

A Deterministic Annealing Approach to Witsenhausen's Counterexample

Mustafa Mehmetoglu
 Dep. of Electrical-Computer Eng.
 UC Santa Barbara, CA, US
 Email: mehmetoglu@ece.ucsb.edu

Emrah Akyol
 Dep. of Electrical Eng.
 USC, CA, US
 Email: eakyol@usc.edu

Kenneth Rose
 Dep. of Electrical-Computer Eng.
 UC Santa Barbara, CA, US
 Email: rose@ece.ucsb.edu

Abstract—This paper proposes an optimization method, based on information theoretic ideas, to a class of distributed control problems. As a particular test case, the well-known and numerically “over-mined” problem of decentralized control and implicit communication, commonly referred to as Witsenhausen’s counterexample, is considered. The key idea is to randomize the zero-delay mappings, which become “soft”, probabilistic mappings to be optimized in a deterministic annealing process, by incorporating a Shannon entropy constraint in the problem formulation. The entropy of the mapping is controlled and gradually lowered to zero to obtain deterministic mappings, while avoiding poor local minima. For the particular test case, our approach obtains new mappings that shed light on the structure of the optimal solution, as well as achieving a small improvement in total cost over the state of the art in numerical approaches to this problem. Proposed method is general and applicable to any problem of similar nature.

I. INTRODUCTION

Witsenhausen’s counterexample (WCE) [1] to a conjecture in decentralized control theory is a deceptively simple stochastic control problem that has remained unsolved for several decades. The fact that the sum of jointly Gaussian random variables is Gaussian yields a simple solution to the centralized linear quadratic Gaussian (LQG) control problems where optimal control mappings are known to be linear in the observations. This observation led to a natural conjecture that linear control mappings remain optimal even in decentralized settings. However, Witsenhausen provided an example of a decentralized LQG control problem wherein non-linear mappings outperform linear mappings. The problem has been viewed as a benchmark in stochastic networked control, see eg. [2] for a detailed treatment.

Beyond its key strategic importance in decentralized control systems, WCE has implications as an “implicit communication” problem. Connections between information theory and WCE were first studied in [3], and extensions of the original problem to vector settings were studied from an information theory perspective, see eg. [4], [5]. The functional properties of the problem have been extensively studied in the literature, see e.g., [1], [6]–[8]

In this paper, we propose an optimization method, derived from information theoretic principles, which can be applied to a class of distributed control problems, and specifically to WCE. There has been a significant amount of prior work on

optimization methods (see e.g. [9]–[13]), which can be clustered into two camps. The first camp is based on optimization of “structured” continuous mappings where a parametric form is assumed and these parameters are optimized. The other camp attacks a related discrete problem with the argument that the discrete problem asymptotically approaches to the original analog problem, at increasingly fine discretization. The approach in this paper belongs to the first camp and employs a powerful nonconvex optimization method, namely deterministic annealing (DA), for optimization process.

DA is based on a statistical physics interpretation of information theory ([14], [15]), see also [16] for an analysis of the relationship between statistical physics and information theory. DA introduces controlled randomization into the optimization process where the expected cost is minimized subject to a constraint on the level of randomness as measured by the Shannon entropy of the system. The resultant Lagrangian functional and parameter are analogous to the “free energy” and the “temperature”, respectively, of a corresponding physical system. The optimization is equivalent to an annealing process that starts by minimizing the cost (free energy) at a high temperature. This minimum is then tracked at successively lower temperatures (lower levels of entropy) as the system typically undergoes a sequence of phase transitions through which the complexity of the solution grows. As the temperature approaches zero, the original cost term dominates the free energy and hard (nonrandom) mappings are obtained.

This paper is organized as follows. In Section II we give the problem definition of WCE. In Section III, we review some of the prior results in the literature. Proposed method is described in Section IV and the experimental results are given in Section V. Discussion and concluding remarks are in Section VI.

