Witsenhausen's counterexample as Assisted Interference Suppression

Pulkit Grover and Anant Sahai*

Wireless Foundations, Department of EECS,
University of California at Berkeley,
CA 94720, USA
E-mail: pulkit@eecs.berkeley.edu
E-mail: sahai@eecs.berkeley.edu
*Corresponding author

Abstract: Motivated by the existence of an implicit channel in Witsenhausen's counterexample, we formulate a vector extension that can be viewed as a toy wireless communication problem "Assisted Interference Suppression" (AIS). Information-theoretic lower and upper bounds (based respectively on ideas from rate-distortion theory and dirty-paper coding) are then derived on the optimal cost and the asymptotic optimal cost is characterised to within a factor of 2 regardless of the problem parameters. Restricted to the scalar problem, it is shown that the new lower bound can be better than Witsenhausen's bound by an arbitrarily large factor.

Keywords: Witsenhausen's counterexample; distributed control; information theory.

Reference to this paper should be made as follows: Grover, P. and Sahai, A. (2010) 'Witsenhausen's counterexample as Assisted Interference Suppression', *Int. J. Systems, Control and Communications*, Vol.

Biographical notes: Pulkit Grover (BS '03, IIT Kanpur, MS '05, IIT Kanpur) is a graduate student at the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley since 2005. He is a member of the Wireless Foundations Center. His research interests are in information theory, computational complexity, distributed control, spectrum sharing and finance.

Anant Sahai (BS '94 UCB, MS '96 MIT, PhD '01 MIT) joined the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley in 2002 as an Assistant Professor. He is a member of the Berkeley Wireless Research Center and the Wireless Foundations Center. His research interests are in wireless communication, spectrum sharing, signal processing, information theory, and distributed control. He is particularly interested in bridging the gulf between information theory and distributed control.

1 Introduction

Distributed control problems have proved challenging for control engineers. Witsenhausen (1968) gives a counterexample showing that even a seemingly simple distributed control problem can be hard to solve. For the counterexample, Witsenhausen chose a two-stage distributed LQG system. The choice is appropriate because LQG systems that are *not* distributed are well understood – control laws linear in the observations are optimal for such systems. For his counterexample, Witsenhausen provided a nonlinear control strategy that outperforms all linear laws, thus arguing that distributed systems might be harder to understand. It is now clear that Witsenhausen's problem itself is quite hard; the optimal strategy and the optimal costs for the problem are still unknown. The non-convexity of the problem makes the search for an optimal strategy hard (Bansal and Basar, 1987; Baglietto et al., 1997; Lee et al., 2001). Discrete approximations of the problem (Ho and Chang, 1980) are even NP-complete (Papadimitriou and Tsitsiklis, 1986).

In order to understand what makes the problem hard, modifications of the problem have been considered that still admit linear optimal solutions. In Bansal and Basar (1987), for example, Bansal and Basar consider a parametrised family of two-stage stochastic control problems. The family includes the Witsenhausen counterexample. Using results from information theory, they show that for this family, whenever the cost function does not contain a product of two decision variables, affine control laws are still optimal. Rotkowitz (2006) shows that affine control laws continue to be optimal for a deterministic variant of the counterexample if the cost function is the induced two-norm instead of Witsenhausen's original expected two-norm.

Ho et al. (1978) view the Witsenhausen problem as a close sibling of Shannon's problem of explicit communication. Mitter and Sahai (1999) build upon this observation to suggest that the counterexample might be hard because there is an *implicit* communication channel from the first controller to the second. The first stage cost can be interpreted as the power that is input to this channel. The second stage cost is the distortion in estimating the state at time 1. Using this interpretation, they propose nonlinear control strategies based on state quantisation. Over a range of possible problem parameters and corresponding quantisation levels, these nonlinear strategies can outperform optimal linear strategies by an arbitrarily large factor. Martins (2006) points out that even in the presence of a non-perfect (but still pretty good) explicit communication link, nonlinear strategies continue to outperform linear ones.

In this paper, we ask the following question: Can the problem actually be simplified by using the observation that it has an implicit communication channel (Mitter and Sahai, 1999)? Experience from information theory suggests that it is often easier to analyse communication problems with asymptotically long vector lengths, because they allow us to avoid the complications of the geometry of finite-dimensional spaces. The asymptotically-long vector length setting also permits the use of the laws of large numbers. However, as pointed out[1] by Tatikonda (2000), in control problems such limits need to be interpreted as wide-area spatial asymptotics rather than the long-delay interpretation favoured in communications theory. This motivates us to propose a vector version of the counterexample in Section 2 that is illustrated in Figure 1(a). Section 3 interprets the vector problem as an information theory problem, shown in Figure 1(b) where the distributed controllers $\underline{\underline{C_1}}$ and $\underline{\underline{C_2}}$ are interpreted as an encoder and a decoder. Section 3.1 provides a toy wireless communication problem, called Assisted Interference Suppression, that can be modelled by the vector Witsenhausen counterexample. Section 3.2 shows that the problem is one of many closely related information theory problems, some of which have been recently addressed in the literature (Costa, 1983; Merhav and Shamai, 2007; Kim et al., 2008).

From a control perspective, our claim that the vector formulation simplifies the problem is counterintuitive because it might seem that the simplest version of the problem must contain the fewest number of variables.[2] In Section 4 we show that the asymptotic vector extension is indeed easier to understand since it permits us to characterise the optimal costs within a constant factor. This constant factor is *uniformly bounded* over all values of the problem parameters k and σ_0^2. This characterisation is performed in three steps. First, a new information-theoretic lower bound to the cost is derived. Random-coding based upper bounds are then provided on the cost by analysing nonlinear control strategies inspired by the information-theoretic techniques of lossy source coding, channel coding and DPC (Costa, 1983). Finally, the ratio of the upper bound to the lower bound is proved to be no greater than 11. Numerical evaluation shows that this factor is actually two in the worst case, and is close to one for most values of the problem parameters. The characterisation is in the spirit of some recent advances in information theory where approximate characterisations of rate-regions within a uniform bound have been obtained for some long standing multiuser problems (Avestimehr et al., 2007; Etkin et al., 2007; Avestimehr, 2008).

Section 4.1 also shows that even the scalar restriction of the new lower bound is an improvement over Witsenhausen's original lower bound for certain parameter values. Proper selection of parameters shows that the ratio of our lower bound to Witsenhausen's bound can be arbitrarily large. Because of this looseness in Witsenhausen's bound, there could not have been analogous results characterising the costs for the scalar problem within a constant factor. Whether such a result is possible for the original scalar problem with the new lower bound is an open question, but Section 5 explains why the assumption of asymptotically long vector lengths is crucial to the particular results derived here.

2 The vector version of Witsenhausen's counterexample

The scalar Witsenhausen problem is generalised to the vector case, and the resulting block-diagram is shown in Figure 1(a). The system is still a two-stage control system. The only change from Witsenhausen's original problem is that the states and the inputs are now vectors of length m. A vector is represented in bold font (e.g., $\mathbf{x}$). The space of possible control strategies, including randomised strategies, is denoted by $\mathcal{S}$, which is the space of all Borel measurable functions on the respective observations of the two controllers. For expository convenience, we also allow for strategies that rely on common randomness shared by the two controllers.[3] As in conventional notation, x is used to denote the state, u the input, and y the observation, and $\hat{v}$ is used to denote the estimate of any random variable v. The problem description is as follows:

- The initial state $\mathbf{x}_0$ is Gaussian, distributed $\mathcal{N}(0, \sigma_0^2 \mathbb{I})$.
- The state transition functions describe the state evolution with time. The state transitions are linear:

$$\mathbf{x}_1 = f_1(\mathbf{x}_0, \mathbf{u}_1) = \mathbf{x}_0 + \mathbf{u}_1, \quad \text{and}$$
$$\mathbf{x}_2 = f_2(\mathbf{x}_1, \mathbf{u}_2) = \mathbf{x}_1 - \mathbf{u}_2.$$

- The outputs observed by the controllers:

$$\begin{aligned} \mathbf{y}_1 &= g_1(\mathbf{x}_0) = \mathbf{x}_0, \quad \text{and} \\ \mathbf{y}_2 &= g_2(\mathbf{x}_1) = \mathbf{x}_1 + \mathbf{w}, \end{aligned} \tag{1}$$

  where $\mathbf{w} \sim \mathcal{N}(0, \sigma_w^2 \mathbb{I})$ is Gaussian distributed observation noise.
  As in Witsenhausen (1968), without loss of generality, we assume $\sigma_w^2 = 1$.

Figure 1 In (a), the m-length vector version of the Witsenhausen counterexample is posed as a control problem. The total cost is given by $\frac{k^2}{m}\|\mathbf{u}_1\|^2 + \frac{1}{m}\|\mathbf{x}_2\|^2$, and the objective is to minimise the expected cost. In (b), the same problem is cast as an information theory problem, with $\underline{\underline{C_1}}$ interpreted as the encoder E with input power constraint P, and $\underline{\underline{C_2}}$ interpreted as the decoder D. The objective is to minimise the average distortion $\mathbb{E}[\|\mathbf{x}_1 - \hat{\mathbf{x}}_1\|^2]$

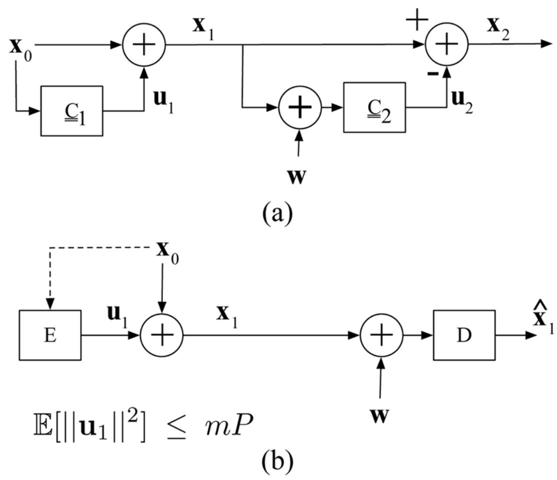

- The control objective is to minimise the expected cost, averaged over the random realisations of $\mathbf{x}_0$ and $\mathbf{w}$, and potentially any randomisation in the control strategy itself. The cost is a quadratic function of the state and the input, given by:

$$\mathcal{C}_1(k^2, \mathbf{x}_1, \mathbf{u}_1) = \frac{1}{m} k^2 \|\mathbf{u}_1\|^2, \quad \text{and}$$
$$\mathcal{C}_2(k^2, \mathbf{x}_2, \mathbf{u}_2) = \frac{1}{m} \|\mathbf{x}_2\|^2$$

where $\| \cdot \|$ denotes the usual Euclidean 2-norm. The cost expressions are normalised by the vector-length m so that they do not necessarily grow with the problem size. The total cost, $\mathcal{C}(\mathbf{x}_1, \mathbf{u}_1, \mathbf{x}_2, \mathbf{u}_2)$ is given by

$$\mathcal{C}(k^2, \mathbf{x}_1, \mathbf{u}_1, \mathbf{x}_2, \mathbf{u}_2) = \mathcal{C}_1(k^2, \mathbf{x}_1, \mathbf{u}_1) + \mathcal{C}_2(k^2, \mathbf{x}_2, \mathbf{u}_2). \tag{2}$$

Given a control strategy $S \in \mathcal{S}$, the expected costs are denoted by $\overline{\mathcal{C}}(k^2, S)$ and $\overline{\mathcal{C}}_i(k^2, S)$ for $i = 1, 2$. We define $\overline{\mathcal{C}}_{\min}(k^2)$ as follows

$$\overline{\mathcal{C}}_{\min}(k^2) := \inf_{S \in \mathcal{S}} \overline{\mathcal{C}}(k^2, S). \tag{3}$$

For notational simplicity, we drop the arguments of the various cost functions when there is no ambiguity in doing so.

- The *information pattern* represents the information available to each controller at each time it takes an action. Following Witsenhausen's notation in Witsenhausen (1971), the information pattern for the vector problem is

$$Y_1 = \{\mathbf{y}_1\}; \quad U_1 = \varnothing; \quad Q_1 = \{q\},$$
$$Y_2 = \{\mathbf{y}_2\}; \quad U_2 = \varnothing; \quad Q_2 = \{q\}.$$

Here Y_i denotes the information about the outputs in equation (1) available at the controller $i \in \{1, 2\}$. Similarly, U_i denotes the information about the previously applied inputs available at the ith controller. $Q_i = q$ denotes the randomness available to the two controllers. Here, q is independent of both $\mathbf{x}_0$ and $\mathbf{w}$. Because it is available to both controllers, the allowed strategies have common randomness.

Note that the second controller does not have knowledge of the output observed or the input applied at the first stage. This makes the information pattern non-classical (or non-nested), and the problem distributed.

We note that for the scalar case of $m = 1$, the problem above reduces to Witsenhausen's original counterexample (Witsenhausen, 1968). Furthermore, because of the diagonal dynamics and diagonal covariance matrices, the optimal linear strategies act on a component-by-component basis. So, even if $m > 1$, the relevant linear strategies are still essentially scalar in nature.

3 Connections with information theory

Figure 1(a) is the vector version for Witsenhausen's counterexample drawn in traditional form with the state evolution forming the backbone of the figure. This is transformed in Figure 1(b) by redrawing the blocks so that the implicit communication channel is conspicuous. The first controller is interpreted as an 'encoder' that modifies the state to enable it to be better communicated to the second controller. The encoder knows the 'interference' $\mathbf{x}_0$. Consider the second cost term

$$\mathcal{C}_2 = \frac{1}{m}\|\mathbf{x}_2\|^2 = \frac{1}{m}\|\mathbf{x}_1 - \mathbf{u}_2\|^2. \tag{4}$$

In order to minimise the expected cost $\overline{\mathcal{C}}_2$, it is optimal to choose $\mathbf{u}_2 = \hat{\mathbf{x}}_1$, the MMSE estimate of $\mathbf{x}_1$, for the second input. The second controller can therefore be interpreted as a 'decoder' that estimates $\mathbf{x}_1$. The cost $\mathcal{C}_2$ is the mean-square error in estimating $\mathbf{x}_1$. We now impose a mean-square constraint on the input $\mathbf{u}_1$,

$$\frac{1}{m}\mathbb{E}[\|\mathbf{u}_1\|^2] \leq P. \tag{5}$$

This is the average power with which the first controller can modify the state. The greater the permitted power P, the smaller we can make the MMSE error $\overline{\mathcal{C}}_2$. We define

$$\overline{\mathcal{C}}_{2,\min}(P) := \inf_{S\in\mathcal{S}:\mathbb{E}[\|\mathbf{u}_1(S)\|^2]\leq mP} \mathbb{E}[\mathcal{C}_2(S)], \tag{6}$$

The following lemma shows that finding $\overline{\mathcal{C}}_{2,\min}(P)$ for all P is equivalent to finding the optimal cost $\overline{\mathcal{C}}_{\min}(k^2)$ for all k.

Lemma 1: *The total minimum cost, $\overline{\mathcal{C}}_{\min}(k^2)$, can be obtained from the optimal tradeoff between P and $\overline{\mathcal{C}}_2$, given by $\overline{\mathcal{C}}_{2,\min}(P)$ for all P. Conversely, given $\overline{\mathcal{C}}_{\min}(k^2)$ for all k, $\overline{\mathcal{C}}_{2,\min}(P)$ can be obtained.*

Proof: The geometric intuition is illustrated in Figure 2. The proof is in Appendix A. □

Figure 2 Shaded regions are the achievable $(\overline{\mathcal{C}}_2, P)$ pairs and the achievable total costs, respectively. The two problems of finding the optimal P and $\overline{\mathcal{C}}_2$ tradeoff and finding $\overline{\mathcal{C}}_{\min}(k^2)$ for all k are equivalent. Given the tradeoff curve (a), to find a point on the minimum cost curve (b) for given k, draw a tangent to the curve in (a) of slope $-\frac{1}{k^2}$. The intercept on the $\overline{\mathcal{C}}_2$ axis gives the minimum total cost. Conversely, given the optimal expected cost $\overline{\mathcal{C}}_{\min}(k^2)$ for all choices of k, draw the line segments given by equation $k^2P + \overline{\mathcal{C}}_2 = \overline{\mathcal{C}}_{\min}(k^2)$ as shown in (a). The supremum over k of all these segments gives the $P - \overline{\mathcal{C}}_2$ tradeoff curve

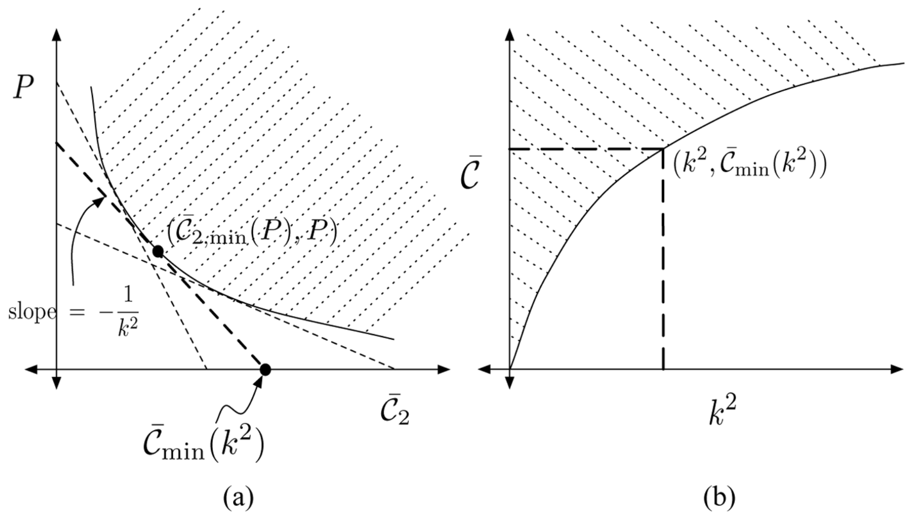

Thus there are two equivalent formulations of the Witsenhausen problem. The P and MMSE tradeoff is reminiscent of tradeoffs in information theory in that an increase in the permitted power P reduces the distortion in estimating $\mathbf{x}_1$.

