On the Scalability of Hierarchical Cooperation for Dense Sensor Networks

Tamer ElBatt
Information Sciences Laboratory
HRL Laboratories, LLC
Malibu, CA 90265, USA
telbatt@hrl.com

ABSTRACT

In this paper we study the problem of information dissemination in dense multi-hop sensor networks characterized by highly correlated sample measurements. In particular, we investigate the benefits, and trade-offs, of exploiting correlations via cooperatively compressing the data as it hops around the network. First, we study two extreme cooperation strategies, namely no cooperation and network-wide cooperation. We show that network-wide cooperation achieves logarithmic growth rate for the transport traffic with the network size whereas the schedule length growth rate remains linear. Next, we analyze a two-phase cooperation strategy which localizes cooperation within regions of the network in an attempt to assess the performance of strategies bounded by the two aforementioned extremes. Finally, we extend two-phase cooperation to a multi-phase hierarchical cooperation strategy where the number of phases depends on the number of nodes and the size of the cooperation set. The rationale behind this strategy is to achieve logarithmic scaling laws at the expense of more complexity in coordinating nodes' cooperation. In addition, hierarchical cooperation opens room for optimizing the transport traffic and schedule length for a given network size.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication*; H.1.1 [**Models and Principles**]: Systems and Information Theory—*Information Theory*

General Terms

Theory, Performance, Algorithms

Keywords

Sensor networks, data compression, spatial correlations, transport traffic, scheduling latency, scaling laws

IPSN'04, April 26–27, 2004,Berkeley,California,USA.

1. INTRODUCTION

Future wireless networks are presumed to accommodate large numbers of embedded devices that operate cooperatively to achieve a pre-specified sensing task. One of the main hurdles towards the realization of this objective is network scalability. It has been shown in [1] that the *peer-to-peer* (or one-to-one) transport capacity of wireless ad hoc networks scales as $O(\sqrt{N})$ where N is the number of nodes per unit area. This, in turn, implies that the per-node throughput scales as $O(\frac{1}{\sqrt{N}})$, and, hence, asymptotically vanishes as the node density grows to infinity. Therefore, it was concluded that design and implementation should be limited to wireless networks accommodating small numbers of nodes. In [2], the authors studied the problem of broadcast communications[1] in multi-hop sensor networks where samples of a random field are recorded at each node in the network and disseminated to all other nodes in order to obtain an estimate of the entire field within a prescribed distortion value. They observed that the scaling laws derived in [1] are based on the assumption that the traffic generated at different nodes in the network is generally independent. Furthermore, they argued that this conclusion may not be generally true for dense sensor networks due to the fact that spatially close sensors may experience correlations among their sample measurements. Thus, they proposed to use classical source codes at respective sensors and then re-encode the data as it hops around the network in order to remove correlations and, hence, reduce the traffic generated by each sensor. However, this may involve a trade-off between transport traffic and transmission scheduling delays that has been illustrated via an example in [2]. In this paper, we shed some light on this trade-off and explore potential avenues for balancing it. In [3], a distributed algorithm for removing correlations among sensor data via computing wavelet transforms has been proposed. However, the scaling laws of the associated traffic and scheduling delays were left as open problems. On the other hand, [4, 5] fall within the scope of distributed source coding that utilizes Slepian-Wolf coding schemes [6] to remove correlations without any communications among sensors. Finally, the many-to-one capacity of data gathering applications over dense sensor networks was characterized in [7, 8] under different sets of assumptions.

Our main contribution in this paper is to show that hierarchical cooperation opens room for achieving logarithmic scaling laws for the transport traffic and schedule length. First, we quantify the scaling laws of the two performance measures under extreme cooperation strategies (no cooperation and network-wide cooperation) along with the set of strategies bounded by those extremes (two-phase cooperation). Motivated by the linear schedule length growth

[1]"It is also referred to as *flooding* or *many-to-many communications*"

rate that persists under all three strategies, we study the class of hierarchical cooperation schemes. Hierarchical cooperation proceeds through two stages, namely cooperation and distribution. Under the cooperation stage, the scheme goes through $log_i N$ phases (i is the size of the cooperation set) whereby members in the same cooperation set compress their respective data according to a network-wide cooperation scheme in each phase. By the completion of the cooperation stage, only a subset of the nodes would have access to the estimate of the entire random field. Therefore, the distribution stage proceeds, again in a hierarchical manner, through $[(log_i N) - 1]$ phases in order to disseminate the missing information to nodes who do not have access to the entire field estimate. This strategy not only achieves sub-linear scaling laws for the transport traffic and schedule length but also creates room for trading one performance measure for the other via controlling the cooperation set size.

The paper is organized as follows: In section 2, the underlying network model is introduced. Afterwards, the traffic and delay scaling laws under the no cooperation, network-wide cooperation and two-phase cooperation strategies are investigated in section 3. This is followed by a detailed description and analysis of the hierarchical cooperation strategy in section 4. Finally, conclusions are drawn in section 5.

