Minimizing Age of Correlated Information for Wireless Camera Networks

Qing He, György Dán, Viktoria Fodor
Department of Network and Systems Engineering, School of Electrical Engineering and Computer Science,
KTH Royal Institute of Technology, Stockholm, Sweden
Email: {*qhe, gyuri, vfodor*}@*kth.se*

***Abstract*—Freshness of information is of critical importance for a variety of applications based on wireless camera networks where multi-view image processing is required. In this study, we propose to jointly optimize the use of communication and computing resources such that information from multiple views is delivered is obtained in a timely fashion. To this end, we extend the concept of age of information to capture packets carrying correlated data. We consider the joint optimization of processing node assignment and camera transmission policy, so as to minimize the maximum peak age of information from all sources. We formulate the multi-view age minimization problem (MVAM) and prove that it is NP-hard. We provide fundamental results including tractable cases and optimality conditions. To solve the MVAM efficiently, we develop a modular optimization algorithm following a decomposition approach. Numerical results show that, by employing our approach, the maximum peak age is significantly reduced in comparison to a traditional centralized solution with minimum-time scheduling.**

I. Introduction

Wireless camera networks acquire, process, and analyze digital images of areas or objects of interest. They form an important building block of future smart cities, and serve a variety of applications, such as surveillance [1], tracking [2], healthcare [3], and intelligent transportation [4].

Many of these applications require multiple cameras with overlapping fields of view (FoV) to monitor a given scene, as doing so can improve the robustness and accuracy of tracking, and enables 3D scene reconstruction. Images from cameras with overlapping FoVs have to be processed jointly, and under strict delay constraints to enable real-time operation. The emerging paradigm of fog computing could enable meeting the strict delay constraints, as it allows to distribute computation, communication, control and storage to computing nodes close to the cameras, referred to as fog nodes. Still, real-time processing of the visual information in a fog computing enabled wireless camera system requires joint optimization of the allocation of computing resources and the transmission scheduling of images from cameras with overlapping FoVs.

A promising metric for quantifying the timeliness of end-to-end data delivery, including queuing and transmission times, is the recently introduced age of information, or simply, age [5]. Age is commonly defined as the time elapsed since the most recently received message was generated (see [6] and references therein), as shown in Figure 1. The average age of information, calculated as the area under the sawtooth curve in Figure 1, normalized by the observation interval,

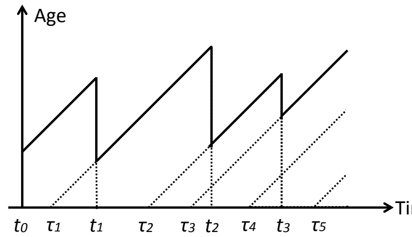

Fig. 1. Age evolution of a source. The i^{th} packet with time stamp τ_i is received by the destination at t_i. The age at t_i equals $t_i - \tau_i$. If no packet is received, the age increases linearly.

was considered in [5], [7] for capturing the average system behavior in a queueing system fed by a single source and by multiple sources, respectively. To characterize the worst case system behavior, the authors of [8] introduced the notion of peak age, which is the maximum value of the age achieved immediately before receiving a new packet (e.g., in Figure 1, the i^{th} peak age of the source is obtained observing the i^{th} peak value in the sawtooth curve), and proposed strategies for queue management. Subsequently, [9] considered optimal link scheduling to drain a given set of packets with respect to the total age. In all these works, the age of information changes upon receiving each individual packet. Nonetheless, when processing requires information from multiple senders, carried in different packets, age of information should change only when all packets carrying correlated information are received. Hence, existing results do not apply for the case of packets carrying correlated information and delivered by multiple transmitters, such as the case of cameras with overlapping FoVs.

In this paper, we adapt the concept of age of information to correlated information, which allows us to apply this metric to wireless camera networks as well as to other systems where an age update is triggered by multiple correlated packets. We consider the joint optimization of fog node assignment and transmission scheduling, so as to minimize the age of the multi-view image data at the fog nodes. We formulate the optimization problem and derive the fundamental results about problem complexity, tractable cases, and optimality conditions. We propose a heuristic algorithm based on problem decomposition, and use simulations to explore the benefits of the jointly optimized assignment and transmission strategy in improving the freshness of information in wireless camera networks.

The rest of the paper is organized as follows. In Section II, we define the system model and formulate the problem, followed by the complexity analysis in Section III. In Section IV, we present structural results, which motivate the optimization algorithm in Section V. Numerical results are provided in Section VI. In Section VII, we conclude the paper.

