Minimizing Age-of-Information in Heterogeneous Multi-Channel Systems: A New Partial-Index Approach

Yihan Zou, Kwang Taik Kim, Xiaojun Lin, and Mung Chiang
Electrical & Computer Engineering, Purdue University
{zou59,kimkt,linx,chiang}@purdue.edu

ABSTRACT

We study how to schedule data sources in a wireless time-sensitive information system with multiple heterogeneous and unreliable channels to minimize the total expected Age-of-Information (AoI). Although one could formulate this problem as a discrete-time Markov Decision Process (MDP), such an approach suffers from the curse of dimensionality and lack of insights. For single-channel systems, prior studies have developed lower-complexity solutions based on the Whittle index. However, Whittle index has not been studied for systems with multiple heterogeneous channels, mainly because indexability is not well defined when there are multiple dual cost values, one for each channel. To overcome this difficulty, we introduce new notions of partial indexability and partial index, which are defined with respect to one channel's cost, given all other channels' costs. We then combine the ideas of partial indices and max-weight matching to develop a Sum Weighted Index Matching (SWIM) policy, which iteratively updates the dual costs and partial indices. The proposed policy is shown to be asymptotically optimal in minimizing the total expected AoI, under a technical condition on a global attractor property. Extensive performance simulations demonstrate that the proposed policy offers significant gains over conventional approaches by achieving a near-optimal AoI. Further, the notion of partial index is of independent interest and could be useful for other problems with multiple heterogeneous resources.

CCS CONCEPTS

• **Networks → Network performance analysis**; **Mobile networks**; • **Theory of computation → Scheduling algorithms**.

KEYWORDS

Age-of-Information, Whittle index, restless bandits, Markov decision processes, heterogeneous channels.

ACM Reference Format:
Yihan Zou, Kwang Taik Kim, Xiaojun Lin, and Mung Chiang. 2021. Minimizing Age-of-Information in Heterogeneous Multi-Channel Systems: A New Partial-Index Approach. In *The Twenty-second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '21), July 26–29, 2021, Shanghai, China.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3466772.3467030

MobiHoc '21, July 26–29, 2021, Shanghai, China

ACM ISBN 978-1-4503-8558-9/21/07…$15.00
https://doi.org/10.1145/3466772.3467030

1 INTRODUCTION

Many emerging wireless applications (e.g., real-time control in robotics systems, and data collection for IoT applications) rely on timely status updates from information sources [1, 4]. In these applications, oftentimes only the information with the latest timestamp is valuable to the receiver, while out-dated packets have little value. These applications have motivated a growing body of literature in optimizing the Age-of-Information (AoI), which is defined as the elapsed time of the last-received information packet since it was generated (at the source). Intuitively, AoI captures the freshness of information from the data source's perspective, and is considered a more useful metric for time-sensitive information systems than packet-level delays [15].

In this paper, we are interested in minimizing AoI for a wireless system with multiple heterogeneous sources and channels. This is a difficult setting that still lacks effective solutions in the literature. Many existing work on minimum-AoI scheduling policies study only a *single-source* system [10, 15]. For multiple sources, most of the existing work assumes that data sources are transmitting in a *single shared channel* [6–9, 14, 16]. Further, most of these studies assume the channel to be reliable, with only a few extension to the case of a single unreliable channel [8, 14, 16]. For studies that do involve multiple channels, a recent article [12] assumes a *homogeneous* channel model, where each user-channel pair has equal ON/OFF probability. Thus, the solutions in these studies cannot be used in wireless systems that exhibit heterogeneous channel condition (e.g., transmission success probability) for each source-channel pair, which is common due to antenna beamforming, frequency selectivity and location-dependent fading [19, 23, 24]. [17] and [13] are the closest work to ours, as they also study heterogeneous multi-channel systems. [17] proposes a scheduling policy for ON/OFF multi-channel systems based on max-age matching. Under a similar setting, [13] proposes a policy that is asymptotically 8-optimal in minimizing the total weighted age. However, [17] and [13] assume that the ON/OFF states of all channels are known before the scheduling decisions are made. This assumption, combined with the setting that the number of channels are large, ensures that with high probability each source sees at least one ON channel. In this way, the impact of unreliable channels can be absorbed by an event with negligible probability in their analysis. In contrast, we are interested in a model where the channel states are unknown when scheduling decisions are made. Therefore, the question remains open on *how to design a provably optimal scheduling policy to minimize AoI in time-sensitive information systems with multiple heterogeneous and unreliable channels.*

One of the key obstacles in deriving the optimal scheduling policy under multiple heterogeneous sources and channels is the complexity of the associated Markov decision problem. Note that such AoI optimization problems (regardless of the channel conditions) are often formulated as Markov Decision Processes (MDP) or Restless Multi-armed Bandits (RMAB), which in theory can be optimally solved by value iteration [2, 3]. However, this approach suffers from *the curse of dimensionality* and lack of insights. Therefore, it is highly desirable to develop low-complexity and near-optimal solutions. For single-channel systems, policies based on Whittle index [21], whose complexity does not grow with the number of sources, have been found to exhibit good performance. Further, they are known to be asymptotically optimal when the number of sources and the channel capacity both grow to infinity [6, 7, 14, 16, 20]. However, to the best of our knowledge, there have been no such Whittle index policies for systems with multiple heterogeneous channels/resources. Part of the difficulty is that Whittle's notion of "indexability" [21] is not well-defined when there are multiple heterogeneous channels. Specifically, in [21], a project is indexable if there is a single threshold for the channel cost, above which the optimal action of the project will be passive (i.e., not to consume the channel resource). Thus, Whittle indexability critically relies on the assumption that there is only one dual cost for either a single channel or a single group of homogeneous channels. For heterogeneous multi-channel systems, each channel naturally has a different dual cost. The optimal action of the project will also depend on all channel costs. As a result, one cannot even define such a threshold or index.

In this paper, we propose a new Whittle-like scheduling policy for heterogeneous and unreliable multi-channel systems. Similar to [21], we first formulate the MDP for the system, and decompose the problem into per-source sub-problems using Lagrange relaxation [3] (Section 2). However, to overcome the difficulty of Whittle indexability as mentioned above, we introduce the new notions of *partial indexability* and partial index, which are defined with respect to the cost of one channel, given the costs of all other channels (see Section 3 for detailed definitions). Then, we propose a low-complexity Whittle-like scheduling policy, which we call the Sum Weighted Index Matching (SWIM) policy, by computing a maximum-weighted matching (MWM) between the sources and channels, where the weight between each source-channel pair is the above-defined partial index. Our key contribution in Section 3 is to identify a precise-division condition, under which the SWIM policy is asymptotically optimal, under a technical assumption on a global attractor property (which has also been used in the literature [5, 18]). To the best of our knowledge, our work is the first in the literature to extend the concept of indexability to heterogeneous multi-channel settings. We note that both the notion of partial indexability and the SWIM policy are very general, and can be applied to various large-scale MDP problems with multiple heterogeneous channels. We then verify in Section 4 that our AoI problem indeed satisfies the partial indexability and precise-division property. Our simulation results in Section 5 shows that applying the SWIM policy to our AoI problem produces significant performance gains over conventional approaches, and achieves a near-optimal average AoI.

We note that in the RMAB literature there is also a line of work on multi-action bandits [5, 18]. However, we emphasize that "actions" and "channels" are very different, because the multiple actions in [5, 18] are still applied to a *single resource*. This is the reason why [5] can still define a Whittle index based on the (single) dual cost associated with the resource. In contrast, multiple heterogeneous channels correspond to *multiple resources* and multiple dual costs. Therefore, the techniques in [5] cannot be directly applied to our setting with heterogeneous channels.

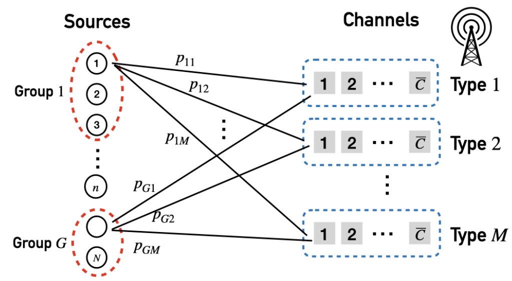

Figure 1: Uplink of a heterogeneous, unreliable multi-channel system with a base station and N data sources.

2 MODEL AND PROBLEM FORMULATION

We consider a wireless system where a base station (BS) is scheduling N data sources or sensors on multiple channels for timely status updates in the uplink (Fig. 1). Each source corresponds to one sender node on the left in Fig. 1. Note that each source may experience different channel conditions due to their locations. As a result, sources may have different preferences on the set of communication channels. To model such heterogeneity, we assume that the sources are divided into G groups. Let $\mathcal{N}_g$ be the set of sources in group g. Then, the set of all sources is $\mathcal{N} = \bigcup_{g=1}^{G} \mathcal{N}_g$. The sources $n \in \mathcal{N}_g$ in the same group g experience the same condition on each channel. We consider a discrete-time system where time is indexed by $t \in \mathcal{T}$. We assume that the transmission from the source to the BS takes one time slot.

