

Redundant Data Transmission in Control/Estimation Over Wireless Networks: Extended Version

Alexandre Rodrigues Mesquita¹, João Pedro Hespanha¹, Girish Nair²

Abstract—In wireless networks the probability of successful communication can be significantly increased by transmitting multiple copies of a same packet. Communication protocols that exploit this by dynamically assigning the number of transmitted copies of the same data can significantly improve the control performance in a networked control system with only a modest increase in the total number of transmissions. In this paper we develop techniques to design communication protocols that exploit multiple packets transmissions while seeking a balance between stability/estimation performance and communication rate. An average cost optimality criterion is employed to obtain optimal protocols. Optimal protocols are also obtained for networks whose nodes are subject to limited computation.

I. INTRODUCTION

Packet losses in wireless networks have a critical role in determining the performance of networked control systems (NCS). Losses due to fading can generally be mitigated through the use of *diversity*, i.e., the transmission of redundant signals through mostly independent channel realizations. In general, this leads to an overuse of communication resources, but in NCSs it is possible to use redundant transmissions judiciously so as to reap the benefits of diversity with limited communication overload.

Diversity schemes include time-domain, frequency-domain and space redundancy [1]. In a time diversity scheme, multiple instances of a same signal are sent at different time instants. This scheme is particularly suitable for mobile nodes, that exhibit relatively short coherence time. In the frequency diversity scheme, multiple versions of the signal are spread over a wide spectrum. The space diversity scheme consists of transmitting the signal over different propagation paths, which is typically achieved by the use of multiple antennas. In higher layers, path diversity is also possible by sending packets through multiple routes. Although many diversity schemes are dynamically exploited in data networks (see e.g. [2], [3]), where transmissions are scheduled according to the network status, these techniques do not take into account nor benefit from the dynamical nature of NCSs.

This paper is concerned with diversity techniques specifically applied to NCSs. The results presented are independent of the scheme used for diversity and simply assume that a number of independent redundant channels are available for data transmission. For simplicity data drops in the different channels are assumed i.i.d. Our focus is on deciding how

many redundant copies of a packet should be transmitted at each sampling time and what benefits can be drawn from this. We initially show how diversity scheduling can improve the stability properties of a NCS. Next, we design scheduling techniques that optimize a criterion that involves the conflicting objectives of high control/estimation performance and low transmission rate.

The adopted NCS architecture is depicted in Fig. 1, which considers the case of a single sensor and a controller. At each time step, the sensor sends measurements to the controller with a certain redundancy degree. As discussed above, we assume that the network drops packets with a given probability (i.i.d.) and that some acknowledgement mechanism is available so that the sensor knows which measurements were received by the controller. Along with packet redundancy, we also consider the possibility of not sending packets at a given time instant, which may prevent unnecessary transmissions when the control/estimation behavior is satisfactory. This idea has also been explored in [4], [5].




Fig. 1. NCS architecture

In the first part of the paper we focus our attention on the control of a simple scalar unstable process. For such process it is well known that mean-square instability arises whenever the drop probability rises above a certain threshold. Moreover, no matter how small the drop probability is, some statistical moments of the process' state will always be unbounded. It turns out that the use of redundant transmissions can be used to stabilize any given statistical moments for any probability of drop. Surprisingly, we show that this can be achieved with no significant increase in communications by a judicious use of redundant transmissions.

In the second part we consider a general linear time-invariant process with a certainty equivalence control on the same NCS of Fig. 1. The controller constructs estimates of the process state using the measurements transmitted by the sensor, which uses a redundant transmissions policy that minimizes the combined average cost of the estimation error in the controller and the number of packets sent. First, we

¹ Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara, CA 93106

² Department of Electrical and Electronic Engineering, University of Melbourne, VIC 3010, Australia

consider the ideal case in which the sensor can reconstruct the state estimates available to the controller from the acknowledgement information. Secondly, and motivated by the fact that in some applications sensors have very limited computation capabilities, we find policies that minimize the same cost but base their decisions solely on the number of consecutive transmission failures.

For simplicity of presentation, we consider NCSs with full local state measurements and no network delays. However, the results obtained can be readily extended to the case of partial state measurements, multiple nodes, and delays in the network by following the procedure in [6].

Overall, our results suggest that a very small redundancy level is actually sufficient for drastic performance gains. In fact, our results led us to the belief that a redundancy degree higher than two is seldom necessary in practical applications, which would imply that the implementation of diversity schemes does not demand an extremely expensive infra-structure.

