Redundant Data Transmission in Control/Estimation Over Lossy Networks

Alexandre R. Mesquita[a], João P. Hespanha[a], Girish N. Nair[b]

[a] *Center for Control, Dynamical Systems and Computation, University of California, Santa Barbara, CA 93106*

[b] *Departament of Electrical and Electronic Engineering, University of Melbourne, VIC 3010, Australia*

Abstract

In lossy networks the probability of successful communication can be significantly increased by transmitting multiple copies of a same message through independent channels. In this paper we show that 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. We develop techniques to design communication protocols that exploit the transmission of multiple packets while seeking a balance between stability/estimation performance and communication rate. An average cost optimality criterion is employed to obtain a number of optimal protocols applicable to networks with different computational capabilities. We also discuss stability results under network contention when multiple nodes utilize these protocols.

Key words: Networked Control Systems; Estimation theory; Wireless communications.

1 Introduction

In the new generations of wireless communications (3G and 4G), channel adaptive techniques are able to provide large improvements under almost any performance metric. These techniques utilize an adaptive allocation of communication resources as the channel conditions change with time. This paper explores a similar idea in the context of networked control systems (NCS). In addition to compensating for the uncertainty generated by the channel, we aim at allocating communication resources to compensate for the uncertainty being generated within the controlled process. These adaptive techniques are well-suited for NCSs because they permit an increase in the reliability of communication without increasing the transmission delays, to which NCSs have low tolerance.

Adaptation can be achieved by adjusting the transmit power, by adaptive coding (an example of which is changing the quantization coarseness) and by *diversity* schemes, which correspond to the transmission of redundant signals through mostly independent channel real-

---

Email addresses: `mesquita.alexandre@gmail.com` (Alexandre R. Mesquita[a]), `hespanha@ece.ucsb.edu` (João P. Hespanha[a]), `gnair@unimelb.edu.au` (Girish N. Nair[b]).

izations. Diversity schemes may involve using multiple time slots, frequency slots, antennas or network paths [17]. While many diversity schemes are dynamically exploited in data networks by scheduling transmissions according to the network status (see e.g. [7,20]), these techniques do not take into account nor benefit from the dynamical nature of NCSs. In this paper we focus on diversity schemes and show that something as simple as transmitting multiple independent copies of the same packet can provide significant performance gains in the context of NCSs.
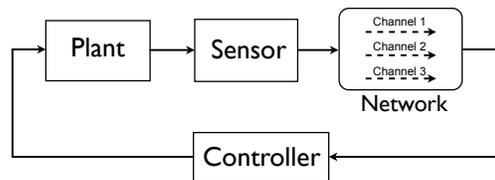
Fig. 1. NCS architecture

The adopted NCS architecture is depicted in Fig. 1, which considers the case of a single sensor and a controller. We assume that, by means of some diversity scheme, a number of independent redundant channels is available for data transmission. These are erasure channels with independent identically distributed (iid)

packet drops. At each time step, the sensor sends measurements to the controller with a certain level of redundancy. Further, by means of an acknowledgement mechanism, the sensor knows which measurements were received by the controller. 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.

The basic intuition behind this work is that one can use redundancy to increase the probability of a successful transmission whenever the estimation error in the controller becomes large. On the other hand, at time instants for which the control performance is satisfactory, one may send only one packet or not send data at all, which would save communication resources. This adaptive behavior is desirable for NCSs because it improves the reliability of transmissions without relying on error correction schemes that induce delay in the transmissions. Indeed, if a packet containing some measurement data is dropped at a time instant, it is generally more important from a control point of view to guarantee that the measurements at the next time instant are delivered, rather than to retransmit the old information that was dropped previously.

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, and delays in the network by using a Kalman filter collocated with the sensor and transmitting optimal estimates as suggested by Xu in [19]. He and several authors (e.g. [14,6,10]) have explored the idea of saving communication resources in a NCS by scheduling transmissions in a judicious way. However, these works do not consider the possibility of redundant transmissions.

Our approach can be regarded as part of the real-time rate-distortion theory introduced in [16] and further pursued in many papers including [3,11,5]. While these works share with ours a Markov Decision theory framework, their focus is mostly on the encoding problem. However, due to the nature of NCSs, data packets are typically small to the point that no practical advantage would be achieved by exploiting different encoders, that spread information among many packets. In this way, our work is distinct in that we consider erasure channels and take packets as the basic information unit.

In the first part of the paper we focus our attention on the problem of stabilizing a discrete-time linear time-invariant process with a certainty equivalence control. For such process it is well know 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 redundant transmissions can be used to stabilize any

given statistical moments for any probability of drop. Surprisingly, we show that, by a judicious use of redundant transmissions, this can be achieved with no significant increase in the average communication rate.

In the second part, motivated by the observation that the redundant channels are rarely used and yet their availability provides significant performance gains, we discuss the possibility of having multiple nodes sharing the redundant channels. This is a natural practice to maximize network efficiency but one must be aware that transmissions at the different feedback loops become interdependent. Nevertheless, we provide protocols for which mean-square stability is preserved for an arbitrary number of nodes sharing the same set of channels.

In the third part, we investigate optimal redundant transmission protocols for the NCS. In this setting, the controller constructs estimates of the process state using the measurements transmitted by the sensor. The sensor, in turn, uses a redundant transmissions protocol that minimizes the weighted average cost of the estimation error in the controller and the average communication rate. First, we consider the ideal case in which the sensor has enough computational power to 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 computational capabilities, we find protocols that minimize the same cost but base their decisions solely on the number of consecutive transmission failures.

2   A linear NCS with redundant transmissions

Throughout the paper we consider a linear time-invariant plant with

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

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 with positive definite covariance matrix Σ. The pair (A, B) is assumed to be controllable.