II. System Model and Problem Formulation

We consider a wireless camera network that consists of a set of cameras $\mathcal{C} = \{1, 2, \ldots, C\}$, a set of scenes $\mathcal{S} = \{1, 2, \ldots, S\}$, which we also refer to as sources, and a set of fog (computing) nodes $\mathcal{N} = \{1, 2, \ldots, N\}$. The cameras have overlapping FoV, and each scene is covered by multiple cameras. We denote by $\mathcal{C}(s)$ the set of cameras that have a view of scene s, and we assume that cameras monitoring the same scene capture images simultaneously. The cameras send the images to their respective serving fog node for processing, and we define

$$l_{cn} = \begin{cases} 1 & \text{if camera } c \text{ is served by node } n; \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

To allow the processing of multi-views, the cameras $\mathcal{C}(s)$ that cover scene s need to transmit their images to the same fog node, hence we have

$$\sum_{n\in N} \prod_{c\in\mathcal{C}(s)} l_{cn} = 1, \ \forall s \in \mathcal{S}. \tag{2}$$

The images are delivered to the fog nodes via a shared wireless channel. To determine whether or not a subset of the cameras are allowed to transmit simultaneously, we follow the physical model, aka, the signal-to-interference-and-noise ratio (SINR) model. For a subset of cameras $\mathcal{g} \subseteq \mathcal{C}$ to be able to transmit together, given SINR thresholds must be met at each fog node serving the cameras,

$$\begin{aligned} \mathrm{SINR}_n(c, \mathcal{g}) \triangleq & \frac{P_c G_{cn}}{\sum_{\iota\in\mathcal{g},\iota\neq c} P_\iota G_{\iota n} + \sigma_n^2} \geq \gamma_c, \\ & \forall c \in \mathcal{g}, \ n \in \{\mathcal{N} : l_{cn} = 1\}, \end{aligned} \tag{3}$$

where P_c is the transmit power of camera c and G_{cn} is the channel gain between camera c and fog node n, incorporating the effects of path loss, shadowing and fading, and σ_n^2 is the noise variance.

We consider that time is slotted and is divided into subsequent scheduling cycles. In each scheduling cycle, we schedule the transmission of the images queued at the beginning of the cycle. We denote by U_{ci} the i^{th} image in the queue of camera c. The images that arrive during the cycle are queued by the cameras under the first-come-first-served (FCFS) discipline, and will be scheduled in the next cycle. The transmission rate of a camera-fog node pair is one image per time slot if the SINR threshold is met.

Let us denote by t_0 the starting time of the current scheduling cycle, and by K_c the number of queued images of camera c at t_0. Furthermore, we denote by t_j the time at the end of the j^{th} slot of the scheduling cycle. Since the image data are queued using FCFS and the cameras monitoring the same scene capture images simultaneously, the i^{th} image in the queue of each camera $c \in \mathcal{C}(s)$ carries time stamp τ_{si}. We refer to the set of images $\{U_{ci}, \forall c \in \mathcal{C}(s)\}$, as the i^{th} image block of s.

Let us denote by a_{s0} the initial age of scene s at time t_0. Due to the requirement of multi-view processing, the information of scene s will not be updated until all the images of an image block are delivered to the fog node n serving cameras $\mathcal{C}(s)$. Therefore, at time t_j the age of a source s is defined as

$$a_{sj} = \begin{cases} t_j - \tau_{si} & \text{if all the images of the } i^{th} \text{ image} \\ & \text{block of } s \text{ have been delivered to} \\ & \text{node } n \text{ exactly by } t_j; \\ a_{s,j-1} + 1 & \text{otherwise.} \end{cases} \tag{4}$$

Note that the age calculation in (4) differs from the case where the age is updated upon the delivery of each packet, as considered in [5]–[12]. This also affects the definition of the peak age, as the peak ages of a scene are attained immediately before the last image of an image block is delivered to the fog node. To express the peak ages, let us denote by t_{ci} the time when image U_{ci} is delivered, and by $T_{ci} = t_{ci} - t_0$ the number of time slots before U_{ci} is delivered, and observe that T_{ci} is a positive integer in $[1, \sum_{c\in\mathcal{C}} K_c]$. Then by defining $\tau_{s0} = t_0 - a_{s0}$, we can express the i^{th} peak age of s as $\alpha_s^i = \max_{c\in\mathcal{C}(s)} t_{ci} - \tau_{s,i-1} = t_0 + \max_{c\in\mathcal{C}(s)} T_{ci} - \tau_{s,i-1}$.