II. MOMENT STABILIZATION USING PACKET REDUNDANCY

In this section, we use a first order process to show how packet redundancy can be used to improve the stability properties of a networked control process in the presence of communication drops. We consider the following unstable scalar process

$$x(k+1) = ax(k) + u(k) + w(k) , \quad (1)$$

where $|a| > 1$, $x(k) \in \mathbb{R}$ is the state at an integer time k , $u(k)$ is the control variable at the same instant of time and $w(k)$ is a zero-mean Gaussian white noise process with variance σ^2 . The controller and a sensor that measures the state $x(k)$ are connected through a network that drops packets independently of each other, with probability $p \in (0, 1)$. In order to improve the probability that the measurement $x(k)$ reaches the controller, the sensor may transmit multiple copies of this message in independent packets.

For simplicity, we consider a deadbeat controller of the type:

$$u(k) = \begin{cases} -ax(k), & \text{no drop at time } k \\ 0, & \text{drop at time } k . \end{cases} \quad (2)$$

A convenient fact about this control strategy is that the statistical moments of the state can be easily computed from the number of consecutive transmission failures, where a *transmission failure* is characterized by the failure of all the tentatives of transmitting $x(k)$ at time k . Let us denote by $l(k)$ the number of consecutive transmission failures that occurred before time k . We are interested in designing protocols that determine how many identical packets to send at time k as a function of how many consecutive transmission failures occurred before time k . Such protocols can be specified by a function v that maps the number of consecutive drops $l(k)$ to the number of packets to send. For example,

if we use the identity function $v(l) = l$, then $l(k)$ identical packets will be sent at time k (note that we do not exclude the possibility of sending zero packets).

Under the assumption of independent drops, $l(k)$ can be written as an infinite Markov chain with transition probabilities

$$\Pr(l(k+1) = l(k) + 1 \mid l(k)) = p^{v(l(k))}, \quad k \geq 0 \quad (3)$$

$$\Pr(l(k+1) = 0 \mid l(k)) = 1 - p^{v(l(k))}, \quad k \geq 0 . \quad (4)$$

The stationary probabilities $\mu(l)$ for this Markov chain must therefore satisfy

$$\mu(l+1) = p^{v(l)} \mu(l) = p^{\sum_{m=0}^l v(m)} \mu(0), \quad l \geq 0 \quad (5)$$

$$\sum_{l=0}^{\infty} \mu(l) = 1 , \quad (6)$$

which allows us to conclude that

$$\mu(0) = \left(1 + \sum_{l=1}^{\infty} p^{\sum_{m=0}^{l-1} v(m)} \right)^{-1} \quad (7)$$

and, for $l > 1$,

$$\mu(l) = p^{\sum_{m=0}^{l-1} v(m)} \left(1 + \sum_{m=1}^{\infty} p^{\sum_{n=0}^{m-1} v(n)} \right)^{-1} . \quad (8)$$

Notice that $\mu(0)$ is well defined as long as there exists a constant $L > 0$ such that $v(l) \geq 1, \forall l \geq L$. Under this condition one can also verify that the chain is aperiodic and recurrent (see e.g. [7, Chap. 8]). Therefore, we can apply [7, Thm. 14.3.3] to conclude that

$$\lim_{k \rightarrow \infty} \mathbb{E} [|x(k)|^2] = \sum_{l=0}^{\infty} \mathbb{E} [|x(k)|^2 \mid l(k) = l] \mu(l) . \quad (9)$$

In view of (1) and (2), we have that

$$\mathbb{E} [|x(k)|^2 \mid l(k)] = \mathbb{E} \left[\left(\sum_{m=0}^l a^m w(k-m) \right)^2 \right] , \quad (10)$$

which can be used in (9) to obtain

$$\mathbb{E} [|x(k)|^2] \rightarrow \sigma^2 \mu(0) \left(1 + \sum_{l=1}^{\infty} p^{\sum_{m=0}^{l-1} v(m)} \sum_{m=0}^l |a|^{2m} \right) \quad (11)$$

as $k \rightarrow \infty$. The following theorem then follows from the ratio test:

Theorem 1: The second moment of $x(k)$ is bounded if

$$\lim_{l \rightarrow \infty} |a|^2 p^{v(l)} < 1 . \quad (12)$$

Conversely, the second moment of $x(k)$ is unbounded if $\lim_{l \rightarrow \infty} |a|^2 p^{v(l)} > 1$.

We conclude from here that mean-square stability can be achieved for any unstable pole a and any drop probability $p < 1$ by a proper choice of the redundant packet transmission protocol that specifies the function $v(l)$. In fact, all that

is needed is to select $v(l)$ sufficiently large for large values of l :

$$\lim_{l \rightarrow \infty} v(l) > \frac{2 \log |a|}{-\log p} \quad (13)$$

Analogously, the condition

$$\lim_{l \rightarrow \infty} v(l) > \frac{q \log |a|}{-\log p} \quad (14)$$

can be shown to guarantee boundedness of the q -th moment.

