

Optimizing Timely Coverage in Communication Constrained Collaborative Sensing Systems

Jean Abou Rahal*, Gustavo de Veciana*, Takayuki Shimizu**, and Hongsheng Lu**

*Dept. of Electrical and Computer Engineering, The University of Texas at Austin, TX

**Toyota Motor North America, Inc., InfoTech Labs, USA, 465 Bernardo Avenue, Mountain View, CA 94043, USA,

*{jeanabourahal,deveciana}@utexas.edu

**{takayuki.shimizu,hongsheng.lu}@toyota.com

Abstract—We consider a collection of distributed sensor nodes periodically exchanging information to achieve real-time situational awareness in a communication constrained setting, e.g., collaborative sensing amongst vehicles to enable safety-critical decisions. Nodes may be both consumers and producers of sensed information. Consumers express interest in information about particular locations, e.g., obstructed regions and/or road intersections, whilst producers provide updates on what they are currently able to see. Accordingly, we introduce and explore optimizing trade-offs between the coverage and the space-time average of the “age” of the information available to consumers. We consider two settings that capture the fundamental character of the problem. The first addresses selecting a subset of producers which optimizes a weighted sum of the coverage and the average age given that producers provide updates at a fixed rate. The second addresses the minimization of the weighted average age achieved by a fixed subset of producers with possibly overlapping coverage by optimizing their update rates. The former is shown to be submodular and thus amenable to greedy optimization while the latter has a non-convex/non-concave cost function which is amenable to effective optimization using tools such as the Frank-Wolfe algorithm. Numerical results exhibit the benefits of context dependent optimization information exchanges among obstructed sensing nodes in a communication constrained environment.

I. INTRODUCTION

In the near future it is envisaged that there will be several disruptions and challenges to the automotive and wireless industries. Amongst these, an intriguing and challenging one will be the emergence of automated cars (also UAVs, robots, etc.) with the ability to collaboratively navigate through complex environments. In order to enable such functionality, it is expected that nodes will collaborate by sharing sensed information, e.g., cars share their views of an obstructed environment. The aim is to achieve a high degree of “real-time situational-awareness”, i.e., to detect/recognize and then effectively track dynamic objects in their vicinity so as to enable safety critical decisions. To that end, it is expected that vehicles will not only rely on on-board multimodal sensing, but also share (raw or fused) sensing information with each other, with the clear goal of facilitating coordination. This may require the transport of substantial volumes of data within and among cars, as well as to the network edge and/or cloud. The optimization of information sharing in a communication constrained system will thus be a fundamental problem underlying such systems. The focus of this

paper is on the modeling and analysis of this problem and its implications in practical collaborative sensing systems.

A key step in this direction is to identify appropriate/usable metrics to assess how well an information sharing policy is performing. This involves at least two concerns. On the one hand, one is interested in *coverage*, i.e., the fraction of the relevant region that a set of collaborating nodes (producers) will be able to track. On the other hand, for dynamic environments, one is interested in the timeliness of the available information across space, e.g., the Age of Information (AoI), when sensors periodically share what they see. Intuitively high coverage is achieved by ensuring that all sensors disseminate their (possibly redundant) information to all the other relevant nodes while the minimization of age may involve giving some well positioned sensors/nodes a higher update rate than others or leveraging overlaps among sensors’ coverage sets. Additionally it is of interest to ensure consumers have a higher awareness of the on-goings in close proximity, e.g., it is more critical for a car to have fresh information of on-goings in its neighborhood rather than receiving frequent updates about distant locations. Roughly speaking optimizing the “timely coverage” for a collaboratively sensing system requires modeling the relative value of each sensor’s updates, e.g., in terms of all of overall coverage, importance, and timeliness.

We focus on four major intertwined questions:

- 1) *Can one provide a model and metrics to evaluate the timeliness/coverage trade-offs achieved by a given information sharing policy for a set of collaborative sensors?*
- 2) *Assuming that a fixed number of producers are sharing information at a fixed rate regarding their possibly overlapping coverages, can one determine the best subset to participate in information exchanges?*
- 3) *Assuming a subset of sensors is chosen, can one jointly optimize their update rates to minimize the weighted average (space-time) age over their coverage set?*
- 4) *How do optimized information sharing policies compare to simple policies as a function of the sensor/node density, i.e., inherent overlap, and system communication capacity?*

Related Work. A key motivating application for this work is collaborative sensing in support of automated vehicles. The basic idea is to facilitate real-time exchanges of sensor information among vehicles and/or road side units to enhance their “situational awareness” in obstructed and dynamic environments, see e.g., [1]–[3]. The recent work in [4], [5] is unique in that

This work was partially supported by Toyota Motor North America, Inc., InfoTech Labs and NSF Grant ECC-1809327.

it uses stochastic geometry to model and analyze collaborative sensing coverage in obstructed environments as a function of the penetration of vehicles with sensing capabilities.

When addressing real-time situational awareness, it is key that the decision-making nodes have access to timely information. The modeling and delivery of timely information has recently received substantial attention, see e.g., [6], [7]. The newly proposed metric, age of information (AoI) became popular since it better represents the information freshness compared to the traditional delay metric. AoI has been extensively studied in the literature, see e.g., [8]–[10]. The work in [11] is perhaps the closest to this paper in that it addresses the issue of optimizing the overall AoI by carefully choosing sensors' update rates and allocating network resources. However, by contrast with these works, in this paper, we model and explore the impact that updates from *multiple independent* sensors will have on the AoI, as well as trade-offs between coverage and timeliness.

Many instances of coverage and sensor selection problems, e.g., [12] are known to have submodular characteristics which in turn are amenable to greedy approximations, see e.g., [13]. To our knowledge, this paper is distinct from previous work in that it introduces and addresses a new fundamental trade-off between coverage and age of the available information for collaborative sensing systems.

