

Clock Synchronization in Highly Dynamic Networks: An Information Theoretic Approach

Ofer Feinerman*

Amos Korman[†]

Abstract

We consider the *External Clock Synchronization* problem in dynamic sensor networks. Initially, sensors obtain inaccurate estimations of an external time reference and subsequently collaborate in order to synchronize their internal clocks with the external time. For simplicity, we adopt the *drift-free* assumption, where internal clocks are assumed to tick at the same pace. Unlike most works on clock synchronization that assume static networks, this paper focuses on an extreme case of highly dynamic networks. Specifically, we assume a non-adaptive scheduler adversary that dictates in advance an arbitrary, yet *independent*, meeting pattern. Such meeting patterns fit, for example, with short-time scenarios in highly dynamic settings, where each sensor interacts with only few other arbitrary sensors.

We propose an extremely simple clock synchronization algorithm that is based on weighted averages, and prove that its performance on any given independent meeting pattern is $O(1)$ -competitive with that of the best possible algorithm, which operates without any resource or computational restrictions, and knows the meeting pattern in advance. In particular, when all distributions involved are Gaussian, the performances of our scheme coincide with the optimal performances. Our proofs rely on an extensive use of the concept of Fisher information. We use the Cramér-rao bound and our definition of a *Fisher Channel Capacity* to quantify information flows and to obtain lower bounds on collective performance. This opens the door for further rigorous quantifications of information flows within collaborative sensors.

This paper should be considered as a regular submission.

*The Shlomo and Michla Tomarin Career Development Chair, The Weizmann Institute of Science, Rehovot, Israel. E-mail: ofer.feinerman@weizmann.ac.il. Supported by the Clore Foundation, the Israel Science Foundation (FIRST grant no. 1694/10) and the Minerva Foundation.

[†]CNRS and University Paris Diderot, Paris, 75013, France. E-mail: amos.korman@liafa.univ-paris-diderot.fr. Supported in part by the ANR project DISPLEXITY, and by the INRIA project GANG.

1 Introduction

1.1 Background and Motivation

Representing and communicating information is a main interest of theoretical distributed computing [35]. However, such studies often seem disjoint from what may be the largest body of work regarding coding and communication: Information theory [8, 39]. Perhaps the main reason for this stems from the fact that theoretical distributed computing studies are traditionally concerned with noiseless models of communication, in which the content of a message that passes from one node to another is not distorted. This reliability in transmission relies on an implicit assumption that error-corrections is guaranteed by a lower level protocol that is responsible for implementing communication. Indeed, when bandwidth is sufficiently large, one can encode a message with a large number of error-correcting bits in a way that makes communication noise practically a non-issue.

In some distributed scenarios, however, distortion in communication is unavoidable. One example concerns the classical problem of *clock synchronization*, which have attracted a lot of attention from both theoreticians in distributed computing [2, 26, 29, 34, 41], as well as practitioner engineers [7, 11, 12, 14, 18, 37, 40], see [28, 38, 44, 46] for comprehensive surveys. In this problem, processors need to synchronize their internal clocks (either among themselves only or with respect to a global time reference) relying on relative time measurements between clocks. Due to unavoidable unknown delays in communication, such measurements are inherently noisy. Furthermore, since the source of the noise is the delays, error-correction does not seem to be of any use for reducing the noise. The situation becomes even more complex when processors are mobile, preventing them from reducing errors by averaging repeated measurements to the same processors, and from contacting reliable processors. Indeed, the clock synchronization problem is particularly challenging in the context of wireless sensor networks and ad hoc networks which are typically formed by autonomous, and often mobile, sensors without central control.

Distributed computing models which include noisy communication call for a rigorous comprehensive study that employs information theoretical tools. Indeed, a recent trend in the engineering community is to view the clock synchronization problem from a signal processing point of view, and adopt tools from information theory (e.g., the Cramér-Rao bound) to bound the affect of inherent noise [6, 7, 18, 25], see [46] for a survey. However, this perspective has hardly received any attention by theoreticians in distributed computing that mostly focused on worst case message delays [2, 4, 26, 29], which do not seem to be suitable for information theoretic considerations. In fact, very few works on clock synchronisation consider a system with random delays and analyse it following a rigorous theoretical distributed algorithmic type of analysis. An exception to that is the work of Lenzen et al. [27], but also that work does not involve information theory. In this current paper, we study the clock synchronization problem through the purely theoretical distributed algorithmic perspective while adopting the signal processing and information theoretic point of view. In particular, we adopt tools from Fisher Information theory [42, 47].

We consider the *external* version of the problem [9, 12, 14, 32, 34, 44] in which processors (referred to as sensors hereafter) collaborate in order to synchronize their clocks with an external *global clock*. Informally, sensors initially obtain inaccurate estimates¹ of a global (external) time $\tau^* \in \mathbb{R}$ reference, and subsequently collaborate to align their internal clocks to be as close as possible to the external clock. To this end, sensors communicate through uni-directional pairwise interactions that include inherently *noisy measurements* of the relative deviation between their internal clocks and, possibly, some complementary information. To focus on the problems occurred by the initial inaccurate estimations of τ^* and the noise in the communication we restrict our attention to *drift-free* settings [2, 29], in which all clocks tick at the same rate. This setting essentially reduces the problem to the problem of estimating τ^* . See, e.g., [16, 43, 45] for works on estimation in the engineering community.

¹Traditional protocols like NTP [31] and TEMPO [17] use an external standard like GPS (Global Positioning System) or UTC (Universal Time) to synchronize networks. However, the use of of such systems poses a high demand for energy which is usually undesired in sensor networks. Hence, works in sensor networks typically assume that one *source* processor obtains an accurate estimate of the global time reference and essentially governs the synchronization of the rest of the sensors [34]. Here, we generalize this framework by assuming that each processor may initially have a different estimate quality of the global time reference, and our goal is to investigate what can be achieved given the qualities of initial estimations.

With very few exceptions [10, 23], almost all works on clock synchronization (and estimation) considered static networks. Indeed, the construction of efficient clock synchronization algorithms for dynamic networks is considered as a very important and challenging task² [38, 44]. This paper addresses this challenge by considering highly dynamic networks in which sensors have little or no control on who they interact with. Specifically, we assume a non-adaptive scheduler adversary that dictates in advance a meeting-pattern for the sensors. However, the adversary we assume is not unlimited. Specifically, to ease the analysis, we restrict the adversary to provide *independent-meeting patterns* only, in which it is guaranteed that whenever a sensor views another sensor, their transitive histories are disjoint³. Although they are not very good representatives of communication in static networks, independent meeting patterns fit very well with highly stochastic communication patterns as well as short-time scales, in which each sensor observes only few other arbitrary sensors (see more discussion in Section 2.1). Given such an adversarial meeting-pattern, we are concerned with minimizing the deviation of each internal clock from the global time.

As our objective is to model small and simple sensors, we are interested in algorithms that employ elementary internal computations and economic use of communication. We use competitive analysis to evaluate the performances of algorithms, comparing them to the best possible algorithm that knows the whole meeting pattern in advance and operates under the most liberal version of the model that allows for unrestricted resources in terms of memory and communication capacities, and individual computational ability.

1.2 Our contribution

Lower bounds on optimal performance. We first consider algorithm **Opt**, the best possible algorithm operating on the given independent meeting pattern. We note that specifying **Opt** seems challenging, especially since we do not assume a prior distribution on the starting global time, and hence the use of Bayesian statistics seems difficult. Fortunately, for our purposes, we are merely interested in lower bounding the performances on that algorithm. We achieved that by relating the smallest possible variance of a sensor at a given time to the largest possible *Fisher Information (FI)* of the sensor at that time. This measure quantifies the sensor’s current knowledge regarding the relative deviation between its local time and the global time. We provide a recursive formula to calculate J_a , the FI at sensor a , for any sensor a . Specifically, initially, the FI at a sensor is the Fisher information in the distribution family governing its initial deviation from the global time (see Section 2.4 for the formal definitions). When sensor a observes sensor b , the FI at a after this observation (denoted by J'_a) satisfies:

$$J'_a \leq J_a + \frac{1}{\frac{1}{J_b} + \frac{1}{J_N}}, \quad (1)$$

where J_N is the Fisher Information in the noise distribution related to the observation. To obtain this formula we prove a generalized version of the *Fisher information inequality* [42, 47]. Relying on the *Cramér-Rao bound* [8], this formula is then used to bound the corresponding variance under algorithm **Opt**. Specifically, the variance of the internal clock of sensor a is at least $1/J_a$.

Equation 1 provides immediate bounds on the convergence time. Specifically, the inequality sets a bound of J_N for the increase in the *FI* per interaction. In analogy to Channel Capacity as defined by Shannon [8] we term this upper bound as the *Fisher Channel Capacity*. Given small $\epsilon > 0$, the convergence time $T(\epsilon)$ is defined as the minimal number of observations required by the typical sensor until its variance drops below ϵ^2 (see Section 2.4 for the formal definition). Let J_0 denote the median initial Fisher Information of sensors. Based on the Fisher Channel Capacity we prove the following.

Theorem 1.1. *Assume that $J_0 \ll 1/\epsilon^2$ for some small $\epsilon > 0$, then $T(\epsilon) \geq (\frac{1}{\epsilon^2} - J_0)/J_N \approx 1/\epsilon^2 J_N$.*

²For example, dynamic meeting patterns prevent the use of classical external clock synchronization algorithms (e.g., [31, 34]) that are based on one or few *source* sensors that obtain accurate estimation of the global time and govern the synchronization of other sensors.

