

Demystifying the Witsenhausen Counterexample

Q. In my graduate-level controls course, the instructor mentioned the “Witsenhausen counterexample,” which seems to have something to do with decentralized control and communication. But I must confess that I did not understand what was going on or why it was important. Could someone at *IEEE Control Systems Magazine* provide a clear and simple explanation? Thanks.

Pulkit and Anant: Witsenhausen’s counterexample to a conjecture in decentralized control theory is a deceptively simple decentralized stochastic control problem that has remained unsolved for the last four decades [1] (see “Hans S. Witsenhausen”). This problem, shown in Figure 1, has a scalar state that evolves in discrete time according to linear dynamics. The state is constrained by two controllers, each of which acts only once. The “weak” controller C_w operates first after seeing the initial plant state X_0 . The resulting control signal U_w modifies the state X_0 by means of $X_1 = X_0 + U_w$. The “blurry” controller C_b acts next after having observed a noisy version $X_1 + Z$ of the state X_1 . The second control signal U_b modifies the state X_1 by means of $X_2 = X_1 - U_b$. Both the initial state X_0 and the observation noise Z are independent zero-mean Gaussian random variables with variances σ_0^2 and one, respectively. The goal is to design control laws that minimize the expected quadratic cost $E[k^2 U_w^2 + X_2^2]$, where k^2 is a weight that trades off control costs with terminal state costs.

For this problem, the optimal strategy for the second controller C_b

is to optimize set U_b to be the minimum mean squared error estimate—the conditional expectation—of X_1 given its noisy observation $X_1 + Z$. This choice minimizes the second term $E[X_2^2]$ in the cost. We thus need to find only the first controller C_w since that is what determines the distribution of X_1 . If we were to constrain ourselves to a linear controller C_w , then X_1 would be Gaussian. In that case, the optimal choice for C_b would also be linear. But surprisingly, in his 1968 paper [1], Witsenhausen provides a nonlinear strategy for the problem that outperforms all linear strategies.

In Witsenhausen’s strategy, the weak controller quantizes the initial state X_0 by forcing X_1 to be the closer of either $+\sigma_0$ and $-\sigma_0$. This “discreteness” of X_1 greatly reduces the second stage cost $E[X_2^2]$. Furthermore, general-

izing from two quantization points to an infinite number of regularly spaced quantization points, the resulting nonlinear control strategies outperform optimal linear strategies by arbitrarily large factors [2]. These nonlinear strategies do significantly better if two conditions are met, a) the observation noise variance is small compared to the initial system state variance, and b) “an ounce of prevention is worth a pound of cure,” that is, the weight parameter k^2 on the weak controller’s cost is much smaller than one. Such small weights on the weak controller’s input U_w^2 allow for investing in a moderately large input U_w to greatly reduce the estimation error X_2^2 .

At first glance, all this might seem to be an isolated curiosity at best. Why is it important? It turns out that the Witsenhausen counterexample is notable because of its links with computational

Hans S. Witsenhausen

The following biography of Hans S. Witsenhausen appeared in *IEEE Transactions on Automatic Control* in 1968.

Hans S. Witsenhausen (M66) was born in Frankfurt/Main, Germany, on May 6 1930. He received the I.C.M.E. degree in electrical engineering in 1953 and the degree of Licencie en Sciences in mathematical physics in 1956, both from the Université Libre de Bruxelles, Brussels, Belgium. He received the S.M. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1964 and 1966, respectively. From 1957 to 1959 he was engaged in problem analysis and programming at the European Computation Center, Brussels. From 1960 to 1963 he was a senior engineer at the Research and Computation Division of Electronic Associates, Inc., Princeton, N.J., where he worked on analog and hybrid computer techniques and on systems analysis problems. From 1963 to 1965 he was associated with the Electronic Systems Laboratory and the Lincoln Laboratory at M.I.T. During 1965–1966 he was a fellow of the Fannie and John Hertz Foundation. He joined Bell Telephone Laboratories, Inc., Murray Hill, N.J., in 1966, engaged in the study of control in the presence of uncertainty. Dr. Witsenhausen is a member of Sigma Xi and the International Association for Analog Computation.

complexity and information theory. However, to appreciate these links as well as the problem, it is important to first understand the strategic place of Witsenhausen's counterexample within the general research program for decentralized control.