Lyapunov techniques. However, in this paper we focus on the interplay between control performance and communication cost.

From (14), we can see that to achieve stability one may require a protocol that, at times, sends a large number of packets, which would require a large communication rate. It is important thus to study the communication requirements of these protocols. To this effect let us derive the expected communication rate for a given function $v(l)$. We assume that the packet sizes are constant and sufficiently large so that the controller receives $x(k)$ with negligible quantization loss. In this case, the expected asymptotic transmission rate is

$$\bar{R} := \lim_{k \rightarrow \infty} \mathbb{E}[v(l(k))] = S\mu(0) \left(1 + \sum_{l=1}^{\infty} v(l) p^{\sum_{m=0}^{l-1} v(m)} \right), \quad (15)$$

where S denotes the size of a single packet.

Interestingly, one can find stabilizing protocols with rate \bar{R} arbitrarily close to S , which corresponds to a single packet per time unit. Consider for example the case of

$$v(l) = \begin{cases} 1 & \text{for } l \leq N \\ M & \text{for } l > N \end{cases}, \quad (16)$$

where M is a redundancy degree satisfying (14). A Taylor expansion of (17) around $p = 0$ leads to

$$\bar{R} = S(1 + O(p^{N+1})), \quad (17)$$

that is, the expected transmission rate is order p^{N+1} larger than S . Thus, \bar{R} can be made arbitrarily close to S by increasing N . Evidently, the larger we make N the larger the moments will be. This relationship between average transmission rate and control performance is investigated in the following sections.

We note that such a strong result as (17) would no longer hold when packet sizes increase with l , as indicated by some preliminary results in the case where quantization is necessary before transmitting $x(k)$.

III. A GENERAL NCS

In the remainder of the paper we consider a more general NCS and study the effect of using packet redundancy in the control performance of the system (as opposed to just its stability). In this architecture our goal is to stabilize a linear time-invariant process

$$x(k+1) = Ax(k) + Bu(k) + w(k) \quad (18)$$

where $x \in \mathbb{R}^n$ denotes the state of the process, $u \in \mathbb{R}^{n_1}$ the control input, and $w(k) \in \mathbb{R}^n$ an n -dimensional zero-mean Gaussian white noise process. As in the previous section, we assume that the whole state $x(k)$ can be measured by a sensor which communicates with the controller through a network that drops packets independently of each other, with a probability $p \in (0, 1)$. We assume that we use a certainty equivalence control law of the form

$$u(k) = K\hat{x}(k) \quad (19)$$

where the matrix K is chosen such that $A + BK$ is Schur and $\hat{x}(k)$ is an optimal estimate of $x(k)$ based on the measurements that successfully reached the controller up to time k . In particular,

$$\hat{x}(k) := \mathbb{E}[x(k) | x(s), s < k, s \in \mathcal{T}_{\text{success}}] \quad (20)$$

where $\mathcal{T}_{\text{success}}$ denotes the set of times at which the sensor succeeded in transmitting the measured state to the controller. This optimal state estimate can be computed recursively using

$$\hat{x}(k+1) = \begin{cases} A\hat{x}(k) + Bu(k) & \text{if } k \notin \mathcal{T}_{\text{success}} \\ Ax(k) + Bu(k) & \text{if } k \in \mathcal{T}_{\text{success}} \end{cases} \quad (21)$$

From (18) and (21), we conclude that the estimation error $e(k) := \hat{x}(k) - x(k)$ evolves according to the following dynamics:

$$e(k+1) = \begin{cases} Ae(k) - w(k) & \text{if } k \notin \mathcal{T}_{\text{success}} \\ -w(k) & \text{if } k \in \mathcal{T}_{\text{success}} \end{cases} \quad (22)$$

The closed-loop dynamics (18)–(19) can be expressed in terms of this error using

$$x(k+1) = (A + BK)x(k) + BK e(k) + w(k). \quad (23)$$

In our certainty equivalence control scheme, we assume that minimizing the estimation error will also minimize the control performance. Note that, even though the separation principle does not hold in the most general case, it does hold for the case presented in section V (see e.g. [8], [9]) under the assumption of perfect acknowledgement.

IV. OPTIMAL COMMUNICATION PROTOCOL

Our goal now is to determine an optimal policy that decides when to send multiple copies of the same packet and how many copies to send. This policy should be optimal in the sense that it achieves a desirable trade off between the conflicting objectives of keeping small the estimation error $e(k)$, that drives the closed-loop dynamics (23), while achieving this with a minimal amount of communication.