³Another informal way to view such patterns is that they guarantee that, given the global time, whenever a sensor views another sensor, their local clocks are independent; see Section 2.1 for a formal definition.

A constant-competitive compact algorithm. We propose a simple clock synchronization algorithm and prove that its performance on any given independent meeting pattern is highly competitive with that of the optimal one. That is, estimations of global time at each sensor remain unbiased throughout the execution and the variance at any given time is $O(1)$ -competitive with the best possible variance. In contrast to the optimal algorithm that may be based on transmitting complex functions in each interaction, and on performing complex internal computations, our simple algorithm is based on far more basic rules. First, transmission is restricted to a single *accuracy* parameter. Second, using the noisy measurement of deviation from the observed sensor, and the accuracy of that sensor, the observing sensor updates its internal clock and accuracy parameter by careful, yet elementary, weighted-averaging procedures.

Our weighted-average algorithm is designed to maximize the flow of Fisher information in interactions. This is proved by showing that the accuracy parameter is, at all times, both representative of the reciprocal of the sensor’s variance and close to the Fisher Information upper bound. In short, we proved the following.

Theorem 1.2. *There exists a simple weighted-average based clock synchronization algorithm which is $O(1)$ -competitive (at any sensor and at any time).*

A direct corollary of Theorem 1.2 implies that if all distributions involved are Gaussians, then the performances of our algorithm coincide with the optimal one, for each sensor and at any time. A second corollary implies that for populations whose size goes to infinity, the performances of our algorithm converge to those of the optimal one, as time goes to infinity.

We note that our algorithm does not require the use of sensor identities and can thus be also employed in *anonymous* networks [1, 13], yielding the same performances.

2 Preliminaries

2.1 The Model

We consider a collection of n sensors that collaborate in order to synchronize their internal clocks with an external global clock reference. We consider a finite set \mathcal{F} of sufficiently smooth (see definition in Section 2.4), probability density distributions (*pdf*) centered at zero. One specific distribution among the *pdfs* in \mathcal{F} is the *noise* distribution, referred to as $N(\eta)$. Each sensor a is associated with a distribution $\Phi_a(x) \in \mathcal{F}$ which governs the initialization deviation of its internal clock from the global time as described in the next paragraph. Depending on the specific model, we assume that sensor a knows various properties of Φ_a . In the most restricted version we consider, sensor a knows only the variance of Φ_a and in the most liberal version, we assume that a knows the full description of Φ_a . Execution is initiated when the global time is some $\tau^* \in \mathbb{R}$, chosen by an adversary.

Local clocks. Each sensor a is initialized with a local clock $\ell_a(0) \in \mathbb{R}$, randomly chosen according to $\Phi_a(x - \tau^*)$, independently of all other sensors. That is, the initial local time at sensor a is distributed around τ^* , and this distribution is governed by Φ_a . We stress that sensor a does not know the value τ^* and from its own local perspective the execution started at time $\ell_a(0)$. Sensors rely on both social interactions and further environmental cues⁴ to improve their estimates of the global time. In-between such events sensors are free to perform “shift” operations to adjust their local clocks. To focus on the problems occurred by the initial inaccurate estimations of τ^* and the noise in the communication we restrict our attention to *drift-free* settings [2, 29], in which all clocks tick at the same rate, consistent with the global time.

Opinions. The drift-free assumption reduces the external clock-synchronization problem to the problem of estimating τ^* . Indeed, recall that local clocks are initialized to different values but progress at the same rate. Because sensor a can keep the precise time since the beginning of the execution, its deviation from the global time

⁴In order for the model to include environmental cues, one or more of the sensors can be taken to represent the global clock. The initial times of these sensors are chosen according to highly concentrated distributions, Φ_a , around τ^* and remain fixed thereafter.

can be corrected had it known the difference between, $\ell_a(0)$, the initial local clock of a , and τ^* , the global time when the execution started. Hence, one can view the goal of sensor a as estimating τ^* . That is, without loss of generality, we may assume that all shifts performed by sensor a throughout the execution are shifts of its initial position $\ell_a(0)$ aiming to align it to be as close as possible to τ^* . Taking this perspective, we associate with each sensor an *opinion* variable x_a , initialized to $x_a(0) := \ell_a(0)$, and the goal of a is to have its opinion be as close as possible to τ^* . We view the opinion x_a as an *estimator* of τ^* , and note that initially, due to the properties of Φ_a , this estimator is unbiased, i.e., $\text{mean}(x_a(0) - \tau^*) = 0$. It is required that at any point in the execution, the opinion x_a remains an unbiased estimator of τ^* , and the goal of a is to minimize its Mean Square Error (MSE).

Due to this simple relation between internal clocks and opinions, in the remaining of this paper, we shall adopt the latter perspective and concern ourselves only with optimizing the opinions of sensors as estimators for τ^* , without discussing further the internal clocks.

Rounds. For simplicity of presentation, we assume that the execution proceeds in discrete steps, or rounds. We stress however that the rounds represent the order in which communication events occur (as determined by the meeting-pattern, see below), and do not necessarily correspond to the actual time. Given an algorithm A , the opinion maintained by the algorithm at round t (where t is a non-negative integer) at sensor a is denoted by $x_a(t, A)$. As mentioned, the algorithm aims to keep this value as close as possible to τ^* . When A is clear from the context, we may omit writing it and use the term $x_a(t)$ instead.

In each round $t \geq 1$, each sensor may first choose to shift (or not) its opinion, and then, if specified in the meeting pattern, it observes another specified sensor, thus obtaining some information. To summarize, in each round, a sensor executes the following consecutive actions: (1) Perform internal computation; (2) Perform an opinion-shift: $x_a(t) = x_a(t-1) + \Delta(x)$; and (3) Observe (or not) another sensor. For simplicity, all these three operations are assumed to occur instantaneously, that is, in zero time.

Mobility and adversarial independent meeting patterns. In cases where sensors are embedded in a Euclidian space, distances between positioning of sensors may impact the possible interactions. To account for physical mobility, and be as general as possible, we assume that an oblivious adversary controls the meeting pattern. That is, the adversary decides (before the execution starts), for each round, which sensor observes which other sensor.

A model that includes an unlimited adversary that controls the meeting pattern in in some sense too general⁵. To ease the analysis, in this paper, we restrict the adversary to provide only *independent* meeting patterns, in which the set of sensors in the transitive history of each observing sensor is disjoint from the one of the observed sensor. More formally, given a pattern of meetings \mathcal{P} , sensor a and round t , we first recursively define the set of *relevant* sensors of a at time t , denoted by $\mathcal{R}_a(t, \mathcal{P})$. At time zero, we define $\mathcal{R}_a(0, \mathcal{P}) := a$, and at round t , $\mathcal{R}_a(t, \mathcal{P}) := \mathcal{R}_a(t-1, \mathcal{P}) \cup \mathcal{R}(b, t-1, \mathcal{P})$ if a observes b at time $t-1$ (otherwise $\mathcal{R}_a(t, \mathcal{P}) := \mathcal{R}_a(t-1, \mathcal{P})$). A meeting pattern \mathcal{P} is called *independent* if whenever some sensor a observes a sensor b at some time t , then $\mathcal{R}_a(t-1, \mathcal{P}) \cap \mathcal{R}(b, t-1, \mathcal{P}) = \emptyset$. Note that an independent meeting pattern guarantees that given τ^* , the internal clocks of two interacting sensors a and b are independent. However, given τ^* and the internal clock of a , the internal clock of b and the relative time measurement between them are dependent (this point is explained in further details in Section 3).

Note that independent-meeting patterns are not very good representatives of communication in static networks. Indeed, in such patterns a sensor will not contact the same sensor twice, which contradicts many natural communication schemes in static networks. We note, however, that in some cases, a sequence of multiple consecutive observations between sensors can be compressed into a single observation of higher accuracy thus reducing the

⁵For example, in our model we assume that sensors are anonymous but we compare such algorithm to the best possible algorithm that knows the identities of sensors and the whole meeting pattern in advance. In such case, an arbitrary interaction pattern can match all sensors such that interactions occur only within pairs. As the sensors are anonymous, they cannot distinguish this case from other, more uniform, meeting patterns and hence, cannot be expected to act as efficiently as algorithms with identified sensors. Some limitations on the adversarial interaction network are therefore required.

dependencies between observations, and possibly converting a dependent meeting pattern into an independent one⁶. On the other hand, independent meeting patterns fit very well with highly stochastic short-time scales communication patterns, in which each sensor observes only few other arbitrary sensors. In this sense, such patterns can be considered as representing an extreme case of dynamic systems.

Because sensors have no control of when their next interaction will occur, or if it will occur at all, we require that estimates at each sensor be as accurate as possible at *any* point in time. This requirement is stronger than the liveness property that is typically required from distributed algorithms [24].

Convergence time. Consider a meeting pattern \mathcal{P} . Given small $\epsilon > 0$, the *convergence time* $T(\epsilon)$ of an algorithm A is defined as the minimal number of observations made by the typical sensor until its variance is less than ϵ^2 . More formally, let ρ denote the first round until we have more than half of the population satisfying $\text{var}(X_a(t, A)) < \epsilon^2$. For each sensor a , let $R(a)$ denote the number of observations made by a until time ρ . The convergence time $T(\epsilon)$ is defined as the median of $R(a)$ over all sensors a . Note that $T(\epsilon)$ is a lower bound on ρ , since $\rho \geq R(a)$ for every sensor a .

