2009 American Control Conference
Hyatt Regency Riverfront, St. Louis, MO, USA
June 10-12, 2009

The online ouija board : a testbed for multi-party control of dynamical systems

Jérôme Barral Robert Wilson Cédric Langbort

***Abstract*— We introduce the *online ouija board*, a multi-player server-based game in which a team of agents must coordinate their control actions in real-time, so as to drive a token across a board and spell as many words as possible in a given time. This ouija game presents several of the typical features of multi-party control systems, namely: (i) it is *a networked control system*, since messages between the players' individual computers and the server are affected by asynchronism, delays, and possible packet drops; (ii) it is a *team theoretic/ distributed decision problem*, since different players have authority over different inputs, and individual choices influence the information available to other players, and (iii) it is a *distributed design problem*, since, when no communication is allowed, each player control law must be chosen independently, with access to a partial description of the token's dynamics.**

In this paper, we propose a simple model of the ouija board which, while assuming away the complications due to (i) and (ii), allows us to focus on the distributed design aspect of the problem mentioned in (iii). We show that simple control strategies exist, which require players to know the token's position and their own actuation direction, but nothing about their teammates' directions or input values. We then compare this simple strategy to the choices made by actual human players in the ouija game, and discuss the role that team communication may play in these choices.

I. INTRODUCTION

A traditional ouija board is printed with letters and sometimes words or numbers. A group of people play by placing their hands on a token and moving it around the board (sometimes unconsciously). In this way, the board spells out messages [6].

Our online multiplayer game, which is accessible and playable at [7], is inspired by the traditional Ouija board. Unlike in a traditional game of ouija however, players are all aware of the words to be spelled and have the goal of spelling as many target words as possible in a limited time. Figure 1 is a capture of the screen with which each player is presented during the game. The board and its token are in the upper left area and the control panel is in the upper right area, with the slanted line representing the constant direction along which a player can exert a force on the token. Some information about the game is posted in the lower right area, and the chat box under the board allows players to communicate with each other, when permitted by the game. By restricting the players' ability to communicate through the chat box, we can study how communication affects the choice of control strategy and performance. This will be discussed in more details in Section IV.

What makes this a difficult (but fun) task is that, although

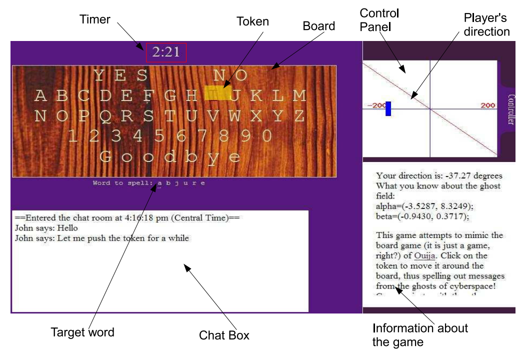

Fig. 1. The ouija board online game screen as seen by a player.

players have a common goal and can all observe the exact position of the token, each one can only apply a force in a *single*, randomly attributed, direction. In addition, each player initially ignores whether other players are present and, if so, how many, and in which direction they can actuate. In other words, the motion of the token from one letter to the next must be planned in a distributed way by a group of players who, individually, have only partial information about how the group can affect it.

Our online ouija board thus combines the features of distributed/team theoretic decision problems [4], [5] (since non-classical information patterns may arise if no side communication is allowed) with those of networked control systems [3] (since observation and actuation are possibly affected by asynchronism and network delays). In addition, playing the game with restricted information requires solving a *distributed design problem* [2], since players must construct their control inputs individually, with no knowledge of the token's dynamics. We believe that the online ouija board thus provides an ideal environment for investigating these issues and developing algorithms for general multi-party control problems, such as multi-vehicle command and control.

In addition to introducing the game, the goal of this paper is to show that, for a simple model of the online ouija board game, which neglects network effects and noise, there exist satisfactory control laws that can be constructed in a distributed way (i.e., which can be implemented by players knowing only their own actuation direction), and to compare them empirically with strategies used by human players.