2. NETWORK MODEL

In this paper we limit our attention to one-dimensional sensor networks consisting of N stationary nodes which communicate only via the wireless medium and are uniformly spaced along a horizontal straight line of unit length. We assume that nodes are indexed in an ascending order from left to right. Extending the results to two-dimensional grids lies out of the scope of this paper and is a subject of ongoing research. We assume that all nodes are equipped with omni-directional antennas that radiate energy according to an isotropic pattern. Unless stated otherwise, nodes use fixed and equal transmission power, which translates to the same transmission range (r), where $\frac{1}{(N-1)} \leq r \leq \frac{2}{(N-1)}$, i.e. nodes $(m-1)$ and $(m+1)$ are one-hop neighbors of an arbitrary node m. All nodes are assumed to share the same frequency band, and time is divided into equal size slots and each transmission fits exactly in a single slot.

Each sensor is assumed to record periodic samples of the sensed field, scalar quantize, encode and transmit them such that the sensed field can be reconstructed at all nodes in the network up to a certain level of distortion. We assume that successive samples taken by the same sensor are *temporally uncorrelated* and, hence, we focus on the set of samples recorded by all sensors at a given time instant and drop the time index [2]. On the other hand, sensor measurements are assumed to be *spatially correlated* according to a stationary one-dimensional random process $S(y)$, where $S(y)$ is a real-valued random variable representing the field value at location $0 \leq y \leq 1$. This is motivated by the fact that the vast majority of physical phenomena are analog such that sensors are better modeled as continuous rather than discrete sources. Moreover, we assume that the random process $S(y)$ has the property that the correlation between samples increases as the sensors get dense. We assume that the reading of sensor m, denoted S_m, is quantized by a fixed quantizer q(.) subject to a constraint on the average distortion per sample (i.e. $\frac{1}{N}\sum_{m=1}^{N} E(d(S_m, X_m)) \leq D$) where $q(S_m) = X_m$, $d(.,.)$ is a distortion measure and D is a prescribed constraint on distortion. Thus, the minimum number of bits required to represent the output of the quantizer of node m is given by the entropy $H(X_m)$ of the associated discrete random variable X_m [11]. Finally, we focus on the sensor broadcast problem addressed in [2, 3], where each sensor wishes to disseminate an approximation of its sample to all other nodes in the network. This traffic pattern may arise under a wide variety of application scenarios ranging from *collaborative video surveillance* where individual cameras display images from other cameras in the area of interest to *data gathering applications* where the collector node (sink) is an unmanned aerial vehicle (UAV) that deploys a sensor network onto a geographical field and then intermittently covers that field for the purpose of receiving the aggregated data.

3. COOPERATIVE DATA COMPRESSION IN DENSE SENSOR NETWORKS

In this section, we determine, with the aid of basic information theory, the scaling laws of the transport traffic and scheduling delays associated with the sensor broadcast problem under three cooperation strategies. We employ the notion of transport traffic (measured in bit-meters), introduced in [1], to quantify the traffic volume associated with the sensor broadcast problem under various strategies. A network is said to transport one bit-meters when a single bit has been forwarded a distance of one meter towards its destination. On the other hand, the scheduling delay is measured in the number of slots needed to complete the information dissemination task.

3.1 No Cooperation

Under this strategy, each sensor transmits a quantized version of its sample to its neighbors. A node who receives a sample of a neighbor is supposed to rebroadcast it blindly without re-encoding it. Accordingly, the minimum number of bits generated by node 1 would be given by $H(X_1)$ which should be forwarded over $(N-1)$ hops to reach all other nodes. Similarly, node 2 generates $H(X_2)$ bits which, also, go over $(N-1)$ hops. This process is repeated for all nodes until node N goes through the same procedure. This, in turn, explains the non-cooperative nature of this strategy, where the notion of cooperation in our context implies the role that each sensor node plays in re-encoding the data of other sensors as it hops around the network.

In order to quantify the transport traffic (TT) scaling law, we assume that each node sends only one sample per transmission, i.e. it does not include multiple samples from different sensors (generated and forwarded) in the same transmission[2]. Accordingly, the TT under this strategy would scale linearly with the network size as given by,

$$TT(No\,Coop) = \frac{(N-1)}{(N-1)} \sum_{j=1}^{N} H(X_j)$$

$$= O(N) \;\; bit.meters \qquad (1)$$

On the other hand, the minimum schedule length (SL) is given by $\frac{NT}{NTPS}$, where NT is the total number of transmissions needed to complete the dissemination task and NTPS is the maximum number of non-conflicting transmissions per slot. As illustrated earlier, the sample of each node has to be forwarded over $(N-1)$ hops to reach all other nodes. Therefore, NT would be given by $N(N-1)$, i.e. $O(N^2)$. On the other hand, NTPS depends solely on the interference model. In this paper, we adopt an interference model that is widely employed in the multi-hop packet radio networks literature

[2]"Even if multiple samples are transmitted in a single slot via increasing the link data rate, it can be shown that the scaling laws derived in this section would still hold."