Decentralized control systems, in contrast to centralized control systems, have multiple controllers that are collaboratively trying to control a system by taking actions based on their individual observations. The observations of one controller may not be available to the other controllers—that is what makes a control problem decentralized. Consider the decentralized control system shown in Figure 2. This system could be, for example, a set of robots assembling a car, a group of UAVs flying over hostile territory, or a set of sensors and actuators controlling the temperature throughout a building. The designer must craft an ensemble of control laws for the task at hand by mathematically modeling the dynamics and capturing the task objective in a cost function.

So what could possibly go wrong? Well, evaluating the cost function itself can sometimes be fundamentally uncomputable in the sense of Turing's halting problem or Gödel's undecidability. This uncomputability can *actually happen* for nonlinear hybrid control systems! Putting such pathological cases aside, even for finite-horizon discrete-time systems the challenge is one of practical uncomputability or complexity. Brute-force search is impossible over the uncountable space of control laws, and discretization of the search space may lead to exponential-time (or worse!) algorithms.

The only hope is to exploit problem structures that we can understand. One such structure is that of linear-quadratic-Gaussian (LQG) problems with linear dynamics, quadratic cost functions, and Gaussian distributions for initial conditions and noises. The fact that the sum of jointly Gaussian random variables is Gaussian turns out to imply that dynamic programming is easy when the LQG control




FIGURE 1 Witsenhausen's counterexample. The state evolves in discrete-time $X_0 \rightarrow X_1 \rightarrow X_2$. The objective in this problem (a) is to minimize the weighted sum of the control cost $\mathbb{E}[U_w^2]$ at the first stage and the error $\mathbb{E}[X_2^2] = \mathbb{E}[(X_1 - U_b)^2]$ in estimating X_1 at the second stage. (b) is a cartoon depiction of the problem. A hypothetical discrete-time inverted pendulum is being controlled by a "weak" controller C_w , which operates first and has perfect observations, but a costly control signal, and a "blurry" controller C_b , which operates second and has noisy observations, but has zero weighting on its control signal. The goal is to maneuver the system state close to zero in two time steps while keeping the control cost low.

problem is centralized. The optimal centralized control law is linear in the observations and can be found efficiently using Riccati equations. This success leads to a natural conjecture that linear control laws remain optimal even in decentralized settings. This is the conjecture to which Witsenhausen provided a counterexample.

DECENTRALIZED CONTROL CAN BE HARD

Unfortunately, the problem of finding an optimal control law for the Witsenhausen counterexample is nonconvex, and hence convex-optimization-based

search techniques do not help in finding an optimal solution. An optimal law does exist, as shown by Witsenhausen [1, Sect. 3]. But Witsenhausen's proof is nonconstructive, and the general difficulty in finding an optimal controller can be better appreciated through the problem's links with computational complexity theory.

It is shown in [3] that obtaining optimal control laws for a discrete version of the Witsenhausen counterexample is NP-complete. The class of NP-complete problems includes the traveling salesman problem and the graph coloring problem. Moreover,

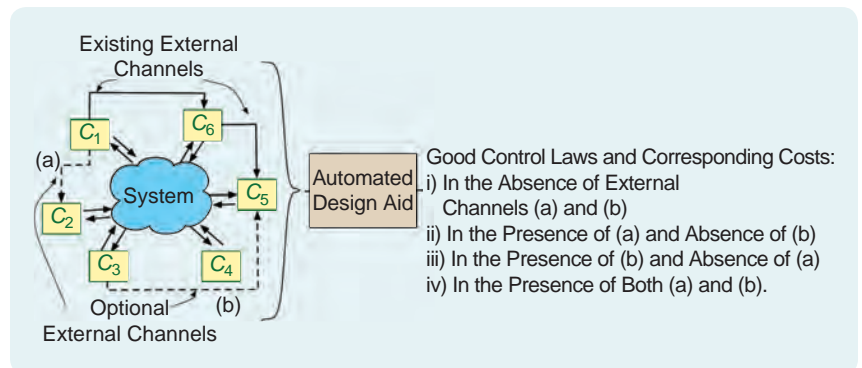


FIGURE 2 A decentralized control system. Each controller takes actions based on its own partial observations of the system. External channels may be available to allow communication between the controllers, but these channels are added by the designer at additional cost.

these problems are known to be equally hard in the sense that obtaining an efficient polynomial-time solution to one will yield efficient solutions for all. It is thus widely accepted that there is no hope of obtaining efficient solutions for any of these problems. The hardness result for Witsenhausen's counterexample thus shatters the prospects for a general way to find *optimal* decentralized control laws for more realistic systems.

