

Effective Communication: When to Pull Updates?

Pouya Agheli*, Nikolaos Pappas[†], Petar Popovski[‡], and Marios Kountouris*

*Communication Systems Department, EURECOM, Sophia Antipolis, France

[†]Department of Computer and Information Science, Linköping University, Linköping, Sweden

[‡]Department of Electronic Systems, Aalborg University, Aalborg, Denmark

Email: pouya.agheli@eurecom.fr, nikolaos.pappas@liu.se, petarp@es.aau.dk, marios.kountouris@eurecom.fr

Abstract—We study a pull-based communication system where a sensing agent updates an actuation agent using a query control policy, which is adjusted in the evolution of an observed information source and the usefulness of each update for achieving a specific goal. For that, a controller decides whether to pull an update at each slot, predicting what is probably occurring at the source and how much effective impact that update could have at the endpoint. Thus, temporal changes in the source evolution could modify the query arrivals so as to capture important updates. The amount of impact is determined by a *grade of effectiveness* (GoE) metric, which incorporates both freshness and usefulness attributes of the communicated updates. Applying an iterative algorithm, we derive query decisions that maximize the long-term average GoE for the communicated packets, subject to cost constraints. Our analytical and numerical results show that the proposed query policy exhibits higher effectiveness than existing periodic and probabilistic query policies for a wide range of query arrival rates.

I. INTRODUCTION

The effectiveness problem in communication systems deals with whether or not a message conveyed by a sender, leads to a desirable impact at the receiver for achieving a specific goal. This has been first articulated in [1], and has recently been revived under the prism of goal-oriented semantic communication [2], [3]. In cyber-physical systems with interacting sensing entities and actuation/monitoring agents, a message ought to be generated and communicated *if* it can potentially have the desired impact in the system. In this setting, the communication goal determines the *grade of effectiveness* a message has according to its usefulness or importance in fulfilling a set of attributes required for achieving that goal. This approach has the potential to enable system scalability and judicious use of resources by avoiding the acquisition, processing, and transport of information that turns out to be ineffective, irrelevant, or useless.

In current networked intelligent systems, information transfer, e.g., in the form of status update packets, over the network is mainly done using a *push-based* communication model. Packets are pushed toward the endpoint according to the source’s decision, regardless of what and when the endpoint

actually needs. On the other hand, we have the *pull-based* communication framework, where the endpoint requests and controls the type of the generated information and its arrival time [4]–[8]. To this end, existing query policies overlook the evolution of the source (what is probably happening on the sensing side) and the expected importance or usefulness of the updates at the time of pulling them. As a result, query-based policies, even if they are aligned to the receiver’s objectives, do not necessarily result in high system effectiveness. To address this challenge and to bridge the pull-based with the push-based world, we introduce a novel policy where the endpoint pulls updates based on the statistics of the source and the previously received updates. Thanks to this policy, the endpoint timely adapts its query instances according to the received updates, hence being able to capture updates (e.g., critical source events, alarms, novelty, etc.) that have high impact for achieving the subscribed goal. In a way, our policy can be seen as an adaptive query policy where requests tend to align to the most useful or important system-wise source realizations.

This paper falls within the realm of pull-based communication methods, but with a query control policy, owing to which the query controller timely decides to pull updates following the source changes and the updates’ expected usefulness. Our work leverages the concept of Just-in-Time systems [9] and extends prior work on query age of information (QAoI) [4]–[8] into the systems where queries arrive at the right times to request the source realizations (updates) with high usefulness and significant effect at the endpoint. In this regard, a class of optimal policies is derived to maximize the long-term expected effectiveness of the update packets a sensing agent sends to an acting agent, subject to a communication cost. Our analytical and simulation results show that the proposed query policy outperforms existing baseline approaches in terms of effectiveness grade. Furthermore, we demonstrate that the solution converges to a threshold-based query control scheme, where the query controller can timely decide whether to pull an update or not by merely looking up a table with the obtained threshold boundaries for a given goal.

II. SYSTEM MODEL

We consider an end-to-end pull-based communication system in which update packets are generated via a *sensing agent* (SA) and transmitted to an *actuation agent* (AA) for taking appropriate actions at the endpoint and fulfilling the subscribed

The work of P. Agheli and M. Kountouris has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101003431). The work of N. Pappas is supported by the VR, ELLIIT, and the European Union (ETHER, 101096526), and the work of P. Popovski was supported by the Villum Investigator Grant “WATER” from the Velux Foundation, Denmark.

goal. The packet generation is done in the event of receiving a *query* from a query controller at the AA. Here, the packets contain status updates of a sensor observing an information source/event or payload data from the application layer towards the network. We assume the system operates in a slotted manner with time slots n , for $n \in \mathbb{N}$. At the n -th time slot, the query arrival indicator is denoted by $\alpha_n = \{0, 1\}$, with $\alpha_n = 1$ indicating the arrival of a query, and $\alpha_n = 0$, otherwise. Thus, the set of slots at which queries arrive is defined as $\mathcal{N} = \{n \mid \forall n : \alpha_n = 1\}$. The update and query channels are assumed to be *error-prone* and *error-free*, respectively, and the whole duration of receiving a query and the consequent update acquisition and its communication is normalized into one time slot. Also, we consider x_n the update communicated at slot n , which can be seen as some form of semantic representation.