[9, 10], whereby a collision arises whenever multiple transmissions are heard by a receiver in the same slot. Otherwise, a transmission is deemed successful if it is the only one heard by the receiver. Accordingly, broadcast transmissions of nodes who are more than two-hops away are considered non-conflicting and may share the same time slot. This, in turn, guarantees that: i) No simultaneous transmission/reception could arise at a node and ii) No multiple transmissions to the same receiver may co-exist in the same slot. For one-dimensional networks, NTPS would be $\lceil \frac{2N}{3} \rceil$ when each node broadcasts its own sample to its left and right neighbors. On the other hand, when a node forwards a sample of another node received through one of its neighbors, its broadcast would yield one useful sample transfer to the other neighbor. Hence, NTPS would be given by $\lceil \frac{N}{3} \rceil$ under this scenario. Notice that the former scenario occurs only once for each node, whereas the latter one arises several times throughout the sample forwarding process. This, in turn, yields $\lceil \frac{N}{3} \rceil \leq NTPS \leq \lceil \frac{2N}{3} \rceil$ and, hence, SL(No Coop) scales as $O(N)$. Finally, we conclude that both, transport traffic and schedule length of the non-cooperative dissemination strategy, grow *linearly* with the size of the network. In the next section, we explore candidate strategies for achieving sub-linear growth rate for the transport traffic and assess the price paid in terms of the schedule length growth rate.

3.2 Network-wide Cooperation

In this section, we analyze two candidate network-wide cooperation schemes. First, we consider *sequential cooperation* where each node takes a turn in a round-robin fashion to encode its sample given the samples of left nodes received through its left neighbor. Accordingly, node 1 generates $H(X_1)$ (since it has no left neighbors). Next, node 2 sends back $H(X_2|X_1)$ (denoted "2|1" in Figure 1) to 1 and then sends a joint version of the two samples, i.e. $H(X_1, X_2)$ (denoted "$1, 2$" in Figure 1), to node 3 . The scheme proceeds in the same manner with node m sending $H(X_m|X_1, .., X_{(m-1)})$ to node $(m-1)$ followed by $H(X_1, X_2, .., X_m)$ to node $(m+1)$ until $m = N$ as shown in Figure 1 for $N = 4$. Notice that the encoded sample of node m sent to node $(m-1)$ should be propagated in the left direction such that each node gets the samples of all other nodes. Thus, the TT generated by the sequential strategy is given by,

$$TT(Seq\, Coop) = \frac{(N-1)}{(N-1)}\, H(X_1, X_2, ..., X_N) \qquad (2)$$

$$\leq O(N) \qquad (3)$$

Notice that the first equality follows from the chain rule for entropies. On the other hand, equality in (3) holds when the sensor nodes are sufficiently far to render their sample measurements independent. This, in turn, confirms that network-wide cooperation attempts to exploit correlations among sensor samples whenever they exist. Correlations may arise among sensor samples due to the following reasons: i) Deploying dense sensor networks dictated by reliability and network connectivity constraints and/or ii) Measuring the random field with adaptive spatial resolution depending on the sensor activation strategy and the spatial process bandwidth which may be time varying. Next, we need to characterize the scaling law of the joint entropy of N quantized random variables in (2). This problem has been addressed by Marco et al. in [7] for a stationary Gaussian random process and a scalar quantizer with uniform step size and infinite number of levels. Furthermore, it was assumed that the distortion measure is mean square error (MSE) and the autocorrelation function of the spatial process $S(y)$ is exponential and is given by $R_s(y) = e^{-y^2}$. For this setup, they showed that $H(X_1, X_2, ..., X_N)$ scales as $O(log\ N)$ as $N \to \infty$. Thus, we conclude that the TT for the sequential cooperation strategy grows only *logarithmically* with the network size, that is well below the linear scaling law achieved by the no cooperation strategy.

Next, we quantify the price paid for sequential cooperation, namely the SL growth rate. According to Figure 1, it is evident that node 1 participates in 1 transmission and $(N-1)$ receptions, node 2 participates in 1 transmission and $(N-2)$ receptions, and node $(N-1)$ participates in 1 transmission and 1 reception. Hence, NT would be given by $\frac{(N-1)(N+2)}{2}$, i.e. $O(N^2)$.