To make progress, we thus relax the demand for optimality by accepting suboptimal strategies that are "good enough." Because of its difficulty, Witsenhausen's counterexample has been a popular benchmark for testing techniques that obtain "good enough" solutions to hard nonconvex decentralized control problems. The research community uses it to explore heuristics, such as those based on neural networks [4], hierarchical search [5], and game theory [6].

HOW CAN WE DEEM A SOLUTION TO BE "GOOD ENOUGH?"

While relaxing the requirement of optimality, if we lower our standards

too much, we run the risk of obtaining strategies that are locally optimal but significantly worse than the globally optimal ones. While it is generally felt that solutions obtained through the heuristic search methods in [4]–[6] are probably close to the global optimum, and subsequent information-theoretic investigations provide further indirect evidence that supports this belief [7, Sect. 4.3], these heuristics search within only a subset of control laws. Theoretical "goodness" cannot be ensured without a *guaranteed gap from optimality*. Unfortunately, the approaches in [4]–[6] do not provide such guarantees.

The study of guaranteed approximately optimal strategies is a central topic within computational complexity theory. In addition, guarantees on approximation quality have been obtained [8]–[10] in information theory for longstanding communication problems of roughly the same vintage as the Witsenhausen counterexample. These bounds are obtained using insights from a deterministic channel model that further abstracts the Gaussian noise model [9], [10]. As an

illustration, one such deterministic model is shown in Figure 3(a) for the multiple-access channel, which is a traditional abstraction for the uplink from mobiles to the base-station in cellular systems. In these deterministic models, binary expansions are used to justify the abstraction of a real-valued signal as an ordered list of bits, where the most-significant bits correspond to the strong parts of the signal at higher signal-to-noise ratios. The least-significant bits are affected by noise and are therefore corrupted at the receiver. If two transmissions occur simultaneously, all the bits representing the weaker signal, as well as the lower order bits of the strong signal, are mangled in the collision. This deterministic simplification offers insights into the nature of the contention for the underlying communication resources and thereby facilitates the design of interference-mitigation strategies for wireless networks. The resulting strategies can be shown to deliver data rates within a constant number of bits of the channel capacity uniformly over all problem parameters [8]–[10]. Finding the exact channel capacity for these multiuser communication problems has proved to be an extremely hard unsolved problem, much like decentralized control problems.

UNDERSTANDING IMPLICIT COMMUNICATION TO GUARANTEE A GAP FROM OPTIMALITY

As shown in [7] and [11], a parallel information-theoretic analysis helps find strategies that attain within a guaranteed factor of the optimal cost for the Witsenhausen counterexample, as well as for many other decentralized-control problems in the counterexample's neighborhood. The relevant gap here is multiplicative since the optimal costs themselves can decrease to zero as $k \rightarrow 0$. That the analysis required is information theoretic may seem surprising since nothing in the Figure 1 is




FIGURE 3 Deterministic models. (a) is a deterministic model for a multiple-access channel, where two users are trying to send their messages to a single receiver. The bits b_i and b'_i denote the binary representations of the transmitted signals of the two users. The observation noise corrupts the bits at and below the noise level. (b) is the deterministic model for Witsenhausen's counterexample. The initial state $X_0 = (b_1, b_2, b_3, b_4, b_5)$ is known at the weak controller, and thus the control input $U_w = (b'_1, b'_2, b'_3)$ can depend on X_0 . Choosing $b'_1 = b_3, b'_2 = b_4,$ and $b'_3 = b_5$, the weak controller can force the lower order bits of state X_1 to zero, effectively making X_1 a coarse quantization of X_0 . Since the blurry controller knows that the lower order bits of X_1 are zero, the observation noise does not corrupt its estimate of those bits.

labeled as a channel, and there is thus no designated means of information transfer between the controllers. Moreover, there is no apparent message to be communicated!

Is information theory just delivering some specific mathematical inequalities or is it providing generally useful intuition? To answer this question, we need to think more deeply about the Witsenhausen counterexample. It turns out that the problem allows a seemingly Rube-Goldberg-esque possibility, namely, the weak controller can communicate to the blurry controller through the *implicit* channel provided by the modifiable system state itself. The word *implicit* is used because the plant's state is presumably not explicitly designed to be a communication pathway. This sort of implicit communication is called "signaling" in the control and economics literature and, as a possibility, it is ubiquitous in decentralized control systems. For example, a UAV could signal to others following it by wiggling its motion and honeybees actually do similar signaling to communicate at the hive! But the bees, unlike the weak controller in the counterexample, presumably have identifiable messages to communicate, namely, the locations of food. This lack of an externally identifiable *message* makes the implicit communication in the Witsenhausen counterexample even more intriguing.