Contributions of this paper. Given a set of sensors generating periodic updates (at possibly different rates) regarding their coverage sets we define and characterize the weighted (space-time) average age for the information exchanged. To our knowledge this is the first work addressing the “timely coverage” for a set of collaborative sensors.

We then explore the resource allocation and performance trade-offs in such systems. In particular we formulate two possible frameworks. The first captures a trade-off between minimizing the weighted average age and maximizing coverage of the spatial information requested by the consumers when sensors have a fixed update rate. In this setting we study how to select the best subset of sensors to achieve good coverage but at the same time reduce the weighted age through redundancy in sensor's independent updates. We show that this age-coverage trade-off optimization problem has a submodular structure which allows efficient greedy optimization algorithms. In the second setting we fix the subset of sensors, e.g., all that are available or those selected in the first setting, which now act as producers of information, and explore the benefits of jointly optimizing their update rates towards minimizing the weighted average age. When producers have non-overlapping coverage sets, we show that their optimal update rates are proportional to the square root of their coverage's weight. However, more generally, the weighted age minimization problem has a non-convex/non-concave structure, but explore the use of the Frank-Wolfe gradient method to show the potential benefits of joint update rate optimization.

A numerical evaluation of the benefits of these various approaches from the point of view of age and coverage is conducted, exhibiting the possible advantages that resource allocation in a collaborative sensing setting should play, particularly in congested environments with limited communication resources.

II. SYSTEM MODEL

We shall begin by formally describing our model for a collaborative sensing system along with the associated notation.

A. Sensor coverage sets, consumers, producers and weighting measure

Without loss of generality we consider a set of sensors V in a given overall region $R \subset \mathbb{R}^2$. Sensors are indexed by their locations $v \in V$ and the coverage of sensor v in a given environment is denoted by a subset $C_v \subseteq R$. Given a subset of sensors $X \subseteq V$, we denote X 's overall coverage by $C(X) := \bigcup_{v \in X} C_v$.

The coverage sets for a subset of sensors X , i.e., $(C_v, v \in X)$ induce a partition of their overall coverage set $C(X)$ which we denote by $\mathcal{P}^X = \{P_i^X, i = 1, \dots, |\mathcal{P}^X|\}$. Each subset of the partition P_i^X is such that each location $x \in P_i^X$ can be seen by the same unique subset of sensors $V_i^X \subseteq V$, i.e., such that $x \in P_i^X$ if and only if $x \in C_v$ for all $v \in V_i^X$. It should be clear that if $i \neq j$ then $P_i^X \cap P_j^X = \emptyset$. Further it should be clear that $\bigcup_{i=1}^{|\mathcal{P}^X|} P_i^X = C(X)$, thus we have a partition of $C(X)$. In fact, assuming it is nonempty, if we further include an additional set $R \setminus C(X)$ corresponding to the locations which are not covered by X , we get a partition of the overall region R . It is also possible that the coverage sets of two or more sensors intersect on a set of measure zero. For simplicity, and to avoid unnecessary burdens, we assume that all elements of the partition have non-zero measure.




Fig. 1: Three sensors observing their environments.

Fig.1 illustrates three sensors $X = \{v_1, v_2, v_3\}$ which for simplicity have each a disc coverage set. Sensors v_1 and v_2 have overlapping coverage regions. The figure also exhibits the four subsets in the induced partition, $\mathcal{P}^X = \{P_1^X, P_2^X, P_3^X, P_4^X\}$. The set of sensors associated with subset P_4^X is $V_4^X = \{v_1, v_2\}$. **Definition 1. (Weight measure)** We let $w(\cdot)$ denote a measure on the region R which for any measurable* set $C \subseteq R$, assigns a weight $w(C)$.

For example, if $w(\cdot)$ corresponds to the area of a set, then $w(C(X))$ denotes the area covered by the sensors in X , and if normalized, $w(C(X))/w(R)$ represents the coverage they provide of region R . The weight measure provides a flexible means to encode the importance of certain regions, e.g., in a vehicular setting, road intersections may be more important than other locations. It also provides a means to capture the relative importance of a region from the perspective of information sharing. Note that in general the weight measure

*For simplicity we shall suppress unnecessarily technical details, but in principle we need $w(\cdot)$ to be measurable with respect to the sets in the σ -field generated by the sensors coverage sets.

could be continuous or discrete. In the latter case we envisage that discrete locations correspond to *anchor points* which based on known geometry of the environment are known to have higher importance, e.g., intersections for incoming vehicles or locations obstructed by other vehicles.

Without loss of generality we suppose each sensor node is simultaneously a consumer and producer which may broadcast timely updates regarding what it is able to see. A consumer indicates its interest for information regarding a particular location or region through its spatial interest measure. In turn the sum of the consumers' spatial interest measures, captures the aggregate interest of consumers.

Definition 2. (Consumer's spatial interest) A consumer $v \in V$ indicates interest in timely information about the environment via its spatial interest measure $w_v()$ on R .

Definition 3. (Aggregated spatial interest) The aggregated spatial interest weight $w()$ is given by the sum of the consumers' spatial interest measures, i.e., $w() = \sum_{v \in V} w_v()$ on R .