All authors are with the Department of Aerospace Engineering, University of Illinois at Urbana-Champaign, 306 Talbot Laboratory MC-236, 104 S. Wright Street Urbana IL 61801. Email: {jbarral2, langbort}@illinois.edu, millstadtf@gmail.com

We also make some observations regarding the influence of communication on a team's performance, based on the data obtained in small scale experiments.
In Section II, we present the ouija board in more details. In Section III, we develop a control strategy for the token based on orthogonal projections, and investigate its performance in different situations (simple two-players case, with bounded inputs, with delay, N players). Finally, in Section IV, we present some experimental data.

II. Game description and implementation

Due to implementation constraints (limited speed of the server) and to make sure the game runs smoothly, control inputs for each user are sent to the server (and stored in the database) once per second. The token's new position is computed only when the *team leader* (defined as the player who connected first to the game) submits his control input. Players apply their control input through the control panel by moving the sliding bar to the right or to the left. *Figure 2* depicts the process occurring when players submit their control input to the server depending on their status (team leader or regular player). When a regular player submits his input, it is stored in the database and the server sends back the last computed position for the token. When the leader sends his input, the server retrieves the most recent control inputs for the rest of the players from the database and uses these to determine the new token position, which is then stored in the database as the current position, and sent back to the team leader.

More precisely, if we assume that the delays between the times at which the team leader and other players submit their inputs to the server is less than one second, we can decompose the update process into successive decision phases. Each phase starts and ends when the leader connects to the server, and every player sends a single control input per phase. At phase k, the token position is updated according to

$$\vec{x}_S(k) = \vec{x}_S(k-1) + \sum_{i=1}^{n} u_i(k) \cdot \vec{v_i} + \vec{n}(k), \qquad (1)$$

where $\vec{x}_S(k-1)$ is the position of the token saved on the server/database at the end of decision phase $k-1$, N is the number of players, $u_i(k)$ is the control input submitted by player i at phase k, and applied along its own direction $\vec{v_i}$ (attributed randomly to players when they connect to the server and assumed to be of unit norm). $\vec{n}$ is some random noise, treated as the input of some extra, dummy player.
Occasionally, communication with the server will be interrupted by packet losses. Depending on when the loss occurs, a player's control input may not be updated, resulting in the use of stale data in Equation (1), or in the inability of the server to calculate a new position for the token. In practice, packet losses affect game play on the scale of seconds and thus are rarely noticeable.

However, even in the absence of network degradations, the implementation described above in Equation (1) will introduce delays.
To understand why, consider the timeline pictured in Figure 3, in which d_{12} is the delay between the times at which player 2 and player 1 (assumed to be the leader) send their input to the server.

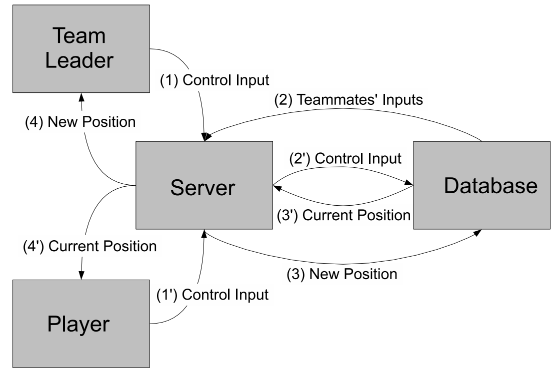

Fig. 2. Chronology of the communication between players and the server/database.

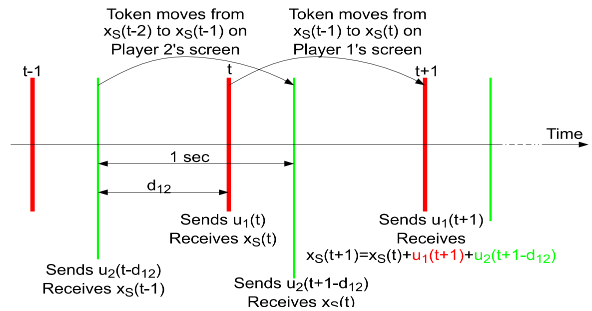

Fig. 3. Timeline of events in a decision phase for two players.