In parallel with its acquisition, the n -th update is evaluated based on its usefulness for satisfying the goal and is attributed a rank of importance (value) at the source level in the form of a meta-value v_n , $\forall n \in \mathcal{N}$ [10, Section III-A]. For the sake of generality, we assume that v_n is a random variable (r.v.) that follows a discrete-time Markov process with finite state space $\mathcal{V} = \{\nu_i \mid i = 1, 2, \dots, |\mathcal{V}|\}$ and transition matrix $\mathbf{P}_v = [p_{ij}]_{|\mathcal{V}| \times |\mathcal{V}|}$ with $p_{ij} = \Pr(v_{n+1} = \nu_j \mid v_n = \nu_i)$. From an effectiveness viewpoint, a successfully received update has a specific impact at the receiver or decision-making side. In the proposed model, we consider that the goal is perfectly known at both SA and AA, thus the same framework is employed to measure the usefulness of updates at both ends.¹

A. Query Control Policies

We consider that the query controller can apply two different types of policies to pull updates from the SA in a pull-based system, as follows.

1) *Effect-agnostic policy*: Under this policy, the queries arrive with a *controlled query rate*, pursuing a specific schedule or a stochastic process, e.g., Poisson, binomial, and Markovian [5]–[8]. Therefore, the existing effect-agnostic policies pose *aleatoric* uncertainty associated with the nature of random updates by ignoring what is probably occurring at the source during the time of pulling those updates.

2) *Effect-aware policy*: With this policy, the AA tries to infer or predict the probable state of the source and the expected usefulness of an update at the time of the decision, hence adapting its query instances and pulling updates to the right time slots. This can be done via building or updating a model of the source's process at the AA. Owing to this policy, albeit packet generation and communication is requested by queries, what happens at the source is also considered. Therefore, *epistemic* uncertainty arises with the effect-aware policy, as decisions are done based on probabilistic predictions, instead of accurate knowledge. This uncertainty could be reduced or harnessed using for instance machine learning or prediction methods. We delve into this query control problem in Section IV.

¹The analysis can be easily extended to the case where the goal is not shared, but the AA, based on the received updates, may learn or estimate the probability distribution of the updates' usefulness at the endpoint.

III. GRADE OF EFFECTIVENESS METRIC

To model effectiveness, we advocate for a metric that involves two system-level attributes: *freshness* and *usefulness* of the successfully received updates. Freshness describes how obsolete an update gets as time passes from the instant it is correctly received at the AA. This attribute is commonly quantified via the age of information (AoI) metric or its variants. Besides, the usefulness signifies the rank of importance, as defined in Section II, each correctly received packet offers at the endpoint.

With the above explanation in mind, we propose the grade of effectiveness (GoE) metric for measuring the amount of impact an update has at the endpoint.

A. GoE Formulation

The GoE for the update packet communicated at the n -th time slot is denoted by $\text{GoE}_n \in \mathbb{R}_0^+$ and modeled in the form of a composite function $f : \mathbb{R}_0^+ \times \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ of the AoI, called Δ_n for the n -th time slot, and the usefulness of that update, i.e., v_n , to satisfy the subscribed goal. Thus, we can write²

$$\text{GoE}_n = f(g_\Delta(\Delta_n), g_v(v_n)) \quad (1)$$

where $g_\Delta : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is a non-increasing *penalty* function, and $g_v : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ shows a non-decreasing *utility* function. In (1), $\Delta_n = n - m$ with $\Delta_0 = 1$, and $m = \max_{i: \alpha_i(1-\epsilon_i)=1, i \leq n} i$. Also, $\epsilon_n = g_\epsilon(d(x_n, \hat{x}_n))$ depicts the update discrepancy, where \hat{x}_n indicates the received update at the AA, $g_\epsilon : \mathbb{R}_0^+ \rightarrow \{0, 1\}$ is the mapping function to the Boolean space, and $d : \mathbb{R} \rightarrow \mathbb{R}_0^+$ is a distance function. In (1), if we overlook the usefulness of updates or their freshness, the GoE metric turns into the penalty definition for the query AoI (QAoI), i.e., $\text{GoE}_n = f(g_\Delta(\Delta_n))$, or the utility formulation for the *value of information* (VoI), i.e., $\text{GoE}_n = f(g_v(v_n))$, respectively.