Heterogeneous and Unreliable Channels: As shown in Fig. 1, the BS is capable of communicating in multiple channels at each time. Depending on the frequency, modulation, and beam-forming schemes used, the channels may have similar or different qualities. To model such heterogeneity, we divide the channels also into $M > 1$ types. We assume that each type $m \in \mathcal{M} = \{1, \ldots, M\}$ of channels has $\overline{C}$ identical instances (which we refer to in the future as "channels"). As we explain below, all sources in a given group g sees the same channel quality in the $\overline{C}$ channel instances of a given type m. We adopt the standard collision channel model [14] as follows. At each time, the BS can schedule at most one source to transmit update packets on each channel. The channel is potentially *unreliable*, due to wireless channel fading. In contrast to most existing work in the literature, we consider *heterogeneous source-channel conditions*. Specifically, we assume that each transmission from source $n \in \mathcal{N}_g$ on a channel of type $m \in \mathcal{M}$ succeeds with probability $p_{gm} \in (0, 1]$, independently from all other transmissions. We denote the *channel quality vector* for group g as $\vec{p}_g = [p_{g1}, \ldots, p_{gM}]^T$. With slight abuse of notation, we denote $p_{nm} = p_{g(n),m}$ where $g(n)$ is the group containing n.

Packet Generation: To focus our discussion on the effect of multiple heterogeneous channels, we adopt the generate-at-will model as [9, 16]. Specifically, whenever a source is scheduled for transmission, it can generate a fresh update. In this work, we use age-of-information (or simply age) to measure the information freshness, which is defined as the elapsed time of the last-received information packet since it was generated (at the source). Denote $h_n(t)$ as the age of source n at time t. If the transmission is successful, the age of this source reduces to 1. If the source is not scheduled for transmission, or if the transmission fails, the age increases by 1. Then, the AoI evolution of source n can be written as

$$h_n(t+1) = \begin{cases} 1, & \text{successful update from } n; \\ h_n(t)+1, & \text{otherwise.} \end{cases} \tag{1}$$

Intuitively, to avoid wasting channel resources, the BS should only schedule each source on at most one channel instance. Let $u_n(t)$ be the decision variable at time t such that $u_n(t) = m$ if source n is scheduled to transmit on channel of type m, and $u_n(t) = 0$ if the source is not scheduled for transmission. In summary, we have the constraints that $\sum_{n=1}^{N} \mathbb{1}\{u_n(t) = m\} \le \overline{C}$ for all channel type m, and $\sum_{m=1}^{M} \mathbb{1}\{u_n(t) = m\} \le 1$ for all source n.

2.1 MDP-based Formulation

Now, we can formulate the average AoI minimization problem for the above heterogeneous and unreliable multi-channel system as an MDP. Let $\overline{\mathcal{S}}(t) \triangleq \{h_1(t), h_2(t), \dots, h_N(t)\} \in \mathbb{N}_+^N$ be the system state at time t. Denote the action space of the entire system as $\overline{\mathcal{U}} \triangleq \{0, 1, \dots, M\}^N$. (Recall that action 0 denotes no scheduled transmission and action $m \in \mathcal{M}$ denotes the scheduled channel type). A policy $\boldsymbol{\pi}$ maps from the system state $\overline{\mathcal{S}}(t)$ to the action in $\overline{\mathcal{U}}$. The state transition probability of source n when it is passive is

$$\mathbb{P}\{h_n(t+1) = d+1 | h_n(t) = d, u_n(t) = 0\} = 1. \tag{2}$$

The state transition probabilities when source n is scheduled on a channel of type m are

$$\begin{aligned} &\mathbb{P}\{h_n(t+1) = d+1 | h_n(t) = d, u_n(t) = m\} = 1 - p_{nm}, \\ &\mathbb{P}\{h_n(t+1) = 1 | h_n(t) = d, u_n(t) = m\} = p_{nm}. \end{aligned} \tag{3}$$

We can define the T-horizon average AoI and the long-term average AoI of the system under policy $\boldsymbol{\pi}$ as

$$H_{\boldsymbol{\pi}}^{(T)} = \frac{1}{TN} \sum_{t=1}^{T} \sum_{n=1}^{N} \mathbb{E}[h_n^{\boldsymbol{\pi}}(t)], \text{ and } \overline{H}_{\boldsymbol{\pi}} \triangleq \limsup_{T\to\infty} H_{\boldsymbol{\pi}}^{(T)}, \tag{4}$$

respectively, where T is the length of time horizon, and $h_n^{\boldsymbol{\pi}}(t)$ is the AoI of source n at time t under policy $\boldsymbol{\pi}$. The objective of the MDP is to minimize the long-term average system AoI in (4), i.e.,

$$\min_{\boldsymbol{\pi} \in \overline{\mathcal{U}}^T} \limsup_{T\to\infty} \frac{1}{TN} \sum_{t=1}^{T} \sum_{n=1}^{N} \mathbb{E}\left[h_n^{\boldsymbol{\pi}}(t)\right]. \tag{5}$$

In theory, the above MDP can be solved optimally as an infinite-horizon average cost per stage problem using relative value iteration [2]. However, this approach suffers from the curse of dimensionality and lack of insights for the solution structure. Hence, many efforts have been focusing on developing low-complexity solutions.

2.2 Decomposition Using Lagrange Relaxation

For lower-complexity solutions, two representative approaches in the literature are based on the relaxed problem and index policies. In this section, we will discuss how they are related to a Lagrange relaxation of the MDP, and the challenges of applying these existing approaches to our setting with multiple heterogeneous channels.

We first introduce the relaxed problem. Denote $u_{nm}^{\boldsymbol{\pi}}(t) \triangleq \mathbb{1}\{u_n(t) = m\}$, i.e., the indicator variable that source n is scheduled on channel type m at time t under policy $\boldsymbol{\pi}$. The MDP formulated in Section 2.1 can be equivalently written in the following optimization form:

$$\begin{aligned} \underset{\boldsymbol{\pi}}{\text{minimize}} \quad & \limsup_{T\to\infty} \frac{1}{TN} \sum_{t=1}^{T} \sum_{n=1}^{N} \mathbb{E}\left[h_n^{\boldsymbol{\pi}}(t)\right] \\ \text{subject to} \quad & \sum_{n=1}^{N} u_{nm}^{\boldsymbol{\pi}}(t) \le \overline{C}, \quad \forall m \in \mathcal{M}, t \in \mathcal{T}, && \text{(6a)} \\ & \sum_{m=1}^{M} u_{nm}^{\boldsymbol{\pi}}(t) \le 1, \quad \forall n \in \mathcal{N}, t \in \mathcal{T}, && \text{(6b)} \\ & u_{nm}^{\boldsymbol{\pi}}(t) \in \{0,1\}, \quad \forall t \in \mathcal{T}. && \text{(6c)} \end{aligned}$$

Following Whittle's approach [21], we relax the instantaneous constraint (6a) to an average constraint, and obtain the *relaxed problem*

$$\begin{aligned} \underset{\boldsymbol{\pi}}{\text{minimize}} \quad & \limsup_{T\to\infty} \frac{1}{TN} \sum_{t=1}^{T} \sum_{n=1}^{N} \mathbb{E}\left[h_n^{\boldsymbol{\pi}}(t)\right] \\ \text{subject to} \quad & \limsup_{T\to\infty} \frac{1}{T} \mathbb{E}[\sum_{t=1}^{T} \sum_{n=1}^{N} u_{nm}^{\boldsymbol{\pi}}(t)] \le \overline{C}, \quad \forall m \in \mathcal{M}, && \text{(7a)} \\ & \text{(6b), (6c)} && \text{(7b)} \end{aligned}$$

Next, we use Lagrange relaxation in [3, Chapter 6]. Specifically, we introduce a dual cost λ_m to each of (7a), and decouple the relaxed problem of (7) into N *sub-problems*, i.e., $\forall n \in \mathcal{N}$,

$$\begin{aligned} \underset{\boldsymbol{\pi}}{\text{minimize}} \quad & \limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\big[h_n^{\boldsymbol{\pi}}(t) + \sum_{m\in\mathcal{M}} \lambda_m u_{nm}^{\boldsymbol{\pi}}(t)\big] && \text{(8a)} \\ \text{subject to} \quad & \sum_{m=1}^{M} u_{nm}^{\boldsymbol{\pi}}(t) \le 1, \quad \forall t \in \mathcal{T}, && \text{(8b)} \\ & u_{nm}^{\boldsymbol{\pi}}(t) \in \{0,1\}, \quad \forall t \in \mathcal{T}, \forall m \in \mathcal{M}. && \text{(8c)} \end{aligned}$$

It is easy to see that, given channel costs $\vec{\lambda} = [\lambda_m, \forall m \in \mathcal{M}]$, each sub-problem (8) is an average cost per stage problem [2] for optimizing the long-term average AoI of source n plus the costs of its channel use, which by itself is a decoupled MDP with state space $\mathcal{S} = \mathbb{N}_+$, action space $\mathcal{U} = \{0, 1, \dots, M\}$, and transition probabilities in (2) and (3).

With this dual decomposition, the relaxed problem can be solved by iteratively solving all independent sub-problems (8) given the current $\vec{\lambda}^{(k)}$ (denote the resulting decisions by $u_{nm}^{\boldsymbol{\pi}(k)}(t)$) and updating the dual costs $\vec{\lambda}^{(k)}$ by dual gradient ascent in the k-th iteration [11], i.e., for all $m \in \mathcal{M}$

$$\lambda_m^{(k+1)} = \left[\lambda_m^{(k)} + \rho \cdot \left(\mathbb{E}\left[\sum_{n\in\mathcal{N}} u_{nm}^{\boldsymbol{\pi}(k)}(t)\right] - \overline{C}\right)\right]^+ \tag{9}$$

where $\rho > 0$ is the step size, and $[x]^+ = \max\{x, 0\}$. When $\vec{\lambda}^{(k)}$ converges, the corresponding solution of the relaxed problem is known to provide us with a lower bound for the objective of the original MDP (6) [20]. However, this solution does not always provide a feasible scheduling decision, because, in the real system, resource constraints (6a) must be met at all time, instead of just in the average sense as in (7a). Moreover, before the primal-dual iteration converges, the average constraint may be violated severely, resulting in poor policy performance.

