Information storage and transfer in the synchronization process in locally-connected networks

Rommel V. Ceguerra*, Joseph T. Lizier*†, Albert Y. Zomaya*
*School of Information Technologies, The University of Sydney, NSW 2006, Australia
†Max Planck Institute for Mathematics in the Sciences, Inselstraße 22, 04103 Leipzig, Germany
Email: lizier@mis.mpg.de

***Abstract*—Synchrony is a phenomenon where a group of weakly interacting oscillators demonstrate the ability of mutual entrainment, organizing themselves to reach a highly ordered state. The study of synchrony has been hindered by the inherent nature of complex systems requiring the system be studied at a global level. To create new insights here, we use information-theoretical techniques to view the synchronization process as a distributed computation, and to measure the information storage and transfer at each node in the system at each time step. These application-independent measures provide results that can be compared to other systems, and also study the local dynamics of synchronization within the system. This has produced novel insights, including that the computation of the synchronized state appears to be completed much more quickly than application-specific measures (such as the order parameter) would indicate. We also found distinct differences in the computational properties of various nodes, allowing us to develop a theory of the computational dynamics of the synchronization process.**

I. Introduction

Synchrony has been independently observed and studied in many different fields as seemingly unrelated phenomena for many years [1], e.g. in swarms of flashing fireflies, clusters of pacemaker cells in our own heart, and electrons in a superconductor [2]. The ability of these non-intelligent, and even inanimate, systems to exhibit apparently intelligent behavior has intrigued researchers [2], but only recently has the study of all these phenomena been unified, making synchrony a universal concept in complex systems science.

The understanding of synchrony, the conditions where it is exhibited and the process by how it is achieved, have wide ranging possible applications. Quantum synchrony has already found applications in lasers, superconductivity, and Josephson junctions allowing extremely sensitive magnetometers. In medicine, epilepsy has been found to be caused by excessive synchronous neuronal activity [3], and furthering the understanding of why these neurons begin to fire simultaneously could lead to better treatments or even cures. The ubiquity of distributed synchronization processes in biological systems (as opposed to centralized mechanisms) is theorized to be due to its increased robustness to damage [2]; taking heed of this, the tendency for technology to mimic natural systems is reflected in the trend towards distributed self-organizing systems in I.T. and robotics. Gaining further understanding of this phenomenon could lead to much more efficient and reliable implementations of these technologies.

The success of the Kuramoto model [4], [5] has been a significant driver of the interest in synchrony. Importantly, (for most connection topologies) this model exhibits a phase transition from unsynchronized behavior to robust synchronization as the coupling strength between the elements is increased. Yet while the behavior of some special topologies (e.g. fully-connected systems) can be identified by tractable analytic solutions, in general other topologies yield intricate behavior, and it is non-trivial to determine whether a given system will support synchronization (aside from running a full simulation).

A particular hindrance to the study of complex systems is that by their inherent nonlinear nature, the entire system must be considered simultaneously, thus little is known about the relative importance of individual nodes within a system. This is particularly true for synchrony, since as described above the behavior of most systems is entirely non-trivial. Indeed, most work so far studying synchrony has taken the holistic approach, analyzing properties of the system as a whole and their effect on synchronizability. In this way the actual synchronization process is somewhat taken as a black box with the initial state of the system as an input and synchrony as an output. We will take a much more atomistic approach and thus study the *local* dynamical processes occurring during the synchronization process at an individual level.

To do this, we use an emerging technique in complex systems science of the perspective of analyzing the distributed computation undertaken by a system in its unfolding dynamics. Stemming from information theory, *local information dynamics* [6]–[10] focuses on quantifying how information is stored and transferred within a distributed system, and in particular how these operations vary in time and across the system. In this paper, we seek to study the information dynamics in synchronization processes in locally-connected networks.

The information-theoretic basis of the approach means that information dynamics can view any multivariate dynamical process as a distributed computation, regardless of the underlying governing dynamics. As such, insights gained regarding the information dynamics of synchronization processes can be directly compared to the distributed computation undertaken by other complex systems (e.g. neural networks), possibly

revealing unifying processes in what are currently thought as unrelated phenomena. This is simply not possible with application-specific measures such as the traditional order parameter o representing how well a system is synchronized.

We will utilize the capabilities of information dynamics to analyze a system achieving synchrony, exploring the role of individual components of the system in the synchronization process from a distributed computation perspective.

We begin in Section II by describing the phenomenon of synchrony and the Kuramoto model used for exploration here. In Section III we then introduce the measures used here to quantify information storage and transfer in these systems, specifically the active information storage and transfer entropy. We then apply such analysis to the synchronization model, and describe the results in Section IV. Interestingly, these measures reveal that the distributed computation appears to be complete in a much shorter time period than application-specific measures (i.e. the order parameter o) would suggest. We also show that during the phase transition to synchronized behavior, information storage falls while transfer increases, reflecting decreased unilateral behavior and increased co-ordination between nodes. We explore how the information measures vary across the network, finding relationships between topological and information properties of nodes, and allowing us to posit a qualitative theory on how the information flow in the system brings about the synchronized state.

II. Synchrony

The most general definition of synchrony is *the mutual entrainment of rhythms of oscillators by their weak interactions* [1]. The most common end result of synchronization is synchronous motion, where the frequency and/or phase of every object in the system is equal, called frequency locking and phase locking respectively. Not all systems which exhibit synchronous motion however exhibit *synchrony*, thus we need to make the distinction between these two terms. To exhibit synchrony, each oscillator must be autonomous and able to oscillate independently with the complete absence of interactions with the rest of the system. Thus in a system where two objects are oscillating synchronously, yet one object is completely dependent on another to oscillate, the system does not exhibit synchrony, as total dependence implies hierarchy. Similarly, the interaction must not be strong enough such that the system can be defined as unified. Thus a system where two objects are absolutely limited in their movements such that it is only possible for them to oscillate in unison, while demonstrating synchronous motion, does not exhibit synchrony. An example is if two pendulums are mechanically joined by a completely rigid link, the pendulums will always move as one, each completely dependent on the other with no individuality, a trivial process.

