





Decentralized Control of Distributed Cloud Networks with Generalized Network Flows

Yang Cai , *Graduate Student Member, IEEE*, Jaime Llorca , *Member, IEEE*,
Antonia M. Tulino , *Fellow, IEEE*, and Andreas F. Molisch , *Fellow, IEEE*

Abstract—Emerging distributed cloud architectures, e.g., fog and mobile edge computing, are playing an increasingly important role in the efficient delivery of real-time stream-processing applications (also referred to as augmented information services), such as industrial automation and metaverse experiences (e.g., extended reality, immersive gaming). While such applications require processed streams to be shared and simultaneously consumed by multiple users/devices, existing technologies lack efficient mechanisms to deal with their inherent multicast nature, leading to unnecessary traffic redundancy and network congestion. In this paper, we establish a unified framework for distributed cloud network control with generalized (mixed-cast) traffic flows that allows optimizing the distributed execution of the required packet processing, forwarding, and replication operations. We first characterize the *enlarged* multicast network stability region under the new control framework (with respect to its unicast counterpart). We then design a novel queuing system that allows scheduling data packets according to their current destination sets, and leverage Lyapunov drift-plus-penalty control theory to develop the *first fully decentralized, throughput- and cost-optimal algorithm for multicast flow control*. Numerical experiments validate analytical results and demonstrate the performance gain of the proposed design over existing network control policies.

Index Terms—Distributed computing, routing, flow control, distributed control, cost optimal control.

I. INTRODUCTION

APPLICATIONS receiving the most recent attention, such as system automation (e.g., smart homes/factories/cities, self-driving cars) and metaverse experiences (e.g., multiplayer gaming, immersive video, virtual/augmented reality) are characterized by intensive resource consumption and real-time interactive requirements, where multiple users/devices simultaneously consume information that results from the real-time processing of a variety of live multimedia streams [3]. We refer to this general class of services as augmented information (AgI) services.

The high computation and low latency requirements of AgI services are fueling the evolution of network architectures

toward widespread deployments of increasingly distributed computation resources, leading to what is referred to as *distributed cloud/computing networks*, including fog and mobile edge computing (MEC) [4]. In contrast to traditional architectures, where there is a clear separation between data processing at centralized data centers and data transmission between remote data centers and end users, distributed cloud networks are evolving toward tightly integrated compute-communication systems, where the ubiquitous availability of computation resources enables reduced access delays and energy consumption, thus providing better support for next-generation delay-sensitive, compute-intensive AgI services.

In addition, the increasing amount of real-time multi-user interactions present in AgI services, where media streams can be shared and simultaneously consumed by a large number of users/devices, is creating the need to efficiently support multicast traffic. While, in general, one can describe four types of network flows, i.e., unicast (packets intended for a unique destination), multicast (packets intended for multiple destinations), broadcast (packets intended for all destinations), and anycast (packets intended for any node in a given group), we focus our attention on multicast flows as the most general class, given that the other three types either become special cases of or can be transformed into multicast flows.

To maximize the benefit of distributed cloud networks for the delivery of next-generation multicast AgI services, two fundamental flow control problems need to be addressed:

- **packet processing:** *where to process data packets by the required service functions and how much computation resource to allocate*
- **packet forwarding and replication:** *how to route data packets through the required sequence of service functions, where to replicate data packets in order to satisfy the demand from multiple destinations, and how much network resource to allocate*

In addition, the above processing, routing (including forwarding and replication), and resource allocation problems must be addressed in an online manner, in response to unknown time-varying network conditions and service demands.

A. Related Work

With the advent of software defined networking (SDN) and network function virtualization (NFV), network services can be deployed as a sequence of software functions or service function chains (SFCs) that can be flexibly interconnected and elastically executed at distributed cloud locations [4]. A

Part of this work was presented at the 2021 IEEE ICC [1]. A longer version of this paper with more detailed derivations and proofs is given in [2].

Y. Cai and A. F. Molisch are with the Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA 90089, USA (e-mail: yangcai@usc.edu; molisch@usc.edu).

J. Llorca is with the Electrical and Computer Engineering Department, New York University, Brooklyn, NY 11201 USA (e-mail: jllorca@nyu.edu).

A. M. Tulino is with the Electrical and Computer Engineering Department, New York University, Brooklyn, NY 11201 USA, and also with the Department of Electrical Engineering, Università degli Studi di Napoli Federico II, Naples 80138, Italy (e-mail: atulino@nyu.edu; antoniamaria.tulino@unina.it).