II. PRELIMINARIES AND PROBLEM DEFINITION

Let \mathbb{R} , $\mathbb{P}\{\cdot\}$ and $\mathbb{E}\{\cdot\}$ denote the set of real numbers, probability and expectation operations, respectively. $\mathbb{E}\{\cdot|\cdot\}$ is the conditional expectation, $H(\cdot)$ and $H(\cdot|\cdot)$ are the entropy and conditional entropy. $\nabla_x f$ denote the partial derivative of f with respect to x . Upper case letters are used to denote random variables and lower case letters for their realizations.

The problem setting is given in Figure 1. The source X_0 and noise N are independent zero-mean Gaussians with variances σ_X^2 and 1, respectively. The two controllers $g : \mathbb{R} \rightarrow \mathbb{R}$ and




Fig. 1. The problem setting. Witsenhausen showed a nonlinear $g(\cdot)$ minimizes the cost $\mathbb{E}\{k^2 g^2(X_0) + X_2^2\}$.

$h : \mathbb{R} \rightarrow \mathbb{R}$ aim to minimize the cost

$$D = \mathbb{E}\{k^2 g^2(X_0) + X_2^2\} \quad (1)$$

where $X_2 = X_1 - h(X_1 + N)$ is estimation error and $X_1 = X_0 + g(X_0)$. The given constant k^2 is a weight to trade off the control cost $\mathbb{E}\{g^2(X_0)\}$ with the estimation cost $\mathbb{E}\{X_2^2\}$. We also define $f(x) = g(x) + x$.

Since the output X_2 is the estimation error of X_1 , the second controller that minimizes the mean squared error distortion is given in closed form as

$$h(Y) = \mathbb{E}\{X_1|Y\}. \quad (2)$$

where $Y = X_1 + N$. We refer to $g(\cdot)$ as the encoder and to $h(\cdot)$ as the decoder for obvious information theoretic reasons.

III. PRIOR RESULTS

Since the decoder is in closed form, research on this problem has been focused on finding the optimal encoder g . Witsenhausen has derived some properties of $g(\cdot)$, including the property of symmetry about the origin. For a given set of problem parameters ($k = 0.2$, $\sigma_X = 5$) Witsenhausen provided the following solution (that outperforms any affine solution): $f(x) = \sigma_X \text{sgn}(x)$ where $\text{sgn}(\cdot)$ is the signum function. This type of solution is referred to as a “1-step” solution, since the function consists of a single “step” in the positive half of real line (due to symmetry around the origin, positive half is enough to describe a given solution). Further improvements has been found by various researchers that utilize 2.5, 3, 3.5 and 4-step solutions (an $x.5$ step solution means there is a step that straddles the origin, thus half of it is on the positive portion of real line). Moreover, the latter solutions made improvements by using slightly sloped steps rather than constant ones. Some of the prior results appeared in the literature are given in Table I, including the best result to date.

TABLE I
MAJOR PRIOR RESULTS

Solution	Cost
Optimal Affine Solution	0.961852
1-step, Witsenhausen [1]	0.404253
2-step, [9]	0.190
Sloped 2.5 - step, [10]	0.1701
Sloped 3.5 - step, [11]	0.1673132
Sloped 3.5 - step, [12]	0.1670790
Sloped 4 - step, [13]	0.16692462

IV. PROPOSED METHOD

In this section we describe a method based on the concept of deterministic annealing. The cost is optimized by searching for the encoding function g within a parametric class of functions where each function is specified by (a) a partition of space and (b) a parametric local model for each partition. The crucial idea in DA is to introduce randomness into the optimization process, wherein the deterministic partition of space is replaced by a random partition, i.e. every input point is associated with each one of regions in probability. During the minimization of the cost, the Shannon entropy of the system is controlled, gradually lowered, and a deterministic partition (encoder) is obtained at the limit of zero entropy.