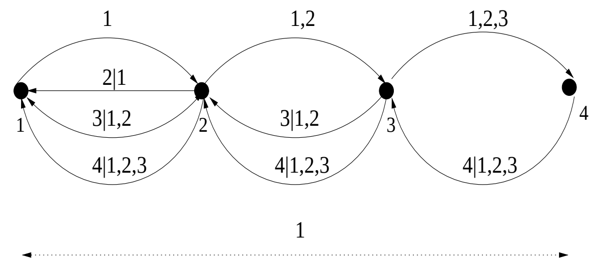

Figure 1: Sequential cooperation over a 4 node sensor network

Notice that the sequential strategy enforces certain order on nodes' communication since node m has to wait for the joint samples of nodes 1,2,... $(m-1)$ in order to generate its own conditional and joint samples. This significantly limits the spatial reuse of slots which causes NTPS to be one and SL(Seq Coop) to scale as $O(N^2)$, well above the linear growth rate achieved by the no cooperation strategy. This considerable degradation may suggest that network-wide cooperation reduces the transport traffic growth rate at the expense of longer schedule lengths. However, we show next that this is not generally the case. The key observation that led to this conclusion is two-fold. First, there is no need to have dedicated transmissions for individual conditional samples to send them back to left nodes as shown in Figure 1. Instead, conditional samples could be sent collectively, rather than individually, in a more efficient way. Second, slot reuse plays an important role in minimizing the schedule length. These observations are illustrated with the aid of the following two examples. In the first example, given the sequential cooperation strategy in Figure 1, eliminate the dedicated intermediate conditional sample transmissions propagated to left nodes. Instead, gather the joint samples of all nodes at node N, through cooperation in the forward direction. Afterwards, start sending collective conditional samples in the reverse direction at higher data rates, compared to the transmission rate of individual conditional samples, in order for each transmission to fit in a single slot. We refer to this strategy as *forward/reverse cooperation* since nodes cooperate in the forward direction first and then in the reverse direction as shown in Figure 2. This strategy entertains a logarithmic growth rate for the transport traffic similar to the sequential cooperation strategy[3]. Furthermore, it consumes $(N-1)$ transmissions in the forward direction and $(N-1)$ transmissions in the reverse direction, i.e. NT scales as $O(N)$. On the other hand, NTPS turns out to be one since, again, this policy does not exploit spatial reuse of slots. Accordingly, SL(Forward/Reverse Coop) scales as $O(N)$, well below the quadratic growth rate associated with the sequential cooperation scheme.

In the second example, we investigate the impact of slot reuse on the scaling laws of the forward/reverse cooperation strategy. No-

[3]"If the broadcast nature of omni-directional antennas is exploited, the transport traffic of forward/reverse cooperation can be further reduced."

tice that the transmissions of nodes $m + 3j$, where j takes integer values, can share the same slot and, hence, cooperation in the forward direction consumes $\sum_{j=1}^{\lceil \frac{N}{3} \rceil} 3j$ slots in this case. This is attributed to the fact that cooperation is limited to reuse clusters (of size 3 nodes)[4] and is repeated over those clusters in order to achieve network-wide cooperation. On the other hand, cooperation in the reverse direction remains unchanged (i.e. $O(N)$) since it involves propagating conditional samples throughout the entire the network. Accordingly, NT would scale as $O(N^2)$ for the forward/reverse cooperation strategy with slot reuse. Moreover, NTPS scales as $O(N)$ based on arguments similar to those used in the previous section. Therefore, we notice that SL still scales as $O(N)$, even when slot reuse is exploited. This result stems from the fact that slot reuse does not only increase NTPS from one to $O(N)$, but it also increases NT from $O(N)$ to $O(N^2)$ in order to achieve network-wide cooperation in the forward direction.

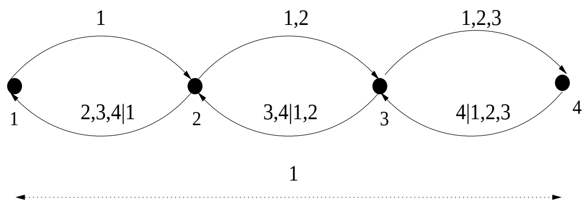

Figure 2: Forward/Reverse cooperation over a 4 node sensor network

Finally, we conclude that there exists network-wide cooperation schemes that reduce the transport traffic growth rate from linear to logarithmic, while preserving the linear growth rate for the schedule length. Motivated by the scaling laws of the extreme cooperation strategies, we need to investigate the trends over the space of strategies bounded by those extremes. This is achieved with the aid of two-phase cooperation discussed in the next section.

3.3 Two-Phase Cooperation

The essence of this strategy is to localize cooperation within regions of the network, where nodes cooperate in compressing each other's data , and beyond those regions no cooperation is performed. This constitutes a simple approach for trading traffic for scheduling delays via introducing the cooperation set size i as a degree of freedom. Accordingly, a one-dimensional network of N nodes is partitioned into $\frac{N}{i}$ cooperation sets, each accommodating i nodes as shown in Figure 3 for $i = 2$. As the name suggests, this strategy proceeds through two phases. In the first phase, members of each set cooperatively compress their sample data according to the forward/reverse cooperation strategy described in the previous section. Once this is done, any node in an arbitrary set would have the sample measurements of all other nodes in its set, however, it would lack the sample measurements of nodes in other sets. Hence, the role of the second phase is to exchange the sample measurements among various cooperation sets. This is achieved via inter-set exchange among representative nodes in respective cooperation sets (e.g. nodes 1,$(i+1)$,$(2i+1)$,...,$(N-i+1)$) in a non-cooperative manner as described in section 3.1^5. This non-cooperative exchange should be followed by data distribution within each set in order to disseminate the sample measurements gathered at the representative nodes to other members in their respective sets.