To this effect, we consider the following average cost (AC) minimization criterion

$$J(\pi, e_0) := J_{\text{est}}(\pi, e_0) + \lambda J_{\text{com}}(\pi, e_0) \quad (24)$$

where

$$J_{\text{est}}(\pi, e_0) := \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{e_0}^{\pi} \left[\sum_{k=0}^{N-1} e(k+1)' Q e(k+1) \right] \quad (25)$$

$$J_{\text{com}}(\pi, e_0) := \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{e_0}^{\pi} \left[\sum_{k=0}^{N-1} v(k) \right] \quad (26)$$

where λ is a positive scalar, Q a positive definite matrix (typically $Q = I$), $v(k) \geq 0$ is the number of packets sent at time k , and $\mathbb{E}_{e_0}^{\pi}$ denotes the expectation given a policy π and an initial state $e(0) = e_0$. We consider policies π to be functions that map $e(k)$ to $v(k)$. For technical reasons, we restrict our attention to policies π that send M packets whenever $\|e(k)\|$ exceeds some pre-specified constant L , i.e.,

$$\|e(k)\| > L \Rightarrow \pi(e(k)) = M. \quad (27)$$

The criterion in (24) is a weighted sum of two terms: the first term $J_{\text{est}}(\pi, e_0)$ penalizes a time-averaged expected quadratic estimation error (the variance, if $Q = I$), whereas the second term $J_{\text{com}}(\pi, e_0)$ penalizes the average communication rate, measured in terms of the number of messages sent per unit of time. The constant λ allows one to adjust the relative weight of the two terms. As $\lambda \rightarrow 0$, communication is not penalized, whereas as $\lambda \rightarrow \infty$, communication is heavily penalized. Intermediate values of λ will yield Pareto-optimal compromises between the two conflicting criteria.

The number $v(k)$ of packets sent at time k essentially controls the probability of a successful transmission. In particular, since drops are assumed independent,

$$P(k \notin \mathcal{T}_{\text{success}}) = p^{v(k)}. \quad (28)$$

It is therefore convenient to imagine that policies π are actually directly controlling this probability. Redefining

$$\pi(k) := p^{v(k)}, \quad (29)$$

we can thus re-write

$$J_{\text{com}}(\pi, e_0) := \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{e_0}^{\pi} \left[\sum_{k=0}^{N-1} \frac{\log \pi(k)}{\log p} \right]. \quad (30)$$

The set of admissible control actions is therefore

$$\Pi(e) := \begin{cases} \{p^v : v \in \{0, \dots, M\}\} & \text{if } \|e\| \leq L \\ \{p^M\} & \text{if } \|e\| > L. \end{cases} \quad (31)$$

We denote the set of all control policies by Δ (this includes time-variant and random policies). The set Δ_0 of *stationary* policies is the set of measurable functions b such that $b(e) \in \Pi(e)$ for all $e \in \mathbb{R}^n$.

A policy π^* is said to be *AC-optimal* if

$$J(\pi^*, e) = \inf_{\pi \in \Delta} J(\pi, e) =: J^*(e), \quad \forall e \in \mathbb{R}^n, \quad (32)$$

and J^* is called the *optimal AC-function*.

Using (22), we can rewrite the cost function (24) as follows

$$J(\pi, e) = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_e^{\pi} \sum_{k=0}^{N-1} c(e(k), b(k)) \quad (33)$$

where $b(k) = \pi(e(k))$,

$$c(e, b) = be' A' Q A e + \lambda \log_p b^{-1} + \text{trace } Q \Sigma, \quad (34)$$

and Σ denotes the covariance matrix of $w(k)$.

To formulate the main result of this section we need to define the transition probability measure $P(Y|e, b) = \Pr\{e(k+1) \in Y | e(k) = e, b(k) = b\}$. From (22), we have that

$$P(dy|e, b) = (1-b)f(y) + bf(y - Ae), \quad (35)$$

where f is the p.d.f. of the normal distribution with zero mean and covariance Σ .

Assumption 1: a) The constant M in (27) is chosen sufficiently large so that

$$p^M < \frac{1}{\rho(A)^2}, \quad (36)$$

where $\rho(A)$ denotes the spectral radius of the matrix A . b) The constant L in (27) is chosen so that there exist scalars $r \geq L$ and $\alpha < 1$ such that

$$\nu(C_r) \leq \frac{\alpha (\lambda_{\min}(Q)r^2 - p^M \text{trace}(H\Sigma))}{\lambda_{\min}(Q)(r^2 - L^2) + \alpha \lambda_{\max}(H)(1 - p^M)L^2} \quad (37)$$

where C_r denotes the open ball with radius r centered at the origin in \mathbb{R}^n , $\nu(\cdot)$ denotes the measure corresponding to the the density f (the multi-variable normal distribution with zero mean and covariance Σ), and H is the unique positive definite solution of

$$p^M A' H A - \alpha H = -Q. \quad (38)$$

Remark 1: Assumption 1(a) guarantees that $p^{M/2}A$ is Schur. Because of this, for α sufficiently close to 1, (38) has a unique positive definite solution. Although the condition in Assumption 1(b) is not very restrictive, we conjecture that it is actually not necessary for the result in Theorem 2. ■

The next theorem states the existence of a solution to the AC-optimality problem.