Communication. We assume that sensors are anonymous and hence, in particular, they do not know who they observe. Conversely, for the sake of lower bounds, we allow a much more liberal setting, in which sensors have unique identifiers, know who they interact with, and in fact, are given the whole pattern of meeting in advance.

When a sensor a observes another sensor b at some round t , the information transferred in this interaction contains a *passive* component and, possibly, a complementary *active* one. The passive component is a noisy relative deviation measurement between their opinions:

$$\tilde{d}_{ab}(t) = x_b(t) - x_a(t) + \eta,$$

where the additive noise term, η , is chosen from the noise probability distribution $N(\eta) \in \mathcal{F}$ whose variance is known to the sensors. (Note that this measurement is equivalent to the relative deviation measurement between the sensors' current local times because all clocks tick at the same pace.)

2.2 Compact algorithms

Our reference for evaluating performances is algorithm **Opt** which operates under the most liberal version of our model, which carries no restrictions on memory, communication capacities or internal computational power, and provides the best possible estimators at any sensor and at any time (we further assume that sensors acting under **Opt** know the meeting pattern in advance). In general, algorithm **Opt** may use complex calculations over very wasteful memories that include detailed distribution density functions, and possibly, accumulated measurements. Our main goal is to identify an algorithm whose performance is highly competitive with that of **Opt** but wherein communication and memory are economically used, and the local computations simple. Indeed, when it comes to applications to tiny and processors, simplicity and economic use of communication are crucial restrictions.

An algorithm is called *compact* if the internal state of each sensor a is some real number $y_a \in \mathbb{R}$, and the internal computations that a sensor can perform consist of a constant number of basic arithmetic operations, namely: addition, subtraction, multiplication, and division.

2.3 Competitive analysis

Fix a finite family \mathcal{F} of smooth *pdf*'s centered at zero (see the definition for smoothness in the next paragraph), and fix an assignment of a distribution $\Phi_a \in \mathcal{F}$ to each sensor a . For an algorithm A and an independent meeting pattern \mathcal{P} , let $X_a(t, A, \mathcal{P})$ denote the random variable indicating the opinion of sensor a at round t . Let $\text{mean}(X_a(t, A, \mathcal{P}))$ and $\text{var}(X_a(t, A, \mathcal{P}))$ denote, respectively, the mean and variance of $X_a(t, A, \mathcal{P})$, where

⁶For example, if sensors have unique identities and sensor a observes sensor b several times in a row, and it is guaranteed that sensor b did not change its state during these observations, then these observations can be treated by a as a single, more accurate, observation of b .

these are taken over all possible random initial opinions, communication errors, and possibly, coins flipped by the algorithm. Note that the unbiased assumption requires that $\text{mean}(X_a(t, A, \mathcal{P})) = \tau^*$. An algorithm A is called λ -competitive, if for *any* independent pattern of meetings \mathcal{P} , *any* sensor a , and at *any* time t , we have: $\text{var}(X_a(t, A, \mathcal{P})) \leq \lambda \cdot \text{var}(X_a(t, \mathbf{Opt}, \mathcal{P}))$.

2.4 Fisher information and the Cremér-Rao bound

The Fisher information is a standard way of evaluating the amount of information that a set of random measurements holds about an unknown parameter τ of the distribution from which these measurements were taken. We provide here some definitions for this notion; for more information the reader may refer to [8, 47].

A single variable probability distribution function (*pdf*) Φ is called *smooth* if it satisfies the following conditions, as stated by Stam [42]: (1) $\Phi(x) > 0$ for any $x \in \mathbb{R}$, (2) the derivative Φ' exists, and (3) the integral $\int \frac{1}{\Phi(y)} (\Phi'(y))^2 dy$ exists, i.e., $\Phi'(y) \rightarrow 0$ rapidly enough for $|y| \rightarrow \infty$. Note that, in particular, these conditions hold for natural distributions such as the Gaussian distribution. Recall that we consider a finite set \mathcal{F} of smooth one variable *pdfs*, one of them being the noise distribution $N(\eta)$, and all of which are centered at zero.

For a smooth *pdf* Φ , let $J_\Phi^\tau := \int \frac{1}{\Phi(y)} (\Phi'(y))^2 dy$ denote the Fisher information in the parameterized family $\{(\Phi(x, \tau))\}_{\tau \in \mathbb{R}} = \{(\Phi(x - \tau))\}_{\tau \in \mathbb{R}}$ with respect to τ . In particular, let $J_N = J_N^\tau$ denote the Fisher information in the parameterized family $\{N(\eta - \tau)\}_{\tau \in \mathbb{R}}$. More generally, consider a multi-variable *pdf* family $\{(\Phi(z_1 - \tau, z_2 \dots z_k))\}_{\tau \in \mathbb{R}}$ where τ is a translation parameter. The Fisher information in this family with respect to τ is defined as: $J_\Phi^\tau = \int \frac{1}{\Phi(z_1 - \tau, z_2 \dots z_k)} \left[\frac{d\Phi(z_1 - \tau, z_2 \dots z_k)}{d\tau} \right]^2 dz_1, dz_2 \dots dz_k$ (if the integral exists). As previously noted [47], since τ is a translation parameter, Fisher information is both unique (there is no freedom in choosing the parametrization) and independent of τ .

The Fisher information derives its importance by association with the Cremér-Rao inequality [8]. This inequality lower bounds the variance of the best possible estimator of τ^* by the reciprocal of the Fisher information that corresponds to the random variables on which this estimator is based.

Theorem 2.1. [The Cremér-Rao inequality] *Let \hat{X} be any unbiased estimator of $\tau^* \in \mathbb{R}$ which is based on a multi-variable sample $\bar{z} = (z_1, z_2 \dots z_k)$ taken from $\Phi(z_1 - \tau^*, z_2 \dots z_k)$. Then $\text{var}(\hat{X}) \geq 1/J_\Phi^\tau$.*

Initial Fisher-tightness : To define the initial Fisher-tightness parameter Δ_0 , we first define the *Fisher-tightness* of a single variable smooth distribution Φ centered at zero, as $\Delta(\Phi) = \text{var}(\Phi) \cdot J_\Phi^\tau$. Note that, by the Cramér-Rao bound, $\Delta(\Phi) \geq 1$ for any such distribution Φ . Moreover, equality holds if Φ is Gaussian [8]. Recall that \mathcal{F} is the finite collection of the smooth distributions containing the distributions Φ_a governing the initial opinions of sensors. The *initial Fisher-tightness* Δ_0 is the maximum of the Fisher-tightness over all distributions in \mathcal{F} and the noise distribution. Specifically, let $\Delta_0 = \max\{\Delta(\Phi) \mid \Phi \in \mathcal{F}\}$. Since \mathcal{F} contains finitely many distributions, Δ_0 is a constant. Moreover, observe that if the distributions in \mathcal{F} are all Gaussians then $\Delta_0 = 1$.

3 Technical difficulties

It is known for a single sensor, one can associate weights to samples so that a weighted-average procedure can fuse them optimally [30]. The proof therein relies on the assumption that all probability distributions are Gaussians whose functional forms, indeed, depend on their second moments only. Our setting is more complex, since it includes arbitrary differentiable *pdf*'s and multiple distributed sensors whose relative opinions constantly change.

The extension to multiple mobile sensors adds another dimension to the problem. One difficulty lies in the fact that the partial knowledge held by each sensor is relative (*e.g.*, an estimation of the deviation between the sensor's opinion and τ^*) and hence may require the sensors to carefully fuse perspectives other than their own. This difficulty is enhanced, as the sensors constantly shift their opinions. Indeed, for compact algorithms, where memory is restricted to a single parameter, storing the sum of previous shifts in the memory of a sensor is possible, but could drastically limit the degrees of freedom for encoding other information. On the other hand, without

encoding previous shifts, it is not clear how sensor a should treat information it had received from b .

In addition, compression of memory and communication appears to be detrimental. Indeed, maintaining and communicating highly detailed memories can, in some cases, significantly improve a sensor’s assessment of the target value (e.g., see Fig. 1 in Appendix A). However, maintaining a high degree of detail requires storing an arbitrary number of *pdf* moments which may grow with every interaction. Hence, it is not clear how to compress the information into few meaningful parameters while avoiding the accumulation of errors and runaway behavior.

Several technical difficulties arise when attempting to bound the performances of different algorithms. In natural type of algorithms, sensors’ memories can be regarded as maintaining *pdf*s that summarize their knowledge regarding their deviation from the target value τ^* . One of the analysis difficulties corresponds to the fact that the *pdf* held by a sensor at round t depends on many previous deviation measurements in a non-trivial way, and hence the variance of a realization of the *pdf* does not necessarily correspond to the variance of the sensors’ opinion, when taking into account all possible realizations of all measurements. Hence, one must regard each *pdf* as a multi-variable distribution. A second problem has to do with dependencies. The independent meeting pattern guarantees that the memory *pdf*’s of two interacting sensors are independent, yet, given the *pdf* of the observing sensor, the *pdf* of the observed sensor and the deviation measurement become dependent. Such dependencies make it difficult to track the evolution of a sensor’s accuracy of estimation over time. Indeed, to tackle this issue, we had to extend the Fisher information inequality [36, 42, 47] to a multi-variable dependent convolution case.

