DISTRIBUTED CHANGE DETECTION IN LARGE SCALE SENSOR NETWORKS THROUGH THE SYNCHRONIZATION OF PULSE-COUPLED OSCILLATORS

Yao-Win Hong and Anna Scaglione

School of Electrical and Computer Engineering
Cornell University, Ithaca, NY 14853
yh84@cornell.edu; anna@ece.cornell.edu

ABSTRACT

This paper proposes the use of a distributed synchronization mechanism, which locks in phase the pulse-coupled oscillators, to rapidly alert the nodes of a sensor network of a change detected by a group of the sensors. By encoding into an abrupt variation of the phase their positive detection of a change, the nodes force all other nodes to reach a new synchronization equilibrium. Therefore, the information about the change is implicitly encoded in the phase transitions. While the local detection problem at each sensor can be is addressed using the standard change detection algorithms, the interesting aspect of this work is the unconventional way through which the nodes broadcast their information to each other and fuse their decisions. The main advantages of the proposed method is the scalability and low complexity of the fusion algorithm.

1. INTRODUCTION

As sensor technology evolves it becomes apparent that the bottle-neck in using pervasive sensor networks is in the distribution of the sensor information. In the classical setting of distributed change detection problems [2], the nodes are designed either to transmit their detected information to the fusion center, or to base its own decision on the collective information from other nodes. Both of these methods require the information exchange through point-to-point communication links which creates bottle-necks due to the congestion problem. In fact, the centralized model of a fusion center polling the sensors suffers from the intrinsic limitations of the MAC channel, whose aggregate capacity grows as $O(\log(N))$, where N is the number of nodes in the network, resulting in a vanishing per node throughput. Hence, unless the sensors cooperate to send information, the network cannot rapidly convey the information that produces an alarm to the fusion center on the MAC channel. Perfect cooperation would eventually eliminate the need of "fusion" altogether. Unfortunately, cooperation does not appear to be a much easier task since multi-hop networks have an aggregate capacity which grows only as $O(\sqrt{N})$ [1]. Thus, even if the nodes detect exactly the same measurement and agree on what data to share through point-to-point links, the network will still go into congestion if no mechanism reduces the redundancy in the information exchange. It goes without mentioning that all these problems arise even if scheduling and routing issues were solved optimally. But even if the bandwidth were infinitely large and the latency was not a concern, the complexity required to orchestrate the information exchange will grow enormously. Therefore, overlaying a standard communication network architecture over sensor networks is nonapplicable for systems that require low latency, such as those that have to deliver emergency alarms. This suggests that it is necessary to find alternative architectures where the intercommunications and transmissions to the fusion centers are intrinsically unreliable, yet the network itself is capable of fusing the data and conveying the information reliably.

The architecture we propose in this paper is a significant departure from the conventional communication network. The structure of the transmitter is very simple: no routing, no multiple access, only a very simple "pulse position modulation" mechanism. In particular, we assume that the nodes can transmit only through the emission of pulses with constant amplitude (no power control). The information of the sensor data and the interaction among nodes can only be encoded in the timing of the pulse emission. The approach of our fusion technique relies on one classic nonlinear dynamical system problem — the synchronization of pulse coupled oscillators [3]. In this paper, we propose a simple strategy that jointly considers the decision fusion and broadcasting of the detection information. Note that, in general, any non-centralized data fusion algorithm would entail several iterations of data exchange among the nodes where the information exchange is a nonlinear function of the current data. This is to say that cooperative fusion falls into the class of problems that are referred to as *nonlinear dynami-*

This work is supported in part by the National Science Foundation under grant CCR-0347514 "SGER: Scalable cooperative communications in adaptive large scale networks inspired by natural swarms" and ONR Contract N00014-00-1-0564.

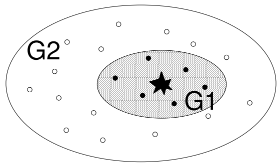

Fig. 1. The detection of a local change in the environment.

cal problems. Our method relies upon the fact that the convergence of the network dynamics has been proven theoretically [3]. In this work, we demonstrate the features of the system numerically while leaving to future work the thorough analysis of the performance. The major advantages of the strategy lies in the scalability and low complexity of the fusion algorithm.

2. PROBLEM SETUP

Consider a network of N wireless sensor nodes randomly distributed in a specified area. The goal of the sensors is to detect an abrupt change in the local environment, such as the attack from enemy forces on a battlefield or the intrusion in secure facilities, and to distribute the information of such changes to the entire network serving as an alarm. As shown in Fig. 1, the change of the environment is often localized to a certain region where only a portion of the sensors are able to detect it. However, this information may be desirable at all locations within the network. Our goal is to develop a simple strategy enabling all the nodes in the network to detect the change in the environment.