Taking this implicit communication perspective, we arrive at the information-theoretic interpretation of the counterexample illustrated in Figure 4. In this view, the quantization in the strategies of [1, 2] allows the controller to more reliably *communicate* the now discrete state X_1 through the noisy implicit communication channel to the blurry controller. The counterexample requires smudging the traditionally sharp distinction between the message and the messenger, both of which are the state X_1 . Even so, the multipoint quantization strategies of




FIGURE 4 An information-theoretic interpretation of Witsenhausen's counterexample. The weak controller C_w is interpreted as an encoder operating under an average power constraint P . The blurry controller is interpreted as a decoder attempting to minimize the MMSE error $\mathbb{E}[(X_1 - U_b)^2]$, where U_b is the MMSE-estimate of X_1 . The problem of minimizing $k^2P + \text{MMSE}$ is equivalent to Witsenhausen's counterexample. This interpretation highlights the role of implicit communication in the counterexample, where the weak controller tries to communicate the state X_1 to the blurry controller while using as little power as possible.

[2] are easily understood using the deterministic model shown in Figure 3(b). As shown in [7] and [11], multipoint quantization strategies, complemented by linear strategies, always attain a cost within a factor of eight of optimal for all problem parameters [11], providing a performance guarantee for the Witsenhausen counterexample.

ARE LINEAR STRATEGIES GOOD ENOUGH WITH THE RIGHT ARCHITECTURAL CHANGES?

In the real world, the control designer has architectural choices too. It is tempting to think that perhaps the complexity here is a Gordian knot—an artificial consequence of a bad control system architecture. Suppose the designer could connect the weak controller to the blurry one using a *perfect* and *instantaneous* explicit communication link. The weak controller could then send X_0 over this explicit link, and the linear strategy $U_w = 0$, $U_b = -X_0$ would be trivially optimal because the attained cost would be zero. This general principle holds whenever there is no temptation to signal [12]. However, perfect and instantaneous communication channels do not exist. With realistic imperfect external channels, sometimes even

nonlinear strategies that ignore the external channel can still outperform the best linear strategies [13].

The deterministic model for Witsenhausen's counterexample gives us an intuitively compelling understanding of the counterexample that explains why linear strategies are not the right thing to do even when an external channel connects the two controllers. A signal over the external channel that is linear in the observations of the weak controller predominantly communicates the most-significant bits of the system state. But these are precisely the bits that are most robust to observation noise and are therefore known reliably at the second controller even without using the external channel. It is thus wasteful to use linear strategies over an imperfect external channel. Rather, it is much smarter to use such external channels in a nonlinear way, namely, by communicating the *most-important of the state's lower order bits* that are battered by noise in the implicit channel.

One moral of the story so far is that familiarity with linear strategies cannot act as the sole justification for their use. The fields of communication and circuits have both embraced nonlinear strategies despite initial reluctance. Classical control of linear systems, on the other hand, relies steadfastly on linear strategies. Witsenhausen's counterexample unequivocally argues against this reliance, since it shows that decentralized control of linear systems needs an understanding of nonlinear strategies. Furthermore, the nonlinearity of provably good solutions, based on quantization, as well as believed-to-be optimal solutions (shown in Figure 5) is structured. The fact that these solutions are structured suggests that more general decentralized control problems might be tractable as well.

IN SUMMARY

Witsenhausen's counterexample emerges as a minimalist toy problem occupying a key strategic position because it exposes links between




FIGURE 5 Approximately optimal quantization-based nonlinear strategies for the Witsenhausen counterexample [11]. “Slopy” quantization strategies are believed to be almost optimal (albeit without proof!). The simplicity of these strategies offers hope that good practical nonlinear strategies can be found for more complex problems.

decentralized control, information theory, and computational complexity. The need for nonlinear control strategies for the counterexample is a symptom of the possibility of implicit signaling through the plant among controllers in decentralized systems. But while understanding the role of implicit signaling may be useful for addressing larger systems, it is but one key to the puzzle of decentralized control.