The controller and a sensor that measures the full state $x(k)$ are connected through a network that drops packets independently of each other, with probability $p \in (0, 1)$. The state $x(k)$ is assumed to be transmitted with negligible quantization error. In order to "adjust" the probability that the measurement $x(k)$ reaches the controller, the sensor may transmit multiple copies of this message through independent channels (see Fig. 1). The transmitter is equipped with a feedback channel that allows it to know which packets are dropped. We denote by $\ell(k)$ the number of consecutive transmission failures that occurred immediately before time k, where a *transmission*

failure is characterized by the failure of all the attempts to transmit $x(k)$ at time k.

We are interested in designing protocols that determine how many identical packets to send at time k as a function of $x(k)$ and $\ell(k)$. We denote by $v(k)$ the number of redundant packets transmitted at time k. Under our assumption of independent drops, the probability of a transmission failure at time k is equal to $p^{v(k)}$.

We adopt a certainty equivalence control law of the form

$$u(k) = K\hat{x}(k) \tag{2}$$

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 before time k. In particular,

$$\hat{x}(k) := \mathrm{E}\left[x(k)\middle|x(s),\ s < k,\ s \in \mathcal{T}_{\text{success}}\right] \tag{3}$$

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

$$\hat{x}(k+1) = \begin{cases} A\hat{x}(k) + Bu(k) & \text{if } k \notin \mathcal{T}_{\text{success}} \\ Ax(k) + Bu(k) & \text{if } k \in \mathcal{T}_{\text{success}}. \end{cases} \tag{4}$$

Subtracting (1) from (4), we conclude that the estimation error $e(k) := \hat{x}(k) - x(k)$ evolves according to the dynamics:

$$e(k+1) = \begin{cases} Ae(k) - w(k) & \text{if } k \notin \mathcal{T}_{\text{success}} \\ -w(k) & \text{if } k \in \mathcal{T}_{\text{success}}. \end{cases} \tag{5}$$

The closed-loop dynamics (1)–(2) can be expressed in terms of this error using

$$x(k+1) = (A+BK)x(k) + BKe(k) + w(k). \tag{6}$$

In the following sections we investigate stability and optimal estimation in this setting. From (6), our certainty equivalence control guarantees a bounded covariance for $x(k)$ if $e(k)$ has bounded covariance. One can verify that this stability property is preserved if we replace the static gain K by other stabilizing linear time-invariant controllers.

3 Moment Stabilization Using Packet Redundancy

In this section we investigate the stability properties of redundant transmission protocols that can be specified by a static law π that maps the number $\ell(k)$ of consecutive transmission failures to the number $v(k)$ of packets

to send, i.e., $v(k) = \pi(\ell(k))$. For example, if we use the identity function $\pi(l) = l$, then $\ell(k)$ identical packets will be sent at time k. In this paper, we do not exclude the possibility of sending zero packets.

Theorem 1 shows that mean-square stability can be achieved for any system matrix A and any drop probability $p < 1$ by a suitable choice of the redundant packet transmission protocol that specifies the function $\pi(\cdot)$.

Theorem 1 *Let the spectral radius of the matrix A be denoted by a. The covariance of $x(k)$ is bounded when the following limit exists and satisfies*

$$\lim_{l \to \infty} a^2 p^{\pi(l)} < 1 \ . \tag{7}$$

Conversely, the covariance of $x(k)$ is unbounded when $\lim_{l\to\infty} a^2 p^{\pi(l)} > 1$.

In view of Theorem 1, all that is needed to guarantee stability is to select $\pi(\ell)$ sufficiently large for large values of ℓ, namely, we need to have

$$\lim_{l \to \infty} \pi(l) > \frac{2\log a}{-\log p} \ . \tag{8}$$

Proof. In view of (6) and the fact that $A + BK$ in (6) is Schur, it is sufficient to verify the boundedness of the covariance matrix for $e(k)$. We assume that $a \geq 1$, since otherwise the process would always be stable regardless of drops. From (5), we have that

$$Q_{\ell(k)} := \mathrm{E}\left[e(k)e(k)' \mid \ell(k)\right] = \mathrm{E}\Big[\left(\sum_{m=0}^{\ell(k)} A^m w(k-m)\right)$$
$$\cdot \left(\sum_{m=0}^{\ell(k)} A^m w(k-m)\right)'\Big] = \sum_{m=0}^{\ell(k)} A^m \Sigma A'^m \ . \tag{9}$$

The process $\ell(k)$ is itself a Markov chain that takes values in the set of non-negative integers, with transition probability given by

$$\Pr\left(\ell(k+1) = l|\ell(k)\right) = \begin{cases} p^{\pi(\ell(k))} & l = \ell(k) + 1 \text{ (drop)} \\ 1 - p^{\pi(\ell(k))} & l = 0 \text{ (no drop)} \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

Let μ denote the stationary probabilities of the infinite Markov chain $\ell(k)$. By definition, we have

$$\mu(l) = \sum_{m=0}^{\infty} \Pr(\ell(k+1) = l|\ell(k) = m)\mu(m) \ . \tag{11}$$

From (10) we see that there is only one nonzero term in the above sum when $l > 0$:

$$\mu(l) = p^{\pi(l-1)}\mu(l-1) \ . \qquad (12)$$

Iterating (12) and using the fact that $\sum_{l=0}^{\infty} \mu(l) = 1$, we obtain

$$\mu(l) = p^{\sum_{m=0}^{l-1} \pi(m)}\mu(0), \ l \geq 1 \qquad (13)$$

$$\mu(0) = \left(1 + \sum_{l=1}^{\infty} p^{\sum_{m=0}^{l-1} \pi(m)}\right)^{-1} \ . \qquad (14)$$

Here we only have to consider the case $\lim_{l\to\infty} \pi(l) \geq 1$ since, in the case $\lim_{l\to\infty} \pi(l) = 0$, the process would reach an open-loop state with positive probability and become unstable. This assumption guarantees that the series in (14) converges and that μ is well defined.