In some practical settings, a consumer's v interest may be limited to a smaller region, say R_v . For example, a self-driving car with a response time of t_{interest} moving at a speed s would primarily care about what is happening in a region $s \cdot t_{\text{interest}}$ around it. Thus R_v might be modelled as a rectangle of length $2 \cdot s \cdot t_{\text{interest}}$ and width typically covering the road width. Fig.2 illustrates the coverage of a sensor v (green region), obstructed by neighbouring vehicles (red region behind the vehicles), as well as its rectangular region of interest R_v . In this case, v 's spatial interest measure would primarily be supported by the red regions. Assuming that sensor v 's location on the road is (x_v, y_v) , where x_v and y_v stand for the x-y coordinates of v in 2-D, and the origin 0 is at the center of R , then R_v is defined as, $R_v := R \cap \left(\left[-\frac{w_{\text{road}}}{2}, \frac{w_{\text{road}}}{2} \right] \times \left[x_v - s \cdot t_{\text{interest}}, x_v + s \cdot t_{\text{interest}} \right] \right)$, where w_{road} denotes the width of the road.




Fig. 2: Region of interest and coverage of a consumer.

B. Network capacity, sensor updates, and age of information

We shall assume that each producer sensor v generates periodic updates regarding its coverage set C_v at a rate r_v updates per second, i.e., the update interval is $\frac{1}{r_v}$. The updates are either broadcast to the other sensors or shared with a central controller. For any subset $X \subseteq V$ we let $\mathbf{r}(X) = (r_v : v \in X)$ denote the vector of update rates for the sensors in X . The delay for sensor v to access the communication medium and transmit its update is assumed to be exactly, or at most, d_v . Thus if only one sensor can access the communication medium at a time the fraction of time sensor v holds the medium is d_v over the update interval $1/r_v$, i.e., $d_v r_v$. The selected update rates for a set of sensors X must then satisfy a ‘‘capacity’’ constraint,

$$\sum_{v \in X} d_v r_v \leq 1, \quad (1)$$

ensuring the medium is not overlooked. Note in practice, depending on the character of the medium, one would require a back off $\sum_{v \in X} d_v r_v \leq 1 - \epsilon$ for some $\epsilon > 0$ to ensure the deadlines d_v are met. The back off will depend on the details of channel access and/or scheduler. For simplicity we will neglect the ϵ term in the sequel.

A natural metric that captures the freshness of the received updates is the age of information (AoI) available at the receiver. Fig.3 exhibits the time-varying AoI at the receiver for such a periodic update process with rate r_v for which the update transmission delay is no more than d_v . In the sequel, we will consider both *average* age and the probability that the age exceeds a pre-specified threshold at a typical time from sensor v . For example the average age is given by,

$$\text{average age of sensor } v = d_v + \frac{1}{2r_v}. \quad (2)$$

To keep things simple we will assume the channel access and update transmission delay, or possibly an upper bound on this quantity, is the same for all sensors, i.e., $d_v = d$ for all $v \in V$.




Fig. 3: Time-varying age of information.

However since producers may be covering overlapping region sets, consumers may be receiving, at different times, the same update from multiple sensors, which results in a reduction in the age of the available information at the consumer end. We define a generic age function for such regions as follows.

Definition 4. (Age of regions with overlapping sensor updates) Suppose the producers X are transmitting updates at rates $\mathbf{r}(X) = (r_v, v \in X)$. Recall that X induces a partition where the locations in P_i^X are covered by a set of sensors V_i^X . The age of P_i^X depends on the update rates of these sensors, i.e., $\mathbf{r}(V_i^X)$. With a slight abuse of notation we will define an **age function** which captures a proxy for the age (e.g., average or probability of exceeding a threshold) of the set P_i^X as

$$\text{age}(P_i^X) = \text{age}(\mathbf{r}(V_i^X)), \quad (3)$$

with the intention of emphasizing its dependence on the associated sensors' update rates.

In the next section we shall explore the characteristics of the age as a function of the number of sensors and their associated update rates.

Definition 5. (Weighted age for a set of sensors) Given a weight measure $w()$ on R , a set of sensors X , and sensor update rates $\mathbf{r}(X)$, the **weighted age** of the coverage set $C(X)$ associated with X is given by,

$$a(X, \mathbf{r}(X)) := \sum_{i=1}^{|\mathcal{P}^X|} w(P_i^X) \text{age}(P_i^X) = \sum_{i=1}^{|\mathcal{P}^X|} w(P_i^X) \text{age}(\mathbf{r}(V_i^X)).$$

If $w()$ corresponds to the area, we say $a(X, \mathbf{r}(X))$ is the **area weighted age** of $C(X)$. If the weighted age is divided by $w(C(X))$ it will be referred to as the **normalized weighted age** or the **spatial age**.

III. CHARACTERIZING THE AGE FUNCTION

In this section we define and characterize the properties of two possible age functions, as introduced in Definition 4.

A. Definition and computation of the age function

We first consider two simple examples which motivate the more general formulation. Recall that the age function depends on a vector of update rates of sensors which see a given region. As discussed in the previous section, if updates from only a single sensor at rate r_1 are available, then the *average* age, denoted age_1 depends on the scalar r_1 and is given by,

$$\text{age}_1(r_1) = d + \frac{1}{2r_1}.$$

Let A_1 be a random variable denoting the age of the saw-tooth function when viewed at a random time (see Fig. 3). Given the saw-tooth linear age growth, it should be clear that $A_1 \sim d + \frac{1}{r_1}U_1$, where $U_1 \sim \text{Unif}[0, 1]$. We now define two performance age functions, the average age, age_1 , and the γ -age violation, age_2 , given by,

$$\text{age}_1(r_1) = \mathbb{E}[A_1] \quad \text{and} \quad \text{age}_2(r_1) = \mathbb{P}(A_1 > \gamma),$$

where $\gamma \geq 0$ is a target age one would not wish to exceed.