While DA is founded on information theoretic principles, it is motivated by statistical physics and the analogy is emphasized herein to provide further intuition. The entropy-constrained Lagrangian cost is in fact the free energy of a corresponding physical system, with the Lagrange parameter playing the role of “temperature” in the system. At high temperature, where entropy is maximum, the system effectively has a single local model that dominates the entire space. During the “annealing” process, where temperature is gradually lowered, the system goes through “phase transitions” which correspond to increase in the number of local models. At zero temperature, the Lagrangian reduces to the original cost function whose minimum is achieved by a deterministic encoder, thus at this stage the desired solution is obtained. In the rest of this section, we present detailed derivation of the proposed method.

The parametric encoder functions are specified by local models denoted as $g_m(x_0)$ and the space partitions \mathbb{R}_m , where $m = \{1, 2, \dots, M_{max}\}$, such that the encoder can be written as

$$g(x_0) = g_m(x_0) \quad \text{for } x_0 \in \mathbb{R}_m \quad (3)$$

While noting that local models can be in any parametric form, for the particular case of Witsenhausen counterexample we use affine local models given by

$$g_m(x_0) = a_m x_0 + b_m. \quad (4)$$

Remark 1: Piecewise affine mapping as we use in this model had been conjectured to be optimal¹ in [10], [11] based on the numerical results. However, recently it is shown that the optimal $f(\cdot)$ must have an analytic left inverse [7], disproving the conjecture since the left inverse of a piecewise affine function cannot be analytic. Nevertheless, it is understood that the minimal cost can be approached arbitrarily close by piecewise affine functions [7]. Hence, for numerical purposes, the piecewise assumption in our general model does not introduce any loss.

To derive a DA based approach, we introduce randomness into the optimization process by defining association probabilities

$$p(m|x_0) = \mathbb{P}\{g(x_0) = g_m(x_0)\} = \mathbb{P}\{x_0 \in \mathbb{R}_m\}. \quad (5)$$

¹The existence of an optimal solution was shown in [1].

for each x_0, m . Accordingly, the system has a Shannon entropy $H(X_0, M) = H(X_0) + H(M|X_0)$. Since the first term is a constant determined by source, we remove it for convenience and define

$$H \triangleq H(M|X_0) = \mathbb{E}\{\log p(M|X_0)\}. \quad (6)$$

where H measures the average level of uncertainty in the partition of space.

Remark 2: By employing the so-called mass constrained DA approach [17], we could equivalently minimize the mutual information $I(M; X_0)$ instead of maximizing the conditional entropy, obtaining direct relation with rate-distortion theory, see [14] for details.

In order to minimize the cost at a specified level of randomness, we define the Lagrangian

$$F = D - TH \quad (7)$$

referred to as (Helmholtz) “free energy” and the Lagrange parameter T called “temperature”, to emphasize the intuitively compelling analogy to statistical physics. The algorithm starts at high temperature, where minimization of (7) effectively maximizes the entropy. Accordingly, the association probabilities are uniform and all models are identical, or effectively, there is a single distinct local model. As the temperature is decreased, a bifurcation point is reached where the current solution is no longer a minimum, such that there exist a better solution with the local models divided into two or more groups. Intuitively, the current solution becomes a saddle point and a slight perturbation of local models will trigger the discovery of the new solution with increased number of effective local models. Such bifurcations are referred to as “phase transitions” and the corresponding temperatures are called “critical temperatures”. See [15] for phase transition analysis in the general DA setting. In order to trigger phase transitions, we always keep a duplicate for each local model and perturb them at each temperature. Until the critical temperature is reached, they will merge back during free energy optimization, but will split at phase transitions. At lower values of T , randomness is traded for reduction in D . In the limit $T = 0$, minimizing F corresponds to minimizing D directly, which produces a deterministic mapping. Therefore, in the practical algorithm we start at a high value of T and gradually lower it while minimizing F at each step.

A brief sketch of the algorithm can be given as follows.

- 1) Start at high temperature, single model.
- 2) Duplicate local models.
- 3) Minimization of F .
 - a) Optimize $p(m|x_0)$ for all m, x_0 .
 - b) Optimize $a_m(x_0)$ and $b_m(x_0)$, for all m , using gradient descent.
 - c) Optimize $h(\cdot)$.
 - d) Go to Step 4 if F has converged, go to (a) otherwise.
- 4) Stop if desired or lower temperature and go to Step 2.