A system exhibiting synchrony must undertake a non-trivial dynamical process of synchronization to achieve synchronous motion. In the total dependence and limited motion scenarios above, one need only examine one component of the system to determine the behavior of the entire system. The interest in synchrony lies in the complex interactions of all components of the system to come to a global agreement without hierarchy.

A classic example an audience clapping in unison. Without a single leader, individuals are influenced by what they hear and also influence others, and all quickly settle into a common rhythm. This displays all the hallmarks of synchrony: each person has their own normal clapping speed, no one is told explicitly to follow another person, and the audience begins in a disordered state but undertakes a complex process of mutual influence to achieve global synchronization.

Many models have been developed to demonstrate and study synchrony based on the various real systems observed to exhibit it [1], [11], but none as widely used and successfully studied [12] as the Kuramoto model [4], [5]. This model is general and flexible enough to be adapted to many different contexts, simple to simulate, and yet complex enough to exhibit non-trivial behavior. The model consists of a population of P phase oscillators, coupled in a network topology described by the adjacency matrix A_{XY}.[1] Each oscillator (or node in the network) has a natural frequency ω_X (distributed by a probability distribution function $g(\omega)$) for the time derivative $\dot{\theta}_X(t)$ of its phase $\theta_X(t) = [0, 2\pi)$. In isolation, each node X will intrinsically oscillate at ω_X. In the network however, coupling between neighbors in A_{XY} influences the dynamics of phase differentials $\dot{\theta}_X(t)$ of a node X at time t as governed by the following equation:

$$\dot{\theta}_X(t) = \omega_X + K \sum_{Y=1}^{P} A_{XY} \sin\left(\theta_Y(t) - \theta_X(t)\right). \quad (1)$$

Here K represents the *coupling strength* between neighbors in A_{XY}. The linearly-summed coupling terms $\sin\left(\theta_Y(t) - \theta_X(t)\right)$ ensure that neighbors Y with a phase $\theta_Y(t)$ leading that of X prompt faster oscillation of X, while neighbors with a lagging phase prompt slower oscillation of X. Kuramoto also gave a convenient way to measure the coherence of the system at an instant t in time, the *complex order parameter* given by:

$$o(t)e^{i\phi(t)} = \frac{1}{P} \sum_{X=1}^{P} e^{i\theta_X(t)}. \quad (2)$$

Intuitively, if the phase of each node is plotted on the unit circle, the distance of the centroid of the population from the center of the unit circle would be the *order parameter* $o(t)$, with average phase of the population being $\phi(t)$. An incoherent system would have the population uniformly distributed around the circle so $o = 0$, and where the nodes are clustered on top of each other and moving as one, i.e. synchronized, $o = 1$.

Crucially, the general behavior of the model is fairly consistent despite changes in the definition of $g(\omega)$ or (network topology) A_{XY}, or even the addition of noise [12]. Below a critical coupling K_C the system is *incoherent*, no nodes

[1] Here we use binary, symmetric values for A_{XY}, though in general one could use any $A_{XY} \in \mathbb{R}$ to describe a directed network with weighted links.

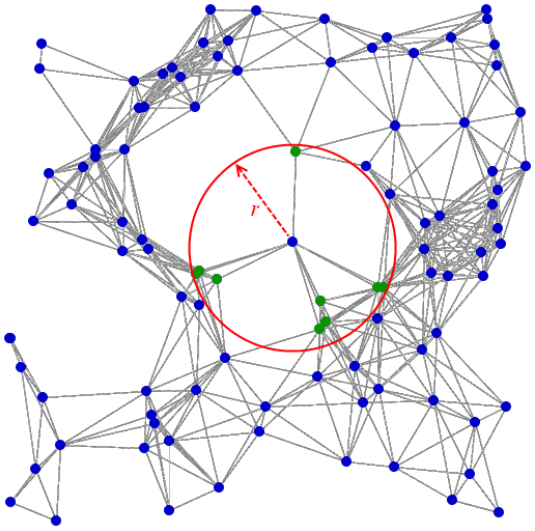

Fig. 1. Sample instance of the scattered, locally-connected network topology

are able to form a synchronized group. As K increases past K_C a phase transition occurs and the system exhibits *partial coherence*, as more *rogue nodes* are able to join the coherent group until it comprises the entire population at K_P (for finite P), and the system is completely *coherent*.[2]

The model we are using to define the adjacency matrix A_{XY} is a locally-connected topology derived from wireless sensor networks [13], or "unit disk graph" [14]. It has been shown that locally-connected topologies have a tendency to hinder synchronization compared to other random topologies with long range connections [15]–[17].

The topology is generated by first scattering nodes uniformly in $[0, 1]$ on a two dimensional plane; this is analogous to the deployment of a wireless sensor network in the field. Every node has the same communication range r, and every node within that distance is bi-directionally coupled, thus $A_{XY} = A_{YX}$ with $A_{XY} = 1$ for a symmetrically coupled pair and $A_{XY} = 0$ for uncoupled. As we see in Fig. 1 this generally results in a roughly lattice-like structure, but with highly connected clusters and outlying groups barely connected to the main body.

In our implementation of the model, there are $P = 100$ nodes, the initial phases are uniformly distributed in $[0, 2\pi)$, and the native frequencies use a Gaussian distribution for $g(\omega)$ with $\mu = 3$ and $\sigma = 1$. We use the Euler method [12] to numerically approximate the governing equations at varying values of dt to generate $N = 1000$ discrete time steps n (observations) in the relevant time interval including the entire transient stage and approximately the same time interval of the settled state. Note that continuous time t in the model is not the same as the discrete time step n (instead $t = n.dt$).

