Information Freshness for Monitoring and Control over Wireless Networks

Vishrant Tripathi

PhD Affiliation: Massachusetts Institute of Technology, Electrical Engineering and Computer Science

Homepage: https://www.mit.edu/~vishrant

Thesis: Information Freshness for Monitoring and Control over Wireless Networks

Advisor: Eytan Modiano, MIT

Brief Biography: Vishrant Tripathi obtained his PhD from the EECS department at MIT, working with Prof. Eytan Modiano at the Lab for Information and Decision Systems (LIDS). He is currently working on building efficient data center networks at Google. His research interests primarily lie in the optimization of resources in resource constrained networked systems. The main applications of his work are in multi-agent robotics, federated learning, edge computing, cloud infrastructure, and monitoring for IoT. More recently, he has also been working on software defined networking and next-generation wireless networks. In 2022, he won the Best Paper Runner Up Award at ACM MobiHoc.

1. INTRODUCTION

Monitoring and control of dynamical systems are fundamental and well-studied problems. Many emerging applications involve performing these tasks over communication networks. Examples include: sensing for IoT, control of robot swarms, real-time surveillance, and environmental monitoring by sensor networks. Such systems typically involve multiple agents collecting and sending information to a central entity where data is stored, aggregated, analyzed, and then possibly used to send back control commands. Due to the dramatic improvements both in on-device and edge computing, and in wireless communication over the past two decades, there has been a rapid growth in the size and scale of such networked systems.

The central focus of my thesis [1] is *understanding how to optimize the flow of information in networks, in order to achieve real-time monitoring and control.* We make contributions in three directions.

First, we consider the optimization of general cost functions of Age of Information (AoI). Here, we develop computationally efficient scheduling algorithms for optimizing information freshness in both single-hop and multi-hop wireless networks. We further develop an online learning formulation when the cost functions of AoI are unknown and propose a new online learning algorithm for this setting called Follow-the-Perturbed-Whittle-Index.

Second, we consider weighted-sum AoI minimization. In this setting, we study how correlation impacts information freshness. We also propose a near-optimal distributed scheduling protocol called Fresh-CSMA for AoI minimization, that has provable performance guarantees.

Third, we apply our theoretical results to problems in multi-agent robotics and monitoring – both via simulations and practical system implementations. We use simulations to demonstrate significant performance improvements in the collection of time-varying occupancy grid maps using multiple robots via the Whittle Index framework. Further, to demonstrate the benefits of our theoretical contributions, we build a real system (WiSwarm) for mobility tracking using a swarm of UAVs, communicating with a central controller over WiFi. Our experimental results show that, when compared to the standard IEEE 802.11 MAC layer + TCP/UDP, our system can *reduce AoI by a factor of* ***109x/48x*** *and improve tracking accuracy by a factor of* ***4x/6x****, respectively.*

1.1 Age of Information

In multi-agent robotics applications, such as tracking in adversarial environments or search-and-rescue missions, robots often need to exchange high fidelity data in real-time. This is the main motivation behind our work. Traditional wireless networking solutions optimize for standard performance metrics such as throughput and delay. However, these tend to perform poorly when used out-of-the-box for applications that require real-time performance. A metric called the **Age of Information** (AoI), that measures information freshness of general systems [2, 3, 4], has gained popularity in the networking community over the last decade.

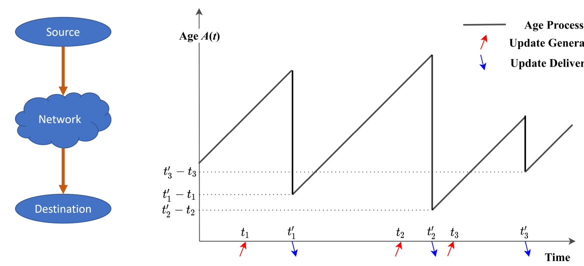

Figure 1: Illustration of the AoI evolution for sample generation and delivery processes. The first update is generated at the source at time t_1 and is delivered to the destination at time t'_1. The destination now has information about the source that is $t'_1 - t_1$ old, so AoI drops to $A(t'_1) = t'_1 - t_1$.