B. The Pull-Based System's GoE

Since the update acquisition and communication is done on the condition of receiving queries in the pull-based system, $\text{GoE}_n^{(\text{pull})}$ becomes limited to the slots at which the queries arrive, i.e., $\forall n \in \mathcal{N}$. From (1), we can write

$$\text{GoE}_n^{(\text{pull})} = \text{GoE}_n \mid_{\alpha_n=1} \times \mathbb{1}(\alpha_n = 1) \quad (2)$$

where $\mathbb{1}(\cdot)$ denotes the indicator function.

IV. QUERY CONTROL

In this section, we first formulate the query control problem for the effect-aware policy, and we then propose a solution for the defined decision/control problem.

A. Problem Formulation

The objective is to maximize the long-term expected GoE via controlling the query arrivals, subject to a constraint on the average communication cost C_{\max} , which cannot be surpassed.

²The GoE metric can be seen as a particular case of the Semantics of Information (SoI) metric [2], [11] and can also be defined in various forms based on the scenario, including that utilized for active fault detection in [12].

In this sense, a class of optimal policies, named π^* , is derived by solving an optimization problem as follows

$$\begin{aligned} \mathcal{P}_1 : \max_{\pi} \limsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{n=1}^N \text{GoE}_n^{(\text{pull})} \mid \text{GoE}_0^{(\text{pull})} \right] \\ \text{s.t. } \limsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{n=1}^N g_c(\alpha_n c_n) \right] \leq C_{\max} \end{aligned} \quad (3)$$

where $g_c : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is a non-decreasing function, and c_n indicates the communication cost at the n -th time slot.

We cast \mathcal{P}_1 into an *infinite-horizon* constrained Markov decision process (CMDP) based on the following definitions:

- I) *States*: The state at the n -th slot depicts the GoE and is denoted with a tuple $S_n = (\Delta_n, v_n)$. Without loss of generality, we consider the AoI to be truncated by the maximum value of Δ_{\max} , which is large enough to represent excessive staleness and meet $g_{\Delta}(\Delta_{\max} - 1) \leq (1 + \varepsilon)g_{\Delta}(\Delta_{\max})$ with accuracy ε . Given this, S_n is a member of the state space $\mathcal{S} = \{\sigma_i \mid i = 1, 2, \dots, |\mathcal{S}|\}$, which is countable and finite with $|\mathcal{S}| = \Delta_{\max} \cdot |\mathcal{V}|$.
- II) *Actions*: With regard to Section II, we define the action space \mathcal{A} with two possible outcomes: $\alpha_n = 1$ for pulling an update, and $\alpha_n = 0$ for keeping silent.
- III) *Transition probabilities*: The transition probability from S_n to S_{n+1} under the action α_n is $P(S_n, \alpha_n, S_{n+1}) = \Pr((\Delta_{n+1}, v_{n+1}) \mid (\Delta_n, v_n), \alpha_n)$, $\forall n$, where we define
 - $\Pr((1, \nu_j) \mid (\Delta_n, \nu_i), \alpha_n) = p_{ij} \alpha_n (1 - p_{\varepsilon}^{(n)})$, $\forall i, j$,
 - $\Pr((\min\{\Delta_n + 1, \Delta_{\max}\}, \nu_i) \mid (\Delta_n, \nu_i), \alpha_n) = \alpha_n p_{\varepsilon}^{(n)} + (1 - \alpha_n)$, $\forall i$.
Otherwise, we have $P(S_n, \alpha_n, S_{n+1}) = 0$. Here, $\nu_i, \nu_j \in \mathcal{V}$, and $p_{\varepsilon}^{(n)} = \Pr(\varepsilon_n = 1)$.
- IV) *Rewards*: The reward of going from S_n to S_{n+1} in the course of the action α_n is $R(S_n, \alpha_n, S_{n+1}) = \text{GoE}_{n+1}^{(\text{pull})}$.

With the above definitions in mind, we give the following proposition that states that the expected sum of GoE in (3) is the same for all initial states, hence there exists an optimal stationary policy for the defined problem.

Proposition 1. *The modeled CMDP pursues the weak accessibility (WA) condition.*

Proof. We can partition the space set \mathcal{S} into two subsets $\mathcal{S}_a = \{S_n \in \mathcal{S} \mid S_n = (1, \nu_n), \forall \nu_n \in \mathcal{V}\}$ and $\mathcal{S}_b = \mathcal{S} - \mathcal{S}_a$. The primary subset contains all states whose $\Delta_n = 1$, hence $|\mathcal{S}_a| = |\mathcal{V}|$. The latter subset includes the rest of the states with $\Delta_n \geq 2$, and $|\mathcal{S}_b| = (\Delta_{\max} - 1) \cdot |\mathcal{V}|$. In this regard, all states of \mathcal{S}_b are transient under any stationary policy, while every state of an arbitrary pair of two states in \mathcal{S}_a is accessible from the other state. Given this and considering the WA condition as in [13, Definition 4.2.2], the modeled CMDP is weakly accessible.