[2] With a *distribution* of natural frequencies rather than a common value, it is impossible for perfect synchrony to occur ($o = 1$). Instead, we describe effective synchronization as a *coherent* state, where $o \to 1$ and remains stable, because the coherent group encompasses the entire population.

III. Information Dynamics

Information dynamics refers to the storage, transfer and modification of information by the elements or agents undertaking a distributed computation, and in particular focuses on how these operations vary in time and space [6]–[9]. The measures for these operations are built on information theory [18], and provide a common, non-linear, application-independent language in which to analyze and design complex computation. In this paper, we will focus on measuring information storage and transfer in synchronization processes, where the elements compute if and how they will synchronize their frequency and phases. In this section, we describe the measures used to quantify these operations.

A. Information storage

Information storage refers to the amount of information in the past of a variable that is relevant to predicting its future. The *active information storage* measures the stored information that is *currently in use* in computing the next state of a variable [8]. Specifically, the *local active information storage* for variable X is the local (or un-averaged) mutual information between its semi-infinite past $x_n^{(k)} = \{x_{n-k+1}, \ldots, x_{n-1}, x_n\}$ (as $k \to \infty$) and its next state x_{n+1} at time step $n + 1$:

$$a_X(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_{n+1} \mid x_n^{(k)})}{p(x_{n+1})}. \tag{3}$$

Finite-k estimates are represented as $a_X(n, k)$. The active information storage is the average over time: $A_X = \langle a_X(n) \rangle_n$. Both average and local measures are given in bits. The local active information storage $a_X(n+1)$ is the stored information that is currently in use by variable X in computing its next state x_{n+1} at time $n + 1$.[3] The average value A_X will always be positive, but is limited by the average entropy $H_X = \langle -\log_2 p(x_{n+1}) \rangle_n$ (i.e. the amount of information in the next state x_{n+1}).

The dynamics of the local active information storage have been studied in Cellular Automata (CAs) [8], finding that blinkers and regular domains were the dominant information storage structures there. The average active information storage was also studied for Random Boolean Networks (RBNs) [20], where it was shown to be maximized close to the critical state in the order-chaos phase transition of RBNs.

B. Information transfer

Information transfer is formulated by the *transfer entropy* [21] (TE) as the information provided by a source about a destination's next state that was not contained in the past of the destination. Specifically, the *local transfer entropy* [7] from a source Y to a destination X is the local mutual information

[3] This contrasts with the (predictive information form of the) *excess entropy* [19], which measures the average *total* information stored by the variable (which will be at some point in the future, though not necessarily at the next state). The active information storage, measuring only that which is currently in use, represents a portion of or lower bound to the excess entropy [8], [10]. We focus on the active information here since it yields an immediate contrast in the relative contributions of storage and transfer to each computation of the next state of a variable (see [9]).

between the previous state of the source y_n and the next state of the destination x_{n+1}, *conditioned* on the semi-infinite past of the destination $x_n^{(k)}$ (as $k \to \infty$):

$$t_{Y \to X}(n+1) = \lim_{k \to \infty} \log_2 \frac{p(x_{n+1} \mid x_n^{(k)}, y_n)}{p(x_{n+1} \mid x_n^{(k)})}. \quad (4)$$

Again, the transfer entropy is the (time) average $T_{Y \to X} = \langle t_{Y \to X}(n) \rangle$, while finite-$k$ estimates are represented as $t_{Y \to X}(n+1, k)$. The TE can be measured for any two time series, but only represents information transfer when measured on a causal link [22] and in the limit $k \to \infty$ [7]. Importantly, the TE properly measures a directed, dynamic flow of information, unlike previous inferences with the mutual information which measure correlations only. Also, in measuring the contribution of the source in the context of the past of the destination, it specifically juxtaposes the information transfer from that source against the storage of the destination. That is to say, conditioning on the past of the destination removes any stored information from being considered as transfer [9]. In a similar fashion to A_X, the average TE $T_{Y \to X}$ is always positive but limited by the average entropy H_X.

The dynamics of the local TE have been studied in (CAs) [7], finding that gliders and domain walls were the dominant information transfer agents as had long been postulated. The average TE was also studied for Random Boolean Networks (RBNs) [20], where it was shown to be maximized close to the critical state in the order-chaos phase transition of RBNs.

IV. Results and discussion

In this section, we present the results of our information dynamics measurements on the Kuramoto model with locally-connected nodes. While all of the following results are taken from a single randomized instance of the model, the general nature of all of the main results described were verified on at least 5 repeated experiments with different randomized instances from the same parameters.

The information-theoretic measures were computed on the time-series of differentials $\dot{\theta}_X$ generated by our model for each node X, over the N transient time steps here. In this manner, we calculate how much of the change in phase at each time step for a given node can be predicted from its own past (A_X), and how much can be predicted from each of its neighbors that was not in its past ($T_{Y \to X}$). Specifically, the active information storage is computed for each node in the network; we refer to the average of the local values $a_X(n)$ across all nodes X for a given time step n as $\overline{a_X(n)}$, and the average across all nodes (for all time steps) as $\overline{A_X}$. Furthermore, the transfer entropy is computed for each connected pair of nodes, using $\dot{\theta}_X$ as the destination time series and the differences $\theta_Y - \theta_X$ as the source time series. We refer to the average of the local values $t_{Y \to X}(n)$ across all edges $Y \to X$ for a given time step n as $\overline{t_{Y \to X}(n)}$, and the average across all edges (for all time steps) as $\overline{T_{Y \to X}}$. Since the data set is continuous-valued, the measures were computed using kernel estimation [21], [23], with a unit step kernel with width of 1 standard deviation per variable. Also, while large values of past history k are desirable for each measure, through experimentation we found that increasing values k larger 2 had little effect, indicating that there is limited conditional reach of the past on the future phase differentials. For the $N = 1000$ discrete time steps for each run of the model, using $k = 2$ and kernel width of 1 standard deviation per variable satisfies a desire for at least 10 observations on average per joint tuple $(x_{n+1}, x_n^{(k)}, y_n)$ (meeting the suggestion for adequate sampling in [24]).