In order to optimize the worst case performance, we formulate the multi-view age minimization (MVAM) problem as that of minimizing the maximum peak age of all sources,

$$\underset{\{T_{ci}\in Z^+\}}{\text{minimize}} \quad \max_{s\in\mathcal{S}, \ c\in\mathcal{C}(s), \ i=1,\ldots,K_c} \alpha_s^i \tag{5a}$$

$$\text{subject to} \quad (2), \ (3), \ \text{and}$$

$$\alpha_s^i = t_0 + \max\nolimits_{c\in\mathcal{C}(s)} T_{ci} - \tau_{s,i-1}, \tag{5b}$$

$$1 \leq T_{c1} < T_{c2} < \cdots < T_{c,K_c} \quad \forall c \in \mathcal{C}, \tag{5c}$$

$$\mathrm{SINR}_n(c, \mathcal{g}_j) \geq \gamma_c, \forall c \in \mathcal{g}_j, T_j = 1, \ldots, \sum_{c\in\mathcal{C}} K_c,$$

$$\mathcal{g}_j = \{c \in \mathcal{C} : T_{ci} = T_j, i = 1, \ldots, K_c\}. \tag{5d}$$

Observe that solving MVAM requires joint optimization of the camera to fog node assignment, and of the transmission schedule of the cameras to their serving fog nodes. We remark that at the optimum of (5), if the solution uses $T < \sum_{c\in\mathcal{C}} K_c$ slots, then $\mathcal{g}_j$ are empty sets for $j = \{T+1, \ldots, \sum_{c\in\mathcal{C}} K_c\}$. Hence in the formulation, T can be overestimated without loss of optimality. We summarize the key notation in Table I.

III. MVAM Complexity

Theorem 1. *The MVAM, as defined in* (5)*, is NP-hard.*

Proof: The decision problem of the MVAM is to determine whether or not there exists a solution such that the maximum peak age is no more than a given value. We show that the decision problem is NP-complete by constructing a

TABLE I

Notation	Description
$\mathcal{S}$	the set of scenes
$\mathcal{C}$	the set of cameras
$\mathcal{N}$	the set of fog nodes
$\mathcal{C}(s)$	The subset of cameras monitoring scene s
$\mathcal{G}$	A subset of cameras that can transmit together
t_0	The initial time
t_j	The time corresponds to the end of the j^{th} time slot
K_c	The number of queued images in camera c
U_{ci}	The i^{th} image of camera c
τ_{si}	The time stamp carried by U_{ci}, $\forall c \in \mathcal{C}(s)$
a_{s0}	The initial age of s at t_0
a_{sj}	The age of s at t_j
α_s^i	The i^{th} peak age of s
T_{ci}	The number of slots before U_{ci} is delivered

polynomial-time reduction from the 3-satisfiability (3-SAT) problem, which is known to be NP-complete [13].

For an instance of the 3-SAT problem, denote by $\mathcal{S}$ and $\mathcal{S}'$ the set of its positive literals and its negative literals, respectively. The set of clauses is denoted by $\mathcal{D}$. We construct an MVAM instance with $|\mathcal{S}|+|\mathcal{D}|$ scenes and the same number of fog nodes. Each scene is covered by two cameras. Based on the index, we map the members of $\mathcal{S}$ and $\mathcal{S}'$ (i.e., the literals) to the cameras associated with the first $|\mathcal{S}|$ scenes. Thus, the i^{th} scene is monitored by the two cameras representing literals s_i and $\bar{s}_i$. The j^{th} clause in $\mathcal{D}$, i.e., d_j, is mapped to the two cameras with a view of the $(|\mathcal{S}|+j)^{th}$ scene. We denote these two cameras by d_j' and d_j''. The initial ages of the first $|\mathcal{S}|$ scenes and the remaining $|\mathcal{D}|$ scenes are a_0 and a_0+1, respectively. Each camera contains one image to be delivered.

The transmit power of all cameras is uniformly set to one. For cameras s_i and $\bar{s}_i$, $i=1,2,\ldots,|\mathcal{S}|$, the SINR threshold is $\gamma_c=2$. For cameras d_j' and d_j'', $j=1,2,\ldots,|\mathcal{D}|$, the SINR threshold is $\gamma_c=\frac{1}{3}$. The noise power is $\sigma^2=0.5$ at the fog nodes 1 to $|\mathcal{S}|$, and $\sigma^2=1$ at the remaining $|\mathcal{D}|$ nodes. The channel gains between the cameras and the nodes are as follows.