The minimization of free energy in step 3 is done by iteratively optimizing the association probabilities, local model parameters and decoder until the decrease in F becomes insignificant.

We now derive the optimal association probabilities. Setting $\nabla_{p(m|x_0)} F = 0$, we have

$$D_m(x_0) + T \log p(m|x_0) + T = 0 \quad (8)$$

where

$$D_m(x_0) = \mathbb{E}\{k^2 g_m^2(x_0) + X_2^2\}. \quad (9)$$

From (8) we have

$$p(m|x_0) \propto e^{-[D_m(x_0) - T]/T}. \quad (10)$$

Finally we normalize to obtain

$$p(m|x_0) = \frac{e^{-[D_m(x_0)]/T}}{\sum_m e^{-[D_m(x_0)]/T}} \quad \forall x_0, m. \quad (11)$$

The optimal association probabilities take the form of the Gibbs distribution, which emerge at equilibrium in statistical physics. A fundamental principle of statistical physics is that the minimum of the free energy determines the distribution at thermal equilibrium, at which point the system is governed by the Gibbs distribution.

Some intuition to the annealing process can be obtained by observing how (11) evolves as the temperature is decreased. At the limit of high temperature, (11) yields uniform distributions and every input x is equally associated with all $g_m(\cdot)$. As T is lowered, the distributions become more discriminating. As $T \rightarrow 0$, we obtain *hard* associations, i.e. every x is fully associated with the local model that contributes least to the distortion (the one with minimum $D_m(x_0)$).

The local model parameters a_m, b_m can be optimized by any optimization method such as line search or gradient descent. The optimum decoder is given in (2).

V. EXPERIMENTAL RESULTS

We give the details of the annealing process with its phase transitions in Figure 2. At high temperature, there is only one local model, which corresponds to a 1-step solution. As the temperature is lowered, the system undergoes phase transitions that increase the number of local models. Each phase transition reveals a new step for the encoder. One can observe that the phase transitions effectively generate the entire class of n -step solutions - an important advantage of the proposed method. In order to generate a solution for a particular n , one needs to run the method until the desired number of steps (i.e. local models) is obtained, and then decrease the temperature without growing the model size.

In this work we calculated the solutions up to $n = 5$, whereas noting that more steps possibly exist. The cost obtained for 3, 4 and 5-step solutions are given in Table II. In addition to the improvement over the previously reported costs, our algorithm revealed a fifth step in the solution. Although the improvement in the cost is very small with this new step, we note the theoretical importance pertaining to the structure

