\sigma)}Bu(\zeta_\sigma)d\sigma. \tag{16}$$

Recall that τ_s and u_s were assumed to be right-continuous. Furthermore ζ_t is continuous by Proposition 1. It follows that $x_{s+\delta} \to x_s$ as $\delta \downarrow 0$. Thus x_s is right-continuous.

Fix $\epsilon > 0$ and assume that $x_s = x$ and $u_s = u$ (where x and u are simply vectors, not processes). By the inverse property, $u(\zeta_{\tau_s}) = u_s = u$ almost surely. Right-continuity of u_s and ζ_t implies that for $\delta > 0$ sufficiently small,

$$x_{s+\delta} = e^{A(\tau_{s+\delta}-\tau_s)}x + \int_{\tau_s}^{\tau_{s+\delta}} e^{A(\tau_{s+\delta}-\sigma)}Bud\sigma \\ + O((\tau_{s+\delta}-\tau_s)\epsilon) \tag{17}$$

Furthermore, the second term on the right can be written as $\int_0^{\tau_{s+\delta}-\tau_s} e^{A\sigma}Bud\sigma$.

Let $\hat{\tau}_s$ be an independent realization of the process τ_s. Define the noisy state matrices $\mathcal{A}_\delta$ and $\mathcal{B}_\delta$ by

$$\mathcal{A}_\delta = e^{A\hat{\tau}_\delta}, \quad \mathcal{B}_\delta = \int_0^{\hat{\tau}_\delta} e^{A\sigma}d\sigma B. \tag{18}$$

Since $\tau_{s+\delta} - \tau_s$ and $\hat{\tau}_\delta$ have the same distribution, it follows that

$$x_{s+\delta} \overset{d}{=} \mathcal{A}_\delta x + \mathcal{B}_\delta u + O(\hat{\tau}_\delta \epsilon), \tag{19}$$

where $\overset{d}{=}$ means that the two sides have the same distribution. Furthermore, $\mathcal{A}_\delta$, $\mathcal{B}_\delta$, and the $O(\hat{\tau}_\delta\epsilon)$ term and are independent of $\mathcal{F}_s$.

Since x_s and u_s are right-continuous, the integral in (15) can be approximated simply:

$$\int_s^{s+\delta} \left(x_r^T Q x_r + u_r^T R u_r\right) dr = \\ \delta\left(x^T Q x + u^T R u\right) + O(\epsilon\delta). \tag{20}$$

With the local approximations defined, the cost-to-go can now be derived. Note that $J_S(x) = x^T \Phi x$, so define $Y_S = \Phi$. For $s + \delta \le S$, assume inductively that $J_{s+\delta}(x) = x^T Y_{s+\delta} x + O(\epsilon\delta)$ for some positive semidefinite matrix $Y_{s+\delta}$. Combining (15), (19), and (20) gives

$$J_s(x) = \min_u \delta\left(x^T Q x + u^T R u\right) + \\ \mathbb{E}\left[(\mathcal{A}_\delta x + \mathcal{B}_\delta u + O(\hat{\tau}_\delta\epsilon))^T Y_{s+\delta}(\mathcal{A}_\delta x + \mathcal{B}_\delta u + O(\hat{\tau}_\delta\epsilon))\right] \\ + O(\epsilon\delta). \tag{21}$$

The proof of Lemma 1 shows that $\mathbb{E}\left[\hat{\tau}_\delta^2\right]$, $\mathbb{E}\left[\hat{\tau}_\delta \mathcal{A}_\delta\right]$, and $\mathbb{E}\left[\hat{\tau}_\delta \mathcal{B}_\delta\right]$ are all $O(\delta)$ functions, whenever $\mathrm{dom}(\psi)$ contains 0 and the eigenvalues of $-2A$. Thus, the $O(\hat{\tau}_\delta\epsilon)$ terms can be absorbed into the $O(\epsilon\delta)$ term. Because $\mathcal{A}_\delta$ and $\mathcal{B}_\delta$ are independent of $\mathcal{F}_s$, the expected value can be expanded as

$$\mathbb{E}\left[(\mathcal{A}_\delta x + \mathcal{B}_\delta u)^T Y_{s+\delta}(\mathcal{A}_\delta x + \mathcal{B}_\delta u)\right] = \\ x^T \mathbb{E}\left[\mathcal{A}_\delta^T Y_{s+\delta}\mathcal{A}_\delta\right] x + 2u^T \mathbb{E}\left[\mathcal{B}_\delta^T Y_{s+\delta}\mathcal{A}_\delta\right] x \\ + u^T \mathbb{E}\left[\mathcal{B}_\delta^T Y_{s+\delta}\mathcal{B}_\delta\right] u. \tag{22}$$