4 Lower bounds on the variance of **Opt**

In this section we provide lower bounds on the performances of algorithm **Opt** over a fixed independent pattern of meeting \mathcal{P} . Note that we are interested in bounding the performances of **Opt** and not in specifying its instructions. Identifying the details of **Opt** may still be of interest, but it is beyond the scope of this paper.

For simplicity of presentation, we assume that the rules of **Opt** are deterministic. We note, however, that our results can easily be extended to the case that **Opt** is probabilistic. For simplicity of notations, since this section deals only with algorithm **Opt** acting over \mathcal{P} , we use variables, such as the opinion $X_a(t)$ and the memory $Y_a(t)$ of sensor a , without parametrizing them by neither **Opt** nor by \mathcal{P} .

Under algorithm **Opt**, we assume that each sensor holds initially not only the variance of Φ_a , but the precise functional form of the distribution Φ_a (recall, Φ_a is centered at zero). In addition, we assume that sensors have unique identifiers and that each sensor knows the whole pattern \mathcal{P} in advance. Moreover, we assume that each sensor a knows for each other sensor b , the *pdf* Φ_b governing b ’s initial opinion. All this information is stored in one designated part of the memory of a .

Since **Opt** does not have any bandwidth constrains, we may assume, without loss of generality, that whenever some sensor a observes another sensor b , it obtains the whole memory content of b . Since **Opt** is deterministic, its previous opinion-shifts can be extracted from its interaction history, which is, without loss of generality, encoded in its memory⁷. Hence, when sensor a observes sensor b at some round t , and receives b ’s memory together with the noisy measurement $\tilde{d}_{ab}(t) = x_b(t) - x_a(t) + \eta$, sensor a may extract all previous opinion-shifts of both itself and b , treating the measurement $\tilde{d}_{ab}(t)$ as a noisy measurement of the deviation between the initial opinions, i.e., $\tilde{d}_{ab}(0) = x_b(0) - x_a(0) + \eta$. In other words, to understand the behavior of **Opt** at round t , one may assume that sensors never shift their opinions until round t , when they use all memory they gathered to shift their opinion in the best possible manner⁸. It follows that apart from the designated memory part that all sensors share, the memory $M_a(t)$ of sensor a at round t contains the initial opinion $X_a(0)$ and a collection $Y_a(t-1) := \{\tilde{d}_{bc}(0)\}_{bc}$ of

⁷In case **Opt** is probabilistic, previous shifts can be extracted from the memory plus the results of coin flips which may be encoded in the memory of the sensor as well.

⁸This observation implies, in particular, that previous opinion-shifts of sensors do not affect subsequent estimators in a way that may cause a conflict (a conflict may arise, e.g., when optimizing one sensor at one time necessarily makes estimators at another sensor, at a later time, sub-optimal), hence algorithm **Opt** is well-defined.

relative deviation measurements between initial opinions. That is, $M_a(t) = (X_0(t), Y_a(t-1))$. This multi-valued memory variable $M_a(t)$ contains all the information available to a at round t . In turn, this information is used by the sensor to obtain its opinion $X_a(t)$ which is required to serve as an unbiased estimator of τ^* .

4.1 The Fisher Information of sensors

We now define the notion of the Fisher Information associated with a sensor a at round t . This definition will be used to bound from below the variance of $X_a(t)$ under algorithm **Opt**.

Consider the multi-valued memory variable $M_a(t) = (X_0(t), Y_a(t-1))$ of sensor a that at round t . Note that $Y_a(t-1)$ is independent of τ^* . Indeed, once the adversary decides on the value τ^* , all sensors' initial opinions are chosen with respect to τ^* . Hence, since sensors memories contains only relative deviations between opinions, the memories by themselves do not contain any information regarding τ^* . In contrast, given τ^* , the random variables $Y_a(t-1)$ and $X_a(0)$ are, in general, dependent. Furthermore, in contrast to $Y_a(t-1)$, the value of $X_a(0)$ depends on τ^* , as it is chosen according to $\Phi_a(x - \tau^*)$. Hence, $M_a(t)$ is distributed according to a *pdf* family $\{(m_a(t), \tau)\}$ parameterized by a translation parameter τ . Based on $M_a(t)$, the sensor produces an unbiased estimation $X_a(t)$ of τ^* , that is, it should hold that: $\text{mean}(X_a(t) - \tau^*) = 0$, where the mean is taken with respect to the distribution of the random multi-variable $M_a(t)$.

Definition: The *Fisher Information (FI)* of sensor a at round t , termed $J_a(t)$, is the the Fisher information in the parameterized family $\{(m_a(t), \tau)\}_{\tau \in \mathbb{R}}$ with respect to τ .

By the Cramér-rao bound, the variance of any unbiased estimator used by the sensor a at round t is bounded from below by the reciprocal of the *FI* of sensor a at that time. That is, we have:

Lemma 4.1. $\text{var}(X_a(t)) \geq 1/J_a(t)$.

4.2 An upper bound on the Fisher Information $J_a(t)$

Lemma 4.1 implies that lower bounds on the variance of the opinion of a sensor can be obtained by bounding from above the corresponding *FI*. To this end, our next goal is to prove the following recursive inequality.

Theorem 4.2. *The FI of sensor a under algorithm **Opt** satisfies: $J_a(t+1) \leq J_a(t) + 1/(\frac{1}{J_b(t)} + \frac{1}{J_N})$.*

Proof. Consider the case that at round t , sensor a observes sensor b . After the interaction, the random multi-variable $Y_a(t)$ is composed of: (1) the random variable $\tilde{D}_{ab}(0) := X_b(0) - X_a(0) + N$, corresponding to the noisy deviation measurement between the initial opinions of a and b , and (2) the relative deviation measurements in both $Y_a(t-1)$ and $Y_b(t-1)$. We now aim at calculating the FI $J_a(t+1)$ available to sensor a at time $t+1$, with respect to the parameter τ . This is the FI with respect to τ , in the multi-variable $M_a(t+1) = (X_a(0), Y_a(t)) = (X_a(0), \tilde{D}_{ab}(0), Y_a(t-1), Y_b(t-1))$, where $X_a(0)$ is distributed according to $\Phi_a(x - \tau^*)$. Taking $\tilde{X}_b(0) = X_a(0) + \tilde{D}_{ab}(0) = X_b(0) + N$, this latter FI becomes the same as the FI in the random variables: $(X_a(0), \tilde{X}_b(0), Y_a(t-1), Y_b(t-1))$. Since the meeting pattern is independent, then given the environment value τ^* , the random multi-variable $(X_a(0), Y_a(t-1))$ is independent of the random multi-variable $(\tilde{X}_b(0), Y_b(t-1))$. By the additivity property of the Fisher information with respect to independent random multi-variables (see [42]), the *FI* $J_a(t+1)$ therefore equals the *FI* $J_a(t)$ (which is the FI in the random multi-variable $(X_a(0), Y_a(t-1))$) plus the FI $\tilde{J}_b(t)$ in the random multi-variable $(\tilde{X}_b(0), Y_b(t-1))$, both with respect to τ . That is, we have:

$$J_a(t+1, A) = J_a(t) + \tilde{J}_b(t). \quad (2)$$

Let us now focus on the rightmost term in Equation 2 and calculate $\tilde{J}_b(t)$. Given that the target value is some τ , the distribution of $(\tilde{X}_b(0), Y_b(t-1))$ can be described by the following convolution:

$$f_{\tilde{X}_b(0), Y_b(t-1)}[(\tilde{x}_b(0), y_b(t-1)) | \tau] = \int f_{X_b(0), Y_b(t-1)}[\tilde{x}_b(0) - \eta, y_b(t-1) | \tau] N(\eta) d\eta. \quad (3)$$

Observe that the right hand side of Equation 3 is a convolution of the distribution of $(X_b(0), Y_b(t-1))$ with the noise distribution N , where the convolution occurs with respect to the random variable $X_b(0)$. Our goal now is to bound the Fisher information in this convolution with respect to τ .

The Fisher information inequality [42, 47] bounds the Fisher Information of convolutions of single-variable distributions. Essentially, the theorem says that if x, y and τ are real values, $K(x-\tau), R(x-\tau)$ and $Q(x-\tau)$ are parameterized families and $K = R \otimes Q$, then $J(K) \leq 1/(\frac{1}{J(R)} + \frac{1}{J(Q)})$. To apply this inequality to Equation 3, we generalize it to distribution with multiple variables, where only one of them being convoluted. We rely on the fact that the random variable $Y_b(t-1)$ does not depend on τ^* (recall, it contains only relative deviation measurements). This fact turns out to be sufficient to overcome the potential complication rising from the fact that given the environmental value τ^* , the random variable $X_b(0)$ and the random multi-variable $Y_b(t-1)$ are no longer independent. In Appendix B we prove Lemma B.1 which extends the Fisher information inequality to our multi-variable (possibly dependent) convolution case, enabling to prove the inequality $\tilde{J}_b(t) \leq 1/(\frac{1}{J_b(t)} + \frac{1}{J_N})$. Together with Equation 2, we obtain the required recursive inequality for the FI. This completes the proof of the theorem. \square

5 A constant-competitive compact algorithm

We define a compact algorithm, termed **ALG**, and prove that its performances are highly-competitive with those of **Opt**. In this algorithm, each sensor a stores in its memory a single parameter $c_a \in \mathbb{R}$ that represents its *accuracy* regarding the quality of its current opinion with respect to τ^* . The initial accuracy of sensor a is set to $c_a(0) = 1/\text{var}(\Phi_a)$. When sensor a observes sensor b at some round t , it receives $c_b(t)$ and $\tilde{d}_{ab}(t)$, and acts as follows. Sensor a first computes the value $\hat{c}_b(t) = c_b(t)/(1 + c_b(t) \cdot \text{var}(N))$, a reduced accuracy parameter for sensor b that takes measurement noise into account, and then proceeds as follows:

Algorithm ALG

- **Update opinion:** $x_a(t+1) = x_a(t) + \frac{\tilde{d}_{ab}(t) \cdot \hat{c}_b(t)}{c_a(t) + \hat{c}_b(t)}$.
- **Update accuracy :** $c_a(t+1) = c_a(t) + \hat{c}_b(t)$.