As the WA condition holds for the modeled CMDP, it can be concluded that the optimal expected GoE remains the same for all initial states, i.e., $\text{GoE}_n^{(\text{pull})}$, $\forall n$, is independent of $\text{GoE}_0^{(\text{pull})}$ [13, Proposition 4.2.3]. Moreover, Proposition 1 confirms the existence of an optimal stationary policy π^* for \mathcal{P}_1 . This

optimal policy is *unchain* [13, Proposition 4.2.6]. To solve the problem, we first relax its constrained form via defining a dual problem, then we proceed with proposing an algorithm.

B. Dual problem

We convert the constrained form of \mathcal{P}_1 to an unconstrained one via writing the Lagrange function $\mathcal{L}(\mu; \pi)$ as follows

$$\begin{aligned} \mathcal{L}(\mu; \pi) = \limsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{n=1}^N \left(\text{GoE}_n^{(\text{pull})} - \mu g_c(\alpha_n c_n) \right) \right] \\ + \mu C_{\max} \end{aligned} \quad (4)$$

where $\mu \geq 0$ indicates the Lagrange multiplier. Then, we can summarize the Lagrange dual problem to be solved as

$$\mathcal{P}_2 : \inf_{\mu \geq 0} \max_{\pi} \mathcal{L}(\mu; \pi) \quad (5)$$

$:= h(\mu)$

in the form of an *unconstrained* MDP, where $h(\mu) = \mathcal{L}(\mu; \pi_{\mu}^*)$ shows the Lagrange dual function. Here, $\pi_{\mu}^* : \mathcal{S} \rightarrow \mathcal{A}$ appears as the μ -optimal policy and is derived from the dual problem for a given μ , as $\pi_{\mu}^* = \arg \max_{\pi} \mathcal{L}(\mu; \pi)$.

Since \mathcal{S} has finite states, which ensures the growth condition as in [14], and $\text{GoE}_n \geq 0, \forall n$, from (1), \mathcal{P}_1 and \mathcal{P}_2 converge to the same optimal values [14, Corollary 12.2]. Therefore, we can find the class of optimal policies π^* after applying an *iterative algorithm* approach [15].

C. Iterative Algorithm

The iterative approach is illustrated in Algorithm 1 with two *inner* and *outer* loops for deriving the μ -optimal policy, i.e., π_{μ}^* , and the optimal Lagrange multiplier, i.e., μ^* , respectively.

1) *Computing π_{μ}^** : In the inner loop, with a given μ from the outer loop, the query policy is iteratively updated taking an optimal action which maximizes the *expected utility* (value) $V_{\pi_{\mu}^*}^{(t)}(s)$ for the state $s \in \mathcal{S}$ at the t -th, $t \in \mathbb{N}$, iteration. Under the form of the value iteration approach for the unichain policy MDPs [16], the optimal value function is derived from Bellman's equation [17], as follows

$$V_{\pi_{\mu}^*}^{(t)}(s) = \max_{\alpha \in \mathcal{A}} \sum_{s' \in \mathcal{S}} P(s, \alpha, s') \left[R_{\mu}(s, \alpha, s') + V_{\pi_{\mu}^*}^{(t-1)}(s') \right]. \quad (6)$$

Consequently, the optimal policy for $s \in \mathcal{S}$ is updated by

$$\pi_{\mu}^*(s) \in \arg \max_{\alpha \in \mathcal{A}} \sum_{s' \in \mathcal{S}} P(s, \alpha, s') \left[R_{\mu}(s, \alpha, s') + V_{\pi_{\mu}^*}^{(t-1)}(s') \right] \quad (7)$$

where $R_{\mu}(s, \alpha, s') := R(s, \alpha, s') - \mu g_c(\alpha c)$ is the *net* reward.

The inner loop stops once the stopping/convergence criterion $\max_{s \in \mathcal{S}} |V_{\pi_{\mu}^*}^{(t+1)}(s) - V_{\pi_{\mu}^*}^{(t)}(s)| \leq \varepsilon_v$ is satisfied, where ε_v indicates the convergence accuracy. As the query policies are unichain with aperiodic transition matrices, the above convergence criterion is reached for some finite iterations [16, Theorem 8.5.4].

2) *Computing μ^** : In order to find the optimal Lagrange multiplier $\mu^{(l)}$ at the l -th, $l \in \mathbb{N}$, step of the outer loop, according to the updated π_{μ}^* from the inner loop, we apply

Algorithm 1: Solution for deriving π^* and μ^* .