$$G_{cn}=\begin{cases} 1 & \text{if } c=s_i \text{ or } \bar{s}_i,\ n=i,\ i=1,\ldots,|\mathcal{S}|; \\ 1 & \text{if } c=d_j' \text{ or } d_j'',\ n=|\mathcal{S}|+j,\ j=1,\ldots,|\mathcal{D}|; \\ \frac{1}{|\mathcal{S}|-1} & \text{if } c=s_i \text{ or } \bar{s}_i,\ i=1,\ldots,|\mathcal{S}|,\ n=|\mathcal{S}|+1, \\ & \ldots,|\mathcal{S}|+|\mathcal{D}|, \text{ and } c \text{ is not one of the} \\ & \text{three literals of clause } d_{n-|\mathcal{S}|}; \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Below we show that the 3-SAT instance is satisfiable if and only if the above constructed MVAM instance has a feasible solution for the objective value a_0+2. By construction, we observe that, in a feasible solution, the literal cameras s_i and $\bar{s}_i$ will be served by the fog node i, and the clause cameras d_j' and d_j'' will be served by the fog node $|\mathcal{S}|+j$. Due to the SINR constraint, cameras s_i and $\bar{s}_i$ cannot transmit simultaneously, and consequently, at least two time slots are needed to deliver all images. Therefore, the minimum objective value for the MVAM instance is at least a_0+2. This is achieved if and only if a feasible solution with the following properties exists.

- All images are delivered in two time slots.
- Cameras s_i and $\bar{s}_i$ transmit in different time slots, and hence $|\mathcal{S}|$ literal cameras transmit simultaneously in each time slot.
- All the $2|\mathcal{D}|$ clause cameras transmit in the first time slot, together with $|\mathcal{S}|$ literal cameras. Due to the choice of the channel gains and transmit powers, for each clause, at least one camera representing one of its three literals transmit in the first time slot. Otherwise, suppose it is not the case for a clause, then the $|\mathcal{S}|$ literal cameras that transmit in the first time slot would all interfere with the two corresponding clause cameras. This, together with the interference they receive from each other, would result in an SINR $=\frac{1}{\frac{|\mathcal{S}|}{|\mathcal{S}|-1}+1+1}<\frac{1}{3}$.

Based on the above analysis, we establish a mapping from the solution to that of the corresponding 3-SAT problem by setting all the literals represented by the cameras transmitting in the first time slot to be true, and the others to be false. Due to the third property, for the 3-SAT instance, every clause has at least one TRUE literal, and hence the 3-SAT instance is satisfiable. On the contrary, if the 3-SAT instance is satisfiable, then we map its solution to the corresponding MVAM problem by scheduling all the cameras representing the clauses and the TRUE literals in the first time slot, and the rest of the literal cameras in the second time slot. By doing so, we obtain a feasible solution of the MVAM instance with the objective value of a_0+2. Therefore, the 3-SAT instance is satisfiable if and only if there is a feasible solution for the corresponding MVAM with the objective value of a_0+2. Hence the decision problem of the MVAM is at least as hard as the 3-SAT problem. Since the 3-SAT problem is NP-complete, we conclude that the MVAM is NP-hard. ■

IV. Structural Results

Since the MVAM is hard in general, it is of interest to identify tractable cases. We first consider an "ideal" case where the network allows all cameras to transmit together.

Theorem 2. *Any MVAM instance in which all cameras are able to transmit simultaneously can be recognized and solved in polynomial-time.*

Proof: To recognize the case defined in this theorem, given an MVAM instance, we calculate the SINR value for each camera-node pair, with the assumption that all the cameras are active, i.e., in the denominator of (3), the interference takes into account the signals from all other cameras. The computational complexity in this step is $O(CN)$. For each $s\in\mathcal{S}$, we set the binary variable $\iota_{sn}=1$ if and only if for each camera $c\in\mathcal{C}(s)$, the SINR between c and n exceeds its threshold. If for each $s\in\mathcal{S}$, there exists $n\in\mathcal{N}$ such that $\iota_{sn}=1$ holds, then we conclude that for this MVAM instance, all cameras can transmit together. This step takes time $O(SN)$, and the solution of the fog-node assignment task is derived by setting $l_{cn}=\iota_{sn}$, $\forall c\in\mathcal{C}(s)$.

The optimal transmission strategy for this case is straightforward, that is, in each time slot, all cameras transmit together as long as they have images in their queues. By (5), the optimal objective equals $t_0 + \max_{s\in\mathcal{S},i\in\mathcal{K}_s}(i - \tau_{s,i-1})$, where $\mathcal{K}_s = \{1,\ldots,K_c\}$, $c \in \mathcal{C}(s)$. The computation can be done in polynomial-time and hence the result. ■

Next, we consider the opposite case in which only one camera is allowed to transmit in a time slot, i.e., a time division multiple access (TDMA) transmission policy. This case corresponds to severely interference-limited networks, e.g., when the cameras and nodes are densely located, and hence each camera causes significant interference to all nodes. We first establish a theoretical result concerning the transmission schedule.

