Version Age in Multi-Process Gossip Networks under Distributed Sensing

Anuj Meena
IIT Kharagpur, India
anujmeena225@gmail.com

Amitalok J. Budkuley
IIT Kharagpur, India
amitalok@ece.iitkgp.ac.in

***Abstract*—We study the problem of maintaining version freshness of information, measured as version age of information (VAoI), for multiple *concurrent* stochastic processes jointly observed and communicated over a network of heterogeneous nodes in a gossip-like manner. Unlike prior work that assumes a single process with exogenous version-generation rates, we model each process evolution as a continuous-time Markov chain (CTMC) and the sensing activity of hybrid nodes (with sensing and communication capabilities) as finite-capacity Poisson processes. This formulation captures the asynchronous and probabilistic nature of real-world monitoring systems, where sensing, version generation, and communication operate at distinct timescales.**

We derive the effective version-generation rate at each sensor as a function of the underlying process dynamics via CTMC steady-state analysis. Building on this, we extend the stochastic hybrid systems (SHS) framework to jointly model multiple processes across multi-sensor networks, obtaining steady-state VAoI expressions for arbitrary node subsets. We further consider networks with queuing constraints on a single link to a remote monitor, deriving VAoI using an M/M/1 model (with possible finite buffer) that incorporates service delays, scheduling policies, and packet drops. Numerical results validate our analysis. Our work generalizes prior single-process VAoI models and provides a unified foundation for information freshness in hybrid sensing–communication networks.

I. Introduction

Motivation: Consider multiple software services evolving independently while a remote monitor seeks to maintain up-to-date copies. Updates are discovered through distributed sensors and propagated through a communication network. Because sensing opportunities, process changes, and transmissions occur asynchronously, the monitor may lag behind by different numbers of versions across services. Version Age of Information (VAoI), first formulated in [1] (see also, [2]), captures this version-based staleness by measuring the number of versions by which the remote monitor lags behind each application (relative to the latest state of each application) rather than merely the time since the last update.

A case for Multi-process, Multi-sensor VAoI: Prior VAoI studies, including Yates [1], *et al.* [3], and Khojastepour *et al.* [4], model version-generation as exogenous processes with prescribed rates. The former considers a single-sensor setting, while the latter two extend to multi-sensor systems with in-network exchange of version packets. However, they do not model how versions arise from the evolution of the underlying monitored processes.

These exogenous assumptions mask key challenges: (i) stochastic process evolution at timescales distinct from sensing, (ii) correlated version generation due to multiple sensors observing the same process, and (iii) competition among multiple process flows for shared communication resources. Addressing these requires joint modeling of process evolution, sensing, and dissemination in multi-process, multi-sensor systems with hybrid sensing–communication nodes and shared resources, which is the focus of this work.

Related work under AoI: The VAoI model naturally extends the broader body of research on *age of information (AoI)*, a notion first formalized in [5] (see also [6]), which measures the timeliness of updates (rather than just their version count as in VAoI). Numerous generalizations of AoI have been studied, including (but not limited to) peak age [7], age penalties [8], incorrect age [9], and, more recently, semantically-inspired metrics [10]. These works have broadly explored the role of sampling policies, queueing disciplines, and resource sharing in maintaining information freshness across general networks. Several ingredients central to our study, such as modeling processes as continuous-time Markov chains (CTMCs) [5], [11], accounting for multi-source sampling correlations [12], and analyzing gossip-based information exchange [1], are classic and have appeared before in these lines of work. To our knowledge, however, these ingredients have not been unified within a single VAoI framework prior to this work.

Our work: Unlike prior VAoI formulations that assume exogenous version-generation processes, we derive version creation endogenously from sampled CTMC dynamics, thereby coupling process evolution, sensing, dissemination, and queueing within one analytical model. Each process is modeled as a CTMC, and version generation is induced by finite-rate sensing when state transitions are detected. The resulting updates are disseminated through gossip and delivered via a queueing interface, yielding our unified model (subsuming single-process VAoI works as special cases).

Contributions: We summarize the key contributions below:

- we develop in Sec. II a process-driven multi-process VAoI framework in which versions emerge from CTMC process evolution and stochastic sensing rather than being prescribed exogenously.
- using an SHS formulation, we characterize steady-state

The work of A. Meena and A. J. Budkuley was partially supported by Dept. of Science and Technology, Govt. of India, through grant MTR/2023/001412.

