Implicit coordination in two-agent team problems with continuous action sets. Application to the Witsenhausen cost function

Achal Agrawal, Florian Danard, Benjamin Larrousse, Samson Lasaulce

► To cite this version:

Achal Agrawal, Florian Danard, Benjamin Larrousse, Samson Lasaulce. Implicit coordination in two-agent team problems with continuous action sets. Application to the Witsenhausen cost function. 2015 European Control Conference (ECC), Jul 2016, Linz, Austria. 10.1109/ECC.2015.7330808 . hal-01272497

HAL Id: hal-01272497

<https://hal.science/hal-01272497>

Submitted on 15 Feb 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Implicit Coordination in Two-Agent Team Problems with Continuous Action Sets. Application to the Witsenhausen Cost Function.

Achal Agrawal, Florian Danard, Benjamin Larrousse, and Samson Lasaulce
L2S (CNRS – Supélec– Univ. Paris Sud 11), France
{agrawal, danard, larrousse, lasaulce}@lss.supelec.fr

Abstract—The main contribution of this paper is the characterization of the limiting average performance of a system involving two agents who coordinate their actions which belong to a *continuous* set. One agent has complete and noncausal knowledge of the sequence of i.i.d. realizations of a random state X_0 , which is called the system state and affects the common team payoff. For the other agent, two scenarios in terms of observation assumptions are considered: in the first scenario, the other agent has a strictly causal knowledge of the system state while in the second scenario it has no direct knowledge of the state at all. However, in both scenarios, the less informed agent can always observe a noisy and strictly causal version of the actions taken by the (most) informed agent. There exists no dedicated communication channel between the two agents, and thus, the informed agent can only communicate via its actions which in turn affect the common payoff, hence the term implicit communication. Thus, there is a tradeoff to be found for the informed agent between communicating information about the incoming realizations of the system state and maximizing the payoff at the current stage. We use this general framework and apply it to a specific cost function, namely the Witsenhausen cost function. Although the problem tackled differs from the famous Witsenhausen’s counterexample, the authors believe interesting new connections which help to understand the corresponding open problem might be established over time. A numerical analysis is conducted to assess the Witsenhausen’s cost for two sub-optimal classes of strategies.

I. INTRODUCTION

In this paper, we consider a team with two agents who are trying to maximize their common payoff over a long time period, i.e., composed of many time-slots. At every time-slot or stage $i \in \{1, \dots, T\}$, Agent k , $k \in \{1, 2\}$ chooses its action $x_k \in \mathcal{X}_k$, where $\mathcal{X}_k$ is a continuous set. The instantaneous payoff function $w(x_0, x_1, x_2)$ depends on the realization of the random variable X_0 with realizations $x_0 \in \mathcal{X}_0$. The set $\mathcal{X}_0$ is also continuous and realizations of X_0 are assumed to be i.i.d.. A problem with the same information structure was addressed for the first time in [1]. Therein, the assumptions made are as follows: it is assumed that at any time Agent 1 knows the past, current, and future realizations of X_0 perfectly, whereas Agent 2 only observes the actions of Agent 1 in a strictly causal manner. Such a scenario has been extended in a couple of papers cited further. However, all of them treated the case where the action sets as well as the system state set were discrete and finite, i.e., $\forall k \in \{0, 1, 2\}, |\mathcal{X}_k| < \infty$. Reference [1] treated the case of an information structure in which Agent 2 has perfect observation and showed that the average

performance characterization is equivalent to finding the appropriate information constraint. The information constraint basically gives us the limit for the communication that can take place via actions in such a scenario. In [2] this result was generalized to the case where Agent 2 has imperfect observations. While all these contributions assume a strictly causal knowledge of the system state x_0 at Agent 2, the case where this assumption is relaxed was first presented in [3] and treated rigorously in [4]. The main contribution of the present paper is to generalize this approach of finding the limiting performance to the case of continuous action sets, which is an important case for control problems as many designs involve continuous controllers. A second contribution is to consider the Witsenhausen cost function [5] as a common cost function to be minimized by the 2-agent team under the two mentioned scenarios in terms of information structures; this establishes for the first time a link between [1][2][4] and [5]. Indeed, although the Witsenhausen problem can be seen as a one-shot coordination problem, whereas we consider a long-term coordination problem here, the idea of joint control-communication strategies is present in both formulations. As it will be seen, characterizing the feasible performance of the long-term coordination problems amounts to determining a certain information constraint. Although information constraints normally appear when (large) sequences intervene, it has to be noted that one-shot problems closely related to the Witsenhausen problem have been solved by introducing an information constraint; this is the case, for instance, for the Gaussian Test Channel (GTC) [6]. The approach we adopt has connections with that of [7] where probability distributions which minimize the cost function are used. However, in the latter the authors restrict their attention to what modifications render the (one-shot) Witsenhausen problem simpler to solve, and do not tackle the general framework of long-term implicit communication.