Input: Known parameters $N \gg 1$, $c_n, \forall n$, ε_μ , C_{\max} , η , states \mathcal{S} , and actions \mathcal{A} . Initial values $l \leftarrow 1$, $\mu^{(0)} \leftarrow 0$, $\mu^- \leftarrow 0$, $\mu^+ \gg 1$, $\pi_\mu^- \leftarrow 0$, and $\pi_\mu^+ \leftarrow 0$. The form of g_c .

- 1 Initialize $\pi_\mu^*(s), \forall s \in \mathcal{S}$, via running $\text{Utility}(\mu^{(0)})$.
- 2 if $\mathbb{E}\left[\sum_{n=1}^N g_c(\alpha_n c_n)\right] \leq NC_{\max}$ then goto 10.
- 3 while $|\mu^+ - \mu^-| \geq \varepsilon_\mu$ do
 - Step l : \triangleright Outer loop (Bisection search)
 - 4 Update $\mu^{(l)} \leftarrow \frac{\mu^+ + \mu^-}{2}$ and $\pi_\mu^* \leftarrow \text{Utility}(\mu^{(l)})$.
 - 5 if $\mathbb{E}\left[\sum_{n=1}^N g_c(\alpha_n c_n)\right] \geq NC_{\max}$ then
 - 6 $\mu^- \leftarrow \mu^{(l)}$, and $\pi_\mu^- \leftarrow \text{Utility}(\mu^-)$.
 - 7 else $\mu^+ \leftarrow \mu^{(l)}$, and $\pi_\mu^+ \leftarrow \text{Utility}(\mu^+)$.
 - 8 Reset $l \leftarrow l + 1$.
- 9 if $\mathbb{E}\left[\sum_{n=1}^N g_c(\alpha_n c_n)\right] < NC_{\max}$ then
 - $\pi_\mu^*(s) \leftarrow \eta \pi_\mu^-(s) + (1 - \eta) \pi_\mu^+(s), \forall s \in \mathcal{S}$.
- 10 return $\mu^* = \mu^{(l)}$ and $\pi^*(s) = \pi_\mu^*(s), \forall s \in \mathcal{S}$.

Function $\text{Utility}(\mu)$:

Input: Known parameters $N \gg 1$, ε_v , states \mathcal{S} , and actions \mathcal{A} . Initial values $t \leftarrow 1$, $\pi_\mu^*(s) \leftarrow 0$, and $V_{\pi_\mu^*}^{(0)}(s) \leftarrow 0, \forall s \in \mathcal{S}$.

Iteration t : \triangleright Inner loop (Value iteration approach)

- 11 for state $s \in \mathcal{S}$ do
 - 12 compute $V_{\pi_\mu^*}^{(t)}(s)$ from (6), then update $\pi_\mu^*(s)$ as in (7) with the optimal action from $V_{\pi_\mu^*}^{(t)}(s)$.
 - 13 if $\max_{s \in \mathcal{S}} |V_{\pi_\mu^*}^{(t+1)}(s) - V_{\pi_\mu^*}^{(t)}(s)| \geq \varepsilon_v$ then
 - 14 step up $t \leftarrow t + 1$, and goto 11.
 - 15 return π_μ^* .

the so-called bisection method with the stopping criterion $|\mu^+ - \mu^-| \leq \varepsilon_\mu$ and the search accuracy ε_μ , as depicted in Algorithm 1. From (4) and (5), the increase of μ continuously increases the dual function $h(\mu)$ while decreasing the net reward $\text{GoE}_n^{(\text{pull})} - \mu g_c(\alpha_n c_n)$ and the query arrival rate. Thus, we search for the smallest value of μ that satisfies the communication cost C_{\max} . Since $\text{GoE}_n^{(\text{pull})}$ is independent of $\mu g_c(\alpha_n c_n)$, one can verify that $h(\mu)$ is a Lipschitz continuous function of μ with the Lipschitz constant of

$$\left| C_{\max} - \limsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{n=1}^N g_c(\alpha_n c_n) \right] \right|.$$

Thus, the outer loop converges to the optimal multiplier after finite iterations [18, pp. 294]. The optimal value is attained based on a simple non-randomized stationary policy or a mix of two non-randomized policies with a mixing probability η , which can be obtained such that $\mathbb{E}\left[\sum_{n=1}^N g_c(\alpha_n c_n)\right] = NC_{\max}$ [19].

Algorithm 1 has at most $\mathcal{O}(2LT\Delta_{\max}^2|\mathcal{V}|^2)$ arithmetic operations, where L and T indicate the step sizes of the outer

and inner loops, respectively. The increase of the states and the iteration size of either loop increase the algorithm's complexity. Nevertheless, this complexity is manageable in real-world scenarios owing to the following items:

- Δ_{\max} is set to a small value since the effectiveness of an update at the endpoint saturates to a close to zero value past a level of staleness, making that update not useful.
- As the usefulness of the updates can be normalized, i.e., $\nu_i \in [0, 1], \forall i \in \mathcal{V}$, a large outcome space is not necessarily needed for \mathcal{V} .
- The communication cost is usually fixed, thus reducing the bisection search interval, hence L , after some trials.