Lemma 3. *Consider an instance of the MVAM with TDMA. There exits an optimal transmission schedule where the images in an image block, i.e., U_{ci}, $\forall c \in \mathcal{C}(s)$, are delivered in consecutive time slots.*

Proof: To prove the lemma, suppose Ω is an optimal solution in which images of an image block are not delivered in consecutive time slots. Denote by U^{λ}_{ci}, $\lambda = 1,\ldots,|\mathcal{C}(s)|$, the λ^{th} image of U_{ci}, $\forall c \in \mathcal{C}(s)$, delivered in Ω. We construct a new solution by moving all U^{λ}_{ci}, $\lambda = 1,\ldots,|\mathcal{C}(s)|-1$, right before $U^{|\mathcal{C}(s)|}_{ci}$ and shifting the other images in between earlier in time. By doing so, the i^{th} peak age of s remains unchanged and the other peak ages (of s as well as the other sources) either remain unchanged or decrease. Hence the maximum peak age of the new solution is not higher than that of the previous solution. Repeating the operation for all image blocks that are not delivered in consecutive time slots we obtain a schedule Ω' that satisfies the lemma, and has the age less than or equal to that in Ω, which proves the lemma. ■

Theorem 4. *Any MVAM instance in which only TDMA is allowed can be recognized and solved in polynomial-time.*

Proof: To recognize a TDMA instance, we only need to verify that no two cameras can transmit together. In the worst case, this can be done in $O(C^2N^2)$. For the camera assignment, the solution is derived by following an approach similar to the one in the proof of Theorem 2, where we need to verify that for each camera $c \in \mathcal{C}$, there exists at least one fog node n, such that the SNR is above the threshold γ_c.

To construct an optimal transmission schedule, recall that by Lemma 4 there is an optimal solution in which the images in an image block are delivered consecutively. We now construct a schedule Ω, in which the image blocks are scheduled in ascending order of their time stamps τ_{si}, $\forall s \in \mathcal{S}, i = 1,\ldots,K_c$, $c \in \mathcal{C}(s)$. The transmission order of the images in each image block is arbitrary. To show that Ω is optimal, we provide an indirect proof. Assume that Ω is not optimal. Then there is an optimal solution Ω_1, in which there exist two adjacent image blocks B_1 and B_2 with time stamps τ_1 and τ_2, respectively, not transmitted in the defined order. That is, if $\tau_1 > \tau_2$, then B_1 is transmitted before B_2 in Ω_1. Denote by T_0 the time when the transmission of B_1 starts, and by $|B_1|$ and $|B_2|$ the number of images in B_1 and B_2, respectively. By (5b) and because of TDMA, the achieved peak ages by the two blocks are $\alpha^1 = T_0 + |B_1| - \tau_1$ and $\alpha^2 = T_0 + |B_1| + |B_2| - \tau_2$. Assume now that we swap the transmission of B_1 and B_2, obtaining a schedule Ω_2. Clearly, Ω_2 is feasible as it does not violate the FCFS discipline, and the peak ages of B_1 and B_2 change to $\alpha^3 = T_0 + |B_2| - \tau_2$ and $\alpha^4 = T_0 + |B_2| + |B_1| - \tau_1$. Since $\tau_1 > \tau_2$, it can be easily verified that α^2 is maximal one among these four peak ages. Hence the maximal peak age in Ω_2 is either equal to or less than that in Ω_1. Repeating the process for all adjacent image blocks that are not delivered in the defined order (cf. bubble sorting), after a finite number of steps, we obtain the transmission schedule Ω. According to the above analysis, the maximum peak age of Ω cannot be greater than that of Ω_1. This contradicts the assumption that Ω is not optimal, and hence the conclusion.

To construct the optimal transmission schedule, the bottleneck is to sort τ_{si}. The computational complexity is hence $O(SK\log(SK))$, where $K = \max_{c\in\mathcal{C}} K_c$. As all the steps can be done in polynomial-time, the theorem follows. ■

Note that by Theorem 3 and its proof, the optimal transmission strategy is determined solely by the values of the time stamps but not the number of images in each image block, i.e., $|\mathcal{C}(s)|$. Motivated by this observation, following the same flow of arguments, we extend the result of the optimal transmission schedule to a scenario in which only cameras monitoring the same scene are possible to transmit together.