For single-channel systems ($M = 1$), Whittle's index policy in the literature overcomes these drawbacks by producing a scheduling decision that is always feasible, and is near-optimal [3, 14, 20, 21]. However, Whittle index or indexability have not been defined for multiple heterogeneous channels. Note that in the RMAB literature, the indexability is defined based on the following: for each state of a project, there exists a *scalar* price threshold such that, when the price is above (or below) that threshold, the resource is not used (or used). However, as we have shown above, the sub-problem (8) in our model is parameterized by multiple channel costs with distinct values. Obviously, the decision to use each channel type m depends on not only the cost λ_m of this type, but also the costs of other channel types. As a result, there is no longer a single threshold that can divide the spaces of cost vectors into one where the resource is used, and the opposite one where the resource is not used. Next, we overcome this difficulty by introducing the new notions of *partial indexability* and *partial index.*

3 PARTIAL INDEXABILITY AND ASYMPTOTICALLY OPTIMAL POLICIES

In this section, we will propose a powerful framework to design asymptotically optimal scheduling policies, which generalizes the notion of indexability to heterogeneous multi-channel settings. Specifically, we introduce a new notion of *partial indexability*, which are defined with respect to the cost of one channel, given the costs of the others. Partial indexability and the corresponding partial index then allow us to develop a near-optimal policy for heterogeneous multi-channel systems, which is a key contribution of our work.

Our proposed solution framework in this section is based on only the relaxed-problem formulation in Section 2.2. Note that the formulation of the relaxed problem in Section 2.2 can be applied to any MDP with the cost function given by $h(\cdot)$. Thus, our methodology not only applies to the AoI minimization problem in this paper, but also other large MDP problems with multiple heterogeneous channels (or resources). In that sense, the applicability of our proposed framework in this section is beyond the current problem. Thus, although we still use the notions of "sources/channels" in this section, they could be easily applied to more general notions of "projects/resources" as in the typical Whittle-index literature [20].

3.1 Partial Indexability

We first focus on the sub-problem (8) with a given vector $\vec{\lambda} = [\lambda_1, \ldots, \lambda_M]^T$ of costs for all channels. As we mentioned in Sec. 2.2, the MDP of each sub-problem is an infinite-horizon average cost per stage problem with countably infinite state space [2]. Since the sub-problems of all sources n in the same group g are identical and independent, next we can write the Bellman Equation of the sub-problem (8) for each group g as

$$f^g(s) + J^* = \min_{u\in\{0,\ldots,M\}} \left[g_u^g(s,\vec{\lambda}) + \sum\nolimits_{d\in\mathcal{S}} p_{sd}^g(u) f^g(d)\right], \quad (10)$$

where $f^g(\cdot)$ is the optimal relative value function, J^* is the optimal average cost. Here, to keep our notations general, we have used $g_u(s,\vec{\lambda}) = C_{gs}^u + \lambda_u$ to denote the stage cost in (8a) at state $s \in \mathcal{S}$ under action $u \in \mathcal{U} = \{0\} \bigcup \mathcal{M}$, and $p_{sd}^g(u)$ to denote the transition probability from state s to state d by taking action u. For the AoI minimization problem, $C_{gs}^u = h(s)$ and $p_{sd}^g(u)$ specializes to the transition probabilities in (2) and (3).

Next, we define the partial indexability and partial index that generalize Whittle's index [21]. Given the cost vector $\vec{\lambda}$, let

$$\mu_u^g(s,\vec{\lambda}) \overset{\Delta}{=} g_u^g(s,\vec{\lambda}) + \sum\nolimits_{d\in\mathcal{S}} p_{sd}^g(u) f^g(d)$$

denote the expected cost-to-go from state s under action u, assuming that the optimal policy is used in the future. We first define the following concepts that are analogous to Whittle's notations [21].

Definition 3.1 (Passive Set). *Given the cost vector $\vec{\lambda}$, the set of passive states for channel-type m is*

$$\mathcal{P}_m^g(\vec{\lambda}) \overset{\Delta}{=} \{s \in \mathcal{S} | \mu_m^g(s,\vec{\lambda}) > \min_{u\neq m, u\geq 0} \mu_u^g(s,\vec{\lambda})\}. \quad (11)$$

In other words, if the current state of a source $n \in \mathcal{N}_g$ is $s \in \mathcal{P}_m^g(\vec{\lambda})$, the solution to the relaxed problem under $\vec{\lambda}$ will not schedule source n on channel-type m. Let $\vec{\lambda}_{-m}$ denote the cost vector of all channels except for channel type m. We now fix all channel costs $\vec{\lambda}_{-m}$ except that of type m, but vary the channel cost of type m to λ'_m. Let the new cost vector be $\vec{\lambda}' = [\lambda'_m, \vec{\lambda}_{-m}]$. We define the partial indexability as follows.

Definition 3.2 (Partial Indexability). *Given the cost vector $\vec{\lambda}$, the sub-problem* (8) *is partially indexable (or indexable as abbr.) if, for all $m \in \mathcal{M}$, the size of the passive set $|\mathcal{P}_m^g(\vec{\lambda}')|$ increases monotonically to the entire state space as λ'_m increases from* 0 *to* ∞ *(while fixing other channels' costs $\vec{\lambda}_{-m}$).*

If the sub-problem (8) is partially indexable, then for each state s, there is a largest value of λ'_m such that the passive set $\mathcal{P}_m^g(\vec{\lambda}')$ no longer includes the state s. We refer to this value of λ'_m as the partial index, as defined below.

Definition 3.3 (Partial Index). *Given channel vector $\vec{p}$ and cost vector $\vec{\lambda}$, the partial index (or index, as abbr.) $I_m^g(s,\vec{\lambda}_{-m})$ of state $s \in \mathcal{S}$ for channel type $m \in \mathcal{M}$ is defined as the supremum of cost λ'_m such that the expected cost-to-go from state s for using channel type m is no larger than that under any other actions, i.e.,*

$$I_m^g(s,\vec{\lambda}_{-m}) \overset{\Delta}{=} \left[\sup\{\lambda'_m | \mu_m^g(s,\vec{\lambda}') \leq \mu_k^g(s,\vec{\lambda}'), \forall k\geq 0\}\right]^+. \quad (12)$$

In addition, we define the index for passive action ($m = 0$) as

$$I_0^g(s,\vec{\lambda}) \overset{\Delta}{=} \left[\sup\{\lambda' | \mu_0^g(s,\vec{\lambda}) + \lambda' \leq \mu_k^g(s,\vec{\lambda}), \forall k \in \mathcal{M}\}\right]^-, \quad (13)$$

where $[x]^- \overset{\Delta}{=} \min\{x, 0\}$.

Similar to Whittle policy, partial indexability allows us to characterize the urgency of each state by its indices, based on which an efficient solution for the original problem can be derived. However, in contrast to standard Whittle indexability, partial indexability is defined given *all the channel costs other than channel type m*. Like Whittle indexability, verifying such partial indexability is nontrivial, and often requires significant work. We will show how to verify partial indexability for the AoI minimization problem in Section 4.

Next, we are interested in designing a Whittle-like policy that can utilize partial indices. For single-channel system, the Whittle index policy simply picks the project with the highest index. However, such a simple decision will not work for multi-channel systems anymore, because each source is also restricted to transmit on one channel at a time. Intuitively, to respect the capacity constraints (6a) and (6b) for each channel and each source, the decision should involve some matching between sources and channels. The goal of the next section is to establish this matching formally.

3.2 Max-Weight Matching of Partial Indices

Motivated by Whittle's index policy, we aim to schedule a group of users with higher partial indices, while satisfying the resource constraints on each channel. The problem can be naturally formulated as a Maximum Weighted Matching problem based on partial indices (MWM-PI). Define the graph $\mathcal{R} \overset{\Delta}{=} (\mathcal{N} \bigcup \mathcal{M}, \mathcal{E})$, where $\mathcal{E}$ is the set of all source-channel-type pairs. We define the problem of MWM-PI as follows,

$$\underset{y_{nm} \in \{0,1\}}{\text{maximize}} \quad \sum\nolimits_{n \in \mathcal{N}} \sum\nolimits_{m=0}^{M} w_{nm} y_{nm} \tag{14a}$$

$$\text{subject to} \quad \sum\nolimits_{n \in \mathcal{N}} y_{nm} \le \overline{C}, \quad \forall m \in \mathcal{M}, \tag{14b}$$

$$\sum\nolimits_{m \in \mathcal{M}} y_{nm} \le 1, \quad \forall n \in \mathcal{N} \tag{14c}$$

where y_{nm} is the binary decision to schedule source n on channel type m, $w_{nm} \overset{\Delta}{=} I_m^{g(n)}(s_n, \vec{\lambda}_{-m})$ is the edge weight given by the partial index in (12) and (13), and $g(n)$ is the group index for user n. We then schedule the sources according to $u_{nm}(t) = y_{nm}$.