Fix an independent meeting pattern. First, algorithm **ALG** is designed such that at all times, the opinion is preserved as an unbiased estimator of τ^* and the accuracy, $c_a(t)$, remains equal to the reciprocal of the current variance of the opinion $X_a(t, \mathbf{ALG})$. Indeed, the following lemma is proven in Appendix C.

Lemma 5.1. *At any round t and for any sensor a , we have: (1) the opinion $X_a(t, \mathbf{ALG})$ serves as an unbiased estimator of τ^* , and (2) $c_a(t) = 1/\text{var}(X_a(t, \mathbf{ALG}))$.*

We are now ready to analyze the competitiveness of algorithm **ALG**, by relating the variance of a sensor a at round t to the corresponding *FI*, namely, $J_a(t)$. Recall that Lemma 4.1 gives a lower bound on the variance of algorithm **Opt** at a sensor a , which depends on the corresponding *FI* at the sensor. Specifically, we have: $\text{var}(X_a(t, \mathbf{Opt})) \geq 1/J_a(t)$. Initially, the *FI* $J_a(0)$ at a sensor a equals the Fisher information in the parameterized family $\Phi_a(x-\tau)$ with respect to τ , and hence is at most the initial accuracy $c_a(0)$ times Δ_0 . In Equation A-3 (see Appendix D) we show that the gain in accuracy following an interaction is always at least as large the corresponding upper bound on the gain in Fisher information as given in Theorem 4.2, divided by the initial Fisher-tightness. That is: $c_a(t+1) - c_a(t) \geq \left(1/(\frac{1}{J_b(t)} + \frac{1}{J_N})\right) / \Delta_0$. Informally, this property of **ALG** can be interpreted as maximizing the Fisher information flow in each interaction up to an approximation factor of Δ_0 . By induction (see proof in Appendix D), we obtain the following.

Lemma 5.2. *At every round t , we have $c_a(t) \geq J_a(t) / \Delta_0$.*

Lemmas 4.1, 5.1 and 5.2 can now be combined to yield: $\text{var}(X_a(t, \mathbf{ALG})) \leq \Delta_0 \cdot \text{var}(X_a(t, \mathbf{Opt}))$. This establishes Theorem 1.2. \square

Note that Theorem 1.2 states that **ALG** is constant-competitive at any sensor and at any time, as the initial Fisher-tightness Δ_0 is a constant. We now aim to identify those cases where **ALG** performs even better. One such case is when the distributions in \mathcal{F} as well as the noise distribution $N(\eta)$ are all Gaussians. In this case $\Delta_0 = 1$ and Theorem 1.2 therefore states that the variance of **ALG** equals that of **Opt**, for any sensor at any time. Another case is when the noise is Gaussian and both the population size n and the round t go to infinity. In this case, analyzed in Appendix E, as time increases, the performances of **ALG** become arbitrarily close to those of **Opt**.

6 The Fisher Channel Capacity and convergence times

For a fixed independent meeting pattern, the FI $J_a(t)$ at a sensor a and round t was defined in Section 4.1 with respect to algorithm **Opt**. We note that this definition applies to any algorithm A as long as it is sufficiently smooth so that the corresponding Fisher informations are well-defined. This quantity $J_a(t, A)$ would respect the same recursive inequality as state in Theorem 4.2, that is, we have: $J_a(t+1, A) \leq J_a(t, A) + \frac{1}{\frac{1}{J_b(t, A)} + \frac{1}{J_N}}$. This directly implies the following:

$$J_a(t+1, A) - J_a(t, A) \leq J_N. \quad (4)$$

The inequality above sets a bound of J_N for the increase in *FI* per round. In analogy to Channel Capacity as defined by Shannon [8] we term this upper bound as the *Fisher Channel Capacity*.

The restriction on information flow as given by the Fisher Channel Capacity can be translated into lower bounds for convergence time of algorithm **Opt** (and hence also apply for any algorithm). Recall, ρ is the first round until we have more than half of the population satisfying $\text{var}(X_a(t)) < \epsilon^2$. By Lemma 4.1, a sensor, a , with variance smaller than ϵ^2 must have a large *FI*, specifically, $J_a(\rho) \geq 1/\epsilon^2$. To get some intuition on the convergence time, assume that the number of sensors is odd, and let J_0 denote the median initial *FI* of sensors (this is the median of the *FI*, J_{Φ_a} , over all sensors a), and assume $J_0 \ll 1/\epsilon^2$. By definition, more than a half of the population have initial Fisher information at most J_0 . By the Pigeon-hole principle, at least one sensor has an *FI* of, at most, J_0 at $t = 0$ and, at least, $1/\epsilon^2$ at $t = \rho$. Theorem 1.1 follows by the fact that, by Equation 4, this sensor could increase its *FI* by, at most, J_N in each observation.

7 Conclusion

We provide a fresh approach to the study of clock synchronizing, following a purely theoretic distributed algorithmic type of study and employing techniques from information theory. We have focused on arbitrary, yet independent, meeting patterns, and demanded the performances of each sensor to be as high as possible at any point in the execution. We have established lower bounds on the performances of algorithm **Opt**, the best possible clock synchronization algorithm operating under the most liberal version of our model. We have identified algorithm **ALG**, an extremely simple algorithm whose performances are constant-competitive with those of **Opt**. Moreover, under Gaussian conditions, the accuracies of sensors under **ALG** precisely equal those of **Opt**.

Algorithm **ALG** is based on storing and communicating a single *accuracy* parameter that complements noisy deviation measurements, and on internal computations and update rules that are based on *weighted-average* operations. Our proofs rely on an extensive use of the concept of Fisher information. We use the Cramér-rao bound and our definition of a *Fisher Channel Capacity* to quantify information flows and to obtain lower bounds on best possible performance. This opens the door for further rigorous quantifications of information flows within collaborative sensors.

Our information theoretic approach allowed us to tackle the clock synchronization problem in dynamic networks. In this initial work, we focus on independent meeting patterns which can be considered as representing short times scales in highly dynamic scenarios. Our hope is that studying such extreme dynamic cases will help to provide tools and insights for future work dealing with other dynamic scenarios.

This work is further relevant to the problem of collective approximation of environmental values by biological groups [22].

References

- [1] D. Angluin. *Local and Global Properties in Networks of Processors* STOC 82-93, 1980.
- [2] H. Attiya, A. Herzberg, and S. Rajsbaum. *Optimal Clock Synchronization under Different Delay Assumptions*. SIAM J. Comput. 25(2), 369-389, 1996.
- [3] Z. Bar-Yossef, T. S. Jayram, R. Kumar, D. Sivakumar. *Information Theory Methods in Communication Complexity*. IEEE Conference on Computational Complexity, 93-102, 2002.
- [4] S. Biaz, and J. L. Welch. *Closed form bounds for clock synchronization under simple uncertainty assumptions*. Inf. Process. Lett. 80(3), 151-157, 2001.
- [5] N. M. Blachman. *The convolution inequality for entropy powers*. IEEE Transactions on Information Theory 11(2), 267-271, 1965.
- [6] Q. Chaudhari, E. Serpedin, and Y.C. Wu, *Improved estimation of clock offset in sensor networks*. Proc. IEEE Int. Conf. Communications (ICC), 2009.
- [7] Q. Chaudhari, E. Serpedin, and K. Qaraqe. *On minimum variance unbiased estimation of clock offset in distributed networks*. IEEE Trans. Inform. Theory, vol. 56, no. 6, 2010.
- [8] T. M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2nd edition, 2006.
- [9] F. Cristian. *Probabilistic Clock Synchronization*. Distributed Computing 3(3): 146-158, 1989.
- [10] D. Dolev, J. Halpern, B. Simons, and R. Strong. *Dynamic fault-tolerant clock synchronization*. Journal of the ACM, 42(1), 143–185, 1995.
- [11] J. Elson, L. Girod, and D. Estrin. *Fine-Grained Network Time Synchronization Using Reference Broadcasts*. ACM SIGOPS Operating Systems Review, 36, 147–163, 2002.
- [12] J. Elson and K. Romer. *Wireless Sensor Networks: A New Regime for Time Synchronization*. Proc. First Workshop on Hot Topics In Networks (HotNets-I), Princeton, New Jersey. Oct. 2002.
- [13] Of. Feinerman, B. Haeupler, and A. Korman. *Breathe before speaking: efficient information dissemination despite noisy, limited and anonymous communication*. PODC, 114-123, 2014.
- [14] C. Fetzer and F. Cristian. *Integrating External and Internal Clock Synchronization*. Real-Time Systems 12(2): 123-171, 1997.
- [15] A. El Gamal and Y. Kim. *Network Information Theory*. Cambridge University Press, 709pp, 2012.
- [16] J. Gubner. *Distributed Estimation and Quantization* IEEE Transactions on Information Theory, 39(4), 1993.
- [17] R. Gusella and S. Zatti. *TEMPO - A Network Time Controller for a Distributed Berkeley UNIX System*. IEEE Distributed Processing Technical Committee Newsletter 6, No S12-2, 1984.
- [18] D. Jeske. *On the maximum likelihood estimation of clock offset*. IEEE Trans. Commun., vol. 53, no. 1, 2005.
- [19] S. Kar and J. M.F. Moura. *Distributed Consensus Algorithms in Sensor Networks With Imperfect Communication: Link Failures and Channel Noise* IEEE Transactions on Signal Processing 57(1), 355–369, 2009.
- [20] D. Kempe, A. Dobra, J. Gehrke. *Gossip-Based Computation of Aggregate Information*. FOCS 2003: 482-49.