Since $\pi(l) \geq 1$ for l large enough, there is a positive probability that the Markov chain $\ell(k)$ will return to the state $l = 0$ for all k larger than some constant L. This implies that the chain is aperiodic and positive recurrent in the language of [13]. In view of the ergodic theorem in [13, Thm. 14.3.3], these properties allow us to use the invariant probability μ when we take the limit

$$\lim_{k\to\infty} \mathrm{E}\left[e(k)e(k)'\right] = \lim_{k\to\infty} \mathrm{E}\left[Q_{\ell(k)}\right] = \sum_{l=0}^{\infty} \mu(l)Q_l \ . \qquad (15)$$

From this and (9), we conclude that

$$\lim_{k\to\infty} \mathrm{E}\left[e(k)e(k)'\right] = \sum_{l=0}^{\infty} \mu(l) \sum_{m=0}^{l} (A^m \Sigma A'^m) \ . \qquad (16)$$

For any submultiplicative matrix norm $\|\cdot\|$, (16) gives

$$\|\lim_{k\to\infty} \mathrm{E}\left[e(k)e(k)'\right]\| \leq \|\Sigma\| \sum_{l=0}^{\infty} \mu(l) \sum_{m=0}^{l} (\|A^m\|^2) \ . \qquad (17)$$

By Cauchy's root test, this series is convergent if

$$\lim_{l\to\infty} \mu(l)^{\frac{1}{l}} \|A^l\|^{2/l} = a^2 \lim_{l\to\infty} p^{\pi(l)} < 1 \ , \qquad (18)$$

where the equality comes from (13) and the fact that $a = \lim_{k\to\infty} \|A^k\|^{1/k}$. This gives the first part of the theorem. The second part can be deduced by pre- and post-multiplying (16) by the eigenvector corresponding to a and then using the root test to conclude divergence of the resulting series. ∎

For the stability of higher moments, one can obtain conditions analogous to (8) with a similar proof.

Remark 1 *One factor that may alter our results is the presence of failures in the acknowledgement mechanism. While there may be many ways to deal with this issue, such as using redundancy in the acknowledgement channel, one can always adopt the conservative approach of using positive acknowledgements. This approach preserves stability since it sends necessarily more packets than in the case of perfect acknowledgements.*

From (8), we can see that to achieve stability one may require a protocol that, at times, sends a large number of packets, which seems to require a large communication rate. To verify that this is not the case, we investigate the expected communication rate for a given function $\pi(\cdot)$. We assume that the packet size is a constant with value 1 and that it is sufficiently large so that the controller receives $x(k)$ with negligible quantization loss.

Theorem 2 *Suppose that $a^2 p^M < 1$ for some integer M. Then, for every integer $N \geq 0$, there exists a protocol with $\pi(l) \leq M$ that stabilizes the covariance of $x(k)$ with an expected asymptotic transmission rate:*

$$\bar{R} := \lim_{L\to\infty} \frac{1}{L} \sum_{k=0}^{L-1} \mathrm{E}[\pi(\ell(k))] = O(1/N) \ , \qquad (19)$$

which can be made arbitrarily small by choosing N sufficiently large.

While we can obtain an arbitrarily small communication rate, the larger we make N the larger the error covariance will be. This relationship between average transmission rate and control performance is the subject of Sections 5 and 6.

Proof. Consider the protocol

$$\pi(l) = \begin{cases} 0 & \text{for } l < N \\ M & \text{for } l \geq N \ . \end{cases} \qquad (20)$$

From Theorem 1, this is a stabilizing protocol. As in the proof of Theorem 1, we use the ergodicitiy of the chain to compute the expected asymptotic transmission rate as $\bar{R} = \sum_{l=0}^{\infty} \mu(l)\pi(l)$. Then, (13)-(14) gives

$$\bar{R} = \mu(0)\left(v(0) + \sum_{l=1}^{\infty} \pi(l) p^{\sum_{m=0}^{l-1} v(m)}\right) \ . \qquad (21)$$

Substituting (20), we have that

$$\bar{R} = \frac{M}{N(1 - p^M) + 1} \ , \qquad (22)$$

which gives the result in the theorem. ∎

A transmission rate as slow as (19) may not be possible when the quantization errors cannot be neglected. This is because, when many consecutive failures take place, one may need to increase the packet size to maintain a "reasonable" quantization error in spite of an exponential growth in the state.

4   Redundant Protocols with Multiple Nodes

Theorem 2 indicates that redundant protocols may need to utilize redundant channels quite rarely. This suggests that some of these channels could be shared by other feedback loops without compromising performance. In this section we explore the situation in which multiple processes share the redundant channels.

Consider the case where S feedback loops (also referred here as nodes) share M iid channels. We say that a protocol that determines the number $v_i(k)$ of redundant transmissions used by the node i as a function of the number of consecutive failures $\ell_i(k)$ is *L-drop triggered* if it satisfies the following property:

$$\mathcal{S}(k) := \big\{ i \in \{1, \ldots, S\} : \ell_i(k) > L \big\} \neq \emptyset$$
$$\Rightarrow \sum_{i \in \mathcal{S}(k)} v_i(k) = M \ , \qquad (23)$$

i.e. the nodes $\{i : \ell_i(k) > L\}$ that experience more than L consecutive failures share all the M channels among themselves.

L-drop triggered protocols do not necessarily share the M channels fairly. However, the condition (23) does require these protocols to efficiently utilize the M channels without collisions. Condition (23) also forbids nodes with less than L consecutive failures to transmit any messages when some loops have suffered more than L failures. This would be possible either for a protocol that assigns the $v_i(k)$ in a centralized fashion or for a decentralized protocol that takes advantage of the *capture effect* [21] so that, even if multiple nodes access the same channel, one of them is still able to successfully use it.

We are interested in evaluating under what conditions L-triggered protocols stabilize the estimation error in the mean-square sense. Somewhat surprisingly, we show next that stability can be achieved independently of the number of processes S in the network. Let A_i be the system matrix of the i-th process and a_i its spectral radius. Then we define $a = \max_{i \leq S} \{a_i\}$.