The paper is structured as follows. Section II provides the proposed problem formulation. It explains that characterizing the feasible set of expected common payoffs amounts to characterizing implementable joint probability distributions. Section III provides, for the two information structures considered, the two information constraints which allows one to characterize the implementable distributions. The Gaussian case is provided as a special instance, which establishes a connection with the dirty-paper coding problem [8]. In Section IV-A, we discuss the Witsenhausen cost function in

context to our problem. Section IV-B describes the numerical analysis for the special instance of payoff function (which equals minus the Witsenhausen cost function) and provide numerical results. Section V concludes the paper.

II. PROBLEM STATEMENT

Consider two agents, Agent 1 and Agent 2, who want to coordinate through their actions $x_1 \in \mathcal{X}_1$ and $x_2 \in \mathcal{X}_2$. The problem is said to be distributed in the sense that each agent can only control one variable of their common payoff function $w(x_0, x_1, x_2)$. The action set for both agents $\mathcal{X}_1, \mathcal{X}_2$ as well as the set of system states $\mathcal{X}_0$ are continuous sets. The realizations of the system state are assumed to be i.i.d. and generated from a random variable X_0 whose probability density function is denoted by $f_0(x_0)$. We shall use the notation $f_V(v)$ or $f(v)$ to refer to the probability density function of the generic continuous random variable V . The control strategies of Agents 1 and 2 are sequences of functions which are respectively defined by:

$$\begin{cases} u_i & : & \mathcal{X}_0^T & \rightarrow & \mathcal{X}_1 \\ v_i^a & : & \mathcal{X}_0^{i-1} \times \mathcal{Y}^{i-1} & \rightarrow & \mathcal{X}_2 \\ v_i^b & : & \mathcal{Y}^{i-1} & \rightarrow & \mathcal{X}_2 \end{cases} \quad (1)$$

where $T \geq 1$ is the total number of stages over which the agents are assumed to interact, $\mathcal{Y}$ is the observation set of Agent 2 and the superscripts a or b correspond to the two considered scenarios in terms of observation structure. The control strategy for Agent 1 u_i basically means that it knows the realizations of the system state for all T beforehand, and uses that information to choose its actions; note that the methodology used in this paper can also be exploited under less restrictive knowledge assumptions at Agent 1. The merit of the assumptions made for Agent 1 is that it allows one to make progress in the direction of quantifying the relationship between agents' observation capabilities and reachable performance, which is not well understood. Additionally, there already exist applications for which it is relevant: coordination between robots when a leader knows the trajectory in advance; distributed power control in wireless networks; robust image watermarking. For Agent 2, we consider two different control strategies v_i^a and v_i^b . The control strategy of *scenario a* assumes that Agent 2 observes the past realizations of the system state $x_0(1), \dots, x_0(i-1)$ as well as $y(1), \dots, y(i-1)$. The control strategy of *scenario b* is only based on the latter sequence and seems to be more in line with a possible information structure of a long-term version of the Witsenhausen problem. In any case, it is less demanding in terms of information assumptions. The observations $y(1), \dots, y(T)$ are assumed to be generated by a memoryless channel whose transition probability is denoted by γ and verifies a Markov condition $f_{Y|X_0, X_1, X_2}(y|x_0, x_1, x_2) = \gamma(y|x_1)$. The additive white Gaussian noise channel $Y = X_1 + Z$ is an intensively used model which verifies this condition.

The instantaneous team payoff function is denoted by $w(x_0, x_1, x_2)$. Since X_0 is not deterministic we shall be

considering the expected payoff

$$\mathbb{E}_f[w(X)] = \int_{x \in \mathcal{X}} w(x_0, x_1, x_2) f(x_0, x_1, x_2) dx_0 dx_1 dx_2 \quad (2)$$

where $X = (X_0, X_1, X_2)$, $x = (x_0, x_1, x_2)$, and $\mathcal{X} = \mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2$. In the sequel, we will also denote by $W(f)$ the above expected payoff i.e., $W(f) = \mathbb{E}_f[w(X)]$. What matters for the expected payoff is function f which characterizes the possible correlations among the three random variables X_0, X_1 , and X_2 . This correlation precisely measures the degree to which the agents can coordinate with each other and the system state. To understand the relationship between the agents strategies (1) and the expected payoff (2), let us define the notion of implementable distributions.