- [21] R. Koetter and F. R. Kschischang. *Coding for errors and erasures in random network coding*. IEEE Transactions on Information Theory 54 no. 8, 3579–3591, 2008.
- [22] A. Korman, E. Greenwald and O Feinerman. *Confidence Sharing: an Economic Strategy for Efficient Information Flows in Animal Groups*. PLOS Computational Biology, 10(10), 2014.
- [23] F. Kuhn, C. Lenzen, T. Locher, and R. Oshman. *Optimal gradient clock synchronization in dynamic networks*. PODC, 430-439, 2010.
- [24] L. Lamport. *Proving the Correctness of Multiprocess Programs*. IEEE Transactions on Software Engineering (2): 125–143. 1977.
- [25] M. Leng and Y. Wu. *On joint synchronization of clock offset and skew for wireless sensor networks under exponential delay*. in Proc. IEEE ISCAS. 2010.
- [26] C. Lenzen, T. Locher, and R. Wattenhofer. *Tight Bounds for Clock Synchronization*. J. of the ACM, 57(2), 2010.
- [27] C. Lenzen, P. Sommer, and R. Wattenhofer. *PulseSync: An Efficient and Scalable Clock Synchronization Protocol*. ACM/IEEE Transactions on Networking (TON), March 2014.
- [28] C. Lenzen, T. Locher, P. Sommer, and R. Wattenhofer. *Clock Synchronization: Open Problems in Theory and Practice*. Proc. 36th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM), January 2010.
- [29] J. Lundelius and N. Lynch. *An Upper and Lower Bound for Clock Synchronization*. Information and Control 62, 190-204 (1984).
- [30] J. M. McNamara and A. I. Houston. *Memory and the efficient use of information..* Journal of theoretical biology 125(4), 385–395, 1987.
- [31] D.L. Mills. *Internet time synchronization: the network time protocol*. IEEE Transactions of Communications 39 (10) 1482–1493, 1991.
- [32] D.L. Mills. *Improved algorithms for synchronizing computer network clocks*. IEEE/ACM Trans. Networks 3, 3, 1995.
- [33] R. Ostrovsky and B. Patt-Shamir. *Optimal and Efficient Clock Synchronization Under Drifting Clocks*. PODC 1999, 3-12.
- [34] B. Patt-Shamir and S. Rajsbaum. *A theory of clock synchronization*. STOC 1994, 810-819.
- [35] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- [36] O. Rioul. *Information theoretic proofs of entropy power inequalities*. IEEE Transactions on Information Theory, Vol. 57, No. 1, pp. 33-55, 2011.
- [37] M. Sichitiu and C. Veerarittiphan. *Simple, Accurate Time Synchronization for Wireless Sensor Networks*. In Proc. IEEE Wireless Communications and Networking Conference (WCNC), 2003.
- [38] F. Sivrikaya and B. Yener. *Time synchronization in sensor networks: a survey*. IEEE Network 18(4): 45-50, 2004.
- [39] C. Shannon. *A Mathematical Theory of Communication*. Bell System Technical Journal 27(3), 379–423, 1948.

- [40] R. Solis, V. Borkar, and P. R. Kumar. *A New Distributed Time Synchronization Protocol for Multihop Wireless Networks*. In Proc. 45th IEEE Conference on Decision and Control (CDC), 2006.
- [41] T. K. Srikanth and S. Toueg. *Optimal Clock Synchronization*. J. ACM, 34(3), 626–645, 1987.
- [42] A. J. Stam. *Some inequalities satisfied by the quantities of information of Fisher and Shannon*. Inform. and Control, 2, 101-112. 1959.
- [43] L. Xiao, S. Boyd and S. Lall. *A scheme for robust distributed sensor fusion based on average consensus*. Proc of the 4th international symposium on Information processing in sensor networks (IPSN), 2005.
- [44] B. Sundararaman, U. Buy, and A. D. Kshemkalyani. *Clock synchronization for wireless sensor networks: a survey*. Ad Hoc Networks 3, 281–323, 2005.
- [45] R. Viswanathan, and P.K. Varshney. *Distributed detection with multiple sensors I. Fundamentals*. Proceedings of the IEEE, 1997.
- [46] Y.C. Wu, Q. M. Chaudhari, and E. Serpedin. *Clock Synchronization of Wireless Sensor Networks*. IEEE Signal Process. Mag. 28(1): 124-138, 2011.
- [47] R. Zamir. *A proof of the Fisher Information inequality via a data processing argument*. IEEE Trans Inf Theory, 482–491, 2003.

APPENDIX

A An example indicating that storing and communicating detailed assessments of target opinion is, in some cases, extremely beneficial

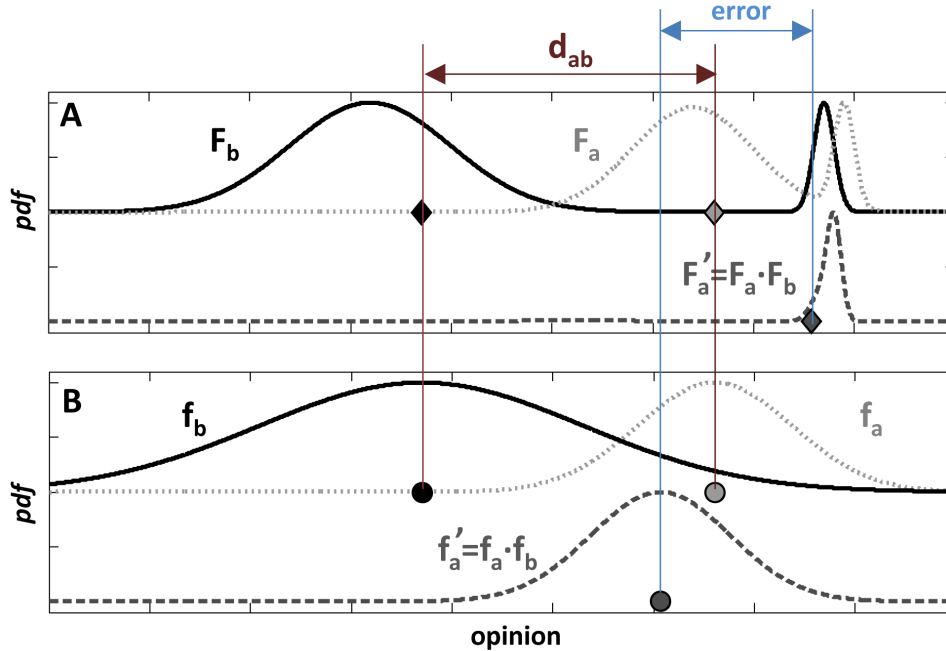


Figure 1: The apparent importance of storing and communicating highly detailed opinions. **A.** The memories of sensor *a* (dotted line) and sensor *b* (full line) are represented by a *pdf*'s (capital *F*'s) that summarize their full information regarding the target. The sensors opinions (marked by the light gray and black diamonds) are set at the mean of their corresponding *pdf*'s and deviate by d_{ab} . Following an interaction, in which sensor *a* obtains this relative opinion deviation as well as the full memory state of sensor *b*, it optimally updates memory to the product of these two distributions (dashed line). Sensor *a* then moves to the mean of its updated distribution (dark gray diamond). **B.** The sensors are initiated as in **A**. However, before they interact, their memories are compressed into Gaussians (lowercase *f*'s) that agree in mean and variance with their previous assessments. Note that since the mean values have not changed, the initial opinions of the sensors in both panels are identical and the deviation between them remains d_{ab} . Following the interaction, sensor *a* has shifted to an opinion whose deviation from the optimal opinion is marked by the 'error' arrow. Sensor *a* also gains less information when compared to the case in which compression had been avoided (compare variances of dashed curves in **A** and in **B**). For simplicity, all measurements in this figure were taken to be noiseless.

B Extending the Fisher inequality

The Fisher information inequality [42] (see also [5, 36, 47]) applies for three one-variable distribution families $r(z)$, $p_1(x_1)$, and $p_2(x_2)$ parameterized by μ such that r is a convolution of p_1 and p_2 , that is, $r(z) = \int p_1(z-t) \cdot p_2(t) dt$. The theorem gives an upper bound of the Fisher information J_r^μ of the family $r(z - \mu)$ (with respect to μ) based on the Fisher information $J_{p_1}^\mu$ and $J_{p_2}^\mu$ of the families $p_1(x_1 - \mu)$, and $p_2(x_2 - \mu)$, respectively. Specifically, the theorem states that: $(\alpha_1 + \alpha_2)^2 J_r^\mu \leq \alpha_1^2 J_{p_1}^\mu + \alpha_2^2 J_{p_2}^\mu$, for any two real numbers α_1 and α_2 . This in particular implies that $1/J_r^\mu \geq 1/J_{p_1}^\mu + 1/J_{p_2}^\mu$.