Age-of-Information (AoI) is an end-to-end metric that characterizes *how old the information is from the perspective of the destination.* Consider a destination receiving time-stamped updates from a source over a network. Let $\tau(t)$ be the time-stamp of the *latest update* received at the destination by time t. The AoI associated with this source-destination pair is then defined as $A(t) := t - \tau(t)$. The AoI increases linearly with time when no updates are delivered, representing the information getting older. At the moment a *fresher* update from the source is received at the destination, the value of $\tau(t)$ increases and the AoI reduces to the delay of the received update. This evolution of the AoI metric with time is illustrated in Fig. 1.

Typically, AoI represents a measure of distortion between the state of the system that is expected at the monitor based on past updates and the actual current state of the system. Thus, a larger age corresponds to the monitor having a higher uncertainty about the current state of the system being observed. This, in turn, means that ensuring a low average AoI can lead to higher monitoring accuracy or better control performance.

While AoI is a proxy for measuring the cost of having out-of-date information, it may not properly reflect the impact of stale information on system performance. It turns out that even for simple monitoring and control applications linear AoI can be an inaccurate metric to track accuracy or overall system performance. This has motivated interest in using general, possibly non-linear cost functions of AoI that reflect the cost of delayed information more accurately.

1.2 Relation to Monitoring and Control

Through a simple example, we establish why non-linear functions of AoI are key to performing monitoring tasks over wireless networks.

Consider a linear dynamical system that is being observed over a costly wireless channel (see Fig. 2). The source to be monitored evolves as follows:

$$\boldsymbol{x}(t+1) = \boldsymbol{G}\boldsymbol{x}(t) + \boldsymbol{w}(t), \tag{1}$$

where $\boldsymbol{G}$ describes the source dynamics and $\boldsymbol{w}(t) \sim \mathcal{N}(0, \Sigma)$ is i.i.d. Gaussian noise in every time-slot.

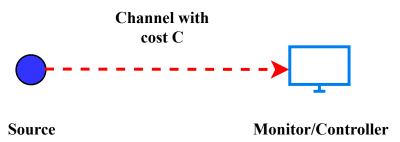

Figure 2: Single source monitoring over a wireless channel.

The monitor can, in every time-slot, decided to observe the state of the source exactly and pay an observation cost C or estimate the current state of the source based on prior observations. The goal of the monitor is to minimize the sum of the monitoring error and observation cost, averaged over time.

Now consider a scheduling policy π that specifies whether the monitor should sample the source at time-slot t or not, for every time-slot. This decision is represented by the indicator variable $u(t)$. The optimization problem faced by the monitor can then be formulated as:

$$\arg\min_{\pi} \left(\limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \left(\mathbb{E}\left[||\boldsymbol{x}(t) - \hat{\boldsymbol{x}}(t)||^2 \right] + Cu(t) \right) \right). \tag{2}$$

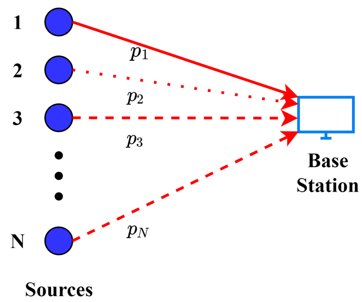

Figure 3: N sources transmitting updates to a base station over a wireless channel, with different reliabilities.

In our work, we show that this problem can be converted to an equivalent problem of the form

$$\arg\min_{\pi} \left(\limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \left(f(A(t)) + Cu(t) \right) \right), \tag{3}$$

where $A(t)$ is the AoI at the monitor of information regarding the source and $f(\cdot)$ is a monotone increasing function given by $f(h) = \sum_{k=0}^{h} Tr(\boldsymbol{G}^{k^T} \boldsymbol{G}^k \Sigma)$

2. THESIS OVERVIEW

Next, we provide a high-level overview of the technical contributions of the works that make up [1].