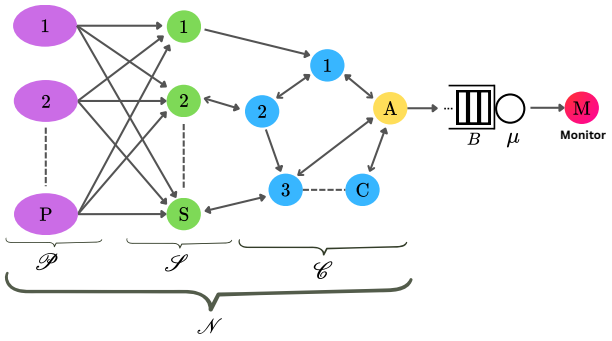

Fig. 1: Remote monitoring architecture with multiple processes, sensors, and communication nodes forwarding versioned updates to a monitor.

VAoI for arbitrary node subsets in gossip networks Thm. 1, and extend the result to queue-constrained remote monitoring in Thm. 2.

Paper organization: The remainder of the paper is organized as follows. In Sec. II, we describe the system model and formulate our problem formulation. We then present our key analytical results in Sec. III-A, and simulations results in Sec. III-B. We present proofs of our main results in Sec. IV. Finally, we make concluding remarks in Sec. V.

Notation: Let $\mathbb{R}$ and $\mathbb{N}$ denote the sets of real numbers and natural numbers, respectively, and for $n \in \mathbb{N}$ define $[n] := \{1, 2, \ldots, n\}$. Random variables are denoted by uppercase letters (e.g., X), realizations by lowercase letters (e.g., x), and random vectors by boldface (e.g., $\mathbf{X} = (X_1, \ldots, X_n)$). Unless stated otherwise, sets are represented by calligraphic letters (e.g., $\mathcal{A}$, $\mathcal{B}$). In some cases, script-style symbols (e.g., $\mathscr{P}$) are used to denote key fixed sets, while upright Roman fonts (e.g., P) are used to denote fixed constants. Sensing rates and communication network link rates are denoted by $\lambda \in \mathbb{R}_+$. Matrices are denoted by uppercase letter. However, the sensing probability matrix is $\Pi \in [0,1]^{n \times n}$, with entries $\pi_{i,j}$. Version age is denoted by Δ for a single process and $\boldsymbol{\Delta}$ as a vector over multiple processes. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a directed graph with vertex set $\mathcal{V}$ and edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$.

II. System Model and Problem Formulation

We consider a continuous-time remote monitoring system comprising multiple time-evolving processes observed by distributed sensors and relayed through a communication network to a remote monitor, as illustrated in Fig. 1.

Processes. Let P denote the number of processes, indexed by $\mathscr{P} = [\mathrm{P}]$. Each process $p \in \mathscr{P}$ evolves independently as a continuous-time Markov chain (CTMC), which is assumed to be irreducible and positive recurrent, with state $Z_p(t) \in \mathcal{Z}_p, \quad t \geq 0$, and with generator matrix $\Sigma_p = [\Sigma_p(i,j)]_{i,j \in \mathcal{Z}_p}$ and steady-state distribution $\boldsymbol{\sigma}_p$ satisfying $\boldsymbol{\sigma}_p^\top \Sigma_p = 0$ and $\mathbf{1}^\top \boldsymbol{\sigma}_p = 1$. We denote the collection of all process parameters as $\{\boldsymbol{\sigma}_p, \Sigma_p\}_{p \in \mathscr{P}}$. These processes generate distinct streams that are subsequently sensed by sensors. The sensors primarily detect *changes of state*, which are then propagated through the network using appropriately defined *versioned* packets. No version increments are made in created packets if CTMC state is unchanged.

Sensors. The system has S sensors indexed by $\mathscr{S} = [\mathrm{S}]$, each with a total sensing capacity $\overline{\lambda}_s > 0$. Sensor s initiates sensing opportunities according to a Poisson process of rate $\overline{\lambda}_s$. At each opportunity, it selects, *randomly* and *independently*, process p to sense according to a fixed sensing probability matrix $\Pi = [\pi_{p,s}], \quad \sum_{p \in \mathscr{P}} \pi_{p,s} = 1, \ \forall s \in \mathscr{S}$, and hence, fixed *effective sensing rate* of process p, $\overline{\lambda}_{p,s} = \pi_{p,s} \overline{\lambda}_s$. We retain Π for modeling flexibility and assume uniform sensing allocation unless stated otherwise, where $\pi_{p,s} = 1/\mathrm{P}$, giving $\overline{\lambda}_{p,s} = \overline{\lambda}_s/\mathrm{P}$. Events occur on independent and continuous-time Poisson clocks, so simultaneous sensing events occur with probability zero (almost surely).