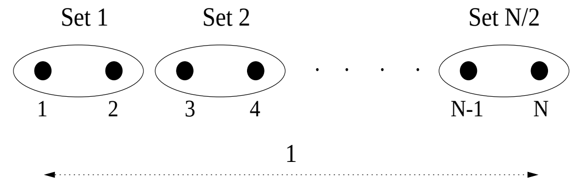

Figure 3: Example of a sensor network with cooperation set size $i = 2$

In the first phase, cooperation takes place within $\frac{N}{i}$ sets simultaneously. This could be achieved assuming that neighboring nodes at the edges of two neighboring sets (e.g. nodes i and $(i+1)$) do not interfere throughout their communication tasks (possibly using directional antennas). Accordingly, the transport traffic generated in phase 1 is given by,

$$TT(Phase\ 1) = [\frac{(i-1)}{(N-1)} * H(X_1, X_2, ..., X_i)] * \frac{N}{i}$$

$$= O(H(X_1, X_2, ..., X_i)) \quad (4)$$

On the other hand, the transport traffic generated in phase 2 would be given by,

$$TT(Phase\ 2) = O(\frac{N}{i} * H(X_1, X_2, ..., X_i))$$

$$+O(\frac{i}{N} * H(X_{i+1}, X_{i+2}, ..., X_N | X_1, X_2, ..., X_i) * \frac{N}{i}) \quad (5)$$

Notice that the first term corresponds to the non-cooperative data exchange among the representative nodes of various cooperation sets, whereas the second term represents the transport traffic associated with distributing the sample measurements to other members within each set. The summation of (4) and (5) yields the transport traffic under the two-phase cooperation strategy,

$$TT(Two\ Phase\ Coop) =$$

$$O(H(X_1, X_2, ..., X_N) + \frac{N}{i} * H(X_1, X_2, ..., X_i)) \quad (6)$$

The scaling law in (6) can be easily verified via observing that the special cases of $i = 1$ and $i = N$ simply reduce to the no cooperation and network-wide cooperation extremes respectively. Moreover, it is evident that the transport traffic monotonically increases as the cooperation set size (i) decreases. This is attributed to the fact that the first term in (6) does not depend on i and constitutes the minimum amount of transport traffic when $i = N$. For $1 < i < N$, the two terms in (6) persist and hence the transport traffic would be greater than the first term. Finally, it is straightforward to show that as i decreases the second term increases.

In the remaining of this section, we determine the behavior of the schedule length under the two-phase cooperation strategy. It is evident that phase 1 consumes $O(i)$ slots to be completed. In addition, phase 2 consumes $O(\frac{N}{i})$ throughout the non-cooperative sample exchange task and $O(i)$ for distributing the gathered sample measurements throughout each set. Therefore, the schedule length for this policy, SL(Two-phase Coop), scales as $O(i + \frac{N}{i})$ slots. Although the linear growth rate with N still persists, the parameter i provides a degree of freedom for optimizing the schedule length for a given network size. The question that remains unanswered

[4] "A reuse cluster is defined as a group of nodes where no slots are reused within this group."

[5] "We assume that the transmission power of representative nodes could be raised in order to reach each other directly over a single-hop."

is: How does SL vary with i? For a given N, the non-linear dependence of SL on i suggests that this function has an extremum which turns out to be a minimum at $i^* = c\sqrt{N}$ where c is a constant. For $i < i^*$, SL decreases as i increases until it reaches the minimum at i^*. For $i > i^*$, SL increases with i until the two-phase cooperation scheme reduces to network-wide cooperation at $i = N$. Accordingly, with proper choice of the parameter i, we can control the relative importance of TT and SL. For instance, the best operating point for SL is around its minimum. However, if the associated TT is excessive, then we should start increasing i in order to reduce the generated traffic. Thus, we conclude that two-phase cooperation attempts to strike a balance between minimizing traffic and minimizing scheduling delays. In the next section, we extend two-phase cooperation to a hierarchical $log_i N$-phase cooperation strategy in an attempt to achieve sub-linear growth rates for the transport traffic and schedule length.

