



HAL
open science

Stochastic Alignment Processes

Amos Korman, Robin Vacus

► **To cite this version:**

| Amos Korman, Robin Vacus. Stochastic Alignment Processes. 2021. hal-03124213v1

HAL Id: hal-03124213

<https://hal.science/hal-03124213v1>

Preprint submitted on 29 Jan 2021 (v1), last revised 20 Dec 2022 (v4)

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.

Stochastic Alignment Processes

Amos Korman and Robin Vacus

January 29, 2021

Abstract

The tendency to align to others is inherent to social behavior, including in animal groups, and flocking in particular. Here we introduce the *Stochastic Alignment* Problem, aiming to study basic algorithmic aspects that govern alignment processes in unreliable stochastic environments. Consider n birds that aim to maintain a cohesive direction of flight. In each round, each bird receives a noisy measurement of the average direction of others in the group, and consequently updates its orientation. Then, before the next round begins, the orientation is perturbed by random drift (modelling, e.g., the affects of wind). We assume that both noise in measurements and drift follow Gaussian distributions. Upon receiving a measurement, what should be the orientation adjustment policy of birds if their goal is to minimize the average (or maximal) expected deviation of a bird's direction from the average direction? We prove that a distributed *weighted-average* algorithm, termed \mathcal{W}^* , that at each round balances between the current orientation of a bird and the measurement it receives, maximizes the social welfare. Interestingly, the optimality of this simple distributed algorithm holds even assuming that birds can freely communicate to share their gathered knowledge regarding their past and current measurements. We find this result surprising since it can be shown that birds other than a given i can collectively gather information that is relevant to bird i , yet not processed by it when running a weighted-average algorithm. Intuitively, it seems that optimality is nevertheless achieved, since, when running \mathcal{W}^* , the birds other than i somehow manage to collectively process the aforementioned information in a way that benefits bird i , by turning the average direction towards it. Finally, we also consider the game-theoretic framework, proving that \mathcal{W}^* is the only weighted-average algorithm that is at Nash equilibrium.

Keywords: Noisy communication, Kalman filter, Flocking, Biological distributed algorithms, Distributed signal processing, Clock Synchronization, Weighted-average algorithms.

1 Introduction

1.1 Background and Motivation

Reaching agreement, or approximate agreement, is fundamental to many distributed systems, including, e.g., computer networks, mobile sensor systems, animal groups, and neural networks [11, 8, 28, 23, 27, 7]. In the natural world, one of the most beautiful manifestations of approximate agreement happens when flocking birds (or, e.g., schooling fish) manage to maintain a cohesive direction of movement by constantly aligning themselves to others [28, 27, 7]. In this context, as well as in multiple other contexts, such as during clock synchronization [26, 25], the space in which the approximate agreement process occurs is continuous, measurements are noisy and the output needs to be maintained over time despite drift.

This paper introduces the *Stochastic Alignment* problem, aiming to capture some of the basic algorithmic challenges involved in reaching approximate agreement under such stochastic conditions. Informally, the problem considers a group of n agents positioned on the real line, aiming to be located as close as possible to one another. Initially, agents' positions are sampled from a Gaussian distribution around 0. In each round, each agent receives a noisy measurement of its current deviation from the average position of others. Then, governed by the rules of its algorithm, each agent performs a *move* to re-adjust its position. Subsequently, before the next round begins, the position of each agent is perturbed following random drift. Both noises in measurements and random drifts are governed by Gaussian distributions. We are mostly interested in the following questions:

- Which re-adjustment rule should agents adopt if their goal is to minimize the maximal (or average) expected distance of an agent from the center of mass?
- Could further communication between agents (e.g., by sharing measurements) help?
- What would be the impact on the global alignment when each agent aims to minimize its own distance from the center of mass?

Importantly, we assume that agents are unaware of the actual value of their current positions, and of the realizations of the random drifts, and instead, must base their movement decisions only on noisy measurements of relative positions. This lack of global “sense of orientation” prevents the implementation of the trivial distributed protocol in which all agents simply move to a predetermined point, say 0.

One trivial algorithm is the “fully responsive” protocol, where in each round, each agent moves all the way to its current measurement of the average position of others. This alignment protocol was assumed in various models that consider alignment, including in the celebrated flocking model by Vicsek et al. [27]. When drift is large, measurement noise is negligible, and the number of agents is large, this protocol is expected to be highly efficient. However, when measurement noise is non-negligible, it is expected that incorporating past measurements could enhance the cohesion, even though drift may have changed the configuration considerably.

Perhaps the simplest algorithms that take such past information into account are *weighted-average* algorithms; By weighing the current position against the measured position in a non-trivial way, such algorithms can potentially exploit the fact that the current position implicitly encodes information from past measurements. Indeed, in a centralized setting, when a single agent aims to estimate a fixed target relying on noisy Gaussian measurements, a weighted-average algorithm is known to be optimal [1]. However, here the setting is more complex since it is distributed, and the objective goal is to estimate (and get closer to) the center of mass, which is a function of the agents decisions.

1.2 The Stochastic Alignment problem.

We consider n agents located on the real line¹ \mathbb{R} . Let $I = \{1, \dots, n\}$ be the set of agents. We denote by $\theta_i^{(t)} \in \mathbb{R}$ the position of Agent i at round t , where it is assumed that initially agents are normally distributed

¹Depending on the application, the actual domain may be bounded, or periodic. For example, when modeling directions, the domain is $[-\pi, \pi]$ and when modeling clock synchronization, the domain may be $[0, T]$ for some phase duration T . Since we are interested in the cases where agents are more or less aligned, approximating an interval domain with the real line is not expected to reduce the generality of our results.

around 0, with variance σ_0^2 , that is, for each agent i ,

$$\theta_i^{(0)} \sim \mathcal{N}(0, \sigma_0^2).$$

Execution proceeds in discrete rounds. At round t , each agent i receives a noisy measurement of the deviation from the current average position of all other agents. Specifically, denote the average of the positions of all agents except i by:

$$\langle \theta_{-i}^{(t)} \rangle = \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \theta_j^{(t)}.$$

Let $\bar{\theta}_i^{(t)} = \langle \theta_{-i}^{(t)} \rangle - \theta_i^{(t)}$ denote the *stretch* of Agent i . At any round t , for every $i \in I$, a noisy measurement of the stretch of Agent i is sampled:

$$Y_i^{(t)} = \bar{\theta}_i^{(t)} + N_{m,i}^{(t)}, \quad (1)$$

where $N_{m,i}^{(t)} \sim \mathcal{N}(0, \sigma_m^2)$. In response, Agent i makes a move $d\theta_i^{(t)}$ and may update its memory state (if it has any). Finally, the position of Agent i at the next round is obtained by adding a drift:

$$\theta_i^{(t+1)} = \theta_i^{(t)} + d\theta_i^{(t)} + N_{d,i}^{(t)}, \quad (2)$$

where $N_{d,i}^{(t)} \sim \mathcal{N}(0, \sigma_d^2)$. All random perturbations $(N_{m,i}^{(t)})_{i \in I}$ and $(N_{d,i}^{(t)})_{i \in I}$ are mutually independent, and we assume that $\sigma_m, \sigma_d > 0$.

The *cost* of Agent i at a given time t is the absolute value of its expected stretch at that time ², i.e.,

$$C_i^{(t)} := \mathbb{E} \left(\left| \bar{\theta}_i^{(t)} \right| \right).$$

Note that the cost depends on the algorithm used by i but also on the algorithms used by others. As these algorithms will be clear from the context, we typically omit mentioning them in notations.

Definition 1 (Optimality). We say that an algorithm is *optimal* if, for every $i \in \{1, \dots, n\}$ and every round t , no algorithm can achieve a strictly smaller cost $C_i^{(t)}$.

Weighted-average algorithms. Perhaps the simplest algorithms that one may consider are *weighted-average* algorithms. Such an algorithm is characterized by a responsiveness parameter $\rho^{(t)}$ for each round t , indicating the weight given to the measurement at that round. Formally, an agent i following the *weighted-average* algorithm $\mathcal{W}(\rho^{(t)})$ at round t , sets

$$d\theta_i^{(t)} = \rho^{(t)} Y_i^{(t)}. \quad (3)$$

Full communication model. When executing a weighted-average algorithm, an agent bases its decisions solely on its own measurements. A main question we ask is whether, and if so to what extent, can performances be improved if agents could communicate with each other to share their measurements. In order to study the impact of communication, we compare the performances of the best weighted-algorithm to the performances of the best algorithm in a full-communication setting, where agents are free to share their measurements with all other agents at no cost. In the case that agents have identities, this setting is essentially equivalent to the following centralized setting: Consider a *master* agent that is external to the system. The master agent receives, at any round t , the stretch measurements of all agents, i.e., the collection $\{Y_j^{(t)}\}_{j=1}^n$, where these measurements are noisy in the same manner as described in Eq. (1). Analyzing these measurements at round t , the master agent then instructs each agent i to move by a quantity $d\theta_i^{(t)}$. After moving, the agents are subject to drift, as described in Eq. (2). Note that the master agent is unable to “see”

²Another natural cost measure is the expected deviation from the average position of all agents (including the agent), i.e., $C_i^{(t)'} := \mathbb{E} \left(\left| \frac{1}{n} \sum_{j=1}^n \theta_j^{(t)} - \theta_i^{(t)} \right| \right)$. These two measures are effectively equivalent. Indeed, $C_i^{(t)'} = \frac{n-1}{n} C_i^{(t)}$, thus an algorithm minimizing one measure will also minimize the other.

the positions of the agents, and its information regarding their locations is based only on the measurements it gathers from the agents, and on the movements it instructs. The goal of the master agent is to minimize the (average or maximal) cost of agents, per round. In particular, an algorithm is said to be optimal in the centralized setting if it satisfies Definition 1.

At a first glance, it may appear that weighted-average algorithms may be sub-optimal in the centralized setting. This is because it can be shown that the measurements made by Agent i contain strictly less information about the agent's relative position to the center of mass than the information contained in the collection of all measurements $\{Y_j^{(t)}\}_{j \neq i}$. Indeed, although the $\{Y_j^{(t)}\}_{j \neq i}$ measurements are not centered around the stretch $\bar{\theta}_i^{(t)}$ of Agent i , they still contain useful information. For example, it can be shown that

$$-\sum_{\substack{j=1 \\ j \neq i}}^n Y_j^{(t)} = \bar{\theta}_i^{(t)} - \sum_{\substack{j=1 \\ j \neq i}}^n N_{m,j}^{(t)}, \quad (4)$$

thus representing an additional “fresh” estimation of $\bar{\theta}_i^{(t)}$. Therefore, it may appear that in the centralized setting, the stretch of agents could potentially be reduced by letting the master agent process all measurements.

1.3 Our results

Weighted-average algorithms. We first investigate weighted-average algorithms in which all agents have the same responsiveness ρ , that furthermore remains fixed throughout the execution (see Eq. (3)). The proof of the following theorem is deferred to Appendix A.1.

Theorem 2. *Assume that all agents execute $\mathcal{W}(\rho)$, for a fixed $0 \leq \rho \leq 1$. Then for every $i \in \{1, \dots, n\}$ and every $t \in \mathbb{N}$, the stretch $\bar{\theta}_i^{(t)}$ is normally distributed, and*

$$\lim_{t \rightarrow +\infty} \text{Var} \left(\bar{\theta}_i^{(t)} \right) = \frac{\frac{n}{n-1}(\rho^2 \sigma_m^2 + \sigma_d^2)}{1 - (1 - \frac{n}{n-1}\rho)^2}, \quad (5)$$

with the convention that $\lim_{t \rightarrow +\infty} \text{Var} \left(\bar{\theta}_i^{(t)} \right) = +\infty$ if the denominator $1 - (1 - \frac{n}{n-1}\rho)^2 = 0$.