Communication Network. Sensed updates propagate through a communication network modeled as a directed graph $\mathscr{G}_c = (\mathscr{C}, \mathscr{E}_c)$ of C relay nodes in $\mathscr{C} = [\mathrm{C}]$ with link rates $\overline{\lambda}_{j,k} > 0$ for each $(j,k) \in \mathscr{E}_c$. The *aggregator* node $A \in \mathscr{C}$ collects updates and forwards them to the *monitor* M through an M/M/1 queue of service rate μ. The monitor is *not considered part* of the in-network node set, as its queue-based reception model introduces heterogeneous latency.

Augmented Network. We define the augmented node set $\mathscr{N} := \mathscr{P} \cup \mathscr{S} \cup \mathscr{C}$, and the augmented directed graph $\mathscr{G}_{\texttt{AUG}} = (\mathscr{N}, \mathcal{E}_{\texttt{AUG}})$ that captures both sensing and communication links. For any node $k \in \mathscr{N}$, denote its in- and out-neighborhoods by $\mathcal{N}_k^{\downarrow} := \{j : (j,k) \in \mathcal{E}_{\texttt{AUG}}\}$ and $\mathcal{N}_k^{\uparrow} := \{j : (k,j) \in \mathcal{E}_{\texttt{AUG}}\}$ resp..

Version Generation and Propagation. Each sensor $s \in \mathscr{S}$ generates a new version of process $p \in \mathscr{P}$ only when it detects a state change in the underlying CTMC $Z_p(t)$. The effective version generation rate is $\lambda_{p,s}^{(p)} = f(\overline{\lambda}_{p,s}, \boldsymbol{\sigma}_p, \Sigma_p)$, where the explicit form of $f(\cdot)$ is derived later in Lemma 1. Note that $\lambda_{p,s}^{(p)} \leq \overline{\lambda}_{p,s}^{(p)}$ (under any given (σ_p, Σ_p) and any Π).

Once generated at the sensors, version information spreads through the network via *decentralized gossip interactions*. Versions at any node j about process p at any time t, denoted as $N_j^{(p)}(t)$, propagate via *gossip interactions*: for each directed link $(i,j) \in \mathcal{E}_{\texttt{AUG}}$, node j updates as

$$N_j^{(p)}(t^+) = \max\{N_j^{(p)}(t), N_i^{(p)}(t)\}, \quad \forall p \in \mathscr{P}, \qquad (1)$$

where t^+ is the time instant immediately *after* t, and triggered by Poisson events with rate $\lambda_{i,j}^{(p)} = \frac{\overline{\lambda}_{i,j}}{\mathrm{P}}$. Uniform activation ensures that each process experiences identical link statistics (the extension to non-uniform activation is straightforward).

Remark: Note that the gossip rule (1) specifies *what happens when an interaction occurs*, but not *how often* such events take place. The latter is specified through the link rate, given by $\overline{\lambda}_{i,j}^{(p)}$ for process $p \in \mathscr{P}$, which sets the stochastic event rate for that link in the augmented graph.

Version Age. For each process $p \in \mathscr{P}$, define the freshest version index among sensors as $N_0^{(p)}(t) := \max_{s \in \mathscr{S}} N_s^{(p)}(t)$. *Remark:* As is standard in VAoI (and related age metrics), we assume that the global latest version $N_0^{(p)}(t)$ of each process

p acts as a benchmark, to help make calculation of version age at nodes meaningful.

We can now define, for any node $k \in \mathscr{N} \setminus \mathscr{P}$, the version age of process p: $\Delta_k^{(p)}(t) := N_0^{(p)}(t) - N_k^{(p)}(t)$, and hence, for a subset $\mathcal{S} \subseteq \mathscr{N} \setminus \mathscr{P}$, $\Delta_{\mathcal{S}}^{(p)}(t) := \min_{j\in\mathcal{S}} \Delta_j^{(p)}(t)$. From the update rule, it follows that $\Delta_{S\cup\{j\}} := \min(\Delta_S, \Delta_j)$, for $j \in \mathscr{N} \setminus \mathscr{P}$, highlighting how gossip interactions merge information freshness across nodes.