Note that MWM-PI is based on the current set of prices $\vec{\lambda}$. As we present next, the outcome of the MWM-PI will also guide us in updating the prices $\vec{\lambda}$. This idea leads to the proposed Sum Weighted Index Matching (SWIM) policy in Algorithm 1. Specifically, Line 1 initializes the system parameters. Lines 3-5 compute the scheduling decision for time t by solving the MWM-PI problem. Lines 6-7 correspond to the transmission phase of the update packets. Line 8 updates each channel type's cost for the next time $t+1$ as a weighted average (by the parameter β) of the previous channel cost and the optimal dual cost associated with (14b) at time t.

Algorithm 1: Sum Weighted Index Matching (SWIM)

1 At $t = 0$: Initialize parameters $N, M, \overline{C}, \beta$, and $\vec{\lambda}(1)$;
2 At time $t \ge 1$:
3 Compute partial indices $\vec{I}^{g(n)}(t) = [I_m^{g(n)}(s_n(t), \vec{\lambda}_{-m}(t))]$ for every source $n \in \mathcal{N}$, given current cost $\vec{\lambda}(t)$;
4 Solve MWM-PI in (14) with $w_{nm} \leftarrow I_m^{g(n)}(s_n(t), \vec{\lambda}_{-m}(t))$, and obtains the scheduling decision $\vec{y}(t)$;
5 Schedule sources according to $\vec{u}(t) \overset{\Delta}{=} [u_{nm}(t)] = \vec{y}(t)$;
6 Wait for updates from scheduled sources on all channels;
7 Broadcast an ACK message to indicate all successful updates;
8 Update channel cost as $\lambda_m(t+1) \leftarrow (1-\beta)\lambda_m(t) + \beta v_m(t)$, where $v_m(t)$ is the optimal dual variable associated with (14b) for channel type m in the MWM-PI problem at time t.

Remark. Clearly, Algorithm 1 is a generalization of Whittle's index policy. In fact, in the single-channel case, the MWM-PI reduces to Whittle's policy. The critical difference is that, in heterogeneous multi-channel systems, source n's index for channel type m depends on other channels' costs $\vec{\lambda}_{-m}$, whose optimal value also needs to be found. To address this difficulty, Algorithm 1 uses adaptive updates to approach the optimal channel costs in Line 8.

3.3 Fluid Analysis and Asymptotic Optimality

In the literature, the optimality of Whittle index policies is often shown using a fluid limit argument, by considering the regime of a large-scale system. Specifically, [20] shows that the difference between the state distribution under the Whittle index policy and the steady-state distribution under the optimal policy for the relaxed problem (7) diminishes to zero, when $N, \overline{C} \to \infty$ and $\alpha = \overline{C}/N$ is kept constant. Similarly, in this section, we will focus on such a fluid limit. We will show that the fixed point of the MWM-PI problem is equivalent to that of the relaxed problem (i.e., when dual gradient descent on $\vec{\lambda}$ converges and when the steady-state distribution is reached). Since the optimal solution for the relaxed problem at the fixed point is a lower bound for the original MDP (6), the above-mentioned equivalence relationship is essential for establishing the asymptotic optimality of our proposed SWIM policy later.

We first define the fluid limit model of the relaxed problem and its fixed point as follows. For any group g, let z_{gs} be the fraction of sources of group g that is in state s, with $\sum_{s \in \mathcal{S}} z_{gs} = 1$. Thus, $\vec{z}_g \overset{\Delta}{=} [z_{gs}, s \in \mathcal{S}]$ denotes the state distribution vector of group g. Given the current cost vector $\vec{\lambda}$, we assume that the distribution under the relaxed policy has reached the steady state. Let $x_{gs}^u \in [0,1]$ be the fraction of sources of state s in group g that is assigned to channel u by the relaxed policy $\boldsymbol{\pi}_{\text{rel}}$. We use $(\vec{x}, \vec{z}^*, \vec{\lambda}^*)$ to denote a fluid fixed point of the relaxed problem at steady state (i.e., when the dual gradient ascent on $\vec{\lambda}$ converges). Similar to the fluid analysis in [18], we can verify that, at the fixed-point channel cost $\vec{\lambda}^*$, $(\vec{x}, \vec{z}^*)$ also solves the following fluid problem (which is a linear program (LP))

$$\underset{\vec{x}, \vec{z}}{\text{minimize}} \quad \sum\nolimits_{g \in \mathcal{G}} \sum\nolimits_{s \in \mathcal{S}} \sum\nolimits_{u=0}^{M} z_{gs} C_{gs}^u x_{gs}^u \tag{15a}$$

subject to

$$\sum\nolimits_{g \in \mathcal{G}} \sum\nolimits_{s \in \mathcal{S}} z_{gs} x_{gs}^u \le \overline{C}, \quad \forall u \in \mathcal{M}, \tag{15b}$$

$$\sum\nolimits_{u \in \mathcal{M}} x_{gs}^u \le 1. \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S}, \tag{15c}$$

$$\sum_{u \ge 0} \sum_{d \in \mathcal{S}} z_{gs} x_{gs}^u p_{sd}^g(u) = \sum_{u' \ge 0} \sum_{d \in \mathcal{S}} z_{gd} x_{gd}^{u'} p_{ds}^g(u'), \forall s \in \mathcal{S}, \forall g \in \mathcal{G}, \tag{15d}$$

where C_{gs}^u and $p_{sd}^g(u)$ are defined in (10) (recall that $u = 0$ corresponds to passive). Thus, the primal and dual variables $(\vec{x}, \vec{z}^*, \vec{\lambda}^*)$ will satisfy the KKT conditions of (15). Similar to [18, Lemma 4.3], it can be shown that the optimal solution of the fluid problem (15), denoted as $V^*(\vec{x}, \vec{z}^*, \vec{\lambda}^*)$, is a lower bound for the original MDP.

For Algorithm 1, we can similarly define its fluid limit and fixed point as follows. Suppose that the steady state is reached. Denote the corresponding state distribution, channel cost vector and decision vector as $\vec{z}', \vec{\lambda}'$ and $\vec{y}$, respectively. Recall that MWM-PI is based on a set of dual costs $\vec{\lambda}$, and the edge weight is computed by $w_{gs}^m = I_m^g(s, \vec{\lambda}_{-m})$ for $m \in \mathcal{M}$, and $w_{gs}^0 = I_0^g(s, \vec{\lambda})$. At steady state, $(\vec{y}, \vec{z}', \vec{\lambda}')$

must solve the following fluid problem

$$\underset{\vec{y}}{\text{maximize}} \quad \sum_{g\in\mathcal{G}}\sum_{s\in\mathcal{S}}\sum_{u=0}^{M} z_{gs}w_{gs}^{u}y_{gs}^{u} \tag{16a}$$

$$\text{subject to} \quad \sum_{g\in\mathcal{G}}\sum_{s\in\mathcal{S}} z_{gs}y_{gs}^{u} \le \overline{C}, \quad \forall u \in \mathcal{M}, \tag{16b}$$

$$\sum_{u=0}^{M} y_{gs}^{u} = 1, \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad , \tag{16c}$$

$$\text{replace } \vec{x} \text{ by } \vec{y} \text{ in (15d).} \tag{16d}$$

Denote the Lagrange multiplier associated with (16b) as ν_u (define $\nu_0 = 0$). At the fixed point, $\vec{\lambda}' = \vec{\nu}$ must hold. Thus, we denote such $(\vec{y}, \vec{z}', \vec{\lambda}')$ as the fixed point of Algorithm 1, which should also satisfy the KKT conditions of (16).

Ideally, our goal is to show that the fixed point of the relaxed problem is identical to that of Algorithm 1, and thus they produce the same near-optimal objective value. However, we need a slightly stronger condition than partial indexability, as follow.

Definition 3.4 (Precise Division). *Given state space $\mathcal{S}$ and channel costs $\vec{\lambda}$, suppose that the sub-problem* (10) *is partially indexable with index $I_m^g(s, \vec{\lambda}_{-m})$ for $s \in \mathcal{S}$. We say that the preference for channel-type m is precisely divisible by its partial-index $I_m^g(s, \vec{\lambda}_{-m})$, if the following holds: for all $s \in \mathcal{S}$ and $m \ge 1$,*

(i) If $I_m^g(s, \vec{\lambda}_{-m}) = \lambda_m$, then $\mu_m^g(s, \vec{\lambda}) \le \mu_u^g(s, \vec{\lambda}), \forall u \ne m, u \ge 0$.
(ii) If $I_m^g(s, \vec{\lambda}_{-m}) > \lambda_m$, then $\mu_m^g(s, \vec{\lambda}) < \mu_u^g(s, \vec{\lambda}), \forall u \ne m, u \ge 0$.
(iii) Otherwise, there exists $u \ne m, u \ge 0$ s.t. $\mu_m^g(s, \vec{\lambda}) > \mu_u^g(s, \vec{\lambda})$.

Note that Definition 3.4 implies partial indexability in Definition 3.2. To see this, note that given $\vec{\lambda}_{-m}$, for any state $s \in \mathcal{S}$, its partial index $I_m^g(s, \vec{\lambda}_{-m})$ is independent of λ_m. Thus, as λ_m increases, we transition from $I_m^g(s, \vec{\lambda}_{-m}) > \lambda_m$ (i.e., using channel type m per Definition 3.4-(ii)) to $I_m^g(s, \vec{\lambda}_{-m}) < \lambda_m$ (i.e., not using channel type m per Definition 3.4-(iii)). In other words, as λ_m increases, $\mathcal{P}_m^g(\vec{\lambda})$ increases monotonically to the entire state space $\mathcal{S}$. On the other hand, Definition 3.4 is stronger than partial indexability because it states that this transition occurs precisely at $I_m^g(s, \vec{\lambda}_{-m}) = \lambda_m$.