Theorem 2: If Assumption 1 holds, then:

- 1) There exist a constant $\varrho^* \geq 0$, a continuous function φ^* and a stationary policy $\pi^* \in \Delta_0$ such that the triplet $(\varrho^*, \varphi^*, \pi^*)$ satisfies the average cost optimality equation (ACOE):

$$\begin{aligned} \varrho^* + \varphi^*(e) &= \min_{b \in \Pi(e)} \left[c(e, b) + \int \varphi^*(y) P(dy|e, b) \right] \\ &= c(e, \pi^*(e)) + \int \varphi^*(y) P(dy|e, \pi^*(e)), \quad \forall e \in \mathbb{R}^n; \end{aligned} \quad (39)$$

- 2) π^* is AC-optimal and ϱ^* is the optimal AC-function.

Proof. The proof is based on [10, Thm. 2.5]. We start by proving Assumption 2.3 of [10]. Define $h_\pi(e) := 1 - \pi(e)$. Then, we have from (35) that

$$P(C|e, \pi(e)) \geq h_\pi(e)\nu(C) . \quad (40)$$

By Assumption 1(a), there exists α close to 1 and H positive definite such that (38) holds. Define $V(e) := \gamma e' H e + 1_{C_r}(e)$, where

$$\gamma = \frac{\alpha p^M}{\lambda_{\min}(Q)(r^2 - L^2) + \alpha \lambda_{\max}(H)(1 - p^M)L^2} . \quad (41)$$

Next we show that V satisfies the Assumption 2.3(b) of [10], which in our formalism can be written as follows.

$$\int V(y)P(dy|e, \pi(e)) \leq \alpha V(e) + h_\pi(e) \int V(y)d\nu(y) . \quad (42)$$

This condition can be understood as a Lyapunov-Foster condition that is satisfied uniformly on the set of policies. To verify that (42) indeed holds, we define

$$\begin{aligned} \Delta V(e) &:= \int V(y)P(dy|e, \pi) - h_\pi(e) \int V(y)d\nu(y) \\ &= \int \pi V(y)f(y - Ae)dy \leq \pi \gamma e' A H A e + \pi \gamma \text{trace}(H\Sigma) \\ + \pi \nu(C_r) &= \frac{\pi}{p^M} \gamma e' (\alpha H - Q)e + \pi \gamma \text{trace}(H\Sigma) + \pi \nu(C_r) , \end{aligned} \quad (43)$$

where the last equality comes from (38) and the dependence of π on e was omitted to simplify the notation. From this it follows that

$$\begin{aligned} \Delta V(e) - \alpha V &\leq \gamma e' \left(\alpha \left(\frac{\pi}{p^M} - 1 \right) H - Q \right) e + \\ &\quad \pi \gamma \text{trace}(H\Sigma) + \pi \nu(C_r) - \alpha 1_{C_r}(e) . \end{aligned} \quad (44)$$

In the region $\|e\| < L$, we can upper bound the right-hand side of (44) by using the fact that $\pi(e) \leq 1$ and $\|e\|^2 \leq L^2$ to obtain

$$\begin{aligned} \Delta V - \alpha V &\leq \alpha \gamma \left(\frac{1}{p^M} - 1 \right) \lambda_{\max}(H)L^2 - \frac{\gamma}{p^M} \lambda_{\min}(Q)L^2 \\ &\quad + \gamma \text{trace}(H\Sigma) + \nu(C_r) - \alpha \leq 0 , \end{aligned} \quad (45)$$

where the last inequality is obtained by using (37) and (41).

In the region $\|e\| \geq L$, we have that $\pi(e) = p^M$, which implies that the right-hand side of (44) decreases strictly with $\|e\|$ in this region, except for $\|e\| = r$, where it jumps. Thus, it only remains to investigate if (42) is satisfied at $\|e\| = r$. For $\|e\| = r$, the following upper bound can be obtained

$$\begin{aligned} \Delta V - \alpha V &\leq -\gamma e' Q e + \gamma p^M \text{trace}(H\Sigma) + p^M \nu(C_r) \\ &\leq -\gamma \lambda_{\min}(Q)r^2 + \gamma p^M \text{trace}(H\Sigma) + p^M \nu(C_r) \leq 0 , \end{aligned} \quad (46)$$

where again the last inequality follows from (37) and (41). Therefore, we have verified (42).