The control input $u_1(k)$ of the leader in decision phase k is computed for a token correctly assumed to be in position $x_S(k)$. However, control input $u_2(k)$, *which will actuate on position* $x_S(k)$, is computed using, at best, the position that was communicated to player 2 the last time it connected to the server, namely, $x_S(k-1)$. This implementation thus induces a one-step delay for every player except the leader. The effect of such delays on the stability of the system under a particular control law is considered in section III-C. Random delays are the subject of our current work.

III. A communication-less control law

Even when network effects can be neglected and all players can be assumed to have access to the same state observations at each decision period, the design of individual control strategies can still be complicated in the absence of inter-players communication. The main reason for this fact is that the control input $u_i(k)$ of player i must depend *only* on the token's position and player i's direction.
In absence of any other information, it may seem natural for player i to try and move the token towards the closest point to the target letter that is accessible to him, namely,

the orthogonal projection of the target onto its direction $\vec{v}_i$. This motivates using the control law :

$$u_i(k) = \alpha_i(\vec{x}(k) - \vec{x}_t) \cdot \vec{v}_i \quad (2)$$

where $\vec{x}_t$ represents the position of the current target letter.

A. Simple two players case

In this section we consider the simple case of two players applying control law (2) with $\alpha_1 = \alpha_2 = -1$, and neglect noise, server's delays, and dropped packets that can occur in the real game. The closed-loop token's dynamics then reduces to:

$$\vec{x}(k+1) = \vec{x}(k) - \sum_{i=1}^{2} (\vec{x}(k) - \vec{x}_t) \cdot \vec{v}_i$$

In order to compute the convergence rate to every target point and, in turn, deduce the number of steps needed to spell a given sequence of words, it is enough to study the following system:

$$\vec{x}(k+1) = \vec{x}(k) - \overrightarrow{\Pi_1(\vec{x}(k))} - \overrightarrow{\Pi_2(\vec{x}(k))} \quad (3)$$

where $\overrightarrow{\Pi_i(\vec{x}(k))}$ is the orthogonal projection of $\vec{x}(k)$ onto the direction $\vec{v}_i$. This amounts to shifting the origin to the target point.

Proposition 1. *Given a system whose dynamics are described by (3), and assuming that $\vec{v}_1$ and $\vec{v}_2$ are* not *colinear, an upper bound on the number of steps required to enter an ϵ-ball around the origin from any initial position $\vec{x}(0)$, is given by*

$$n = \left\lceil \frac{\log\left(\frac{\epsilon}{\|\vec{x}(0)\|}\right)}{\log(\delta)} \right\rceil \quad (4)$$

where $\delta = \max(|\cos(\theta_2 - \theta_1)|, |\cos(\theta_2 + \theta_1)|)$ and θ_i is the angle between vectors $\vec{x}(0)$ and $\vec{v}_i$.

Proof: Note that, if directions were globally known, the target could be reached exactly in a single step.

The proof of Proposition 1 is divided in two parts. First we prove asymptotic convergence to the origin and then we compute the bound (4).

On Figure 4, $\alpha_1(k)$ and $\alpha_2(k)$ are the coordinates of vector $\vec{x}(k) - \vec{x}_t$ in the basis defined by $\vec{v}_1$ and $\vec{v}_2$. The control law (3) yields $\alpha_1(k+1) = -\alpha_2(k)\vec{v}_1 \cdot \vec{v}_2$ and $\alpha_2(k+1) = -\alpha_1(k)\vec{v}_1 \cdot \vec{v}_2$, such that

$$|\alpha_1(k+1)| \leq |\alpha_2(k)| \text{ and } |\alpha_2(k+1)| \leq |\alpha_1(k)|.$$

Hence, $\alpha_1^2 + \alpha_2^2$ decreases at every step and is a Lyapunov function for our system.