Corollary 5. *If no two cameras $c \in \mathcal{C}(s)$ and $c' \in \mathcal{C}(s')$, $s, s' \in \mathcal{S}$, $s \neq s'$, can transmit simultaneously then in the optimal solution the images are delivered in ascending order of their time stamps, and images of an image block are delivered in the same or consecutive time slots, depending on whether or not the corresponding cameras can transmit simultaneously.*

We can use this result for devising a polynomial-time algorithm for MVAM instances where all cameras monitoring the same scene can transmit together.

Corollary 6. *Any instance of the MVAM in which only cameras $c \in \mathcal{C}(s)$, $\forall s \in \mathcal{S}$ are capable of transmitting together, can be solved in polynomial-time.*

Proof: To assign the optimal node to each camera, a similar process as for the TDMA case can be followed. For each pair of $c \in \mathcal{C}$ and $n \in \mathcal{N}$, since $\mathcal{C}(s)$ are able to transmit together, we calculate the SINR value in which the interference is the sum of the signals from $\mathcal{C}(s)\setminus\{c\}$. Then the steps in the proof of Theorem 2 are used to derive the fog node assignment. As for the case defined in the corollary, only $\mathcal{C}(s)$, $\forall s \in \mathcal{S}$, are compatible sets, i.e., no two cameras $c \in \mathcal{C}(s)$ and $c' \in \mathcal{C}(s')$, $s, s' \in \mathcal{S}$, $s \neq s'$, can transmit simultaneously, and thus by Corollary 5, the optimal schedule follows directly. ■

Finally, we provide a general optimality condition of the transmission schedule that applies to all MVAM instances.

Theorem 7. *Given a transmission schedule, let us call the minimum time stamp of the images that are delivered in a time slot as the slot time stamp. Then for any instance of the MVAM, there exists an optimal schedule in which the slot time* *stamps are non-decreasing.*

Proof: The proof is indirect, based on swapping the transmission of cameras in adjacent time slots, but we omit the proof due to lack of space. ■

V. Correlated Maximum Age First (CMAF) Algorithm

Inspired by the above structural results, in what follows we propose an efficient heuristic algorithm for the MVAM. The *correlated maximum age first (CMAF)* algorithm is based on a decomposition of the MVAM into a camera to node assignment, and for a given assignment it computes a transmission schedule.

A. Camera-node Assignment Algorithms

The *CMAF* uses two polynomial-time camera-node assignment algorithms. The first algorithm is based on the observation that in order to obtain the minimum peak age, in (5), T_{ci}, $\forall c \in \mathcal{C}$, $i = 1, \ldots, K_c$, should be as low as possible. This can be achieved if as many as possible cameras are active in each time slot. Therefore, the first algorithm aims at finding an assignment that maximizes the number of cameras that can transmit simultaneously.

It is worth noting that different from the classical maximum link activation problem [14], where only the SINR condition is considered, to support multi-view processing, a feasible assignment solution must satisfy (2). In view of this, we construct the weighted bipartite graph $\mathcal{G} = (\mathcal{S}, \mathcal{N}, \mathcal{E})$. To create the edge set $\mathcal{E}$ of the bipartite graph, we first check the "feasibility" of connecting each s and n in the graph. To do so, for each camera-node pair we calculate the SNR value (i.e., assuming TDMA), and we add an edge connecting s and n if for all $c \in \mathcal{c}(s)$ the SNR value exceeds its respective SINR threshold γ_c. Next, in order to construct an assignment solution that maximizes the number of cameras transmitting simultaneously, we define the weight $w_{sn} = \prod_{c \in \mathcal{c}(s)} \min(1, \text{SINR}_n(c, \mathcal{C})/\gamma_c)$ for each edge $(s, n) \in \mathcal{E}$. Intuitively, a higher value of w_{sn} implies that more cameras can be activated simultaneously. Based on the weighted bipartite graph, we construct an assignment by selecting the node $n = \text{argmax}_n \{w_{sn}, n \in N\}$ as the serving node of cameras $c \in \mathcal{c}(s)$. Ties are broken by selecting the node with smaller index. We refer to this algorithm as the *SINR-based assignment* algorithm.