3.1 Assisted Interference Suppression (AIS)

The vector Witsenhausen problem can further be interpreted as a toy wireless communication problem. Figure 3 illustrates this interpretation, which we refer to as "Assisted Interference Suppression" (AIS). In AIS, the observation noise $\mathbf{w}$ is instead interpreted as a vector of Gaussian symbols that transmitter Tx sends to the receiver Rx in presence of interference $\mathbf{x}_0$ (from the interferer I). The interference vector $\mathbf{x}_0$ is known *non-causally* (prior to transmission) at the 'helper' H (similar to the formulation in, for example (Devroye et al., 2006; Jovicic and Viswanath, 2006)). The helper attempts to wirelessly suppress the effect of the interference at the receiver Rx. The objective is to minimise the mean-square error in estimating $\mathbf{w}$ at Rx under an average power constraint P on the $\mathbf{u}_1$ sent by the helper.

Figure 3 Assisted Interference Suppression problem, a model for which is given in Figure 4. The transmitter Tx sends uncoded Gaussian symbols $\mathbf{w}$ to receiver Rx in the presence of interference $\mathbf{x}_0$ from the interferer I. The helper H has non-causal knowledge of $\mathbf{x}_0$. It attempts to suppress this interference at receiver Rx (by modifying it using $\mathbf{u}_1$) so that the receiver Rx can better estimate the symbols sent by the transmitter Tx

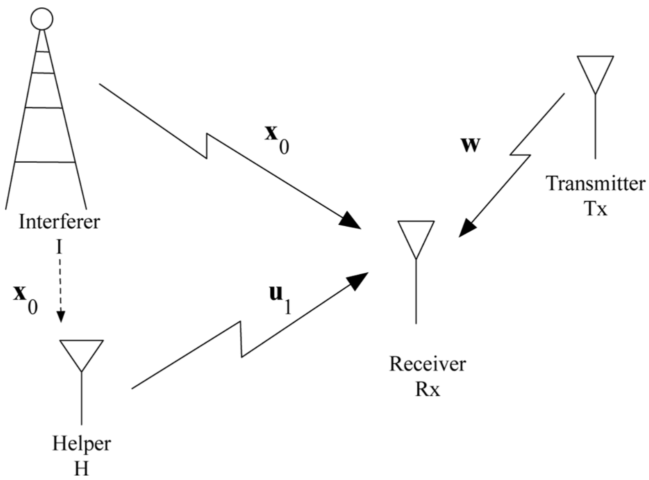

The signal received at Rx is given by

$$\begin{aligned}\mathbf{y} &= \mathbf{x}_0 + \mathbf{u}_1 + \mathbf{w} \\ &= \mathbf{x}_1 + \mathbf{w}.\end{aligned} \tag{7}$$

Let $\hat{\mathbf{w}}$ denote the MMSE estimate of $\mathbf{w}$. Then,

$$\begin{aligned}\hat{\mathbf{w}} &= \mathbb{E}[\mathbf{w} \mid \mathbf{y}] \\ \text{and} \quad \hat{\mathbf{x}}_1 &= \mathbb{E}[\mathbf{x}_1 \mid \mathbf{y}] \\ \text{Therefore,} \quad \hat{\mathbf{x}}_1 + \hat{\mathbf{w}} &= \mathbb{E}[\mathbf{y} \mid \mathbf{y}] = \mathbf{y}.\end{aligned}$$

We conclude that

$$\mathbb{E}[\|\hat{\mathbf{x}}_1 - \mathbf{x}_1\|^2] = \mathbb{E}[\|\hat{\mathbf{w}} - \mathbf{w}\|^2], \tag{8}$$

and thus, minimising the LHS under a power constraint P on $\mathbf{u}_1$, which is the goal of the vector Witsenhausen problem, is equivalent to minimising the RHS under the same constraint, which is AIS.

3.2 Related problems in information theory

Table 1 shows four related problems discussed in the information-theoretic literature and compares their features with AIS. All these problems are special cases of the same block-diagram shown in Figure 4. Because these problems are closely related, one expects that AIS (and hence the vector Witsenhausen problem) may also be amenable to an information-theoretic analysis.

Table 1 A comparison of various information-theoretic problems

Problem	*Reconstruct/hide*	*Performance measures*	*Power constraint(s) on*	*H has M?*	*Solved?*
AIS	Uncoded signal $\hat{\mathbf{w}}$	$\mathbb{E}[\|\mathbf{w} - \hat{\mathbf{w}}\|^2]$	$\mathbf{u}_1$	No	No
DPC (Costa, 1983)	Message $\widehat{M}$	R, P_e	$(\mathbf{u}_1 + \mathbf{w})$	Yes	Yes
Distrib. DPC (Kotagiri and Laneman, 2008)	Message $\widehat{M}$	R, P_e	$\mathbf{u}_1;\mathbf{w}$	No	No
State amplif. (Kim et al., 2008)	Message $\widehat{M}$, 'state' $\hat{\mathbf{x}}_0$	R, P_e, $\mathbb{E}[\|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|^2]$	$(\mathbf{u}_1 + \mathbf{w})$	Yes	Yes
State masking (Merhav and Shamai, 2007)	Message $\widehat{M}$, hide $\mathbf{x}_0$	R, P_e, min $I(\mathbf{x}_0, \mathbf{y}_2)$	$(\mathbf{u}_1 + \mathbf{w})$	Yes	Yes

Costa's original DPC problem (Costa, 1983) addresses the problem of communicating a message M (lying in a set $\{1, 2, \ldots, 2^{mR}\}$ where R is the rate of communication) reliably (with small error probability P_e) across a channel with additive interference and additive white Gaussian noise. The interference vector $\mathbf{x}_0$ is assumed to be known noncausally at the transmitter, that is, the switch in Figure 4 is closed and H and Tx can fully cooperate. The Tx-H pair communicates a message M to the receiver by modifying the interference $\mathbf{x}_0$ using the average power constrained input $\mathbf{u}_1 + \mathbf{w}$. Surprisingly, it turns out that asymptotically, the maximum rate R is the same as that for a channel with no interference.

Recently, significant advances have been made in understanding long-standing open problems in information theory by simplifying them using *deterministic models* (Avestimehr et al., 2007). These models can sharpen the intuition for designing strategies for multi-user communication problems. For example, an approximate rate region has been found to within one bit for the two-user *interference channel* (Etkin et al., 2007), where two transmitter-receiver pairs transmit simultaneously in the same band, causing interference at each other's receiver. Similarly, achievable rates for the relay channel (Avestimehr, 2008) have been characterised to within a constant number of bits,[4] where this constant is independent of the noise variances and the transmit power constraints.

Figure 4 A generic information-theoretic block-diagram that can be used to represent each of the four problems shown in Table 1. For $\mathbf{z} = 0$, no message M, and $\mathbf{w}$ Gaussian, the block-diagram represents the vector Witsenhausen problem (or the AIS problem) with the reconstruction $r = \hat{\mathbf{x}}_1$ (or equivalently, $r = \hat{\mathbf{w}}$). When $\mathbf{w}$ is a coded signal, $r = \widehat{M}$, and the objective is to maximise the message rate while keeping the probability of error $P_e = P(M \neq \widehat{M})$ small, the block-diagram represents Costa's Dirty-Paper Coding (DPC) problem (Costa, 1983) if the switch is closed, and the power constraint is on the power of the sum $\mathbf{w} + \mathbf{u}_1$. It represents distributed DPC (Kotagiri and Laneman, 2008) under the same conditions if the switch is open and constraints are on the individual powers of $\mathbf{w}$ and $\mathbf{u}_1$. With the switch closed and a sum-power constraint, if the objective is to communicate M and *estimate* state $\mathbf{x}_0$, it is the state amplification problem (Kim et al., 2008). If the objective were instead to communicate M and *hide* state $\mathbf{x}_0$, it is the state masking problem

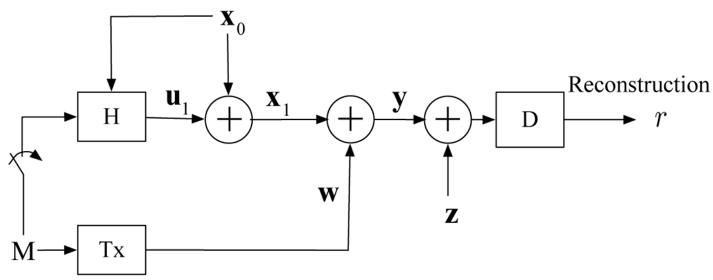

Source: Merhav and Shamai (2007)

The DPC result can be explained intuitively using the deterministic model shown in Figure 5. The transmitter Tx first performs a partial cleaning on the channel by zero-forcing the 'low order bits' of the interference vector, which corresponds to using a power much smaller than the potentially high interference power. The desired message is now encoded into these low order bits. Interestingly, this deterministic interpretation suggests that a DPC scheme would work even if the helper that cleans the channel and the transmitter that transmits the message are different – a situation that can be thought of as distributed DPC. If the transmitter and the helper have equal power, the helper can clean some space for use by the transmitter, who can now communicate reliably at capacity for the power constraint in the space cleaned up. The total power required is thus at most twice the power required for communicating with zero interference. This suggests that a distributed DPC implementation suffers a capacity-loss of at most half a bit.

The distributed DPC problem was addressed by Kotagiri and Laneman (2008) as a special case of the multiple-access channel with partial state information at some encoders. The authors provide upper bounds and lower bounds on achievable rates for the second transmitter, given constraints on the average powers of P and σ_w^2 on the two transmitters. There is, however, a subtle distinction between this problem and AIS. In distributed DPC, the second user is interested in maximising its rate, and can code its message on the channel. In AIS, the objective is to minimise the distortion, and the transmissions $\mathbf{w}$ are uncoded Gaussian symbols.

Figure 5 Deterministic model for Dirty-Paper Coding in which real-valued addition is simplified to bitwise modulo-two addition by essentially dropping any 'carry bits'. The bits I_i represent the interference signal. The transmitter zero-forces the low order bits of the interference using this knowledge, and then encodes the information bits B_i to take their place. This suggests that Dirty-Paper Coding can be implemented in a distributed manner, since in this model the zero-forcing of the interference bits does not require the knowledge of the information bits, and encoding the desired information bits does not require knowledge of the interference

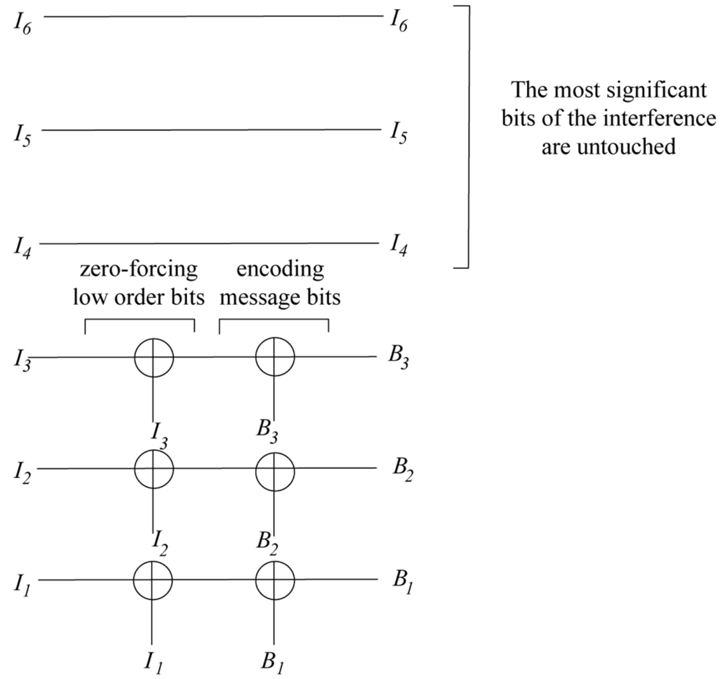

Yet another related pair of problems are *state amplification* (recently solved by Kim et al. (2008)) and *state masking* (recently solved by Merhav and Shamai (2007)). For the system in Figure 4, the objective in Kim et al. (2008) is to reconstruct the message *and* the original interference $\mathbf{x}_0$ at the receiver. The authors characterise the tradeoff between the rate R and the mean-square error in estimating $\mathbf{x}_0$ under an average power constraint. The optimal strategy splits its power – a part of it is used to amplify the state, and the rest of it is used to communicate the message by DPCing against the amplified state. The state-masking problem of Merhav and Shamai (2007) is the opposite – the objective is to *minimise* the information about $\mathbf{x}_0$ that can be obtained from $\mathbf{y}_2$. The optimal strategy here turns out to attenuate[5] the state by using a part of the power, while using the rest of the power to communicate by DPCing against the decreased interference. The crucial distinction between these problems and the vector Witsenhausen problem is that for state-amplification and state-masking the objective is to reconstruct/hide $\mathbf{x}_0$. In contrast, for the vector Witsenhausen problem, the objective is to reconstruct $\mathbf{x}_1$. The helper H gets to modify what is to be estimated. Interestingly, the fact that the helper does not know the message is what seems to make the problem hard in Table 1.

4 Characterisation of the optimal costs for the vector Witsenhausen problem to within a constant factor

This section characterises the asymptotically optimal costs for the vector Witsenhausen problem (in the limit $m \to \infty$) within a constant factor for all values of problem parameters k and σ_0^2.

Theorem 1: *For the problem as stated in Section 2 with $\sigma_w^2 = 1$, in the limit of $m \to \infty$, the optimal expected cost $\overline{\mathcal{C}}_{\min}(k^2)$ for the vector Witsenhausen problem satisfies*

$$\frac{1}{\gamma_1} \min\left\{k^2, k^2\sigma_0^2, \frac{\sigma_0^2}{\sigma_0^2+1}\right\} \le \overline{\mathcal{C}}_{\min}(k^2) \le \min\left\{k^2, k^2\sigma_0^2, \frac{\sigma_0^2}{\sigma_0^2+1}\right\}. \tag{9}$$

Alternatively, $\overline{\mathcal{C}}_{\min}(k^2)$ satisfies

$$\begin{aligned} &\inf_{P\ge 0} k^2P + \left(\left(\sqrt{\kappa(P)} - \sqrt{P}\right)^+\right)^2 \le \overline{\mathcal{C}}_{\min}(k^2) \\ &\le \gamma_2 \inf_{P\ge 0} k^2P + \left(\left(\sqrt{\kappa(P)} - \sqrt{P}\right)^+\right)^2, \end{aligned} \tag{10}$$

where $(\cdot)^+$ is shorthand for $\max(\cdot, 0)$ and

$$\kappa(P) = \frac{\sigma_0^2}{\sigma_0^2 + 2\sigma_0\sqrt{P} + P + 1}. \tag{11}$$

The factors γ_1 and γ_2 are no more than 11 (Numerical evaluation shows that $\gamma_1 < 4.45$, and $\gamma_2 < 2$).

Proof: The proof proceeds in three steps. Section 4.1 states a new information-theoretic lower bound on the expected cost that is valid for all vector lengths. This provides the expressions in equation (10). An upper bound is then derived in Section 4.2 by providing three schemes, and taking the best performance among the three. This provides the expressions in equation (9). The first scheme (providing the k^2 in equation (9)) is a randomised nonlinear controller that we call the Joint Source-Channel Coding (JSCC) scheme. A linear scheme that zero-forces $\mathbf{x}_0$ by using $\mathbf{u}_1 = -\mathbf{x}_0$ achieves the second term $k^2\sigma_0^2$. The third term of $\frac{\sigma_0^2}{\sigma_0^2+1}$ is achieved by another trivial linear scheme using $\mathbf{u}_1 = 0$ and performing an MMSE estimation for $\mathbf{x}_1$ on observing $\mathbf{y}_2$. Figure 6 partitions the (k^2, σ_0^2) parameter space into three different regions, showing which of the three upper bounds is the tightest for various values of k^2 and σ_0^2. It is interesting to note that the nonlinear JSCC scheme is required only in the small-k large-σ_0^2 regime. A similar figure in Baglietto et al. (1997, Figure 1) for the scalar problem shows that the same regime is interesting there as well.

A 3-D plot of the ratio between the upper and lower bounds for varying k^2 and σ_0^2 is shown in Figure 7. The figure shows that the ratio is bounded by a constant γ_1, numerically evaluated to be 4.45, and attained at $k^2 = 0.5$ and $\sigma_0^2 = 1$. The figure also shows that for most of the (k^2, σ_0^2) parameter space, the ratio is in fact close to 1 so the upper and lower bounds are almost equal there. This asymptotic characterisation can be further tightened by improving the upper bound using a balanced combination of DPC and linear control described in Section 4.3 and detailed in Appendix D.8. Numerical evaluation of this ratio leads us to conclude that $\gamma_2 < 2$, as is illustrated in Figure 8. The worst ratio of 2 is achieved along $\sigma_0^2 = \frac{\sqrt{5}-1}{2}$, the golden ratio, and k small.

Figure 6 The plot maps the regions where each of the three schemes (JSCC, zero-forcing $\mathbf{x}_0$, and zero input) perform better than the other two. For large k, zero input performs best. For small k and small σ_0^2, the cost of zero-forcing the state is small, and hence the zero-forcing scheme performs better than the other two. For small k but large σ_0^2, the nonlinear JSCC cost is the smallest amongst the three

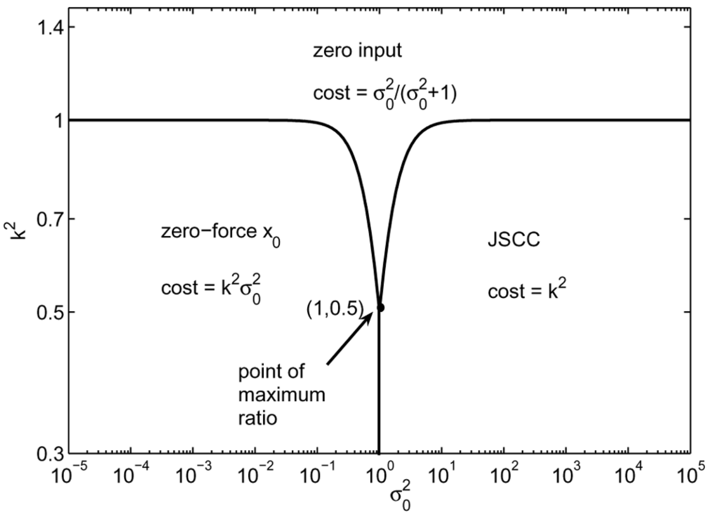

Finally, Appendix E complements the plots by giving an explicit proof that the ratio of the upper and lower bounds is always smaller than 11. □

4.1 A lower bound on the expected cost

Witsenhausen (1968, Section 6) derived the following lower bound on the optimal costs for the scalar problem.