Definition 1 (Implementability): Let $s \in \{a, b\}$ be the assumed information structure. The probability density function $\bar{f}(x_0, x_1, x_2)$ is implementable if there exists a pair of control strategies (u_i, v_i^s) such that as $T \rightarrow +\infty$, we have for all $(x_0, x_1, x_2) \in \mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2$,

$$\frac{1}{T} \sum_{i=1}^T \int_{y \in \mathcal{Y}} f_{X_0 X_1 X_2 Y, i}(x_0, x_1, x_2, y) \rightarrow \bar{f}(x_0, x_1, x_2) \quad (3)$$

where $f_{X_0 X_1 X_2 Y, i} = \gamma \times f_{X_1, X_2 | X_0, i} \times f_0$ is the joint distribution induced by (u_i, v_i^s) at stage i .

Note that since the expectation value of the payoff is a linear operator with respect to the distribution f , the time averaged expected payoff $\bar{W}$ is reachable if and only if the corresponding distribution $\bar{f}$ is implementable. In Section III, we shall characterize the set of reachable or feasible average payoffs under the information structure given by (1), which is equivalent to characterizing the set of implementable distributions.

III. PERFORMANCE ANALYSIS: LIMITING PERFORMANCE CHARACTERIZATION

A. General case

In the case of finite alphabets $|\mathcal{X}_i| < \infty$, $i \in \{0, 1, 2\}$, it has been shown in the cases which have been treated so far [1], [2], [3], [4] that characterizing the set of implementable (mass) probability distributions amounts to determining a certain information constraint. For instance in the case of *discrete* sets and *perfect observation* of [1] ($Y = X_1$), the necessary and sufficient condition for a joint probability mass distribution $Q(x_0, x_1, x_2)$ to be implementable is that

$$H_Q(X_0) + H_Q(X_2) - H_Q(X_0, X_1, X_2) \leq 0 \quad (4)$$

where H_Q is the discrete entropy function under a fixed joint distribution (see Section [9] for the different expressions of the entropy used in this section). A well-known reasoning in information theory [9], and intensively used in control when communication problems are involved, is to use the information constraint derived in the discrete case and just replace the discrete entropy function with the differential entropy. It can be proved that this reasoning is perfectly valid if considered continuous variables are Gaussian (see e.g., [10] for a recent reference). For coordination problems

such as the one under investigation, imposing the agents' actions to be Gaussian is generally suboptimal. Elaborating further, if we replace the discrete entropy function with the differential entropy we obtain

$$h_f(X_0) + h_f(X_2) - h_f(X_0, X_1, X_2) \leq 0 \quad (5)$$

where h_f is the differential entropy under the fixed joint distribution f . It turns out that this condition can be shown to be non-necessary in general, indicating that the transition from the discrete case to the continuous case needs some special care in the problem under investigation. To convince the reader, let us recall one of the Cantor's theorems (see e.g., [11]). There exists a bijective map from $\mathbb{R}^T$ to $\mathbb{R}$. Therefore, a possible control strategy for Agent 1 might be as follows. On the first stage, Agent 1 maps or encodes the whole sequence of states $(x_0(1), \dots, x_0(T)) \in \mathbb{R}^T$ into a single action $x_1(1) \in \mathbb{R}$. Since Agent 2 observes this action perfectly, it can decode it perfectly and is thus informed of the sequence of states as well. This would mean that from stage $i = 2$, the two agents can correlate their actions in an arbitrary manner with the system state; in particular they can choose the pair (or one of the pairs) which maximizes w at a given stage $i \geq 2$. This means that any probability density function $\bar{f}_{X_0 X_1 X_2}$ can be implemented (asymptotically), contradicting the fact that any implementable distribution has necessarily to verify the continuous counterpart of (4) which is (5). This apparent contradiction comes from the fact that expressing (4) in the continuous case with differential entropies relies on assumptions which need to be specified rigorously for the problem. Indeed, the information constraint can be shown to be necessary and sufficient for implementability within some classes of random variables. One of the broadest classes which is known is provided in [12]. It turns out that if one wants to define a probability measure on the Cantor set, one does not fall into this broad class which is specified below. Let's first give the definition of a field in probability theory.

Definition 2 (field): Let $(\Omega, \mathcal{B})$ be a measurable space. We call field $\mathcal{F}$ a collection of subset of Ω such that :

$$\Omega \in \mathcal{F}$$

$$\text{if } F \in \mathcal{F} \text{ then } F^c \in \mathcal{F}$$

$\mathcal{F}$ is stable under finite union

A set A of a field $\mathcal{F}$ is called an atom if and only if the only subsets which are also member of the field are the set itself and the empty set.