D. Threshold-Based Query Control Model

Finding the optimal policy π^* from Algorithm 1 gives us a threshold criterion Ω_{th} for every CMDP's state, following which maximizes the GoE of the system subject to the communication cost. With Ω_{th} in hand, the query controller timely decides to pull an update or not based on the current AoI, i.e., Δ_n , and the meta-value of the latest correctly received update at the query time, i.e., v_m for $m = \max_{i: \alpha_i(1-\varepsilon_i)=1, i \leq n} i$.

In this sense, the optimal action α_n^* for the n -th time slot can be derived according to two alternative options, as given below.

Option I: The value of Δ_{th} is a function of v_m , where

$$\Omega_{\text{th}} = \Delta_{\text{th}}^{(v_m)}, \text{ hence } \alpha_n^* = \mathbb{1}(\Delta_n \geq \Delta_{\text{th}}^{(v_m)} | v_m). \quad (8)$$

Option II: The level of v_{th} depends on Δ_n . Thus, we have

$$\Omega_{\text{th}} = v_{\text{th}}^{(\Delta_n)}, \text{ hence } \alpha_n^* = \mathbb{1}(v_m \leq v_{\text{th}}^{(\Delta_n)} | \Delta_n). \quad (9)$$

V. SIMULATION RESULTS

We assess the performance of the proposed effect-aware query control policy within $N = 1000$ time slots and compare it with three existing effect-agnostic query arrival models, namely (i) *periodic* model, and stochastic models following (ii) *binomial* and (iii) *Markovian* process. For the latter model, we have a Markov chain with two states of "pulling an update" and "keeping silent", in which the self-transition probability of the latter state is 0.95, while the one for the primary state relies on the query rate. The default query rate for the effect-agnostic policies is 0.8, unless otherwise stated. Furthermore, we equally divide the interval $[0, 1]$ into $|\mathcal{V}| = 10$ levels and initialize the importance set $\mathcal{V} = \{0, 0.11, 0.22, 0.33, 0.44, 0.56, 0.67, 0.78, 0.89, 1\}$, each indicating a normalized rank of importance with $p_{ii} = p_{ij} = \frac{1}{|\mathcal{V}|}, \forall \nu_i, \nu_j \in \mathcal{V}$. The maximum acceptable AoI is $\Delta_{\max} = 10$, and the probability of the update discrepancy is $p_\epsilon^{(n)} = 0.2, \forall n$. Also, for Algorithm 1, we set $\varepsilon_v = \varepsilon_\mu = 10^{-3}$ and $\eta = 0.5$.

For performance evaluation, we define a *net* GoE (NGoE) metric, which incorporates the GoE and the cost, as follows

$$\begin{aligned} \text{NGoE}_n^{(\text{pull})} &= \exp(\text{GoE}_n^{(\text{pull})} - g_c(\alpha_n c_n)) \\ &= \exp(-v_n \Delta_n - c_0 \alpha_n) \end{aligned} \quad (10)$$

for the n -th slot, $\forall n \in \mathcal{N}$, where the exponential form is arbitrarily used to guarantee positiveness. In (10), without loss

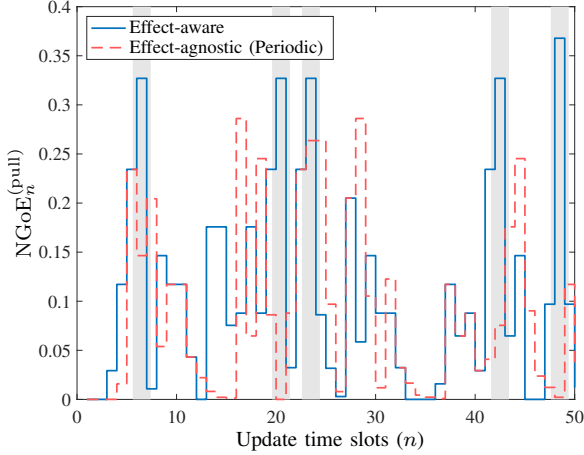


Fig. 1. The effectiveness performance of different query control policies within $N = 50$ slots.




Fig. 2. The CDF of the average NGoE provided over $N = 1000$ time slots.

of generality, we employ a *linear* form for g_Δ , g_v , and g_c , and a *multiplication* form for f according to (1). Besides, we consider *uniform* cost $c_n = c_0, \forall n$, and unless otherwise specified, we initialize $c_0 = 0.5$ and $C_{\max} = 0.4$.