Theorem 2 (Witsenhausen's lower bound)**:** *The optimal cost for the scalar Witsenhausen counterexample is lower bounded by*

$$\overline{\mathcal{C}}_{\min}^{\text{scalar}}(k^2) \geq \frac{1}{\sigma_0} \int_{-\infty}^{+\infty} \phi\left(\frac{\xi}{\sigma_0}\right) V_k(\xi) \mathrm{d}\xi, \tag{12}$$

where $\phi(t) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{t^2}{2})$ *is the standard Gaussian density and*

$$V_k(\xi) := \min_a [k^2(a-\xi)^2 + h(a)], \tag{13}$$

Figure 7 Ratio of the upper bound in equation (9) to the lower bound in equation (10) for varying σ_0 and k. The ratio is upper bounded by 4.45. This shows that the proposed schemes achieve performance within a constant factor of optimal for the vector Witsenhausen problem in the limit of long vector lengths. Notice the ridges along the parameter values where we switch from one control strategy to another in Figure 6 (see online version for colours)

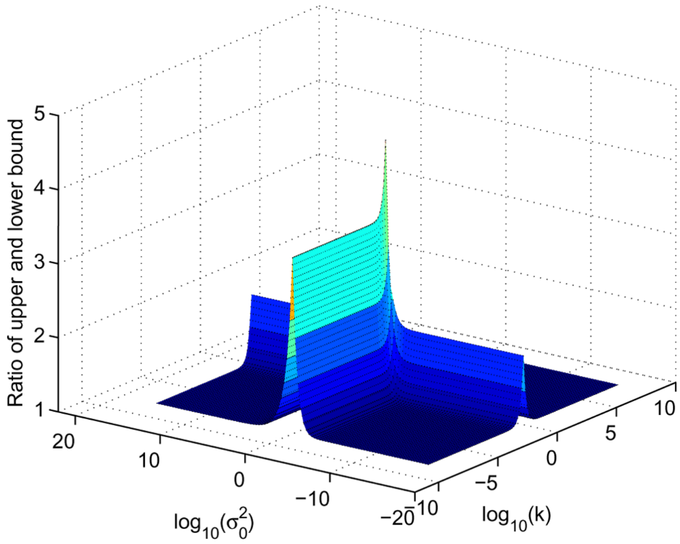

Figure 8 Ratio of the performance of the combined DPC/linear scheme of Section 4.3 (analysed in Appendix D.8) to the lower bound of equation (10) as σ_0 and k vary. Relative to Figure 7, this new scheme has a maximum ratio of two attained on the ridge of $\sigma_0^2 = \frac{\sqrt{5}-1}{2}$ and small k. Also, the ridge along $k = 1$ is reduced as compared to Figure 7. It is eliminated for small σ_0^2, while its asymptotic peak value of about 1.29 is attained at $k \approx 1.68$ and large σ_0^2 (see online version for colours)

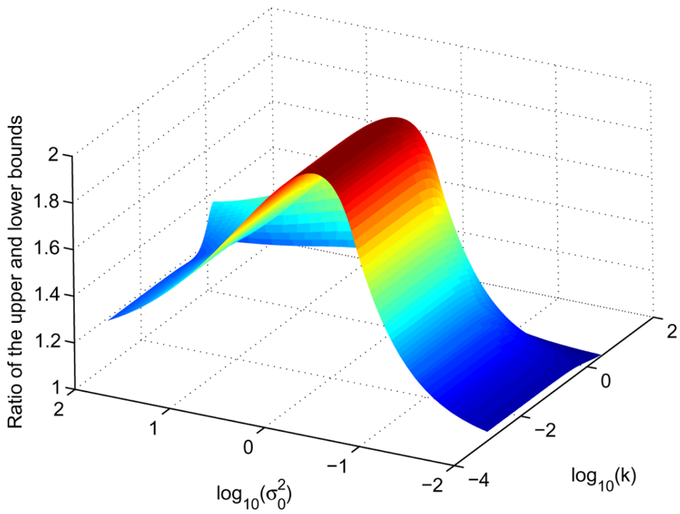

where

$$h(a) := \sqrt{2\pi}a^2\phi(a)\int_{-\infty}^{+\infty}\frac{\phi(y)}{\cosh(ay)}\mathrm{d}y. \tag{14}$$

However, Witsenhausen's scalar-specific proof of this lower bound does not generalise to the vector case. The following theorem provides a new (and simpler to work with) lower bound that is valid for all vector lengths.

Theorem 3 (Lower bound to the vector problem)**:** *For all $m \geq 1$, and all strategies S, given an average power P of $\mathbf{u}_1$, the second stage cost, $\overline{\mathcal{C}}_2(S)$ is lower bounded by*

$$\overline{\mathcal{C}}_2(S) \geq \overline{\mathcal{C}}_{2,\min}(P) \geq \left(\left(\sqrt{\kappa(P)} - \sqrt{P}\right)^+\right)^2, \tag{15}$$

where $\kappa(P)$ is the function of P given by equation (11).

Equivalently, the optimal total cost is lower bounded by

$$\overline{\mathcal{C}}_{\min}(k^2) \geq \inf_{P\geq 0} k^2P + \left(\left(\sqrt{\kappa(P)} - \sqrt{P}\right)^+\right)^2. \tag{16}$$

Proof: See Appendix B. □

Figure 9 plots Witsenhausen's lower bound from Witsenhausen (1968) and compares it with the lower bound of Theorem 3. A particular sequence of $k = \frac{100}{n^2}$ and $\sigma_0^2 = 0.01n^2$ is chosen to visually demonstrate that for this sequence of problem parameters, in the limit of $n \to \infty$, the ratio of the bounds diverges to infinity. Thus, we conclude that prior to this work, it was not possible to provide a uniform (over problem parameters) characterisation of the optimal cost to within a constant factor for the scalar problem. Whether such a characterisation is now possible with this new lower bound is an open question, although we indicate in Section 5 that a further tightening of the bound in Theorem 3 is probably needed in order to obtain such a characterisation.

4.2 An upper bound on the asymptotic expected cost

In Theorem 1, the upper bound is a minimum of three terms. This section describes a nonlinear strategy that asymptotically (in the vector length) attains the cost of k^2 given by the first term of equation (9). We call the strategy the JSCC scheme. The proof uses a randomised code that exploits common randomness.

This is a quantisation-based control strategy and is illustrated in Figure 10, where '$+$' denotes the JSCC quantisation points. The quantisation points are generated randomly according to the distribution $\mathcal{N}(0, (\sigma_0^2 - P)\mathbb{I})$. This set of quantisation points is referred to as the *codebook*. Given a particular realisation of the initial state $\mathbf{x}_0$, the first controller essentially[6] finds the point $\mathbf{x}_1$ in the codebook closest to $\mathbf{x}_0$. The input $\mathbf{u}_1 = \mathbf{x}_1 - \mathbf{x}_0$ then drives the state to this point. The number of quantisation points is chosen carefully – there are sufficiently many of them to ensure that the required average power of $\mathbf{u}_1$ is close to P, but not so many that there could be confusion at the second controller. With this careful choice, we show in Appendix C that on average the state $\mathbf{x}_1$ can be recovered perfectly in the limit $m \to \infty$ as long as the input power $P > \sigma_w^2 = 1$. Thus, asymptotically, $\overline{\mathcal{C}}_2 = 0$, $\overline{\mathcal{C}}_1 = \sigma_w^2 = 1$ and the total cost approaches k^2.

Figure 9 Plot of the two lower bounds on the optimal cost as a function of n, with $k_n = \frac{100}{n^2}$, $\sigma_{0,n} = 0.01n^2$ on a log-log scale for comparing the two lower bounds. The figure shows that the vector lower bound derived here is tighter than Witsenhausen's scalar lower bound in certain cases (see online version for colours)

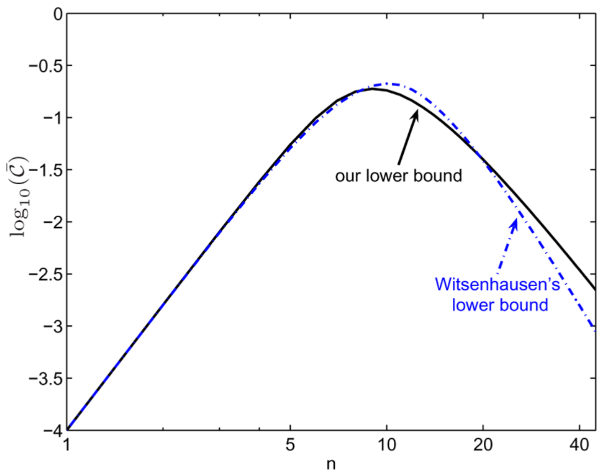

Figure 10 Geometric representation of the Joint Source-Channel Coding scheme of Section 4.2 and the Dirty-Paper Coding scheme of Section 4.3 for the parameter $\alpha = 1$. The grey shell contains the typical $\mathbf{x}_0$ realisations. The JSCC scheme quantises to points inside this shell. The DPC scheme for $\alpha = 1$ quantises the state to points outside this shell (for the same input power), making it robust to larger observation noise variances

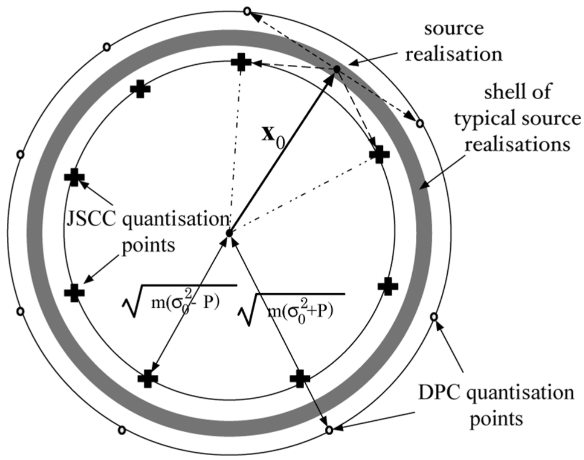

4.3 *Improved upper bound using Dirty-Paper Coding and optimal linear scheme*

In Theorem 1, for analytical simplicity, we used the JSCC scheme and two trivial linear schemes for getting an upper bound on the cost. For numerical evaluation, it is clear that optimal linear schemes can be used instead. This could certainly improve the cost ratio in the regime where the two linear schemes outperform the JSCC scheme in Figure 6. However, the key small-k large-σ_0^2 region is not clear because JSCC outperforms the two linear schemes there.

For a given value of σ_0^2, the cost for the optimal linear scheme is Mitter and Sahai (1999, equation (1))

$$\min_a k^2 a^2 \sigma_0^2 + \frac{(1+a)^2 \sigma_0^2}{1 + (1+a)^2 \sigma_0^2}. \tag{17}$$

The ratio of the optimal linear cost to the k^2 cost for the JSCC scheme is therefore:

$$\inf_a a^2 \sigma_0^2 + \frac{(1+a)^2 \frac{1}{k^2}}{\frac{1}{\sigma_0^2} + (1+a)^2}. \tag{18}$$

Now let $k \to 0$ and $\sigma_0^2 \to \infty$. If a is close to 0, the second term is unbounded. If a is close to -1, the first term is unbounded. For any other value of a, both terms are unbounded. Thus for any sequence of (k, σ_0^2) such that $k \to 0$ and $\sigma_0^2 \to \infty$, the ratio diverges to infinity. Thus an improvement over the JSCC scheme is needed to improve the upper bound in the interesting small-k large-σ_0^2 region.

Such an improvement is obtained by using a nonlinear control strategy based on the concept of DPC (Costa, 1983). DPC techniques (Costa, 1983) can also be thought of as performing a (possibly soft) quantisation. The quantisation points are chosen randomly in the space of realisations of $\mathbf{x}_1$ according to the distribution $\mathcal{N}(0, (P + \alpha^2 \sigma_0^2)\mathbb{I})$. For $\alpha = 1$ the quantisation is hard and a pictorial representation is given in Figure 10, with '$\circ$' denoting the DPC quantisation points. Given the vector $\mathbf{x}_0$, the first controller finds the quantisation point $\mathbf{x}_1$ closest to $\mathbf{x}_0$ and again uses $\mathbf{u}_1 = \mathbf{x}_1 - \mathbf{x}_0$ to drive the state to the closest point. For $\sigma_0^2 > \sigma_w^2 = 1$, we show in Appendix D that asymptotically, $\overline{\mathcal{C}}_2 = 0$, and that this scheme performs better than JSCC.

For $\alpha \neq 1$, the transmitter does not drive the state all the way to a quantisation point. Instead, the state $\mathbf{x}_1 = \mathbf{x}_0 + \mathbf{u}_1$ is merely correlated with the quantisation point, given by $\mathbf{v} = \mathbf{x}_0 + \alpha \mathbf{u}_1$. With high probability, the second controller can decode the underlying quantisation point, and using the two observations $\mathbf{y} = \mathbf{x}_0 + \mathbf{u}_1 + \mathbf{w}$ and $\mathbf{v} = \mathbf{x}_0 + \alpha \mathbf{u}_1$, it can estimate $\mathbf{x}_1 = \mathbf{x}_0 + \mathbf{u}_1$. This scheme has $\overline{\mathcal{C}}_2 \neq 0$, but when k is moderate, the total cost can be lower than that for DPC with $\alpha = 1$. Appendix D describes this strategy and analyses its performance in detail. Interestingly, for $\alpha \neq 1$, the DPC scheme turns out to be a vector extension of the 'neural schemes' (Baglietto et al., 1997) or the soft quantisation schemes developed in Lee et al. (2001). This correspondence is explained through Figure 11. Since the schemes in Baglietto et al. (1997), Lee et al. (2001) were developed using numerical optimisation techniques, Figure 11 helps us understand why they outperform hard quantisation based schemes.

Minor further improvements can be obtained by using a combination scheme that divides its power into two parts: a linear part and a part dedicated to DPC. The linear component is used first to reduce the variance in $\mathbf{x}_0$ by scaling it down in a manner reminiscent of state-masking (Merhav and Shamai, 2007). The remaining power is used to DPCing against the resulting reduced interference. Appendix D.8 provides the details of this combination strategy. As shown in Figure 8, using the combination scheme, the approximation ratio γ_2 is 2.

Figure 11 DPC scheme is a vector extension of the 'soft quantisation' schemes found by numerical optimisation techniques in Baglietto et al. (1997), Lee et al. (2001). These scalar soft quantisation schemes can be interpreted as a sequence of three operations. First, as shown in the upper part of (a), the initial state x_0 is scaled by a constant α. The resulting 'shadow state' αx_0 is then quantised to the nearest quantisation point. The input u_1 that would force αx_0 to the quantisation point is then used as the *actual* input at time 1. The resulting output x_1 will be sloped as a function of x_0, as shown in the lower part of (a), and will take values in intervals on the y-axis. The same sequence of operations – scaling down by α, quantising the resulting shadow state $\alpha \mathbf{x}_0$, and then using the $\mathbf{u}_1$ required for quantising the shadow state as the actual input – yields a continuum of points where $\mathbf{x}_1$ can take values. These values resemble soccer-ball style caps over a sphere, in much the same way as intervals on the real line (see online version for colours)

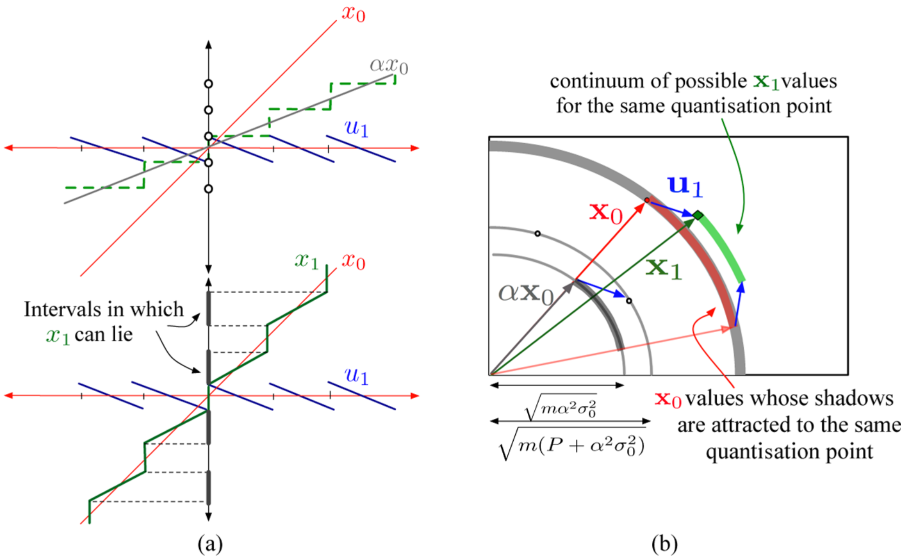

The performance of the various schemes is compared in Figure 12 by finding tradeoffs between the power P and $\overline{C}_2$ for fixed $\sigma_0^2 = 2.5, 0.75$, and 0.25 respectively. The plot on the top-left that corresponds to $\sigma_0^2 = 2.5$ also shows Witsenhausen's lower bound (obtained by appealing to Lemma 1 and drawing tangents to Witsenhausen's lower bound). The plot shows that this lower bound can exceed the vector *upper bound* of the combination ('DPC + linear') scheme. We conclude that Witsenhausen's lower bound is not a valid bound for the vector problem. This also implies that using scalar schemes for the vector problem is strictly suboptimal. Also shown in Figure 12 is the performance attained by the DPC scheme and the balanced combination of DPC and linear control. For $\sigma_0^2 = 0.75$, the combination scheme performs better than both the pure DPC scheme and the optimal linear scheme. Figure 13 illustrates for what values of P the combination scheme is useful over the pure DPC and the pure linear scheme. At low power, DPC performs better than the linear scheme, and all the power is dedicated to the DPC scheme. As the power increases, the fraction η dedicated to the linear scheme increases until the point when all the power is dedicated to the linear scheme.