If all agents run $\mathcal{W}(\rho)$, then the extent to which they are aligned with each-other asymptotically is captured by $\text{Var}(\rho) := \lim_{t \rightarrow +\infty} \text{Var} \left(\bar{\theta}_i^{(t)} \right)$. Indeed, for every i , since $\bar{\theta}_i^{(t)}$ is normally distributed,

$$\lim_{t \rightarrow +\infty} \mathbb{E} \left(\left| \bar{\theta}_i^{(t)} \right| \right) = \sqrt{\frac{2}{\pi} \text{Var}(\rho)}.$$

The minimal value of this is achieved when taking $\text{argmin}_\rho \text{Var}(\rho)$ as the responsiveness parameter. The proof of the following theorem is deferred to Appendix A.1.

Theorem 3. *The weighted-average algorithm that optimizes group variance among all weighted-average algorithms $\mathcal{W}(\rho)$ (that use the same responsiveness parameter ρ at all rounds) is $\mathcal{W}(\rho^*)$, where*

$$\rho^* = \frac{\sigma_d \sqrt{4\sigma_m^2 + \left(\frac{n}{n-1}\sigma_d\right)^2} - \frac{n}{n-1}\sigma_d^2}{2\sigma_m^2}. \quad (6)$$

When n is large, Eq (6) becomes

$$\rho^* \approx \frac{\sigma_d \sqrt{4\sigma_m^2 + \sigma_d^2} - \sigma_d^2}{2\sigma_m^2}.$$

For example, if $\sigma_m \gg \sigma_d$, then $\rho^* \approx 0$. However, if $\sigma_m \ll \sigma_d$ then $\rho^* \approx 1$. Interestingly, if $\sigma_m = \sigma_d$ then $\rho^* \approx \frac{\sqrt{5}-1}{2}$, which is highly related to the golden ratio. Moreover, for large n , the minimal $\text{Var}(\rho)$ is

$$\text{Var}(\rho^*) = \frac{1}{2}\sigma_d \left(\sqrt{4\sigma_m^2 + \sigma_d^2} + \sigma_d \right). \quad (7)$$

Note that when the measurements are perfect, i.e., $\sigma_m = 0$, we have $\text{Var}(\rho^*) = \sigma_d^2$, which is the best achievable value that an agent can hope for, since no strategy can overcome the drift-noise.

The impact of communication. Our next goal is to understand whether, and if so, to what extent, can performances be improved if further communication between agents is allowed. For this purpose, we compare the performances of $\mathcal{W}(\rho^*)$ to the performances of the best algorithm in a centralized (full-communication) setting.

A natural candidate for an optimal algorithm in the centralized setting is the “meet at the center” algorithm. This algorithm first obtains, for each agent, the best possible estimate of the distance from the agent’s position to the center of mass $\langle \theta \rangle$, based on all measurements, and then instructs the agent to move by this quantity (towards the estimated center of mass). However, it is not immediate to figure out the distances to the center of mass, and furthermore, quantify the performances of this algorithm. To this end, we adapt the celebrated Kalman filter tool, commonly used in statistics and control theory [1], to our setting. Solving the Kalman filter system associated with the centralized version of our alignment problem, we obtain an estimate of the relative distance of each agent i from the center of mass (based on all measurements). To describe these estimates we first define the following.

Definition 4. We inductively define the sequence $(\alpha_t)_{t=0}^\infty$. Let $\alpha_0 = n\sigma_0^2/(n-1)$, and for every integer t , let

$$\alpha_{t+1} = \frac{\sigma_m^2 \alpha_t}{\frac{n}{n-1} \alpha_t + \sigma_m^2} + \frac{n}{n-1} \sigma_d^2.$$

Definition 5. For every integer t , let

$$\rho_\star^{(t)} = \frac{\alpha_t}{\frac{n}{n-1} \alpha_t + \sigma_m^2}.$$

At each round t , the Kalman filter returns an estimate of the relative distance of each agent i from the center of mass which turns out to be $\frac{n-1}{n} \rho_\star^{(t)} \left(Y_i^{(t)} - \frac{1}{n-1} \sum_{j \neq i} Y_j^{(t)} \right)$. As guaranteed by the properties of the Kalman filter, these estimates minimize the expected sum of square-errors, which can be translated to our desired measure of minimizing the agents’ costs. The “meet at the center” algorithm is given by Algorithm 1 below, and the following theorem stating its optimality is proved in Section 3.

Algorithm 1: Meet at the center

```

1 foreach round  $t$  do
2   Consider all measurements at round  $t$ ,  $\{Y_j^{(t)} \mid 1 \leq j \leq n\}$ ;
3   foreach agent  $i$  do
4     Set  $d\theta_i^{(t)} = \frac{n-1}{n} \rho_\star^{(t)} \left( Y_i^{(t)} - \frac{1}{n-1} \sum_{j \neq i} Y_j^{(t)} \right)$ ;      /* Output an estimate of  $\langle \theta \rangle - \theta_i$  */
5   end
6 end

```

Theorem 6. *Algorithm 1 is optimal in the centralized setting.*

Quite remarkably, another solution the follows from the Kalman filter estimations is in the form of a weighted-average algorithm. The proof of the following theorem is given in Section 3.

Theorem 7. *The weighted-average algorithm $\mathcal{W}^* := \mathcal{W}(\rho_\star^{(t)})$ is optimal in the centralized setting.*

Note that by the strong definition of optimality (Definition 1), for any given i , no algorithm in the centralized setting can achieve a better cost for agent i (for any round t). We find this optimality result surprising since, as mentioned, agents other than a given i can collectively gather information that is relevant to agent i , yet not processed by it when running a weighted-average algorithm (see Eq. (4)). Intuitively, it seems that optimality is nevertheless achieved, since, when running \mathcal{W}^* , the agents other than i somehow manage to collectively process the aforementioned information in a way that benefits agent i , by shifting the center of mass towards it.

In contrast to $\mathcal{W}(\rho^*)$, Algorithm \mathcal{W}^* uses a different responsiveness $\rho_\star^{(t)}$ at each round t . We next argue that the sequences (α_t) and $(\rho_\star^{(t)})$ converge. Not surprisingly, at the limit, we recover the optimal responsiveness as stated in Theorem 3. The proof of the following claim is deferred to Appendix A.2.




Figure 1: Position of the center of mass of a group with $n = 3$ agents over time, when \mathcal{W}^* is used (red), and when “meet at the center” is used (blue), while both algorithms face the same randomness, in both measurement noise and drift, and initialization of positions. Parameters are $\sigma_m = \sigma_d = 1$.

Claim 8. *The sequence α_t converges to $\alpha_\infty := \lim_{t \rightarrow +\infty} \alpha_t = \frac{1}{2} \left(\sigma_d \sqrt{4\sigma_m^2 + \left(\frac{n}{n-1}\sigma_d\right)^2} + \frac{n}{n-1}\sigma_d^2 \right)$.*

Moreover, $\lim_{t \rightarrow +\infty} \rho_\star^{(t)} = \rho_\star$.

Note that in the centralized setting, once we have an optimal algorithm A , we can derive another optimal algorithm B but simply shifting all agents, at each round t , by a fixed quantity λ_t . Indeed, such shifts do not influence the relative positions between the agents. Conversely, we prove in Appendix B.2 that all (optimal) deterministic algorithms in the centralized setting, are in fact, shifts of one another, though, we stress that shifts λ_t are not necessarily the same for all rounds t . In particular, Algorithm \mathcal{W}^* can be obtained by adding the shift $\lambda_t = \frac{1}{n}\rho_\star^{(t)} \sum_{i=1}^n Y_i^{(t)}$ to the agents in Algorithm 1 (see Appendix B.3). Figure 1 depicts the trajectory of the center of mass of the group, when \mathcal{W}^* is used, and when the “meet at the center” algorithm is used, while facing the same randomness instantiations.

Game theory. Finally, we investigate the game theoretic setting, in which each agents aims to minimize its own deviation from the center of mass, and identify the specific role played by \mathcal{W}^* also in this setting.

Definition 9. We say that a profile of algorithms $(A_i)_{i \in I}$ for the agents is a *strong Nash equilibrium* if for every $i \in I$, and for every round t , no algorithm A'_i for Agent i yields a smaller cost $\mathcal{C}_i^{(t)}$, if each other agent j keeps using A_j .

The proof of the following theorem is deferred to Appendix A.3.

Theorem 10. *Algorithm \mathcal{W}^* is a (symmetric) strong Nash-equilibrium. Moreover, if all agents are restricted to execute weighted-average algorithms, then \mathcal{W}^* is the only strong Nash equilibrium.*

1.4 Related works

Kalman filter and noisy self-organizing systems. The Kalman filter algorithm is a prolific tool commonly used in control theory and statistics, with numerous technological applications [1]. This algorithm receives a sequence of noisy measurements, and produces estimates of unknown variables, by estimating a joint probability distribution over the variables for each time step. In this paper, we used the Kalman filter to investigate the centralized setting, where all relative measurements are gathered at one master agent that processes them and instructs agents how to move. Fusing relative measurements of multiple agents is often referred to in the literature as *distributed Kalman filter*, see surveys in [4, 21, 22]. However, there, the typical setting is that each agent invokes a separate Kalman filter to process its measurements, and then the resulted outputs are integrated using communication. Moreover, works in that domain often assume that

observations are attained by static observers (i.e., the agents are residing on a fixed underlying graph) and that the measured target is external to the system. In contrast, here we consider a self-organizing system, with mobile agents that measure a target (center of mass) that is a function of their positions.

The study of *flocking* was originally based on computer simulations, rather than on rigorous analysis [7, 2, 27]. In recent years, more attention has been given to such self-organizing processes by control theoreticians [20, 18], physicists [28], and computer scientists [5]. Instead of considering all components of flocking (typically assumed to be attraction, repulsion, and alignment), here we focus on the alignment component, and the ability to reach cohesion while avoiding excessive communication.

Another related self-organization problem is *clock synchronization*, where the goal is to maintain a common notion of time in the absence of a global source of real time. The main difficulty is handling the fact that clocks count time at slightly different rates and that messages arrive with some delays. Variants of this problem were studied in wireless network contexts, mostly by control theoreticians [24, 26, 25]. A common technique uses oscillating models [19, 14]. A recent trend in the engineering community is to study the clock synchronization problem from a signal processing perspective, while adopting tools from information theory [31]. However, so far this perspective hardly received any attention by distributed computing theoreticians.

Distributed computing studies on stochastic systems with noisy communication. The problems of consensus and clock synchronization were also extensively studied in the discipline of theoretical distributed computing [8, 16, 11]. However, the corresponding works almost exclusively assume that, although the processes themselves may be faulty, the communication is nevertheless reliable, that is, not subject to noise. Moreover, the typical setting is adversarial, and, for example, very few studies in this discipline consider the clock synchronisation problem with random delays [17, 10].