Theorem 3 *Suppose that M channels are available and denote by p the dropout probability for each independent channel and assume that $a^2 p^M < 1$. Then, for any number of feedback loops S, there always exists an integer L such that any L-drop triggered protocol stabilizes the estimation error in the mean-square sense.*

Proof. We start by assuming that $L > 2S$ and $a > 1$. For Q_ℓ given in (9), we have that $\|Q_\ell\| \leq \|\Sigma\|(1-a^2)^{-1}(1 - a^{2\ell})$ and $\mathrm{E}[e(k)e(k)'] = \mathrm{E}[Q_{\ell(k)}]$. Then, to prove mean-square stability, it suffices to prove the boundedness of $E[W(\ell(k))]$, where the Lyapunov function $W(\ell)$ is defined as $W(\ell) = a^{2V(\ell)}$ and $V(\ell) := \max_{j \leq S} \ell_j$. To this purpose, we evaluate $W^+(l) := E[W(\ell(L)) \mid \ell(0) = l]$.

By noticing that $V(\ell)$ may increase by as much as L in L time steps, we can write

$$W^+(l) \leq \beta(l) a^{2(V(l)+L)} + (1 - \beta(l)) \cdot a^{4L}$$
$$= a^{2L} \beta(l) W(l) + (1 - \beta(l)) \cdot a^{4L} \ , \qquad (24)$$

where $\beta(l)$ is the probability that $\{V(\ell(L)) > 2L\}$ given the initial condition l. Denoting by P_L the maximum value of $\beta(l)$ over l, we have from (24) that

$$W^+(l) \leq a^{2L} P_L W(l) + a^{4L} \ . \qquad (25)$$

Taking expected values on both sides with respect to l and using the time-homogeneity of the chain, we conclude that

$$\mathrm{E}\left[W(\ell((k+1)L))\right] \leq a^{2L} P_L \, \mathrm{E}[W(\ell(kL))] + a^{4L} \quad (26)$$

Therefore, $\mathrm{E}[W(\ell(k))]$ is bounded provided that

$$a^2 P_L^{1/L} < 1 \ . \qquad (27)$$

Now we compute an upper bound P_L on the probability of the event $Z := \{V(\ell(L)) > 2L \mid \ell(0) = l\}$. The worst case occurs when $l_j > L, \forall j$, since then all nodes have to transmit successfully at least once to prevent the event Z from happening. Since, for an L-drop triggered protocol, each node can transmit successfully only once in L time steps (unless $\max_j \ell_j$ becomes less than L), Z, when $l_j > L, \forall L$, corresponds to the occurrence of less than S successful transmissions in L time steps. As we only care about the total number of transmissions regardless of which nodes transmit, the probability of less than S transmissions in L time steps is maximized for the L-drop triggered protocol in which one node takes over all M channels in a given time step (thus transmitting with probability $1 - p^M$). Therefore, we have that

$$P_L = \sum_{i < S} \Pr\{i \text{ nodes transmit in } L \text{ steps}\}$$
$$= \sum_{i < S} \binom{L}{i} (1 - p^M)^i \, p^{M(L-i)}$$
$$\leq (S - 1) \binom{L}{S-1} p^{M(L-S+1)} \ , \qquad (28)$$

where the inequality is obtained by majorizing each element of the sum (using the fact that $S < L/2$ to majorize the binomial). Taking the limit $L \to \infty$, we have

$$\lim_{L \to \infty} P_L^{\frac{1}{L}} \le \lim_{L \to \infty} \binom{L}{S-1}^{\frac{1}{L}} p^{M(L-S+1)/L} = p^M \ . \tag{29}$$

Since $a^2 p^M < 1$, one can always choose L large enough to satisfy (27). This proves the theorem for $a > 1$. For $a = 1$, we can use a dominance argument noting that there exists $\bar{a} > 1$ such that $\bar{a}^2 p^M < 1$. ∎

For scenarios in which capacity is degraded by collisions, a stability result will depend on the specifics of the communication system. As in (27) in the proof of Theorem 3, stability will depend on the quantity $\lim_{L \to \infty} P_L^{1/L}$, which is closely related to the notion of anytime reliability in [15].

5 Optimal Communication Protocols

In the last two sections, we have seen that it is possible to stabilize one or several processes in the mean-square sense, with very few communication resources. However, this may lead to large error covariances. Our goal now is to determine an optimal protocol for the single node case that decides how many copies of the same packet to send. This protocol should be optimal in the sense that it achieves an optimal trade off between the conflicting objectives of keeping small the estimation error $e(k)$ that drives the closed-loop dynamics (6) while achieving this with a minimal amount of communication.

To compute the optimal communication protocol, we start by noting that the estimation error $e(k)$ can be seen as a Markov Decision Process in the sense of [9], with transition probabilities controlled by the number $v(k)$ of redundant packets. In particular,

$$e(k+1) = \begin{cases} Ae(k) - w(k) & \text{with prob. } p^{v(k)} \\ -w(k) & \text{with prob. } 1 - p^{v(k)}. \end{cases} \tag{30}$$

We consider protocols that decide how many packets to send based on the current value of the estimation error $e(k)$. Such protocols can be viewed as functions π that define $v(k) = \pi(e(k))$, $\forall k$. The state estimation error $e(k)$ can be calculated by the sensor using (4) and $\ell(k)$.

Our objective is the minimization of the following average cost (AC) criterion

$$J(\pi, e_0) := J_{\text{est}}(\pi, e_0) + \lambda J_{\text{com}}(\pi, e_0) \tag{31}$$

where

$$J_{\text{est}}(\pi, e_0) := \lim_{N \to \infty} \frac{1}{N} \mathrm{E}_{e_0}^\pi \left[\sum_{k=0}^{N-1} e(k)' Q e(k) \right] \tag{32}$$

$$J_{\text{com}}(\pi, e_0) := \lim_{N \to \infty} \frac{1}{N} \mathrm{E}_{e_0}^\pi \left[\sum_{k=0}^{N-1} v(k) \right] \tag{33}$$

where λ is a positive scalar, Q a positive definite matrix and $\mathrm{E}_{e_0}^\pi$ denotes the expectation given a protocol π and an initial state $e(0) = e_0$.