4. HIERARCHICAL COOPERATION

4.1 Analysis

In this section, we explore hierarchical cooperation as a potential avenue for circumventing the linear schedule length growth rate hurdle that persists throughout the aforementioned cooperation schemes. It proceeds through two stages in a hierarchical manner, namely the *cooperation stage* and the *distribution stage*. At one hand, the cooperation stage goes through $k = log_i N$ phases as shown in Figure 4. On the other hand, the distribution stage goes through $(k-1)$ phases as shown in Figure 5. Next, we illustrate the operation of this strategy with the aid of an example. Consider a one-dimensional sensor network consisting of $N = 8$ nodes. In phase P=1 of the cooperation stage, the network is partitioned into $\frac{N}{i}$ cooperation sets, each accommodating i nodes as shown in Figure 4 for $i = 2$. Members of each set cooperatively compress their sample data similar to phase 1 of the two-phase cooperation scheme described earlier. At the beginning of phase $1 < P < k$, any node in an arbitrary set would have the sample measurements of all other nodes in its phase $(P-1)$ set as shown in Figure 4. However, it would lack the sample measurements of nodes in other sets. Thus, a single sensor node is elected from each cooperation set in phase $(P-1)$ as a representative node of that set. Accordingly, at the beginning of phase P, there will be $\frac{N}{i^{(P-1)}}$ representative nodes who need to disseminate the sample measurements gathered during phase $(P-1)$. At this point, the N node sensor network at the beginning of phase 1 reduces to a network of $\frac{N}{i^{(P-1)}}$ nodes at the beginning of phase P. Furthermore, we assume that representative nodes can raise their transmission power appropriately in each phase in order to construct respective cooperation sets. At the beginning of phase $P = 2$ in Figure 4, there are four representative nodes (1,3,5,7) such that node 1 has access to the samples of its phase $P = 1$ set members, namely node 2 and node 3 has access to node's 4 sample and so on. Likewise, at the beginning of phase $P = 3$, there would be two representative nodes, namely 1 and 5, who have access to samples from their phase $P = 2$ set members. Hence, the cooperation stage proceeds recursively through $k = log_i N$ phases until i nodes have access to an estimate of the entire field (nodes 1 and 5 in the example at hand). At this point, the role of the distribution phase prevails in order to distribute the entire field estimate to the nodes whose access is limited so far to a partial view of the entire field (i.e. nodes 2,3,4,6,7,8 in the example). As shown in Figure 5, the distribution stage proceeds recursively in a hierarchical manner from phase 1 up to phase $(k-1)$. Throughout the first distribution phase, node 1 distributes its field estimate to node 3 and node 5 distributes its own estimate to node 7. In a similar fashion, node 1 distributes its estimate to node 2, node 3 to node 4, node 5 to node 6, and node 7 to node 8 throughout the second distribution phase. Thus, by the completion of the distribution stage, all N nodes would have access to the entire field estimate which is the prime objective of the information dissemination application under investigation in this paper.

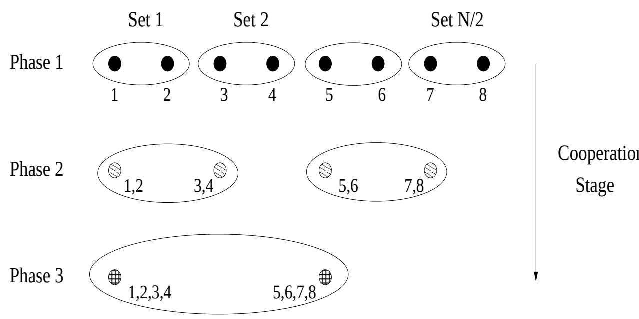

Figure 4: Example of the hierarchical cooperation stage for $N = 8$ and $i = 2$

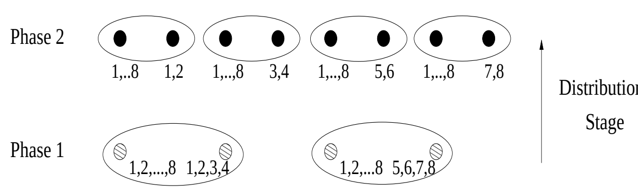

Figure 5: Example of the hierarchical distribution stage for $N = 8$ and $i = 2$

Next, we characterize the scaling law of the transport traffic induced by the hierarchical cooperation scheme. First, we focus on the cooperation stage where cooperation in phase P takes place within each of the $\frac{N}{i^P}$ cooperation sets simultaneously. The transport traffic generated throughout the cooperation stage is given by,

$$TT(Coop\ Stage) = \frac{(i-1)}{(N-1)} *$$

$$[\sum_{j=1,(i+1),\ldots}^{(N-i+1)} H(X_j, X_{j+1}, ..., X_{j+i-1})$$

$$+i \sum_{j=1,(i^2+1),\ldots}^{(N-i^2+1)} H(X_j, X_{j+1}, ..., X_{j+i^2-1})$$

$$\vdots$$

$$+i^{k-2} \sum_{j=1,(i^{(k-1)}+1),\ldots}^{(N-i^{(k-1)}+1)} H(X_j, X_{j+1}, ..., X_{j+i^{(k-1)}-1})$$

$$+i^{k-1} \sum_{j=1,(i^k+1),\ldots}^{(N-i^k+1)} H(X_j, X_{j+1}, ..., X_{j+i^k-1})] \quad (7)$$

Notice that the first term in (7) represents the TT generated throughout phase 1 in the cooperation stage, the P-th term represents the