In recent years, more attention has been given in the distributed computing discipline to study stochastic processes under noisy communication. Following [9], such processes were studied in [3, 6, 12], under the assumption that each message is a bit that can be flipped with some small probability. A model concerning a group of individuals that aim to estimate an external signal relying on noisy real-valued measurements was studied in [15]. Despite the differences between models, similarly to our paper, the findings in [15] emphasize the effectiveness of performing a weighted-average between the current opinion of the individual and the sample it receives. Nevertheless, we note that in the context of the model in [15], this result is less surprising since it does not involve any drift, and since each individual communicates with only one individual at a round.

2 Preliminaries: An Introduction to the Kalman filter algorithm

We follow the notations of [30]. Denote by $\mathcal{N}(\mu, \Sigma)$ the multivariate normal distribution with mean vector $\mu \in \mathbb{R}^n$ and co-variance matrix $\Sigma \in \mathbb{R}^{n \times n}$, and by I the identity matrix.

2.1 Definition of the discrete linear filtering problem

Time proceeds in discrete rounds. The problem is to estimate a vector $x_t \in \mathbb{R}^n$ at each round t . Informally, in each round t , a vector of measurements for x_t is given as input, and the output is an estimation \hat{x}_t of x_t . The vector x_t is updated by some known linear transformation that is subject to noise. Then next round starts with new measurements for the updated vector x_{t+1} , and so forth.

Formally, consider round t . Given matrices A_t, B_t, H_t, Q_t and $R_t \in \mathbb{R}^{n \times n}$, the measurement vector z_t is given by

$$z_t = H_t x_t + v_t, \tag{8}$$

where v_t is a normally distributed noise, $v_t \sim \mathcal{N}(0, R_t)$. The update at the end of the t 'th round is given by:

$$x_{t+1} = A_t x_t + B_t u_t + w_t, \tag{9}$$

where u_t is an arbitrary quantity, which is called here *move*, known by the Kalman filter, and w_t is another noise factor, called here *drift*, which is distributed by $w_t \sim \mathcal{N}(0, Q_t)$. The definition of u_t can depend on the estimation of the Kalman filter at round t . The noise vectors w_t and v_t that perturb the process and the measurements are assumed to be independent of each-other, and independent across rounds.

2.2 The Kalman filter estimator

We define \hat{x}_t the estimate of x_t *after* the measurements at round t were obtained (Eq. (8)) and before the update of round t occurs (Eq. (9)). Let P_t denote the error co-variance matrix associated with this estimate. Specifically,

$$P_t = \mathbb{E} [(x_t - \hat{x}_t)(x_t - \hat{x}_t)^\top].$$

We add the superscript “-” to these notations (for example P_t^-), to denote the same quantities *before* the measurement obtained at round t . So, in a sense, round t can be divided into the following four consecutive time steps: (a) the filter produces an estimation \hat{x}_t^- of x_t , (b) a new measurement vector of x_t is obtained, (c) the filter produces an estimation \hat{x}_t of x_t given the new measurement, and (d) x_t is updated to x_{t+1} .

Measurement update. In order to produce (c), the filter incorporates the measurement in (b) to the estimation in (a). Specifically, the filter first computes a quantity called the “Kalman gain”:

$$K_t = P_t^- H_t^\top (H_t P_t^- H_t^\top + R_t)^{-1}. \quad (10)$$

Then, it produces the following estimate, as required in step (c):

$$\hat{x}_t = \hat{x}_t^- + K_t (z_t - H_t \hat{x}_t^-). \quad (11)$$

The new error co-variance matrix can then be written as:

$$P_t = (I - K_t H_t) P_t^-. \quad (12)$$

Time update. The estimation for (a) for the following round is then given by:

$$\hat{x}_{t+1}^- = A_t \hat{x}_t + B_t u_t. \quad (13)$$

Finally, the new error co-variance matrix that will be used in the Kalman gain corresponding to the $t+1$ 'th round (Eq. (10)) is:

$$P_{t+1}^- = A_t P_t A_t^\top + Q_t. \quad (14)$$

2.3 Optimality

For the system described in Section 2.1, the Kalman filter produces two estimates at each round. We are mainly interested in the first of these estimates, \hat{x}_t^- , that is, the one corresponding to step (a), which is established before the t 'th measurement. We shall use the well-known fact that these estimates are optimal, in the sense that they minimize the expected mean square distance to x_t . To state this more formally, we need to define a general estimator for such a system.

Definition 11. Consider the system described in Section 2.1. An *estimator* at round t is given a sequence of measurements $(z_s)_{s \leq t-1}$ and a sequence of moves $(u_s)_{s \leq t-1}$, and produces an estimation \tilde{x}_t for x_t . Note that with this definition, the Kalman filter restricted to the outputs \hat{x}_t^- is an estimator.

Definition 12. An estimator at round t is said to be *optimal* if for every sequence of $t-1$ measurements $(z_s)_{s \leq t-1}$ and every sequence of $t-1$ moves $(u_s)_{s \leq t-1}$, it minimizes

$$\mathbb{E} \left(\sum_{i=1}^n (x_{t,i} - \tilde{x}_{t,i})^2 \mid (z_s, u_s)_{s \leq t-1} \right).$$

where the expectation is taken with respect to the noise in the measurements and the drifts, the initialization of positions, and the coins tosses by the estimator (in case it is probabilistic). It is said to be *optimal w.r.t the i th coordinate* if it minimizes $\mathbb{E} ((x_{t,i} - \tilde{x}_{t,i})^2 \mid (z_s, u_s)_{s \leq t-1})$.

The following is a well-known result, see e.g., [1].

Theorem 13. *For every round t , the estimator \hat{x}_t , given by the Kalman filter, is optimal. Moreover, this is the only deterministic optimal estimator.*

At this point, we note that optimality (as stated in Definition 12) implies optimality for each coordinate.

Corollary 14. *For every round t , and for every $i \in \{1, \dots, n\}$, the estimator $\hat{x}_{t,i}$, produced by the Kalman filter, is optimal w.r.t the i th coordinate. Moreover, this is the only deterministic optimal estimator w.r.t this coordinate.*

Proof. Fix $i \in \{1, \dots, n\}$. Consider an alternative estimator $\tilde{x}_{t,i} \neq \hat{x}_{t,i}$ at round t , and a sequence $(z_s, u_s)_{s \leq t-1}$ of measurements and updates. If $\mathbb{E}((x_{t,i} - \tilde{x}_{t,i})^2 \mid (z_s, u_s)_{s \leq t-1}) \leq \mathbb{E}((x_{t,i} - \hat{x}_{t,i})^2 \mid (z_s, u_s)_{s \leq t-1})$, then $\mathbb{E}((x_{t,i} - \tilde{x}_{t,i})^2 + \sum_{i=2}^n (x_{t,i} - \hat{x}_{t,i})^2 \mid (z_s, u_s)_{s \leq t-1}) \leq \mathbb{E}(\sum_{i=1}^n (x_{t,i} - \hat{x}_{t,i})^2 \mid (z_s, u_s)_{s \leq t-1})$, contradicting Theorem 13. \square

Next, we show that if one wishes to choose the move u_{t-1} in order to minimize x_t , then, whenever possible, the best choice is to set u_{t-1} such that the Kalman filter would produce $\hat{x}_t = 0$ (see Eq. (13)).

Proposition 15. *Fix $i \in \{1, \dots, n\}$ and an integer $t \geq 1$, and consider a sequence of moves $(u_s)_{s \leq t-2}$, and a sequence of measurements $(z_s)_{s \leq t-1}$. If there exists a move u_{t-1} at round $t-1$ such that the Kalman filter produces $\hat{x}_{t,i} = 0$ on input $(z_s, u_s)_{s \leq t-1}$, then for every other move u'_{t-1} ,*

$$\mathbb{E}(x_{t,i}^2 \mid (z_s)_{s \leq t-1}, (u_s)_{s \leq t-2}, u'_{t-1}) \geq \mathbb{E}(x_{t,i}^2 \mid (z_s)_{s \leq t-1}, (u_s)_{s \leq t-1})$$

with equality if and only if the corresponding estimate $\hat{x}_{t,i}'$ produced by the Kalman filter is equal to 0.

Proof. Consider a round t , and a history $\mathcal{H}_t = \{(z_s)_{s \leq t-1}, (u_s)_{s \leq t-2}\}$ of $t-1$ measurements and $t-2$ moves (if $t=1$, then we consider that $(u_s)_{s \leq t-2}$ is an empty sequence). We assume that there exists a move u_{t-1} at round $t-1$ such that the Kalman filter produces $\hat{x}_{t,i} = 0$ on input $(z_s, u_s)_{s \leq t-1}$. Consider an alternative move u'_{t-1} at round $t-1$, and denote by $\hat{x}_{t,i}'$ the estimation produced by the Kalman filter estimator on input $(\mathcal{H}_t, u'_{t-1})$. Our goal is to show that

$$\mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u'_{t-1}) \geq \mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u_{t-1}), \quad (15)$$

with equality in Eq. (15) if and only if $\hat{x}_{t,i}' = 0$.

Let c be the i th coordinate of the vector $B_t(u_{t-1} - u'_{t-1})$. By Eq. (9), c is the difference between the position of Agent i in the beginning of round t if it moves by u_{t-1} instead of moving by u'_{t-1} (conditioning on having the same drift at the end of round $t-1$). As a consequence, we have

$$\mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u'_{t-1}) = \mathbb{E}((x_{t,i} - c)^2 \mid \mathcal{H}_t, u_{t-1}), \quad (16)$$

and, by Eq. (13),

$$\hat{x}_{t,i}' = \hat{x}_{t,i} - c. \quad (17)$$

First, let us consider the case where $\hat{x}_{t,i} = \hat{x}_{t,i}' = 0$. In this case, by Eq. (17), $c = 0$, and so, by Eq. (16),

$$\mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u'_{t-1}) = \mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u_{t-1}).$$

Next, let us consider the case that $\hat{x}_{t,i}' \neq 0$. By Eq. (17), $c \neq \hat{x}_{t,i}$. We thus have

$$\begin{aligned} \mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u'_{t-1}) &= \mathbb{E}((x_{t,i} - c)^2 \mid \mathcal{H}_t, u_{t-1}) && \text{(Eq. (16))} \\ &> \mathbb{E}((x_{t,i} - \hat{x}_{t,i})^2 \mid \mathcal{H}_t, u_{t-1}) && \text{(by Corollary 14 and because } c \neq \hat{x}_{t,i}\text{)} \\ &= \mathbb{E}(x_{t,i}^2 \mid \mathcal{H}_t, u_{t-1}), && \text{(because } \hat{x}_{t,i} = 0\text{)} \end{aligned}$$

which concludes the proof. \square

3 Solving the Alignment problem

Letting $\mathbf{1}$ be the matrix whose all coefficients are equal to 1, we denote

$$\mathcal{M}(a, b) = b\mathbf{1} + (a - b)I,$$

the matrix having diagonal coefficients equal to a , and other coefficients equal to b .