Suppose there are in fact updates from two sensors covering a given set in the partition, e.g., as shown in Fig.1. Sensors v_1 and v_2 are providing updates of region P_4^X with transmission delay at most d and update rates $\mathbf{r} = (r_1, r_2)$. Without loss of generality, assume $r_1 \geq r_2$. As shown in Fig.4, the dashed and dotted saw-tooth curves correspond to the updates of the two sensors. Assuming that the phases of the saw-tooth curves are randomly distributed, it is easy to see that the average age at a typical point in time is given by the minimum of the two curves,

$$\text{age}_1(\mathbf{r}) = \mathbb{E}[\min[A_1, A_2]] = d + \frac{1}{r_1} \left(\frac{1}{2} - \frac{1}{6} \frac{r_2}{r_1} \right),$$

where $A_1 \sim d + \frac{1}{r_1}U_1$ and $A_2 \sim d + \frac{1}{r_2}U_2$, and where U_1, U_2 are uniformly distributed and assumed to be i.i.d., and A_1, A_2 are the ages of the updates from Sensors 1 and 2 observed at a random time. The reduction in age due to redundancy in the sensors' updates is clear. The probability of γ -age violation shares similar properties as the average age.




Fig. 4: Age of partition P_4^X is the minimum of both age functions $\text{age}(r_1)$ and $\text{age}(r_2)$.

Definition 6. (Age functions) Consider a region observed by n sensors generating periodic updates at rates $\mathbf{r} = (r_1, \dots, r_n)$ and with associated transmission delays d such that $d \leq \frac{1}{r_i}$ for all $i = 1, \dots, n$. Assuming the phases of the sensors' updates are independent and uniformly distributed[†] then the **average age** and **γ -age violation** functions of locations in the region are

$$\text{age}_1(\mathbf{r}) = \mathbb{E}[A] \quad \text{and} \quad \text{age}_2(\mathbf{r}) = \mathbb{P}(A \geq \gamma),$$

where $A = \min_{i=1, \dots, n} A_i$ and $A_i \sim d + \frac{1}{r_i}U_i$ and $U_i \sim \text{Unif}[0, 1]$ are independent of each other.

Theorem 1: (Characterization of the age functions) Consider a region observed by n sensors which generate periodic updates at rates $\mathbf{r} = (r_1, \dots, r_n)$ such that without loss of generality, $r_1 \geq r_2 \geq \dots \geq r_n$, then the average age function is

$$\begin{aligned} \text{age}_1(\mathbf{r}) = \mathbb{E}[A] &= d + \int_d^{d + \frac{1}{r_1}} \prod_{j=1}^n z_j(x) dx \\ &= d + \frac{1}{r_1} \left[\sum_{k=1}^n (-1)^{k-1} \frac{c(k, \mathbf{r})}{k(k+1)} \frac{1}{r_1^k} \right], \end{aligned} \quad (4)$$

where $z_i(x) = 1 - r_i(x - d)$, for $i = 1, \dots, n$, and where $c(k, \mathbf{r}) = r_1 \left[\sum_{i_2, \dots, i_k} r_{i_2} \dots r_{i_k} \right]$, for $k = 2, \dots, n$, and $i_2 < i_3 < \dots < i_k$, and $c(1, \mathbf{r}) = r_1$. The γ -age violation function is

$$\text{age}_2(\mathbf{r}) = \mathbb{P}(A \geq \gamma) = \begin{cases} 1, & \text{if } 0 \leq \gamma \leq d, \\ \prod_{i=1}^n z_i(\gamma), & \text{if } d < \gamma \leq d + \frac{1}{r_1}, \\ 0, & \text{if } d + \frac{1}{r_1} < \gamma. \end{cases} \quad (6)$$

The proof of Theorem 1 is relegated to the Appendix. As we will see, these age functions are somewhat complex, thus we will take some time to characterize their properties.

B. Properties of the age functions

The following further corollary characterizes two basic characteristics of the age functions.

Corollary 1: (Properties of the age functions). Suppose that $\mathbf{r} = (r_i : i = 1, \dots, n)$, where $r_i = r$, then the average age function is given by,

$$\text{age}_1(\mathbf{r}) = d + \frac{1}{n+1} \frac{1}{r}, \quad (7)$$

and the γ -age violation function is given by,

$$\text{age}_2(\mathbf{r}) = \begin{cases} 1, & \text{if } 0 \leq \gamma \leq d, \\ (1 - r(\gamma - d))^n, & \text{if } d < \gamma \leq d + \frac{1}{r}, \\ 0, & \text{if } d + \frac{1}{r} < \gamma. \end{cases} \quad (8)$$

Suppose now that \mathbf{r} is such that $r_1 \geq r_2 \geq \dots \geq r_n$ and let $\bar{\mathbf{r}} = (\bar{r}_i : i = 1, \dots, n)$, where $\bar{r}_i = \bar{r} = \frac{1}{n} \sum r_i$. Then for $j = 1, 2$, we have that

$$\text{age}_j(\bar{\mathbf{r}}) \geq \text{age}_j(\mathbf{r}) \geq \text{age}_j(n\bar{r}). \quad (9)$$

The proof of this corollary is left out of the paper for space constraints. This corollary characterizes the decrease in the age's behaviour as the number n of sensors with a fixed update rate r grows. Indeed both the average/violation age functions decrease (as convex functions) to lower bounds d and 0 respectively. So no matter how many sensors send updates about a location

[†]Coordination and optimization of phase offsets was considered impractical.

one can not reduce the average age below the transmission delay d . The corollary also suggests that if sensors with equal update rates view the same location, it is preferable, in order to minimize the age, to replace them by a single sensor and shift the resources to it.

IV. SENSOR SELECTION: AGE-COVERAGE TRADEOFFS

In this section, we consider a setting where sensors (producers) generate updates at the *same fixed* rate r to a central observer e.g., base station which then broadcasts to consumers, or possibly broadcast updates directly to consumers in the region R . We assume that due to capacity constraints at most k sensors can be active. Based on Eq. (1) and for equal transmission delays d and fixed sensor update rate r , it should be clear that $k \leq N \leq \lfloor \frac{1}{rd} \rfloor$. The goal is to select a subset of producers which realizes a compromise between ensuring coverage of consumers' aggregated spatial interest while minimizing the weighted age of the updates consumers would see.