Figure 12 We normalise $\overline{C}_2$ by the maximum possible $\overline{C}_2$ that is attained when $P = 0$ since in the limit of $P \to 0$, the lower bound and the performance of the linear scheme match. For large σ_0^2 (top figure), the DPC scheme performs better than the linear scheme, and the optimal combination of linear and DPC scheme reduces to the DPC scheme. For small σ_0^2 (bottom figure), the purely linear scheme outperforms DPC. For $\sigma_0^2 = 0.75$, we see that neither scheme is completely dominant for all values of $\overline{C}_2$. The combination scheme performs better than either for some values of $\overline{C}_2$. In the top figure, we also plot Witsenhausen's lower bound on the tradeoff, obtained by drawing tangents to Witsenhausen's lower bound on total cost at 22 different points. Observe that for small $\overline{C}_2$, Witsenhausen's lower bound exceeds the upper bound achieved by the DPC scheme, showing that Witsenhausen's lower bound is not valid for the vector problem (see online version for colours)

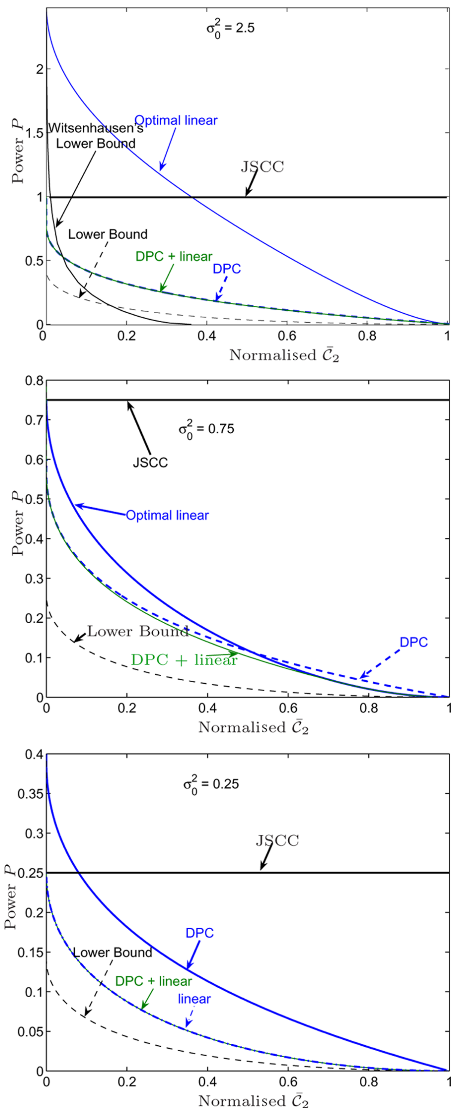

Figure 13 η denotes the fraction of power that is dedicated to the linear scheme. In the optimal combination of DPC/linear control, η is close to 1 for low P, implying that all the power is dedicated to the linear part. As η increases, the power is shared between DPC and linear control, until at high power, when a pure DPC-based approach performs best (see online version for colours)

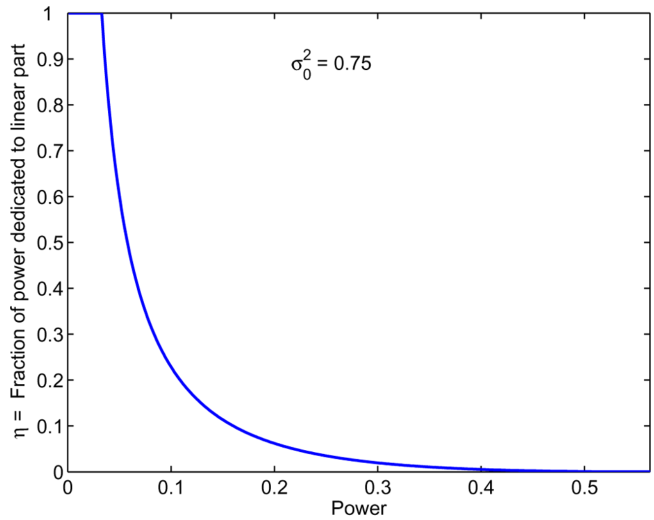

5 Discussions and conclusions

5.1 The case of finite vector lengths

It is important to note that our characterisation of the optimal cost within a uniformly bounded factor is proved only in the limit of infinite vector length. For finite vector lengths, the probability of error in reconstructing $\mathbf{x}_1$ is non-zero, and hence whether the results would continue to apply is an open question. The scalar case, which corresponds to vector length 1, is of particular interest.

It seems to us that even our new lower bound, together with the scalar quantisation-based schemes of Mitter and Sahai (1999), would not suffice to provide a constant-factor guarantee. Consider the interesting regime of $k \to 0$ and $\sigma_0^2 \to \infty$. Figure 6 shows that in this region the vector JSCC scheme outperforms the two linear schemes for the vector problem. In the optimisation of the total cost, a natural heuristic strategy is to equate the costs of the two stages, since the total cost can converge to zero only as fast as the larger of the two terms. We denote the bin size for scalar quantisation by B. In the limit $\sigma_0^2 \to \infty$, the first stage cost scales like to $k^2 \frac{B^2}{12}$, because the distribution of x_0 conditioned on it falling in any particular bin is approximately uniform. The second stage cost is the mean-square error introduced due to incorrect decoding. The error event corresponds to the noise realisation making x_1 look like it is in a different quantisation bin. The probability of this event is approximately $e^{-\frac{B^2}{8}}$. Any incorrect decoding incurs a mean-square error of at least B^2. Equating the two costs, $e^{-\frac{B^2}{8}} B^2 \approx \frac{k^2 B^2}{12}$, suggests that the optimal bin size B is approximately $\sqrt{16 \ln(\frac{2\sqrt{3}}{k})}$. The ratio of the costs for the scalar quantisation scheme and the JSCC scheme is thus approximately $\frac{k^2B^2}{6k^2} = \frac{B^2}{6}$, which diverges since $\ln\left(\frac{1}{k}\right) \to \infty$ as $k \to 0$.

5.2 Conclusions and further work

Using tools from information theory, the asymptotically optimal cost for the vector version of Witsenhausen's counterexample is characterised to within a factor of two for all parameter values. Sections 4.2 and 4.3 provide costs that are attained by randomised JSCC and DPC strategies. However, within the collection of deterministic control strategies over which the randomisation is being performed, there exists a deterministic strategy that attains a cost no larger than the average. So our characterisation holds even when randomised strategies are not allowed.

Since such a constant-factor result is not known for the original scalar problem, we conclude that the counterexample is indeed simplified by considering the vector extension with asymptotically long vector lengths. The results also reaffirm the notion that implicit communication is central to Witsenhausen's counterexample. From an information-theoretic perspective, there are three main remaining issues: closing the gap between the upper and lower bounds, understanding what happens for finite-length vectors, and showing how to exploit known DPC codes to get reasonably good explicit nonlinear control strategies.

The tools we develop in this work might be useful in understanding general distributed control problems. After all, this is a positive result. So it would be interesting to consider a more realistic version of Witsenhausen's counterexample where the first controller receives a noisy observation of the state $\mathbf{x}_0$, and there are also quadratic costs associated with both $\mathbf{x}_1$ and $\mathbf{u}_2$. It might be possible to characterise the asymptotic costs within a constant factor for all these problems as well.

Acknowledgements

We thank Aaron Wagner and Kristen Woyach for interesting discussions, and Gireeja Ranade and the anonymous reviewers for careful reading and suggestions for improvement. The ideas contained herein have also been enhanced by the experience of the second author in co-teaching a special topics course with David Tse. We gratefully acknowledge the support of the National Science Foundation (CNS-403427 and CCF-729122) and Sumitomo Electric.

References

Avestimehr, A. (2008) *Wireless Network Information Flow: A Deterministic Approach*, PhD Thesis, UC Berkeley, Berkeley, CA.

Avestimehr, A., Diggavi, S. and Tse, D. (2007) 'A deterministic approach to wireless relay networks', *Proc. Allerton Conference on Communications, Control and Computing*, Monticello, IL, USA.

Baglietto, M., Parisini, T. and Zoppoli, R. (1997) 'Nonlinear approximations for the solution of team optimal control problems', *Proceedings of the IEEE Conference on Decision and Control CDC*, San Diego, CA, USA, pp.4592–4594.

Bansal, R. and Basar, T. (1987) 'Stochastic teams with nonclassical information revisited: When is an affine control optimal?', *IEEE Transactions on Automatic Control*, Vol. 32, pp.554–559.

Boyd, S. and Vandenberghe, L. (2004) *Convex Optimization*, Cambridge University Press, Cambridge, UK.

Costa, M. (1983) 'Writing on dirty paper', *IEEE Trans. Inform. Theory*, Vol. 29, pp.439–441.

Cover, T.M. and Thomas, J.A. (1991) *Elements of Information Theory*, 1st ed., Wiley, New York.

Dembo, A. and Zeitouni, O. (1998) *Large Deviations Techniques and Applications*, Springer, New York, NY, USA.

Devroye, N., Mitran, P. and Tarokh, V. (2006) 'Achievable rates in cognitive radio channels', *IEEE Trans. Inform. Theory*, Vol. 52, No. 5, pp.1813–1827.

Durrett, R. (2005) *Probability: Theory and Examples*, 1st ed., Brooks/Cole, Belmont, CA.

Etkin, R., Tse, D. and Wang, H. (2007) 'Gaussian interference channel capacity to within one bit', *Submitted to IEEE Transactions on Information Theory*, Vol. 54, No. 12, pp.5534–5562.

Gallager, R.G. (1971) *Information Theory and Reliable Communication*, John Wiley, New York, NY.

Gamal, A.E. and Cover, T. (1980) 'Multiple user information theory', *IEEE Trans. Inform. Theory*, Vol. 68, pp.1466–1483.

Ho, Y.C. and Chang, T. (1980) 'Another look at the nonclassical information structure problem', *IEEE Transactions on Automatic Control*, Vol. 25, No. 3, pp.537–540.

Ho, Y.C., Kastner, M.P. and Wong, E. (1978) 'Teams, signaling, and information theory', *IEEE Trans. Automat. Contr.*, Vol. 23, pp.305–312.

Jovicic, A. and Viswanath, P. (2006) 'Cognitive radio: an information-theoretic perspective', *Proceedings of the 2006 International Symposium on Information Theory*, Seattle, WA, pp.2413–2417.

Kay, S.M. (1998) *Fundamentals of Statistical Signal Processing: Detection Theory*, Vol. 2, Prientice-Hall PTR, New Jersey, USA.

Kim, Y.H., Sutivong, A. and Cover, T.M. (2008) 'State amplification', *IEEE Trans. Inform. Theory*, Vol. 54, pp.1850–1859.

Kotagiri, S. and Laneman, J. (2008) 'Multiaccess channels with state known to some encoders and independent messages', *EURASIP Journal on Wireless Communications and Networking*, Vol. 2008, pp.1–14.

Lee, J.T., Lau, E. and Ho, Y.C.L. (2001) 'The Witsenhausen counterexample: a hierarchical search approach for nonconvex optimization problems', *IEEE Trans. Automat. Contr.*, Vol. 46, No. 3, pp.382–397.

Martins, N.C. (2006) 'Witsenhausen's counter example holds in the presence of side information', *Proceedings of the 45th IEEE Conference on Decision and Control CDC*, San Diego, CA, USA, pp.1111–1116.

Merhav, N. and Shamai, S. (2007) 'Information rates subject to state masking', *IEEE Trans. Inform. Theory*, Vol. 53, pp.2254–2261.

Mitter, S.K. and Sahai, A. (1999) 'Information and control: witsenhausen revisited', in Yamamoto, Y. and Hara, S. (Eds.): *Learning, Control and Hybrid Systems: Lecture Notes in Control and Information Sciences*, Springer, New York, NY, Vol. 241, pp.281–293.

Papadimitriou, C.H. and Tsitsiklis, J.N. (1986) 'Intractable problems in control theory', *SIAM Journal on Control and Optimization*, Vol. 24, pp.639–654.

Rotkowitz, M. (2006) 'Linear controllers are uniformly optimal for the witsenhausen counterexample', *Proceedings of the 45th IEEE Conference on Decision and Control CDC*, San Diego, CA, USA, pp.553–558.

Simon, M. (2002) *Probability Distributions Involving Gaussian Random Variables*, Springer, New York, NY, USA.

Tatikonda, S. (2000) *Control under Communication Constraints*, PhD Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Witsenhausen, H.S. (1968) 'A counterexample in stochastic optimum control', *SIAM Journal on Control*, Vol. 6, pp.131–147.

Witsenhausen, H.S. (1971) 'Separation of estimation and control for discrete time systems', *Proceedings of the IEEE*, Vol. 59, pp.1557–1566.

Notes

[1]Tatikonda used a similar long-vector formulation with parallel dynamics but joint control to show that the sequential-rate distortion function was indeed the right choice to bound performance in the special case when only the state is penalised (Tatikonda, 2000).

[2]Witsenhausen himself states "There does not appear to exist any counterexample involving fewer variables . . . " than the one presented in Witsenhausen (1968).

[3]The common randomness can be viewed as allowing averaging between different strategies at different times. As we use long vector lengths, such performance can in fact be asymptotically attained by cutting the vector into segments and using different strategies on different segments. Thus, all asymptotic results in this paper continue to hold even without the common-randomness assumption.

[4]At large power, the rates scale logarithmically in power. Therefore, in this regime, characterising the rate within a constant number of bits is equivalent to characterising the power required to achieve a specified rate within a constant factor. So the results in this work are similar in spirit to these recent results in information theory.

[5]The deterministic-model perspective on amplification and attenuation is that they both perform a bit-shift operation on the interference.

[6]For the purpose of simplifying the exposition, the proof actually includes a few exponentially-rate cases when this is not done.

Appendix

A Proof of Lemma 1

We note that in the definition of $\overline{\mathcal{C}}_{2,\min}(P)$ in equation (6), the inequality constraint can in in fact be replaced by equality. Suppose a strategy S has average input power $P-\epsilon$ for some $\epsilon > 0$. Using a randomised strategy, the first controller can add a random vector (known at the second controller) of power ϵ independent of the the input $\mathbf{u}_1$ chosen according to strategy S. The second controller can simply subtract this random component and attain the performance of S.

We first show that we can obtain $\overline{\mathcal{C}}_{\min}(k^2)$ given $\overline{\mathcal{C}}_{2,\min}(P)$ for all P.

$$\begin{aligned}\overline{\mathcal{C}}_{\min}(k^2) &= \inf_{S\in\mathcal{S}} \frac{1}{m} k^2 \mathbb{E}[\|\mathbf{u}_1\|^2] + \overline{\mathcal{C}}_2(S) \\ &= \inf_{P\geq 0} \inf_{S\in\mathcal{S}:\frac{1}{m}\mathbb{E}[\|\mathbf{u}_1\|^2]=P} k^2 P + \overline{\mathcal{C}}_2(S)\end{aligned}$$

$$= \inf_{P \geq 0} k^2 P + \inf_{S \in \mathcal{S}: \frac{1}{m}\mathbb{E}[\|\mathbf{u}_1\|^2] = P} \overline{\mathcal{C}}_2(S)$$

$$= \inf_{P \geq 0} k^2 P + \overline{\mathcal{C}}_{2,\min}(P). \tag{19}$$

Thus $\overline{\mathcal{C}}_{\min}(k^2)$ can be obtained from $\overline{\mathcal{C}}_{2,\min}(P)$. Now we show that we can find $\overline{\mathcal{C}}_{2,\min}(P)$ if we know $\overline{\mathcal{C}}_{\min}(k^2)$. We first need the following lemma.

Lemma 2: *$\overline{\mathcal{C}}_{2,\min}(P)$ is convex in P.*

Proof: For any P_1 and P_2, we want to show that $\overline{\mathcal{C}}_{2,\min}(\lambda P_1 + (1-\lambda)P_2)$ is no greater than $\lambda\overline{\mathcal{C}}_{2,\min}(P_1) + (1-\lambda)\overline{\mathcal{C}}_{2,\min}(P_2)$.

Consider strategies S_1 and S_2 that operate at power P_1 and P_2 respectively such that $\overline{\mathcal{C}}_2(S_i) < \overline{\mathcal{C}}_{2,\min}(P_i) + \epsilon$. A randomised strategy S_r that chooses S_1 with probability λ and S_2 with probability $1-\lambda$ achieves cost that is $\lambda\overline{\mathcal{C}}_{2,\min}(P_1) + (1-\lambda)\overline{\mathcal{C}}_{2,\min}(P_2) + \epsilon$. Also, the power required by S_r is clearly the average power $\lambda P_1 + (1-\lambda)P_2$. Thus,

$$\overline{\mathcal{C}}_{2,\min}(\lambda P_1 + (1-\lambda)P_2) \leq \lambda\overline{\mathcal{C}}_{2,\min}(P_1) + (1-\lambda)\overline{\mathcal{C}}_{2,\min}(P_2) + \epsilon.$$

The convexity follows because ϵ can be made as small as desired.