Now we observe that, because $\|\overrightarrow{v}_1\|$=1,

$$\begin{aligned}\vec{x}(k+1) \cdot \vec{v}_1 &= \vec{x}(k) \cdot \vec{v}_1 - \vec{x}(k) \cdot \vec{v}_1 \\ &\quad - (\vec{x}(k) \cdot \vec{v}_2)(\vec{v}_2 \cdot \vec{v}_1)\end{aligned}$$

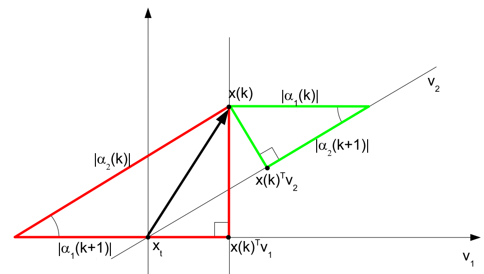

Fig. 4. proof of Proposition 1

i.e.

$$\begin{aligned}cos(\theta_1(k+1))\|\vec{x}(k+1)\| &= -\|\vec{x}(k)\|cos(\theta_2(k)) \\ &\quad cos(\theta_1(k) - \theta_2(k))\end{aligned}$$

From (3) this yields $cos(\theta_1(k+1)) = \pm cos(\theta_2(k))$. Similarly one can show that $cos(\theta_2(k+1)) = \pm cos(\theta_1(k))$. By inspection of all the possible cases, we can conclude that $cos(\theta_1(k) - \theta_2(k)) = cos(\theta_1(0) \pm \theta_2(0))$ for all k.
Hence, by defining

$$\delta = \max(|\cos(\theta_2 - \theta_1)|, |\cos(\theta_2 + \theta_1|)$$

we obtain: $\|\vec{x}(k)\| \leq \delta^k \|\vec{x}(0)\|$.

B. Bounded input case

Proposition 1 showed that a communication-less control law exists that can drive the token from any point to an ϵ-neighborhood of any other point in finite time. To use this result in the context of the ouija game, we need to modify our bound slightly to include the fact that:

- inputs sent by every player are bounded in magnitude by some constant $\mathbf{u}$,
- a letter is considered reached by the server when the center of the token enters a square region of side 2ϵ around the letter (instead of the ball considered in Proposition 1).

To this end, we now consider the case of a limited magnitude input and assume that the input of player i is now given by

$$u_i(k) = \begin{cases} \overrightarrow{\Pi_i(\vec{x}(k))} \cdot \vec{v}_i & \text{if } \|\overrightarrow{\Pi_i(\vec{x}(k))}\| \leq \mathbf{u} \\ \mathbf{u} & \text{if } \overrightarrow{\Pi_i(\vec{x}(k))} \cdot \vec{v}_i > \mathbf{u} \\ -\mathbf{u} & \text{if } \overrightarrow{\Pi_i(\vec{x}(k))} \cdot \vec{v}_i < -\mathbf{u} \end{cases} \quad (5)$$

Proposition 2. *In the case described above, where two players apply the projection strategy* with bounded inputs, *the target area around the point $\overrightarrow{x}_t = (x_t, y_t)$ is reached in a number of steps*

$$n = \left\lceil \frac{\sqrt{4+a^2}}{h(2+a^2)}(x_t + a y_t) \right\rceil + \left\lceil \frac{\sqrt{(x_0 - x_t - \varepsilon)^2 + (y_0 - y_t - \varepsilon)^2}}{\sqrt{3\frac{\mathbf{u}^2}{4}}} \right\rceil,$$

where a is defined as the slope of the line orthogonal to the direction of vector $\overrightarrow{v}_2$, assuming $\overrightarrow{v}_1$ is horizontal.

Proof: This proposition is proved in details in [9].