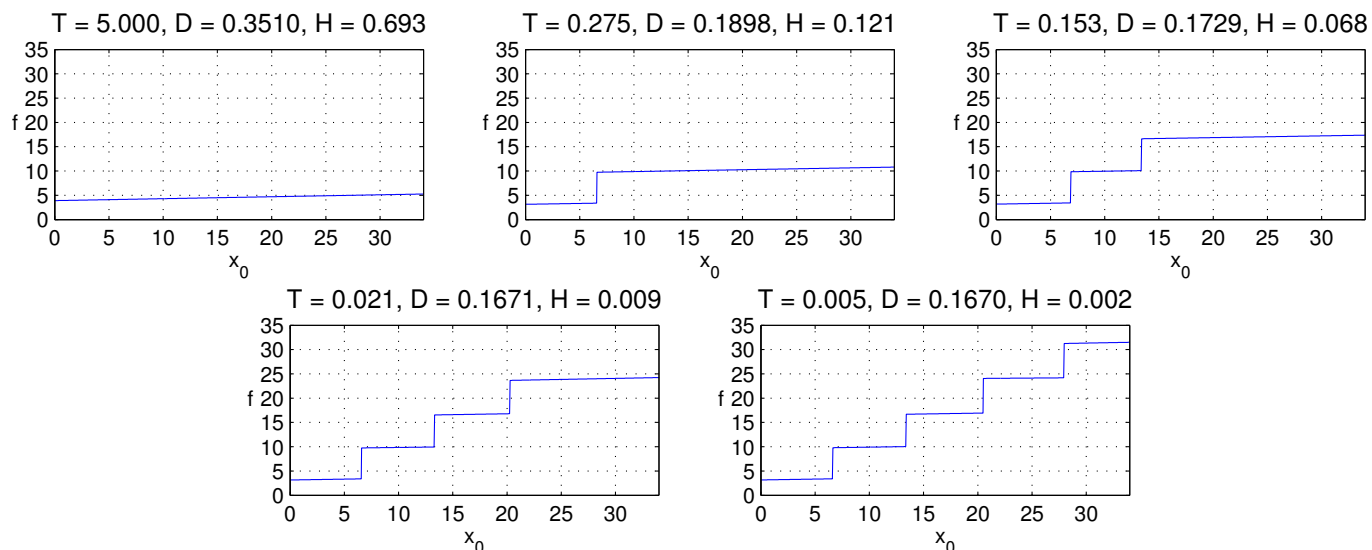


Fig. 2. Encoder during various phases of the annealing process.

of the solution. The 5-step and 4-step encoders are given in Figure 3.

TABLE II
OUR RESULTS

Solution	Cost
3 - step	0.16694471
4 - step	0.16692319
5 - step	0.16692291




Fig. 3. 5-step solution.

We also compare our results to the best previous result in [13]. The difference between the two encoders is plotted in Figure 5. There are three main differences: i) Our best solution has 5 steps. ii) Each step is exactly linear as illustrated in Figure 4. iii) The step boundaries differ slightly as illustrated in Figure 5.

There are several advantages of the method proposed here. i) We perform optimization process in the original, analog domain, without discretization. This approach yields the linearity




Fig. 4. Comparison view of the third step in our 4-step solution (straight line) and the 4-step solution in [13].

of the *steps* as illustrated in Figure 4. ii) We employ a powerful non-convex optimization tool, DA. [13] uses “noisy channel relaxation” [18] (NCR) originally developed for vector quantizer design purposes. NCR offers improvement over greedy techniques thanks to the ad hoc relaxation it employs to avoid local optima, however, it has been outperformed by DA which is derived from basic principles, in a variety of optimization settings, including most relevantly in the optimization of zero-delay source-channel mappings [19].

Remark 3: The costs here are calculated according to (1). For integrations, we used the same numerical methods as described in [13].

Note: Matlab code for our calculations of the total cost, including our decision functions can be found in [20].

VI. DISCUSSION

In this paper, we proposed an optimization method for distributed control problems, whose solutions are known to be non-linear. As an example we showed the effectiveness of the




Fig. 5. The difference between the 5-step solution proposed in this paper ($g(x)$) and the 4-step solution in [13] ($g^*(x)$).

proposed method on the very well-known benchmark problem known as the Witsenhausen's counter-example. We note again that although our numerical results pertain to a particular problem, our approach is general and applicable to any problem of this nature (see e.g. [21] for some variations of Witsenhausen's counter-example). For example, when controllers have side information correlated with the source -i.e., when the problem setting resembles the classical Wyner-Ziv like problems in information theory- the cost surface becomes riddled with locally optimum points, see e.g. [22] for a control theoretic analysis of such a setting. The greedy competitor approaches get trapped in a local optimum while deterministic annealing mitigates this obstacle in this setting, our preliminary analysis can be found in [23].

ACKNOWLEDGMENTS

This work is supported by the NSF under grants CCF-1118075 and CCF-1016861.