3.1 Rephrasing the alignment problem as a linear filtering problem

First, we write our equations in matricial form to allow us to apply the Kalman filter straightforwardly. Let

$$\bar{\theta}^{(t)} = \left(\bar{\theta}_1^{(t)}, \dots, \bar{\theta}_n^{(t)} \right), \quad d\theta^{(t)} = \left(d\theta_1^{(t)}, \dots, d\theta_n^{(t)} \right) \quad \text{and} \quad Y^{(t)} = \left(Y_1^{(t)}, \dots, Y_n^{(t)} \right),$$

and

$$N_m^{(t)} = \left(N_{m,1}^{(t)}, \dots, N_{m,n}^{(t)} \right) \quad \text{and} \quad N_d^{(t)} = \left(N_{d,1}^{(t)}, \dots, N_{d,n}^{(t)} \right).$$

Measurement rule. We recall the equation giving the measurement of Agent i at time t :

$$Y_i^{(t)} = \bar{\theta}_i^{(t)} + N_{m,i}^{(t)}.$$

We simply rewrite this equation using vectors:

$$Y^{(t)} = \bar{\theta}^{(t)} + N_m^{(t)}, \tag{18}$$

where, by definition, $N_m^{(t)} \sim \mathcal{N}(0, \sigma_m^2 I)$.

Update rule. We recall the update equation of the stretch, which follows from Eq. (25):

$$\begin{aligned} \bar{\theta}_i^{(t+1)} &= \bar{\theta}_i^{(t)} - d\theta_i^{(t)} - N_{d,i}^{(t)} + \left(\langle \theta_{-i}^{(t+1)} \rangle - \langle \theta_{-i}^{(t)} \rangle \right) \\ &= \bar{\theta}_i^{(t)} - d\theta_i^{(t)} - N_{d,i}^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}} \left(d\theta_j^{(t)} + N_{d,j}^{(t)} \right). \end{aligned} \tag{19}$$

We define the matrix

$$M_n = \mathcal{M} \left(-1, \frac{1}{n-1} \right).$$

Let $\tilde{N}_d^{(t)} = M_n N_d^{(t)}$. It follows from these definitions and Eq.(19) that

$$\bar{\theta}^{(t+1)} = \bar{\theta}^{(t)} + M_n d\theta^{(t)} + \tilde{N}_d^{(t)}. \tag{20}$$

By definition, $N_d^{(t)} \sim \mathcal{N}(0, \sigma_d^2 I)$, so by Claim 26, $\tilde{N}_d^{(t)} \sim \mathcal{N}(0, Q)$ where

$$Q = \sigma_d^2 \cdot M_n I M_n^\top = \sigma_d^2 M_n^2.$$

At this point, we note that Equations (18) and (20) correspond to Equations (8) and (9), with $A_t = I$, $B_t = M_n$, $H_t = I$, $v_t = N_m^{(t)}$, $w_t = \tilde{N}_d^{(t)}$, $R_t = \sigma_m^2 I$ and $Q_t = Q = \sigma_d^2 M_n^2$.

Let $\hat{\theta}_t^-, \hat{\theta}_t^+$ denote the estimates of the stretch produces by the Kalman filter, before and after the measurement at round t , respectively.

Definition 16. We say that an algorithm for the alignment problem is *Kalman-perfect* if it always produces a sequence of moves $(d\theta^{(t)})_{t \geq 0}$, such that for every integer $t \geq 1$, the estimates $\hat{\theta}_t^-$ by the Kalman filter corresponding to this process is equal to 0.

The following proposition follows directly from Proposition 15.

Proposition 17. *If there exists a Kalman-perfect algorithm for the alignment problem, then this algorithm is optimal in the centralized setting (in the sense of Definition 1. Moreover, any other optimal (deterministic) algorithm is Kalman-perfect.*

3.2 Applying the Kalman filter

Recall that we denote K_t as the Kalman gain at round t , and by P_t^- and P_t the error co-variance matrices before and after the measurement at round t , respectively. Recall also that at round 0, i.e., at the initialization stage, the agents are normally distributed around 0. For technical reasons, we define the Kalman filter estimate at round 0 to be zero, that is, $\hat{\theta}_0^- = 0$.

Measurement update. In our case, the Kalman gain (Eq. (10)) writes

$$K_t = P_t^- (P_t^- + \sigma_m^2 I)^{-1}.$$

We have the following expression for the estimate (Eq. (11)):

$$\hat{\theta}_t = \hat{\theta}_t^- + K_t (Y^{(t)} - \hat{\theta}_t^-).$$

Eventually, the update equation for the error co-variance is (Eq. (12)):

$$P_t = (I - K_t) P_t^-.$$

Time update. We have (Eq. (13))

$$\hat{\theta}_{t+1}^- = \hat{\theta}_t + M_n d\theta^{(t)},$$

and (Eq. (14))

$$P_{t+1}^- = P_t + \sigma_d^2 \cdot M_n^2.$$

3.3 Computing the Kalman gain and the error co-variance matrix

Recall the sequences $(\alpha_t)_{t \geq 0}$ and $(\rho_\star^{(t)})_{t \geq 0}$ introduced in Definitions 4 and 5.

Lemma 18. For every $t \in \mathbb{N}$,

$$P_t^- = \mathcal{M} \left(\alpha_t, \frac{-\alpha_t}{n-1} \right) = -\alpha_t M_n \quad \text{and} \quad K_t = \frac{-\frac{n-1}{n} \alpha_t M_n}{\frac{n}{n-1} \sigma_m^2 + \alpha_t} = -\rho_\star^{(t)} M_n.$$

Proof. We prove the first part of the claim by induction, and prove that for every round t , the second part of the claim (regarding K_t) follows from the first part (regarding P_t^-).

By Claim 26,

$$P_0^- = \sigma_0^2 M_n^2.$$

By Claim 27,

$$M_n^2 = \mathcal{M} \left(\frac{n}{n-1}, -\frac{n}{(n-1)^2} \right) = \frac{-n}{n-1} M_n.$$

Therefore,

$$P_0^- = -\frac{n}{n-1} \sigma_0^2 M_n,$$

and so the first part of the claim holds at round 0 since $\alpha_0 = \frac{n}{n-1} \sigma_0^2$.

Now, let us assume that the first part of the claim holds for some $t \in \mathbb{N}$. It follows that

$$P_t^- + \sigma_m^2 I = \mathcal{M} \left(\alpha_t + \sigma_m^2, \frac{-\alpha_t}{n-1} \right).$$

By Claim 28, since $\sigma_m > 0$, we have

$$(P_t^- + \sigma_m^2 I)^{-1} = \frac{\mathcal{M} \left(\alpha_t + \sigma_m^2 - (n-2) \cdot \frac{\alpha_t}{n-1}, \frac{\alpha_t}{n-1} \right)}{\left(\alpha_t + \sigma_m^2 + \frac{\alpha_t}{n-1} \right) \left(\alpha_t + \sigma_m^2 - (n-1) \cdot \frac{\alpha_t}{n-1} \right)} = \frac{\mathcal{M} \left(\frac{\alpha_t}{n-1} + \sigma_m^2, \frac{\alpha_t}{n-1} \right)}{\sigma_m^2 \left(\frac{n}{n-1} \alpha_t + \sigma_m^2 \right)}.$$

By Claim 27 again, we can compute the ‘‘Kalman gain’’:

$$\begin{aligned} K_t = P_t^- (P_t^- + \sigma_m^2 I)^{-1} &= \frac{\mathcal{M} \left(\frac{\alpha_t^2}{n-1} + \sigma_m^2 \alpha_t - (n-1) \cdot \frac{\alpha_t^2}{(n-1)^2}, \frac{-\alpha_t^2}{(n-1)^2} - \frac{\sigma_m^2 \alpha_t}{n-1} + \frac{\alpha_t^2}{n-1} - (n-2) \cdot \frac{\alpha_t^2}{(n-1)^2} \right)}{\sigma_m^2 \left(\frac{n}{n-1} \alpha_t + \sigma_m^2 \right)} \\ &= \frac{\mathcal{M} \left(\sigma_m^2 \alpha_t, \frac{-\sigma_m^2 \alpha_t}{n-1} \right)}{\sigma_m^2 \left(\frac{n}{n-1} \alpha_t + \sigma_m^2 \right)} \\ &= \frac{-\alpha_t M_n}{\frac{n}{n-1} \alpha_t + \sigma_m^2}, \end{aligned}$$

This proves that the second part of the claim holds at round t .

Next, we compute the error co-variance matrix after the measurement:

$$\begin{aligned}
P_t &= (I - K_t) \cdot P_t^- = \frac{\mathcal{M}\left(\frac{\alpha_t}{n-1} + \sigma_m^2, \frac{\alpha_t}{n-1}\right)}{\frac{n}{n-1}\alpha_t + \sigma_m^2} \cdot P_t^- \\
&= \frac{\mathcal{M}\left(\frac{\alpha_t^2}{n-1} + \sigma_m^2\alpha_t - (n-1) \cdot \frac{\alpha_t^2}{(n-1)^2}, -\frac{\alpha_t^2}{(n-1)^2} - \frac{\sigma_m^2\alpha_t}{n-1} + \frac{\alpha_t^2}{n-1} - (n-2) \cdot \frac{\alpha_t^2}{(n-1)^2}\right)}{\frac{n}{n-1}\alpha_t + \sigma_m^2} \\
&= \frac{\mathcal{M}\left(\sigma_m^2\alpha_t, \frac{-\sigma_m^2\alpha_t}{n-1}\right)}{\frac{n}{n-1}\alpha_t + \sigma_m^2} \\
&= -\frac{\frac{n-1}{n}\sigma_m^2\alpha_t}{\alpha_t + \frac{n-1}{n}\sigma_m^2} M_n.
\end{aligned}$$

Eventually, we compute the error co-variance matrix at round $t+1$, before the measurement:

$$P_{t+1}^- = P_t + \sigma_d^2 M_n^2 = P_t - \frac{n}{n-1} \sigma_d^2 M_n.$$

Plugging in the expression of P_t , we get

$$\begin{aligned}
P_{t+1}^- &= -\frac{\frac{n-1}{n}\sigma_m^2\alpha_t}{\alpha_t + \frac{n-1}{n}\sigma_m^2} M_n - \frac{n}{n-1} \sigma_d^2 M_n \\
&= -\left(\frac{\frac{n-1}{n}\sigma_m^2\alpha_t}{\frac{n-1}{n}\sigma_m^2 + \alpha_t} + \frac{n}{n-1} \sigma_d^2\right) M_n \\
&= -\alpha_{t+1} M_n,
\end{aligned}$$

which concludes the induction proof. \square

3.4 Meet at the center

Our next goal is to prove that Algorithm 1 is optimal. We then explain why this algorithm is called “meet at the center”.

Theorem 19. *Algorithm 1 is optimal in the centralized setting.*

Proof. Our goal is to prove that Algorithm 1 is Kalman-perfect, that is, for every integer $t \geq 1$, the Kalman filter associated with the moves produced by Algorithm 1 gives the estimate $\hat{\theta}_t^- = 0$. This would conclude the proof of the theorem, by Proposition 17.

For this purpose, assume that all agents run Algorithm 1. We prove by induction that for every integer $t \geq 0$, the Kalman filter produces the estimate $\hat{\theta}_t^- = 0$.

The base case, where $t = 0$, holds since we assumed that the Kalman filter estimates zero at round zero, i.e., $\hat{\theta}_0^- = 0$. Next, let us assume that $\hat{\theta}_t^- = 0$ for some integer $t \geq 0$. We have by definition,

$$\hat{\theta}_{t+1}^- = \hat{\theta}_t^- + M_n d\theta^{(t)}, \tag{21}$$

and

$$\hat{\theta}_t = \hat{\theta}_t^- + K_t \left(Y^{(t)} - \hat{\theta}_t^-\right) = K_t Y^{(t)}, \tag{22}$$

where the second equality in Eq. (22) is by induction hypothesis.