Fig. 1 depicts a 50-slot snapshot of the status update system according to the arrived queries and the updates' NGoE, comparing the normalized GoE of the proposed effect-aware with the periodic effect-agnostic query control policy. The latter policy is assumed to have 7-slot period intervals. We observe that our effect-aware policy enables capturing updates with the highest NGoE, of which the horizons are highlighted. However, in some slots, the effect-aware policy does not pull usefulness updates due to the probabilistic uncertainties of the CMDP problem, while the periodic one can catch them by chance.

Fig. 2 presents the cumulative distribution function (CDF) of the average NGoE provided by applying the effect-aware and effect-agnostic policies over $N = 1000$ time slots. It is shown that the effect-aware policy highly boosts the effectiveness of the system, thanks to its prediction of the updates' usefulness


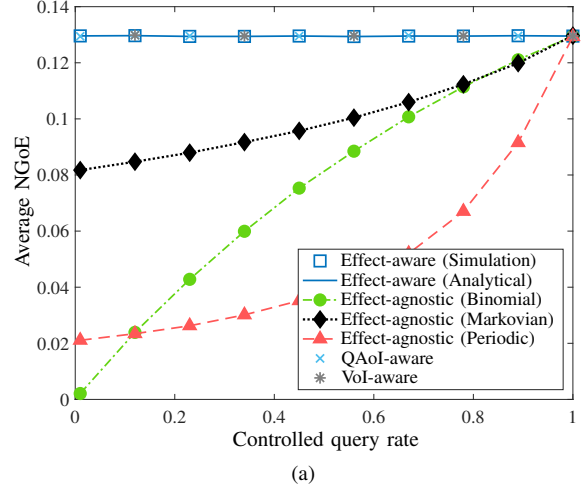


Fig. 3. The interplay between average NGoE and query rate for different query control policies and cost coefficients (a) $c_0 = 0.1$ and (b) $c_0 = 0.5$.

and through pulling significant updates, compared to the effect-agnostic policies. Specifically, the effect-aware policy increases the effectiveness by 91%, 47%, and 149% on average, as compared to the binomial, Markovian, and periodic models, respectively. This comes at the cost of 16% on average higher transmission rate for the effect-aware policy than the others.

Figs. 3(a) and 3(b) show the interplay between average NGoE and controlled query rate for $c_0 = 0.1$ and $c_0 = 0.5$, respectively, for effect-aware and effect-agnostic policies, as well as for two special cases relevant to QAoI and VoI within $N = 1000$ time slots. In the QAoI-aware case, decision policies are obtained so as to maximize the long-term expected QAoI regardless of the usefulness or importance of the updates, whereas in the VoI-aware case, decision policies depend on the expected VoI, without taking freshness into account. We can see that the effect-aware policy outperforms the other effect-agnostic policies for all query rates and under both cost values. Comparing Fig. 3(a) with Fig. 3(b), we infer that, for $c_0 = 0.1$, the special cases offer almost the same performance as the




Fig. 4. The expected utility obtained in each iteration of Algorithm 1.

effect-aware policy considering both freshness and usefulness attributes. This is because all policies can pull the same updates, thus having identical performance. However, for $c_0 = 0.5$, the effect-aware policy outperforms both special cases for all query rates, even using fewer resources, owing to its effect-aware query control. Also, since $\Delta_n \geq 1$, and $0 \leq v_n \leq 1, \forall n$, bypassing the freshness attribute in the GoE metric leads to a significant performance drop. Besides, the increase of the cost decreases the offered average NGoE, and that at a faster speed for the effect-agnostic policies. Indeed, not useful or irrelevant transmissions under the effect-agnostic policies result in higher waste of resources and larger gap between the achieved NGoE.

To study the convergence of the iterative approach for solving the CMDP problem regarding to Algorithm 1, we plot Fig. 4, which illustrates the value (expected utility) obtained in each iteration for different cost coefficients. It can be seen that the policy convergences to its final value after 130, 133, and 137 iterations, sequentially, for the cost coefficients $c_0 = 0.1$, $c_0 = 0.5$, and $c_0 = 1$. We can also observe that the value provided in each iteration decreases by increasing the cost, in line with the formulation of the CMDP's rewards in Section IV-A.

Finally, Table I demonstrates the threshold boundaries for the threshold-based query control model discussed in Section IV-D for cost coefficients $c_0 = 0.5$ and $c_0 = 1$. The term "Silent" corresponds to $\alpha_n = 0$, and "Pull" indicates $\alpha_n = 1$. Either Option I or Option II could be applied to find Ω_{th} based on $\Delta_{th}^{(v_m)}$ or $v_{th}^{(\Delta_n)}$, respectively. As in Table I, the query controller pulls updates in case $v_m \leq 0.44, \forall m$, regardless of the current AoI. Also, updates are always pulled if $\Delta_n \geq 4, \forall n$, independent of the degree of usefulness. For the above conditions, we thus have fixed $\Delta_{th} = 4$ and $v_{th} = 0.44$ under any cost coefficient for the primary and latter options, respectively. In order to derive the optimal action for the other conditions (CMDP's states), the communication cost plays a key role, hence resulting in variable threshold metrics. As the cost increases, the controller should pull updates merely under more *critical* conditions so as to increase the system's NGoE.