C. Delayed case

We now assume that Player 1 is the team leader defined in the presentation of the game (section 2) and is not delayed, while Player 2 is delayed by d steps ($d \geq 1$). The corresponding closed loop is then:

$$\vec{x}(k+1) = \vec{x}(k) - (\vec{x}(k) - \vec{x}_t) \cdot \vec{v}_1 - (\vec{x}(k-d) - \vec{x}_t) \cdot \vec{v}_2$$

Proposition 3. *There exists an increasing function γ of d, such that if $(\vec{v_1} \cdot \vec{v_2})^2 \geq \gamma(d)$, then the system is stable.*

Proof: In the delayed case, the dynamics of the token (when the target point is set as the origin) is described by

$$\begin{bmatrix} \vec{x}(k+1) \\ \vec{x}(k+2) \\ \cdot \\ \cdot \\ \vec{x}(k+1+d) \end{bmatrix} = A \begin{bmatrix} \vec{x}(k) \\ \vec{x}(k+1) \\ \cdot \\ \cdot \\ \vec{x}(k+d) \end{bmatrix}$$

where $A \in \mathbb{R}^{2(d+1)\times 2(d+1)}$ and can be shown to have characteristic polynomial :

$$\lambda^d(\lambda^{d+2} - \lambda^{d+1} + \lambda - v)$$

where $v := (\vec{v_1} \cdot \vec{v_2})^2$. For $v = 1$, we can compute the roots of this polynomial exactly as

$\lambda_j = exp\left(i\dfrac{\pi + 2\pi j}{d+1}\right)$ for $j = 0...d$ (the $d+1^{th}$ roots of -1),
$\lambda_{d+1} = 1$,
$\lambda_i = 0$ for $i = d+2...2d+1$.

We study how these roots vary as v is decreased from 1. The roots $\{\lambda_i\}_{i=d+2}^{2d+1}$ are always equal to zero and are stable, so we focus on the solutions of

$$F(\lambda, v) = \lambda^{d+2} - \lambda^{d+1} + \lambda - v = 0.$$

By the implicit function theorem, for all $j = 0...d+1$, there exists a smooth function $\lambda_j(v)$ and $\epsilon > 0$ such that

$$\begin{cases} F(\lambda_j(v), v) = 0 \,\forall\, v \in [1-\epsilon, 1] \\ \lambda_j(1) = \lambda_j \end{cases}$$

provided that $\dfrac{\partial F(\lambda_j, 1)}{\partial \lambda} \neq 0$.
In our case, $\dfrac{\partial F(\lambda_j, 1)}{\partial \lambda} = (d+2)\lambda^{d+1} - (d+1)\lambda^d + 1 \neq 0 \forall j$.
Hence functions $\lambda_j(v)$ exist and

$$\frac{d\lambda_j}{dv} = -\frac{\dfrac{\partial F}{\partial v}}{\dfrac{\partial F}{\partial \lambda_j}} = \frac{1}{(d+2)\lambda^{d+1} - (d+1)\lambda^d + 1}.$$

From this we can compute

$$\frac{d|\lambda_j|^2}{dv} = \lambda_j \frac{d\bar{\lambda}_j}{dv} + \bar{\lambda}_j \frac{d\lambda_j}{dv}$$

and we get

$$\left.\frac{d|\lambda_{d+1}|^2}{dv}\right|_{v=1} = \frac{1}{2} + \frac{1}{2} = 1 > 0$$

and for $j = 0...d$

$$\frac{d|\lambda_j|^2}{dv} = 2Re\left[-\frac{exp\left(i\dfrac{\pi+2\pi j}{d+1}\right)}{(d+1)\left(1 + exp\left(-id\dfrac{\pi+2\pi j}{d+1}\right)\right)}\right]$$
$$= Re\left[-\frac{exp\left(i\dfrac{\pi+2\pi j}{d+1}\right)\left(1 + exp\left(id\dfrac{\pi+2\pi j}{d+1}\right)\right)}{(d+1)\left(1 + \cos\left(d\dfrac{\pi+2\pi j}{d+1}\right)\right)}\right]$$
$$= \frac{1 - \cos\left(\dfrac{\pi+2\pi j}{d+1}\right)}{(d+1)\left(1 + \cos\left(d\dfrac{\pi+2\pi j}{d+1}\right)\right)} > 0.$$