The criterion in (31) is a weighted sum of two terms: the first term $J_{\text{est}}(\pi, e_0)$ penalizes a time-averaged expected quadratic estimation error, 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 $\lambda \to 0$, communication is not penalized, whereas as $\lambda \to \infty$, communication is heavily penalized. Intermediate values of λ will yield Pareto-optimal compromise solutions between the two conflicting criteria.

Defining the one-step cost

$$c(e, v) := e'Qe + \lambda v \ , \tag{34}$$

we can rewrite the cost function (31) as follows

$$J(\pi, e_0) = \lim_{N \to \infty} \frac{1}{N} \mathrm{E}_{e_0}^\pi \sum_{k=0}^{N-1} c(e(k), v(k)) \ . \tag{35}$$

For technical reasons, we restrict the optimization problem to the set of protocols that satisfy

$$\pi(e) \in \Pi(e) := \begin{cases} \{0, \ldots, M\} & \text{if } \|e\| < L \\ \{M\} & \text{if } \|e\| \ge L \ , \end{cases} \tag{36}$$

where M denotes the maximum number of redundant packets possible and $L > 0$ is a constant. This restriction ensures that M packets are sent when $\|e\|$ grows as large as L, but our results allow L to be arbitrarily large. The set of protocols that obey (36) is denoted by Γ.

A protocol $\pi^* \in \Gamma$ is said to be *AC-optimal* if

$$J(\pi^*, e_0) = \inf_{\pi \in \Gamma} J(\pi, e_0) =: J^*(e_0), \ \forall e \in \mathbb{R}^n, \tag{37}$$

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

Next we provide a brief review of results in average cost optimization. The reader is refered to [9] for an in depth treatment. For functions $\varphi : \mathbb{R}^n \to \mathbb{R}$, we denote by

$T^{\bar{v}}\varphi(\bar{e})$ the expected value of $\varphi(e(k+1))$ given $e(k) = \bar{e}$ and $v(k) = \bar{v}$, i.e.,

$$T^{\bar{v}}\varphi(\bar{e}) := \mathrm{E}[\varphi(e(k+1)) \mid e(k) = \bar{e}, \ v(k) = \bar{v}] \ , \quad (38)$$

and by $\Delta^{\bar{v}}\varphi(e)$ we denote the expected one-step variation of the function φ:

$$\Delta^{\bar{v}}\varphi(e) = T^{\bar{v}}\varphi(e) - \varphi(e) \ . \quad (39)$$

It follows readily from (30) that

$$T^{\bar{v}}\varphi(e) = \int_{\mathbb{R}^n} \varphi(e)[(1 - p^{\bar{v}})f(e)de + p^{\bar{v}}f(e - A\bar{e})de] \ , \quad (40)$$

where f is the probability density of the normal distribution with zero mean and covariance Σ.

For a protocol π, suppose that there exist a function φ and a constant ϱ such that

$$\varrho \geq c(e, \pi(e)) + \Delta^{\pi(e)}\varphi(e) \ . \quad (41)$$

It follows from [13, Thm. 14.2.2] that $J(\pi, e) \leq \varrho$, i.e, ϱ is an upper bound on the average cost achieved by the protocol π.

Bellman's equation for the average cost problem takes the form

$$\varrho^* = \min_{v \in \Pi(e)} [c(e, v) + \Delta^v \varphi^*(e)]$$
$$=: c(e, \pi^*(e)) + \Delta^{\pi^*(e)}\varphi^*(e), \ \forall e \in \mathbb{R}^n \ , \quad (42)$$

where φ^* is known as the relative value function and the average cost is given by the constant ϱ^*. Under suitable conditions, (42) implies that π^* is AC-optimal [1, Thm. 5.1].

A possible solution method for (42) is the value iteration algorithm below. Let s_N be defined as follows:

$$s_N(e) := \min_{v \in \Pi(e)} [c(e, v) + T^v s_{N-1}(e)] \ , \quad (43)$$

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

$$j_N := s_N(z) - s_{N-1}(z) \ \text{and} \ \varphi_N(e) := s_N(e) - s_N(z) \ . \quad (44)$$

Then, the value iteration algorithm is said to converge if

$$\lim_{N \to \infty} j_N \to \varrho^* \ \text{and} \ \lim_{N \to \infty} \varphi_N(e) \to \varphi^*(e), \ \forall e, \quad (45)$$

for some φ^* and ϱ^* satisfying (42).

Our next theorem states the existence of a solution to (42) and of an AC-optimal protocol as long as one has enough channels to stabilize the process.

Theorem 4 *Suppose that the number of channels M in (36) is sufficiently large so that*

$$a^2 p^M < 1 \ . \quad (46)$$

Then:

(1) *There exists a constant $\varrho^* \geq 0$, a continuous function φ^* and a protocol $\pi^* \in \Gamma$ that satisfy equation (42).*
(2) *π^* is AC-optimal and ϱ^* is the optimal AC-function (which happens to be constant).*
(3) *The given value iteration algorithm converges.*

The proof of this is based on results from [18] and involves finding a common Lyapunov function for all protocols in Γ. Due to space constraints, we provide this proof in [12, Theorem 19]. Furthermore, in this reference we show that the optimal protocol given by the theorem is optimal even if time-varying and randomized protocols are considered.