By Lemma 18,

$$K_t = -\rho_\star^{(t)} M_n.$$

Note that, by definition of Algorithm 1,

$$d\theta^{(t)} = -\frac{n-1}{n} \rho_\star^{(t)} M_n Y^{(t)} = \frac{n-1}{n} \hat{\theta}_t^-.$$

Finally, by the aforementioned equations, we rewrite Eq. (21) as:

$$\hat{\theta}_{t+1} = \left(I + \frac{n-1}{n} M_n \right) \hat{\theta}_t = \left(I + \frac{n-1}{n} M_n \right) K_t Y^{(t)} = - \left(M_n + \frac{n-1}{n} M_n^2 \right) \rho_\star^{(t)} Y^{(t)}.$$

Since $M_n^2 = \frac{-n}{n-1} M_n$, we have $\hat{\theta}_{t+1} = 0$. This concludes the induction, and completes the proof of the theorem. \square

Next, we wish to explain why we refer to Algorithm 1 as the “meet at the center” algorithm. Let $i \in \{1, \dots, n\}$. Since $\hat{\theta}_{t,i}$ is produced by the Kalman filter, it is the optimal estimate of the stretch $\bar{\theta}_i^{(t)}$ of agent i , given the previous measurements and moves. Next, recall that $\langle \theta^t \rangle$ denotes the center of mass of all agents, and that $\langle \theta^t \rangle - \theta_i^{(t)} = \frac{n-1}{n} \bar{\theta}_i^{(t)}$. Therefore, $d\theta_i^{(t)} = \frac{n-1}{n} \hat{\theta}_{t,i}$ is the optimal estimate of $\langle \theta^t \rangle - \theta_i^{(t)}$ (given the previous measurements and moves). Consequently, instructing Agent i to move by $d\theta_i^{(t)}$, would make this agent be located as close as possible to the center of mass. This observation justifies that Algorithm 1 consists in “meet at the center”.

3.5 The weighted-average algorithm \mathcal{W}^\star is optimal in the centralized setting

Our next goal is to prove Theorem 7.

Theorem 7 (restated). The weighted-average algorithm \mathcal{W}^\star is optimal in the centralized setting.

Proof. (The proof follows the same line of arguments as the proof of Theorem 19.) Our goal is to prove that Algorithm \mathcal{W}^\star is Kalman-perfect, that is, for every integer $t \geq 1$, the Kalman filter associated with the moves produced by Algorithm \mathcal{W}^\star gives the estimate $\hat{\theta}_t = 0$. This would conclude the proof of the theorem, by Proposition 17.

For this purpose, assume that all agents run Algorithm \mathcal{W}^\star . We prove by induction that for every integer $t \geq 0$, the Kalman filter produces the estimate $\hat{\theta}_t = 0$.

The base case, where $t = 0$, holds since we assumed that the Kalman filter estimates zero at round zero, i.e., $\hat{\theta}_0 = 0$. Next, let us assume that $\hat{\theta}_t = 0$ for some integer $t \geq 0$, and consider $t+1$. We have by definition,

$$\hat{\theta}_{t+1} = \hat{\theta}_t + M_n d\theta^{(t)}, \tag{23}$$

and

$$\hat{\theta}_t = \hat{\theta}_t + K_t \left(Y^{(t)} - \hat{\theta}_t \right) = K_t Y^{(t)}, \tag{24}$$

where the second equality in Eq. (24) is by induction hypothesis. By Lemma 18, $K_t = -\rho_\star^{(t)} M_n$. Hence, Eq. (23) rewrites

$$\hat{\theta}_{t+1} = M_n \left(-\rho_\star^{(t)} Y^{(t)} + d\theta^{(t)} \right).$$

Finally, by the definition of \mathcal{W}^\star , we have $d\theta^{(t)} = \rho_\star^{(t)} Y^{(t)}$, so $\hat{\theta}_{t+1} = 0$, concluding the induction step. \square

4 Discussion and Future Work

In our setting, the cost of an individual is defined as its expected distance from the average position of other agents at steady state. Minimizing this quantity is equivalent to minimizing the expected distance from the average position of all agents (see Footnote 2). Another interesting measure is the expected *diameter* of the group, defined as the maximal distance between two agents, at steady state. It would not come as a surprise if Algorithm \mathcal{W}^\star would turn out to be optimal also with respect to this measure, however, analyzing its expected diameter would require handling further dependencies between agents, and therefore remains for future work.

Finally, we conclude with a philosophical remark concerning the dichotomy between conformity and individuality. When individuals have a priory conflicted preferences it is natural to assume that the individual responsiveness to the group would be moderated [29]. Without such preferences, when the goal is to purely

conform, existing models in collective behavior typically assume that whenever an individual receives a measurement of the group’s average it tries to align itself with it as much as it can [13, 27]. However, we show here that due to noise, each individual should actually moderate its social responsiveness, by weighing its current direction in a non-trivial manner. This insight suggests that the dichotomy between conformity and individuality (manifested here as persistency) might be more subtle than commonly perceived.

Acknowledgment. We would like to thank Yongcan Cao for helpful discussion regarding related work on distributed Kalman filter.

References

- [1] Brian DO Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2012.
- [2] Ichiro Aoki. A simulation study on the schooling mechanism in fish. *NIPPON SUISAN GAKKAISHI*, 48(8):1081–1088, 1982.
- [3] Lucas Boczkowski, Ofer Feinerman, Amos Korman, and Emanuele Natale. Limits for rumor spreading in stochastic populations. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 49:1–49:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [4] Yongcan Cao, Wenwu Yu, Wei Ren, and Guanrong Chen. An overview of recent progress in the study of distributed multi-agent coordination. *IEEE Transactions on Industrial informatics*, 9(1):427–438, 2012.
- [5] Bernard Chazelle. Natural algorithms. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 422–431. SIAM, 2009.
- [6] Andrea E. F. Clementi, Luciano Gualà, Emanuele Natale, Francesco Pasquale, Giacomo Scornavacca, and Luca Trevisan. Consensus vs broadcast, with and without noise (extended abstract). In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 42:1–42:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [7] Iain D Couzin, Jens Krause, Richard James, Graeme D Ruxton, and Nigel R Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.
- [8] Rui Fan and Nancy A. Lynch. Gradient clock synchronization. *Distributed Comput.*, 18(4):255–266, 2006.
- [9] Ofer Feinerman, Bernhard Haeupler, and Amos Korman. Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication. *Distributed Comput.*, 30(5):339–355, 2017.
- [10] Ofer Feinerman and Amos Korman. Clock synchronization and estimation in highly dynamic networks: An information theoretic approach. In Christian Scheideler, editor, *Structural Information and Communication Complexity - 22nd International Colloquium, SIROCCO 2015, Montserrat, Spain, July 14-16, 2015, Post-Proceedings*, volume 9439 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2015.
- [11] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382, 1985.
- [12] Pierre Fraigniaud and Emanuele Natale. Noisy rumor spreading and plurality consensus. *Distributed Comput.*, 32(4):257–276, 2019.

- [13] Aviram Gelblum, Itai Pinkoviezky, Ehud Fonio, Abhijit Ghosh, Nir Gov, and Ofer Feinerman. Ant groups optimally amplify the effect of transiently informed individuals. *Nature communications*, 6:7729, 2015.
- [14] Yao-Win Hong and Anna Scaglione. A scalable synchronization protocol for large scale sensor networks and its applications. *IEEE Journal on Selected Areas in Communications*, 23(5):1085–1099, 2005.
- [15] Amos Korman, Efrat Greenwald, and Ofer Feinerman. Confidence sharing: An economic strategy for efficient information flows in animal groups. *PLoS Computational Biology*, 10(10), 2014.
- [16] Christoph Lenzen, Thomas Locher, and Roger Wattenhofer. Tight Bounds for Clock Synchronization. In *Journal of the ACM, Volume 57, Number 2, New York, NY, USA*, January 2010.
- [17] Christoph Lenzen, Philipp Sommer, and Roger Wattenhofer. Pulsesync: An efficient and scalable clock synchronization protocol. *IEEE/ACM Trans. Netw.*, 23(3):717–727, 2015.
- [18] Shukai Li, Xinzhi Liu, Wansheng Tang, and Jianxiong Zhang. Flocking of multi-agents following a leader with adaptive protocol in a noisy environment. *Asian Journal of Control*, 16(6):1771–1778, 2014.
- [19] Renato E Mirollo and Steven H Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662, 1990.
- [20] Reza Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on automatic control*, 51(3):401–420, 2006.
- [21] Reza Olfati-Saber. Distributed kalman filtering for sensor networks. In *2007 46th IEEE Conference on Decision and Control*, pages 5492–5498. IEEE, 2007.
- [22] Wei Ren, Randal W Beard, and Ella M Atkins. A survey of consensus problems in multi-agent coordination. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 1859–1864. IEEE, 2005.
- [23] Craig W Reynolds. *Flocks, herds and schools: A distributed behavioral model*, volume 21. ACM, 1987.
- [24] Osvaldo Simeone, Umberto Spagnolini, Yeheskel Bar-Ness, and Steven H Strogatz. Distributed synchronization in wireless networks. *IEEE Signal Processing Magazine*, 25(5):81–97, 2008.
- [25] Fikret Sivrikaya and Bülent Yener. Time synchronization in sensor networks: a survey. *IEEE network*, 18(4):45–50, 2004.
- [26] Bharath Sundararaman, Ugo Buy, and Ajay D Kshemkalyani. Clock synchronization for wireless sensor networks: a survey. *Ad hoc networks*, 3(3):281–323, 2005.
- [27] Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.*, 75:1226–1229, Aug 1995.
- [28] Tamás Vicsek and Anna Zafeiris. Collective motion. *Physics Reports*, 517(3):71 – 140, 2012. Collective motion.
- [29] Ashley JW Ward, James E Herbert-Read, David JT Sumpter, and Jens Krause. Fast and accurate decisions through collective vigilance in fish shoals. *Proceedings of the National Academy of Sciences*, 108(6):2312–2315, 2011.
- [30] Greg Welch, Gary Bishop, et al. An introduction to the kalman filter, 1995.
- [31] Yik-Chung Wu, Qasim Chaudhari, and Erchin Serpedin. Clock synchronization of wireless sensor networks. *IEEE Signal Processing Magazine*, 28(1):124–138, 2010.

A More Results Regarding Weighted-average Algorithms

A.1 An optimal weighted-average algorithm (at steady state)

Our goal in this section is to prove Theorems 2 and 3. We start with the following observation.

Lemma 20. *For any t , and whatever the positions of the agents, we have*

$$\sum_{i=1}^n \bar{\theta}_i^{(t)} = 0.$$

Proof.

$$\sum_{i=1}^n \bar{\theta}_i^{(t)} = \sum_{i=1}^n \left(\langle \theta_{-i}^{(t)} \rangle - \theta_i^{(t)} \right) = \frac{1}{n-1} \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \theta_i^{(t)} - \sum_{i=1}^n \theta_i^{(t)} = \frac{1}{n-1} \cdot (n-1) \sum_{i=1}^n \theta_i^{(t)} - \sum_{i=1}^n \theta_i^{(t)} = 0.$$

□

Next, we compute how the stretch $\bar{\theta}_i^{(t)}$ of Agent i changes when all agents perform weighted-average moves with constant responsiveness parameter ρ .