Now, define $g(\cdot)$ as the conjugate function (Boyd and Vandenberghe, 2004, p.91) of $\overline{\mathcal{C}}_{2,\min}(\cdot)$,

$$g(z) := \sup_{z}(zP - \overline{\mathcal{C}}_{2,\min}(P_2)). \tag{20}$$

Since $\overline{\mathcal{C}}_{2,\min}(\cdot)$ is convex, it is the conjugate function of its conjugate function, $g(\cdot)$ (Boyd and Vandenberghe, 2004, p.94). Thus, we can obtain $\overline{\mathcal{C}}_{2,\min}(\cdot)$ from $g(\cdot)$. Now, observe that

$$\overline{\mathcal{C}}_{\min}(k^2) = -g(-k^2). \tag{21}$$

Therefore, we can obtain $\overline{\mathcal{C}}_{2,\min}(\cdot)$ from $\overline{\mathcal{C}}_{\min}(\cdot)$.

B Derivation of the lower bound on the cost for vector Witsenhausen problem

In this section, we derive a lower bound on the cost for the vector Witsenhausen problem. Since the bound is valid for any vector length m, it is also valid for $m = 1$. This bound is needed because the techniques of the lower bound in Witsenhausen (1968), Section 6 do not generalise to $m > 1$.

First, a simple lemma.

Lemma 3: *For any three vector random variables A, B and C,*

$$\sqrt{\mathbb{E}[d(B,C)]} \geq |\sqrt{\mathbb{E}[d(A,C)]} - \sqrt{\mathbb{E}[d(A,B)]}|, \tag{22}$$

where $d(A,B) = \|A - B\|^2$.

Proof: Using the triangle inequality on Euclidian distance, □

$$\sqrt{d(B,C)} \geq \sqrt{d(A,C)} - \sqrt{d(A,B)}. \tag{23}$$

Similarly,

$$\sqrt{d(B,C)} \geq \sqrt{d(A,B)} - \sqrt{d(A,C)}. \tag{24}$$

Thus,

$$\sqrt{d(B,C)} \geq |\sqrt{d(A,C)} - \sqrt{d(A,B)}|, \tag{25}$$

Squaring both sides,

$$d(B,C) \geq d(A,C) + d(A,B) - 2\sqrt{d(A,C)}\sqrt{d(A,B)}. \tag{26}$$

Taking the expectation on both sides,

$$\begin{aligned}\mathbb{E}[d(B,C)] \geq\ & \mathbb{E}[d(A,C)] + \mathbb{E}[d(A,B)] \\ & - 2\mathbb{E}[\sqrt{d(A,C)}\sqrt{d(A,B)}].\end{aligned} \tag{27}$$

Now, using the Cauchy-Schwartz inequality (Durrett, 2005, p.13),

$$\left(\mathbb{E}[\sqrt{d(A,C)}\sqrt{d(A,B)}]\right)^2 \leq \mathbb{E}[d(A,C)]\mathbb{E}[d(A,B)]. \tag{28}$$

Using equations (27) and (28),

$$\begin{aligned}\mathbb{E}[d(B,C)] &\geq \mathbb{E}[d(A,C)] + \mathbb{E}[d(A,B)] - 2\sqrt{\mathbb{E}[d(A,C)]}\sqrt{\mathbb{E}[d(A,B)]} \\ &= \left(\sqrt{\mathbb{E}[d(A,C)]} - \sqrt{\mathbb{E}[d(A,B)]}\right)^2.\end{aligned}$$

Taking square-roots on both the sides completes the proof.

Substituting $\mathbf{x}_0$ for A, $\mathbf{x}_1$ for B, and $\mathbf{u}_2$ for C in Lemma 3, we get

$$\sqrt{\mathbb{E}[d(\mathbf{x}_1, \mathbf{u}_2)]} \geq \sqrt{\mathbb{E}[d(\mathbf{x}_0, \mathbf{u}_2)]} - \sqrt{\mathbb{E}[d(\mathbf{x}_0, \mathbf{x}_1)]}. \tag{29}$$

We wish to lower bound $\mathbb{E}[d(\mathbf{x}_1, \mathbf{u}_2)] = \mathbb{E}[\|\mathbf{x}_1 - \mathbf{u}_2\|^2]$. The second term on the RHS is $\sqrt{mP}$. Therefore, it now suffices to lower bound the first term on the RHS of equation (29). To that end, we will interpret $\mathbf{u}_2$ as an estimate for $\mathbf{x}_0$. Now,

$$\begin{aligned}\mathbf{y}_2 &= \mathbf{x}_1 + \mathbf{w} \\ &= \mathbf{x}_0 + \mathbf{u}_1 + \mathbf{w}.\end{aligned}$$

This is an implicit AWGN channel with 'input' $\mathbf{x}_1$, noise $\mathbf{w}$, and output $\mathbf{y}_2$. The input power, which is the power of $\mathbf{x}_1$, is at most $P_{ch} = P + \sigma_0^2 + 2\sqrt{P\sigma_0^2}$ (attained when $\mathbf{x}_0$ and $\mathbf{u}_1$ are perfectly aligned). The channel capacity can be upper bounded by $\overline{C}$, the maximum mutual information at power P_{ch} (Cover and Thomas, 1991, p.242).

$$\overline{C} = \frac{1}{2}\log_2\left(1 + \frac{P_{ch}}{\sigma_w^2}\right). \tag{30}$$

Now, $\mathbb{E}\left[\|\mathbf{x}_0 - \mathbf{u}_2\|^2\right]$ can be lower bounded by the distortion-rate function $D(R) = \sigma_0^2 2^{-2R}$ (Cover and Thomas, 1991, pp.344–346) (for the Gaussian source that generates $\mathbf{x}_0$) evaluated at rate R equal to $\overline{C}$.

$$\begin{aligned} D(\overline{C}) &= \sigma_0^2 2^{-2\overline{C}} \\ &= \sigma_0^2 \frac{\sigma_w^2}{P_{ch} + \sigma_w^2} \\ &= \frac{\sigma_0^2\sigma_w^2}{\sigma_0^2 + P + 2\sqrt{P}\sigma_0 + \sigma_w^2} = \kappa(P). \end{aligned}$$

Thus, $\mathbb{E}\left[\|\mathbf{x}_0 - \mathbf{u}_2\|^2\right] \geq m\kappa(P)$.

If $\kappa(P) > P$, using $\mathbb{E}\left[\|\mathbf{x}_0 - \hat{\mathbf{x}}_1\|^2\right] \geq m\kappa(P)$ and $\mathbb{E}\left[\|\mathbf{x}_0 - \mathbf{x}_1\|^2\right] \leq mP$ in equation (29),

$$\mathbb{E}\left[\|\mathbf{x}_1 - \hat{\mathbf{x}}_1\|^2\right] \geq m\left(\sqrt{\kappa(P)} - \sqrt{P}\right)^2 \tag{31}$$

where we use $\mathbf{u}_2 = \hat{\mathbf{x}}_1$. If $\kappa(P) \leq P$, we lower bound $\mathbb{E}\left[\|\mathbf{x}_1 - \hat{\mathbf{x}}_1\|^2\right]$ by zero. The lower bound follows.

C The joint source-channel scheme

We now describe in detail the randomised JSCC scheme and characterise its average performance, averaged over the realisations of $(q, \mathbf{x}_0, \mathbf{w})$, where q denotes the common randomness. We assume $\sigma_0^2 > \sigma_w^2$, because for $\sigma_w^2 \geq \sigma_0^2$, we can just always force the state to zero by choosing $\mathbf{u}_1 = -\mathbf{x}_0$ and paying a lower cost of $k^2\sigma_0^2$. Because the scheme borrows from both Gaussian source coding (Gallager, 1971, Ch. 9) and Additive White Gaussian Noise (AWGN) channel coding (Cover and Thomas, 1991, pp.241–245), slight modifications to the textbook proofs are needed to put the two together.

In the following, $\mathbb{S}_1$ denotes the m-dimensional sphere centred at zero with radius $2\sqrt{m\sigma_0^2}$, and $\mathbb{S}_2$ denotes a sphere centred at zero with radius $6\sqrt{m\sigma_0^2}$. Also, $\mathbb{1}_{\{A\}}$ denotes the indicator function of an event A in the relevant space. We first describe the randomised strategy (the encoding and the decoding) and then analyse its performance.

C.1 Codebook construction and encoding

The encoding is performed at the first controller $\underline{\underline{C_1}}$. The strategy has a single parameter δ. A list $\mathbb{Q}$ of $2^{mR}+1$ quantisation points $\{\mathbf{x}_q(0), \mathbf{x}_q(1), \ldots, \mathbf{x}_q(2^{mR})\}$ is chosen by drawing the quantisation points iid in $\mathbb{R}^m$ randomly from the distribution $\mathcal{N}(0, (\sigma_0^2 - P)\mathbb{I})$, where the operating 'rate' R and the power P satisfy the pair of equalities

$$R = \mathcal{R}(P) + \frac{\delta}{2} = \frac{1}{2}\log_2(\sigma_0^2/P) + \frac{\delta}{2} \tag{32}$$

$$C(P) = \frac{1}{2}\log_2\left(1 + \frac{\sigma_0^2 - P}{\sigma_w^2}\right) = R + \frac{\delta}{2}, \tag{33}$$

for small $\delta > 0$ where $\mathcal{R}(\cdot)$ is the rate-distortion function for a Gaussian source of variance σ_0^2 (Cover and Thomas, 1991, p.345), and $C(\cdot)$ is the capacity of an AWGN channel with input power constraint $\sigma_0^2 - P$. That this pair of equalities has a solution with $P < \sigma_0^2$ is shown in Appendix C.4. The probability space therefore consists of three independent random variables, the common randomness $\mathbb{Q}$, the initial state $\mathbf{x}_0$ and the noise $\mathbf{w}$.

The encoding now proceeds in three steps.

Step 1: In the case that $\mathbb{Q}$ is *pathological* so that the first quantisation point $\mathbf{x}_q(0) \notin \mathbb{S}_1$, the encoder just uses $\mathbf{u}_1 = -\mathbf{x}_0$ to push the state to zero.

Step 2: If $\mathbf{x}_0 \notin \mathbb{S}_1$, the encoder uses $\mathbf{u}_1 = -\mathbf{x}_0$ to force $\mathbf{x}_0$ to zero.

Step 3: If $\mathbf{x}_q(0) \in \mathbb{S}_1$ and $\mathbf{x}_0 \in \mathbb{S}_1$, the encoding is performed by finding the quantisation point $\mathbf{x}_1 \in \mathbb{S}_2$ such that

$$\mathbf{x}_1(\mathbf{x}_0) = \underset{\mathbf{x} \in \mathbb{Q} \cap \mathbb{S}_2}{\arg\min} \|\mathbf{x}_0 - \mathbf{x}\|, \tag{34}$$

and using $\mathbf{u}_1 = \mathbf{x}_1 - \mathbf{x}_0$ to drive the state to $\mathbf{x}_1$.

C.2 Decoding

The decoder $\underline{\underline{C_2}}$ is assumed to know the realisation of $\mathbb{Q}$ because of the shared common randomness.

Case 1: If $\mathbf{x}_q(0) \notin \mathbb{S}_1$, the second controller decodes to zero and applies $\mathbf{u}_2 = 0$.

Case 2: If $\mathbf{x}_q(0) \in \mathbb{S}_1$, the second controller examines the noisy observation $\mathbf{y}_2 = \mathbf{x}_1 + \mathbf{w}$. It decodes to the quantisation point $\hat{\mathbf{x}}_1$ given by

$$\hat{\mathbf{x}}_1(\mathbf{y}_2) = \underset{\mathbf{x} \in \mathbb{Q} \cap \mathbb{S}_2}{\arg\min} \|\mathbf{y}_2 - \mathbf{x}\|, \tag{35}$$

and applies the control $\mathbf{u}_2 = \hat{\mathbf{x}}_1$.

C.3 Performance analysis

We now show that the average cost of the encoding and the decoding can be made arbitrarily close to $k^2\sigma_w^2$ and zero respectively by choosing m large enough and δ small enough. We first need the following lemma.

Lemma 4: *Let* $\mathbf{x} \sim \mathcal{N}(0, \sigma^2 \mathbb{I})$ *be an* m*-dimensional random vector and let* A^m *be sets such that* $\lim_{m\to\infty} \Pr(A^m) = 0$*. Then,*

$$\lim_{m\to\infty} \frac{1}{m} \mathbb{E}\left[\|\mathbf{x}\|^2 \mathbb{1}_{\{A^m\}}\right] = 0. \tag{36}$$

Further, if $A_L^m = \{\|\mathbf{x}\|^2 \geq mL\}$ *for some fixed* $L > \sigma_0^2$*, then* $\Pr(A_L^m)$ *and* $\mathbb{E}[\|\mathbf{x}\|^2 \mathbb{I}_{\{A_L^m\}}] \to 0$ *exponentially in* m*.*

Proof: First, observe that $T_m = \frac{\|\mathbf{x}\|^2}{\sigma^2}$ is a Chi-square random variable with m degrees of freedom with mean m and variance $2m$ (Simon, 2002, p.14). Thus,

$$\frac{1}{\sigma^4} \mathbb{E}\left[\|\mathbf{x}\|^4\right] = \mathbb{E}\left[T_m^2\right] = \mathrm{Var}(T_m) + (\mathbb{E}\left[T_m\right])^2 = 2m + m^2. \tag{37}$$

By the Cauchy-Schwartz inequality (Durrett, 2005, p.13),

$$\begin{aligned}\left(\frac{1}{m}\mathbb{E}\left[\|\mathbf{x}\|^2 \mathbb{1}_{\{A^m\}}\right]\right)^2 &\leq \frac{\mathbb{E}\left[\|\mathbf{x}\|^4\right]}{m^2} \mathbb{E}\left[\mathbb{1}^2_{\{A^m\}}\right] \\ &= \frac{1}{m^2}\mathbb{E}\left[\|\mathbf{x}\|^4\right] \Pr(A^m) \\ &= \left(\frac{2}{m} + 1\right)\sigma^4 \Pr(A^m) \end{aligned} \tag{38}$$

which converges to zero as $m \to \infty$ since $\Pr(A^m) \to 0$.

Now for the second part, since the elements of $\mathbf{x}$ are iid Gaussian random variables, by Cramer's theorem (Dembo and Zeitouni, 1998, p.27), the probability of the event $A_L^m = \left\{\frac{1}{m}\sum_{i=1}^m x_i^2 > \sigma^2 + \epsilon\right\}$ converges to zero exponentially in m for any $\epsilon > 0$. From equation (38), $\mathbb{E}\left[\|\mathbf{x}\|^2 \mathbb{1}_{\{A_L^m\}}\right] \to 0$ exponentially in m as well.

Let $\mathcal{E}_1 := \{\mathbf{x}_q(0) \notin \mathbb{S}_1\}$ denote the event of pathological $\mathbb{Q}$, $\mathcal{E}_2 := \{\mathbf{x}_0 \notin \mathbb{S}_1\} \cap \mathcal{E}_1^c$ and $\mathcal{E}_3 := \{\mathbf{x}_0 \in \mathbb{S}_1\} \cap \mathcal{E}_1^c$, where A^c denotes the complement of the event A. Since the choice of $\mathbb{Q}$ is independent of $\mathbf{x}_0$, $\mathcal{E}_1$ is independent of the event $\{\mathbf{x}_0 \in \mathbb{S}_1\}$.

The total expected cost is given by

$$\mathbb{E}\left[\mathcal{C}\right] = \mathbb{E}\left[\mathcal{C}|\mathcal{E}_1\right]\Pr(\mathcal{E}_1) + \mathbb{E}\left[\mathcal{C}|\mathcal{E}_2\right]\Pr(\mathcal{E}_2) + \mathbb{E}\left[\mathcal{C}|\mathcal{E}_3\right]\Pr(\mathcal{E}_3) \tag{39}$$

$$= \mathbb{E}\left[\mathcal{C}_1 + \mathcal{C}_2|\mathcal{E}_1\right]\Pr(\mathcal{E}_1) + \mathbb{E}\left[\mathcal{C}_1 + \mathcal{C}_2|\mathcal{E}_2\right]\Pr(\mathcal{E}_2) + \mathbb{E}\left[\mathcal{C}_1 + \mathcal{C}_2|\mathcal{E}_3\right]\Pr(\mathcal{E}_3). \tag{40}$$

We will now upper bound each of the three terms in equation (39).

Cost for $\mathcal{E}_1$: $\mathbb{E}\left[\mathcal{C}|\mathcal{E}_1\right]\Pr(\mathcal{E}_1)$.

Consider the pathological event $\mathcal{E}_1$. By the independence of $\mathbf{x}_0$ and $\mathcal{E}_1$,

$$\mathbb{E}\left[\mathcal{C}_1|\mathcal{E}_1\right]\Pr(\mathcal{E}_1) = k^2\sigma_0^2 \Pr(\mathcal{E}_1).$$

$\Pr(\mathcal{E}_1)$ converges to zero by Lemma 4. Therefore, for any $\epsilon_1 > 0$ there exists $m_1(\epsilon_1, \delta)$ such that for all $m \geq m_1(\epsilon_1, \delta)$, $\mathbb{E}\left[\mathcal{C}_1 \,|\, \mathcal{E}_1\right] \Pr(\mathcal{E}_1) < \epsilon_1$.

For the decoding cost, since the decoder correctly decodes to zero,

$$\mathbb{E}\left[\mathcal{C}_2 \,|\, \mathcal{E}_1\right] \Pr(\mathcal{E}_1) = 0.$$

Thus the first term in equation (39) is smaller than ϵ_1 for $m \geq m_1(\epsilon_1, \delta)$.

Cost for $\mathcal{E}_2$: $\mathbb{E}\left[\mathcal{C} \,|\, \mathcal{E}_2\right] \Pr(\mathcal{E}_2)$.

$$\mathbb{E}\left[\mathcal{C}_1 \,|\, \mathcal{E}_2\right] \Pr(\mathcal{E}_2) = \frac{k^2}{m} \mathbb{E}\left[\|\mathbf{x}_0\|^2 \mathbb{1}_{\{\mathcal{E}_2\}}\right],$$

which converges to zero as $m \to \infty$ by Lemma 4.