The second algorithm is based on the observation that in the age calculation in (4), or equivalently, in (5b), small time stamps result in large peak ages. Thus, it is intuitively preferable to schedule "old" images as soon as possible, and the assignment should facilitate doing so. Consequently, we construct an assignment such that cameras containing images with small time stamps can transmit together. For a given MVAM instance, we first construct the weighted bipartite graph $\mathcal{G} = (\mathcal{S}, \mathcal{N}, \mathcal{E})$, as before. For each $s \in \mathcal{S}$, we define the weight $w'_s = \frac{\tau}{\tau_{s1}}$, where $\tau = \min_{s \in \mathcal{S}} \tau_{s1}$. Thus, the weight is the scaled reciprocal of the minimum time stamp of all images taken for s. Finally, we update the weight of each edge $(s, n) \in \mathcal{E}$ to $w_{sn} = w'_s w_{sn}$. Given $\mathcal{G}$, we construct the assignment by selecting the maximum weight links, as in the case of *SINR-based assignment*. We refer to this as the *age-aware assignment* algorithm.

B. Transmission Scheduling Algorithm

Motivated by the structural results in Section IV, to deliver the images in a timely fashion, the cameras with old images should be scheduled first. We thus propose a greedy strategy for the transmission schedule, which in each time slot chooses a camera group such that the oldest image is delivered together with as many other images as possible.

The algorithm works as follows. In each time slot, the camera group is initially empty. The algorithm sorts the cameras in ascending order of the time stamps of the images at the head of their FCFS queues, and adds the camera with the lowest time stamp to the camera group. It then iterates through the ordered list of cameras, and adds one camera at a time. In each step, denoted by $\mathcal{g}'$ the camera group with the new added camera c'. If $\text{SINR}_n(c, \mathcal{g}') \geq \gamma_c$, $\forall c \in \mathcal{g}'$, then the camera c' is kept; otherwise, c' is removed from the group. The algorithm schedules the computed camera group for transmission, after which it continues with the next time step, until all queues are empty.

C. CMAF Algorithm and its Optimality

The proposed *CMAF* algorithm uses the *SINR-based assignment* and *age-aware assignment* algorithms for computing two camera-node assignments. For both assignments, it executes the greedy scheduling algorithm described in Section V-B, and calculates the obtained maximum peak ages. The algorithm then chooses the camera-node assignment that results in lower maximum peak age of information.

We further verify the rationale of the *CMAF* algorithm by applying it to the tractable cases identified in Section IV.

Lemma 8. *The* CMAF *algorithm achieves the global optimum for all MVAM instances defined in Theorems 2 and 3, and in Corollary 6.*

Proof: For the three tractable cases, i.e., MVAM with compatible $\mathcal{C}$, MVAM with TDMA, and MVAM with compatible $\mathcal{c}(s)$ only, by construction, both *SINR-based assignment* and *age-aware assignment* algorithms provide an optimal camera-node assignment. Together with the *transmission scheduling algorithm*, the CMAF gives the same result as the one we derived in the respective proof of the three theorems. Hence the conclusion. ■

VI. Numerical Study

A. Network Setting

We consider a camera network monitoring an area of 100×100 meters. The network area is divided into 16 sub-areas, each occupying 25×25 meters and consisting of one scene. The number of cameras that cover one scene is uniformly chosen on $[2, 6]$. For each scene s, the cameras $\mathcal{c}(s)$ are uniformly distributed in the respective sub-area. In

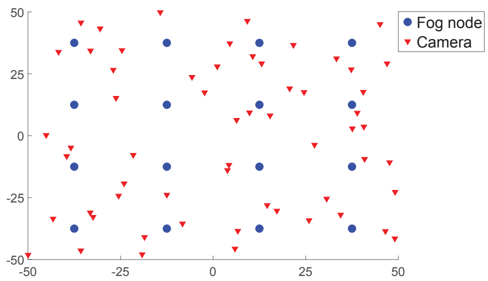

Fig. 2. A sample network with $N = 16$ nodes and $C = 59$ cameras.

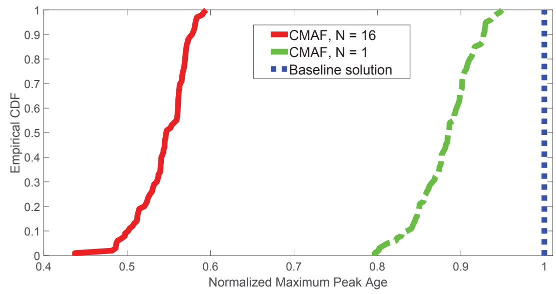

Fig. 3. Empirical CDFs of maximum peak ages for networks with $N = 1,\ 16$.

the network, we deploy N fog nodes, where $N \in \{1, 16\}$, representing a centralized and a distributed network architecture, respectively. The fog node(s) are located in the geometric center after splitting the network area into N sub-areas. In Figure 2, we show the topology of a sample network with $N = 16$ nodes and $C = 59$ cameras.