2.1 Real-Time Monitoring and Control

In [5, 6], we utilize the equivalence between optimizing monitoring error/control cost and functions of AoI to formulate a general model involving multiple sources, where only one source can send updates at any given time due to interference constraints (see Fig. 3). We measure monitoring and control performance using general nonlinear functions of AoI $f_i(A_i(t))$ and formulate the following optimization problem over the space of wireless feasible scheduling policies.

$$\arg\min_{\pi\in\Pi} \frac{1}{T} \mathbb{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{N} f_i(A_i^{\pi}(t)) \right]. \tag{4}$$

We develop a novel scheduling policy based on the Whittle Index approach that solves this problem and is close to optimal.

Our work [5] is the first to look at general AoI cost functions and their scheduling with multiple sources. We were also *the first to show that the Whittle Index approach can be exactly optimal for a finite size asymmetric scheduling problem.*

2.2 Online Learning in Wireless Networks

Continuing along this line of work, we ask a natural question - *what if the cost functions of AoI that we want to optimize are unknown, time-varying and possible adversarial.* This question is motivated by settings where we do not know the dynamics of the systems to be monitored or controlled beforehand, and where the underlying dynamics can also change over time. A typical example is mobility tracking, where target positions and velocities can be unknown, time-varying and adversarial.

Now, not only do we need to optimize scheduling, but

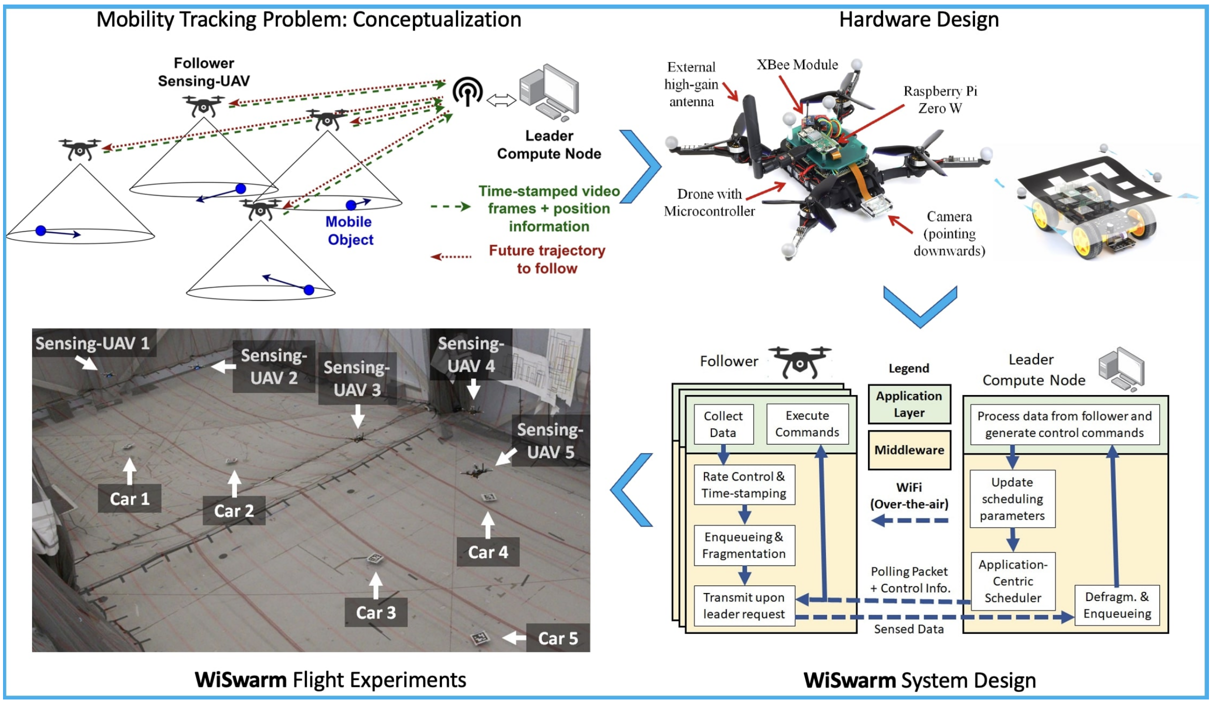

Figure 4: *WiSwarm*: From conceptualization to multi-UAV flight experiments

also learn and adapt to changes in the wireless network. In [7], we formulate an online learning framework that answers this question. We develop a novel algorithm called *Follow the Perturbed Whittle Index* which incorporates the Whittle Index into an online learning scheme while providing *strong regret guarantees*. This was an especially challenging problem since online learning for the class of optimization problems we were interested in had not been looked at before in literature. *We were the first to develop a computationally efficient way to perform online learning for this general class of problems, called Restless Multi-Armed Bandits (RMABs).*