Lemma 21. *Assume that all agents execute $\mathcal{W}(\rho)$, for some $0 \leq \rho \leq 1$. Let*

$$\mathcal{E}_j^{(t)} = \rho N_{m,j}^{(t)} + N_{d,j}^{(t)}.$$

Then for every $i \in \{1, \dots, n\}$ and every $t \in \mathbb{N}$,

$$\bar{\theta}_i^{(t+1)} = \left(1 - \frac{n}{n-1} \rho \right) \bar{\theta}_i^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \mathcal{E}_j^{(t)} - \mathcal{E}_i^{(t)}.$$

Proof. The stretch of Agent i at the $t+1$ st round is given by:

$$\begin{aligned} \bar{\theta}_i^{(t+1)} &= \langle \theta_{-i}^{(t+1)} \rangle - \theta_i^{(t+1)} && \text{by definition} \\ &= \langle \theta_{-i}^{(t+1)} \rangle - \left(\theta_i^{(t)} + d\theta_i^{(t)} + N_{d,i}^{(t)} \right) && \text{by (2)} \\ &= \left(\langle \theta_{-i}^{(t+1)} \rangle - \langle \theta_{-i}^{(t)} \rangle \right) + \left(\langle \theta_{-i}^{(t)} \rangle - \theta_i^{(t)} \right) - d\theta_i^{(t)} - N_{d,i}^{(t)} \\ &= \left(\langle \theta_{-i}^{(t+1)} \rangle - \langle \theta_{-i}^{(t)} \rangle \right) + \bar{\theta}_i^{(t)} - d\theta_i^{(t)} - N_{d,i}^{(t)}. \end{aligned} \tag{25}$$

Let us break down the first term:

$$\langle \theta_{-i}^{(t+1)} \rangle - \langle \theta_{-i}^{(t)} \rangle = \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \left(d\theta_j^{(t)} + N_{d,j}^{(t)} \right) = \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \left(\rho \left(\bar{\theta}_j^{(t)} + N_{m,j}^{(t)} \right) + N_{d,j}^{(t)} \right),$$

where the second equality is because $d\theta_j^{(t)} = \rho Y_j^{(t)} = \rho \left(\bar{\theta}_j^{(t)} + N_{m,j}^{(t)} \right)$. By Lemma 20, $\sum_{\substack{j=1 \\ j \neq i}}^n \bar{\theta}_j^{(t)} = -\bar{\theta}_i^{(t)}$, so we can rewrite the last equation as

$$\langle \theta_{-i}^{(t+1)} \rangle - \langle \theta_{-i}^{(t)} \rangle = \frac{-\rho}{n-1} \bar{\theta}_i^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \left(\rho N_{m,j}^{(t)} + N_{d,j}^{(t)} \right) = \frac{-\rho}{n-1} \bar{\theta}_i^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \mathcal{E}_j^{(t)}. \tag{26}$$

Plugging Eq. (26) into Eq. (25) gives

$$\bar{\theta}_i^{(t+1)} = \left(1 - \frac{\rho}{n-1} \right) \bar{\theta}_i^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \mathcal{E}_j^{(t)} - d\theta_i^{(t)} - N_{d,i}^{(t)}. \tag{27}$$

By assumption, $d\theta_i^{(t)} = \rho Y_i^{(t)} = \rho \left(\bar{\theta}_i^{(t)} + N_{m,i}^{(t)} \right)$. Plugging this expression into Eq. (27) gives the result. □

Now we can prove that the stretch of each agent is normally distributed at every round, and compute its variance.

Lemma 22. *Assume that all agents execute $\mathcal{W}(\rho)$, for some $0 \leq \rho \leq 1$. Then, for every $i \in \{1, \dots, n\}$ and every $t \in \mathbb{N}$, the stretch $\bar{\theta}_i^{(t)}$ is normally distributed. Moreover, $\mathbb{E}(\bar{\theta}_i^{(t)}) = 0$, and*

$$\text{Var}(\bar{\theta}_i^{(t+1)}) = \left(1 - \frac{n}{n-1}\rho\right)^2 \text{Var}(\bar{\theta}_i^{(t)}) + \frac{n}{n-1}(\rho^2\sigma_m^2 + \sigma_d^2).$$

Proof. We prove that the stretch is normally distributed with mean 0 by induction on t . By construction, for every i , $\bar{\theta}_i^{(0)}$ is normally distributed with mean 0. Let us assume that $\bar{\theta}_i^{(t)}$ is normally distributed with mean 0 for some round t , and consider round $t+1$. Recall that Lemma 21 gives

$$\bar{\theta}_i^{(t+1)} = \left(1 - \frac{n}{n-1}\rho\right) \bar{\theta}_i^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \mathcal{E}_j^{(t)} - \mathcal{E}_i^{(t)}. \quad (28)$$

Since by definition, for every j , $\mathcal{E}_j^{(t)}$ is normally distributed around 0, and by induction $\bar{\theta}_i^{(t)}$ is normally distributed around 0, then $\bar{\theta}_i^{(t+1)}$ is also normally distributed around 0. This concludes the induction.

Moreover, note that $\text{Var}(\mathcal{E}_j^{(t)}) = \rho^2\sigma_m^2 + \sigma_d^2$, so

$$\text{Var}\left(\frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \mathcal{E}_j^{(t)} - \mathcal{E}_i^{(t)}\right) = \frac{1}{(n-1)^2} \sum_{\substack{j=1 \\ j \neq i}}^n \text{Var}(\mathcal{E}_j^{(t)}) + \text{Var}(\mathcal{E}_i^{(t)}) = \frac{n}{n-1}(\rho^2\sigma_m^2 + \sigma_d^2).$$

Hence, by Eq. (28),

$$\text{Var}(\bar{\theta}_i^{(t+1)}) = \left(1 - \frac{n}{n-1}\rho\right)^2 \text{Var}(\bar{\theta}_i^{(t)}) + \frac{n}{n-1}(\rho^2\sigma_m^2 + \sigma_d^2),$$

which concludes the proof. \square

Before proving Theorem 2, we need a small technical result, which we prove next for the sake of completeness.

Claim 23. *Let $a, b \geq 0$. Consider the sequence $\{u_n\}_{n=0}^\infty$ defined by letting $u_0 \in \mathbb{R}$ and for every integer n , $u_{n+1} = au_n + b$. If $a < 1$, then $\{u_n\}_{n=0}^\infty$ converges and $\lim_{n \rightarrow +\infty} u_n = b/(1-a)$. If $a = 1$, and $b > 0$, then $\lim_{n \rightarrow +\infty} u_n = +\infty$.*

Proof. First, consider the case that $a < 1$. Let $\lambda = b/(1-a)$. Consider the sequence defined by $v_n = u_n + \lambda$. We have

$$v_{n+1} = u_{n+1} + \lambda = au_n + b + \lambda = au_n + (a-1)\lambda + \lambda = au_n + a\lambda = a(u_n + \lambda) = av_n.$$

Since $0 \leq a < 1$, $\lim_{n \rightarrow +\infty} v_n = 0$, and so $\lim_{n \rightarrow +\infty} u_n = -\lambda = b/(1-a)$.

Now, if $a = 1$, then we have for every $n \in \mathbb{N}$, $u_n = u_0 + nb$. If $b > 0$ then $\lim_{n \rightarrow +\infty} u_n = +\infty$. \square

The next theorem follows directly from Lemma 22 and Claim 23.

Theorem 2 (restated). *Assume that all agents execute $\mathcal{W}(\rho)$, for a fixed $0 \leq \rho \leq 1$. Then for every $i \in \{1, \dots, n\}$ and every $t \in \mathbb{N}$, the stretch $\bar{\theta}_i^{(t)}$ is normally distributed, and*

$$\lim_{t \rightarrow +\infty} \text{Var}(\bar{\theta}_i^{(t)}) = \frac{\frac{n}{n-1}(\rho^2\sigma_m^2 + \sigma_d^2)}{1 - (1 - \frac{n}{n-1}\rho)^2},$$

with the convention that $\lim_{t \rightarrow +\infty} \text{Var}(\bar{\theta}_i^{(t)}) = +\infty$ if the denominator $1 - (1 - \frac{n}{n-1}\rho)^2 = 0$.

Proof. We apply Lemma 22. Hence, by Claim 23, $\text{Var}(\bar{\theta}_i^{(t)})$ converges to the limit as stated. Note that the variance is infinite if 1) the responsiveness is equal to 0, in which case the drift adds up endlessly, or 2) the responsiveness is equal to 1 and $n = 2$, in which case the agents “swap” at each round, producing the same result. \square

Theorem 3 (restated). The weighted-average algorithm that optimizes group variance among all weighted-average algorithms $\mathcal{W}(\rho)$ (that use the same responsiveness parameter ρ at all rounds) is $\mathcal{W}(\rho^*)$, where

$$\rho^* = \frac{\sigma_d \sqrt{4\sigma_m^2 + \left(\frac{n}{n-1}\sigma_d\right)^2} - \frac{n}{n-1}\sigma_d^2}{2\sigma_m^2}.$$

Proof. Consider the function

$$\text{Var}(\rho) = \frac{\frac{n}{n-1}(\rho^2\sigma_m^2 + \sigma_d^2)}{1 - \left(1 - \frac{n}{n-1}\rho\right)^2}.$$

Note that this function evaluates to $+\infty$ when $\rho = 0$, or when $\rho = 1$ and $n = 2$. This function can be rewritten as

$$\text{Var}(\rho) = \frac{\rho^2\sigma_m^2 + \sigma_d^2}{2\rho - \frac{n}{n-1}\rho^2}.$$

One can compute the derivative in a straightforward manner:

$$\text{Var}'(\rho) = \frac{2\rho\sigma_m^2 \left(2\rho - \frac{n}{n-1}\rho^2\right) - (\rho^2\sigma_m^2 + \sigma_d^2) \left(2 - 2\frac{n}{n-1}\rho\right)}{\left(2\rho - \frac{n}{n-1}\rho^2\right)^2}.$$

We have that

$$\begin{aligned} \text{Var}'(\rho) = 0 &\iff 2\rho\sigma_m^2 \left(2\rho - \frac{n}{n-1}\rho^2\right) - (\rho^2\sigma_m^2 + \sigma_d^2) \left(2 - 2\frac{n}{n-1}\rho\right) = 0 \\ &\iff \rho^2\sigma_m^2 \left(2 - \frac{n}{n-1}\rho\right) - (\rho^2\sigma_m^2 + \sigma_d^2) \left(1 - \frac{n}{n-1}\rho\right) = 0 \\ &\iff \rho^2\sigma_m^2 - \sigma_d^2 \left(1 - \frac{n}{n-1}\rho\right) = 0. \end{aligned}$$

This equation has a unique solution in the interval $[0, 1]$:

$$\text{Var}'(\rho) = 0, \rho \in [0, 1] \iff \rho = \frac{\sigma_d \sqrt{4\sigma_m^2 + \left(\frac{n}{n-1}\sigma_d\right)^2} - \frac{n}{n-1}\sigma_d^2}{2\sigma_m^2}.$$

We check that this corresponds to a minimum to conclude the proof. \square

A.2 The asymptotic behavior of \mathcal{W}^*

The goal of this section is to prove Claim 8.

Claim 8 (restated). The sequence α_t converges to $\alpha_\infty := \lim_{t \rightarrow +\infty} \alpha_t = \frac{1}{2} \left(\sigma_d \sqrt{4\sigma_m^2 + \left(\frac{n}{n-1}\sigma_d\right)^2} + \frac{n}{n-1}\sigma_d^2 \right)$.