The following lemma extends the Fisher information inequality to the case where the distributions p_1 and r are composed of multiple, not necessarily independent, variables, where the convolution with p_2 takes place over one of the variables of p_1 .

Lemma B.1. *Let $\{p_1(x_1 - \tau, \bar{x}_3)\}_{\tau \in \mathbb{R}}$ and $\{p_2(x_2 - \tau)\}_{\tau \in \mathbb{R}}$ be two pdf families with a translation parameter τ such that x_1 and x_2 are real variables, \bar{x}_3 is a vector of multiple real valued variables and $J_{p_1(x_1 - \tau, \bar{x}_3)}^\tau$ and $J_{p_2(x_2 - \tau)}^\tau$ are the corresponding Fisher information with respect to τ . Let $r(z - \tau, \bar{x}_3) = \int p_1(t - \tau, \bar{x}_3) \cdot p_2(z - t) dt$ be the convolution of p_1 and p_2 . Then the Fisher information in the family $\{r(z - \tau, \bar{x}_3)\}_{\tau \in \mathbb{R}}$ with respect to τ satisfies:*

$$\frac{1}{J_{r(z - \tau, \bar{x}_3)}^\tau} \geq \frac{1}{J_{p_1(x_1 - \tau, \bar{x}_3)}^\tau} + \frac{1}{J_{p_2(x_2 - \tau)}^\tau}.$$

Proof. We start by using the definition of r as a convolution over p_2 and the first variable of p_1 :

$$r(z - \tau, \bar{x}_3) = \int p_1(t - \tau, \bar{x}_3) \cdot p_2(z - t) dt.$$

We can insert the density function $p(\bar{x}_3)$ to rewrite the right hand side as:

$$\begin{aligned} & \int p_1(t - \tau | \bar{x}_3) \cdot p(\bar{x}_3) \cdot p_2(z - t) dt \\ &= p(\bar{x}_3) \int p_1(t - \tau | \bar{x}_3) \cdot p_2(z - t) dt. \end{aligned}$$

Implying that:

$$r(z - \tau | \bar{x}_3) = \int p_1(t - \tau | \bar{x}_3) \cdot p_2(z - t) dt.$$

We now define the distributions $R(z) = r(z - \tau | \bar{x}_3)$ and $P_1(t) = p_1(t - \tau | \bar{x}_3)$ so that the previous equation becomes:

$$R(z) = \int P_1(t) \cdot p_2(z - t) dt,$$

for which we apply the original Lemma as first proved by Stam [42] to deduce that for any two real numbers α_1 and α_2 , we have:

$$(\alpha_1 + \alpha_2)^2 J_{R(z - \mu)}^\mu \leq \alpha_1^2 \cdot J_{P_1(x_1 - \mu)}^\mu + \alpha_2^2 \cdot J_{p_2(x_2 - \mu)}^\mu.$$

Note that $J_{P_1(x_1 - \mu)}^\mu$ is well defined since, for a given \bar{x}_3 , $P_1(x_1 - \mu)$ is proportional to $p_1(x_1 - \mu, \bar{x}_3)$ and the Fisher information integral of $p_1(x_1 - \mu, \bar{x}_3)$ converges when integrating over all possible values of \bar{x}_3 . This implies (see [36]) that the Fisher information in the convolution $R(z - \mu)$ is well defined and the equation above holds. We now multiply both sides of the equation by $p(\bar{x}_3)$ and integrate over \bar{x}_3 , to obtain:

$$(\alpha_1 + \alpha_2)^2 \int J_{R(z - \mu)}^\mu p(\bar{x}_3) d\bar{x}_3 \leq \alpha_1^2 \int J_{P_1(x_1 - \mu)}^\mu p(\bar{x}_3) d\bar{x}_3 + \alpha_2^2 \int J_{p_2(x_2)}^\mu p(\bar{x}_3) d\bar{x}_3. \quad (\text{A-1})$$

Plugging in the definitions for Fisher information and $R(z)$, the integral on the left hand side becomes:

$$\begin{aligned}
\int J_{R(z-\mu)}^\mu p(\bar{x}_3) d\bar{x}_3 &= \int J_{r(z-\mu-\tau|\bar{x}_3)}^\mu p(\bar{x}_3) d\bar{x}_3 \\
&= \int \int \frac{1}{r(z-\mu-\tau|\bar{x}_3)} \left(\frac{dr(z-\mu-\tau, \bar{x}_3)}{d\mu} \right)^2 dz p(\bar{x}_3) d\bar{x}_3 \\
&= \int \int \frac{1}{r(z-\mu-\tau|\bar{x}_3)p(\bar{x}_3)} \left(\frac{d[r(z-\mu-\tau|\bar{x}_3)p(\bar{x}_3)]}{d\mu} \right)^2 dz d\bar{x}_3 \\
&= \int \int \frac{1}{r(z-\mu-\tau, \bar{x}_3)} \left(\frac{dr(z-\mu-\tau, \bar{x}_3)}{d\mu} \right)^2 dz d\bar{x}_3 \\
&= \int \int \frac{1}{r(z-\mu-\tau, \bar{x}_3)} \left(\frac{dr(z-\mu-\tau, \bar{x}_3)}{d\tau} \right)^2 dz d\bar{x}_3 \\
&= \int \int \frac{1}{r(\tilde{z}-\tau, \bar{x}_3)} \left(\frac{dr(\tilde{z}-\tau, \bar{x}_3)}{d\tau} \right)^2 d\tilde{z} d\bar{x}_3 \\
&= J_{r(z-\tau, \bar{x}_3)}^\tau,
\end{aligned}$$

where we used $\tilde{z} = z - \mu$ and the fact that \bar{x}_3 is independent of τ .

Similarly, the integral over the first term on the right hand side of Equation A-1 gives $J_{p_1(x_1-\tau, \bar{x}_3)}^\tau$. The last term is:

$$\int J_{p_2(x_2-\mu)}^\mu p(\bar{x}_3) d\bar{x}_3 = J_{p_2(x_2-\mu)}^\mu \int p(\bar{x}_3) d\bar{x}_3 = J_{p_2(x_2-\mu)}^\mu = J_{p_2(x_2-\tau)}^\tau,$$

by normalization of the distribution \bar{x}_3 .

Finally, Equation A-1 translates to:

$$(\alpha_1 + \alpha_2)^2 J_{r(z-\tau, \bar{x}_3)}^\tau \leq \alpha_1^2 \cdot J_{p_1(x_1-\tau, \bar{x}_3)}^\tau + \alpha_2^2 \cdot J_{p_2(x_2-\tau)}^\tau,$$

for any real α_1 and α_2 . Setting $\alpha_1 = J_{p_2(x_2-\tau)}^\tau$ and $\alpha_2 = J_{p_1(x_1-\tau, \bar{x}_3)}^\tau$, we finally obtain:

$$\frac{1}{J_{r(z-\tau, \bar{x}_3)}^\tau} \geq \frac{1}{J_{p_1(x_1-\tau, \bar{x}_3)}^\tau} + \frac{1}{J_{p_2(x_2-\tau)}^\tau},$$

as desired. □

C Proof of Lemma 5.1

Fix an independent meeting pattern \mathcal{P} and a *pdf* assignment, $\Phi_a \in \mathcal{F}$, for each sensor a . Let us now prove the first part of the lemma, namely, that the opinion $X_a(t)$ at any sensor a and round t serves as an unbiased estimator for τ^* . The claim holds at time zero, and assume by induction that it holds at round t . Now consider the case that sensor a observes another sensor b at round t . The opinion of u after the interaction, becomes:

$$x_a(t+1) = x_a(t) + \frac{\tilde{d}_{ab}(t) \cdot \hat{c}_b(t)}{c_a(t) + \hat{c}_b(t)} = \frac{x_a(t)c_a(t) + x_b(t)\hat{c}_b(t)}{c_a(t) + \hat{c}_b(t)} + \frac{\eta \cdot \hat{c}_b(t)}{c_a(t) + \hat{c}_b(t)}. \quad (\text{A-2})$$

By induction, $X_a(t)$ and $X_b(t)$ are both unbiased estimators of τ^* . Recall now that the noise distribution $N(\eta)$ is centered around zero. Moreover, observe that at round t , the accuracy at each sensor a , namely $c_a(t)$, is

deterministically defined (given the fixed pattern of meeting, and the assignment of *pdf*'s to the sensors). In particular, at round t , the accuracy s $c_a(t)$, $c_b(t)$ as well as $\hat{c}_b(t)$ are all fixed constants. Equation A-2 therefore implies the following.