It follows by (18) and Lemma 1 that the expected values on the right of (22) can be written as

$$\mathbb{E}\left[\mathcal{A}_\delta^T Y_{s+\delta}\mathcal{A}_\delta\right] = Y_{s+\delta} + \delta F(Y_{s+\delta}) + O(\delta^2) \\ \mathbb{E}\left[\mathcal{B}_\delta^T Y_{s+\delta}\mathcal{A}_\delta\right] = \delta B^T G(Y_{s+\delta})^T + O(\delta^2) \\ \mathbb{E}\left[\mathcal{B}_\delta^T Y_{s+\delta}\mathcal{B}_\delta\right] = \delta B^T H(Y_{s+\delta})B + O(\delta^2).$$

Assume without loss of generality that $\delta \le \epsilon$. Plugging the above expected values into (22) gives

$$J_s(x) = x^T Y_{s+\delta} x + \delta x^T\left(Q + F(Y_{s+\delta})\right)x + \\ \min_u \delta\left[2u^T B^T G(Y_{s+\delta})^T x + u^T(R + B^T H(Y_{s+\delta})B)u\right] \\ + O(\epsilon\delta).$$

Note that since $Y_{s+\delta}$ is positive semidefinite, it follows from the form in Lemma 1 that $H(Y_{s+\delta})$ is also positive semidefinite. Similarly, $R + B^T H(Y_{s+\delta})B$ is positive definite. Thus, the unique minimizing input is given by $u = K_{s+\delta}x$, with $K_{s+\delta}$ given in (13). Plugging in the optimal input shows that $J_s(x) = x^T Y_s x + O(\epsilon\delta)$, where Y_s is the positive semidefinite matrix defined by

$$\begin{aligned} Y_s = {} & Y_{s+\delta} + \delta(Q + F(Y_{s+\delta})) \\ & - \delta G(Y_{s+\delta})B(R + B^T H(Y_{s+\delta})B)^{-1}B^T G(Y_{s+\delta})^T. \end{aligned}$$

Letting ϵ and δ go to zero shows that $J_s(x) = x^T Y_s x$, where Y_s solves the backward differential equation (12). ■

B. Proof of Lemma 1

Assume that the Jordan decomposition of A has q blocks, and for each eigenvalue λ_i, the corresponding Jordan block has size $M_i \times M_i$. The matrix exponential can be written as

$$e^{At} = \sum_{i=1}^{q} e^{\lambda_i t} \sum_{j=0}^{M_i-1} t^j \Gamma_j^i,$$

for some matrices $\Gamma_j^i \in \mathbb{C}^{n\times n}$ such that $\sum_{i=1}^q \Gamma_0^i = I$. Thus,

$$e^{A^T t} Y e^{At} = \sum_{i,j=1}^{q} e^{(\overline{\lambda}_i+\lambda_j)t} \sum_{k=0}^{M_i-1} \sum_{l=0}^{M_j-1} t^{k+l} {\Gamma_k^i}^* Y \Gamma_l^j. \quad (23)$$

Set $\mu = \overline{\lambda}_i + \lambda_j$. By hypothesis, $-\mu \in \mathrm{dom}(\psi)$, and thus ψ is analytic at $-\mu$. The desired expected value can be computed in terms of the scalars $\mathbb{E}\left[\tau_s^k e^{\mu\tau_s}\right]$. For simpler algebra, set $\beta(z) = -\psi(-z)$. Then from (6), $\mathbb{E}\left[e^{\mu\tau_s}\right] = e^{-s\psi(-\mu)} = e^{s\beta(\mu)}$. Now $\mathbb{E}\left[\tau_s^k e^{\mu\tau_s}\right]$ can be computed as

$$\mathbb{E}\left[\tau_s^k e^{\mu\tau_s}\right] = \frac{d^k}{d\mu^k}\mathbb{E}\left[e^{\mu\tau_s}\right] = \frac{d^k}{d\mu^k} e^{s\beta(\mu)}. \quad (24)$$

The exponential can be expanded as $e^{s\beta(\mu)} = 1 + s\beta(\mu) + s^2 f(\mu, s)$, where $f(\mu, s) = \sum_{j=0}^{\infty} \frac{s^j \beta(\mu)^{j+2}}{(j+2)!}$. Thus, for $k \geq 0$, (24) becomes

$$\mathbb{E}\left[\tau_s^k e^{\mu\tau_s}\right] = \begin{cases} 1 + s\beta(\mu) + O(s^2) & k = 0 \\ s\frac{d^k\beta}{d\mu^k}(\mu) + O(s^2) & k \geq 1. \end{cases} \quad (25)$$