Moreover, $\lim_{t \rightarrow +\infty} \rho_\star^{(t)} = \rho_\star$.

Proof. For $a, b > 0$, define $f_{a,b} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that $f_{a,b}(x) = a \frac{x}{x+a} + b$. Solving $f_{a,b}(\ell) = \ell$ on \mathbb{R}^+ gives $\ell = \frac{1}{2} \left(\sqrt{b} \sqrt{4a+b} + b \right)$. For every $x \in \mathbb{R}^+$:

$$f_{a,b}(x) - f_{a,b}(\ell) = a \left(\frac{x}{x+a} - \frac{\ell}{\ell+a} \right) = a \cdot \frac{x(\ell+a) - \ell(x+a)}{(x+a)(\ell+a)} = a^2 \frac{x-\ell}{(x+a)(\ell+a)}.$$

Thus, and since $x \geq 0$,

$$|f_{a,b}(x) - f_{a,b}(\ell)| = \left| a^2 \frac{x - \ell}{(x + a)(\ell + a)} \right| \leq \frac{a^2}{a(a + \ell)} |x - \ell| = \frac{a}{a + \ell} |x - \ell|. \quad (29)$$

Claim 24. Let $(u_t)_{t \in \mathbb{N}}$ be a sequence defined by $u_0 \in \mathbb{R}^+$ and for every integer t , $u_{t+1} = f_{a,b}(u_t)$. Then, (u_t) converges, and $\lim_{t \rightarrow \infty} u_t = \ell$.

Proof. Let $k = \frac{a}{a + \ell}$. Since $a, b > 0$ then $\ell > 0$, and so $k < 1$. Let us show by induction that for every t , $|u_t - \ell| \leq k^t \cdot |u_0 - \ell|$. This equality is trivial for $t = 0$. Assuming that it holds for some $t \in \mathbb{N}$, we have

$$\begin{aligned} |u_{t+1} - \ell| &= |f_{a,b}(u_t) - f_{a,b}(\ell)| \quad (\text{by definition of } u_t \text{ and } \ell) \\ &\leq k \cdot |u_t - \ell| \quad (\text{by Eq.(29)}) \\ &\leq k^{t+1} \cdot |u_0 - \ell|, \quad (\text{by induction hypothesis}) \end{aligned}$$

concluding the induction. Since $k < 1$, it implies that $\lim_{t \rightarrow \infty} |u_t - \ell| = 0$, and so $\lim_{t \rightarrow \infty} u_t = \ell$. \square

Applying Claim 24 to $(\alpha_t)_{t \in \mathbb{N}}$ with $a = \frac{n-1}{n} \sigma_m^2$ and $b = \frac{n}{n-1} \sigma_d^2$ gives $\lim_{t \rightarrow \infty} \alpha_t = \alpha_\infty$, as stated.

Next, we show that $\lim_{t \rightarrow \infty} \rho_\star^{(t)} = \rho_\star$. By letting t tend to $+\infty$ in Definition 4, we obtain

$$\alpha_\infty = \frac{\frac{n-1}{n} \sigma_m^2 \alpha_\infty}{\frac{n-1}{n} \sigma_m^2 + \alpha_\infty} + \frac{n}{n-1} \sigma_d^2. \quad (30)$$

Doing the same in Definition 5, we get

$$\lim_{t \rightarrow +\infty} \rho_\star^{(t)} = \frac{\frac{n-1}{n} \alpha_\infty}{\frac{n-1}{n} \sigma_m^2 + \alpha_\infty} = \frac{1}{\sigma_m^2} \frac{\frac{n-1}{n} \sigma_m^2 \alpha_\infty}{\frac{n-1}{n} \sigma_m^2 + \alpha_\infty}.$$

By Eq. (30), this gives

$$\lim_{t \rightarrow +\infty} \rho_\star^{(t)} = \frac{1}{\sigma_m^2} \left(\alpha_\infty - \frac{n}{n-1} \sigma_d^2 \right).$$

Plugging in the expression of α_∞ mentioned in the Lemma, we find that

$$\lim_{t \rightarrow +\infty} \rho_\star^{(t)} = \frac{\sigma_d \sqrt{4\sigma_m^2 + \left(\frac{n}{n-1} \sigma_d \right)^2} - \frac{n}{n-1} \sigma_d^2}{2\sigma_m^2} = \rho_\star,$$

which establishes the proof. \square

A.3 Game-theoretic considerations

Theorem 10 (restated). Algorithm \mathcal{W}^* is a (symmetric) strong Nash-equilibrium. Moreover, if all agents are restricted to execute weighted-average algorithms, then \mathcal{W}^* is the only strong Nash equilibrium.

Proof. The fact that Algorithm \mathcal{W}^* is a (symmetric) strong Nash-equilibrium is a direct consequence of Theorem 7 and our definition of optimality (Definition 1).

Let us now prove the uniqueness result. Consider the case that the agents all use some weighted-average algorithm $\mathcal{W}(\rho^{(t)})$, and that this is a strong Nash equilibrium. We shall show that it implies that, for every round t , $\rho^{(t)} = \rho_\star^{(t)}$.

Fix $i \in I$. First, let us investigate the best response for Agent i to the behavior of others. Let

$$\tilde{N}^{(t)} = -N_{d,i}^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \left(\rho^{(t)} N_{m,j}^{(t)} + N_{d,j}^{(t)} \right).$$

By Lemma 21,

$$\begin{aligned}\bar{\theta}_i^{(t+1)} &= \left(1 - \frac{\rho^{(t)}}{n-1}\right) \bar{\theta}_i^{(t)} + \frac{1}{n-1} \sum_{\substack{j=1 \\ j \neq i}}^n \left(\rho^{(t)} N_{m,j}^{(t)} + N_{d,j}^{(t)}\right) - N_{d,i}^{(t)} - d\theta_i^{(t)} \\ &= \left(1 - \frac{\rho^{(t)}}{n-1}\right) \bar{\theta}_i^{(t)} + \tilde{N}^{(t)} - d\theta_i^{(t)}.\end{aligned}$$

Note that $\tilde{N}^{(t)}$ is normally distributed, with mean 0, and that

$$\text{Var}\left(\tilde{N}^{(t)}\right) = \frac{1}{n-1} \left(\rho^{(t)^2} \sigma_m^2 + \sigma_d^2\right) + \sigma_d^2.$$

These equations show that the evolution of the stretch of Agent i can be written as in Eq. (9), with the u_t variables being the moves $d\theta_i^{(t)}$ of the agent. More precisely, in this case, we have $A_t = \left(1 - \frac{\rho^{(t)}}{n-1}\right)$, $B_t = -1$, $Q_t = \frac{1}{n-1} \left(\rho^{(t)^2} \sigma_m^2 + \sigma_d^2\right) + \sigma_d^2$, $H_t = 1$, and $R_t = \sigma_m^2$. (Note that all these are scalars, since the filtering problem corresponding to Agent i is one dimensional.) This implies that this agent faces a discrete linear filtering problem, that we tackle using the Kalman filter algorithm.

For every round t , and for every sequence of moves $(d\theta_i^{(s)})_{s \leq t}$, we can compute the variance of the Kalman filter estimate, before and after the t 'th measurement, as well as the corresponding Kalman gain (see Section 2.2). By definition of the Alignment problem, we have

$$P_0^- = n\sigma_0^2/(n-1).$$

Following the definitions in Section 2.2, we have

$$K_t = \frac{P_t^-}{P_t^- + \sigma_m^2}, \quad (31)$$

and

$$P_{t+1}^- = \left(1 - \frac{\rho^{(t)}}{n-1}\right)^2 \frac{P_t^- \sigma_m^2}{P_t^- + \sigma_m^2} + \frac{1}{n-1} \left(\rho^{(t)^2} \sigma_m^2 + \sigma_d^2\right) + \sigma_d^2. \quad (32)$$

Claim 25. *Fix i , and assume that all agents $j \neq i$ execute $\mathcal{W}(\rho^{(t)})$. The Kalman filter corresponding to $\bar{\theta}_i^{(t)}$ produces $\hat{\theta}_i = 0$ for every round t , if and only if Agent i chooses $d\theta_i^{(t)} = \left(1 - \frac{\rho^{(t)}}{n-1}\right) \frac{P_t^-}{P_t^- + \sigma_m^2} Y_i^{(t)}$ for every round t .*

Proof. We prove the claim by induction on t . Specifically, our goal is to prove that for every round t , the Kalman filter produces $\hat{\theta}_s = 0$ for every round $s \leq t$, if and only if, Agent i chooses

$$d\theta_i^{(s)} = \left(1 - \frac{\rho^{(s)}}{n-1}\right) \frac{P_s^-}{P_s^- + \sigma_m^2} Y_i^{(s)}$$

for every round $s \leq t-1$. This is trivially true for $t=0$, since, it is assumed that the Kalman filter always produces $\hat{\theta}_0 = 0$. Now, let us assume that this holds for some round t , and consider round $t+1$.

As a consequence of the definition of the Kalman filter, by plugging in Eq. (11) in Eq. (13), we have

$$\hat{\theta}_{t+1} = \left(1 - \frac{\rho^{(t)}}{n-1}\right) \left(\hat{\theta}_t + K_t \left(Y_i^{(t)} - \hat{\theta}_t\right)\right) - d\theta_i^{(t)}. \quad (33)$$

The Kalman filter produces $\hat{\theta}_s = 0$ for every round $s \leq t+1$, if and only if, (1) it produces $\hat{\theta}_s = 0$ for every round $s \leq t$ and (2) it produces $\hat{\theta}_{t+1} = 0$. By Eq. (33), this occurs if and only if (1) it produces $\hat{\theta}_s = 0$ for every round $s \leq t$ and (2)

$$\left(1 - \frac{\rho^{(t)}}{n-1}\right) K_t Y_i^{(t)} - d\theta_i^{(t)} = 0. \quad (34)$$

By the induction hypothesis, and the computation of the Kalman gain K_t in Eq. (31), these two conditions hold if and only if Agent i chooses $d\theta_i^{(s)} = \left(1 - \frac{\rho^{(s)}}{n-1}\right) \frac{P_s^-}{P_s^- + \sigma_m^2} Y_i^{(s)}$ for every round $s \leq t$, which concludes the induction proof. \square

By the assumption that this is a strong Nash equilibrium, Algorithm $\mathcal{W}(\rho^{(t)})$ is a best response for Agent i . According to Claim 25, and by proposition 15 this implies that

$$d\theta_i^{(t)} = \left(1 - \frac{\rho^{(t)}}{n-1}\right) \frac{P_t^-}{P_t^- + \sigma_m^2} Y_i^{(t)} \quad (35)$$

for every round t . Because Agent i was assumed to follow Algorithm $\mathcal{W}(\rho^{(t)})$, $d\theta_i^{(t)} = \rho^{(t)} Y_i^{(t)}$ for every round t . Therefore, Eq. (35) rewrites

$$\rho^{(t)} = \left(1 - \frac{\rho^{(t)}}{n-1}\right) \frac{P_t^-}{P_t^- + \sigma_m^2}, \quad (36)$$

which, by rearranging yields:

$$\rho^{(t)} = \frac{P_t^-}{\frac{n}{n-1} P_t^- + \sigma_m^2}. \quad (37)$$

We can use Eq. (36) to simplify the first term in Eq. (32), to get

$$P_{t+1}^- = \left(1 - \frac{\rho^{(t)}}{n-1}\right) \rho^{(t)} \sigma_m^2 + \frac{1}{n-1} \left(\rho^{(t)2} \sigma_m^2 + \sigma_d^2\right) + \sigma_d^2 = \rho^{(t)} \sigma_m^2 + \frac{n}{n-1} \sigma_d^2.$$

Replacing $\rho^{(t)}$ in the last equation by its expression in Eq. (37) gives

$$P_{t+1}^- = \frac{\sigma_m^2 P_t^-}{\frac{n}{n-1} P_t^- + \sigma_m^2} + \frac{n}{n-1} \sigma_d^2. \quad (38)$$

Note that $P_0^- = \alpha_0 = n\sigma_0^2/(n-1)$. Eq. (38) has the same recursion rule as in Definition 4, so we actually have for every round t , $P_t^- = \alpha_t$. Moreover, Eq. (37) matches Definition 5, so we actually have for every round t , $\rho^{(t)} = \rho_\star^{(t)}$, which concludes the proof. \square

B More proofs related to the Centralized Setting

B.1 Useful linear algebra claims

We first recall the following well-known property,

Claim 26. *If $X \sim \mathcal{N}(\mu, \Sigma)$, then for every $c \in \mathbb{R}^n$ and $B \in \mathbb{R}^{n \times n}$,*

$$c + BX \sim \mathcal{N}(c + B\mu, B\Sigma B^\top).$$

In addition, we give two useful results about matrices of the form $\mathcal{M}(a, b)$.