A. Local information dynamics in time

First we examine how the computation unfolds in time as the system attempts to synchronize. Fig. 2 displays the progression of the conventional order parameter $o(t)$, for a coherent system in Fig. 2a and for a critical or partially coherent system in Fig. 2b. As we see in Fig. 2a for the coherent system, the order parameter shows coherence is effectively achieved at $t \approx 3$. In comparison, Fig. 2c shows that the local information transfer $\overline{t_{Y \to X}(n)}$ drops to almost zero *much* earlier, at $t \approx 1$.[4] This difference is even more pronounced in the partially coherent system, where in Fig. 2b $o(t)$ shows the settled state is achieved at $t \approx 10$, but Fig. 2d shows effectively no information transfer after $t \approx 2$. Information storage exhibits a similar early dwindling (not shown), though approaches a finite value rather than zero and reaches this in a slightly longer time period than the transfer.

The result that information transfer would vanish in the coherent state is expected, because the nodes' behavior are then predictable from their own pasts.[5] This does not mean that the nodes no longer have a causal effect on eachother; it is a statement from a computational perspective that the distributed computation is effectively complete and the nodes are executing trivial information storage processes. The more interesting result is the indication that *the distributed computation by the system which makes synchrony possible is complete much earlier than it would appear synchrony is actually achieved*, metaphorically just pointing the system in the direction of synchrony then letting it go to run its course. This is a particularly important result, since it cannot be observed using the conventional application-specific measures of synchrony. These information dynamics measures could thus be useful as early indicators that an explicitly settled state will soon be reached. Further work is required to determine whether the completion of the distributed computation corresponds to any conventionally-understood part of the synchronization process.

B. Information dynamics through incoherent-coherent phase transitions

We then examine in Fig. 3 how the average information dynamics in the network vary through the phase transition between coherent and incoherent systems as the coupling strength K is altered. Fig. 3a shows the conventional order parameter for this data, with incoherent behavior for $K < K_C$

[4] By $t \approx 2$ the value of $\overline{t_{Y \to X}(n)}$ is indistinguishable from the values obtained from the null data sets described in Section IV-B.

[5] For a distribution of ω_X it is less likely to have exactly zero transfer in the coherent (rather than perfectly synchronized) state.