In all cases $\left.\dfrac{d|\lambda_j|^2}{dv}\right|_{v=1} > 0$ which proves the existence of $0 \leq \gamma(d) < 1$ such that

$$\forall\, v \in [\gamma(d), 1],\ |\lambda_j(v)| < |\lambda_j(1)| = 1.$$

To show that γ is an increasing function of d, we consider Jury's test for $F(\lambda, v)$ for different values of d. The conditions for stability for $d = k+1$ being the same as those for $d = k$ except for an additional one, we conclude that the size of the stability region decreases as d increases.

D. N players case

Proposition 4. *In the case of N players applying the control law defined in (2),*

$$\sum_{i=1}^{n} \alpha_j \ \leq 2$$

is a sufficient condition to have convergence of $\vec{x}$ to $\vec{x}_t$

Proof: In the N players case the dynamics is described by

$$\vec{x}(k+1) = \vec{x}(k) - \sum_{i=1}^{n} \alpha_i \overrightarrow{\Pi_i(\vec{x}(k))}$$

where, as before, we have translated the origin to the target. Hence

$$\|\vec{x}(k+1)\|^2 = \|\vec{x}(k)\|^2 + \|\sum_{i=1}^{n} \alpha_i \overrightarrow{\Pi_i(\vec{x}(k))}\|^2$$
$$- 2\vec{x}(k) \cdot \sum_{i=1}^{n} \alpha_i \overrightarrow{\Pi_i(\vec{x}(k))}$$

and, by expanding the inner products,

$$\|\vec{x}_{k+1}\|^2 = \|\vec{x}(k)\|^2 \Bigg[1 + \sum_{i=1}^{n} \alpha_i^2 cos^2(\theta_i(k))$$
$$+ 2\sum_{i\neq j} \alpha_i\alpha_j cos(\theta_i(k))cos(\theta_j(k))cos(\theta_i(k)-\theta_j(k))$$
$$-2\sum_{i=1}^{n} \alpha_i^2 cos^2(\theta_i(k))\Bigg]$$
$$= \|\vec{x}(k)\|^2 \Bigg[1 + \sum_{i=1}^{n}(\alpha_i^2 - 2\alpha_i)cos^2(\theta_i(k))$$
$$+ \sum_{i\neq j} \alpha_i\alpha_j cos^2(\theta_i(k)-\theta_j(k))$$
$$+ \sum_{i\neq j} \alpha_i\alpha_j (cos^2(\theta_i(k)) + cos^2(\theta_j(k)) - 1)\Bigg]$$
$$= \|\vec{x}(k)\|^2 \Bigg[1 - \sum_{i\neq j} \alpha_i\alpha_j$$
$$+ \sum_{i=1}^{n}(\alpha_i^2 - 2\alpha_i + \alpha_i\sum_{i\neq j}\alpha_j)cos^2(\theta_i(k))$$
$$+ \sum_{i\neq j} \alpha_i\alpha_j cos^2(\theta_i(k)-\theta_j(k))\Bigg]$$

Hence, a sufficient condition for convergence to the origin is

$$\alpha_i^2 - 2\alpha_i + \alpha_i\sum_{i\neq j}\alpha_j \ \leq 0 \ \forall\ i$$

$$\text{i.e., } \alpha_i - 2 + \sum_{i\neq j}\alpha_j \ \leq 0$$

$$\text{or } \sum_{i=1}^{n}\alpha_j \ \leq 2.$$

IV. Experiments and the Role of Communication

The projection-based strategy of section III was developed in a simplified context, which disregarded a number of features of the actual game, and under the implicit assumption that players, although having access to partial information about the rest of the team, all implement the same globally selected strategy.
In order to see how this strategy performs in a real game environment, and identify directions in need of further theoretical developments, we implemented it on two artificial agents. We then compared the results of these experiments with two games of ouija involving human players under various communication scenarios. These comparisons point to two elements that are not present in the projection-based strategy, but may be responsible for a large part of in-play team communication, namely: *signaling* and *team coordination towards a joint control strategy*. We use "signaling" in the sense of [8], to mean that actions are chosen by players not to achieve the advertised control goal, but, instead, to transmit information to others.