Remark: For process nodes $p \in \mathscr{P}$, no version index is maintained since they neither sense nor receive updates. Accordingly, $\Delta_p^{(p)}(t) = 0$, as p is the freshest source of its own version, while $\Delta_p^{(p')}(t) = \infty$ for any $p' \neq p$, since p cannot track other processes. Therefore, we restrict our VAoI analysis to nodes (and subsets) in $\mathscr{N} \setminus \mathscr{P}$.

Steady-State Version Age. The long-term *average* version age for process p at node $k \in \mathscr{N} \setminus \mathscr{P}$ is

$$\overline{\Delta}_k^{(p)} := \limsup_{t\to\infty} \mathbb{E}[\Delta_k^{(p)}(t)], \tag{2}$$

The version age $\overline{\Delta}_{\mathcal{S}}^{(p)}$ (also, its vector counterpart $\overline{\boldsymbol{\Delta}}_{\mathcal{S}}$) and $\overline{\Delta}_{\mathcal{M}}^{(p)}$ are similarly defined for subsets $\mathcal{S} \subseteq \mathscr{N} \setminus \mathscr{P}$ and the monitor M respectively. Note that these depend implicitly on $(\boldsymbol{\sigma}_p, \Sigma_p, \Pi, \overline{\lambda}_{p,s})$.

Goal. We seek to characterize the limiting average version age $\overline{\Delta}_{\mathcal{S}}^{(p)}$ (and its vector counterpart $\overline{\boldsymbol{\Delta}}_{\mathcal{S}}$) and $\overline{\Delta}_{\mathcal{M}}^{(p)}$, $\forall p \in \mathscr{P}$.

III. Our Main Results

A. Our analytical contributions

a) Effective version generation.: Before analyzing VAoI, we characterize the effective version-generation rate at sensors.

Lemma 1 (Effective version-generation rate)**.** *For a process $p \in \mathcal{P}$ evolving as a CTMC with generator Σ_p and stationary distribution $\boldsymbol{\sigma}_p$, and for a sensor $s \in \mathcal{S}$ with induced sensing rate $\bar{\lambda}_{p,s}$, the effective version-generation rate by s is*

$$\lambda_{p,s}^{(p)} = \bar{\lambda}_{p,s}\Big(1 - \bar{\lambda}_{p,s}\, \boldsymbol{\sigma}_p^{\top} \operatorname{diag}\big(R_p(\bar{\lambda}_{p,s})\big)\Big),$$

where $R_p(\bar{\lambda}_{p,s}) := (\bar{\lambda}_{p,s} I - \Sigma_p)^{-1}$ is the resolvent of Σ_p. The aggregate effective rate over sensors $\mathscr{S}$ is $\lambda_{\mathscr{S}}^{(p)} = \sum_{s\in\mathscr{S}} \lambda_{p,s}^{(p)}$.

The proof (deferred to [13]) employs techniques for analysing *sampled* CTMCs [14]. Here, $\lambda_{p,s}^{(p)}$ shows the interaction between sensing and process dynamics: fast sensing often yields redundant observations, while slower sensing is more likely to capture state changes for new versions.

b) Steady-state subset average VAoI: Using rates $\lambda_{p,s}^{(p)}$ from Lemma 1, the steady-state average VAoI for any subset $\mathcal{S} \subseteq \mathscr{N} \setminus \mathscr{P}$ follows a recursive structure.

Theorem 1. *Let $p \in \mathscr{P}$, and $\mathcal{S} \subseteq \mathscr{N} \setminus \mathscr{P}$. Then, the steady-state version age $\Delta_{\mathcal{S}}^{(p)}$ for process p satisfies*

$$\overline{\Delta}_{\mathcal{S}}^{(p)} = \frac{\sum_{j\in\mathscr{S}\setminus\mathcal{S}} \lambda_j^{(p)} + \sum_{\substack{j\in\mathscr{N}\setminus\mathcal{P} \\ k\in\mathcal{S}}} \lambda_{j,k}^{(p)}\, \overline{\Delta}_{\mathcal{S}\cup\{j\}}^{(p)}}{\sum_{k\in\mathcal{S}\cap\mathscr{S}} \lambda_k^{(p)} + \sum_{\substack{j\in\mathscr{N}\setminus\mathcal{P} \\ k\in\mathcal{S}}} \lambda_{j,k}^{(p)}}. \tag{3}$$

where $\lambda_k^{(p)} \leq \overline{\lambda}_{p,k}$, is the k-th sensor effective version generation rate.