Assume that each node in the network obtains a sequence of observations from the environment and has the ability to detect the change of the environment given that the change is within its vicinity. Let $(X_n^{(i)})$ be the sequence of observations obtained by node i. We assume that each sample belongs to one of the two hypotheses:

$$\begin{array}{ll} \mathcal{H}_0: & X_n^{(i)} \sim p_{\theta_0}(x) \\ \mathcal{H}_1: & X_n^{(i)} \sim p_{\theta_1}(x) \end{array} \tag{1}$$

where $p_{\theta_0}(x)$ and $p_{\theta_1}(x)$ represents the distribution of the observation, respectively, before and after the change. In general, when an abrupt change occurs in the environment, each node should observe the change with different performances due to the path loss or fading effects. However, for simplicity, we consider the case where the samples observed by each node belongs only to one of the two hypotheses in (1). There is a vast literature on change detection and an wide set of problems and detectors are studied in detail in [4].

Let the network be divided into two groups where the change of distribution, $p_{\theta_0}(x)$ to $p_{\theta_1}(x)$, occurs only for samples observed by the first group of nodes (G1) and that the samples observed by the second group remains to attain the original distribution. We want to both combine the detection obtained by nodes in G1 and, at the same time, to distribute the information to other nodes in the network. The strategy we propose is based on the synchronization of pulse-coupled oscillators as explained in the following section.

3. PULSE-COUPLED OSCILLATORS

Consider a network of sensors each acting as a pulse-coupled oscillator. In our model, each node in the network transmits replicas of a pulse signal $p(t)$ whose emission is controlled by a state variable. Specifically, for each node i, we define a *state function* $x_i(t)$ which increases monotonically from the initial state 0 to the threshold value 1. When at time τ_i, $x_i(\tau_i^-)$ achieves the threshold value 1, the node immediately emits a pulse $p(t)$ and resets the state variable to 0, i.e. $x_i(\tau_i^+) = 0$. If the node is isolated, meaning that no external signals are received from other nodes, then the state variable $x_i(t)$ follows a periodic pattern increasing as a smooth function of time until it reaches the threshold 1 at which point the function is reset to 0. This results in the periodic emission of a pulse with period T, whose duration depends on how long it takes for $x_i(t)$ to rise from zero to one. In particular, between the time where $x_i(\tau_i^+) = 0$ and where $x_i(\tau_i^+ + T) = 1$, the shape of the state variable is captured by the following function:

$$f_i(\phi) = x_i(\tau_i + \phi T), \quad \phi \in [0, 1], \tag{2}$$

which we call the *dynamics* of the oscillator. We shall refer to τ_i as the firing time of node i and to τ_i/T as the *phase* of the oscillator.

The interaction with other nodes can perturb their periodic pattern as it is explained in the following. We assume that the nodes can receive signals only when it is not firing, and that the firing of an oscillator i will cause an constant increase ε in the state function of every other node j, we refer to this increase as *coupling*. Note that in our idealistic model [3], the pulses from other nodes are received instantaneously and coupling is independent of the relative distances in between nodes. Mathematically, the pulse emission of node i at time τ_i^- changes the state variable of node j as follows:

$$x_j(\tau_i^+) = \begin{cases} x_j(\tau_i^-) + \varepsilon & \text{if } x_j(\tau_i^-) + \varepsilon < 1 \\ 0 & \text{otherwise} \end{cases}, \tag{3}$$

which means that either node j emits the pulse at the same time as node i (now and in the future) or node j will anticipate its own firing time because the residual rise necessary for the state variable $x_j(t)$ to reach the threshold 1 is smaller. Thus, only when the nodes emit the pulse simultaneously will they be insensitive to coupling and therefore behave as oscillators.

The model we just described has been used to explain different synchronization phenomena in large biological networks. In particular, it has explained the flashing of fireflies and the firing of pacemaker cells [3], whose periodic firing of pulses results in synchrony. We wish to remark that a more realistic model would incorporate variable levels of coupling, signal propagation delays and noise. The simplified model is equivalent to assuming that: 1) the pulse duration and propagation delay is short compared to the period T; 2) the path loss among node pairs has negligible differences and the nodes transmit the same power; and 3) the SNR is high. With the simplified assumptions above, Mirollo and Strogatz have proven in [3] the following:

Theorem 1 [3, TH. 3.1-3.2]: *The set of initial states, $x_i(0)\ \forall i$, that never result in synchrony has measure zero, if the function f (defined in 2) is smooth, monotonically increasing and concave down.*

By neglecting the propagation delay and assuming that the nodes operate under the same dynamics, every set of *nodes that are mutually synchronized act as a single oscillator with the coupling strength equal to the sum of the couplings of all the nodes in the set*. We consider the set of synchronously firing nodes to be *absorbed* to each other. The contribution of our paper is exploring the convergence towards synchrony as a mechanism to distribute information throughout the network. This paper utilizes as dynamics, [c.f. (2)], the ones that are provided by the so called Peskin's Model [5].