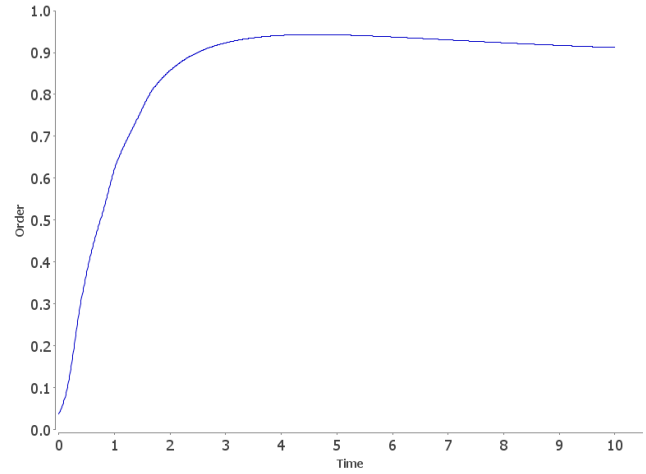

(a) Order $o(t)$ in a coherent system

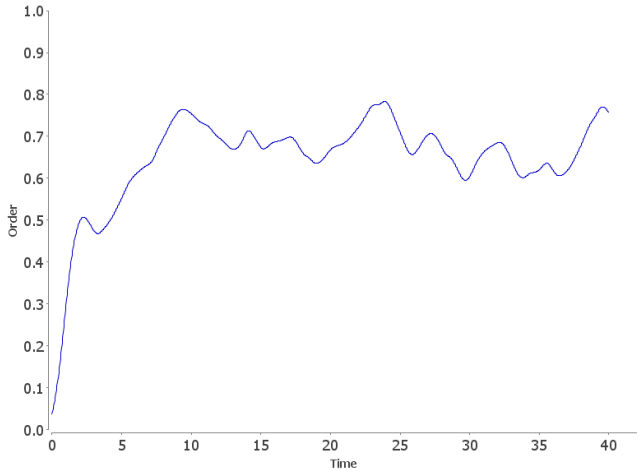

(b) Order $o(t)$ in a partially coherent system

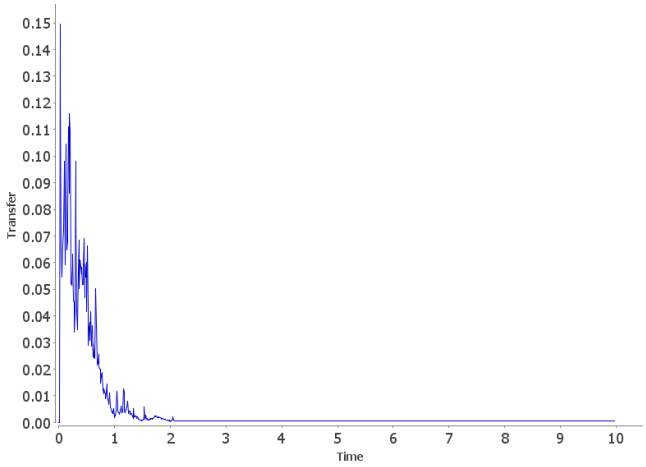

(c) Mean $\overline{t_{Y\to X}(n)}$ of a coherent system

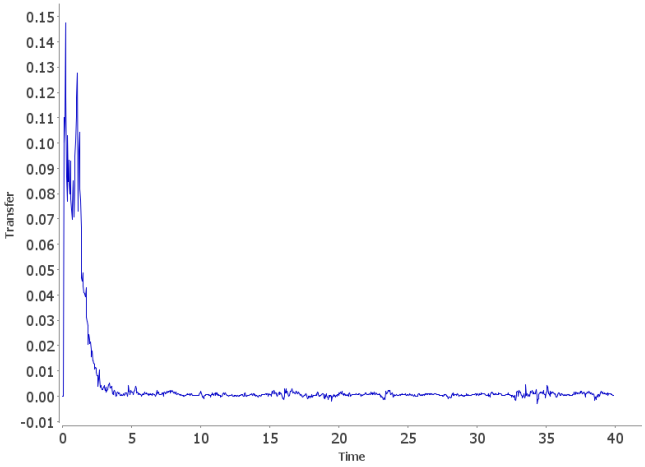

(d) Mean $\overline{t_{Y\to X}(n)}$ of a partially coherent system

Fig. 2. Comparison of the progression of order parameter $o(t)$ and mean local transfer across all edges $\overline{t_{Y\to X}(n)}$ in the system, plotted against model time t. In this system range $r = 0.22$, and $K = 0.2$ for the partially coherent system and $K = 0.6$ for the coherent system

($K_C \approx 0.16$), then extremely unpredictable partial coherence for $K_C < K < K_P$ ($K_P \approx 0.44$),[6] and effectively full coherence for $K > K_P$. Here settled state order is the average over N time steps at some time interval after the transient stage has ended. Error bars indicate the standard deviation of o in the settled state in this time interval, and thus indirectly show fluctuations indicating partial synchrony, thus at $K \approx 0.44$ error bars stabilize indicating coherence. Fig. 3b and Fig. 3c show that in the incoherent phase the information properties of the system are fairly steady, then once $K > K_C$ active information storage starts to decrease and information transfer starts to increase. This reflects the fact that in the partially coherent phase, more nodes have their dynamics more strongly influenced by their neighbors (higher transfer), while fewer nodes exhibit dynamics largely predictable from their pasts (i.e. lower storage). For the coherent phase with $K > K_P$, the settled state order parameter plateaus, so one might expect the level of computation to plateau as well. This is because we already know from Fig. 2 that the local information transfer vanishes (i.e. the computation is complete) once complete synchrony is achieved. An interesting observation however is that the information dynamics continue their respective trends of increasing or decreasing with the level of coupling. In fact, where $K > K_P$ increasing K further slightly accelerates the synchronization process, i.e. synchrony is achieved at an earlier time. It may be this faster synchronization that is reflected in the information transfer continuing to increase.[7]

Importantly also, the computation of the phase differentials at each node is dominated by information storage rather than transfer. This is shown by the larger values of $\overline{A_X}$ than $\overline{T_{Y\to X}}$

[6] The extreme unpredictability in the partial coherence regime means that slight increases in coupling K do not necessarily lead to increased order ω.

[7] Information storage and transfer were maximized near the critical state in order-chaos phase transitions in RBNs in [20]; since this is not a corresponding order-chaos phase transition there is no reason to expect similar results.

in Fig. 3b and Fig. 3c, and we see in Fig. 3b that almost all of the entropy or information $\overline{H_X}$ required to predict the next phase differential is contained in the node's past ($\overline{A_X}$).

Furthermore, we determine the *statistical significance* of these measurements by determining what they would look like if there were no relationship between past and future (for A_X) or no relationship between source and destination (for $T_{Y\to X}$); i.e. under a *null hypothesis* of no such relationship.[8] In Fig. 3b and Fig. 3c the data series with prefix "nh" indicate the distributions under the null hypotheses, validating that the storage and transfer measured here were statistically significant compared to these distributions. The important result is that although the values measured for transfer were low, they still indicate statistically significant relationships.

Finally, note that a similar incoherent-coherent phase transition can be observed by varying the range of interaction r with constant coupling K. Through this phase transition, we observe similar results of decreasing storage and increasing transfer as the system moves from incoherent behavior at low r to coherent behavior at large r (results not shown here).

C. Relationship between information dynamics and network topology

Fig. 4 shows each node in the context of the topology, with the relative magnitudes of their information properties plotted. There are several trends we can observe for this coherent network; importantly, none of these trends are observable in incoherent networks.

Most strikingly, the outermost nodes show high amounts of active information storage A_X, with the highly connected central cluster displaying heavy average incoming transfer $T_X^i = \langle T_{Y\to X}\rangle_Y$ and average outgoing transfer $T_Y^o = \langle T_{Y\to X}\rangle_X$. Interestingly the nodes directly surrounding this central cluster show high outgoing but low incoming transfer. The central cluster appears to be a *computation core*, where most of the important computation in achieving synchrony occurs with a large amount of mutual communication between nodes, and the outer edges of the cluster form a *communication shell*, mostly communicating outwards to the rest of the network. The outermost nodes, being most removed from this computation, show the largest degree of independently predictable dynamics (via large A_X).[9]

We have also observed that bottleneck nodes (e.g. "B" in Fig. 4a) have a tendency towards higher storage than surrounding nodes. This may seem counter-intuitive as one may expect such nodes connecting different clusters to have large

[8] For the transfer $T_{Y\to X}$ for example, one creates a set of surrogate source time series Y' with random permutations of the order of states in the original Y, then computes a distribution of surrogate transfer values $T_{Y'\to X}$ for each source Y' in this set [25], [26]. All probability distributions remain the same, except for those relating the source to the destination (i.e. $p(x_{n+1} \mid x_n^{(k)}, y'_n)$), thus any source-destination relationship is destroyed in the distribution of $T_{Y'\to X}$ while preserving the characteristics of the destination time series X. $\overline{T_{Y'\to X}}$ is plotted as "nhT" in Fig. 3c.