VI. CONCLUSION

We have proposed an effect-aware query control policy for pull-based communication systems, in which the query controller timely decides whether to pull an update, depending on the source's evolution and the updates' effectiveness at the endpoint. We have considered the problem of GoE maximiza-

TABLE I
ILLUSTRATION OF THE THRESHOLD BOUNDARIES FOR THE THRESHOLD-BASED QUERY CONTROL MODEL

		v_m					
		$0 \rightarrow 0.44$	0.56	0.67	0.78	0.89	1
$c_0 = 0.5$	Δ_n	$1 \rightarrow 2$	Silent				
		3	Pull	Pull			Silent
		$4 \rightarrow 10$	Pull		Pull		
$c_0 = 1$	Δ_n	$1 \rightarrow 3$	Silent				
		$4 \rightarrow 10$	Pull		Pull		

tion based on a CMDP with finite state spaces for the AoI and the usefulness rank and have provided an algorithm to find the class of optimal policies. Our results have shown that the effect-aware query policy could provide significant gains in terms of normalized GoE compared to effect-agnostic policies for different query rates and communication costs.

REFERENCES

- [1] C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, 1949.
- [2] M. Kountouris and N. Pappas, "Semantics-empowered communication for networked intelligent systems," *IEEE Comm. Mag.*, vol. 59, no. 6, 2021.
- [3] P. Popovski, O. Simeone, F. Boccardi, D. Gündüz, and O. Sahin, "Semantic-effectiveness filtering and control for post-5G wireless connectivity," *J. IISc*, vol. 100, no. 2, 2020.
- [4] B. Yin, S. Zhang, Y. Cheng, L. X. Cai, Z. Jiang, S. Zhou, and Z. Niu, "Only those requested count: Proactive scheduling policies for minimizing effective age-of-information," in *IEEE INFOCOM*, 2019.
- [5] F. Li, Y. Sang, Z. Liu, B. Li, H. Wu, and B. Ji, "Waiting but not aging: Optimizing information freshness under the pull model," *IEEE/ACM Transactions on Networking*, vol. 29, no. 1, 2021.
- [6] J. Holm, A. E. Kalør, F. Chiariotti, B. Soret, S. K. Jensen, T. B. Pedersen, and P. Popovski, "Freshness on demand: Optimizing age of information for the query process," in *IEEE ICC*, 2021.
- [7] F. Chiariotti, J. Holm, A. E. Kalør, B. Soret, S. K. Jensen, T. B. Pedersen, and P. Popovski, "Query age of information: Freshness in pull-based communication," *IEEE Trans. on Commun.*, vol. 70, no. 3, 2022.
- [8] M. E. Ildiz, O. T. Yavascan, E. Uysal, and O. T. Kartal, "Query age of information: Optimizing AoI at the right time," in *IEEE ISIT*, 2022.
- [9] C. E. Leiserson, "Systolic priority queues," *Proceedings of the Caltech Conference On Very Large Scale Integration*, 1979.
- [10] P. Agheli, N. Pappas, and M. Kountouris, "Semantic filtering and source coding in distributed wireless monitoring systems," *arXiv preprint arXiv:2305.11549*, 2023.
- [11] N. Pappas and M. Kountouris, "Goal-oriented communication for real-time tracking in autonomous systems," in *IEEE ICAS*, 2021.
- [12] G. J. Stamatakis, N. Pappas, A. Fragkiadakis, and A. Traganitis, "Semantics-aware active fault detection in IoT," in *20th WiOpt*, 2022.
- [13] D. Bertsekas, *Dynamic programming and optimal control*. Athena Scientific, 2007, vol. 2.
- [14] E. Altman, *Constrained Markov decision processes*. CRC press, 1999.
- [15] M. Hatami, M. Leinonen, Z. Chen, N. Pappas, and M. Codreanu, "On-demand AoI minimization in resource-constrained cache-enabled IoT networks with energy harvesting sensors," *IEEE Transactions on Communications*, vol. 70, no. 11, 2022.
- [16] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [17] R. Bellman, "On the theory of dynamic programming," *Proceedings of The National Academy of Sciences*, vol. 38, no. 8, 1952.
- [18] G. Wood, *Bisection global optimization methods*, C. A. Floudas and P. M. Pardalos, Eds. Springer Science & Business Media, 2009.
- [19] F. J. Beutler and K. W. Ross, "Optimal policies for controlled Markov chains with a constraint," *Journal of mathematical analysis and applications*, vol. 112, no. 1, 1985.