The state function in the Peskin's model is defined by the following differential equation:

$$\frac{dx_i(t)}{dt} = S_0 - \gamma x_i(t), \quad 0 < x_i(t) < 1 \ \ i = 1, \cdots, N. \quad (4)$$

which is the well-known leaky integrate-and-fire model (IF). S_0 is a constant representing the speed of accumulation when there are no leakage and γ the leakage factor. Therefore, we have

$$f(\phi) = C(1 - e^{-\gamma T \phi}) \quad (5)$$

where $C = 1/(1 - e^{-\gamma T})$ and $T = \gamma^{-1} \ln[S_0/(S_0 - \gamma)]$.

The contribution of our paper is to incorporate the sensor data in the dynamics of each oscillator. By adding to the state variable a perturbation β depending on the sensor data, the synchronization behavior will thus embed the information of the sensor field. Specifically, this information is encoded in the phase at which the synchronization is achieved. Assume, as before, that the coupling between nodes are instantaneous with uniform strength, then we can rewrite the Peskin's model by including the coupling and sensor information as follows:

$$\frac{dx_i(t)}{dt} = S_0 - \gamma x_i(t) + \sum_{j \neq i} \sum_{l} \varepsilon \delta(t - \tau_j^{(l)}) + \beta d(t) \quad (6)$$

where $d(t)$ is a binary function which is 1 where the change is detected and 0 otherwise. $\tau_j^{(l)}$ is the lth firing time of node j and $\delta(\cdot)$ is the Dirac delta function. With this mechanism, a perturbation at a portion ρ of the nodes will effect the synchronized phase which is experienced by all the nodes in the network. Therefore, by properly choosing the constant β, the information of the sensor data can be broadcast throughout the network.

In the case of the change detection problem, assuming that the network is initially synchronized, and an abrupt change is detected at the t_0, we have from (6), for $i \in$ G1,

$$\frac{dx_i(t)}{dt} = S_0 - \gamma x_i(t) + \beta. \quad (7)$$

Hence, all the nodes in G1 experience a sudden change in their phase, moving the firing time from τ_i to $\tau_i + \Delta\phi T$. We call $\Delta\phi$ the *delay* parameter. Note that $\Delta\phi$ and β are one-to-one functions of each other. Since the nodes are no longer synchronized after the imposed phase shift, the pulse-coupling will again synchronize the nodes at a new common phase which differs from that of the original oscillation. We refer to the difference of these two phases as $\Delta\theta$. Interestingly, the phase shift $\Delta\theta$ depends on both ρ and β (see Section 4), and the appropriate choice of β can serve as a tool to broadcast sensor information. In the following, we study through numerical observations the appropriate choice of the phase delay of the oscillator.

4. JOINT DATA FUSION AND BROADCASTING

Consider the case where the offset is chosen to be a delay (in general, the offset may be chosen to anticipate the firing). Define a strategy where each node delays its phase by $\Delta\phi_i$ when a change is detected. Assuming that each node in the set G1 detects the change simultaneously, then the set of nodes will shift its phase by the same amount, and again remain synchronized to other nodes within the set. Since the nodes in G2 do not detect the change and remain at its original phase, we will observe two groups of nodes that are synchronized within the group, but not synchronized among each other. As mentioned previously, the set of nodes that are synchronized are absorbed to each other and act as a single oscillator with sum coupling strength. Therefore, the case reduces to the synchronization problem of two users with non-uniform coupling.

Assume that each individual node in the network receives uniform coupling equal to ε and that the sum coupling of the entire network is $\xi = \varepsilon N$. Define the vari-

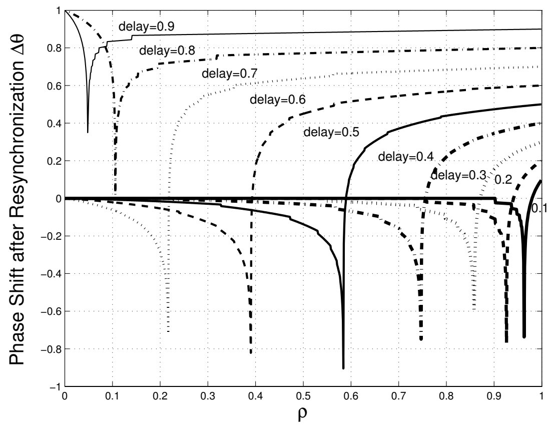

Fig. 2. The phase shift $\Delta\theta$ versus ρ. The legend "delay" represents the variable $\Delta\phi$.