The proof details appear in Sec. IV later. The above expression reveals how the subset VAoI depends on (i) local update rates $\lambda_j^{(p)}$, (ii) cross-node communication rates $\lambda_{j,k}^{(p)}$, and (iii) the age at the extended subset $\mathcal{S} \cup \{j\}$, recursively capturing network information flow. It can be easily verified that the results for the single-process settings in [1], and its extension [4] to multi-node sensing, follow as direct corollaries to our above result.

c) Steady-state average VAoI at monitor.: For the remote monitor node M receiving updates from an *aggregator* $A \in \mathscr{C}$ (equivalently, its network $\mathscr{C}$ *gateway*) over an $M/M/1/B$ queue with service rate $\mu > 0$ under a a first-come-first-served (FCFS) service policy, and utilization $\rho := \lambda_{A,M}^{(p)}/\mu$, the subset limiting version age is characterized next.

Theorem 2. *Given $p \in \mathscr{P}$, and under the aggregator–monitor link modeled as an $M/M/1/B$ queue with FCFS service, the monitor steady-state version age is*

$$\overline{\Delta}_{\mathcal{M}}^{(p)} = \frac{\sum_{j\in\mathscr{S}} \lambda_j^{(p)} + \lambda_{A,M}^{(p)'} \Delta_{M\cup\{A\}}^{(p)}}{\lambda_{A,M}^{(p)'}},$$

where the effective aggregator–monitor rate is

$$\lambda_{A,M}^{(p)'} := \begin{cases} \lambda_{A,M}^{(p)} \left[1 - \frac{(1-\rho)\rho^B}{1-\rho^{B+1}}\right], & \rho < 1, \\ \lambda_{A,M}^{(p)} \left(1 - \frac{1}{B+1}\right), & \rho \nearrow 1. \end{cases}$$

We omit the proof due to space constraints (see [13] for details). The result mirrors the recursive structure of Theorem 1, with the queue replacing the nominal aggregator update process by an effective rate $\lambda_{A,M}^{(p)'}$. This effective rate incorporates the probability that an update successfully traverses the $M/M/1/B$ queue and reaches the monitor. Consequently, monitor freshness depends jointly on network dissemination within the gossip graph and queue-induced congestion effects at the final delivery stage.

Seen together, Theorems 1 and 2 reveal the interplay of sensing, communication, and queuing in determining VAoI. Note that such an interplay is often found in other notions of AoI (cf. [10], [11], for instance), but has not been explored in the context of VAoI. This is the gap our work seeks to fill for a general (CTMC) multi-process setting when distributed sensing is employed. As expected, increasing connectivity or sensing rate reduces version age, while queuing and high utilization ($\rho \approx 1$) significantly deteriorate freshness.

B. Our numerical results

a) Setup and topologies: Two topologies are considered, namely, (i) a ring network (RN), and (ii) a fully connected network (FCN), each with identical node count (Fig. 2). We model two processes, each as 3-state CTMCs with different

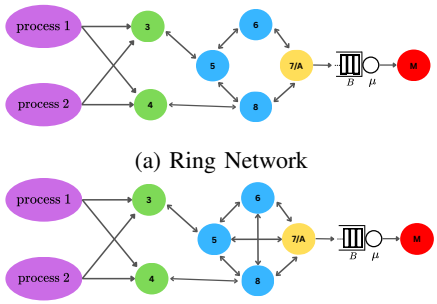

(a) Ring Network

(b) Fully Connected Network

Fig. 2: The two different topologies, namely, ring network (RN) and fully connected network (FCN), for our VAoI study.

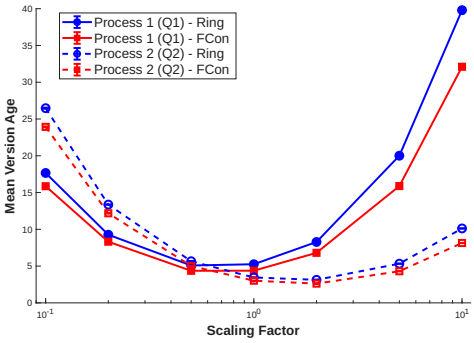

Fig. 3: VAoI study under CTMC transition rate scaling.

transition rates specified by (infinitesimal) generator matrices:

$$Q_1 = Q_2 := \begin{bmatrix} -0.4 & 0.2 & 0.2 \\ 0.3 & -0.6 & 0.3 \\ 0.4 & 0.3 & -0.7 \end{bmatrix}$$