5.1 Suboptimal Protocols

Solving for the optimal protocol in Theorem 4 may be computationally intense for high-dimensional systems. Fortunately, one can exploit the linear structure of the controlled plant to find suboptimal protocols that are computationally tractable. To this purpose, we consider quadratic relative value functions of the form $\bar{\varphi} = e'He$, where H is a positive definite matrix. As in (41), an upper bound on the optimal cost is given by the constant $\bar{\varrho}$ that satisfies

$$\bar{\varrho} \geq \min_{v \in \Pi(e)} [c(e, v) + \Delta^v \bar{\varphi}(e)]$$
$$= \min_{v \in \Pi(e)} [e'Qe + \lambda v + p^v e' A' H A e - e' A' H A e + \mathrm{tr}\, H\Sigma] \ . \quad (47)$$

While one could search numerically for a matrix H for which (47) holds with the smallest possible $\bar{\varrho}$, our numerical examples in Section 7 were obtained by simply taking $H = H_s$, where H_s is the solution to the Lyapunov equation $p^M A' H_s A - H_s + Q = 0$. The motivation for this choice is that $\varphi_s = e'H_se$ approaches the solution to (42) as e grows to infinity (or as λ approaches zero). Alternatively, φ_s could also be obtained as a result of applying just one iteration of the policy iteration algorithm (see [9]) initialized with the protocol $\pi_0(e) \equiv M, \ \forall e$.

From (47), the resulting suboptimal protocol is

$$\pi_s(e) = \arg\min_{v \in \Pi(e)} \{\lambda v + p^v e' A' H_s A e\} \ , \quad (48)$$

The protocol π_s is considerably simpler to compute than π^* and it offers the additional benefit of a smaller (memory-wise) representation.

6   A Simplified Optimal Protocol

In many applications, sensors have limited computational capabilities that could prevent the use of elaborate protocols that require the computation of estimation errors as those considered in Section 5. To address this issue, we can design a simplified protocol that bases its decision rule only on the consecutive number of failures $\ell(k)$ that occurred prior to the k-th sampling time, much like the protocols considered in Section 3.

In general, this approach would lead to a partially observable MDP. Fortunately, due to the special structure of the problem, we show that it is enough to consider a fully observable process. We consider two controlled Markov chains, one whose state is simply $\ell(k)$ and another one whose state is the pair $\big(e(k), \ell(k)\big)$. Let Γ_ℓ denote the set of stabilizing feedback protocols for which $v(k)$ depends only on $\{\ell(s); s \leq k\}$. For $\pi \in \Gamma_\ell$ and the chain $\ell(k)$, define the average cost

$$J_\ell(\pi, \ell_0) = \lim_{N \to \infty} \frac{1}{N} \, \mathrm{E}_{\ell_0}^\pi \sum_{k=1}^{\infty} \bar{c}(\ell(k), v(k)) \ , \qquad (49)$$

where the one-step cost is given by

$$\bar{c}(l, v) = \mathrm{tr}(Q \Sigma_l) + \lambda v \ , \qquad (50)$$

where

$$\Sigma_l := \sum_{m=0}^{l} A'^m \Sigma A^m \ . \qquad (51)$$

Theorem 5 *For every protocol $\pi \in \Gamma_\ell$,*

$$J(\pi, e_0) = J_\ell(\pi, \ell_0), \ \forall e_0, \ell_0 \ , \qquad (52)$$

where J is the original cost in (31). Therefore, an AC-optimal protocol for the cost J_ℓ and the chain $\ell(k)$ is also an AC-optimal protocol for the cost J and the chain $(e(k), \ell(k))$ within the set of protocols Γ_ℓ.

Proof. We can rewrite our cost criterion (35) as

$$J(\pi, e_0) = \lim_{N \to \infty} \frac{1}{N} \, \mathrm{E}_{e_0}^\pi \sum_{k=1}^{\infty} \mathrm{E}_{e_0}^\pi \big[c(e(k), v(k)) \mid \ell(k) \big] \ . \qquad (53)$$

Once a packet is successfully transmitted, the belief (conditional probability) of $e(k)$ is solely given by $f(e)$, through (30), 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 calculate the average costs assuming that $\ell_0 = 0$ and e_0 has distribution $f(\cdot)$. For this initial condition, $e(k)$ is conditionally distributed as $\sum_{m=0}^{\ell(k)} A^m(k)\omega_m$ given $\ell(k)$, where ω_m are i.i.d. variables with density $f(\cdot)$. Then, $\mathrm{E}_{e_0}^\pi \big[c(e(k), v(k)) \mid \ell(k) = l \big] = \bar{c}(l, v)$ and the claim of the theorem follows from (53) and (49). ∎

If we further restrict the set of protocols to be such that $v = M$ for $\ell \geq T$, we no longer need to keep track of the number of consecutive drops $\ell(k)$ when it exceeds T. We can therefore truncate the Markov chain $\ell(k)$ by redirecting the jumps $(\ell = T) \to (\ell = T + 1)$ to $(\ell = T) \to (\ell = T)$. Thus, we have moved from the infinite dimensional problem in Section 5 to an optimization in a finite Markov chain. Using the per-stage cost $\bar{c}$ in (50) and the transition probabilities for ℓ in (10), one can calculate AC-optimal protocols that depend on ℓ only. This could be done either via dynamic programming [2] or via direct optimization, since the average costs can be directly calculated using the stationary distribution as in Section 3. Interestingly, once $\bar{c}$ is known, the complexity of solving for the optimal protocol does not depend on the dimension n of the dynamical system, but only on T and on the size $M + 1$ of the set of control actions.

In general, the designer can select the constants M and T to be small. This is so because the probability of failure p^M becomes indistinguishably small for M larger than some small constant, typically 3. Likewise, the probability of reaching a state $\ell = T$ also decreases exponentially with T. An important consequence of the set of control actions being small is that one can solve for the optimal protocol offline and then use small look-up tables. The same observations also apply to the optimal protocols in the previous section.