Definition 3 (base): A sequence of finite field $\mathcal{F}_n$; $n = 0, 1, \dots$ is called a basis of a field $\mathcal{F}$ if $\mathcal{F}_n \uparrow \mathcal{F}$ and if G_n is a sequence of atoms of $\mathcal{F}_n$ such that $G_n \in \mathcal{F}_n$ and $G_{n+1} \subset G_n$, $n = 0, 1, 2, \dots$ then $\bigcap_{n=1}^{\infty} G_n \neq \emptyset$.

A sequence $\mathcal{F}_n$; $n = 0, 1, \dots$ is called a basis of a measurable space $(\Omega, \mathcal{B})$ if $\mathcal{F}_n$ are a basis of a field $\mathcal{F}$ which generates $\mathcal{B}$: $\mathcal{B} = \sigma(\mathcal{F})$. A field $\mathcal{F}$ is called standard if it has a basis. A measurable space $(\Omega, \mathcal{B})$ is called standard if it can be generated by a standard field i.e. $\mathcal{B}$ has a basis. We can now define the mutual information in a standard space provided by [13].

Definition 4 (Mutual Information): Let $(\Omega, \mathcal{F}, P)$ be a standard probability space and $X \in \mathbb{R}$, $Y \in \mathbb{R}$ two generic random variables: $X : \Omega \rightarrow \mathcal{A}_X$, $Y : \Omega \rightarrow \mathcal{A}_Y$ with $(\mathcal{A}_X, \mathcal{B}_X)$, $(\mathcal{A}_Y, \mathcal{B}_Y)$ two measurable spaces. Let $\mathcal{F}_X = X^{-1}(\mathcal{A}_X)$ and $\mathcal{F}_Y = Y^{-1}(\mathcal{A}_Y)$ the sub- σ -algebra of $\mathcal{F}$ induced by X and Y . Let

$$\mathcal{P}_X = \{A_j\}_{j=1}^{N_X} \subset \mathcal{F}_X \text{ and } \mathcal{P}_Y = \{B_j\}_{j=1}^{N_Y} \subset \mathcal{F}_Y \quad (6)$$

be finite partitions of Ω . With these partitions we associate the following random variables:

$$\begin{aligned} \tilde{X}(\omega) &= j \quad \text{for } \omega \in A_j \quad \text{with } 1 \leq j \leq N_X \\ \tilde{Y}(\omega) &= j \quad \text{for } \omega \in B_j \quad \text{with } 1 \leq j \leq N_Y \end{aligned} \quad (7)$$

The mutual information between $X \in \mathbb{R}$ and $Y \in \mathbb{R}$ is then defined by:

$$i(X; Y) = \sup_{\mathcal{P}_X, \mathcal{P}_Y} I(\tilde{X}; \tilde{Y}) \quad (8)$$

where I is the classical mutual information between two discrete random variables [14]. Similarly, the conditional mutual information is defined by

$$i(X; Y|Z) = i(X; Y, Z) - i(X; Z). \quad (9)$$

The above framework is exploited to prove the following two theorems.

Theorem 1 (Scenario a): Assume that all random variables under use are defined on a standard probability space. Consider a joint probability density distribution $\bar{f}(x_0, x_1, x_2)$ such that $\forall x_0 \in \mathcal{X}_0, \int_{x_1, x_2} \bar{f}(x_0, x_1, x_2) dx_1 dx_2 = f_0(x_0)$. Then, the distribution $\bar{f}$ is implementable if and only if $f(x_0, x_1, x_2, y)$ verifies the following information constraint:

$$i_f(X_0; X_2) \leq i_f(X_1; Y|X_0, X_2) \quad (10)$$

where the arguments of the mutual information $i_f(\cdot)$ are defined from f and $f(x_0, x_1, x_2, y) = \bar{f}(x_0, x_1, x_2)\gamma(y|x_1)$.

Theorem 2 (Scenario b): Assume that all random variables under use are defined on a standard probability space. Consider a joint probability density distribution $\bar{f}(x_0, x_1, x_2)$ such that $\forall x_0 \in \mathcal{X}_0, \int_{x_1, x_2} \bar{f}(x_0, x_1, x_2) dx_1 dx_2 = f_0(x_0)$. Then, the distribution $\bar{f}$ is implementable if and only if $f(x_0, x_1, x_2, y, x'_1)$ verifies the following information constraint:

$$i_f(X_0; X_2) \leq i_f(X'_1; Y, X_2) - i_f(X'_1; X_0, X_2) \quad (11)$$

with $f(x_0, x_1, x_2, y, x'_1) = f_{X'_1|X_0 X_1 X_2}(x'_1|x_0, x_1, x_2)\gamma(y|x_1)\bar{f}(x_0, x_1, x_2)$. X'_1 is an auxiliary variable which helps us exploit the joint typicality between X'_1 and X_1 as well as X'_1 and X_2 to create coding and decoding schemes.