In order to capture the trade-offs between coverage and weighted average age, we shall modify the weighted average age (Definition 5) resulting from selecting sensors X as follows,

$$b^\lambda(X) := \lambda w(C(X)) - \sum_{i=1}^{|\mathcal{P}^X|} w(P_i^X) \text{age}(P_i^X),$$

where λ is a fixed positive parameter which captures the importance of coverage versus the weighted age.

As already specified, since at most k sensors can be selected, our goal is to determine an optimal subset S^* that minimizes the above cost function, i.e.,

$$S^* \in \arg \max_{X \subseteq V} \{ b^\lambda(X) \mid |X| \leq k \}. \quad (10)$$

Such combinatorial problems are NP hard (weighted coverage as shown in [14]), but may satisfy submodularity properties that make greedy approaches quite effective. The following theorem gives such a characterization.

Theorem 2: (Characterization of weighted age) *If age = age₁ and $\lambda \geq d + \frac{2}{3} \frac{1}{r}$ then the weighted average age function $b^\lambda(\cdot)$ satisfies the following properties*

(Monotonicity) *It is monotonically increasing, i.e., if $X \subset Y \subset V$ then $b^\lambda(X) \geq b^\lambda(Y)$.*

(Submodularity) *It is submodular, i.e., if $X \subset Y \subset V$ and $v \notin Y$ then,*

$$b^\lambda(X) - b^\lambda(X \cup \{v\}) \geq b^\lambda(Y) - b^\lambda(Y \cup \{v\}). \quad (11)$$

Similarly if age = age₂ and $\lambda > 1$ then the weighted γ -age violation function $a^\lambda(\cdot)$ is monotonic and submodular.

The proof of this theorem has been left out of the paper due to lack of space.

Although Problem (10) is a complex combinatorial problem, the classical greedy algorithm shown in Algorithm 1 panel requires $O(|V|k)$ function evaluations to determine a subset $S^{(k)}$ which is $1 - 1/e$ constant factor of the optimal [13], i.e.,

$$b^\lambda(S^{(k)}) \geq (1 - \frac{1}{e}) b^\lambda(S^*).$$

There are computationally less costly possibly distributed versions of the algorithm leveraging random sampling. There is a growing line of work to design possibly distributed algorithms with sub-linear cost which have shown to be similarly effective [15]–[17].

Algorithm 1: Greedy submodular optimization [13]

```

Let  $S^{(0)} = \emptyset$ 
for  $i=0, \dots, k-1$  do
     $j \leftarrow \arg \max_j b^\lambda(S^{(i)} \cup \{j\}) - b^\lambda(S^{(i)})$ ;
     $S^{(i+1)} \leftarrow S^{(i)} \cup \{j\}$ ;
end

```

V. OPTIMIZATION OF SENSOR UPDATE RATES

In this section, we consider a setting where the set of active sensors, say without loss of generality V is fixed, but their update rates $\mathbf{r} = (r_v : v \in V)$ can be jointly optimized subject to the communication capacity constraint Eq. (1).

Since in this setting the set of sensors is fixed we shall modify the weighted average age introduced in Definition 5, to be a function solely of the sensor update rates $\mathbf{r} = (r_v : v \in V)$. We let the weighted average age of the covered set $C(V)$ be,

$$a(\mathbf{r}) := \sum_{i=1}^{|\mathcal{P}^V|} w(P_i^V) \text{age}(\mathbf{r}(V_i^V)). \quad (12)$$

In this section we will focus on the case where age = age₁, i.e., on selecting the sensor update rates so as to minimize the weighted age which be formally stated as follows:

Problem 1: (Age minimization)

$$\min_{\mathbf{r}} \{ a(\mathbf{r}) \mid \mathbf{r} \geq 0, \sum_{v \in V} dr_v \leq 1 \}. \quad (13)$$

Proposition 1: (Age minimization for sensors with disjoint coverage) *Suppose the sensor coverage sets $(C_v, v \in V)$ are disjoint then the age minimization Problem 1 is convex and reduces to,*

$$\min_{\mathbf{r}} \{ \sum_{v \in V} w(C_v) (d + \frac{1}{2} \frac{1}{r_v}) \mid \mathbf{r} \geq 0, \sum_{v \in V} dr_v \leq 1 \},$$

whose optimal joint update rates \mathbf{r}^ are given by*

$$r_v^* = \frac{\sqrt{w(C_v)}}{\sum_{u \in V} \sqrt{w(C_u)}} \times \frac{1}{d}.$$

The proof of this proposition follows from standard convex optimization tools and so left out. The solution reveals the first basic insight that for sensors with disjoint coverage set, the age minimizing rate allocation is proportional to the *square-root* of the weight (e.g., area) of the coverage set each sensor is tracking. Thus equal weight sensors, would lead to equal update rate allocations. The general case where sensors have overlapping coverage sets is unfortunately more complex.