REFERENCES

- [1] H.S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147, 1968.
- [2] S. Yüksel and T. Başar, *Stochastic Networked Control Systems: Stabilization and Optimization under Information Constraints*, Springer, 2013.
- [3] R. Bansal and T. Başar, "Stochastic teams with nonclassical information revisited: When is an affine law optimal?," *Automatic Control, IEEE Transactions on*, vol. 32, no. 6, pp. 554–559, 1987.
- [4] P. Grover, Se Yong Park, and A. Sahai, "Approximately optimal solutions to the finite-dimensional Witsenhausen counterexample," *Automatic Control, IEEE Transactions on*, vol. 58, no. 9, pp. 2189–2204, 2013.
- [5] C. Choudhuri and U. Mitra, "On Witsenhausen's counterexample: The asymptotic vector case," in *Information Theory Workshop (ITW), 2012 IEEE*, 2012, pp. 162–166.
- [6] W. McEneaney, S. Han, and A. Liu, "An optimization approach to the Witsenhausen counterexample," in *IEEE Conf. on Dec. and Control*, IEEE, 2011, pp. 5023–5028.
- [7] Y. Wu and S. Verdú, "Witsenhausen's counterexample: A view from optimal transport theory," in *IEEE Conf. on Dec. and Control*, 2011.
- [8] A. Kulkarni and T. Coleman, "An optimizer's approach to stochastic control problems with nonclassical information structures," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, IEEE, 2012, pp. 154–159.
- [9] M. Deng and Y. Ho, "An ordinal optimization approach to optimal control problems," *Automatica*, vol. 35, no. 2, pp. 331 – 338, 1999.
- [10] M. Baglietto, T. Parisini, and R. Zoppoli, "Numerical solutions to the witsenhausen counterexample by approximating networks," *Automatic Control, IEEE Transactions on*, vol. 46, no. 9, pp. 1471–1477, 2001.
- [11] J.T. Lee, E. Lau, and Y.-C. Ho, "The Witsenhausen counterexample: a hierarchical search approach for nonconvex optimization problems," *Automatic Control, IEEE Transactions on*, vol. 46, no. 3, pp. 382–397, 2001.
- [12] N. Li, J. Marden, and J. Shamma, "Learning approaches to the Witsenhausen counterexample from a view of potential games," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, IEEE, 2009, pp. 157–162.
- [13] J. Karlsson, A. Gattami, T. Oechtering, and M. Skoglund, "Iterative source-channel coding approach to witsenhausen's counterexample," in *American Control Conference (ACC), 2011*, IEEE, 2011, pp. 5348–5353.
- [14] K. Rose, "A mapping approach to rate-distortion computation and analysis," *IEEE Transactions on Information Theory*, vol. 40, no. 6, pp. 1939–1952, 1994.
- [15] K. Rose, "Deterministic annealing for clustering, compression, classification, regression, and related optimization problems," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2210–2239, 1998.
- [16] N. Merhav, "Statistical physics and information theory," *Foundations and Trends in Communications and Information Theory*, vol. 6, no. 1–2, pp. 1–212, 2011.
- [17] K. Rose, E. Gurewitz, and G. Fox, "Constrained clustering as an optimization method," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 8, pp. 785–794, 1993.
- [18] S. Gadkari and K. Rose, "Robust vector quantizer design by noisy channel relaxation," *IEEE Transactions on Communications*, vol. 47, no. 8, pp. 1113–1116, 1999.
- [19] M. S. Mehmetoglu, E. Akyol, and K. Rose, "A deterministic annealing approach to optimization of zero-delay source-channel codes," in *Proc. of IEEE Inf. Theory Workshop*, 2013.
- [20] <http://www.scl.ece.ucsb.edu/html/witsen.html>.
- [21] T. Başar, "Variations on the theme of the Witsenhausen counterexample," in *47th IEEE Conference on Decision and Control Proceedings (CDC)*, IEEE, 2008, pp. 1614–1619.
- [22] N.C. Martins, "Witsenhausen's counter example holds in the presence of side information," in *Decision and Control, 2006 45th IEEE Conference on*, 2006, pp. 1111–1116.
- [23] M. Mehmetoglu, E. Akyol, and K. Rose, "A deterministic annealing optimization approach for Witsenhausen's and related decentralized control settings," *arXiv preprint arXiv:1403.5315*, 2014.