We see that the first theorem provides a full characterization of implementable densities in scenario *a*. The second theorem, which relies in part on Gel'fand and Pinsker coding [15], provides a sufficient condition for implementability in scenario *b*; note that, as originally done in [15], we introduce an auxiliary random variable X'_1 to describe the information

constraint. These theorems therefore allow one to know to what extent a team can coordinate under the assumed information structure. In general, to determine the ultimate performance in terms of average payoff, an optimization problem for the functional $W(f)$ has to be solved. The constraints are that: f has to be a density function; its marginal over (x_1, x_2) has to be f_0 ; the density f as defined through the considered theorem has to meet the information constraint. In the next subsection, we apply the derived general result to a special case of probability distributions namely, Gaussian probability density functions. This allows one to exhibit a case where the information constraints can be quite easily expressed and to establish an interesting link with the work by Costa on dirty-paper coding [8].

B. Gaussian case

Here, we assume that all variables which intervene in the information constraints are Gaussian. Agent 2 is assumed to observe the actions of Agent 1 through an additive white Gaussian noise channel: $Y = X_1 + Z$ with $Z \sim \mathcal{N}(0, 1)$. Let σ_0^2 , σ_1^2 , and σ_2^2 respectively denote the variances of X_0 , X_1 , and X_2 . The correlation coefficient between X_i and X_j , $i \neq j$ is denoted by ρ_{ij} . Using these notations and specializing (10) and (11) in the Gaussian case the following results can be proved; proofs are omitted here for obvious space limitations.

Proposition 3 (Information constraint in scenario a):

Fix σ_0^2 . A necessary and sufficient condition for a joint probability density function $\bar{f}_{X_0 X_1 X_2}$ to be implementable is that the variances and correlation coefficients verify the following inequality: $-(\sigma_2^2 \rho_{01}^2 - 2\rho_{01} \rho_{02} \rho_{12} + \sigma_1^2 \rho_{02}^2 + \sigma_0^2 \rho_{12}^2 - \sigma_0^2 \sigma_1^2 \sigma_2^2) \times (\sigma_2^2 \rho_{01}^2 - 2\rho_{01} \rho_{02} \rho_{12} + \sigma_1^2 \rho_{02}^2 + \sigma_0^2 \rho_{12}^2 - \sigma_0^2 \sigma_1^2 \sigma_2^2 + \rho_{02}^2) \leq 0$.

Proposition 4 (Information constraint in scenario b):

Fix σ_0^2 . Let $\alpha_2 \in \mathbb{R}$. A sufficient condition for a joint probability density function $\bar{f}_{X_0 X_1 X_2}$ to be implementable is that the variances and correlation coefficients verify the following inequality: $(-\rho_{01}^2 + \sigma_0^2 + \sigma_0^2 \sigma_1^2) \times (\sigma_2^2 \rho_{01}^2 - 2\rho_{01} + 2\alpha_2 \rho_{12} + \sigma_1^2 \rho_{02}^2 + \sigma_0^2 \rho_{12}^2 - \sigma_0^2 \sigma_1^2 \sigma_2^2) - \sigma_0^2 \sigma_2^2 (\rho_{01}^2 - \sigma_0^2 \sigma_1^2 + \alpha_2^2 \rho_{02}^2 - \alpha_2^2 \sigma_0^2 \sigma_2^2 + \alpha_2^2 \sigma_1^2 \rho_{02}^2 + \alpha_2^2 \sigma_2^2 \rho_{01}^2 + \alpha_2^2 \sigma_0^2 \rho_{12}^2 - 2\alpha_2 \sigma_0^2 \rho_{12} + 2\alpha_2 \rho_{01} \rho_{02} - \alpha_2^2 \sigma_0^2 \sigma_1^2 \sigma_2^2 - 2\alpha_2 \rho_{01} \rho_{02} \rho_{12}) \leq 0$.

For finding the information constraint for scenario b, we have assumed that $X'_1 = X_1 + \alpha_0 X_0 + \alpha_2 X_2$. This choice is inspired by the Costa's dirty paper coding scheme [8]. The parameter α_2 is related to the choice we made for the auxiliary variable in Theorem 2.