Claim 27. *For every $a, b, a', b' \in \mathbb{R}$,*

$$\mathcal{M}(a, b)\mathcal{M}(a', b') = \mathcal{M}(aa' + (n-1)bb', ab' + a'b + (n-2)bb').$$

In particular, $\mathcal{M}(a, b)\mathcal{M}(a', b') = \mathcal{M}(a', b')\mathcal{M}(a, b)$.

Claim 28. *For every $a, b \in \mathbb{R}$ such that $a \neq b$ and $a \neq -(n-1)b$, the matrix $\mathcal{M}(a, b)$ is invertible, and*

$$\mathcal{M}(a, b)^{-1} = \frac{\mathcal{M}(a + (n-2)b, -b)}{(a-b)(a + (n-1)b)}.$$

Proof. Note that $\mathbf{1}^2 = n\mathbf{1}$. Let $A = \mathcal{M}(a, b)$. We have

$$\begin{aligned}
A^2 &= (b\mathbf{1} + (a-b)I)^2 \\
&= b^2\mathbf{1}^2 + 2b(a-b)\mathbf{1} + (a-b)^2I \\
&= (nb + 2(a-b))(b\mathbf{1}) + (a-b)^2I \\
&= (2a + (n-2)b)(b\mathbf{1} + (a-b)I) - (nb + 2(a-b))(a-b)I + (a-b)^2I \\
&= (2a + (n-2)b)A + (a-b)(-nb - 2(a-b) + (a-b))I \\
&= (2a + (n-2)b)A - (a-b)(a + (n-1)b)I.
\end{aligned}$$

Hence

$$A(A - (2a + (n-2)b)I) = -(a-b)(a + (n-1)b)I,$$

from which we conclude (provided that $a \neq b$ and $a \neq (n-1)b$),

$$A^{-1} = \frac{(2a + (n-2)b)I - A}{(a-b)(a + (n-1)b)}.$$

□

B.2 All optimal algorithms are shifts of one another

In this section, we characterize all optimal (deterministic) algorithms for the Alignment problem in the centralized setting. We show that each of these algorithms can be obtained from \mathcal{W}^* , by shifting all the agents by the same quantity λ_t , though we stress that shifts λ_t are not necessarily the same for all rounds t .

Theorem 29. *A deterministic algorithm is optimal in the centralized setting if and only if for every round t , there exists λ_t such that for every $i \in \{1, \dots, n\}$, $d\theta_i^{(t)} = \rho_\star^{(t)}Y_i^{(t)} + \lambda_t$.*

Proof. We have already established that the (deterministic) weighted-average algorithm \mathcal{W}^* is a Kalman-perfect algorithm. Therefore, by Proposition 17, any other deterministic algorithm is optimal in the centralized setting if and only if it is Kalman-perfect. In other words, it is optimal if and only if it produces a sequence of moves such that for every round t , the Kalman-filter estimator operating on the corresponding process yields

$$\begin{aligned}
\hat{\theta}_{t+1}^- = 0 &\iff \hat{\theta}_t + M_n d\theta^{(t)} = 0 \\
&\iff K_t Y^{(t)} + M_n d\theta^{(t)} = 0 \\
&\iff -\rho_\star^{(t)} M_n Y^{(t)} + M_n d\theta^{(t)} = 0 \\
&\iff M_n \left(-\rho_\star^{(t)} Y^{(t)} + d\theta^{(t)} \right) = 0 \\
&\iff -\rho_\star^{(t)} Y^{(t)} + d\theta^{(t)} \in \ker(M_n).
\end{aligned}$$

Writing $\mathbf{1}$ to denote the vector whose coefficients are all equal to 1, we observe that $\mathbf{1} \in \ker(M_n)$. Since $\text{rank}(M_n) = n-1$, $\dim(\ker(M_n)) = 1$, so for every round t ,

$$\begin{aligned}
-\rho_\star^{(t)} Y^{(t)} + d\theta^{(t)} \in \ker(M_n) &\iff \exists \lambda_t \in \mathbb{R}, \quad -\rho_\star^{(t)} Y^{(t)} + d\theta^{(t)} = \lambda_t \cdot \mathbf{1} \\
&\iff \exists \lambda_t \in \mathbb{R}, \quad d\theta^{(t)} = \lambda_t \cdot \mathbf{1} + \rho_\star^{(t)} Y^{(t)}.
\end{aligned}$$

This concludes the proof. □

B.3 Computing the shifts between \mathcal{W}^* and Algorithm 1

In this section, we consider one execution of the process when \mathcal{W}^* is used, and one execution when Algorithm 1 (meet at the center) is used. The variables involved in the execution of \mathcal{W}^* are denoted with $[\cdot]_{\mathcal{W}^*}$, while the variables involved in the execution of Algorithm 1 are denoted with $[\cdot]_{\text{MatC}}$.

We assume that the randomness is the same for both algorithms, that is, the initialization of agents is the same, and for every round t , and every $i \in I$, we have

$$\left[N_{m,i}^{(t)} \right]_{\mathcal{W}^*} = \left[N_{m,i}^{(t)} \right]_{\text{MatC}}, \quad \left[N_{d,i}^{(t)} \right]_{\mathcal{W}^*} = \left[N_{d,i}^{(t)} \right]_{\text{MatC}} \quad \text{and} \quad \left[\bar{\theta}_i^{(0)} \right]_{\mathcal{W}^*} = \left[\bar{\theta}_i^{(0)} \right]_{\text{MatC}}.$$

Claim 30. For every round t , and for every $i \in \{1, \dots, n\}$,

- $\left[\bar{\theta}_i^{(t)} \right]_{\mathcal{W}^*} = \left[\bar{\theta}_i^{(t)} \right]_{\text{MatC}},$
- $\left[d\theta_i^{(t)} \right]_{\mathcal{W}^*} - \left[d\theta_i^{(t)} \right]_{\text{MatC}} = \frac{1}{n} \rho_\star^{(t)} \sum_{i=1}^n \left[Y_i^{(t)} \right]_{\mathcal{W}^*} = \frac{1}{n} \rho_\star^{(t)} \sum_{i=1}^n \left[Y_i^{(t)} \right]_{\text{MatC}}.$

Proof. The proof proceeds by induction. More precisely, we prove the first item in the claim by induction on t . For any round t , we prove that the second item follows from the first one. Then, in the induction step, when proving that the first item regarding time $t+1$ holds, we use the second item regarding the previous time t .

The base case for the first item in the claim, i.e., $\left[\bar{\theta}_i^{(0)} \right]_{\mathcal{W}^*} = \left[\bar{\theta}_i^{(0)} \right]_{\text{MatC}}$ for every $i \in I$, holds by assumption. Next, let us assume that for some round t , we have for every $i \in I$,

$$\left[\bar{\theta}_i^{(t)} \right]_{\mathcal{W}^*} = \left[\bar{\theta}_i^{(t)} \right]_{\text{MatC}}.$$

Since the measurement noises are equal, then by the induction hypothesis, the measurement are also equal, that is, for every $i \in I$,

$$\left[Y_i^{(t)} \right]_{\mathcal{W}^*} = \left[\bar{\theta}_i^{(t)} + N_{m,i}^{(t)} \right]_{\mathcal{W}^*} = \left[\bar{\theta}_i^{(t)} + N_{m,i}^{(t)} \right]_{\text{MatC}} = \left[Y_i^{(t)} \right]_{\text{MatC}} = Y_i^{(t)}.$$

Thus,

$$\begin{aligned} \left[d\theta_i^{(t)} \right]_{\mathcal{W}^*} - \left[d\theta_i^{(t)} \right]_{\text{MatC}} &= \rho_\star^{(t)} Y_i^{(t)} - \frac{n-1}{n} \rho_\star^{(t)} \left(Y_i^{(t)} - \frac{1}{n-1} \sum_{j \neq i} Y_j^{(t)} \right) && \text{(by definition)} \\ &= \frac{1}{n} \rho_\star^{(t)} Y_i^{(t)} + \frac{1}{n} \rho_\star^{(t)} \sum_{j \neq i} Y_j^{(t)} \\ &= \frac{1}{n} \rho_\star^{(t)} \sum_{i=1}^n Y_i^{(t)} := \lambda_t. \end{aligned}$$

Finally,

$$\begin{aligned} \left[\bar{\theta}_i^{(t+1)} \right]_{\mathcal{W}^*} &= \left[\bar{\theta}_i^{(t)} - d\theta_i^{(t)} - N_{d,i}^{(t)} + \frac{1}{n-1} \sum_{j \neq i} \left(d\theta_j^{(t)} + N_{d,j}^{(t)} \right) \right]_{\mathcal{W}^*} && \text{(by Eq. (19))} \\ &= \left[\bar{\theta}_i^{(t)} - N_{d,i}^{(t)} \right]_{\text{MatC}} - \left[d\theta_i^{(t)} \right]_{\mathcal{W}^*} + \frac{1}{n-1} \sum_{j \neq i} \left[d\theta_j^{(t)} \right]_{\mathcal{W}^*} + \left[N_{d,j}^{(t)} \right]_{\text{MatC}} && \text{(1st item of the claim)} \\ &= \left[\bar{\theta}_i^{(t)} - (d\theta_i^{(t)} + \lambda_t) - N_{d,i}^{(t)} + \frac{1}{n-1} \sum_{j \neq i} d\theta_j^{(t)} + \lambda_t + N_{d,j}^{(t)} \right]_{\text{MatC}} && \text{(2nd item of the claim)} \\ &= \left[\bar{\theta}_i^{(t)} - d\theta_i^{(t)} - N_{d,i}^{(t)} + \frac{1}{n-1} \sum_{j \neq i} d\theta_j^{(t)} + N_{d,j}^{(t)} \right]_{\text{MatC}} \\ &= \left[\bar{\theta}_i^{(t+1)} \right]_{\text{MatC}}. && \text{(by Eq. (19))} \end{aligned}$$

This concludes the induction step, and establishes the proof. \square