Claim C.1. *At any round t and for any sensor a , the opinion $x_a(t)$ serves as an unbiased estimator of τ^* .*

Claim C.1 established the first part of the lemma. Let us now turn to prove the second part. This part of the lemma holds for time $t = 0$ by definition of $c_a(0)$. Assume by induction that for any sensor a at round t it holds that $c_a(t) = 1/\text{var}(X_a(t))$ and consider time $t + 1$. We now consider an interaction between two sensors at round t , in which sensor a observes sensor b . The variance of the new opinion of a is:

$$\begin{aligned} \text{var}(X_a(t+1)) &= \text{var}\left(x_a(t) + \frac{\tilde{d}_{ab}\hat{c}_b(t)}{c_a(t) + \hat{c}_b(t)}\right) = \text{var}\left(\frac{x_a(t)c_a(t) + \hat{c}_b(t)(x_b(t) + \eta)}{c_a(t) + \hat{c}_b(t)}\right) \\ &= \frac{c_a^2(t) \cdot \text{var}(X_a(t)) + \hat{c}_b^2(t) \cdot \text{var}(X_b(t) + \eta)}{(c_a(t) + \hat{c}_b(t))^2} \\ &= \frac{c_a^2(t) \cdot \text{var}(X_a(t)) + \hat{c}_b^2(t) \cdot (\text{var}(X_b(t)) + \text{var}(N(\eta)))}{(c_a(t) + \hat{c}_b(t))^2} \\ &= \frac{c_a^2(t) \cdot 1/c_a(t) + \hat{c}_b^2(t) \cdot (1/c_b(t) + \text{var}(N(\eta)))}{(c_a(t) + \hat{c}_b(t))^2} = \frac{c_a^2(t) \cdot 1/c_a(t) + \hat{c}_b^2(t) \cdot 1/\hat{c}_b(t)}{(c_a(t) + \hat{c}_b(t))^2} \\ &= \frac{c_a(t) + \hat{c}_b(t)}{(c_a(t) + \hat{c}_b(t))^2} = \frac{1}{(c_a(t) + \hat{c}_b(t))} = \frac{1}{c_a(t+1)}, \end{aligned}$$

which proves the induction step. This complete the proof of Lemma 5.1.

D Proof of Lemma 5.2

By Lemma 5.1 and the definition of Δ_0 , we have $c_a(0) \geq J_a(0)/\Delta_0$, and hence Lemma 5.2 holds at time 0.

Assume by induction that the lemma holds at round t and consider an interaction at round t when sensor a observes sensor b . Let $c_N = 1/\text{var}(N(\eta))$. By definition of algorithm **ALG**, we have:

$$c_a(t+1) - c_a(t) = \hat{c}_b(t) = \frac{1}{1/c_b(t) + 1/c_N}.$$

By the induction hypothesis applied on sensor b , we have:

$$\frac{1}{1/c_b(t) + 1/c_N} \geq \frac{1}{\frac{\Delta_0}{J_b(t)} + 1/c_N} = \frac{1}{\Delta_0} \cdot \frac{1}{\frac{1}{J_b(t)} + \frac{1}{\Delta_0 \cdot c_N}}.$$

Again, by definition of Δ_0 , we have $\Delta_0 \geq J_N/c_N$. Hence:

$$c_a(t+1) - c_a(t) \geq \frac{1}{\Delta_0} \cdot \frac{1}{\frac{1}{J_b(t)} + \frac{1}{J_N}}. \quad (\text{A-3})$$

This means that the gain in accuracy at sensor a following an observation of sensor b is up to a multiplicative factor of Δ_0 at least as large the corresponding gain in FI of the sensor (operating under **Opt**).

Finally, applying the induction hypothesis for sensor a at round t , we have $c_a(t) \geq J_a(t)/\Delta_0$. Plugging this in Equation A-3, we obtain:

$$c_a(t+1) \geq \frac{1}{\Delta_0} \cdot \left(J_a(t) + \frac{1}{\frac{1}{J_b(t)} + \frac{1}{J_N}} \right) \geq J_a(t+1)/\Delta_0,$$

where the second inequality holds by Theorem 4.2. This completes the proof of the lemma.

E On the performances of ALG at large times

We now investigate the performances of algorithm **ALG** at large times, and show that as time increases, the variance of **ALG** becomes arbitrarily close to zero, and moreover, the performances of **ALG** become closer and closer to those of **Opt**.

The depth $D(\mathcal{P})$ of a given independent meeting pattern \mathcal{P} is defined as the largest round t for which some sensor observes another sensor. For simplicity, we assume synchronous meeting patterns in which at each round each sensor observes another sensor, but our results can be easily extended to the case where the number of total interactions per sensor are all roughly the depth $D(\mathcal{P})$. Note that for any population with n sensors, the depth of an independent synchronous meeting pattern is at most $\log_2 n$. In particular, the depth is finite for populations of a fixed size. Since our goal is to investigate the behavior of **ALG** at large times, whenever we consider a round t , we only inspect populations and corresponding meeting patterns for which the depth is at least t .

In the remaining of this section we fix a family of distributions \mathcal{F} and a noise distribution $N(\eta)$. Given a round t , let $\text{var}_{\text{sup}}(t)$ denote the supremum of $\text{var}(X_a(t, \mathbf{ALG}))$, taken over (1) all possible populations $A_n = \{a_1, a_2, \dots, a_n\}$, for $n = 1, 2, \dots$, (2) all assignments of distributions $\Phi_a \in \mathcal{F}$ to the sensors in A_n , (3) all meeting patterns (with depth at least t), and (4) all sensors $a_i \in A_n$. Our next claim implies that as time increases, the variance of **ALG** becomes arbitrarily close to zero.

Claim E.1. $\lim_{t \rightarrow \infty} \text{var}_{\text{sup}}(t) = 0$.

Proof. For a round t , let $C_{\text{inf}}(t) = 1/\text{var}_{\text{sup}}(t)$. Note that $\text{var}_{\text{sup}}(0)$ is precisely the maximal variance over the distributions in \mathcal{F} . Hence, $C_{\text{inf}}(0)$ is some positive constant (that depends on \mathcal{F} only).

By the definition of $C_{\text{inf}}(t)$ and by Lemma 5.1, it follows that $C_{\text{inf}}(t)$ is the infimum of $c_a(t)$, the accuracy of a sensor a at round t operating under algorithm **ALG**, taken over all possible populations, all assignments of distributions $\Phi_a \in \mathcal{F}$ to sensors, all meeting patterns, and all sensors a .

When sensor a observes sensor b at round t , the gain in accuracy for sensor a is:

$$\frac{1}{1/c_b(t) + v_N} \geq \frac{1}{2} \cdot \min\{c_b(t), \frac{1}{\text{var}(N(\eta))}\}.$$

It follows that $C_{\text{inf}}(t)$ increases in a single round by either at least a multiplicative factor of $3/2$ or by at least an additive constant factor of $1/2\text{var}(N(\eta))$. This implies that $\lim_{t \rightarrow \infty} C_{\text{inf}}(t) = \infty$. The proof of the claim now follows by the definition of $C_{\text{inf}}(t)$. \square

Since algorithm **Opt** is superior over algorithm **ALG**, the same limit property of the variance applies to algorithm **Opt** as well. We now claim that, in fact, if the noise $N(\eta)$ is Gaussian, then the variances in **ALG** and **Opt** go to zero at roughly the same speed.

Given a round t , let $\kappa(t)$ denote the supremum of the fraction $\text{var}(X_a(t, \mathbf{ALG}))/\text{var}(X_a(t, \mathbf{Opt}))$, taken over all possible populations $\{A_n\}_{n=1}^{\infty}$, all assignments of distributions $\Phi_a \in \mathcal{F}$ to sensors a in A_n , all meeting patterns, and all sensors $a \in A_n$. Note that Theorem 1.2 implies that for any t , we have $\kappa(t) \leq \Delta_0$.

Lemma E.2. *If the noise $N(\eta)$ is Gaussian then $\lim_{t \rightarrow \infty} \kappa(t) = 1$.*

Proof. Since the noise is Gaussian we have $\text{var}(N(\eta)) = 1/J_N$. Recall the definition of $C_{\text{inf}}(t)$ from the proof of Claim E.1. Note now that as $C_{\text{inf}}(t)$ becomes larger and larger the gain in accuracy under algorithm **ALG** becomes very close to J_N . Indeed, when sensor a observes sensor b at round t , we have:

$$c_a(t+1) - c_a(t) = \frac{1}{1/c_b(t) + 1/J_N}.$$

Specifically, consider now the case that $c_b(t) > x \cdot J_N$, for some large x . Here, the increase in accuracy at a is some quantity $\Delta J(t, \mathbf{ALG})$, satisfying

$$\frac{1}{1 + 1/x} J_N \leq \Delta J(t, \mathbf{ALG}) \leq J_N.$$

The Cramér-Rao bound and Lemma 5.1 imply that $J_b(t, \mathbf{Opt}) \geq c_b(t)$, and hence, $J_b(t, \mathbf{Opt}) > x \cdot J_N$. This, together with Theorem 4.2, implies that at round t , the increase $\Delta J(t, \mathbf{Opt})$ in Fisher information of a under algorithm **Opt** is some quantity satisfying

$$\frac{1}{1 + 1/x} J_N \leq \Delta J(t, \mathbf{Opt}) \leq J_N.$$

Hence $\Delta J(t, \mathbf{ALG})$ and $\Delta J(t, \mathbf{Opt})$ are the same quantity up to a multiplicative factor of $\frac{1}{1+1/x}$. Finally, since $\lim_{t \rightarrow \infty} C_{\text{inf}}(t) = \infty$ (see the proof of Claim E.1), it follows that x goes to infinity as t goes to infinity. We thus get $\lim_{t \rightarrow \infty} \kappa(t) = 1$, which establishes the proof of the lemma. \square