In the first experiment with human players (which we will refer to as E_1), the chat box placed at the bottom of the screen was deactivated. In turn, each player only had direct access to their own actuation direction and the position of the token at every time. In the second experiment (denoted $E2$), communication through the chat box was enabled. At the beginning of a round, and before the timer started, players could choose and create up to five different messages to be used during the game (for example, "what is your direction?" and "let me push alone" were possibilities, as well as push-buttons sending a player's actuation direction to the chatbox).

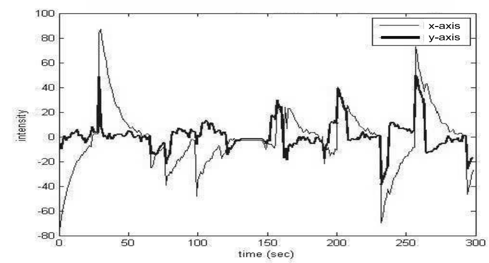

Fig. 5. Two simulated players applying the projection-based strategy. This team spelled 5 complete words in 5 minutes.

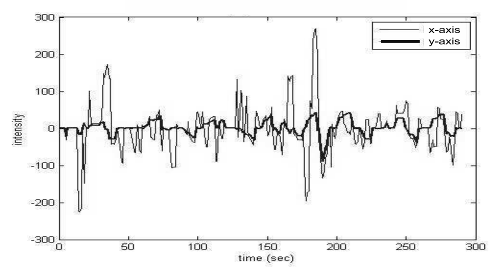

Fig. 6. Two human players without communication. This team spelled 3 complete words in 5 minutes.

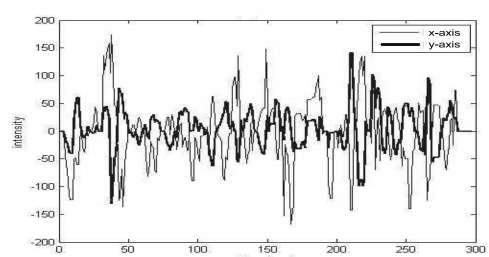

Fig. 7. Two human players with communication. This team spelled 4 complete words in 5 minutes.

The results of these experiments are summarized in Figures 5, 6, and 7. Each curve presents the time history of the $x-$ and $y-$ coordinates of the full control input, i.e.,

$$(u_1\overrightarrow{v_1} + u_2\overrightarrow{v_2}).\overrightarrow{e_x} \text{ and } (u_1\overrightarrow{v_1} + u_2\overrightarrow{v_2}).\overrightarrow{e_y},$$

respectively, as the same sequence of target words is presented to the teams. By looking at these normalized signals, we can compare the input histories of all three experiments, even though players were assigned different actuation directions in each case. The bound $\mathbf{u}$ on an individual player's input intensity was 200 in all cases.

A naked eye observation of the curves (confirmed by a fast Fourier transform of the signals, which is not reproduced here for space reasons) shows that:

- the global input signal of Figure 7 is more oscillatory than the signals of Figures 5 and 6.

- There are several 5 to 10 seconds-long intervals during which the global input of the team does not vary significantly in Figures 5 and 6.

In the case of communication-less human players, this lack of activity may correspond to a learning or signaling episode, during which each player keeps his input constant in the hope that the other one will infer his actuation direction from the token's motion. Another possible reason why a player's input may not vary is that he is unsure where to move the token next.

In order to gain more insight into the reasons for the lack of player's reactivity and be able to distinguish between the kind of episode described above and situations where almost constant inputs are explicitly chosen by the players (as in the case of the experiment of Figure 5), we plan to modify the game environment so as to allow players to comment on their input choices orally in real time.