It follows easily that $\inf_\pi \int h_\pi(e)V(e) > 0$ since $\pi(e) = p^M$ for $\|e\| \geq L$ and therefore Assumption 2.3 of [10] is also satisfied.

By our choice of V , there exists a positive constant δ such that $\sup_{\Pi(e)} c(e, b) < \delta V(e)$, which verifies Assumption 2.1 of [10].

From our definition of $h_\pi(e)$, it follows that the set $\{e : h_\pi(e) > 0\}$ is petite [7, Chap. 5]. Therefore, for each policy π , (42) satisfies a Foster-Lyapunov condition as in [7, Thm. 11.3.4]. From this it follows that the process is positive Harris recurrent for each policy π , which fulfills Assumption 2.2 of [10].

Finally, we verify Assumption 2.4 of [10]. To this effect, let $d(\cdot, \cdot)$ be the discrete metric on $\Pi(e)$. Then, we have

$$\begin{aligned} |c(e, b) - c(\tilde{e}, \tilde{b})| &= b e' A' Q A e - \lambda \log_p b - \tilde{b} \tilde{e}' A' Q A \tilde{e} + \\ &\quad \lambda \log_p \tilde{b} \leq (b - \tilde{b}) e' A' Q A e + \tilde{b} (e - \tilde{e})' A' Q (e - \tilde{e}) \\ &\quad + \kappa_1 d(b, \tilde{b}) \leq \kappa_2(e) d(b, \tilde{b}) + \kappa_3 \|e - \tilde{e}\|^2 \\ &\leq \kappa(e) \left(\max\{d(b, \tilde{b}), \|e - \tilde{e}\|\} \right)^2 , \end{aligned} \quad (47)$$

for some properly chosen nonnegative and finite $\kappa_1, \kappa_2(e), \kappa_3$ and $\kappa(e)$. In a similar way we can find a nonnegative function $\kappa_P(e)$ such that

$$\begin{aligned} \left| \int V(y) \left(P(dy|e, b) - P(dy|\tilde{e}, \tilde{b}) \right) \right| \\ \leq \kappa_P(e) \left(\max\{d(b, \tilde{b}), \|e - \tilde{e}\|\} \right)^2 . \end{aligned} \quad (48)$$

The application of [10, Thm. 2.5] completes the proof. ■

The solution to the ACOE can be obtained using the following value iteration algorithm. Let s_N (which can be seen as a N -th stage cost) and π_N be defined as follows:

$$s_N(e) := \min_{b \in \Pi(e)} \left[c(e, b) + \int s_{N-1}(y)P(dy|e, b) \right] \quad (49)$$

$$=: c(e, \pi_N(e)) + \int s_{N-1}(y)P(dy|e, p_N(e)) , \quad (50)$$

where $s_0 \equiv 0$. Let $z \in \mathbb{R}^n$ be an arbitrary but fixed state. Define a sequence of constants j_N and a sequence of functions $\varphi_N(e)$ as

$$j_N := s_N(z) - s_{N-1}(z) \quad \text{and} \quad \varphi_N(e) := s_N(e) - s_N(z) . \quad (51)$$

Then, the value iteration algorithm is said to converge if

$$j_N \rightarrow \varrho^* \quad \text{and} \quad \varphi_N(x) \rightarrow \varphi^*(x) \quad \text{as} \quad N \rightarrow \infty . \quad (52)$$

Under the conditions of Theorem 2, [10, Thm. 2.6] guarantees that this value iteration algorithm always converges. Further, according to [10, Cor. 2.9], there exists an AC-optimal policy that is an accumulation point of $\{\pi_N(e)\}$.

Given the infinite-dimensionality of the value iteration algorithm, in most cases the protocol should be constructed offline and then a look-up table would be used.

V. A SIMPLIFIED OPTIMAL PROTOCOL

In many applications sensors have limited computational capabilities that could prevent the use of optimal elaborate protocols that require the computation of estimation errors. In addition, solving for the optimal policy may be computationally intense for high-dimensional systems. To address these issues, we can design a simplified protocol that bases its decision rule only the consecutive number of failures $l(k)$ that occurred prior to the k -th sampling time, much like the protocols considered in Sec. II.