As seen through Proposition 4, the value of the parameter α_0 does not play any role in the constraint, showing that only the correlation level between the agents' actions X_1 and X_2 has to be tuned properly. The inequality constraint function of Proposition 4 can be shown to be strictly convex w.r.t. α_2 and the optimum point α_2^* is given by:

$$\alpha_2^* = \frac{\sigma_0^2 \rho_{12} - \rho_{01} \rho_{02}}{\rho_{02}^2 - \sigma_0^2 (\sigma_2^2 - \rho_{12}^2 + \sigma_1^2 \sigma_2^2) + \sigma_1^2 \rho_{02}^2 + \sigma_2^2 \rho_{01}^2 - 2\rho_{01} \rho_{02} \rho_{12}}. \quad (12)$$

When the communication signal-to-noise ratio $\text{SNR} = \frac{\mathbb{E}(X_2^2)}{\mathbb{E}(Z^2)} = \sigma_1^2 \rightarrow \infty$, it is seen that $\alpha_2^* \rightarrow 0$ and the choice $X'_1 = X_1$ is optimal. When $\text{SNR} \rightarrow 0$, we see that $\alpha_2^* \rightarrow \frac{\sigma_0^2 \rho_{12} - \rho_{01} \rho_{02}}{(\rho_{02}^2 - \sigma_0^2 \sigma_2^2) + \sigma_0^2 \rho_{12}^2 + \sigma_2^2 \rho_{01}^2 - 2\rho_{01} \rho_{02} \rho_{12}}$.

IV. APPLICATION TO THE WITSENHAUSEN COST FUNCTION

A. The Witsenhausen cost function. Discussion

Taking the Witsenhausen cost (times minus one) as the payoff function to be minimized over a large number of stages and reducing it to a convex optimization problem, we find the limiting performance in terms of coordination for the two-agent team. However, note that this is significantly different from the original Witsenhausen's counterexample [5] as we optimize over many stages. Witsenhausen's counterexample is also a two-agent team problem, where one of the agents can see the system state perfectly, whereas the second agent only sees a noisy version of the first agent's action. The aim for the team is to have control strategies for each agent which will minimize the expectation value of the cost. The original counterexample has only one stage and has been of great interest since it was proposed in 1968, especially because it was one of the first examples where actions played a dual role of communicating as well as optimizing, rendering the solution of the optimization problem non trivial. The specific payoff function we consider in this section equals minus the Witsenhausen cost function that is,

$$w(x_0, x_1, x_2) = -[k(x_0 - x_1)^2 + (x_1 - x_2)^2], \quad k \geq 0. \quad (13)$$

Note that the notations used here differ from those used in [5].

Although this cost function is inspired by Witsenhausen's Counterexample, there are a few very important differences between the application of our theory to the Witsenhausen's Cost function and the Witsenhausen's counterexample. First, we optimise the average cost over a large time period, unlike in Witsenhausen's Counterexample where the one shot expectation value for the cost is minimised. Second, we assume strictly causal knowledge at Agent 2, whereas this is not the case in the original counterexample. To apply our approach to the original problem, albeit in the repeated case, one will need to consider a new scenario with the causality condition relaxed.

B. Numerical analysis

While the optimization problem is well defined, finding the joint distribution $f_{X_0 X_1 X_2 Y} = \gamma(y|x_1) * f_{X_1 X_2 | X_0} * f_0$ is not a trivial task. Eventhough f_0 is defined by the problem and $\gamma(y|x_1)$ can be generated given the noise model, we still need to find $f_{X_1 X_2 | X_0}$, and since the search space is

over all possible distributions, the computational complexity of such a search is very high. Therefore, in this section we restrict ourselves to two cases which are simpler to handle, complexity wise, but might be sub-optimal.

Both the cases use the results proven in previous sections for the continuous variables. However for the first case, we use a quantizer to discretise $(\mathcal{X}_0, \mathcal{X}_1, \mathcal{X}_2, \mathcal{Y})$ and optimise over them. This approach is inspired by the success of quantizers in finding better solutions for the original Witsenhausen Counterexample. To compare with a strategy most resembling a 'linear control' strategy, we choose the case where all variables are supposed to be gaussian, for which we had simplified the information constraint in terms of variances and covariances in section III-B. Clearly, it would be interesting to compare the two strategies, to see whether like for Witsenhausen Counterexample, non linear (discretisation) strategies outperform 'linear' (and thus continuous) strategies. This argument is used to motivate our choices of simulations, but the similarities are not so straightforward.

The following parameters are taken to be given for the problem: $k = 1, \sigma_0^2 = 25$, and $\mathbb{E}z^2 = 1$ and are common for both simulations (unless specified otherwise). Also, the information constraint considered is for scenario *a*.