Proposition 2: (Characterization of general age minimization problem) *For the general age minimization Problem 1 where coverage sets may overlap, the objective function given in Eq. (12) is a weighted sum of a convex function and a non-convex/non-concave function, and hence belongs to the family of non-convex/non-concave functions.*

It is easy to see this by noting that the average age of a partition as given in Eq. (4) can be re-written as,

$$\text{age}_1(\mathbf{r}) = f(\mathbf{r}) + g(\mathbf{r}),$$

where $f(\mathbf{r}) = d + \frac{1}{2r_1}$ and where

$$g(\mathbf{r}) = - \int_d^{d+\frac{1}{r_1}} z_1(x) \left[1 - \prod_{i=2}^n z_i(x) \right] dx.$$

It is clear that $f(\mathbf{r})$ is convex in r_1 , while $g(\mathbf{r})$ is non-convex/non-concave in \mathbf{r} which can be proved by finding the Hessian of the function $g()$ with respect to \mathbf{r} , $H \in \mathbb{R}^{n \times n}$, and either showing that H has a mix of positive and negative eigenvalues, or that $\mathbf{y}^T H \mathbf{y}$, for all $\mathbf{y} \in \mathbb{R}^{n \times 1}$ can either be positive or negative. Hence, $g()$ has a saddle point. Given that $\text{age}_1(\mathbf{r})$ is part of the objective function in Eq. (12), then the latter belongs as well to the same family of functions. It should be clear by now that in the case of a single sensor v observing a partition and updating at a rate r_v , the age of this partition is convex in r_v and given by $d + \frac{1}{r_v}$. But whenever more than one sensor are observing the same partition, the nature of the function capturing the age of this partition cannot be determined a-priori, and hence it belongs to the family of non-convex/non-concave functions. There exists a family of algorithms that solves for this type of functions, from which we pick the Frank-Wolfe algorithm which provides a $(1 - 1/e)$ approximation guarantee under closed convex constraints [18].

Algorithm 2: Frank-Wolfe Algorithm [19]

```

Let  $\mathbf{r} \in \mathcal{D}$ 
for  $t=0, \dots, T$  do
     $\mathbf{s} \leftarrow \arg \max_{\mathbf{s} \in \mathcal{D}} \langle \mathbf{s}, \nabla (a(\mathbf{r}^{(t)})) \rangle; \gamma := \frac{2}{t+2};$ 
     $\mathbf{r}^{(t+1)} = (1 - \gamma)\mathbf{r}^{(t)} + \gamma\mathbf{s};$ 
end

```

A fundamental goal in non-convex/non-concave optimization is to reach a stationary point [18]. We summarize the FW algorithm as follows. At each iteration t , the algorithm computes the maximal step it can take in the direction of the gradient of the function while satisfying the constraint $\mathbf{s} \in \mathcal{D}$, and then moves in the direction of this maximizer. This process, as explained in [18], intuitively makes sense since we try to find the direction in which we can maximize the improvement in the function value while remaining feasible. Additionally, one key advantage of this algorithm is that it doesn't need to project back into the constraint set, given that it never leaves it.

VI. NUMERICAL RESULTS

We developed a simulation framework to explore the optimization of coverage vs. weighted age trade-offs in collaborative sensing applications which we present in this section.

A. Model

We shall present results for a two-way highway which we model as a rectangular region R of length 1000 m and width 24 m (roughly corresponding to 6 lanes). Vehicles are modeled as $4.8 \text{ m} \times 1.8 \text{ m}$ rectangles with omnidirectional sensors placed in the center (rooftop). The unobstructed coverage set for each sensor is a disc of radius $r = 50 \text{ m}$ with area πr^2 . The coverage area of a sensor does not include the regions off the road, and the only obstruction present are other vehicles blocking its field-of-view. We assume that vehicles are randomly placed in a lane on the road, with spacing at least 10 m between any two sensors

on the same lane. Assuming the speed of a typical vehicle is $s = 60 \text{ km/hour}$, and the reaction time is $t_{\text{interest}} = 5 \text{ sec}$, each sensor v is at the center of its own rectangular *region of interest*, R_v with length $2 \cdot s \cdot t_{\text{interest}} \approx 167 \text{ m}$ and width covering the highway's six lanes. We assume vehicles (consumers) are interested in tracking sensed information in all their region of interest.

B. Communication model

We assume that vehicles are equipped with LIDAR sensors which sample around 1.3 million points per second with a corresponding data rate of 4 Mbytes/sec and update rate ranging between 5 and 30 Hz. Additionally, millimeter-wave (mmWave) technology is used to transmit the producers' updates. We assume that one producer accesses the medium at a time, and broadcasts its update to all the consumers in the system (i.e., all available sensors, including itself). For simplicity we assume an operational bandwidth achieving a data rate of 1.5 Gbps. We combine both the channel access time ($\sim 1 \text{ msec}$) and updates' transmission time into a single deterministic value, d , which we compute according to the previous given values, and find $d = 4.2 \text{ msec}$. We further assume no transmission failures and ignore mmWave blockage.

C. Coverage and spatial average age of the region of interest of a typical sensor

We define the *coverage* of the region of interest R_v of a typical sensor v as the overall percentage of area covered within R_v , which combines the regions that v observes within R_v as well as the updates it receives from active producers observing locations in R_v , that v cannot see. The weighted average age of a particular region in R_v is the product of the weight that sensor v assigns to this region multiplied by the associated age which depends on the number of producers generating updates about this location. The spatial average age of the coverage of a typical consumer v is the normalized weighted average age of updates about regions in R_v received by v via collaboration with other sensors.

We evaluate the performance of three different schemes in terms of both the coverage and the spatial average age of updates of a typical sensor. We assume that the system has N available sensors. The first scheme, denoted *baseline*, selects all N sensors in the region but must satisfy the lower update rate per sensor due to the capacity constraint. One would expect that this technique achieves the best coverage while performing poorly in terms of the spatial average age, given that a lower update rate is assigned per producer. The second and third schemes employ sensor selection, where $k < N$ producers are selected, allowing for a higher per sensor update rates and hence fresher and more frequent updates at the consumer side. Specifically in the second scheme, denoted *aggregated spatial interest*, the weight function is the sum of the consumers' spatial interest which in turn is assumed to be uniform over locations (in the region of interest of a typical consumer) that are either obstructed or not covered. This model is intuitive in that if many consumers share the same obstructed region, then the weight placed there is increased, and thus producers will be favored to cover this region in order to keep its age low. Finally, the third scheme, denoted *uniform aggregated spatial*





Fig. 5: Coverage and spatial average age of the region of interest of a typical sensor as k increases, with $N = 15$ sensors.