[9] The roles of nodes here have echoes to the use of k-shell decomposition in [27] to analyze the roles of various nodes in a disease spreading model (where the nodes at the core of the network were identified as more effective starting points for the disease).

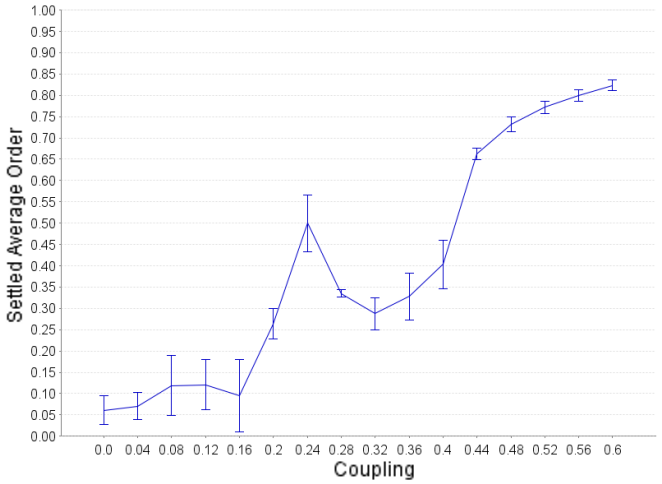

(a) Settled state order parameter σ

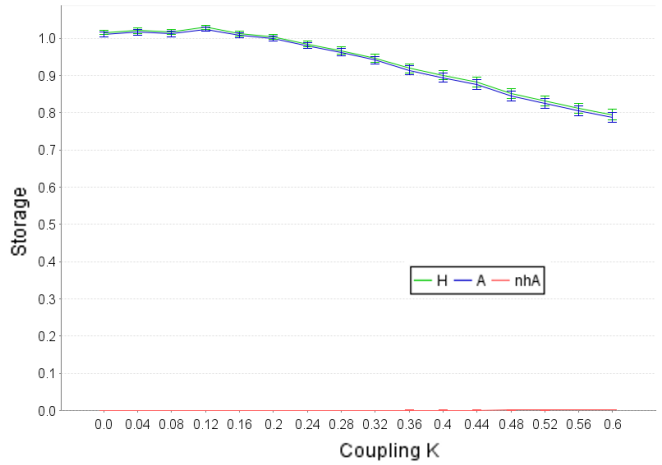

(b) Average entropy $\overline{H_X}$ and storage $\overline{A_X}$ per node

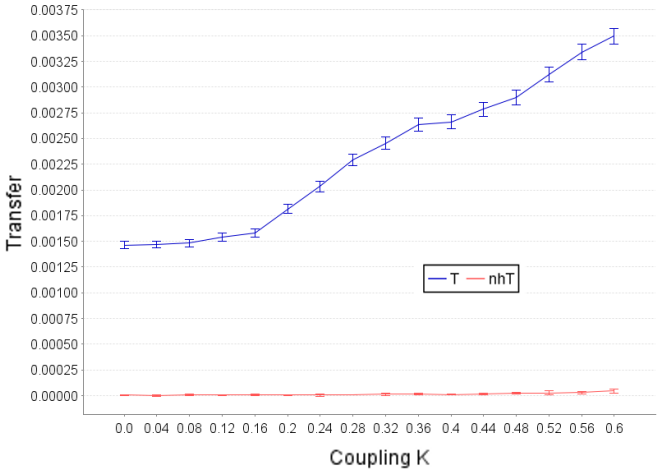

(c) Average transfer per edge $\overline{T_{Y\to X}}$

Fig. 3. Comparison of order parameter and mean information dynamics of the system versus coupling K (with range $r = 0.18$). Error bars indicate (a) standard deviation of fluctuations in o, (b) standard error of the mean over all nodes in the system, (c) standard error over all edges in the system. The plots of $\overline{A_X}$ and $\overline{T_{Y\to X}}$ under null hypotheses of no source-destination relationship are almost indistinguishable from zero.

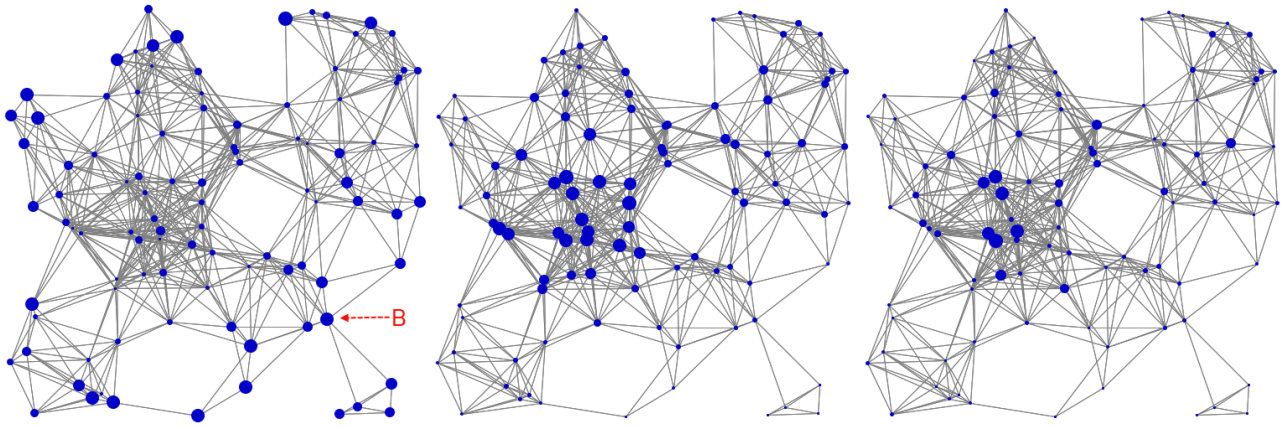

(a) Storage A_X (b) Average outgoing transfer T^o_Y (c) Average incoming transfer T^i_X