Condition 3.5. *The sub-problem* (10) *satisfies the precise division property in Definition 3.4 (which implies partial indexability).*

The next theorem, which is one of our main contributions in this work, establishes the connection between the fixed point of the relaxed problem (7) and the fixed point of Algorithm 1.

Theorem 3.6. *Suppose that Condition 3.5 holds. Then, any fixed-point solutions $\{\vec{x}, \vec{z}^*, \vec{\lambda}^*\}$ of the relaxed problem are equivalent to the fixed-point solutions $\{\vec{y}, \vec{z}', \vec{\lambda}'\}$ of Algorithm 1 at the fluid limit.*

Proof. To prove the equivalence in the theorem, we need to show the statement in both direction. We first prove the "$\Longleftarrow$" direction. Given $(\vec{y}, \vec{z}', \vec{\lambda}')$, our goal is to show that, by letting $\vec{x} = \vec{y}, \vec{z}^* = \vec{z}', \vec{\lambda}^* = \vec{\lambda}'$, the KKT conditions of (16) will imply the KKT conditions of (15). Notice that (15) and (16), and thus their corresponding KKT conditions have very similar forms (see details in our technical report [22]), except for the following conditions regarding optimizing the Lagrangian in (15) and (16), respectively.

(I) *[The relaxed problem* (15)*]:* For each $(g, s) \in \mathcal{G} \times \mathcal{S}$, $[x_{gs}^u]$ corresponds to the optimal decisions for the sub-problem (8), given $\vec{\lambda}^*$. This means that, if $x_{gs}^u > 0$ for some u, it must be true that $\mu_u^g(s, \vec{\lambda}^*) \le \mu_{u'}^g(s, \vec{\lambda}^*), \forall u' \ne u, 0 \le u' \le M$, where $\mu_u^g(s, \vec{\lambda}^*)$ is the expected cost-to-go of action u in (10) under $\vec{\lambda}^*$.
(II) *[The SWIM policy* (16)*]:* Recall that $\vec{\nu}$ is the Lagrange multiplier for (16b), which is equal to $\vec{\lambda}'$ at the fixed point. For each $(g, s) \in \mathcal{G} \times \mathcal{S}$, $[y_{gs}^u]$ should maximize dual objective $\sum_{u=0}^{M} y_{gs}^u(w_{gs}^u - \nu_u)$, subject to the constraint $\sum_{u=0}^{M} y_{gk}^u = 1$. Denote $J^{\max} \triangleq \{u | w_{gs}^u - \nu_u \ge w_{gs}^{u'} - \nu_{u'}, \forall u' \ne u, u \ge 0\}$. Then, we must have $\sum_{u \in J^{\max}} y_{gs}^u = 1$, and $y_{gs}^{u'} = 0$ for $u' \notin J^{\max}$.

To show "$\Longleftarrow$", the key is to show that (II) implies (I). Before we proceed, we first state a corollary and a lemma as follows.

Corollary 3.7. *Suppose that Condition 3.5 holds. For any state $d \in \mathcal{S}$, suppose that there exists one channel-type m such that $I_m^g(d, \vec{\lambda}_{-m}) > \lambda_m$. Then, the other channels $u \ne m$ must have $I_u^g(d, \vec{\lambda}_{-u}) < \lambda_u$.*

Proof. By Definition 3.4-(ii), since $I_m^g(d, \vec{\lambda}_{-m}) > \lambda_m$, we must have $\mu_u^g(d, \vec{\lambda}) > \mu_m^g(d, \vec{\lambda})$ for all channels $u \ne m$. Suppose in contrary that $I_u^g(d, \vec{\lambda}_{-u}) \ge \lambda_u$ for $u \ne m$. Then, we would have $\mu_m^g(d, \vec{\lambda}) \ge \mu_u^g(d, \vec{\lambda})$ by Definition 3.4-(i),(ii), which is a contradiction. ■

Lemma 3.8. *In Condition (II) for the SWIM policy, at least one $u \ge 0$ should satisfy $w_{gs}^u \ge \nu_u$ at the fixed point.*

Corollary 3.7 implies that there can be at most one channel-type m with $I_m^g(d, \vec{\lambda}_{-m}) > \lambda_m$. Lemma 3.8 is also intuitive. Suppose in contrary that $w_{gs}^u < \nu_u$ (i.e., $I_u^g(s, \vec{\lambda}'_{-u}) < \lambda'_u$) for all $u \ge 0$. By (12), it implies that action-0 would have been optimal for s. In that case, $I_0^g(s, \vec{\lambda}') = 0 = \nu_0$ by (13) at the fixed point $\vec{\nu} = \vec{\lambda}'$, which is a contradiction. Thus, Lemma 3.8 must hold (see [22] for details).

Now, suppose that $(\vec{y}, \vec{z}', \vec{\lambda}')$ satisfies (II). To show that $(\vec{x}, \vec{z}^*, \vec{\lambda}^*) = (\vec{y}, \vec{z}', \vec{\lambda}')$ satisfies (I), we divide into two cases (note that Lemma 3.8 implies $\max_{u \ge 0}\{w_{gs}^u - \nu_u\} \ge 0$).

a) (When $\max_{u \ge 0}\{w_{gs}^u - \nu_u\} > 0$) From condition (II), we have $\sum_{u \in J^{\max}} y_{gs}^u = 1$, where $J^{\max}$ contains all actions $u \ge 0$ that attain the maximum of $\{w_{gs}^u - \nu_u\}$. By the definition in (13), channel-0's index is always non-positive. Hence, we must have $0 \notin J^{\max}$. Thus, $y_{gs}^0 = 0$. From Corollary 3.7, there can exist only one $u \ge 1$ such that $w_{gs}^u = I_u^g(s, \vec{\nu}_{-u}) > \nu_u$, in which case $w_{gs}^{u'} = I_{u'}^g(s, \vec{\nu}_{-u'}) < \nu_{u'}, \forall u' \ne u$. Define u^* as the unique index such that $u^* = \arg\max_u\{w_{gs}^u - \nu_u\}$. Thus, $y_{gs}^{u^*} = 1$. By Definition 3.4-(ii), we then have $\mu_{u^*}^g(s, \vec{\nu}) < \mu_{u'}^g(s, \vec{\nu}), \quad \forall u' \ne u^*, 0 \le u' \le M$. According to (10), $x_{gs}^{u^*} = y_{gs}^{u^*} = 1$ and $x_{gs}^{u'} = 0, \forall u' \ne u^*, u' \ge 0$ correspond to the optimal solutions for (8). Hence, (I) holds.
b) (When $\max_{u \ge 0}\{w_{gs}^u - \nu_u\} = 0$) Again, from condition (II), we have $\sum_{u \in J^{\max}} y_{gs}^u = 1$. Thus, $\sum_{u \in J^{\max}} x_{gs}^u = 1$ by our construction. The definition of $J^{\max}$ implies that for all $u \in J^{\max}, u > 0$, we have $w_{gs}^u = I_u^g(s, \vec{\nu}_{-u}) = \nu_u$. Further, for all $u' \notin J^{\max}$, we have $w_{gs}^{u'} = I_{u'}^g(s, \vec{\nu}_{-u'}) < \nu_{u'}$. Thus, from Definition 3.4, every $u \in J^{\max}$ is an optimal action for the sub-problem (8), and every $u' \notin J^{\max}$ is not. The only question is whether the optimal

action for (I) should use action-0 or not. Next, we divide into two sub-cases. If $0 \notin J^{\max}$, then $I_0^g(s, \vec{v}) < 0$. By the definition in (13), $x_{gs}^0 = 0$ must hold for the sub-problem (8). Thus, the decisions $\sum_{u \in J^{\max}} x_{gs}^u = 1$ is optimal for (8). On the other hand, if $0 \in J^{\max}$, then $I_0^g(k, \vec{v}) = 0$. By definition of index (13), for any $\epsilon_{0u} > 0$, we must have $\mu_0^g(s, \vec{v}) - \epsilon_{0u} \leq \mu_u^g(s, \vec{v}), \forall u \geq 1$. Letting $\epsilon_{0u} \to 0$, we then have $\mu_0^g(s, \vec{v}) \leq \mu_u^g(s, \vec{v}), \forall u \geq 1$. Thus, $x_{gs}^0 > 0$ is also optimal for the sub-problem (8). Combining the two sub-cases, the decision $\sum_{u \in J^{\max}} x_{gs}^u = 1$ is always optimal for the sub-problem (8). Thus, condition (I) follows.

Combing the above cases, condition (I) must hold for the fixed point of the SWIM policy. Hence, we have shown the "$\Longleftarrow$" direction. Due to space limit, we omit the proofs for other KKT conditions and for the "$\Longrightarrow$" direction. Readers are referred to our technical report [22] for the complete proof. □

Remark. Theorem 3.6 establishes an important connection between the fixed point of the relaxed problem and the fixed point of Algorithm 1. In contrast to the relaxed problem, the solution for MWM-PI naturally respects the instantaneous resource constraint (6a). Since its fixed point still achieves the optimal performance at the fluid limit, it provides useful guidance for proving the asymptotic optimality of our proposed SWIM policy.

Next, we evaluate the performance of Algorithm 1. We first state the following technical condition called "global attractor" [18].