Discrete case : we quantize all random variables to take nine values: $\mathcal{X}_0 = \mathcal{X}_1 = \mathcal{X}_2 = \mathcal{Y} = \{-\frac{24}{7}\sigma_0, -\frac{16}{7}\sigma_0, -\frac{8}{7}\sigma_0, -\frac{3}{7}\sigma_0, 0, \frac{3}{7}\sigma_0, \frac{8}{7}\sigma_0, \frac{16}{7}\sigma_0, \frac{24}{7}\sigma_0\}$, with $\sigma_0 = 5$. Indeed, as we consider the continuous random variable $\tilde{X}_0 \sim \mathcal{N}(0, \sigma^2)$, we partition uniformly the continuous space that $\tilde{X}_0$ is defined over, so that 99.99 % of the probability mass function lies in the chosen interval. It is well known that considering the interval $[-4\sigma, 4\sigma]$ achieves this. We then calculate the transition probabilities $P(Y|X_1)$ where $\tilde{Y} = X_1 + Z$, Z is supposed to be a Gaussian random variable: $Z \sim \mathcal{N}(0, 1)$, and Y is the discrete random variable that corresponds to $\tilde{Y}$. This problem is computationally simpler as it is easy to calculate entropies for discrete random variables. The optimization problem can be solved using convex optimization algorithms. We search for the joint distribution over $\mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2$ which minimizes the expectation of the Witsenhausen cost function. This approach is similar to the one used in [16], except for the cost function and the information constraint which takes into account the observation noise for Agent 2. It gives us an approximation to the continuous case and valuable ideas as we will explain now.

Results : In Fig 1, for low SNR (-10dB), the probability is almost 1 for $X_1 = 0$. This is logical as $\sigma_1^2 = 0.1$ and thus Agent 1 does not have too much of a choice. For medium SNR (10 dB), we see the probabilities diverging slightly and resembling a Gaussian distribution. The same distribution is observed from SNR = 14 dB onwards as this leads to minimum cost. This can be seen from the graph of expected payoff vs SNR, Fig 2. At 40 dB, one sees a distribution with higher variance but whether it can be a gaussian is tough to speculate given the lack of points.

Fig. 1. Optimal probability distribution of X_1 at different maximum SNRs. SNR=40 dB is the optimal distribution globally

The salient feature of this approach is that it does not suppose any distribution for the variables a priori, and thus finds the optimal distribution, which is not necessarily Gaussian. However, it only searches for finite action alphabets, thus not attaining optimality in the general continuous case.

Gaussian Case : Guessing the optimal distributions to be Gaussian, we find a feasible set of variances and covariances which satisfy the information constraint calculated in Section 3. The feasible set is found by quantizing the search space for all the parameters. The search space are as follows: $\frac{\sigma_1^2}{\sigma_2^2} \in (-10, 40)_{dB}$, $\frac{\sigma_2^2}{\sigma_0^2} \in (-20, 13)_{dB}$, and $\rho_{01}, \rho_{12}, \rho_{02} \in (0, (\sigma_0\sigma_1, \sigma_1\sigma_2, \sigma_2\sigma_0))$. The other constraints of the optimization problem are trivially satisfied since we are taking all distributions to be probability distributions from the beginning.

For a given set of variance and covariance values which satisfy the information constraint, we find the expected payoff by evaluating the integral $\int_{x \in \{\mathcal{X}_0 \times \mathcal{X}_1 \times \mathcal{X}_2\}} f(x)w(x)dx$. We do so using Monte Carlo simulations by randomly generating x_0, x_1 and x_2 (100000 draws) which follow the joint distribution $f(x)$ and averaging over the cost for the randomly generated triplets. We search over all the elements of the feasible set exhaustively to find the optimal joint Gaussian distribution which optimizes the Witsenhausen cost function.

Results : In Fig 2, we see that as $\mathbb{E}\frac{\sigma_1^2}{\sigma_2^2}$ (which could be looked at as Signal to Noise Ratio (SNR)) increases, the expected cost reduces initially and then becomes constant. This is because on the x-axis, we are considering the maximum SNR available, and in both the cases, one sees that after a certain SNR*, the payoff remains constant as the agents choose strategies with SNR*.

We notice that the discrete strategy does better than the continuous gaussian strategy and while this is not conclusive proof, it leads us to suspect that discretisation does better

Fig. 2. Scenario *a*; Plot of Expected payoff vs SNR for both the cases. One sees that the discrete strategy quickly beats the all-gaussian strategy, and achieves a much lower optimal cost

than continuous alphabets, which would be similar to the original Witsenhausen problem where discrete Non-Linear strategies were shown to outperform the best affine strategies. Although this might just be an artefact of gaussian variables being sub-optimal distributions for our problem.