Figures 5, 6, and 7 show that the projection-based strategy control strategy achieves the highest number of correctly spelled words in experiments, even though it does not involve real-time team communication. While this may appear surprising at first, given that communication seems to improve the performance of teams of human players, one should remark that the team of automated agents requires less communication than the human players in experiments E_1 and E_2, since strategies have been chosen for them prior to the beginning of the game, and they thus do not need to agree on a controller. In contrast, the human players (or any team which must choose *and* implement a control strategy during the timed portion of the game) use the chat-box or signaling to communicate not only about the players' private information but also about the choice of a group controller.

In order to rule out the possibility that the strategy played in E_1 is actually the same as the one played in E_2, and that the difference in performance is due to the inability of human players to apply this strategy accurately, we implemented a third experiment in which a human plays with an automated player. At every instant, the human player was shown the input that an automated player would have applied and was instructed to follow it as best as possible. The results, which can be seen on Figure 8, show that a human player do basically as well as an automated one when he is shown exactly what to do. It can thus be hypothesized that the differences between performances are due to differences between strategies.

V. Conclusion and future directions

We have implemented an online multiplayer game, which incorporates the main features of multi-party control problems: distributed design, decision under non-standard information patterns, and network effects. By empirically comparing a simple communication-less control strategy to the control policies used by teams of human players, we could hypothesize that signaling and the need for players to agree on a strategy in real-time contribute an important part to team communication.

In order to study these two processes separately, we plan to modify the setup of our game and, in particular, add an un-timed negotiation round to account for the agreement phase. We also plan to study the interplay between distributed design, communication, and implementation of a control law from a more theoretical viewpoint in the future.

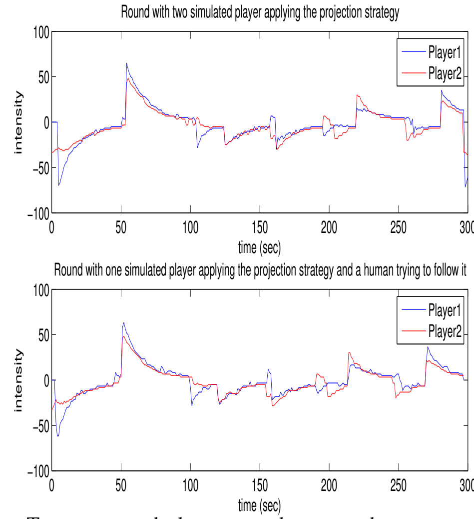

Fig. 8. Two automated players vs a human and an automated player.

References

[1] R. Cogill and S. Lall. An Approximation Algorithm for the Discrete Team Decision Problem. *SIAM J. Control Optim.* 45, 4 (Sep. 2006), 1359-1368.

[2] J.-Ch. Delvenne and C. Langbort. The Price of Distributed Design in Optimal Control. In *Proceedings of the IEEE Conference on Decision and Control*, pp. 3640-3645, 2006.

[3] J. Hespanha, P. Naghshtabrizi, Y. Xu. A Survey of Recent Results in Networked Control Systems. *Proc. of IEEE Special Issue on Technology of Networked Control Systems*, 95(1):138162, Jan. 2007

[4] Y. C. Ho and K. Chu. Team decision theory and information structures in optimal control problems–Part I *IEEE Transactions on Automatic Control*, vol. 17, no. 1, pp. 15 - 22, Feb 1972.

[5] J. Marschak and R. Radner. *Economic Theory of Teams*. Cowles Foundation Monograph 22, Yale University Press, 1972.

[6] The Museum of Talking Boards. `http://www.museumoftalkingboards.com`

[7] The online ouija board website. `http://rwilson4.com/ouija.php`

[8] M. Spence, Job Market Signaling. *The Quarterly Journal of Economics*, Volume 87, Issue 3 (Aug., 1973), 355-374.

[9] Jerome Barral thesis. `http://www.ae.uiuc.edu/ jbarral2/thesis.pdf`