Definition 3.9 (Global attractor). *An equilibrum point $\vec{X}^*$ is a global attractor for a process $X(t)$ if, for any initial point $\vec{X}(0)$, the process $X(t)$ converges to $\vec{X}^*$.*

Next, we assume that a fixed point of Algorithm 1 satisfies the global attractor property. Notice that similar assumption has been made in [5, 18, 20]. As mentioned in [18], in general, it may be difficult to establish analytically that a fixed point is a global attractor for the process; thus, such property is only verified numerically. Our simulation results in Section 5.1 indeed show that such convergence indeed happens for our proposed policy.

Based on this condition, we then show the asymptotic optimality of the SWIM policy. Specifically, we will consider the original MDP (6) in the following r-scaled system: we scale by r both the number of sources in each group, and the number of channels of each type, i.e., $N^r = rN$ and $\overline{C}^r = r\overline{C}$, while keeping $\alpha = \overline{C}^r / N^r$ a constant. The transition probabilities for each source remain unchanged. For such a r-scaled system, we define V_{SWIM}^r as the average cost per stage in the objective of (6) under our proposed SWIM policy. (Note that V_{SWIM}^r in (6) is already averaged by the number of sources N^r.)

Theorem 3.10 (Asymptotic Optimality). *Suppose that Condition 3.5 holds for the sub-problem* (10). *Suppose that a fixed point $(\vec{y}, \vec{z}', \vec{\lambda}')$ of the policy $\boldsymbol{\pi}_{\text{SWIM}}$ in Algorithm 1 is a global attractor according to Definition 3.9. Then, $\boldsymbol{\pi}_{\text{SWIM}}$ is asymptotically optimal in minimizing the average cost per stage. Specifically, we have $\lim_{r\to\infty} V_{\text{SWIM}}^r = V^*$, where V^* is the optimal objective for the fluid relaxed problem* (15).

The proof of Theorem 3.10 follows from the global attractor property and Theorem 3.6. See detailed proof in [22].

4 AOI MINIMIZATION IN HETEROGENEOUS MULTI-CHANNEL SYSTEMS

In this section, we return to the setting of AoI minimization problem described in Section 2.1. Since the results in Section 3 is very general, we only need to verify that Definition 3.2 and Definition 3.4 indeed hold for the AoI setting. Then, the result of Theorem 3.10 and Algorithm 1 can be directly applied. As we will show soon, the verification of the indexability and the precise division property is highly non-trivial for sub-problem (8). Note that for single-channel systems, Whittle indexability has been verified for AoI minimization under the generate-at-will model [16]. However, the approach there is based on directly solving the value function, which appears to be infeasible for our heterogeneous multi-channel setting. Instead, we will develop new structural properties of the value function, based on which we will establish both partial indexability and the precise division property.

With this goal in mind, we assume that the dual costs $\vec{\lambda} = [\lambda_1, \ldots, \lambda_M]^T$ for all channels are given. Since all sub-problems (8) are independent, in the rest of the section, we omit the superscript g of the variables for the sub-problem (10) whenever no ambiguity occurs. As we mentioned in Sec. 2.2, the MDP of each sub-problem is an average cost per stage problem with infinite time horizon and countably infinite state space. For ease of notation, we define $\lambda_0 = 0$ and $p_0 = 0$. The corresponding Bellman Equation (10) for the AoI minimization problem described in Section 2.1 can be written as, for any state (i.e., current AoI) $d \in \mathbb{N}_+$,

$$f(d) + J^* \stackrel{\Delta}{=} \min_{m \in \{0\} \cup \mathcal{M}} \{\lambda_m + (1-p_m)[d + f(d+1)] + p_m f(1)\}. \quad (17)$$

Here we slightly abuse notation, and use $\mu_m(d) \stackrel{\Delta}{=} \mu_m(d, \vec{\lambda})$ to denote each term in the minimization on the RHS of (17) when the parameter $\vec{\lambda}$ is given above. Recall that $\mu_m(d, \vec{\lambda})$ is the expected cost-to-go under channel costs $\vec{\lambda}$ if channel m is selected. Since $d = 1$ is a recurrent state, we can set $f(1) = 0$. Note that Bellman equation (17) cannot be solved in closed form due to multiple heterogeneous channels. Specifically, the optimal action for the current state depends on the value functions of future states, which possibly have different optimal actions. This complex dependency is in sharp contrast to [14, 16], which only consider a single channel.

Even though the exact solution is unavailable, we can still derive useful structure properties from (17). Next, we define a property called "multi-threshold-type" (MTT), and prove that the optimal policy for the sub-problem (8) is indeed MTT.

Definition 4.1 (Multi-threshold-type). *A channel selection policy for the sub-problem* (8) *is MTT if the followings hold:*

(1) (Threshold-based) For any channels $\gamma, \psi \in \{0, 1, \ldots, M\}$ with $p_\gamma > p_\psi$, there exists $H^{\gamma,\psi} \geq 0$ such that $\mu_\gamma(d) \leq \mu_\psi(d)$, for all $d \geq H^{\gamma,\psi}$, and $\mu_\gamma(d) > \mu_\psi(d)$, for all $d < H^{\gamma,\psi}$.
(2) (Ordering of Channels) Suppose two states $d_1 < d_2$. Denote the optimal channels for d_1 and d_2 be $m^*(d_1)$ and $m^*(d_2)$, respectively. Then, $p_{m^*(d_1)} \leq p_{m^*(d_2)}$ must hold.
(3) (Channel Dominance) For any two channel types γ and ψ with $p_\gamma > p_\psi$, if $\frac{\lambda_\gamma}{p_\gamma} < \frac{\lambda_\psi}{p_\psi}$, then ψ is never the optimal channel for any state, i.e., $\mu_\gamma(d) \leq \mu_\psi(d)$ for all state d whenever $\mu_\psi(d) \leq \mu_0(d)$.

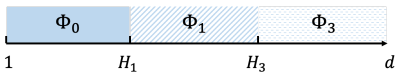

Figure 2: An illustration for MTT policy. Suppose a 3-channel system with $p_1 < p_2 < p_3$ and $\frac{\lambda_1}{p_1} < \frac{\lambda_3}{p_3} < \frac{\lambda_2}{p_2}$ (note that Channel 2 is dominated by Channel 3).

In other words, Condition 1 states that a threshold exists for any pair of channels, such that a better-quality channel is always preferred to the worse one when d is above the threshold. Condition 2 specifies that, as d increases, the optimal decision increasingly prefer more reliable channels. Condition 3 means that if a channel type ψ is less reliable and also "more expensive" than the other channel type γ, it should never be the optimal action. In that case, we say that channel ψ is dominated by channel γ. The next lemma shows that the optimal policy for the sub-problem (8) is MTT.

Lemma 4.2. *Given the cost vector $\vec{\lambda}$, the optimal policy $\boldsymbol{\pi}^*$ satisfying* (17) *for the sub-problem* (8) *is MTT.*

Lemma 4.2 is intuitive because, when the state (i.e., age) is higher, it is more urgent for the source to use a more reliable channel. See detailed proof in our technical report [22].

Fig. 2 illustrates a MTT policy in the state space $d \in \mathbb{N}_+$ of the sub-problem (8) with three channels. Define $\Phi_m \subset \mathbb{N}_+$ as the optimal decision region of m, i.e., m is the optimal channel type for all $d \in \Phi_m$. Thanks to Definition 4.1-(1), Φ_m must be contiguous for all m. We denote $H_m = \min_{d \in \Phi_m} d$ as the threshold for channel type m. Note that the optimal decision regions for some channels (Φ_2 is absent in Fig. 2) may be empty due to channel dominance. Before we proceed to the proof of partial indexability, we first prove the following lemma. Without loss of generality, we assume that all channels have distinct successful probabilities, and their qualities are arranged in an ascending order, i.e., $p_1 < \ldots < p_M$.

Lemma 4.3. *Given $\vec{\lambda}$, suppose $\vec{\lambda}' = [\lambda_1, \ldots, \lambda_m + \Delta, \ldots, \lambda_M]$, $\Delta > 0$. Denote the optimal value functions in* (17) *under $\vec{\lambda}$ (with the optimal policy $\boldsymbol{\pi}$) and under $\vec{\lambda}'$ (with the optimal policy $\boldsymbol{\pi}'$) as $f(\cdot)$ and $f'(\cdot)$, respectively. Then, the difference between two value functions can be upper-bounded by $f'(d) - f(d) < \frac{\Delta}{p_m}, \forall 1 \le m \le M$, and lower-bounded by $f'(d) - f(d) > -\frac{\Delta}{p_{m+1}}, \forall 1 \le m < M$.*

To prove Lemma 4.3, we use the equivalence relationship between the average cost per stage problem and the stochastic shortest path problem [2]. Due to space limit, we refer readers to our technical report [22] for details.

Proposition 4.4. *Given the cost vector $\vec{\lambda}$, the sub-problem* (8) *of heterogeneous multi-channel AoI minimization is partially indexable.*

Proof. To prove the proposition, it suffices to show: (i) If $d \in \mathcal{P}_m(\vec{\lambda})$, then $d \in \mathcal{P}_m(\vec{\lambda}')$ must hold for $\vec{\lambda}'$; and (ii) If $\lambda_m = \infty$, then $\mathcal{P}_m(\vec{\lambda}) = \mathcal{S}$. By Lemma 4.2, the optimal policy for sub-problem (8) is MTT. Then, the passive sets under $\vec{\lambda}$ can be expressed as $\mathcal{P}_m(\vec{\lambda}) = \{1, \ldots, H_m - 1\} \cup \{H_{m+1}, \ldots\}$ if $m < M$, and $\mathcal{P}_M(\vec{\lambda}) = \{1, \ldots, H_M - 1\}$. To show statement (i), it suffices to show that $\mathcal{P}_m(\vec{\lambda}) \subseteq \mathcal{P}_m(\vec{\lambda}')$, which is equivalent to showing that $H_m \le H'_m, \forall m \le M$ and $H_{m+1} \ge H'_{m+1}, \forall m < M$.