For the Gaussian case, the optimal $(\sigma^*, \rho^*) = (\sigma_1^*, \sigma_2^*, \rho_{01}^*, \rho_{02}^*, \rho_{12}^*)$ was found to be (4.6, 5.3, 4.8, 5.1, 4.9) and the optimal expected cost $\mathbb{E}(w)^*$ associated with it to be 18.64. The information constraint is numerically found to be saturated for the optimal point. Note that the optimal correlations ρ^* are not at their maximum values, given by $\sigma_i \sigma_j \geq \rho_{ij}$. This is precisely the information constraint which prevents the correlations to be at their maximum values, thus penalising future communication by Agent 1 and preventing the agents to coordinate better.

Ideally one should be able to do more extensive simulations for the discrete case, so as to get a better idea of the optimal distribution which could then be tackled using an approach similar to the gaussian case. However, the computational complexity prevents us from doing so currently.

V. CONCLUSIONS

We generalized the information constraint for the scenarios discussed in the article from a discrete case to a continuous case, showing equivalence between implementable distributions and reachable payoffs. This result is independent of the cost function, but it depends on the strategies and information structures. Thus we created a general framework for tackling problems with similar information structures, as well as showed a method of generalizing information constraint found for discrete cases to continuous ones. We also found an interesting link between scenario *b* and [8] which needs to be further explored. Since Witsenhausen cost function is an important area for research, we applied our framework to optimize this cost function in our scenario.

While we could not provide simulations which solved the general optimization problem described in the article, we did gain some insights by simplifying the problem and reducing the computational complexity.

Further possibilities of exploration include better simulations with other types of distributions to see if they do better than Gaussian distributions, quantizing the alphabets with more points so as to approach continuity, as well as proving our results for other type of strategies and information structures. To compare a repeated version of Witsenhausen counterexample, one would need performance limit characterisation with a different Information structure, similar to ones discussed in [17].

ACKNOWLEDGMENT

Authors of this paper and particularly A. Agrawal would like to thank project LIMICOS - ANR-12-BS03-0005 for financing the project. We would also like to thank Pierre Gerold for his help in making our simulations faster.

REFERENCES

- [1] O. Gossner, P. Hernandez, and A. Neyman, "Optimal use of communication resources," *Econometrica*, vol. 74, no. 6, pp. 1603–1636, 2006.
- [2] B. Larrousse and S. Lasaulce, "Coded power control: Performance analysis," in *2013 IEEE International Symposium on Information Theory (ISIT'2013)*, Jul. 2013.
- [3] —, "More about optimal use of communication resources," in *International Workshop on Stochastic Methods in Game Theory, Erice, Italy*, Sep. 2013.
- [4] B. Larrousse, S. Lasaulce, and M. Bloch, "Coordination in distributed networks via coded actions with application to power control," *submitted to IEEE Transaction on Information Theory*, 2014.
- [5] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, pp. 131–147, 1968.
- [6] R. Bansal and T. Başar, "Stochastic teams with nonclassical information revisited: When is an affine law optimal?" *IEEE Transactions on Automatic Control*, vol. 32, no. 6, pp. 554–559, June 1987.
- [7] A. A. Kulkarni and T. P. Coleman, "An optimizer's approach to stochastic control problems with nonclassical information structures," in *51st IEEE Conference on Decision and Control (CDC'2012)*, Dec. 2012.
- [8] M. H. Costa, "Writing on dirty paper (corresp.)," *Information Theory, IEEE Transactions on*, vol. 29, no. 3, pp. 439–441, 1983.
- [9] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [10] A. El Gamal and Y.-H. Kim, *Network information theory*. Cambridge University Press, 2011.
- [11] L. A. Steen and J. A. . Seebach, *Counterexamples in Topology*, 1970.
- [12] R. Gray, *Probability, Random Processes, and Ergodic Properties*. SpringerVerlag, 2009.
- [13] A. Wyner, "A definition of conditionnal mutual information for arbitrary ensemble," *Inform. and Control*, pp. 51–59, 1978.
- [14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, 2006.
- [15] S. Gel'Fand and M. S. Pinsker, "Coding for channel with random parameters," *Probl. Contr. Inform. Theory*, vol. 9, no. 1, pp. 19–31, 1980.
- [16] B. Larrousse, A. Agrawal, and S. Lasaulce, "Implicit coordination in 2-agent team problems. application to distributed power allocation," in *IEEE 12th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, May 2014.
- [17] B. Larrousse, S. Lasaulce, and M. Bloch, "Coordination in distributed networks via coded actions with application to power control," *arXiv preprint arXiv:1501.03685*, 2015.