2.3 Routing and Scheduling in General Network Topologies

A second natural question arose from our work in [5] - *how does one optimize information freshness in general network topologies, i.e. multi-hop networks?*

Solving multi-hop AoI optimization has been one of the major challenges in the field, similar to multi-hop delay optimization. It requires the design of both scheduling and routing schemes, and techniques and insights from single-hop optimization do not tend to work well in the multi-hop context. Our work on *Age-Debt* [8, 9], inspired by classical Lyapunov control, is currently the best known general purpose policy for multi-hop networks. It can handle unicast, multicast or broadcast flows; general cost functions; and arbitrary network topologies. *Our work was the first to handle the problem in its full generality and our proposed algorithms currently match or outperforms all prior methods in numerical experiments.*

2.4 Multi-Agent Robotics

Encouraged by our theoretical and numerical results, we wanted to go back to our original motivations and answer the following question - *do the theoretical insights and frameworks developed above apply to real-world multi-agent robotic systems?*

To answer this, we first used simulations. In [10], we look at two problems - multi-agent occupancy grid mapping and multi-agent ride-sharing. We apply the AoI Whittle Index framework to these settings, and show that our approach can a) deliver significantly higher quality and fresher maps in the first problem and b) shorter wait times for riders in the second problem [11]. *This was the first work to directly apply ideas from information freshness to problems in robotics and demonstrate the performance benefits of utilizing networks tailored to application-specific needs.*

Parallel to this, *we designed and implemented a new kind of wireless system for multi-agent robotics* [12]. Fig. 4 described how we formulated a multi-agent mobility tracking problem, designed hardware and software systems using the theoretical insights from our prior work, and then conducted multi-UAV flight experiments using these systems.

Our system, **WiSwarm**, outperforms standard networking solutions by an order of magnitude for the multi-UAV mobility tracking task in real flight experiments [13]. *This verified that our theoretical insights can lead to significant impact in practical applications.*

2.5 Correlated Sources & Distributed Protocols

During our collaborations and discussions with roboticists, we further realized there were two crucial gaps in the AoI literature. First, we didn't know how to handle the monitoring or control of sources that are coupled or send correlated updates. Second, we don't have distributed wireless scheduling policies (which are much easier to implement in practice, and also a part of standards like IEEE 802.11WiFi) with strong information freshness guarantees. Our recent work has directly addressed these open gaps. Our work [14] on information freshness with correlated sources (*which won the Best Paper Runner-Up Award at MobiHoc 2022*) is the first work to provide scheduling policies that take correlation into account. and answers how correlation affects information freshness and how it should be utilized for scheduling. Our algorithm *Fresh-CSMA* [15], motivated by Carrier Sense Multiple Access (CSMA) style protocols, has the strongest known guarantees for distributed protocols in the AoI literature.

3. FUTURE DIRECTIONS

Multi-Agent Robotics and Computational Offloading: We plan to continue working on the development of networking systems for multi-agent robotics, along the lines of [10, 12]. There are recent trends towards computational offloading to create more efficient and scalable robotics systems. However, computation offloading, whether at the edge or the cloud, requires high throughput, low delay and reliable wireless networks to ensure that robots can continue to perform safely and as intended, in real-time, without losing control. *How can we design networking solutions for multi-agent applications such as search-and-rescue in deep subterranean cave networks, large industrial warehouses, and fleets of autonomous vehicles while ensuring safe operation and efficient communication?*

Software Defined Networking (for cloud infrastructure): A large part of traffic engineering, routing and topology optimization for modern data centers happens via Software Defined Networking (SDN), where centralized controllers make decisions to optimize performance based on monitoring data that they receive from the entire network. Herein lies a dilemma. *Fresh and accurate monitoring causes huge overheads, but is necessary to deliver good performance.* Interestingly, this is the same resource optimization vs. information freshness dilemma that we have discussed above. Applying our techniques can provide a very fruitful direction for research, and avenues for collaboration between academia and industry.