This work was supported by the National Science Foundation (NSF) under CNS-1816699 and RINGS-2148315.

number of studies have investigated the problem of joint SFC placement and routing with the objective of either maximizing accepted service requests [5], [6], or minimizing overall resource cost [7], [8], [9]. Nonetheless, these solutions exhibit two main limitations: first, the problem is formulated as a *static* optimization problem without considering the dynamic nature of service requests, a critical aspect in next-generation AgI services; second, due to the combinatorial nature of the problem, the corresponding formulations typically take the form of (NP-hard) mixed integer linear programs, and either heuristic solutions or loose approximation algorithms are developed, compromising the quality of the resulting solution.

More recently, another line of work has addressed the SFC optimization problem in *dynamic* scenarios, where one needs to make joint packet processing and routing decisions in an online manner [10], [11], [12], [13]. The work in [10] leverages Lagrange duality and saddle point theory to design an iterative algorithm that employs global candidate path information to optimize service chain multi-path routing and associated sending rates. The works in [11], [12] employ a generalized cloud network flow model that allows joint control of processing and transmission flows. The work in [13] shows that the cloud network flow control problem (involving processing and transmission decisions) can be reduced to a packet routing problem on a *cloud layered graph* that includes extra edges to characterize the computation operations (i.e., data streams pushed through these edges are interpreted as being processed by a service function). By this transformation, control policies designed for packet routing can be extended to address cloud network control problems.

Dynamic unicast routing is a long-explored problem, with a number of existing algorithms known to maximize network throughput with bounded average delay. In particular, the back-pressure (BP) algorithm [14] is a well-known approach for throughput-optimal routing that leverages Lyapunov drift control theory to steer data packets based on the *pressure* difference (differential queue backlog) between neighbor nodes. In addition, the Lyapunov drift-plus-penalty (LDP) control approach [15] extends the BP algorithm to also minimize network operational cost (e.g., energy expenditure), while preserving throughput-optimality. However, both BP and LDP approaches can suffer from poor average delay performance, especially in low congestion scenarios, where packets can follow unnecessarily long, and sometimes even cyclic, paths for delivery [16]. To address this problem, [17] combines BP and hop-distance based shortest-path routing, which can effectively reduce the average delay.

Going beyond unicast traffic and addressing the design of optimal routing policies for multicast traffic is a much more challenging problem, as the need for packet replications violates the flow conservation law [13]. Despite the large body of existing works on this topic [18], *throughput-optimal least-cost multicast packet routing* remains an open problem. Even under *static* arrivals, the Steiner tree problem, which aims to find the multicast tree (a tree that covers all destinations) with minimum weight, is known to be NP-hard [19]. Many heuristic approaches have been developed to address this problem, such as the Extended Dijkstra's Shortest Path Algorithm (EDSPA)

[20], which delivers multicast packets along a tree formed by the union of the shortest paths from the source to all destinations. However, in addition to their heuristic nature, packets are delivered along fixed paths under these policies, lacking dynamic exploration of route and processing diversity. Considering *dynamic* arrivals becomes a further challenge that requires additional attention. A centralized dynamic packet routing and scheduling algorithm, UMW, was proposed in [21], shown to achieve optimal throughput with mixed-cast network flows. Nonetheless, this design exhibits two limitations: (i) it makes centralized decisions based on global queuing information, incurring additional communication overhead, (ii) it leaves out operational cost minimization, an important aspect in modern elastic network environments.

B. Contributions

The goal of this work is to develop *decentralized* control policies for distributed cloud network flow problems dealing with mixed-cast AgI services. We establish a new multicast control framework that guides the creation of copies of data packets as they travel toward their corresponding destinations. Compared to the state-of-the-art unicast approach [11] that “creates one copy for each destination of a multicast packet upon arrival, and treats them as individual unicast packets”, the proposed policy employs a *joint forwarding and replication* strategy that (i) eliminates redundant transmissions along network links common to multiple copies' routes, and (ii) reduces computation operations by driving computation before replication when beneficial.