Combining (23) and (25) gives the desired expectation:

$$\mathbb{E}\left[e^{A^T\tau_s} Y e^{A\tau_s}\right] = Y +$$

$$s \sum_{i,j=1}^{q} \sum_{k=0}^{M_i-1} \sum_{l=0}^{M_l-1} \frac{d^{j+k}\beta}{d\mu^{j+k}} (\overline{\lambda}_1 + \lambda_2) {\Gamma_k^i}^* Y \Gamma_l^j + O(s^2).$$

The expressions for G and H can be found by combining the argument above with the scalar formula

$$\int_0^t \sigma^k e^{\lambda\sigma} d\sigma = \begin{cases} \begin{array}{c} \left(e^{\lambda t} - 1\right)(-1)^k k! \lambda^{-k-1} + \\ e^{\lambda t} \sum_{i=1}^{k} \frac{k!}{i!} t^i (-1)^{k-i} \lambda^{-k-i-1} \end{array} & \lambda \neq 0 \\ \frac{t^{k+1}}{k+1} & \lambda = 0. \end{cases}$$

■

VI. Conclusion

The time-changed regulator problem in this paper could be extended in numerous ways. For instance, time-varying cost matrices could be incorporated by simply changing Q and R to Q_s and R_s, respectively. Coupling in the cost between x_s and u_s is also simple. Subtleties may arise in infinite horizon, since it is unclear when (12) tends to a steady-state solution. Generalization to output feedback seems challenging.

Lemma 1 and (12) could likely be explained in terms of Itô's formula for Lévy processes. Such an interpretation might be helpful in the extension of the dynamic programming argument to nonlinear systems.

Additionally, most of the arguments should extend with minimal change if the system is also driven by a Brownian motion, independent of the time-change. If, however, they are correlated the results could be more difficult. The problem would also be more interesting if the time-change could be correlated with the state and input history.

References

[1] A. E. D. Veraart and M. Winkel, "Time change," in *Encyclopedia of Quantitative Finance*. Wiley, 2010, vol. 4, pp. 1812–1816.

[2] P. K. Clark, "A subordinated stochastic process model with finite variance for speculative prices," *Econometrica*, vol. 41, pp. 135–155, 1973.

[3] T. Ané and H. Geman, "Order flow, transaction clock, and normality of asset returns," *The Journal of Finance*, vol. 55, no. 5, pp. 2259–2284, 2000.

[4] P. Carr and L. Wu, "Time-changed Lévy processes and option pricing," *Journal of Financial Economics*, vol. 71, pp. 113–141, 2004.

[5] D. M. Eagleman, "Human time perception and its illusions," *Current Opinion in Neurobiology*, vol. 18, no. 2, 2008.

[6] M. Jazayeri and M. N. Shadlen, "Temporal context calibrates interval timing," *Nature Neuroscience*, vol. 13, no. 8, 2010.

[7] T. E. Hudson, L. T. Maloney, and M. S. Landy, "Optimal compensation for temporal uncertainty in movement planning," *PLoS Computational Biology*, vol. 4, no. 7, 2008.

[8] N. M. Freris and P. R. Kumar, "Fundamental limits on synchronization of affine clocks in networks," in *IEEE Conference on Decision and Control*, 2007.

[9] O. Simeone, U. Spagnolini, Y. Bar-Ness, and S. H. Strogatz, "Distributed synchronization in wireless networks," *IEEE Signal Processing Magazine*, vol. 25, no. 5, pp. 81–97, 2008.

[10] C. Lorand and P. H. Bauer, "Stability analysis of closed-loop discrete-time systems with clock frequency drifts," in *American Control Conference*, 2003.

[11] R. Singh and V. Gupta, "On LQR control with asynchronous clocks," in *IEEE Conference on Decision and Control and European Control Conference*, 2011.

[12] S. M. LaValle and M. B. Egerstedt, "On time: Clocks, chronometers, and open-loop control," in *IEEE Conference on Decision and Control*, 2007.

[13] S. G. Carver, E. S. Fortune, and N. J. Cowan, "State-estimation and cooperative control with uncertain time," in *American Control Conference*, 2013.

[14] D. Applebaum, *Lévy Processes and Stochastic Calculus*. Cambridge University Press, 2004.

[15] J. R. Michael, W. R. Schucany, and R. W. Haas, "Generating random variates using transformations with multiple roots," *The American Statistician*, vol. 30, no. 2, pp. 88–90, 1976.

[16] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, 5th ed. Springer, 2002.

[17] F. L. Lewis and V. L. Syrmos, *Optimal Control*, 2nd ed. John Wiley & Sons, 1995.