Remark 2 *When solving an optimal control problem with quadratic costs, the separation principle does not hold in the setting of Section 5 since the communication protocols may depend on $x(k)$. On the other hand, the separation principle does hold in the context of the simplified protocols of this section. To see why, note that, for each fixed simplified protocol, we have a Markov Jump Linear System, for which the separation principle is known to hold [8]. As it turns out that the optimal control and optimal estimator do not depend on the fixed protocol, the separation principle holds for the general problem.*

7   Numerical Examples

It is common to guide autonomous vehicles indoors based on global position information captured by sensors (e.g., cameras) mounted off-board. The results in the paper are exemplified in this context, by studying the problem of estimating the position of a vehicle based on the acceleration provided by on-board sensors and

"occasional" global position measurements from an off-board camera network (with communication through an unreliable camera network). We consider the two dimensional example with $A = [1\ 1; 0\ 1]$, $\Sigma = [0.75\ 0; 0\ 0.75]$, $Q = [1\ 0; 0\ 1]$, $p = 0.15$ and $L = 10$, where x represents the position and the velocity of a vehicle moving on the line, u represents accelerations (due to control and external forces) measured by the inertial sensors and w represents measurement noise in the inertial sensors. The state information transmitted by the camera network is assumed to have much higher precision compared with the one provided by the on-board sensors.

By varying λ from 10^{-5} to 10, we constructed the Pareto frontiers for the protocols described in Sections 5 and 6. The curves are shown in Fig. 2, where protocols are restricted to the different action sets $v \in \{1, 2\}$, $v \in \{1, 2, 3\}$ and $v \in \{0, 1, 2\}$. Some important observations can be deduced from this figure:

(1) The simplified optimal protocol discussed in Section 6 produces protocols that can be quite close to the Pareto-optimal boundary.
(2) Increasing the maximum number of redundant packets beyond 2 hardly improves the Pareto-optimal boundary.
(3) If we 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 (dash-dotted line in Fig. 2).

A phenomenon that commonly arises in multi-objective MDPs is that points on the Pareto frontier do not always correspond to deterministic protocols. This is the case for the Pareto frontier of the simplified protocol, where only the points marked with a cross correspond to deterministic protocols and the lines linking those points correspond to randomized protocols that can be derived from the deterministic ones as explained in [4].

Suboptimal approaches have been considered in Fig. 3. To show the performance improvement that arises from judiciously sending redundant information, we considered also the baseline protocol that always sends one packet per time step. Some observations from this figure are:

(1) The suboptimal protocol π_s in (48) gives a performance remarkably close to the Pareto frontier and is significantly better than the simplified protocols.
(2) Using the trivial protocol $v(k) = 1$, $\forall k$, leads to less communication (x-axis) than the protocols that use $v \in \{1, 2\}$ and $v \in \{1, 2, 3\}$, but this is at the expense of a significantly larger estimation error (y-axis). In fact, based on the results of Section 3, we know that for unstable systems and large drop probabilities, $v(k) \equiv 1$ can lead to instability.
(3) The steep slope of the Pareto frontier at communication cost 1 indicates that a large percentual de-

crease in estimation cost can be obtained with a small percentual increase in communication cost.
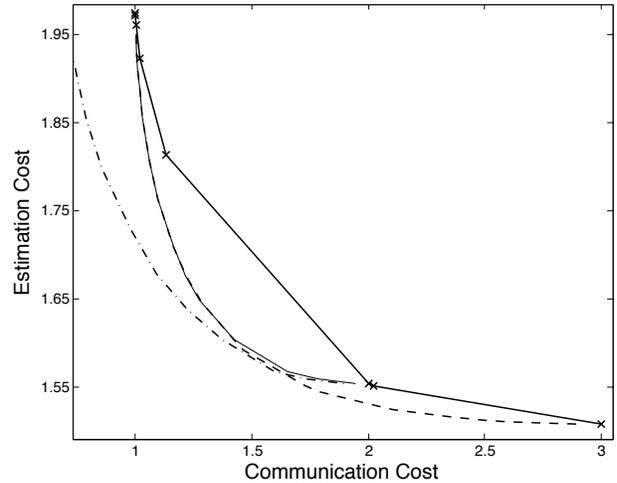
Fig. 2. Pareto Frontiers for: optimal protocol with $v \in \{1, 2\}$ (solid); optimal protocol with $v \in \{1, 2, 3\}$ (dashed); optimal protocol with $v \in \{0, 1, 2\}$ (dash-dotted); simplified optimal protocol with $v \in \{1, 2, 3\}$, $\ell < T = 5$, and $v = 3$ for $\ell \geq T = 5$ (cross).
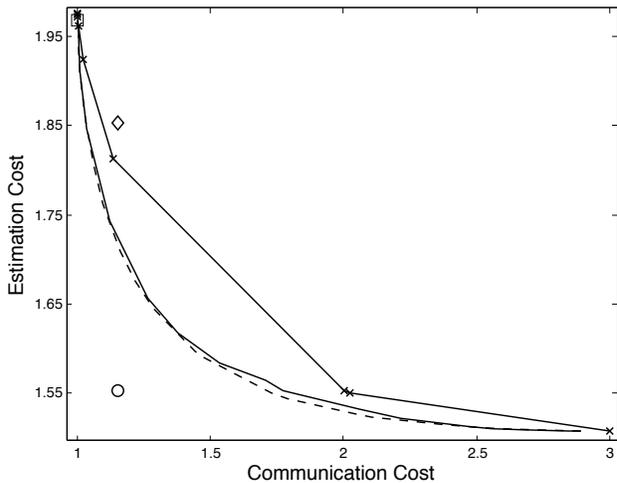
Fig. 3. Optimal vs. suboptimal protocols: optimal protocol with $v \in \{1, 2, 3\}$ (dashed); simplified optimal protocol with $v \in \{1, 2, 3\}$, $\ell < T = 5$, and $v = 3$ for $\ell \geq T = 5$ (cross); suboptimal protocol π_s with $v \in \{1, 2, 3\}$ (solid); protocol $v(k) = 1$ $\forall k$ (square); TCP with no delay (circle); TCP with delayed retransmission (diamond).