We use identical baseline generators to isolate the effect of asymmetric scaling. We model Poisson sensing at nodes, with identical process-sensing rate $\bar{\lambda}_s = 0.5$ at each sensor, as well as for inter-node communication rate $\bar{\lambda}_{i,j} = 0.5$, for distinct pairs of non-process nodes i, j. The sensing matrix $\Pi = \begin{bmatrix} 0.8 & 0.2 \\ 0.4 & 0.6 \end{bmatrix}$, has a preference toward sensing Process 1.

b) Effect of process dynamics: Here we scale the CTMC generator matrices asymmetrically as $Q_1' = \gamma Q_1$ and $Q_2' = \frac{1}{\gamma} Q_2$, with $\gamma \in [0.1, 10]$. As shown in Fig. 3, increasing γ leads to higher VAoI, since faster process dynamics generate updates more rapidly than the network can propagate via versions. The VAoI is minimized and remains stable when γ lies in the range $[0.5, 1]$, where update generation rates are well-matched to network link capacities. Across all settings, as expected, FCN achieves lower VAoI due to its higher connectivity.

c) Effect of sensing correlation: To examine the impact of sensing correlation on VAoI, we vary the sensing probabilities $\alpha := \pi_{11}$ and $\beta := \pi_{21}$ in $[0, 1]$, biasing updates toward Process 1. As shown in Fig. 4a, where we plot the *ratio* of process VAoIs, increasing (α, β) increases the VAoI ratio, demonstrating the effect of such strong bias on version update accumulation at the sensors and congestion

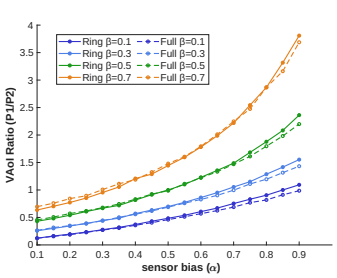

(a) Relative VAoI versus sensing bias (α, β).

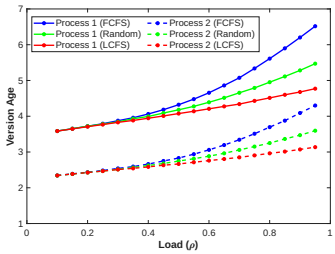

(b) VAoI under different queueing policies as a function of utilization ρ.

Fig. 4: VAoI study under sensing and queueing .

in the network, particularly for the favored process. This effect is more pronounced in the RN topology, which exhibits a steeper increase in VAoI than the FCN topology due to limited connectivity. Results show that aggressive sensing of a single process may prove counter-productive, necessitating balancing of sensing rates with network capacity.

d) Effect of queueing: congestion and scheduling policy: Theorem 2 shows that the monitor VAoI under FCFS is determined by the effective aggregator-to-monitor update rate, which captures queueing effects such as service delays and packet drops. To assess the impact of queue management, we compare FCFS with Random service and LCFS using simulations.

We study an $M/M/1/B$ queue ($B = 5$) before the monitor under FCFS, Random, and LCFS scheduling (Fig. 4b). At low utilization ρ, VAoI remains nearly constant since queueing delays are negligible. As ρ increases, congestion introduces delays and packet drops, causing VAoI to rise. FCFS exhibits the sharpest increase due to the accumulation of stale packets, while Random service moderates this effect by occasionally transmitting fresher updates. LCFS achieves the best performance. Although VAoI increases under all policies at high utilization, its growth under LCFS is substantially slower than under FCFS and Random.

IV. VAoI Analysis: Proof of Theorem 1

Fix a process $p \in \mathscr{P}$, and model the corresponding VAoI dynamics as a stochastic hybrid system (SHS). For reference, a generic SHS is $\tilde{\mathcal{H}} = (\tilde{\mathcal{Q}}, \tilde{\mathcal{X}}, \tilde{t}, \tilde{\mathcal{L}}, \{\tilde{\phi}\}, \{\tilde{\lambda}\})$, where $\tilde{\mathcal{Q}}$ and $\tilde{\mathcal{X}}$ comprise the discrete and continuous state resp. of SHS, evolving with independent variable $\tilde{t}$. Furthermore, $\tilde{\mathcal{L}}$ is the transition set where $\{\tilde{\phi}\}$ and $\{\tilde{\lambda}\}$ are the collection of reset maps and transition rates corresponding to entries in $\tilde{\mathcal{L}}$. See [15] for details. The proof has three parts.