Fig. 4. Relative magnitudes of individual information dynamics at each node over the topological map. In (a), "B" marks a bottleneck node. These results are for a coherent network, with $r = 0.22$, $K = 1.0$.

amounts of transfer to allow these clusters to communicate and synchronize. However, this result suggests that these bridges could be more important in determining the synchronized state than the clusters, as they are less flexible and thus the clusters would have to adapt to them. This supports work showing that increased local synchrony can discourage global synchrony [28], [29], otherwise information would flow freely through these bridges. This *blocking bridge* observation goes against the common assumption that as clusters grow and meet, they merge coming to an agreement with each other [30]. It is possible that when clusters collide, conflicting information from both clusters *may* overload the bridging nodes forcing the clusters to instead adapt to the inflexible minority. This introduces the possibility that there are two types of high storage nodes, those which are intrinsically inflexible but eventually succumb to the influence of the coherent group, and those which are inflexible to their surroundings, forcing influence onto the coherent group.

Finally, we explore the relationship between the information dynamics at each node and the degree of the node in Fig. 5. Importantly, Fig. 5b shows a strongly significant positive correlation of average outgoing transfer T^o_Y to the degree of the node. This was also noted in a similar study of information transfer in cascading failure phenomena [31]. Since high degree nodes have a large diversity of inputs, they presumably have more information available to transfer to other nodes. The result aligns with our interpretation of Fig. 4 above also, since the higher transfer nodes in the center of the network will be those with higher degree. Also, Fig. 5a shows a significant negative correlation of storage to degree. Again, this aligns with our interpretation of Fig. 4, since the outermost nodes with more independent behavior are those with lower degree. There was a weaker positive relationship between degree and incoming transfer (not shown), partly due to its inability to differentiate between the hypothesized computation core and the communication shell among the high degree nodes.

V. Conclusion

Gaining understanding of synchrony is of great consequence due to its fundamental presence in nature, and its possible applications in a wide variety of technological fields. In taking a new perspective of framing the dynamical processes as distributed computation, we have revealed significant insights into the phenomenon which are undetectable by other measures.

Perhaps most interestingly we measured the computation required to achieve synchrony to be effectively complete much earlier than partial or global coherence is achieved (as shown by existing measures). This indicates the non-trivial aspect of the synchronization process to be much shorter than initially thought, and suggests information dynamics could be used for pre-detection of synchrony. We also found information transfer to increase with coupling, even once coherence is achieved, reflecting the acceleration in the synchronization process.

Importantly, we found distinct differences in the information properties of various nodes, and formed a theory of the distributed computation of the synchronization process. Local clusters of synchrony originate in computation cores of highly connected nodes, with high values of outgoing and incoming transfer. This computation core is surrounded by a communication shell of nodes with high outgoing transfer but comparatively low incoming transfer. This shell seems to communicate the computed coherent state to the rest of the system. We also observed "blocking bridge" behavior where bottlenecks tend towards higher storage than surrounding clusters. This suggests the possibility that conflicting information from different clusters could overload the bridge nodes, and while the clusters do appear to merge, this may be because they are forced into agreement with the bridge.

In future work, we will apply these techniques to other network topologies (e.g. fully-coupled or regularly-connected). We will also use these tools to explore the effect of network damage on synchrony, and in particular into the importance of individual nodes to the synchrony of the system as a whole.

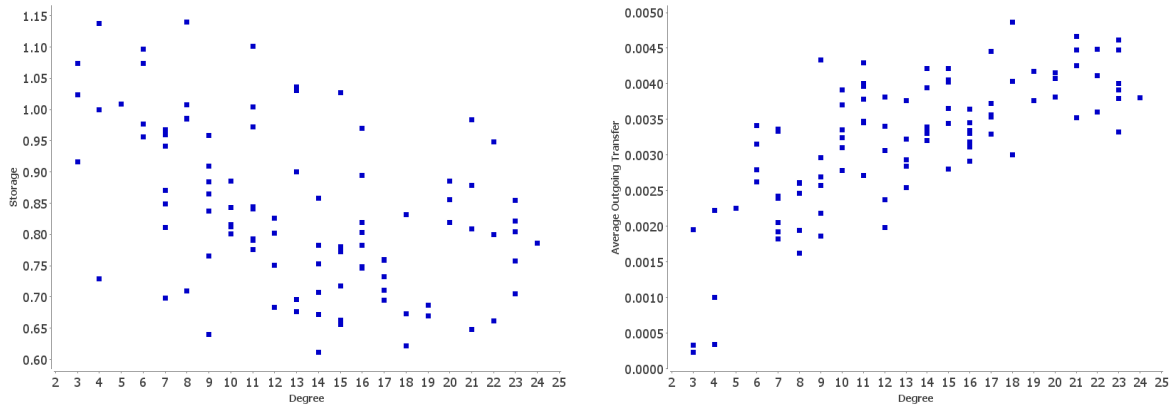

(a) Storage A_X: correlation $= -0.4459$, p-value < 0.0001 (b) Average outgoing transfer T^o_Y: corr. $= 0.7054$, p-value < 0.0001

Fig. 5. Correlation of information dynamics to node degree, and one-sided p-values from t-tests of the statistical significance of these correlations. These results are for a coherent network, with $r = 0.22$, $K = 0.6$.