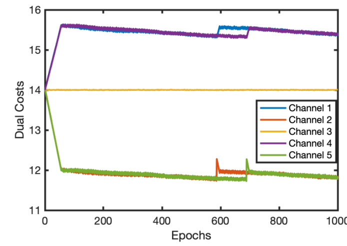

Figure 3: Dual costs update of π_{rel} for the relaxed problem.

We first show $H_m \le H'_m, \forall m \le M$. Suppose in contrary that $H_m > H'_m$. By the definition of H_m, at state H'_m policy $\boldsymbol{\pi}$ must prefer another action $u, 0 \le u < m$, over channel m, i.e., $\mu_u(H'_m, \vec{\lambda}) \le \mu_m(H'_m, \vec{\lambda})$. By (17), this implies that (noting $f(1) = 0$)

$$(p_m - p_u)[H'_m + f(H'_m + 1)] \le \lambda_m - \lambda_u. \quad (18)$$

Similarly, given $\vec{\lambda}'$, policy $\boldsymbol{\pi}'$ must prefer m over other channels at H'_m, i.e., $\mu_m(H'_m, \vec{\lambda}') \le \mu_u(H'_m, \vec{\lambda}')$. By (17), this implies that

$$(p_m - p_u)[H'_m + f'(H'_m + 1)] \ge \lambda'_m - \lambda_u. \quad (19)$$

Recall that $\lambda'_m = \lambda_m + \Delta$. From (18) and (19), we have $f'(H'_m + 1) - f(H'_m + 1) \ge \frac{\Delta}{p_m - p_u} \ge \frac{\Delta}{p_m}$. Clearly, this contradicts with Lemma 4.3 that $f'(d) - f(d) < \frac{\Delta}{p_m}$. Thus, $H_m \le H'_m, \forall 1 \le m \le M$ must hold. Similarly, we can show $H_{m+1} \ge H'_{m+1}, \forall m < M$ using the lower bound in Lemma 4.3. The proof of (ii) is also straightforward. Due to space limit, we refer readers to [22] for the complete proof. □

Proposition 4.5. *Given state space $\mathcal{S}$ and channel costs $\vec{\lambda}$, the sub-problem* (17) *satisfies the precise division property in Definition 3.4.*

The proof is similar to that of Prop. 4.4, and is available in [22].

To summarize this section, we briefly comment on the complexity of SWIM policy. Note that the partial index incurs a higher computational complexity than Whittle's index [21], as it needs to be recomputed for every $\vec{\lambda}$. In practice, it may be more feasible to precompute the partial indices for a quantized subset of $\vec{\lambda}$, and then use linear interpolation to approximate the partial indices in real-time. Note that such precomputed partial indices do not need to be re-calculated even if the user population changes, which is a benefit similar to Whittle's index. For a fixed number of channel types and quantization levels (for $\vec{\lambda}$), the complexity of our solution grows linearly in user population, compared to exponential growth for the value iteration approach. We leave for future work how to further reduce the complexity of the partial index computation.

5 NUMERICAL RESULTS

In this section, we present MATLAB simulation results of our proposed SWIM policy in Algorithm 1. Specifically, we focus on the AoI minimization problem in Section 4 for heterogeneous multi-channel systems. We simulate an information update system with N=50 data sources, which are divided into G=5 groups, with 15, 5, 10, 15 and 5 sources in each group 1 to 5, respectively. Further, we assume that there are M=5 types of channels, each of which is equipped with $\overline{C}$=2 identical instances. To model the heterogeneity and preferences, the channel success probabilities of source-channel pairs

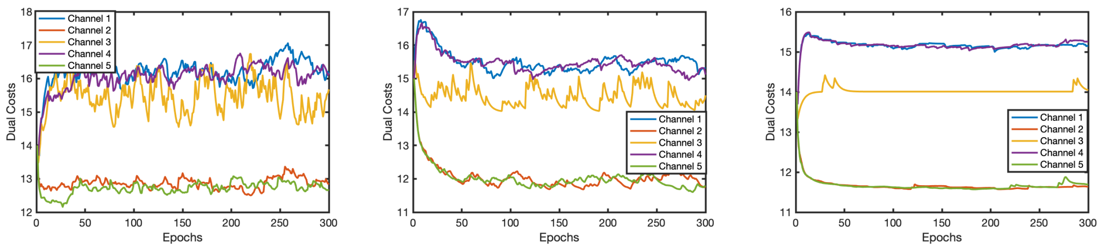

(a) System scale $r = 1$ **(b) System scale $r = 3$** **(c) System scale $r = 10$**

Figure 4: Dual costs update of the proposed index-based policy π_{SWIM} for different system scales.

are different across groups. For simplicity, we assume that Group-g sources prefer Type-g channels, where $g = 1, \dots, 5$. Specifically, the channel quality vector for group-1 is $\vec{p_1} = [0.9, 0.7, 0.5, 0.3, 0.1]$. Then, $\vec{p_g}$ for group $g > 1$ is obtained by circularly right-shifting $\vec{p_1}$ by $g-1$ positions. (Note that the entire system is asymmetric due to the uneven source population in each group.) To compare the scaling performance, we simulate on r-scaled systems that multiplies the number of the sources and channels by r, i.e., $N^r = rN$ and $\overline{C}^r = r\overline{C}$. The simulation time is divided into epochs, each of which consists of 50 discrete time-slots. For a smoother update, we re-compute the channel costs at the end of each epoch based on Line 8 of Algorithm 1 using the averaged $\overline{\nu}_m(t)$ of current epoch.

5.1 Convergence of the Channel-Cost Update

First, we compare the dynamics of the cost updates of the SWIM policy with that of the relaxed policy π_{rel}, and verify that their fixed points indeed match. Fig. 3 illustrates the cost dynamics of all channels for the relaxed problem, with respect to the number of epochs. At each time, all sources independently compute their optimal actions (8) based on their states and current cost $\vec{\lambda}(t)$. At the end of each epoch, the BS performs a dual gradient update according to (9). We simulate for $T_{\text{rel}} = 1000$ epochs for π_{rel} to reach the fixed-point channel costs. In Fig. 3, all channels' costs converge to a small neighborhood of the optimal costs $\vec{\lambda}^*$ of the relaxed problem after $T_{\text{rel}} = 1000$ epochs.

Next, we verify that the channel-cost dynamics under π_{SWIM} approach that of the relaxed problem when the system scale is large. Specifically, we let $\beta = 0.2$ in Algorithm 1. Denote the fixed-point channel cost vector under SWIM policy for system scale r as $\vec{\lambda}_r^{\text{SWIM}}$. We simulate π_{SWIM} for $T_{\text{SWIM}} = 300$ epochs under different system scales. Fig. 4(a)-(c) show the channel-cost dynamics for the system at scale $r = 1$, $r = 3$ and $r = 10$, respectively. Clearly, we can observe that, as the system scale r increases, $\vec{\lambda}_r^{\text{SWIM}}$ approaches very close to the values of $\vec{\lambda}^*$ in Fig. 3. The convergence is more obvious for $r = 10$ (Fig. 4(c)), which confirms the result of Theorem 3.6, i.e., the fixed point solution $\vec{\lambda}'$ of π_{SWIM} is equivalent to the optimal cost $\vec{\lambda}^*$ of the relaxed problem in the fluid limit.

5.2 Average System AoI

Next, we evaluate the average AoI performance of our proposed policy. We will use the solution for the relaxed problem as a performance lower bound for comparison. In addition, we will compare SWIM policy with the following scheduling policies that satisfy the instantaneous constraints in (6).

Rounded Relaxed Policy (RRP). As we discussed in Section 2.2, although the optimal solution for the relaxed problem under $\vec{\lambda}^*$ provides a lower bound for the original AoI minimization problem, π_{rel} may violate the instantaneous hard constraints (6a). To satisfy feasibility, RRP is deduced from π_{rel} with the following modification. (Note that we did not use the priority policy in [18], since it only works for single-channel systems and cannot be applied here.)

1. (Over-subscription) For any channel, if the number of transmitting sources exceeds (or equals to) the number of channel instance $\overline{C}^r$, RRP schedules $\overline{C}^r$ sources uniformly at random;
2. (Under-subscription) Otherwise, RRP schedules additional sources with largest AoI to reach $\overline{C}^r$ total sources for the channel.

Max-Age Matching (MAM) [17]. This policy was originally proposed for systems with multiple ON/OFF channels in [17]. As the name suggests, MAM attempts to serve sources with high AoI values at each time. The MAM scheduler in [17] requires knowledge of whether a channel is ON/OFF. For a fair comparison with our policy that does not require such knowledge, we take all channels with non-zero success probability as being ON. Then, we form a bipartite graph between all pairs of sources and the channel types, with the weights given by the current AoI of the sources. The scheduling problem is then mapped to a bipartite graph matching problem. Note that this policy ignores the exact channel success probability, and thus is expected to have poorer performance.