a) Construction of the VAoI-specific SHS $\mathcal{H}_p$.: For process p, define $\mathcal{H}_p = (\mathcal{Q}, \mathcal{X}, t, \mathcal{L}, \phi, \lambda)$, where (i) $\mathcal{Q} = \{0\}$ trivially for static network topology, (ii) $\mathbf{X}(t) = [\Delta_k^{(p)}(t)]_{k \in \mathscr{N}_{\text{age}}}^{\top} \in \mathcal{X} = \mathbb{N}_0^N$, with $\mathscr{N}_{\text{age}} := \mathscr{N} \setminus \mathscr{P}$, and $N = |\mathscr{N}_{\text{age}}|$. Next, the transition set $\mathcal{L} := \mathcal{L}_{\text{gossip}} \cup \mathcal{L}_{\text{sense-0-reset}} \cup \mathcal{L}_{\text{sense-1-incr}}$ has three parts comprising gossip events, node-reset events

and node increment events (where the first is across pairs of nodes, while the latter two are at specific nodes) given as

$$\begin{aligned}\mathcal{L}_{\text{gossip}} &= \{\ell_{j,k} : j,k \in \mathscr{N}_{\text{age}}, j \neq k\}\\ \mathcal{L}_{\text{sense-0-reset}} &= \{\ell_s^0 : s \in \mathscr{S}\}\\ \mathcal{L}_{\text{sense-1-incr}} &= \{\ell_k^1(s) : s \in \mathscr{S}, k \in \mathscr{N}_{\text{age}} \setminus \{s\}\}.\end{aligned}$$

Next, the collection of reset maps and event rates for each of the entries in $\mathcal{L}$ are given. Reset maps ϕ_ℓ, for each $\ell \in \mathcal{L}$ are:

(i) for $\ell_{j,k} \in \mathcal{L}_{\text{gossip}}$,

$$X_k \leftarrow \min(X_k, X_j), \quad X_i \text{ unchanged for } i \neq k$$

(ii) for $\ell_s^0 \in \mathcal{L}_{\text{sense-0-reset}}$

$$X_s \leftarrow 0, \quad X_i \text{ unchanged for } i \neq s$$

(iii) for $\ell_k^1(s) \in \mathcal{L}_{\text{sense-1-incr}}$

$$X_k \leftarrow X_k + 1, \quad X_i \text{ unchanged for } i \neq k$$

The corresponding event rates λ_ℓ, for each $\ell \in \mathcal{L}$, are

$$\lambda_\ell = \begin{cases}\lambda_{j,k}^{(p)}, & \ell = \ell_{j,k} \in \mathcal{L}_{\text{gossip}},\\ \lambda_{p,s}^{(p)}, & \ell = \ell_s^0 \in \mathcal{L}_{\text{sense-0-reset}},\\ \lambda_{p,s}^{(p)}, & \ell = \ell_k^1(s) \in \mathcal{L}_{\text{sense-1-incr}}.\end{cases}$$

This gives the following correspondence with any SHS $\tilde{\mathcal{H}}$:

$$\mathbf{X}(t) \leftrightarrow \tilde{\mathcal{X}}, \quad q(t) \leftrightarrow \tilde{\mathcal{Q}}, \quad \mathcal{L} \leftrightarrow \tilde{\mathcal{L}}, \quad \phi_\ell \leftrightarrow \tilde{\phi}_\ell, \quad \lambda_\ell \leftrightarrow \tilde{\lambda}_\ell.$$

b) Test function and subset age.: Let $\mathcal{S} \subseteq \mathscr{N} \setminus \mathscr{P}$, and define the test function (cf. [15]) as follows: $\psi(\mathbf{X}(t)) := \Delta_{\mathcal{S}}^{(p)}(t)$.Recall that our goal is to characterize limiting average subset version age in (2). As $\mathcal{H}_p$ is an SHS, the classic Dynkin's formula [15] yields

$$\frac{d}{dt}\,\mathbb{E}[\psi(\mathbf{X}(t))] = \mathbb{E}\big[(L\psi)(\mathbf{X}(t))\big], \tag{4}$$

where the *extended generator* $L\psi$ (characterized in [15, Thm. 1]) simplifies considerably for our SHS $\mathcal{H}_p$ as follows:

$$(L\psi)(\mathbf{X}) = \sum_{\ell \in \mathcal{L}} \lambda_\ell\big(\psi(\phi_\ell(\mathbf{X})) - \psi(\mathbf{X})\big). \tag{5}$$

See [13] for details. Note that at steady state, when the limit exists, we have $\frac{d}{dt}\,\mathbb{E}[\psi(\mathbf{X}(t))] = 0$ which yields the fixed-point condition (cf. (4))

$$0 = \mathbb{E}\big[(L\psi)(\mathbf{X})\big]. \tag{6}$$

c) Solving the equation in (6)*:* We first evaluate $(L\psi)(\mathbf{X})$ in (5) (by decomposing over three subsets of $\mathcal{L}$).