7.1 Redundant Transmissions vs. Retransmissions

The motivation for using redundant transmissions as opposed to retransmissions (such as in TCP) is that NCSs performance may suffer due to the delay accumulated in the transmission, error detection, acknowledgements and retransmission.

To provide some intuition on how the protocols described here compare with TCP we consider two idealized TCP-like protocols. In the first protocol, we consider a TCP-like scheme in which it is always possible to perform one transmission of $x(k)$ and, if necessary, at most one retransmission before the sampling time $k+1$. At time $k+1$, the measurement $x(k+1)$ is available at the sensor for transmission and the old measurement $x(k)$ is discarded (regardless of whether it was successfully transmitted). In the second scheme, a retransmission of $x(k)$ is only received at time $k+1$ so it cannot be used until time $k+2$. Hence, the retransmitted message is only used if the message containing $x(k+1)$ is dropped. We assume that original and retransmitted packets are dropped with i.i.d. probabilities with the same value as in the redundant channels above.

The results are shown in Fig. 3. The delay free TCP scheme outperforms the redundant transmission protocols for a large region of cost combinations. On the other hand, a one-time-step delay is already enough to have the TCP scheme outperformed by the redundant protocols. In practice, many things can go wrong with TCP such as random delays and correlation between drops, which in our framework would imply a higher probability that the retransmitted packets are not used (dropped). Irrespectively of these implementation issues, redundant protocols offer the advantage of flexibility in the choice of different estimation and communication costs combinations by selecting different operating points within the Pareto boundary.

8   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. These results can be readily extended to the case of partial state measurements and delays in the network by following the procedure in [19]. Our results also suggest that redundant channels may be efficiently shared among multiple feedback loops.

The proposed technique has a diminishing returns property in the sense that little additional benefits are obtained by increasing the number of redundant channels beyond two or three. This implies that, as opposed to what one could expect, implementing a large number of channels is not necessary to obtain significant performance gains using our technique.

Future work includes constructing other forms of adap-

tive channel techniques such as varying the transmitter power or using different codes depending on the history of communications between nodes. It is also important to address 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.

References

[1] A. Arapostathis, V.S. Borkar, E. Fernández-Gaucherand, M.K. Ghosh, and S. Marcus. Discrete-time controlled Markov processes with average cost criterion: a survey. *SIAM Journal on Control and Optimization*, 31(2):282–344, 1993.

[2] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1 and 2. Athena Scientific, Belmont, MA, 1995.

[3] V.S. Borkar, S.K. Mitter, A. Sahai, and S. Tatikonda. Sequential source coding: an optimization viewpoint. In *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*, pages 1035–1042. IEEE, 2005.

[4] K. Chatterjee. Markov decision processes with multiple long-run average objectives. In *FSTTCS 2007: Foundations of Software Technology and Theoretical Computer Science*, pages 473–484, December 2007.

[5] P.A. Chou and Z. Miao. Rate-distortion optimized streaming of packetized media. *IEEE Transactions on Multimedia*, 8(2):390–404, 2006.

[6] R. Cogill, S. Lall, and J.P. Hespanha. A Constant Factor Approximation Algorithm for Event-Based Sampling. *American Control Conference, 2007. ACC'07*, pages 305–311, 2007.

[7] D.V. Djonin and V. Krishnamurthy. MIMO Transmission Control in Fading Channels A Constrained Markov Decision Process Formulation With Monotone Randomized Policies. *IEEE Transactions on Signal Processing*, 55(10):5069–5083, 2007.

[8] O.L. do Valle Costa, M.D. Fragoso, and R.P. Marques. *Discrete-time Markov jump linear systems*. Springer Verlag, 2005.

[9] E.A. Feinberg and A. Shwartz. *Handbook of Markov decision processes: methods and applications*. Kluwer Academic Publishers, 2002.

[10] O.C. Imer and T. Basar. Optimal estimation with limited measurements. *International Journal of Systems, Control and Communications*, 2(1):5–29, 2010.

[11] A. Mahajan and D. Teneketzis. Optimal design of sequential real-time communication systems. *IEEE Transactions on Information Theory*, 55(11):5317–5338, 2009.

[12] A.R. Mesquita. *Exploiting Stochasticity in Multi-agent Systems*. PhD thesis, University of California, Santa Barbara, CA, USA, Dec. 2010. Available at `http://www.ece.ucsb.edu/~hespanha/published`.

[13] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Springer, 1996.

[14] K.J. Åström and B.M. Bernhardsson. Comparison of Riemann and Lebesgue sampling for first order stochastic systems (I). In *IEEE Conference on Decision and Control*, volume 2, pages 2011–2016. IEEE; 1998, 2002.

[15] A. Sahai and S. Mitter. The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link. Part I: Scalar systems. *IEEE Transactions on Information Theory*, 52(8):3369–3395, 2006.

[16] S. Tatikonda and S. Mitter. The Sequential Rate Distortion Function and Joint Source-Channel Coding with Feedback. In *Proc. of the Annual Allerton Conf. on Communication, Control and Computing*, volume 41, pages 191–200, 2003.

[17] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[18] O. Vega-Amaya. The average cost optimality equation: a fixed point approach. *Bol. Soc. Mat. Mexicana*, 9(1):185–195, 2003.

[19] Y. Xu. *Communication scheduling methods for estimation over networks*. PhD thesis, University of California, Santa Barbara, CA, USA, Mar 2006.

[20] H.C. Yang and M.S. Alouini. Markov chains and performance comparison of switched diversity systems. *IEEE Transactions on Communications*, 52(7):1113–1125, 2004.

[21] M. Zorzi and R.R. Rao. Capture and retransmission control in mobile radio. *IEEE Journal on Selected Areas in Communications*, 12(8):1289–1298, 1994.