Since the decoder always decodes to a quantisation point inside $\mathbb{S}_2$, which has diameter $6\sqrt{m\sigma_0^2}$,

$$\mathbb{E}\left[\mathcal{C}_2 \,|\, \mathcal{E}_2\right] \Pr(\mathcal{E}_2) \leq 36\sigma_0^2 \Pr(\mathcal{E}_2).$$

By Lemma 4, $\Pr(\mathcal{E}_2)$ also converges to zero.

Thus the second term in equation (39) can be made smaller than ϵ_1 for $m \geq m_2(\epsilon_1, \delta)$.

Cost for $\mathcal{E}_3$: $\mathbb{E}\left[\mathcal{C} \,|\, \mathcal{E}_3\right] \Pr(\mathcal{E}_3)$.

Gallager (1971, pp.471,472) constructs quantisation codebooks for Gaussian vectors by the same random generation as in Appendix C.1, except that the construction in Appendix C.1 has one extra point $\mathbf{x}_q(0)$.

Since $\mathbf{x}_q(0) \in \mathbb{S}_1$, for any initial state realisation $\mathbf{x}_0 \in \mathbb{S}_1$, $\|\mathbf{x}_0 - \mathbf{x}_q(0)\| \leq 4\sqrt{m\sigma_0^2}$, the diameter of $\mathbb{S}_1$. Also, for any $\mathbf{x}_0 \in \mathbb{S}_1$, and any quantisation point $\mathbf{x}_q^* \notin \mathbb{S}_2$, $\|\mathbf{x}_0 - \mathbf{x}_q^*\| \geq 4\sqrt{m\sigma_0^2}$. Thus any $\mathbf{x}_0 \in \mathbb{S}_1$ is encoded to inside $\mathbb{S}_2$. Any $\mathbf{x}_0 \notin \mathbb{S}_1$ is encoded to $0 \in \mathbb{S}_2$, and thus all initial state realisations are encoded to within $\mathbb{S}_2$.

In Gallager (1971, pp.471, 472), it is shown that for a randomly constructed source codebook of 2^{mR} quantisation points, as long as $R = \mathcal{R}(\mathcal{P}) + \delta$ for some $\delta > 0$,

$$\lim_{m \to \infty} \Pr\left(\frac{1}{k^2}\mathcal{C}_1 > P + \epsilon_1\right) \to 0, \tag{41}$$

for all $\epsilon_1 > 0$. Since adding an extra quantisation point can only decrease the distortion,

$$\lim_{m \to \infty} \Pr\left(\frac{1}{k^2}\mathcal{C}_1 > P + \epsilon_1 \,\middle|\, \mathcal{E}_3\right) \to 0. \tag{42}$$

Thus, there exists $m_3(\epsilon_1, \delta)$ such that for all $m > m_3(\epsilon_1, \delta)$, $\Pr\left(\frac{1}{k^2}\mathcal{C}_1 > P + \epsilon_1 | \mathcal{E}_3\right) < \frac{\epsilon_1}{4\sigma_0^2}$. Therefore,

$$\begin{aligned}
\mathbb{E}\left[\mathcal{C}_1|\mathcal{E}_3\right] &= \mathbb{E}\left[\mathcal{C}_1 \,\middle|\, \mathcal{E}_3, \frac{1}{k^2}\mathcal{C}_1 > P + \epsilon_1\right] \Pr\left(\frac{1}{k^2}\mathcal{C}_1 > P + \epsilon_1 \,\middle|\, \mathcal{E}_3\right) \\
&\quad + \mathbb{E}\left[\mathcal{C}_1 \,\middle|\, \mathcal{E}_3, \frac{1}{k^2}\mathcal{C}_1 \le P + \epsilon_1\right] \Pr\left(\frac{1}{k^2}\mathcal{C}_1 \le P + \epsilon_1 \,\middle|\, \mathcal{E}_3\right) \\
&\le 4\sigma_0^2 k^2 \Pr\left(\frac{1}{k^2}\mathcal{C}_1 > P + \epsilon_1 \,\middle|\, \mathcal{E}_3\right) + k^2(P + \epsilon_1) \\
&\le 2k^2\epsilon_1 + k^2 P.
\end{aligned}$$

Thus, $\mathbb{E}\left[\mathcal{C}_1|\mathcal{E}_3\right] \le k^2(P + 2\epsilon_1)$ for all $m > m_3(\epsilon_1, \delta)$.

Now we analyse the cost of decoding under the event $\mathcal{E}_3$. The key observation is that a randomly generated Gaussian codebook also achieves the capacity of a Gaussian channel (Cover and Thomas, 1991, p. 244) for an average power constraint equal to the average power of the codebook. By construction, the average power of the codebook, that is, $\frac{1}{m}\mathbb{E}\left[\|\mathbf{x}_1\|^2\right]$ is $\sigma_0^2 - P$. The channel capacity C for a Gaussian channel with power constraint $\sigma_0^2 - P$ is

$$C = \frac{1}{2}\log_2\left(1 + \frac{\sigma_0^2 - P}{\sigma_w^2}\right). \tag{43}$$

We know that as long as $\log_2\left(2^{mR} + 1\right) < C$, the average error probability $\mathbb{E}\left[P_e(\mathbb{Q})\right]$ for a random Gaussian codebook of $2^{mR} + 1$ codewords converges to zero (Cover and Thomas, 1991, p.244). There exists $m_4(\delta)$ such that for all $m > m_4(\delta)$, $C = R + \frac{\delta}{2} > \log_2\left(2^{mR} + 1\right) + \frac{\delta}{4}$. Now,

$$\begin{aligned}
\mathbb{E}\left[P_e(\mathbb{Q})\right] &\ge \mathbb{E}\left[P_e(\mathbb{Q})|\mathcal{E}_3\right]\Pr(\mathcal{E}_3) &(44)\\
&= \mathbb{E}\left[P_e(\mathbb{Q})|\mathcal{E}_3\right](1 - \Pr(\mathcal{E}_1) - \Pr(\mathcal{E}_2)). &(45)
\end{aligned}$$

Since $\Pr(\mathcal{E}_1)$ and $\Pr(\mathcal{E}_2)$ converge to zero as $m \to \infty$, $1 - \Pr(\mathcal{E}_1) - \Pr(\mathcal{E}_2) \to 1$. Also, $\mathbb{E}\left[P_e(\mathbb{Q})\right] \to 0$. Thus $\mathbb{E}\left[P_e(\mathbb{Q})|\mathcal{E}_3\right] \to 0$ as well.

In case of a decoding error, since $\hat{\mathbf{x}}_1$ and $\mathbf{x}_1$ are both in $\mathbb{S}_2$, they are separated by a distance no more than the diameter $12\sqrt{m\sigma_0^2}$ of $\mathbb{S}_2$. Thus the cost introduced by decoding error is bounded as follows

$$\mathbb{E}\left[\mathcal{C}_2|\mathcal{E}_3\right] \le 144\sigma_0^2\mathbb{E}\left[P_e(\mathbb{Q})|\mathcal{E}_3\right] \tag{46}$$

which goes to zero as $m \to \infty$. Thus there exists $m_6(\epsilon_1, \delta)$ such that for all $m > m_6(\epsilon_1, \delta)$, $\mathbb{E}\left[\mathcal{C}_2|\mathcal{E}_3\right]\Pr(\mathcal{E}_3) \le \epsilon_1$.

Total average cost:

The total average cost is given by

$$\begin{aligned}
\mathbb{E}\left[\mathcal{C}\right] &= \mathbb{E}\left[\mathcal{C}_1 + \mathcal{C}_2|\mathcal{E}_1\right]\Pr(\mathcal{E}_1) + \mathbb{E}\left[\mathcal{C}_1 + \mathcal{C}_2|\mathcal{E}_2\right]\Pr(\mathcal{E}_2) \\
&\quad + \mathbb{E}\left[\mathcal{C}_1|\mathcal{E}_3\right]\Pr(\mathcal{E}_3) + \mathbb{E}\left[\mathcal{C}_2|\mathcal{E}_3\right]\Pr(\mathcal{E}_3) \\
&\le \epsilon_1 + \epsilon_1 + k^2(P + 2\epsilon_1) + \epsilon_1 \le k^2 P + (3 + 2k^2)\epsilon_1,
\end{aligned}$$

for $m > \max\{m_1(\epsilon_1, \delta), m_2(\epsilon_1, \delta), m_3(\epsilon_1, \delta), m_4(\epsilon_1, \delta), m_5(\delta), m_6(\epsilon_1, \delta)\}$. Thus, by letting $m \to \infty$, $\epsilon_1 \to 0$, and the total cost converges to k^2P. In the next section, we show that the required P can be made as small as σ_w^2 in the limit $m \to \infty$.

C.4 Required P for error probability converging to zero

Now we calculate the required P that satisfies equations (32) and (33). Let ξ satisfy $\frac{1}{2}\log_2(1+\xi) = \delta$. Then equations (32) and (33) are satisfied whenever

$$\frac{1}{2}\log_2\left(\frac{\sigma_w^2 + \sigma_0^2 - P}{\sigma_w^2}\right) = \frac{1}{2}\log_2\left(\frac{\sigma_0^2}{P}\right) + \frac{1}{2}\log_2(1+\xi),$$

$$\text{i.e., } \frac{\sigma_w^2 + \sigma_0^2 - P}{\sigma_w^2} = \frac{\sigma_0^2}{P}(1+\xi)$$

$$\text{i.e., } P^2 - P(\sigma_0^2 + \sigma_w^2) + \sigma_w^2\sigma_0^2(1+\xi) = 0. \tag{47}$$

Now, some algebra reveals that equation (47) is satisfied if

$$P = \frac{\sigma_0^2 + \sigma_w^2 - \sqrt{(\sigma_0^2 - \sigma_w^2)^2 - 4\sigma_0^2\sigma_w^2\xi^2}}{2}$$

$$= \sigma_0^2\left(\frac{1 - \sqrt{1 - \frac{4\sigma_0^2\sigma_w^2\xi^2}{(\sigma_0^2 - \sigma_w^2)^2}}}{2}\right) + \sigma_w^2\left(\frac{1 + \sqrt{1 - \frac{4\sigma_0^2\sigma_w^2\xi^2}{(\sigma_0^2 - \sigma_w^2)^2}}}{2}\right),$$

which is along the line segment joining σ_w^2 and σ_0^2, and is hence smaller than σ_0^2. For this P to exist

$$\xi < \frac{\sigma_0^2 - \sigma_w^2}{2\sigma_0\sigma_w}, \quad \text{and} \quad \delta < \frac{1}{2}\log_2\left(1 + \frac{\sigma_0^2 - \sigma_w^2}{2\sigma_0\sigma_w}\right).$$

Also, in the limit $\xi \to 0$ (or equivalently, $\delta \to 0$), P converges to σ_w^2.

D Dirty-Paper Coding (DPC) based schemes

As noted in Section 3.1, the vector Witsenhausen counterexample is similar to the communication problem of multiaccess channels with states known to some encoders (Kotagiri and Laneman, 2008), and so the strategy we propose in this section is also similar to that in Kotagiri and Laneman (2008).

D.1 A pure Dirty-Paper Coding scheme: encoding and decoding

Encoding

In this section, we describe the encoding and decoding of the DPC-based scheme. The scheme has two parameters α and ϵ. The variable P is a function of α and ϵ and can be evaluated from equation (55).

Step 1: Generate a list $\mathbb{Q}$ of $2^{m(T-\epsilon)}$ Gaussian random vectors $\mathbf{v} \sim \mathcal{N}(0, P+\alpha^2\sigma_0^2)$, where

$$T = \frac{1}{2}\log_2\left(\frac{(P+\sigma_0^2+\sigma_w^2)(P+\alpha^2\sigma_0^2)}{P\sigma_0^2(1-\alpha)^2+\sigma_w^2(P+\alpha^2\sigma_0^2)}\right). \tag{48}$$

Step 2: Given $\mathbf{x}_0$, the encoder $\underline{\underline{\mathrm{C_1}}}$ finds a $\tilde{\mathbf{v}} \in \mathbb{Q}$ such that $(\tilde{\mathbf{v}}, \mathbf{x}_0)$ satisfy,

$$\begin{aligned} \left|\left(\frac{1}{m}\sum_{i=1}^{m} x_{0,i}^2\right) - \sigma_0^2\right| &< \epsilon \\ \left|\left(\frac{1}{m}\sum_{i=1}^{m} \tilde{v}_i^2\right) - (P+\alpha^2\sigma_0^2)\right| &< \epsilon \\ \left|\left(\frac{1}{m}\sum_{i=1}^{m} x_{0,i}\tilde{v}_i\right) - \alpha\sigma_0^2\right| &< \epsilon. \end{aligned} \tag{49}$$

If more than one $\tilde{\mathbf{v}} \in \mathbb{Q}$ satisfy equation (49) for the given $\mathbf{x}_0$, then break the tie by choosing any one such $\tilde{\mathbf{v}}$. The control input is $\mathbf{u}_1 = \tilde{\mathbf{v}} - \alpha\mathbf{x}_0$. If no such $\tilde{\mathbf{v}}$ exists, we call the event an *encoding error*, and the chosen $\mathbf{u}_1 = -\mathbf{x}_0$.

Decoding

The decoder $\underline{\underline{\mathrm{C_2}}}$ receives the noisy observation $\mathbf{y}_2 = \mathbf{x}_0 + \mathbf{u}_1 + \mathbf{w}$. The decoding then proceeds in two steps.

Step 1: The decoder finds a $\hat{\mathbf{v}} \in \mathbb{Q}$ such that $(\hat{\mathbf{v}}, \mathbf{y}_2)$ satisfy

$$\begin{aligned} \left|\frac{1}{m}\sum_{i=1}^{m} \hat{v}_i^2 - (P+\alpha^2\sigma_0^2)\right| &< \epsilon \\ \left|\frac{1}{m}\sum_{i=1}^{m} y_{2,i}^2 - (\sigma_0^2+P+\sigma_w^2)\right| &< \epsilon \\ \left|\frac{1}{m}\sum_{i=1}^{m} y_{2,i}\hat{v}_i - (P+\alpha\sigma_0^2)\right| &< \epsilon. \end{aligned} \tag{50}$$

If no such $\hat{\mathbf{v}}$ exists, or more than one such $\hat{\mathbf{v}}$'s exist, the decoder decodes to $\mathbf{u}_2 = \mathbf{y}_2$ and does not continue to *Step 2*.

Step 2: If $\alpha = 1$, the decoder declares $\hat{\mathbf{v}}$ as the decoded codeword and sets $\mathbf{u}_2 = \hat{\mathbf{v}}$.

If $\alpha \neq 1$, the decoder estimates the component $x_{1,i}$ using the column vector ζ_i of length 2, where

$$\begin{aligned} \zeta_{1,i} &= \hat{v}_i, \\ \zeta_{2,i} = y_{2,i} &= u_{1,i} + x_{0,i} + w_i. \end{aligned} \tag{51}$$

The estimate is given by

$$\hat{x}_{1,i} = Q\zeta_i = q_1\zeta_{1,i} + q_2\zeta_{2,i}. \tag{52}$$

where $Q = [q_1, q_2] = K_{x_1\zeta} K_\zeta^{-1}$, where

$$K_{x_1\zeta} = [P + \alpha\sigma_0^2, P + \sigma_0^2],$$

and

$$K_\zeta = \begin{bmatrix} P + \alpha^2\sigma_0^2 & P + \alpha\sigma_0^2 \\ P + \alpha\sigma_0^2 & P + \sigma_0^2 + \sigma_w^2 \end{bmatrix}.$$

The second control $\mathbf{u}_1 = \hat{\mathbf{x}}_1$.

Simple manipulations reveal that the determinant of K_ζ is $\alpha^2(\sigma_0^2 + \sigma_w^2) + (\alpha - 1)^2 P\sigma_0^2 + P\sigma_w^2$, which is strictly positive for all values of α. Thus K_ζ is always invertible.

D.2 *Probability of encoding/decoding error*

Let $\mathcal{E}_1$ denote the event of encoding error, $\mathcal{E}_2$ the event of successful encoding *and* decoding error, and $\mathcal{E}_3$ the event of successful encoding and correct decoding.

Let $v = u_1 + \alpha x_0$ represent an *auxiliary random variable*, where u_1 and x_0 are independent random variables with $u_1 \sim \mathcal{N}(0, P)$ and $x_0 \sim \mathcal{N}(0, \sigma_0^2)$. The variable T in the number of codewords $2^{m(T-\epsilon)}$ is then given by

$$T = I(v; y_2), \tag{53}$$

where $y_2 = x_0 + u_1 + w$. With this interpretation, the encoding is based on finding $\tilde{\mathbf{v}}$ that is ϵ-jointly typical[7] with $\mathbf{x}_0$. The joint-typicality conditions (El Gamal and Cover, 1980),

$$\left| -\frac{1}{m}\log_2\left(f_v(\mathbf{v})\right) - h(v) \right| < \eta_1(\epsilon)$$
$$\left| -\frac{1}{m}\log_2\left(f_{x_0}(\mathbf{x}_0)\right) - h(x_0) \right| < \eta_1(\epsilon)$$
$$\left| -\frac{1}{m}\log_2\left(f_{v,x_0}(\mathbf{v}, \mathbf{x})\right) - h(v, x_0) \right| < \eta_3(\epsilon),$$

for appropriate $\eta_j(\epsilon)$, $j = 1, 2, 3$, are equivalent to the conditions in equation (49). Here $f_v(\cdot)$ represents the pdf of $\mathbf{v}$, and similarly for x_0 and (v, x_0), and $h(\cdot)$ is the differential entropy function.