Fig. 5a shows the total system AoI dynamics under different policies for system scale $r = 7$. First, we can see that the total system AoI (about 1200) under our proposed SWIM policy is very close to the performance lower bound obtained from the relaxed problem (the lowest two curves). This observation verifies our result in Theorem 3.10 on the asymptotic optimality of our proposed policy. In contrast, the AoI of the rounded relaxed policy (RRP) is over 1400, which is about 20% worse than that of SWIM policy. This performance degradation suggests that it may not be efficient to use the solution from the relaxed problem even for medium-scaled systems. Finally, the AoI under the MAM policy is about 2000, which is 65% worse than SWIM policy. The result is not surprising, as the MAM policy simply ignores the channel heterogeneity.

Next, Fig. 5b shows the average total system AoI at different system scales. For all three policies, the average total system AoI scales almost linearly with the system scale r. Again, we observe that our

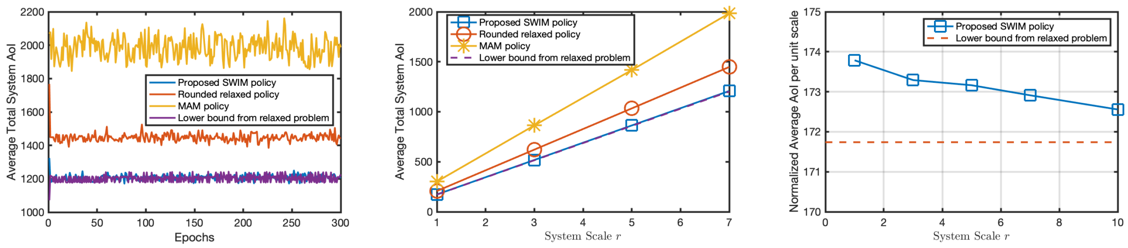

(a) Total system AoI evolution for scale $r = 7$. (b) Average total system AoI v.s. system scales. (c) Normalized AoI v.s. system scales.

Figure 5: Performance comparison among different scheduling policies.

proposed SWIM policy achieves close-to-optimal AoI performance under all simulated system scales. Finally, Fig. 5c shows the normalized average AoI, i.e., the total average system AoI divided by the scale parameter r, under our proposed SWIM policy. Clearly, as r increases, the normalized average AoI of our SWIM policy approaches closer to the lower bound obtained from the relaxed problem. This observation is again consistent with our result in Section 3.3 on the asymptotic optimality of the SWIM policy.

Finally, we report the running time of our algorithm on a MacBook Pro with 2.3 GHz Intel Core i5 CPU and 16GB RAM. For each epoch of 50 time slots and at system scale $r = 1$, the running time for partial-index computation and MWM-computation are $347.93s$ and $0.32s$, respectively. Thus, the average running time per time slot for SWIM policy is $6.96s$, compared to $0.47s$ for RRP policy and $0.07s$ for MAM policy.

6 CONCLUSION

This work studies the problem of minimizing AoI in heterogeneous multi-channel systems. We formulate the problem as an infinite-horizon constrained MDP. Existing results on Whilttle index cannot be applied to such a system with heterogeneous channels. Instead, we introduce a new notion of partial indexability. Then, we propose a new scheduling policy, i.e., SWIM policy, based on partial index and MWM. Under suitable conditions, the SWIM policy asymptotically optimizes the total expected AoI of the system, when the system scale is large. For the first time in the literature, such low-complexity and asymptotically optimal policies are developed for weakly coupled MDP with multiple heterogeneous resources. The simulation results demonstrate near-optimal AoI performance that outperforms other multi-channel scheduling policies in the literature. For future work, we will study more-efficient computation of partial index, and other settings, e.g., with stochastic arrivals.

ACKNOWLEDGEMENT

This work was supported in part by the National Science Foundation under awards CNS-1703014 and CNS-1717493, in part by IEDC grant A293-9-21Fund-10245, and in part by ARL grant W911NF2020221.

REFERENCES

[1] M. A. Abd-Elmagid and H. S. Dhillon. 2019. Average Peak Age-of-Information Minimization in UAV-Assisted IoT Networks. *IEEE TVT* 68, 2 (2019), 2003–2008.
[2] Dimitri P. Bertsekas. 2007. *Dynamic Programming and Optimal Control.* Vol. 2. Athena Scientific Belmont, MA.
[3] John Gittins, Kevin Glazebrook, and Richard Weber. 2011. *Multi-armed bandit allocation indices.* John Wiley & Sons.
[4] Y. Gu, H. Chen, Y. Zhou, Y. Li, and B. Vucetic. 2019. Timely status update in internet of things monitoring systems: An age-energy tradeoff. *IEEE IoT J.* 6, 3 (2019), 5324–5335.
[5] D. J. Hodge and K. D. Glazebrook. 2015. On the asymptotic optimality of greedy index heuristics for multi-action restless bandits. *Adv. App. Prob.* 47, 3 (2015), 652–667.
[6] Yu-Pin Hsu. 2018. Age of information: Whittle index for scheduling stochastic arrivals. In *Proc. IEEE Int. Symp. Inf. Theory.* IEEE, 2634–2638.
[7] Yu-Pin Hsu, Eytan Modiano, and Lingjie Duan. 2019. Scheduling algorithms for minimizing age of information in wireless broadcast networks with random arrivals. *IEEE Trans. Mobile Comput.* (2019).
[8] Igor Kadota and Eytan Modiano. 2019. Minimizing the Age of Information in Wireless Networks with Stochastic Arrivals. In *ACM Mobihoc '19.* 221–230.
[9] Igor Kadota, Abhishek Sinha, Elif Uysal-Biyikoglu, Rahul Singh, and Eytan Modiano. 2018. Scheduling policies for minimizing age of information in broadcast wireless networks. *IEEE/ACM Trans. Netw.* 26, 6 (2018), 2637–2650.
[10] S. Kaul, R. Yates, and M. Gruteser. 2012. Real-time status: How often should one update?. In *Proc. IEEE INFOCOM 2012 - IEEE Conf. on Computer Comm.* 2731–2735.
[11] A. Nedic and A. Ozdaglar. 2008. Subgradient methods in network resource allocation: Rate analysis. In *42nd Ann. Conf. Inform. Sci. and Sys.* 1189–1194.
[12] Zhenzhi Qian, Fei Wu, Jiayu Pan, Kannan Srinivasan, and Ness B. Shroff. 2020. Minimizing Age of Information in Multi-channel Time-sensitive Information Update Systems. In *Proc. IEEE INFOCOM 2020 - IEEE Conf. on Comput. Commun.*
[13] B. Sombabu and S. Moharir. 2020. Age-of-Information Based Scheduling for Multi-Channel Systems. *IEEE Trans. on Wireless Comm.* 19, 7 (2020), 4439–4448.
[14] Jingzhou Sun, Zhiyuan Jiang, Bhaskar Krishnamachari, Sheng Zhou, and Zhisheng Niu. 2019. Closed-form Whittle's index-enabled random access for timely status update. *IEEE Trans. Commun.* 68, 3 (2019), 1538–1551.
[15] Yin Sun, Elif Uysal-Biyikoglu, Roy D Yates, C Emre Koksal, and Ness B Shroff. 2017. Update or wait: How to keep your data fresh. *IEEE Trans. Inf. Theory* 63, 11 (2017), 7492–7508.
[16] Vishrant Tripathi and Eytan Modiano. 2019. A Whittle index approach to minimizing functions of age of information. In *Allerton Conf. 2019.* IEEE, 1160–1167.
[17] V. Tripathi and S. Moharir. 2017. Age of Information in Multi-Source Systems. In *GLOBECOM 2017 - 2017 IEEE Global Communications Conference.* 1–6.
[18] Ina Maria Verloop. 2016. Asymptotically optimal priority policies for indexable and nonindexable restless bandits. *Ann. Appl. Probab.* 26, 4 (2016), 1947–1995.
[19] P. Viswanath, D. N. C. Tse, and R. Laroia. 2002. Opportunistic beamforming using dumb antennas. *IEEE Trans. Inf. Theory* 48, 6 (June 2002), 1277–1294.
[20] Richard R Weber and Gideon Weiss. 1990. On an index policy for restless bandits. *J. Appl. Probab.* (1990), 637–648.
[21] Peter Whittle. 1988. Restless bandits: Activity allocation in a changing world. *J. Appl. Probab.* (1988), 287–298.
[22] Y. Zou, K. Kim, X. Lin, and M. Chiang. 2020. *Minimizing AoI in Heterogeneous Multi-Channel Systems: A New Partial-Index Approach.* Technical Report. Purdue University, https://engineering.purdue.edu/%7elinx/papers.html.
[23] Yihan Zou, Kwang Taik Kim, Xiaojun Lin, Mung Chiang, Zhi Ding, Risto Wichman, and Jyri Hämäläinen. 2019. Low-overhead Multi-antenna-enabled Uplink Random Access for Massive Machine-type Communications with Low Mobility. In *Proc. IEEE Global Commun. Conf. (GLOBECOM).* Waikoloa, HI, 1–6.
[24] Yihan Zou, Kwang Taik Kim, Xiaojun Lin, Mung Chiang, Zhi Ding, Risto Wichman, and Jyri Hämäläinen. 2020. Low-Overhead Joint Beam-Selection and Random-Access Schemesfor Massive Internet-of-Things with Non-Uniform Channel and Load. In *Proc. IEEE Conf. on Comput. Commun. (INFOCOM).* 2126–2135.