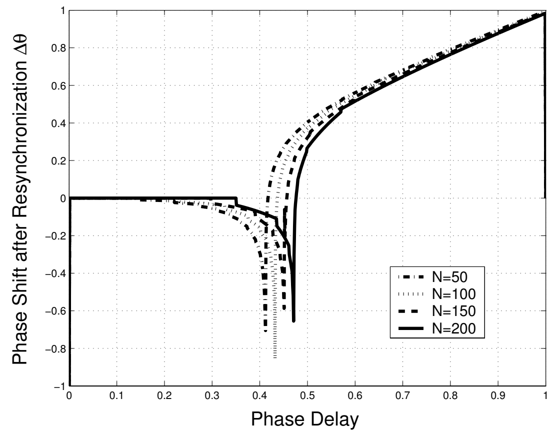

Fig. 3. Fix $\rho = 0.7$. The phase shift after re-synchronization versus the phase delay for $N = 50, 100, 150$ and 200.

able $\rho = \#\{G1\}/N$, where $\#\{G1\}$ represents the number of nodes in G1. If the detection of the emitted pulses can be detected reliably at each node, every initial phase difference $\Delta\phi$ of the two groups can be mapped to a certain phase shift $\Delta\theta$. For the parameters $r = 4.9$, $S_0 = 5$, $N = 100$ and $\xi = 0.1$, we plot numerically, in Fig. 2, the phase shift $\Delta\theta$ against the parameter ρ for phase delay $\Delta\phi = 0.1, 0.2, \cdots, 0.9$. We can see that when the phase delay is small, the pulsing of nodes in G2 will pull the nodes in G1 to fire after the first firing, therefore, zero phase shift may occur. However, when the phase delay is large, the pulsing of G2 will only anticipate the pulsing of G1, while the subsequent pulsing of G1 will pull G2 towards its own phase, resulting in a detectable phase shift when synchronization is achieved again. When the value of ρ is sufficiently large and the delay is appropriately chosen, the phase shift $\Delta\theta$ will be detected serving as a form of broadcast.

For every choice of the delay, there is a region of ρ for which the phase shift saturates towards a constant value and the change is detected. This region is roughly determined by a threshold caused by the nonlinear effect of the system (the deep valley of the curve). For cases where ρ is less than the threshold, the negative feedback of the system eliminates the detection since it is insignificant in the sense that only a small portion of the network detects the change performing, de facto, a fusion of detections. In Fig. 2, we considered the case where all the nodes in G1 agree on a common phase delay, however, we can observe that the thresholding effect of the curves can serve as a way of encoding the reliability of the sensed data at each node. Specifically, for a node that has relatively less reliable data, it delays its phase by a smaller amount than the nodes that have a reliable data. In the case where all the nodes are relatively less reliable, more nodes would be needed for the phase shift to occur at a saturated value, i.e. a detectable value. In this way, the nodes can avoid the false alarms produced by faulty detections at each node. This is a natural way of implementing data fusion through the adjusting the distributed adaptive threshold. In Fig. 3, we show that the curve of the phase shift $\Delta\theta$ over the phase delay has small variations when the scale of the network increases, especially when the phase delay is appropriately chosen. If the saturation level does not depend on the number of nodes then this is a scalable method to convey information about the amplitude of the sensed data. Further work is needed to establish the performance and to analyze the effect of a less idealistic propagation scenario.

5. REFERENCES

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inform. Theory*, vol. IT-46, no. 2, pp. 388–404, Mar. 2000.

[2] Venugopal V. Veeravalli, "Decentralized quickest change detection," *IEEE Trans. Inform. Theory*, vol. 47, no. 4, pp. 1657–1665, May 2001.

[3] Renato E. Mirollo and Steven H. Strogatz, "Synchronization of pulse-coupled biological oscillators," *SIAM Journal on Applied Mathematics*, vol. 50, no. 6, pp. 1645–1662, Dec. 1990.

[4] Michéle Basseville and Igor V. Nikiforov, *Detection of Abrupt Changes: Theory and Applications*, Prentice Hall, 1993.

[5] C. S. Peskin, *Mathematical Aspects of Heart Physiology*, Courant Institute of Mathematical Sciences, New York University, New York, 1975.