Now, by the weak-law of large numbers, for $\mathbf{x}_0$ and $\mathbf{v}$ generated independently, the probability that the two are ϵ-jointly typical is bounded below by $(1 - \epsilon)2^{-m(I(v;x_0)+3\epsilon)}$ for m large Cover and Thomas (1991, p.195). The number of $\mathbf{v}$-codewords is $2^{m(I(v;y_2)-\epsilon)}$. If

$$I(v; y_2) = I(v; x_0) + 5\epsilon, \tag{54}$$

which is equivalent to

$$C(\alpha, P) = \frac{1}{2}\log_2\left(\frac{P(P + \sigma_0^2 + \sigma_w^2)}{P\sigma_0^2(1 - \alpha)^2 + \sigma_w^2(P + \alpha^2\sigma_0^2)}\right) = 5\epsilon, \tag{55}$$

then the average number of $\mathbf{v}$-codewords jointly typical with a typical $\mathbf{x}_0$ increases exponentially in m, and the probability of encoding error $\Pr(\mathcal{E}_1)$ decreases to zero exponentially in m Cover and Thomas (1991, pp.353–356). That equation (55) has a solution is shown in Appendix D.5.

The decoding fails in two cases. If the encoding fails, so might the decoding, but since the error probability of encoding error decreases to zero exponentially, so does the probability of this kind of decoding error.

If the encoding succeeds, the transmitted $\tilde{\mathbf{v}}$ can be decoded correctly as long as the rate is smaller than the mutual information across the $(\tilde{v}, y_2)$ 'channel'. Since the number of $\tilde{\mathbf{v}}$ codewords is $2^{m(I(v;y_2)-\epsilon)}$, the rate for the $\mathbf{v}$-codebook is $R_v = I(v; y_2) - \epsilon$ which is clearly smaller than $I(v; y_2)$. Thus the probability of decoding error conditioned on the success of the encoding converges to zero exponentially in m. Since $\Pr(\mathcal{E}_1^c) \to 1$, the probability $\Pr(\mathcal{E}_2) \to 0$ as well.

D.3 Encoding cost analysis

The average total cost is given by

$$\begin{aligned}\mathbb{E}[\mathcal{C}] &= \mathbb{E}[\mathcal{C}_1 + \mathcal{C}_2]\\ &= \mathbb{E}[\mathcal{C}_1 \mid \mathcal{E}_1]\Pr(\mathcal{E}_1) + \mathbb{E}[\mathcal{C}_1 \mid \mathcal{E}_2 \cup \mathcal{E}_3]\Pr(\mathcal{E}_2 \cup \mathcal{E}_3)\\ &\quad + \mathbb{E}[\mathcal{C}_2 \mid \mathcal{E}_1]\Pr(\mathcal{E}_1) + \mathbb{E}[\mathcal{C}_2 \mid \mathcal{E}_2]\Pr(\mathcal{E}_2) + \mathbb{E}[\mathcal{C}_2 \mid \mathcal{E}_3]\Pr(\mathcal{E}_3).\end{aligned}$$

We will first calculate the encoding costs, followed by the decoding costs.

For the event $\mathcal{E}_1$,

$$\mathbb{E}[\mathcal{C}_1 \mid \mathcal{E}_1]\Pr(\mathcal{E}_1) = \frac{k^2}{m}\mathbb{E}\left[\|\mathbf{x}_0\|^2 \mathbb{1}_{\{\mathcal{E}_1\}}\right] \to 0$$

by Lemma 4, since $\Pr(\mathcal{E}_1) \to 0$. Thus, for any given $\epsilon_1 > 0$, there exists $m_0(\epsilon, \epsilon_1)$ such that for all $m > m_0(\epsilon, \epsilon_1)$, $\mathbb{E}[\mathcal{C}_1 \mid \mathcal{E}_1]\Pr(\mathcal{E}_1) < \epsilon_1$.

For $\mathcal{E}_2 \cup \mathcal{E}_3$, the event of successful encoding, since $\mathbf{u} = \tilde{\mathbf{v}} - \alpha\mathbf{x}_0$ for $(\mathbf{v}, \mathbf{x})$ satisfying equation (49),

$$\begin{aligned}\frac{1}{m}\mathbb{E}\left[\|\mathbf{u}_1\|^2 \mid \mathcal{E}_2 \cup \mathcal{E}_3\right] &= \frac{1}{m}\mathbb{E}\left[\|\mathbf{v}\|^2 + \alpha^2\|\mathbf{x}_0\|^2 - 2\alpha\mathbf{v}^T\mathbf{x}_0 \mid \mathcal{E}_2 \cup \mathcal{E}_3\right]\\ &\le P + \alpha^2\sigma_0^2 + \epsilon + \alpha^2(\sigma_0^2 + \epsilon)\\ &\quad - 2\alpha(\alpha\sigma_0^2) + 2|\alpha|\epsilon = P + (1 + |\alpha|)^2\epsilon.\end{aligned}$$

Thus $\mathbb{E}[\mathcal{C}_1 \mid \mathcal{E}_2 \cup \mathcal{E}_3]\Pr(\mathcal{E}_2 \cup \mathcal{E}_3) \le k^2(P + (1 + |\alpha|)^2\epsilon)$. Thus the total encoding costs are also bounded by $k^2(P + (1 + |\alpha|)^2\epsilon) + \epsilon_1$.

D.4 Decoding cost analysis for $\alpha = 1$

We first concentrate on the easier case $\alpha = 1$.

For the event $\mathcal{E}_1$, since $\mathbf{u} = -\mathbf{x}_0$, and $\mathbf{x}_1 = 0$. There are two cases, the decoder decodes to some erroneous $\hat{\mathbf{v}}$, or the decoder fails to decode and uses $\mathbf{u}_2 = \mathbf{y}_2$. Therefore,

$$\begin{aligned}\mathbb{E}\left[\mathcal{C}_2 \mathbb{1}_{\{\mathcal{E}_1\}}\right] &\leq m^{-1}\mathbb{E}\left[\max\{\|\hat{\mathbf{v}} - 0\|^2, \|\mathbf{y}_2 - 0\|^2\}\mathbb{1}_{\{\mathcal{E}_1\}}\right] \\ &\leq m^{-1}\mathbb{E}\left[(\|\hat{\mathbf{v}}\|^2 + \|\mathbf{y}_2\|^2)\mathbb{1}_{\{\mathcal{E}_1\}}\right] \\ &\overset{(a)}{\leq} \mathbb{E}\left[(P + \alpha^2\sigma_0^2 + \epsilon)\mathbb{1}_{\{\mathcal{E}_1\}}\right] + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] \\ &= (P + \alpha^2\sigma_0^2 + \epsilon)\Pr(\mathcal{E}_1) + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right]\end{aligned}$$

where (a) follows from equation (50) because the decoded $\hat{\mathbf{v}}$ is inside a sphere of radius $\sqrt{m(P + \sigma_0^2 + \epsilon)}$, and under $\mathcal{E}_1$, $\mathbf{y}_2 = \mathbf{w}$. Since $\Pr(\mathcal{E}_1) \to 0$ as $m \to \infty$, the first term goes to zero. Similarly, the second term goes to zero by Lemma 4. Thus there exists $m_1(\epsilon, \epsilon_1)$ such that for all $m > m_1(\epsilon, \epsilon_1)$, $\mathbb{E}[\mathcal{C}_2|\mathcal{E}_1]\Pr(\mathcal{E}_1) < \epsilon_1$.

For the event $\mathcal{E}_2$, again, there are the two cases of erroneous decoding and decoding failure. Thus

$$\begin{aligned}\mathbb{E}\left[\mathcal{C}_2 \mathbb{1}_{\{\mathcal{E}_2\}}\right] &\leq m^{-1}\mathbb{E}\left[\max\{\|\hat{\mathbf{v}} - \tilde{\mathbf{v}}\|^2, \|\mathbf{y}_2 - \tilde{\mathbf{v}}\|^2\}\mathbb{1}_{\{\mathcal{E}_2\}}\right] \\ &\leq m^{-1}\mathbb{E}\left[\|\hat{\mathbf{v}} - \tilde{\mathbf{v}}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right] + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right] \\ &\leq 4(P + \alpha^2\sigma_0^2 + \epsilon)\Pr(\mathcal{E}_2) + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right].\end{aligned}$$

The first term goes to zero because $\Pr(\mathcal{E}_2) \to 0$. The second term goes to zero by Lemma 4. Thus for given $\epsilon_1 > 0$, there exists $m_2(\epsilon, \epsilon_1)$ such that for all $m > m_2(\epsilon, \epsilon_1)$, $\mathbb{E}[\mathcal{C}_2|\mathcal{E}_2]\Pr(\mathcal{E}_2) < \epsilon_1$.

For the event $\mathcal{E}_3$, the encoding is successful and the decoding is correct, therefore $\tilde{\mathbf{v}} = \hat{\mathbf{v}}$, and thus $\mathbb{E}[\mathcal{C}_2|\mathcal{E}_3] = 0$.

D.5 Total costs for $\alpha = 1$

For $\alpha = 1$, for $m > \max\{m_0(\epsilon, \epsilon_1), m_1(\epsilon, \epsilon_1), m_2(\epsilon, \epsilon_1)\}$, $\mathbb{E}[\mathcal{C}_2] \leq 3\epsilon_1$, and $\mathbb{E}[\mathcal{C}_1] \leq k^2(P + 4\epsilon) + \epsilon_1$, and the total (encoding and decoding) cost for $\alpha = 1$ is therefore smaller than

$$k^2(P + 4\epsilon) + 4\epsilon_1. \tag{56}$$

Making m large enough, ϵ_1 can be made as small as desired. Using the fact that this holds for all ϵ, in the following, we show that letting $\epsilon \to 0$, the achievable asymptotic cost is

$$k^2\sigma_0^2 \frac{\sqrt{1 + \frac{4\sigma_w^2}{\sigma_0^2}} - 1}{2}.$$

From equation (55), for $\alpha = 1$, P needs to satisfy $C(1, P) = 5\epsilon$, where

$$C(1, P) = \frac{1}{2}\log_2\left(\frac{P(P + \sigma_0^2 + \sigma_w^2)}{(P + \sigma_0^2)\sigma_w^2}\right). \tag{57}$$

Let ξ be such that $\epsilon = \frac{1}{10}\log_2(1+\xi)$. Then,

$$\frac{1}{2}\log_2\left(\frac{P(P+\sigma_0^2+\sigma_w^2)}{(P+\sigma_0^2)\sigma_w^2}\right) = \frac{1}{2}\log_2(1+\xi)$$
$$\text{i.e. } P^2 + (\sigma_0^2 - \xi\sigma_w^2)P - (1+\xi)\sigma_0^2\sigma_w^2 = 0.$$

Taking the positive root of the quadratic equation,

$$P = (\sigma_0^2 - \xi\sigma_w^2)\frac{\sqrt{1+\frac{4(1+\xi)^2\sigma_w^2\sigma_0^2}{(\sigma_0^2-\xi\sigma_w^2)^2}} - 1}{2}. \tag{58}$$

Now letting ϵ go to zero (and thus $\xi \to 0$) by increasing m to infinity, the required P approaches $\sigma_0^2\frac{\sqrt{1+4\sigma_w^2/\sigma_0^2}-1}{2}$. The asymptotic expected cost for the scheme is, therefore, $k^2\sigma_0^2\frac{\sqrt{1+4\sigma_w^2/\sigma_0^2}-1}{2}$. This expression turns out to be an increasing function in σ_0^2 which is bounded above by $k^2\sigma_w^2$, the cost for the JSCC scheme. Thus even in the special case of $\alpha = 1$, the DPC scheme asymptotically outperforms the JSCC scheme.

D.6 Decoding cost analysis for $\alpha \neq 1$

For the event $\mathcal{E}_1$, similar to the analysis of $\alpha = 1$, there are the two cases of decoding failure and decoding error,

$$\begin{aligned}
\mathbb{E}\left[\mathcal{C}_2\mathbb{1}_{\{\mathcal{E}_1\}}\right] &\leq \frac{1}{m}\mathbb{E}\left[\max\{\|q_1\hat{\mathbf{v}} + q_2\mathbf{w}\|^2, \|\mathbf{y} - 0\|^2\}\mathbb{1}_{\{\mathcal{E}_1\}}\right] \\
&\leq \frac{1}{m}\mathbb{E}\left[\|q_1\hat{\mathbf{v}} + q_2\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] + \frac{1}{m}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] && (59) \\
&= \frac{q_1^2}{m}\mathbb{E}\left[\|\hat{\mathbf{v}}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] + \frac{q_2^2}{m}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] \\
&\quad + \frac{2q_1q_2}{m}\mathbb{E}\left[\hat{\mathbf{v}}^T\mathbf{w}\mathbb{1}_{\{\mathcal{E}_1\}}\right] + \frac{1}{m}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] && (60) \\
&= \frac{q_1^2(P+\sigma_0^2+\epsilon)}{m}\mathbb{E}\left[\mathbb{1}_{\{\mathcal{E}_1\}}\right] + \frac{q_2^2}{m}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right] && (61) \\
&\quad + \frac{2q_1q_2}{m}\mathbb{E}\left[\hat{\mathbf{v}}^T\mathbf{w}\mathbb{1}_{\{\mathcal{E}_1\}}\right] + \frac{1}{m}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_1\}}\right]. && (62)
\end{aligned}$$

As $m \to \infty$, the first term converges to zero since $\Pr(\mathcal{E}_1) \to 0$. The second term and the fourth terms converge to zero by an application of Lemma 4 with $\mathbf{x} = \mathbf{w}$. Looking at the third term in equation (59),

$$\begin{aligned}
m^{-1}\mathbb{E}\left[\hat{\mathbf{v}}^T\mathbf{w}\mathbb{1}_{\{\mathcal{E}_1\}}\right] &= m^{-1}\mathbb{E}\left[\left(\sum_{i=1}^{m}\hat{v}_iw_i\right)\mathbb{1}_{\{\mathcal{E}_1\}}\right] \\
&\overset{(a)}{\leq} m^{-1}\mathbb{E}\left[\|\hat{\mathbf{v}}\|\|\mathbf{w}\|\mathbb{1}_{\{\mathcal{E}_1\}}\right]
\end{aligned}$$

$$\leq m^{-1}\left(\sqrt{P+\sigma_0^2+\epsilon}\right)\mathbb{E}\left[\sqrt{\|\mathbf{w}\|^2 \mathbb{1}_{\{\mathcal{E}_1\}}}\right]$$

$$\overset{(b)}{\leq} m^{-1}\sqrt{P+\sigma_0^2+\epsilon}\sqrt{\mathbb{E}\left[\|\mathbf{w}\|^2 \mathbb{1}_{\{\mathcal{E}_1\}}\right]}$$

where (a) follows from the Cauchy-Schwartz inequality, and (b) follows from Jensen's inequality. By Lemma 4, the third term now converges to zero as well. Therefore, for all $\epsilon_1 > 0$, there exists $m_3(\alpha, \epsilon, \epsilon_1)$ such that for all $m > m_3(\alpha, \epsilon, \epsilon_1)$, $\mathbb{E}\left[\mathcal{C}_2|\mathcal{E}_1\right]\Pr(\mathcal{E}_1) < \epsilon_1$.

Similarly, for the event $\mathcal{E}_2$,

$$\begin{aligned}
\mathbb{E}\left[\mathcal{C}_2 \mathbb{1}_{\{\mathcal{E}_2\}}\right] &\leq m^{-1}\mathbb{E}\left[\max\{\|q_1\hat{\mathbf{v}}+q_2\tilde{\mathbf{v}}\|^2, \|\mathbf{y}-\tilde{\mathbf{v}}\|^2\}\mathbb{1}_{\{\mathcal{E}_2\}}\right] \\
&\leq m^{-1}\mathbb{E}\left[\|q_1\hat{\mathbf{v}}+q_2\tilde{\mathbf{v}}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right] + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right] \\
&\leq m^{-1}\mathbb{E}\left[(q_1^2\|\hat{\mathbf{v}}\|^2 + q_2^2\|\tilde{\mathbf{v}}\|^2 + 2q_1q_2\|\tilde{\mathbf{v}}\|\,\|\hat{\mathbf{v}}\|)\mathbb{1}_{\{\mathcal{E}_2\}}\right] + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right] \\
&\leq (P+\alpha^2\sigma_0^2+\epsilon)(q_1+q_2)^2\Pr(\mathcal{E}_2) + m^{-1}\mathbb{E}\left[\|\mathbf{w}\|^2\mathbb{1}_{\{\mathcal{E}_2\}}\right].
\end{aligned}$$

The first term converges to zero as $\Pr(\mathcal{E}_2)$ converges to zero. The second converges to zero by an application of Lemma 4. Thus there exists $m_4(\alpha, \epsilon, \epsilon_1)$ such that for all $m > m_4(\alpha, \epsilon, \epsilon_1)$, $\mathbb{E}\left[\mathcal{C}_2|\mathcal{E}_2\right]\Pr(\mathcal{E}_2) < \epsilon_1$.

For the event $\mathcal{E}_3$, we need the following lemma.

Lemma 5: *For vectors* $\mathbf{v}$ *and* $\mathbf{x}$*, if* $(\mathbf{v}, \mathbf{x})$ *are* ϵ*-jointly typical as defined by equation* (49)*, then for* $\mathbf{u} = \mathbf{v} - \alpha\mathbf{x}$*,*

$$\left|\sum_{i=1}^{m}\frac{u_{1,i}x_{0,i}}{m}\right| < (|\alpha|+1)\epsilon, \tag{63}$$

and

$$\left|\sum_{i=1}^{m}\frac{u_{1,i}^2}{m}\right| < P + (1+|\alpha|)^2\epsilon. \tag{64}$$

Proof: Since $(\mathbf{v}, \mathbf{x})$ are ϵ-jointly typical, using equation (49), □

$$\left|\sum_{i=1}^{m}\frac{v_i x_{0,i}}{m} - \alpha\sigma_0^2\right| < \epsilon$$

$$\text{i.e.} \left|\sum_{i=1}^{m}\frac{(u_{1,i}+\alpha x_{0,i})x_{0,i}}{m} - \alpha\sigma_0^2\right| < \epsilon$$

$$\text{i.e.} \left|\sum_{i=1}^{m}\frac{u_{1,i}x_{0,i}}{m} + \alpha\sum_{i=1}^{m}\frac{x_{0,i}^2}{m} - \alpha\sigma_0^2\right| < \epsilon$$

$$\text{i.e.} \left|\sum_{i=1}^{m}\frac{u_{1,i}x_{0,i}}{m} + \alpha\left(\sum_{i=1}^{m}\frac{x_{0,i}^2}{m} - \sigma_0^2\right)\right| < \epsilon.$$

By the ϵ-typicality of $\mathbf{x}$ in equation (49), the term in brackets is bounded by ϵ in absolute value. Thus, equation (63) follows and asymptotically, $\mathbf{u}_1$ and $\mathbf{x}_0$ appear statistically uncorrelated.