Analyzing each of the three terms, we first obtain

$$\begin{aligned}(L\psi)(\mathbf{X}) = &\sum_{s \in \mathscr{S}\setminus\mathcal{S}} \lambda_{p,s}^{(p)} - \sum_{s \in \mathcal{S}\cap\mathscr{S}} \lambda_{p,s}^{(p)} \Delta_{\mathcal{S}}^{(p)}\\ &+ \sum_{\substack{j \in \mathscr{N}\setminus\mathscr{P}\\ k \in \mathcal{S}}} \lambda_{j,k}^{(p)}\big(\Delta_{\mathcal{S}\cup\{j\}}^{(p)} - \Delta_{\mathcal{S}}^{(p)}\big).\end{aligned} \tag{7}$$

We now invoke the steady-state condition in (6), substitute for $(L\psi)(\mathbf{X})$ from (7) under the expectation, and after some algebra and re-arrangement of terms (see [13] for details) get the recursion result (3) in Theorem 1.

V. Conclusion

Beyond yielding tractable expressions for steady-state VAoI under general network and queueing configurations, our results reveal how sensing correlations, network connectivity, and service constraints jointly shape informational freshness. Unlike prior VAoI models that assume externally specified version-generation rates, our framework derives version creation from process evolution and sensing. This enables a unified treatment of process dynamics, sensing, gossip dissemination, and queueing within a single analytical framework for VAoI.

References

[1] R. D. Yates, "The age of gossip in networks," in *2021 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2021, pp. 2984–2989.

[2] S. Ioannidis, A. Chaintreau, and L. Massoulié, "Optimal and scalable distribution of content updates over a mobile social network," in *IEEE INFOCOM 2009*. IEEE, 2009, pp. 1422–1430.

[3] B. Buyukates, M. Bastopcu, and S. Ulukus, "Version age of information in clustered gossip networks," *IEEE Journal on Selected Areas in Information Theory*, vol. 3, no. 1, pp. 85–97, 2022.

[4] K. J. Khojastepour, M. Mortaheb, and S. Ulukus, "Age of gossip in networks with multiple views of a source," in *2025 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2025, pp. 1–6.

[5] 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.

[6] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1183–1210, 2021.

[7] M. Costa, M. Codreanu, and A. Ephremides, "Age of information with packet management," in *2014 IEEE International Symposium on Information Theory*. IEEE, 2014, pp. 1583–1587.

[8] C.-H. Tsai and C.-C. Wang, "Distribution-oblivious online algorithms for age-of-information penalty minimization," *IEEE/ACM Transactions on Networking*, vol. 31, no. 4, pp. 1779–1794, 2023.

[9] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Transactions on Networking*, vol. 28, no. 5, pp. 2215–2228, 2020.

[10] M. Salimnejad, M. Kountouris, A. Ephremides, and N. Pappas, "Age of information versions: a semantic view of markov source monitoring," *IEEE Transactions on Communications*, 2025.

[11] 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.

[12] E. Erbayat, A. Maatouk, P. Zou, and S. Subramaniam, "Age of information optimization and state error analysis for correlated multi-process multi-sensor systems," in *Proc. of the Twenty-fifth Intl. Symp. on Theory, Algo. Foundations, and Protocol Design for Mob. Networks and Mobile Comp.*, 2024, pp. 331–340.

[13] A. Meena and A. J. Budkuley, "Version age in multi-process gossip networks under distributed sensing," *ArXiv (Extended draft)*, 2026.

[14] R. Gallager, *Stochastic Processes: Theory for Applications*, ser. Stochastic Processes: Theory for Applications. Cambridge University Press, 2013.

[15] J. P. Hespanha, "Modelling and analysis of stochastic hybrid systems," *IEE Proceedings-Control Theory and Applications*, vol. 153, no. 5, pp. 520–535, 2006.