TT generated throughout phase P, and the last term (i.e. k-th term) represents the TT associated with the k-th cooperation phase. In addition, an arbitrary term P consists of the summation of $\frac{N}{i^P}$ terms which corresponds to the number of cooperation sets in that phase. Accordingly, the first summation consists of $\frac{N}{i}$ terms whereas the last summation collapses to a single term since $N = i^k$. Next, we quantify the transport traffic generated throughout the distribution stage which consists of $(k-1)$ phases,

$$TT(Dist\ Stage) = \frac{(i-1)}{(N-1)} * [i^{k-2} *$$

$$\sum_{j=1,i^{k-1}+1,..}^{N-i^{k-1}+1} H(X_1,..,X_{j-1},X_{j+i^{k-1}},..,X_N|X_j,..,X_{j+i^{k-1}-1})$$

$$\vdots$$

$$+i \sum_{j=1,(i^2+1),...}^{(N-i^2+1)} H(X_1,..,X_{j-1},X_{j+i^2},..,X_N|X_j,..,X_{j+i^2-1})$$

$$+ \sum_{j=1,(i+1),...}^{(N-i+1)} H(X_1,...,X_{j-1},X_{j+i},...,X_N|X_j,...,X_{j+i-1})] \quad (8)$$

In (8), the l-th term represents phase l in the distribution stage. From Figures 4 and 5, it is evident that phase 1 in the cooperation stage matches phase $(k-1)$ in the distribution stage in the sense that they conduct cooperation and distribution respectively among the same set of nodes and they operate over the same number of sets. The same argument applies to phase 2 in the cooperation stage and phase $(k-2)$ in the distribution stage, and more generally, to phases P and $(k-P)$ in respective stages. Thus, summing up the matching terms in (7) and (8) (except for the last term in (7) which has no matching term) yields the TT associated with the hierarchical cooperation strategy as follows,

$$TT(Hier\ Coop) = \frac{(i-1)}{(N-1)} [H(X_1,X_2,...,X_N) * \frac{N}{i}]$$

$$+\frac{i(i-1)}{(N-1)} [H(X_1,X_2,...,X_N) * \frac{N}{i^2}]$$

$$\vdots$$

$$+\frac{i^{k-2}(i-1)}{(N-1)} [H(X_1,X_2,...,X_N) * \frac{N}{i^{k-1}}]$$

$$+\frac{i^{k-1}(i-1)}{(N-1)} [H(X_1,X_2,...,X_N)] \quad (9)$$

The above equation can be reduced to,

$$TT(Hier\ Coop) = \frac{(i-1)}{(N-1)} [k \frac{N}{i} H(X_1,X_2,...,X_N)]$$

$$= O(log_i N\ H(X_1,X_2,...,X_N)) \quad (10)$$

Based on the logarithmic growth rate of $H(X_1,X_2,...,X_N)$, we conclude that the TT under hierarchical cooperation scales faster than network-wide cooperation schemes, yet, it is much slower than the scaling laws associated with no cooperation and two-phase cooperation. This performance degradation, compared to network-wide cooperation, should be weighed against the significant reduction in the SL growth rate achieved by hierarchical cooperation. Therefore, the rest of this section is dedicated to quantifying the behavior of the schedule length under the hierarchical cooperation strategy. It is evident that each phase in the cooperation stage consumes $O(i)$ slots to be completed. Accordingly, the cooperation stage consumes $O(i\ log_i N)$ slots. On the other hand, each phase in the distribution stage consumes $O(i)$ slots to be completed and, hence, the distribution stage consumes $O(i\ log_i N)$ slots. Thus, the schedule length for this strategy, SL(Hier Coop), scales as $O(i\ log_i N)$. This result reveals two key observations: i) Hierarchical cooperation yields a logarithmic growth rate for the schedule length with N, that is well below the linear growth rate achieved by network-wide cooperation or two-phase cooperation and ii) The cooperation set size i constitutes a degree of freedom for trading transport traffic for scheduling delay for a given network size. Clearly, SL increases with the parameter i. On the other hand, (10) suggests that the TT decreases with i. Hence, for a given network size, there is an optimum set size that strikes a balance between these two conflicting objectives.

4.2 Performance Comparison

In this section, we compare the scaling laws associated with various cooperation strategies addressed in this paper and summarized in Table 1. This is of paramount importance to judge their relative performance, their advantages and limitations, and finally potential avenues for extending this work. First, we notice that the no cooperation extreme experiences the worst (linear) scaling laws among all strategies. This is attributed to the fact that this strategy does not exploit the correlations among samples taken at spatially close sensors. At the other extreme, we analyzed the scaling laws of network-wide cooperation. In particular, we studied two examples of network-wide cooperation, namely sequential and forward/reverse cooperation. The former was found to achieve a logarithmic scaling law of the TT at the expense of a quadratic SL growth rate. On the other hand, the latter achieves a logarithmic scaling law of the transport traffic while preserving the linear scaling law for the schedule length. The improvement of the SL growth rate, compared to sequential cooperation, is attributed to sending conditional samples collectively, rather than individually. Thus, we conclude that there exists network-wide cooperation schemes that achieve sub-linear TT scaling law without any degradation in the linear SL scaling law. This, in turn, motivated us to examine the entire space of strategies bounded by the two extremes, via varying the cooperation set size i, using the two-phase cooperation framework. Although the linear SL scaling law still persists under two-phase cooperation, Table 1 shows that varying the parameter i from 1 to N varies the scaling laws between the two extreme strategies. This, in turn, suggests that the parameter i creates room for optimizing the TT and SL for a given network size. Finally, the hierarchical $log_i N$-phase cooperation strategy, which may be viewed as an extension of two-phase cooperation, exhibits the best (poly logarithmic TT and logarithmic SL) scaling laws among all strategies. It has been argued in [2] that the transport traffic associated with a network flow constructed over a two-dimensional grid scales as $O(log\ N)$. However, this conclusion does not account for: i) The traffic volume generated to gather sample measurements of each subnet at nodes on the boundary of a cut and ii) The traffic needed for distributing the network-wide data within each subnet. In this paper, we showed that for one-dimensional networks the transport traffic scales in a logarithmic fashion under network-wide cooperation, poly logarithmically under hierarchical cooperation and our recent results confirm that both strategies scale faster for two-dimensional grids. On the other hand, our recent results reveal that the schedule length follows the same growth rate over one-dimensional and two-dimensional networks. Therefore, the hierarchical cooperation strategy introduced in this paper turns out to be more efficient than the flow constructed in [2] which generates schedules of length $O(\sqrt{N})$ slots. Finally, the parameter i could serve as a vehicle for balancing the trade-off between TT and SL for a given network size. Formulating and solving an optimization problem with these two objectives lies out of the scope of this paper and is a subject of future research.