In general, this would lead to Partially Observable Markov Processes. Fortunately, for this estimation problem, the beliefs for $e(k)$ (probability distributions given the history of $\{l(s); s \leq k\}$) converge almost surely to an invariant set where they are completely determined by $l(k)$, i.e., they do not depend on the history $\{l(j), j < k\}$ or on the previous beliefs. This is because once a packet is successfully transmitted, the belief for $e(k)$ is solely given by $f(e)$, through (35), and it does not depend on any previous beliefs. Hence, the average cost criterion does not depend on the initial belief. Thus, without loss of generality, we can restrict our search for optimal policies to the case where $l_0 = 0$ and e_0 has distribution $f(\cdot)$. This implies that $e(k)$ is distributed according to $\sum_{m=0}^{l(k)} A^m(k)\omega_m$, where ω_m are i.i.d. variables with density $f(\cdot)$.

Thus, we can redefine the per-stage cost as

$$\bar{c}(l, b) = \mathbb{E} [c(e(k), b(k)) \mid l(k) = l] , \quad (53)$$

which can be written as

$$\bar{c}(l, b) = b \text{trace}(A'QA\Sigma_l) + \lambda \log_p b^{-1} + \text{trace} Q\Sigma , \quad (54)$$

where

$$\Sigma_l := \mathbb{E} [e(k)e(k)' \mid l(k) = l] = \sum_{m=0}^l A'^m \Sigma A^m . \quad (55)$$

Moreover, if we restrict the set of policies to be such that $b = p^M$ for $l \geq T$, we can truncate the Markov chain by redirecting the jumps $(l = T) \rightarrow (l = T + 1)$ to $(l = T) \rightarrow (l = T)$ without affecting the optimal policy and the optimal cost. Thus, we have moved from the infinite dimensional problem in Sec. IV to a finite dimensional problem. Using the per-stage cost \bar{c} and the transition probabilities for l that we described in Sec. II, one can calculate AC-optimal policies that depend on l only. This could be done either via dynamic programming or via direct optimization, since the average costs can be directly calculated using the stationary distribution as in Sec. II. Note that the average cost being optimized is indeed the same one as in Sec. IV.

VI. NUMERICAL EXAMPLES

The results in the previous sections were applied to a scalar example with $A = 2$, $\Sigma = 3$, $Q = 1$, $p = 0.15$ and $L = 10$. By varying λ from 0.001 to 200, we constructed the Pareto frontiers shown in Fig. 2.

To show the performance improvement that arises from judiciously sending redundant information, we considered also the base case which always sends one packet per time step. We restricted our policies to a minimum number of transmissions denoted by $\underline{M} \in \{0, 1\}$. Several important observations can be deduced from Fig. 2:

- 1) Using the trivial policy $v(k) = 1, \forall k$, will of course minimize communication (x-axis) when $\underline{M} = 1$, but this is at the expense of a significant larger estimation error (y-axis). In fact, based on the results of Section II, we know that for unstable systems and large drop probabilities, $v(k) \equiv 1$ can lead to instability.
- 2) The policy that uses $\underline{M} = 1$ and $M = 2$ is able to decrease the estimation cost by 30% while increasing the communication cost by only 6%.
- 3) Increasing the maximum number of redundant packets M beyond 2, hardly improves the Pareto-optimal boundary.
- 4) The simplified optimal policy discussed in Section IV produces protocols that can be quite close to the Pareto-optimal boundary.
- 5) A number of the simplified optimal policies are nontrivial, that is, their redundancy degree is not constant. To see that, notice that trivial policies must have an integer communication cost.
just uses the sentence above to comment on these controllers.]
- 6) If one were to allow no transmissions at some time instants (i.e., $v(k) \geq 0$ instead of $v(k) \geq 1$) then one could further improve the optimal Pareto-optimal boundary.

A phenomenon that commonly arises in multi-objective MDPs is that points on the Pareto frontier do not always correspond to deterministic policies. This is the case for the Pareto frontier of the simplified protocol, where only the points marked with a cross correspond to deterministic policies and the lines linking those points correspond to randomized policies that can be derived from the deterministic ones as explained in [11].

Figure 3 illustrates the fact that the use of optimal policies becomes more advantageous as the drop out probability p is increased. To construct this figure we also considered values for p that do not satisfy Assumption 1(b) but we still had convergence of the value iteration algorithm, which strengthens our conjecture that this assumption is not necessary.

VII. CONCLUSIONS AND FUTURE WORK

We introduced new communication protocols for networked control systems that adjust the probability of successful communication by the transmission of redundant packets. We considered protocols that optimize an average cost criterion that seeks to improve the control performance while reducing the transmission rates. Two different types of protocols were proposed, one for nodes with reasonable computational capabilities and a much simpler one that is suitable for nodes with limited computational capabilities.