Similarly, for $\mathbf{u} = \tilde{\mathbf{v}} - \alpha \mathbf{x}_0$,

$$\begin{aligned}\frac{1}{m}\|\mathbf{u}\|^2 &= \frac{1}{m}\left(\|\tilde{\mathbf{v}}\|^2 + \alpha^2\|\mathbf{x}_0\|^2 - 2\alpha\sum_{i=1}^{m}\tilde{v}_i x_{0,i}\right)\\ &\leq (P + \alpha^2\sigma_0^2 + \epsilon) + \alpha^2(\sigma_0^2 + \epsilon) - 2\alpha(\alpha\sigma^2) + 2|\alpha|\epsilon\\ &= P + (1+|\alpha|)^2\epsilon.\end{aligned}$$

Now, for $\mathcal{E}_3$, the normalised mean-square error is given by

$$\begin{aligned}&m^{-1}\mathbb{E}\left[\|\hat{\mathbf{x}}_1 - \mathbf{x}_1\|^2|\mathcal{E}_3\right]\\ &\quad= m^{-1}\mathbb{E}\left[\|Q\zeta - \mathbf{x}_1\|^2\,|\,\mathcal{E}_3\right]\\ &\quad= m^{-1}\mathbb{E}\left[\|q_1\mathbf{v} + q_2\mathbf{y}_2 - \mathbf{x}_1\|^2\,|\,\mathcal{E}_3\right]\\ &\quad= m^{-1}\mathbb{E}\left[\|q_1(\mathbf{u}_1 + \alpha\mathbf{x}_0) + q_2(\mathbf{u}_1 + \mathbf{x}_0 + \mathbf{w}) - \mathbf{x}_1\|^2\,|\,\mathcal{E}_3\right]\\ &\quad= m^{-1}\mathbb{E}\left[\|(q_1 + q_2 - 1)\mathbf{u}_1 + (\alpha q_1 + q_2 - 1)\mathbf{x}_0 + q_2\mathbf{w}\|^2\,|\,\mathcal{E}_3\right]\\ &\quad= m^{-1}((q_1+q_2-1)^2\mathbb{E}\left[\|\mathbf{u}_1\|^2\,|\,\mathcal{E}_3\right] + (\alpha q_1 + q_2 - 1)^2\mathbb{E}\left[\|\mathbf{x}_0\|^2|\mathcal{E}_3\right] + q_2^2\mathbb{E}\left[\|\mathbf{w}\|^2\,|\,\mathcal{E}_3\right]\\ &\qquad+ 2(q_1 + q_2 - 1)(\alpha q_1 + q_2 - 1)\mathbb{E}\left[\mathbf{u}_1^T\mathbf{x}_0\,|\,\mathcal{E}_3\right] + 2(q_1 + q_2 - 1)q_2\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\,|\,\mathcal{E}_3\right]\\ &\qquad+ 2(\alpha q_1 + q_2 - 1)q_2\mathbb{E}\left[\mathbf{x}_0^T\mathbf{w}\,|\,\mathcal{E}_3\right]),\end{aligned} \tag{65}$$

where $(\cdot)^T$ denotes the transpose of the vector. Now,

$$\begin{aligned}m^{-1}\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\right] = 0 &= m^{-1}\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\,|\,\mathcal{E}_3\right]\Pr(\mathcal{E}_3) + m^{-1}\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\,|\,\mathcal{E}_3^c\right]\Pr(\mathcal{E}_3^c)\\ \text{i.e., } m^{-1}\left|\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\,|\,\mathcal{E}_3\right]\right|\Pr(\mathcal{E}_3) &= m^{-1}\left|\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\,|\,\mathcal{E}_3^c\right]\right|\Pr(\mathcal{E}_3^c)\\ &= m^{-1}\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}\mathbb{1}_{\{\mathcal{E}_3^c\}}\right]\\ &\overset{(a)}{\leq} \sqrt{\frac{1}{m^2}\mathbb{E}\left[(\mathbf{u}_1^T\mathbf{w})^2\right]\mathbb{E}\left[\mathbb{1}_{\{\mathcal{E}_3^c\}}\right]}.\end{aligned}$$

where (a) follows from the Cauchy-Schwartz inequality. Now,

$$\begin{aligned}\frac{1}{m^2}\mathbb{E}\left[\left(\mathbf{u}_1^T\mathbf{w}\right)^2\right] &= \frac{1}{m^2}\mathbb{E}\left[\left(\sum_{i=1}^{m}u_{1,i}w_i\right)^2\right]\\ &\overset{(b)}{=} \frac{1}{m^2}\mathbb{E}\left[\sum_{i=1}^{m}u_{1,i}^2 w_i^2\right]\\ &\overset{(c)}{=} \frac{1}{m^2}\sum_{i=1}^{m}\mathbb{E}\left[u_{1,i}^2\right]\sigma_w^2\\ &= \frac{1}{m^2}\mathbb{E}\left[\sum_{i=1}^{m}u_{1,i}^2\right]\sigma_w^2 = \frac{1}{m^2}\mathbb{E}\left[\|\mathbf{u}\|^2\right]\sigma_w^2\end{aligned}$$

where (b) and (c) follow from the fact that $u_{1,i}$ and w_j are uncorrelated, and w_i is zero mean. Now,

$$\begin{aligned}\mathbb{E}\left[\|\mathbf{u}\|^2\right] &= \mathbb{E}\big[\|\mathbf{u}\|^2 \mathbb{1}_{\{\mathcal{E}_1^c\}}\big] + \mathbb{E}\left[\|\mathbf{u}\|^2 \mathbb{1}_{\{\mathcal{E}_1\}}\right] \\ &\leq m(P + (1+|\alpha|)^2\epsilon)\Pr(\mathcal{E}_1^c) + \mathbb{E}\left[\|\mathbf{x}_0\|^2 \mathbb{1}_{\{\mathcal{E}_1\}}\right] \\ &\leq m(P + (1+|\alpha|)^2\epsilon) + \mathbb{E}\left[\|\mathbf{x}_0\|^2\right] \\ &\leq m(P + (1+|\alpha|)^2\epsilon + \sigma_0^2).\end{aligned}$$

Thus,

$$\frac{1}{m^2}\mathbb{E}\left[(\mathbf{u}_1^T\mathbf{w})^2\right]\mathbb{E}\big[\mathbb{1}_{\{\mathcal{E}_3^c\}}\big] \leq \frac{1}{m}(P + (1+|\alpha|)^2\epsilon + \sigma_0^2)\Pr(\mathcal{E}_3^c). \tag{66}$$

Since $\Pr(\mathcal{E}_3^c) \to 0$ as $m \to \infty$, $m^{-1}\mathbb{E}\left[\mathbf{u}_1^T\mathbf{w}|\mathcal{E}_3\right] \to 0$ as $m \to \infty$. Similarly, $m^{-1}\mathbb{E}\left[\mathbf{x}_0^T\mathbf{w}|\mathcal{E}_3\right]$ also converges to zero. Using these results and Lemma 5 for $m^{-1}\mathbb{E}\left[\mathbf{u}_1^T\mathbf{x}_0|\mathcal{E}_3\right]$, there exists $m_5(\alpha, \epsilon, \epsilon_1)$ such that for all $m > m_5(\alpha, \epsilon, \epsilon_1)$, the last three terms in equation (65) are smaller than ϵ_1.

Thus using the ϵ-joint typicality of $\mathbf{v}$ and $\mathbf{x}_0$ equation (49),

$$\begin{aligned}&\frac{1}{m}\mathbb{E}\left[\|\hat{\mathbf{x}}_1 - \mathbf{x}_1\|^2|\mathcal{E}_3\right] \\ &\leq (q_1+q_2-1)^2(P + (1+|\alpha|)^2\epsilon) + (\alpha q_1 + q_2 - 1)^2(\sigma_0^2 + \epsilon) + q_2^2\sigma_w^2 + \epsilon_1 \\ &\leq (q_1+q_2-1)^2 P + (\alpha q_1 + q_2 - 1)^2\sigma_0^2 + q_2^2\sigma_w^2 + \epsilon M(q_1, q_2) + \epsilon_1\end{aligned} \tag{67}$$

where the constant $\epsilon M(q_1, q_2) = \epsilon((q_1+q_2-1)^2(1+|\alpha|)^2 + (\alpha q_1 + q_2 - 1)^2) \to 0$ as $\epsilon \to 0$.

Now suppose the random variables $\mathbf{x}_0' \sim \mathcal{N}(0, \sigma_0^2\mathbb{I})$ and $\mathbf{u}_1' \sim \mathcal{N}(0, P\mathbb{I})$ are independent of $\mathbf{w}' \sim \mathcal{N}(0, \sigma_w^2\mathbb{I})$ and jointly Gaussian with covariance between $\mathbf{u}_1'$ and $\mathbf{x}_0'$ zero. Also define $\mathbf{v}' = \mathbf{u}' + \alpha\mathbf{x}_0'$ and $\mathbf{y}' = \mathbf{x}_0' + \mathbf{u}_1' + \mathbf{w}'$. A simple calculation shows that $\mathbb{E}\left[\|q_1\mathbf{v}' + q_2\mathbf{y}_2' - \mathbf{x}_1'\|^2\right]$ yields the same expression as in equation (67) without the ϵM and the ϵ_1 terms. Observe that the estimation procedure in equation (52) is the MMSE estimation for $\mathbf{x}_1' = \mathbf{x}_0' + \mathbf{u}_1'$ on observing $\mathbf{y}_2'$, where the average error is given by $MMSE_G(\alpha, P)$ is given by Kay (1998, p.382)

$$MMSE_G(\alpha, P) = K_{x_1} - K_{x_1\zeta}K_\zeta^{-1}K_{x_1\zeta}, \tag{68}$$

where,

$$K_{x_1} = \mathbb{E}\left[x_1^2\right] = P + \sigma_0^2. \tag{69}$$

D.7 Total costs for $\alpha \neq 1$

For fixed α and ϵ, if $m > \{m_3(\alpha, \epsilon, \epsilon_1), m_4(\alpha, \epsilon, \epsilon_1), m_5(\alpha, \epsilon, \epsilon_1)\}$, the total cost is smaller than

$$k^2(P + (1+|\alpha|)^2) + MMSE(\alpha, P) + M(q_1, q_2)\epsilon + \epsilon_1, \tag{70}$$

where P satisfies

$$C(\alpha, P) = I(v; y_2) - I(v; x_0) = \frac{1}{2}\log_2\left(\frac{P(P+\sigma_0^2+\sigma_w^2)}{P\sigma_0^2(1-\alpha)^2+\sigma_w^2(P+\alpha^2\sigma_0^2)}\right) = 5\epsilon.$$

By arguments analogous to those in Appendix D.5, the existence of a solution can be proved. Here we concentrate on the case of interest of $\epsilon \to 0$ by letting $m \to \infty$. The condition (55) is then equivalent to

$$P(P+\sigma_0^2+\sigma_w^2) = P\sigma_0^2(1-\alpha)^2+\sigma_w^2(P+\alpha^2\sigma_0^2).$$

Taking the positive root,

$$P = \frac{\sqrt{\sigma_0^2\alpha(2-\alpha)}}{2}\left(\sqrt{1+\frac{4\sigma_w^2}{\sigma_0^2(2-\alpha)^2}}-1\right). \tag{71}$$

By letting $m \to \infty$, we can have $\epsilon_1 \to 0$ and also $\epsilon \to 0$. Optimising the total cost over α, the asymptotic total cost achieved is

$$\min_{\alpha,P} k^2P + MMSE(\alpha, P), \tag{72}$$

for P satisfying equation (71).

D.8 A combination of the linear scheme and the DPC scheme

While the DPC scheme outperforms the linear scheme for high values of σ_0^2, Figure 12 shows that this is not true for low values of σ_0^2. It is natural to think that a combination of the linear scheme and the DPC scheme might perform better than either one alone.

Our approach is to divide the power available and dedicate a part of it to each scheme. A linear term $\mathbf{u}_{1,1} = -\beta\mathbf{x}_0$ constitutes a part of $\mathbf{u}_1$. This term can be thought of as changing the variance of the initial state $\mathbf{x}_0$. To this, we add a vector $\mathbf{u}_{1,2}$ that is dirty-paper coded against $(1-\beta)\mathbf{x}_0$. By Lemma 5, $\mathbf{u}_{1,2}$ is statistically uncorrelated with $(1-\beta)\mathbf{x}_0$, and hence uncorrelated with $\mathbf{u}_{1,1}$ as well. Thus, the total power is $\beta^2\sigma_0^2 + \frac{1}{m}\mathbb{E}\left[\|\mathbf{u}_{1,2}\|^2\right] + \epsilon$, and this is constrained to be smaller than P. The state $\mathbf{x}_1$ is, therefore,

$$\mathbf{x}_1 = (1-\beta)\mathbf{x}_0 + \mathbf{u}_{1,2}, \tag{73}$$

where $\mathbf{u}_{1,2}$ is the input that dirty-paper codes against $(1-\beta)\mathbf{x}_0$. Again, $\epsilon \to 0$ as $m \to \infty$. It turns out that the optimising $\beta \geq 0$, consistent with the intuition that the power in the initial state $\mathbf{x}_0$ should be decreased before using DPC against the state.

E The vector schemes perform within a constant factor of the optimal

The performance of the scheme that zero-forces $\mathbf{x}_0$ and the JSCC scheme is identical for $\sigma_0^2 = 1$, as is evident from Figure 6. Therefore, we consider two different cases: $\sigma_0^2 \leq 1$ and $\sigma_0^2 \geq 1$. In either case, we show that the ratio is bounded by 11. The result can be tightened by a more detailed analysis by dividing the (k, σ_0^2) space into finer partitions. However, we do not present the detailed analysis here for ease of exposition.

Region 1: $\sigma_0^2 \leq 1$.

We consider the upper bound as the minimum of $k^2\sigma_0^2$ and $\frac{\sigma_0^2}{\sigma_0^2+1}$. Consider the lower bound

$$\overline{\mathcal{C}} \geq \min_{P \geq 0} k^2 P + \left(\left(\sqrt{\kappa(P)} - \sqrt{P}\right)^+\right)^2. \tag{74}$$

Now if the optimising power P is greater than $\sigma_0^2/11$, then the first term of the lower bound is greater than $k^2\sigma_0^2/11$. Thus the ratio of the upper bound $k^2\sigma_0^2$ and the lower bound is smaller than 11.

If the optimising $P \leq \frac{\sigma_0^2}{11}$,

$$\begin{aligned} \kappa(P) &= \frac{\sigma_0^2}{(\sigma_0 + \sqrt{P})^2 + 1} \\ &\geq \frac{\sigma_0^2}{(\sigma_0 + \frac{\sigma_0}{\sqrt{11}})^2 + 1} \\ &\overset{(\sigma_0^2 \leq 1)}{\geq} \frac{\sigma_0^2}{(1 + \frac{1}{\sqrt{11}})^2 + 1} \geq 0.37\sigma_0^2 \end{aligned}$$

which is greater than $\sigma_0^2/11 > P$. Thus,

$$\begin{aligned} ((\sqrt{\kappa(P)} - \sqrt{P})^+)^2 &\geq \left(\sqrt{0.37\sigma_0^2} - \sqrt{\frac{\sigma_0^2}{11}}\right)^2 \\ &> 0.094\sigma_0^2 > \frac{\sigma_0^2}{11}. \end{aligned}$$

The lower bound is no smaller than $((\sqrt{\kappa(P)} - \sqrt{P})^+)^2$. Thus, even for $P \leq \frac{\sigma_0^2}{11}$ the ratio of the upper bound $\frac{\sigma_0^2}{\sigma_0^2+1}$ and the lower bound is smaller than 11.

Region 2: $\sigma_0^2 \geq 1$.

The upper bound relevant here is the minimum of k^2 and $\frac{\sigma_0^2}{\sigma_0^2+1}$. Again, looking at equation (74), if $P > \frac{1}{11}$, the ratio of the upper bound k^2 to the lower bound is no more than 11.

Now, if $P \leq \frac{1}{11}$,

$$\kappa(P) \geq \frac{\sigma_0^2}{(\sigma_0 + 1/\sqrt{11})^2 + 1}.$$

Therefore,

$$((\sqrt{\kappa(P)} - \sqrt{P})^+)^2 = \left(\left(\sqrt{\frac{\sigma_0^2}{(\sigma_0 + \frac{1}{\sqrt{11}})^2 + 1}} - \frac{1}{\sqrt{11}}\right)^+\right)^2.$$

For $\sigma_0^2 \geq 1$, the first term on the RHS attains its minima at $\sigma_0^2 = 1$. Evaluated at this point, the term is larger than $\frac{1}{\sqrt{11}}$. Therefore, a bound on the ratio for $P < \frac{1}{11}$ is

$$\begin{aligned}\frac{\sigma_0^2/(\sigma_0^2+1)}{\left(\sqrt{\frac{\sigma_0^2}{(\sigma_0+\frac{1}{\sqrt{11}})^2+1}} - \frac{1}{\sqrt{11}}\right)^2} &\leq \frac{1}{\left(\sqrt{\frac{\sigma_0^2}{(\sigma_0+\frac{1}{\sqrt{11}})^2+1}} - \frac{1}{\sqrt{11}}\right)^2} \\ &\leq \frac{1}{\left(\sqrt{\frac{1}{(1+\frac{1}{\sqrt{11}})^2+1}} - \frac{1}{\sqrt{11}}\right)^2} \\ &\approx 10.56 < 11.\end{aligned}$$

Thus, for $\sigma_0^2 \geq 1$, the ratio is bounded by 11 as well. Therefore, γ_1 and γ_2 are both smaller than 11.