Table 1: Transport Traffic and Schedule Length Scaling Laws

Strategy	TT	SL
No Cooperation	$O(N)$	$O(N)$
Sequential Cooperation	$O(log\ N)$	$O(N^2)$
Forward/Reverse Coop.	$O(log\ N)$	$O(N)$
Two-phase Cooperation	$O(logN+\frac{N}{i}logi)$	$O(i+\frac{N}{i})$
Hier. Cooperation	$O(log_i N.logN)$	$O(i\ log_i N)$

5. CONCLUSIONS

In this paper we studied the trade-off between transport traffic and scheduling delays associated with the information dissemination problem in dense sensor networks where sample measurements are highly correlated. First, we characterized the transport traffic and schedule length scaling laws under no cooperation, network wide cooperation and two-phase cooperation strategies. Network-wide cooperation was found to reduce the transport traffic scaling law from $O(N)$ to $O(log\ N)$. However, the linear schedule length growth rate achieved by all three schemes constitutes a strong motivation for developing more efficient cooperation strategies. Thus, we proposed to extend two-phase cooperation to multiphase hierarchical cooperation that limits cooperation to gradually increasing regions of the network in successive phases. We showed, analytically, that hierarchical cooperation achieves logarithmic scaling laws for the transport traffic and schedule length. Furthermore, it introduces the cooperation set size as a degree of freedom to trade-off between the two conflicting objectives for a given network size. Extending the results of this paper to two-dimensional grids is a potential avenue for future work. Another direction is to employ more realistic interference models where the signal-to-interference-and-noise-ratio (SINR) serves as the criteria for successful reception at any node.

6. REFERENCES

[1] P. Gupta and P.R. Kumar ”The Capacity of Wireless Networks,” *IEEE Transactions on Information Theory*, vol. 46, No. 2, pp. 388-404, March 2000.

[2] A. Scaglione, and S. Servetto ”On the Interdependence of Routing and Data Compression in Multi-hop Sensor Networks,” *Proc. ACM MOBICOM*, Sept. 2002.

[3] S. Servetto ”Distributed Signal Processing Algorithms for the Sensor Broadcast Problem,” *Proc. Conference on Information Sciences and Systems*, March 2003.

[4] D. Neuhoff and D. Marco ”Distributed Encoding of Sensor Data,” *Proc. IEEE Information Theory Workshop*, Oct. 2002.

[5] J. Chou, D. Petrovic and K. Ramchandran ”A Distributed and Adaptive Signal Processing Approach to Reducing Energy Consumption in Sensor Networks,” *Proc. IEEE INFOCOM*, March 2003.

[6] D. Slepian and J.K. Wolf ”Noiseless encoding of correlated information sources,” *IEEE Transactions on Information Theory*, vol. 19, pp. 471-480, July 1973.

[7] D. Marco, E. J. Duarte-Melo, M. Liu and D. Neuhoff ”On the many-to-one transport capacity of a dense wireless sensor network and the compressibility of its data,” *Proc. International Workshop on Information Processing in Sensor Networks*, 2003.

[8] H. El Gamal ”On the Scaling Laws of Dense Wireless Sensor Networks,” *Submitted to IEEE Transactions on Information Theory*, 2003.

[9] I. Cidon and M. Sidi ”Distributed Assignment Algorithms for Multihop Packet Radio Networks,” *IEEE Transactions on Computers*, vol. 38, No. 10, pp. 1353-1361, Oct. 1989.

[10] A. Ephremides and T. Truong ”Scheduling Broadcasts in Multihop Radio Networks,” *IEEE Transactions on Communications*, vol. 38, No. 4, pp. 456-460, April 1990.

[11] T. Cover and J. Thomas, *Elements of Information Theory.* John Wiley & Sons, Inc., New York, 1991.