Federated Learning: Performing federated learning over resource-constrained wireless networks has received significant interest in the networking community over the past few years. The central question here is also one of resource optimization - obtaining gradient information from all users in the network is too involved and time-consuming, so the central aggregator needs to sample gradients from small batches of users at any given time. *How should one go about picking users to guarantee better performance? What happens when the underlying learning tasks and datasets change over time?* We plan to use tools from information freshness and online learning to answer these questions, building on [7].

Representative Papers:

[R1] A Whittle Index Approach to Minimizing Functions of Age of Information. (Allerton 2019)
with E. Modiano.

[R2] Optimizing Age of Information with Correlated Sources. (*Best Paper Runner-Up* at ACM MobiHoc 2022)
with E. Modiano.

[R3] Information Freshness in Multihop Wireless Networks. (ACM/IEEE Transactions on Networking)
with R. Talak, E. Modiano.

[R4] WiSwarm: Time-Sensitive Wireless Networking for a Collaborative Team of UAVs. (IEEE INFOCOM 2023)
with I. Kadota, E. Tal, M. S. Rahman, A. Warren, S. Karaman, E. Modiano.

4. REFERENCES

[1] V. Tripathi, "Information Freshness for Monitoring and Control over Wireless Networks," *PhD Thesis, MIT, https://www.mit.edu/~vishrant/tripathi-vishrant-phd-2023-thesis.pdf*, 2023.

[2] S. Kaul, R. Yates, and M. Gruteser, "Real-Time Status: How Often Should One Update?," in *IEEE INFOCOM*, 2012.

[3] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of Information: An Introduction and Survey," *IEEE JSAC*, 2021.

[4] Y. Sun, I. Kadota, R. Talak, and E. Modiano, *Age of Information: A New Metric for Information Freshness*. Springer Nature, 2022.

[5] V. Tripathi and E. Modiano, "A Whittle Index Approach to Minimizing Functions of Age of Information," in *Allerton*, 2019.

[6] V. Tripathi and E. Modiano, "A Whittle Index Approach to Minimizing Functions of Age of Information," *arXiv preprint arXiv:1908.10438*, 2019.

[7] V. Tripathi and E. Modiano, "An Online Learning Approach to Optimizing Time-Varying Costs of AoI," in *ACM MobiHoc*, 2021.

[8] V. Tripathi and E. Modiano, "Age Debt: A General Framework for Minimizing Age of Information," in *IEEE INFOCOM AoI Workshop*, 2021.

[9] V. Tripathi, R. Talak, and E. Modiano, "Information Freshness in Multihop Wireless Networks," *IEEE/ACM Transactions on Networking*, 2022.

[10] V. Tripathi, L. Ballotta, L. Carlone, and E. Modiano, "Computation and Communication Co-design for Real-Time Monitoring and Control in Multi-Agent Systems," in *IFIP WiOpt*, IEEE, 2021.

[11] "Computation and Communication Co-Design for Real-Time Monitoring and Control in Multi-Agent Systems." https://www.youtube.com/watch?v=gs5lb3YE-ec.

[12] V. Tripathi, I. Kadota, E. Tal, M. S. Rahman, A. Warren, S. Karaman, and E. Modiano, "WiSwarm: Age-of-Information-based Wireless Networking for Collaborative Teams of UAVs," in *IEEE INFOCOM*, 2023.

[13] "WiSwarm: Time-Sensitive Wireless Networking for a Collaborative Team of UAVs." https://www.youtube.com/watch?v=zGjGxK5AXFk.

[14] V. Tripathi and E. Modiano, "Optimizing Age of Information with Correlated Sources," in *ACM MobiHoc, Best Paper Runner-Up*, 2022.

[15] V. Tripathi, N. Jones, and E. Modiano, "Fresh-CSMA: A Distributed Protocol for Minimizing Age of Information," in *IEEE INFOCOM*, 2023.