The proposed approach is based on a novel multicast queuing system that allows formalizing the *packet replication operation*, which defines “where to create copies” and “how to assign destinations to resulting copies”, as a network control primitive. Under the proposed queuing system, each packet is labeled and queued according to its *replication status*, which keeps track of its current destination set. The packet replication operation is then specified by the partition of the destination set of a given packet into the destination sets of each resulting copy. We finally devise a fully decentralized packet processing, forwarding, and replication policy that attains optimal throughput and cost performance (see Theorem 2 for details), as well as a variant achieving sub-optimal performance with polynomial complexity.

Our contributions can be summarized as follows:

- 1) We establish a novel queuing system to accommodate packets according to their replication status that allows supporting packet processing, routing, and replication operations as necessary cloud network control primitives for the delivery of mixed-cast AgI services.
- 2) We characterize the *enlarged* multicast network stability region obtained by including packet replication as an additional control primitive, and quantify the resulting gain with respect to (w.r.t.) its unicast counterpart.
- 3) We devise GDCNC, the first fully decentralized, throughput- and cost-optimal algorithm for distributed cloud network control with mixed-cast network flows.
- 4) We design GDCNC-R, a computational-efficient policy achieving sub-optimal performance with polynomial

complexity by focusing on a subset of effective replication operations.

- 5) We conduct extensive numerical experiments that support analytical results and demonstrate the performance benefits of the proposed design for the delivery of mixed-cast AgI services.

Organization: In Section II, we introduce the model for the “multicast packet routing” problem. In Section III, we define the policy space and present a characterization for the multicast network stability region. Section IV describes the multicast queuing system and defines the problem formulation. In Section V, we devise the GDCNC control policy and analyze its performance, which further motivates the design of GDCNC-R in Section VI. Extensions to the proposed design are discussed in Section VII. Section VIII presents the numerical results, and conclusions are drawn in Section IX.

Notation: Let $\{0, 1\}^D$ denote the set of all D -dimensional binary vectors. We use $\mathbf{0}$ to denote the zero vector, $\mathbf{1}$ the all-ones vector, and b_k the vector with only the k th entry equal to 1 (and 0’s elsewhere). Given a binary vector $q \in \{0, 1\}^D$, $\bar{q} = \mathbf{1} - q$ denotes its complement and $2^q \triangleq \{s : s_k = q_k u_k, u \in \{0, 1\}^D\}$ its power set. The inner product of vector x and y is denoted by $\langle x, y \rangle$. $\mathbb{E}\{z\}$ is the expectation of random variable z , and $\{z(t)\} \triangleq \lim_{T \rightarrow \infty} (1/T) \sum_{t=0}^{T-1} z(t)$ the time average of random process $\{z(t) : t \geq 0\}$. $\mathbb{I}\{\mathcal{A}\}$ denotes the indicator function (equal to 1 if event \mathcal{A} is true, and 0 otherwise), and $|\mathcal{A}|$ the cardinality of set \mathcal{A} .

II. SYSTEM MODEL

A. Cloud Layered Graph

The ultimate goal of this work is to design decentralized control policies for distributed cloud networks equipped with computation resources (e.g., cloud servers, edge/fog computing nodes, etc.) able to host AgI service functions and execute corresponding computation tasks.

While in traditional packet routing problems, each node treats its neighbor nodes as outgoing interfaces over which packets can be scheduled for transmission, a key step to address the AgI service delivery problem, involving packet processing and routing decisions, is to treat co-located computing devices as an additional interface over which packets can be scheduled for processing [11]. As illustrated in [13], the AgI service control problem can be transformed into a packet routing problem on a *layered graph* where cross-layer edges represent computation resources.

Motivated by such a connection and for ease of exposition, in this paper, without loss of generality, we illustrate the developed approach focusing on the *single-commodity, least-cost multicast packet routing* problem. We remark that (i) the optimal decentralized multicast control problem has been open up to now even in traditional communication networks, and (ii) the extension to distributed cloud network control is presented in Section VII-A.

B. Network Model

The considered packet routing network is modeled by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each edge $(i, j) \in \mathcal{E}$

represents a network link supporting data transmission from node i to j , and δ_i^- and δ_i^+ denote the incoming and outgoing neighbor sets of node i , respectively.

Time is slotted. For each link $(i, j) \in \mathcal{E}$, we define: (i) *transmission capacity* C_{ij} as the maximum number of data units (e.g., packets) that can be transmitted in one time slot, and (ii) *unit transmission cost* e_{ij} as the cost (e.g., power consumption) incurred in transmitting one unit of data in one time slot. We note that a linear cost model is used here for ease of exposition, while extensions to a non-linear cost model characterizing power expenditure in a practical wireless scenario is studied in Section VII-B.