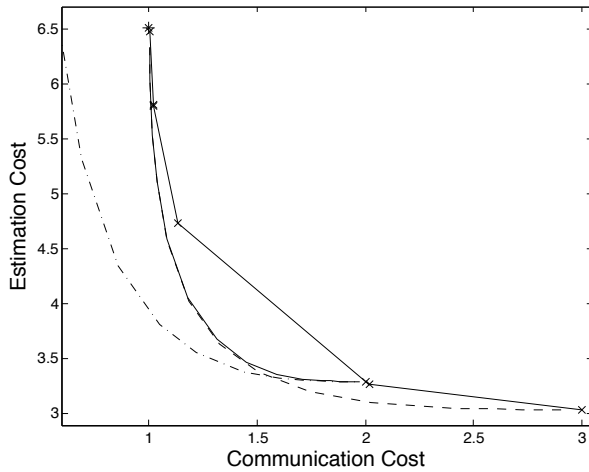


Fig. 2. Pareto Frontiers for: policy $v(k) \equiv 1$ (*); optimal policy with $v \in \{1, 2\}$ (solid); optimal policy with $v \in \{1, 2, 3\}$ (dashed); optimal policy with $v \in \{0, 1, 2\}$ (dash-dotted); simplified optimal policy with $v \in \{1, 2, 3\}$ and $T = 5$ (cross).




Fig. 3. Optimal costs ($\lambda = 2.2$) as a function of the drop probability for: policy $v(k) \equiv 1$ (o); optimal policy with $v \in \{1, 2\}$ (solid); optimal policy with $v \in \{1, 2, 3\}$ (dashed); optimal policy with $v \in \{0, 1, 2\}$ (dash-dotted); simplified optimal policy with $v \in \{1, 2, 3\}$ and $T = 5$ (cross).

measurements, multiple nodes and delays in the network by following the procedure in [6].

Future work includes considering the case when the drops for different packets are not independent of each other. This would be important to study communication faults due to collisions when this type of redundancy strategy is simultaneously employed by different nodes. One should also consider the case in which nodes do not share the same information on what was broadcasted to the network, e.g., acknowledgements are not perfect. The development of new acknowledgement mechanisms would be a valuable approach in this case. In particular, there are cases where nodes can efficiently detect the occurrence of drops through the plant

(as opposed to an acknowledgement signal in the network) as described in [12].

Another interesting variation of this problem would involve the case where packet sizes depend on the estimation error, which would arise when a finite alphabet is used to transmit measurements. In this case one can expect redundant transmissions to be even more beneficial, since high estimation errors would demand high communication rates.

REFERENCES

- [1] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [2] D. Djonin and V. Krishnamurthy, "MIMO Transmission Control in Fading Channels: A Constrained Markov Decision Process Formulation With Monotone Randomized Policies," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 55, no. 10, pp. 5069–5083, 2007.
- [3] H. Yang and M. Alouini, "Markov chains and performance comparison of switched diversity systems," *Communications, IEEE Transactions on*, vol. 52, no. 7, pp. 1113–1125, 2004.
- [4] Y. Xu and J. P. Hespanha, "Optimal communication logics in networked control systems," *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, vol. 4, 2004.
- [5] R. Cogill, S. Lall, and J. Hespanha, "A Constant Factor Approximation Algorithm for Event-Based Sampling," *American Control Conference, 2007. ACC'07*, pp. 305–311, 2007.
- [6] Y. Xu, "Communication scheduling methods for estimation over networks," Ph.D. dissertation, University of California, Santa Barbara, CA, USA, mar 2006, available at <http://www.ece.ucsb.edu/~hespanha/published>.
- [7] S. Meyn and R. Tweedie, *Markov chains and stochastic stability*. Springer, 1996.
- [8] L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. Sastry, "Foundations of Control and Estimation Over Lossy Networks," *PROCEEDINGS-IEEE*, vol. 95, no. 1, p. 163, 2007.
- [9] O. Costa, M. Fragoso, and R. Marques, *Discrete-Time Markov Jump Linear Systems*. Springer Verlag, 2004.
- [10] E. Gordienko and O. Hernandez-Lerma, "Average cost Markov control processes with weighted norms: Value iteration," *Applications Math*, vol. 23, pp. 219–237, 1995.
- [11] K. Chatterjee, "Markov decision processes with multiple long-run average objectives," in *FSTCS 2007: Foundations of Software Technology and Theoretical Computer Science*, December 2007, pp. 473–484. [Online]. Available: <http://chess.eecs.berkeley.edu/pubs/466.html>
- [12] S. Di Cairano, K. Johansson, A. Bemporad, and R. Murray, "Discrete and Hybrid Stochastic State Estimation Algorithms for Networked Control Systems," *LECTURE NOTES IN COMPUTER SCIENCE*, vol. 4981, p. 144, 2008.