The transmit power of the cameras and the noise variance at the fog nodes are uniformly set to 20 dBm and to -100 dBm, respectively. The channel gain follows a distance-based propagation model with a path loss exponent of 4, Rayleigh fading, and log-normal shadowing. The SINR threshold $\gamma_c = -3$ dB. The starting time is $t_0 = 500$. The initial ages a_{s0}, $\forall s \in \mathcal{S}$, are uniformly distributed in $[50, 200]$. Each camera has up to 10 images to be delivered. The time stamps of images capturing scene s are random integers uniformly distributed in $(t_0 - a_{s0}, t_0)$. For each setup, 100 instances are generated.

B. Simulation Results

We apply the proposed *CMAF* algorithm to solve the MVAM instances. For performance comparison, we define the baseline solution as the one derived from the centralized network, i.e., $N = 1$, and following a greedy method of minimum time scheduling [15], i.e., in each time slot, we select the camera with maximal number of images left in queue and pair it with other cameras that it can transmit together with. All results are normalized by the maximum peak age achieved by the corresponding baseline solution. In Figure 3, we present the normalized maximum peak age in form of empirical cumulative distribution functions (CDFs). The numerical results show that *CMAF* obtains a significant peak age reduction; for $N = 1$ it achieves an average improvement of 12% over the baseline solution, and for $N = 16$ the average improvement is 45%. Overall, the results show that the fog architecture and the proposed optimized assignment/transmission strategy result in a synergy that significantly improves the freshness of information.

VII. Conclusions

We have considered the joint optimization of serving node assignment and camera transmission scheduling with respect to age of information in wireless camera networks with fog computing. We have extended the age calculation in the presence of multi-view processing and mathematically formulated the multi-view age minimization problem. Fundamental results including problem complexity, tractable cases, and optimality condition have been derived. An optimization algorithm based on a modular structure has been proposed to solve the problem in polynomial time. Our numerical results show that the optimal assignment and transmission strategy reduces the maximum peak age significantly compared to the traditional centralized approach. Our work has a number of interesting potential extensions, including balancing the number of cameras being served by one fog node and optimizing the length of a scheduling cycle.

References

[1] K. Abas, C. Porto, and K. Obraczka, "Wireless smart camera networks for the surveillance of public spaces," *IEEE Computer*, vol. 47, no. 25, pp. 37–44, 2014.

[2] B. Song et al., "Tracking and activity recognition through consensus in distributed camera networks," *IEEE Trans. Image Process.*, vol. 19, no. 10, pp. 2564–2579, 2010.

[3] J. Ko et al., "Wireless sensor networks for healthcare," *Proc. IEEE*, vol. 98, no. 11, pp. 1947–1960, 2010.

[4] E. Eriksson, G. Dán, and V. Fodor, "Predictive distributed visual analysis for video in wireless sensor networks," *IEEE Trans. Mobile Comput.*, vol. 15, no. 7, pp. 1743–1756, 2016.

[5] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. of IEEE INFOCOM*, 2012, pp. 2731–2735.

[6] A. Kosta, N. Pappas, and V. Angelakis, "Age of information: A new concept, metric, and tool," in *Foundations and Trends in Networking*. Now Publishers, 2017.

[7] R. D. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *Proc. of IEEE ISIT*, 2012, pp. 2666–2670.

[8] M. Costay, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *Proc. of IEEE ISIT*, 2014, pp. 1583–1587.

[9] Q. He, D. Yuan, and A. Ephremides, "Optimal link scheduling that minimizes the age of information in wireless systems," *IEEE Trans. Inf. Theory*, 2017, DOI: 10.1109/TIT.2017.2746751.

[10] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. of IEEE ISIT*, 2013, pp. 66–70.

[11] Y. Hsu, E. Modiano, and L. Duan, "Age of information: Design and analysis of optimal scheduling algorithms," in *Proc. of IEEE ISIT*, 2017, pp. 561–565.

[12] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Trans. Inf. Theory*, pp. 7492–7508, 2017.

[13] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[14] A. Capone, I. Filippini, S. Gualandi, and D. Yuan, "Resource optimization in multi-radio multi-channel wireless mesh networks," in *Mobile Ad Hoc Networking: Cutting Edge Directions*, 2nd ed., S. Basagni, M. Conti, S. Giordano, and I. Stojmenovic, Eds. Wiley and IEEE Press, 2013, pp. 241–274.

[15] V. Angelakis, A. Ephremides, Q. He, and D. Yuan, "Minimum-time link scheduling for emptying backlogged traffic in wireless systems: solution characterization and algorithmic framework," *IEEE Trans. Inf. Theory*, vol. 60, no. 2, pp. 1083–1100, 2014.