AUTHOR INFORMATION

Anant Sahai received the B.S. in 1994, from the University of California, Berkeley, and the S.M. and Ph.D. from the Massachusetts Institute of Technology in 1996 and 2001, respectively. He is an associate professor in the Department of Electrical Engineering and Computer Sciences at

the University of California, Berkeley, where he joined as an assistant professor in 2002. Prior to that, he spent a year as a research scientist at the wireless startup Enuvis in South San Francisco, developing software-radio signal-processing algorithms to enable very sensitive GPS receivers for indoor operation. From 2007 to 2009, he was the treasurer for the IEEE Information Theory Society. His current research interests are at the intersection of information theory and decentralized control, as well as in wireless communication, particularly green radios and dynamic spectrum sharing.

Pulkit Grover received the B.S. in 2003 and the M.S. in 2005 from IIT Kanpur. He received the Ph.D. in 2010 in the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley, where he is currently a postdoc. His research interests include information theory in decentralized control as well as low-power circuits for wireless communications and bioelectronics.

REFERENCES

- [1] H. S. Witsenhausen, “A counterexample in stochastic optimum control,” *SIAM J. Contr.*, vol. 6, no. 1, pp. 131–147, Jan. 1968.
- [2] S. K. Mitter and A. Sahai, “Information and control: Witsenhausen revisited,” in *Learning, Control and Hybrid Systems: Lecture Notes in Control and Information Sciences 241*, Y. Yamamoto and S. Hara, Eds. New York: Springer-Verlag, 1999, pp. 281–293.
- [3] C. H. Papadimitriou and J. N. Tsitsiklis, “Intractable problems in control theory,” *SIAM J. Contr. Optim.*, vol. 24, no. 4, pp. 639–654, 1986.
- [4] M. Baglietto, T. Parisini, and R. Zoppi, “Numerical solutions to the Witsenhausen counterexample by approximating networks,” *IEEE Trans. Automat. Contr.*, vol. 46, no. 9, pp. 1471–1577, 2001.
- [5] J. T. Lee, E. Lau, and Y.-C. Ho, “The Witsenhausen counterexample: A hierarchical search approach for nonconvex optimization problems,” *IEEE Trans. Automat. Contr.*, vol. 46, no. 3, pp. 382–397, 2001.
- [6] N. Li, J. R. Marden, and J. S. Shamma, “Learning approaches to the Witsenhausen counterexample from a view of potential games,” in *Proc. 48th IEEE Conf. Decision and Control*, Shanghai, China, 2009, pp. 157–162.
- [7] P. Grover and A. Sahai, “Vector Witsenhausen counterexample as assisted interference suppression,” *Int. J. Syst., Control Commun. (Special Issue on Information Processing and Decision Making in Distributed Control Systems)*, vol. 2, pp. 197–237, 2010.
- [8] R. Etkin, D. Tse, and H. Wang, “Gaussian interference channel capacity to within one bit,” *IEEE Trans. Inform. Theory*, vol. 54, no. 12, pp. 5534–5562, Dec. 2008.
- [9] A. S. Avestimehr, S. Diggavi, and D. Tse. (2009, July). Wireless network information flow: A deterministic approach. [Online]. Available: <http://arxiv.org/abs/0906.5394>
- [10] A. S. Avestimehr, “Wireless network information flow: A deterministic approach,” Ph.D. dissertation, Univ. California, Berkeley, CA, 2008.
- [11] P. Grover, S. Y. Park, and A. Sahai. (2010). The finite-dimensional Witsenhausen counterexample. [Online]. Available: <http://arxiv.org/abs/1003.0514>
- [12] M. Rotkowitz, “On information structures, convexity, and linear optimality,” in *Proc. 47th IEEE Conf. Decision and Control*, Cancun, Mexico, 2008, pp. 1642–1647.
- [13] N. C. Martins, “Witsenhausen’s counter example holds in the presence of side information,” in *Proc. 45th IEEE Conf. Decision and Control*, San Diego, 2006, pp. 1111–1116.



QED

He bought a book on logic and studied the science of explanations, how to analyze the absolutely true and the relatively true, the proximate causes and the remote causes, how to untangle fallacies and take them apart piece by piece and show mistakes in reasoning. He heard the word “demonstrate” and said to himself: “What do I do when I demonstrate, more than when I reason or prove? How does demonstration differ from other proof?” He looked in Noah Webster’s dictionary and learned that demonstration is “proof beyond the possibility of doubt.”

—C. Sandburg, *Abraham Lincoln, The Prairie Years – 1*, Charles Scribner’s Sons, New York, 1926, pp. 472–473.