ACKNOWLEDGMENT

The simulations and analysis were performed using the cluster of the Advanced Networks Research Group at the School of Information Technologies, The University of Sydney. Figures were generated using *JFreeChart* by Gilbert, and statistical distributions were calculated from *Statistical Distribution Library* by Steinmetz & Warnes. This work was supported by Australian Research Council grant DP1097111.

REFERENCES

[1] A. Pikovsky, M. Rosenblum, and J. Kurths, *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press, 2003.

[2] S. H. Strogatz, *Sync: The Emerging Science of Spontaneous Order*. Hyperion, 2003.

[3] J. Engel and T. A. Pedley, *Epilepsy: A Comprehensive Textbook*. Philadelphia: Lippincott-Raven, 1975.

[4] Y. Kuramoto, "Self-entrainment of a population of coupled non-linear oscillators," in *International Symposium on Mathematical Problems in Theoretical Physics*. Springer Berlin / Heidelberg, 1975, vol. 39, pp. 420–422.

[5] ——, *Chemical Oscillations, Waves, and Turbulence*. Springer-Verlag, 1984.

[6] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, "Detecting non-trivial computation in complex dynamics," in *Proceedings of the 9th European Conference on Artificial Life (ECAL 2007), Lisbon, Portugal*, ser. Lecture Notes in Artificial Intelligence, F. Almeida e Costa, L. M. Rocha, E. Costa, I. Harvey, and A. Coutinho, Eds., vol. 4648. Berlin / Heidelberg: Springer, 2007, pp. 895–904.

[7] ——, "Local information transfer as a spatiotemporal filter for complex systems," *Physical Review E*, vol. 77, no. 2, p. 026110, 2008.

[8] ——, "Local measures of information storage in complex distributed computation," 2010, in preparation.

[9] ——, "Information modification and particle collisions in distributed computation," *Chaos*, vol. 20, no. 3, pp. 037 109–13, 2010.

[10] J. T. Lizier, "The local information dynamics of distributed computation in complex systems," Ph.D. dissertation, 2010.

[11] A. T. Winfree, *The Geometry of Biological Time*. Springer, 1980.

[12] J. A. Acebron, L. L. Bonilla, C. J. Perez Vicente, F. Ritort, and R. Spigler, "The Kuramoto model: A simple paradigm for synchronization phenomena," *Reviews of Modern Physics*, vol. 77, p. 137, 2005.

[13] O. Simeone, U. Spagnolini, Y. Bar-Ness, and S. Strogatz, "Distributed synchronization in wireless networks," *Signal Processing Magazine, IEEE*, vol. 25, no. 5, pp. 81–97, 2008.

[14] B. N. Clark and C. J. Colbourn, "Unit disk graphs," *Discrete Mathematics*, vol. 86, no. 1-3, pp. 165–177, 1990.

[15] F. M. Atay, J. Jost, and A. Wende, "Delays, connection topology, and synchronization of coupled chaotic maps," *Physical Review Letters*, vol. 92, no. 14, p. 144101, 2004.

[16] M. Barahona and L. M. Pecora, "Synchronization in small-world systems," *Physical Review Letters*, vol. 89, no. 5, p. 054101, 2002.

[17] H. Hong, M. Y. Choi, and B. J. Kim, "Synchronization on small-world networks," *Physical Review E*, vol. 65, no. 2, p. 026139, 2002.

[18] D. J. MacKay, *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press, 2003.

[19] J. P. Crutchfield and D. P. Feldman, "Regularities unseen, randomness observed: Levels of entropy convergence," *Chaos*, vol. 13, no. 1, pp. 25–54, 2003.

[20] J. T. Lizier, M. Prokopenko, and A. Y. Zomaya, "The information dynamics of phase transitions in random Boolean networks," in *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI), Winchester, UK*, S. Bullock, J. Noble, R. Watson, and M. A. Bedau, Eds. Cambridge, MA: MIT Press, 2008, pp. 374–381.

[21] T. Schreiber, "Measuring information transfer," *Physical Review Letters*, vol. 85, no. 2, p. 461, 2000.

[22] J. T. Lizier and M. Prokopenko, "Differentiating information transfer and causal effect," *European Physical Journal B*, vol. 73, no. 4, pp. 605–615, 2010.

[23] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis*. Cambridge, MA: Cambridge University Press, 1997.

[24] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns, "Methods for quantifying the informational structure of sensory and motor data," *Neuroinformatics*, vol. 3, no. 3, pp. 243–262, 2005.

[25] P. F. Verdes, "Assessing causality from multivariate time series," *Physical Review E*, vol. 72, no. 2, pp. 026 222–9, 2005.

[26] M. Chávez, J. Martinerie, and M. Le Van Quyen, "Statistical assessment of nonlinear causality: application to epileptic EEG signals," *Journal of Neuroscience Methods*, vol. 124, no. 2, pp. 113–128, 2003.

[27] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse, "Identification of influential spreaders in complex networks," *Nature Physics*, vol. 6, no. 11, pp. 888–893, 2010.

[28] P. N. McGraw and M. Menzinger, "Clustering and the synchronization of oscillator networks," *Physical Review E*, vol. 72, no. 1, p. 015101, 2005.

[29] M. Brede, "Locals vs. global synchronization in networks of non-identical Kuramoto oscillators," *European Physical Journal B*, vol. 62, no. 1, pp. 87–94, 2008.

[30] J. Gómez-Gardeñes, Y. Moreno, and A. Arenas, "Paths to synchronization on complex networks," *Physical Review Letters*, vol. 98, p. 034101, 2007.

[31] J. T. Lizier, M. Prokopenko, and D. J. Cornforth, "The information dynamics of cascading failures in energy networks," in *Proceedings of the European Conference on Complex Systems (ECCS), Warwick, UK*, 2009, p. 54.