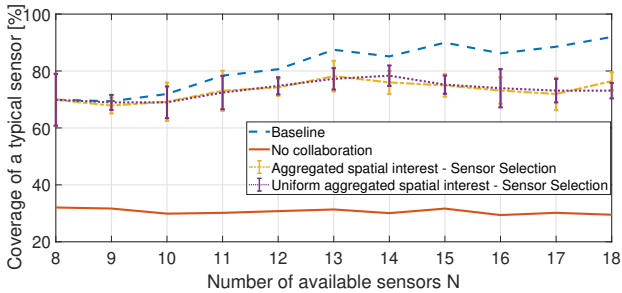


Fig. 6: Coverage and spatial average age of the region of interest of a typical sensor as N increases, with $k = 8$ sensors.

interest, is not consumer oriented, but instead uses an aggregate interest measure which is uniform over the overall region in consideration, R . In our case, this technique will attempt to cover as much of the highway as possible, i.e., spread the selected sensors as much as possible on the map, while trying to keep the unweighted spatial average age as low as possible, which can only be achieved by enhancing producers' coverage overlap. This tension between spreading producers over the map and overlapping their coverage sets is what makes this technique interesting.

Below we verify the robustness of our *aggregated spatial interest* model in achieving a good coverage-age trade-off for a typical consumer. In Fig.5, we fix the number of available sensors in R to $N = 15$ and increase the number of selected producers, k , for both sensor selection schemes, from 1 to N .

The left subfigure of Fig. 5 clearly shows that without collaboration (no information sharing), the coverage of a typical sensor is the lowest ($\sim 30\%$), since sensors do not share information, while the baseline achieves the highest coverage per typical sensor since all N producers are sharing their respective updates with each other. Interestingly, both sensor selection schemes achieve similar coverages as k increases, which shows that the *aggregated spatial interest* scheme performs well in terms of coverage of a typical sensor's region of interest.

The right subfigure of Fig.5 evaluates the spatial average age of updates about the region of interest of a typical sensor. We additionally consider the joint optimization of the update rate allocations across the k sensors selected using the *aggregated spatial interest* technique, in order to further minimize the spatial average age. We refer to this technique as *age minimization*, shown in the solid orange curve. The age of the region solely seen by a sensor v , within R_v , has age 0,

since v is updated about what it sees (without collaboration) at all times. Otherwise, the spatial average age of the covered regions within R_v is a function of the number of producers (different from v) that see these regions. The *baseline* achieves a constant age given that the number of selected sensors is always $N = 15$. While the ages achieved by the other schemes increase, since the number of producers increases, which results in lower update rates allocated per producer, until $k = N = 15$ is reached, where all ages (except for the one achieved by *age minimization*) overlap. Three interesting observations are that (1) for a small number of selected sensors ($k \leq 5$), the non-consumer oriented scheme achieves a better age than the consumer oriented approach, since the spreading of the weights assigned by the consumers is roughly uniform over the map, given the nature of the model in consideration, which gives a slight advantage to the former scheme when the number of producers is low, (2) the consumer oriented sensor selection technique results in an overall better age than the non-consumer oriented scheme (for $k \geq 5$), and (3) for $k = 15$, the *age minimization* technique results in a better age of updates per typical consumer than the *baseline*, with an approximate improvement of 5 msec.

In Fig. 6, we increase the number of available consumers from 8 to 18, and fix the number of selected producers k to 8 for both the *aggregated spatial interest* and *uniform aggregated spatial interest* schemes. In the left subfigure of Fig. 6, both sensor selection schemes have an overall similar coverage per typical sensor demonstrating that the scheme based on aggregating consumers' spatial interests performs well in terms of coverage. As N increases from 8 to 14, the coverages achieved by these two schemes increase from 70% to 80% due to better collaboration between the selected producers, but then slightly decreases ($N \geq 14$) due to more obstructions caused by the

increasing number of available sensors. Considering the spatial age, we clearly see that the *baseline* achieves the worst age given that the number of producers increases, resulting in lower update rates allocated per sensor. Interestingly, the overall age performance of the consumer oriented technique is better than the one achieved by the *uniform aggregated spatial interest* scheme, which shows that the former allows for a selection of a subset of sensors that better meets the consumers' requirements. Finally, we note that the improvement between the *baseline* and the *aggregated spatial interest* is large (~ 20 msec for $N=18$).

We point out that the highway scenario in consideration slightly exhibits the advantage of using a consumer's oriented model, since the consumers' interests are uniform over the highway, due to the nature of the model, and not specific to a particular spot on the map. A more relevant scenario should consider a higher spatial correlation between the obstructed regions of different consumers, which results in a clear improvement in the coverage-age achieved by the *aggregated spatial interest* scheme, e.g., a road intersection where many cars express their interest in the same "unseen" (obstructed) region, or large trucks obstructing the field-of-view of many cars.

VII. CONCLUSION AND FUTURE WORK

In this paper we have exhibited an approach to achieving trade-offs between coverage and timeliness communication constrained collaborative sensing setting wherein spatially distributed sensor nodes can serve dual roles as producers and consumers of sensed information. The proposed framework allows quite a bit of flexibility in terms of capturing the underlying characteristics of information sharing, and suggests the development of a common platform for sharing real-time sensor data in a context dependent manner, i.e., matching nodes' current interest, in order to minimize situational uncertainty so as to enhance vehicular flow and safety. A key aspect is the design of strategies that select producers that attempt to best meet the consumers' coverage and age of updates requirements, subject to communication network capacity constraints. A key part of our future work is to make the sensor selection and update rate optimization more scalable to robustly address heterogeneous and dynamic environments.