We emphasize that in the layered graph, cross-layer edges represent data processing, i.e., data streams pushed through these edges are interpreted as being processed by corresponding service functions, and the capacity and cost of these edges represent the *processing capacity* and *unit processing cost* of the associated computation resources (e.g., cloud/edge servers).

C. Arrival Model

We focus on a multicast application (the extension to multiple applications is straightforward, and details are given in [2, Appendix F]) where each incoming packet is associated with a destination set $\mathcal{D} = \{d_1, \dots, d_D\} \subset \mathcal{V}$, with d_k denoting the k -th destination and $D = |\mathcal{D}|$ the destination set size. At least one copy of the incoming packet must be delivered to every destination in \mathcal{D} . Importantly, we assume that delivering multiple copies of the same packet (containing the same content) to a given destination does not increase network throughput.

Multicast packets originate at the application source nodes $\mathcal{S} \subset \mathcal{V} \setminus \mathcal{D}$. Let $a_i(t)$ denote the number of exogenous packets arriving at node i at time t , with $a_i(t) = 0, \forall i \notin \mathcal{S}$. We assume that the arrival process is i.i.d. over time, with mean arrival rate $\lambda_i \triangleq \mathbb{E}\{a_i(t)\}$ and bounded instantaneous arrivals; the corresponding vectors are denoted by $\mathbf{a}(t) = \{a_i(t) : i \in \mathcal{V}\}$ and $\boldsymbol{\lambda}$.

Remark 1: By properly defining the destination set \mathcal{D} , the above model can capture all four network flow types: (i) unicast and broadcast flows are special cases of a multicast flow, defined by setting $\mathcal{D} = \{d\}$ and $\mathcal{D} = \mathcal{V}$, respectively; and (ii) an anycast flow can be transformed into a unicast flow by creating a *super destination node* connected to all the candidate destinations [21]. Therefore, it suffices to focus on the multicast flow to derive a general solution for the mixed-cast flow control problem.

D. In-network Packet Replication

We now formalize the most important operation for multicast packet routing, namely *in-network packet replication*.

1) *Replication Status:* We assume that each packet is associated with a label that indicates its *current* destination set, i.e., the set of destinations to which a copy of the packet must still be delivered, formally defined as follows.

Definition 1: For a given packet, the *replication status* $q = [q_1, \dots, q_D]$ is a D -dimensional binary vector where the k -th

entry $(k = 1, \dots, D)$ is set to $q_k = 1$ if destination $d_k \in \mathcal{D}$ belongs to its current destination set, and to $q_k = 0$ otherwise.

Three important cases follow: (i) $q = 1$ indicates the status of every newly arriving packet prior to any replication operation, with the entire destination set \mathcal{D} as their current destination set; (ii) $q = b_k$ describes a packet with one destination d_k , which behaves like a unicast packet; (iii) $q = 0$ describes a packet without any destination, which is removed from the system immediately.

2) **Packet Replication and Coverage Constraint:** A replication operation creates copies of a packet and assigns a new destination set, which must be a subset of the original packet's destination set, to each resulting copy. Let $q \in \{0, 1\}^D$ denote the replication status of the original packet; then, the set of all possible replication status of the copies is given by its power set 2^q . To ensure the delivery of a packet to all of its destinations, we impose the **Coverage** constraint on the replication operation: each destination node of the original packet must be present in the destination set of *at least* one of the resulting copies.

3) **Conditions on Replication Operation:** In addition to the **Coverage** constraint, we require the replication operation to satisfy the following **Conditions**:

- a) **Joint forwarding and replication:** Replication is performed only on packets to be transmitted.
- b) **Efficient replication:** The destination sets of the created copies do not to overlap.
- c) **Duplication:** Only two copies are created by one replication operation.
- d) **Non-consolidation:** Co-located packets of identical content are *not* combined; they are kept as separate copies.

It can be shown that these **Conditions** do not reduce the achievable throughput nor increase the minimum attainable cost. **Condition a)** avoids replicating packets that are not scheduled for transmission, which only increases network congestion and should be avoided. **Condition b)** is motivated by the fact that "receiving multiple packets of identical content at the same destination does not increase network throughput", and thus replication should be performed in an efficient manner, i.e., NOT producing copies with overlapping destinations, to alleviate network traffic and associated resource consumption. **Conditions c)** and **d)** are justified in Appendix A.