VIII. APPENDIX

A. Proof of Theorem 1

Given the saw-tooth character of the time varying age per sensor and the randomization of the update phases, the age at a random time is given by $A = \min_{v=1, \dots, n} A_v$, where $A_v = d + \frac{1}{r_v} U_v$ and U_v , for $v = 1, \dots, n$ are i.i.d. $\text{Unif}[0, 1]$ and independent, and $\text{age}_1(\mathbf{r}) = \mathbb{E}[A]$, while $\text{age}_2(\mathbf{r}) = \mathbb{P}(A \geq \gamma)$. It is easy to see that given A_v is a shifted and scaled uniform random variable such that

$$P(A_v > y) = \begin{cases} 1, & \text{if } 0 \leq y \leq d, \\ 1 - r_v(y - d), & \text{if } d \leq y \leq d + \frac{1}{r_v}, \\ 0, & \text{if } d + \frac{1}{r_v} < y, \end{cases}$$

and given that $P(A > y) = \prod_{v=1}^n P(A_v > y)$, and further given that $r_1 \geq r_2 \geq \dots \geq r_n$, we have

$$P(A \geq y) = \begin{cases} 1, & \text{if } 0 \leq y \leq d, \\ \prod_{i=1}^n z_i(y), & \text{if } d < y \leq d + \frac{1}{r_1}, \\ 0, & \text{if } d + \frac{1}{r_1} < y, \end{cases} \quad (14)$$

where $z_i(y) = 1 - r_i(y - d)$ for $i = 1, \dots, n$. Thus the γ -age violation function for γ such that $d < \gamma \leq d + \frac{1}{r_1}$ is given by,

$$\text{age}_2(\mathbf{r}) = \prod_{i=1}^n z_i(\gamma).$$

We can compute the average age function as follows

$$\text{age}_1(\mathbf{r}) = \int_0^\infty P(A > y) dy = d + \int_d^{d + \frac{1}{r_1}} \prod_{j=1}^n z_j(y) dy.$$

REFERENCES

- [1] A. Rauch, F. Klanner, R. Raschhofer, and K. Dietmayer, "Car2x-based perception in a high-level fusion architecture for cooperative perception systems," in *Intelligent Vehicles Symposium (IV), 2012 IEEE*. IEEE, 2012, pp. 270–275.
- [2] S.-W. Kim, Z. J. Chong, B. Qin, X. Shen, Z. Cheng, W. Liu, and M. H. Ang, "Cooperative perception for autonomous vehicle control on the road: Motivation and experimental results," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 5059–5066.
- [3] X. Zhao, K. Mu, F. Hui, and C. Prehofer, "A cooperative vehicle-infrastructure based urban driving environment perception method using a D-S theory-based credibility map," *Optik-International Journal for Light and Electron Optics*, vol. 138, pp. 407–415, 2017.
- [4] Y. Wang, G. de Veciana, T. Shimizu, and H. Lu, "Performance and scaling of collaborative sensing and networking for automated driving applications," in *Proc. IEEE ICC Workshop 5G and Cooperative Autonomous Driving*, May 2018, pp. 1–6.
- [5] Y. Wang, G. de Veciana, T. Shimizu, and H. Lu, "Performance and Scaling of Collaborative Sensing and Networking for Automated Driving Applications," Dec 2018.
- [6] A. Kosta, N. Pappas, V. Angelakis *et al.*, "Age of information: A new concept, metric, and tool," *Foundations and Trends® in Networking*, vol. 12, no. 3, pp. 162–259, 2017.
- [7] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *2012 Proceedings IEEE INFOCOM*. IEEE, 2012, pp. 2731–2735.
- [8] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Minimizing the age of information in broadcast wireless networks," in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2016, pp. 844–851.
- [9] R. D. Yates and S. K. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Transactions on Information Theory*, vol. 65, no. 3, pp. 1807–1827, 2018.
- [10] R. Talak, S. Karaman, and E. Modiano, "Optimizing information freshness in wireless networks under general interference constraints," in *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*. ACM, 2018, pp. 61–70.
- [11] J. Rahal, G. de Veciana, T. Shimizu, and H. Lu, "Optimizing networked situational awareness," in *Proc. IEEE WIOPT RAWNET Workshop*, June 2019, pp. 1–8.
- [12] A. Prasad, S. Jegelka, and D. Batra, "Submodular meets Structured: Finding Diverse Subsets in Exponentially-Large Structured Item Sets," *arXiv e-prints*, Nov 2014.
- [13] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—i," *Mathematical programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [14] U. Feige, "A threshold of $\ln n$ for approximating set cover," *Journal of the ACM (JACM)*, vol. 45, no. 4, pp. 634–652, 1998.
- [15] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, "Distributed submodular maximization: Identifying representative elements in massive data," in *Advances in Neural Information Processing Systems*, 2013, pp. 2049–2057.
- [16] E. Balkanski and Y. Singer, "The adaptive complexity of maximizing a submodular function," in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2018, pp. 1138–1151.
- [17] A. Ene and H. L. Nguyen, "Submodular maximization with nearly-optimal approximation and adaptivity in nearly-linear time," in *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2019, pp. 274–282.
- [18] A. A. Bian, B. Mirzasoleiman, J. M. Buhmann, and A. Krause, "Guaranteed non-convex optimization: Submodular maximization over continuous domains," 2016.
- [19] M. Jaggi, "Revisiting frank-wolfe: Projection-free sparse convex optimization." 2013, pp. 427–435.