Combining the **Coverage** constraint and the above **Conditions** leads to the following important property: each destination node of a packet undergoing duplication (we use "duplication" instead of "replication" in the rest of the paper, e.g., duplication status) must be present in the destination set of exactly one of the two resulting copies.

As illustrated in Fig. 1, the duplication operation process works as follows. Let q denote the duplication status of a packet selected for (transmission) operation. Upon duplication, one copy is transmitted to the corresponding neighbor node (referred to as the *transmitted* copy, of status s), and the other copy stays at the node waiting for future operation (referred to as the *reloaded* copy, of status r). Then, $q = s + r$, and we refer to the pair (q, s) as the *duplication choice* (and $r = q - s$).

Let $\Omega = \{(q, s) : q \in \{0, 1\}^D, s \in 2^q\}$ denote the set of all duplication choices.

Remark 2: We note that the duplication choice (q, q) describes the special case that a packet is transmitted *without duplication*. In particular, note that for a status b_k packet, i.e., a (unicast) packet with one destination, (b_k, b_k) is the *only* duplication choice.

Remark 3: Duplication (q, b_k) automatically takes place when a status q packet with d_k in its current destination set (i.e., $q_k = 1$) arrives at destination d_k , in which case: the status b_k copy departs the network immediately, and the status $q - b_k$ copy stays at node d_k .

III. POLICY SPACES AND NETWORK STABILITY REGION

In this section, we introduce the policy space for multicast packet delivery, based on which we characterize the multicast network stability region.

A. Policy Space

1) **Decision Variables:** We consider a general policy space for multicast packet delivery, encompassing joint packet forwarding and duplication operations, whose associated forwarding and duplication scheduling variables are described by

$$\mu(t) = \{\mu_{ij}^{(q,s)}(t) : (q, s) \in \Omega, (i, j) \in \mathcal{E}\} \quad (1)$$

where $\mu_{ij}^{(q,s)}(t)$ denotes the amount of status q packets selected for (forwarding and duplication) operation, with duplication choice (q, s) and forwarding choice (i, j) , at time t . That is, $\mu_{ij}^{(q,s)}(t)$ status q packets are duplicated, resulting in $\mu_{ij}^{(q,s)}(t)$ status s packets transmitted over link (i, j) , and $\mu_{ij}^{(q,s)}(t)$ status $q - s$ packets reloaded to node i , as illustrated in Fig. 1.

2) **Admissible Policies:** A control policy is *admissible* if the flow variables satisfy:

- a) non-negativity: $\mu(t) \succeq 0$, i.e., $\mu_{ij}^{(q,s)}(t) \geq 0$ for $\forall (i, j) \in \mathcal{E}, (q, s) \in \Omega$.
- b) link capacity constraint: $\sum_{(q,s) \in \Omega} \mu_{ij}^{(q,s)}(t) \leq C_{ij}$ for $\forall (i, j) \in \mathcal{E}$.
- c) generalized flow conservation: for each intermediate node $i \in \mathcal{V} \setminus \mathcal{D}$ and $q \in \{0, 1\}^D$,

$$\sum_{s \in 2^q} \sum_{j \in \delta_i^-} \{\mu_{ij}^{(q+s,q)}(t)\} + \sum_{s \in 2^q} \sum_{j \in \delta_i^+} \{\mu_{ij}^{(q+s,s)}(t)\} + \lambda_i^{(q)} = \sum_{s \in 2^q} \sum_{j \in \delta_i^+} \{\mu_{ij}^{(q,s)}(t)\} \quad (2)$$

where $\lambda_i^{(q)}$ is the mean rate of exogenously arriving packets $a_i^{(q)}(t) = a_i(t) \mathbb{I}\{q = 1\}$, and $\bar{q} = 1 - q$ denotes the complement of q , i.e., the set of destinations not included in q .

- d) boundary conditions: $\mu_{d_k j}^{(q,s)}(t) = 0$ for $\forall d_k \in \mathcal{D}, j \in \delta_{d_k}^+, q$ with $q_k = 1, k \in \{1, \dots, D\}$.

The generalized flow conservation c) can be described as follows. (i) In contrast to the *instantaneous* constraints, a), b), and d), which must hold at each time slot, c) imposes an equality constraint on the *average flow rates* of incoming/outgoing status q packets to/from node i . (ii) As illustrated

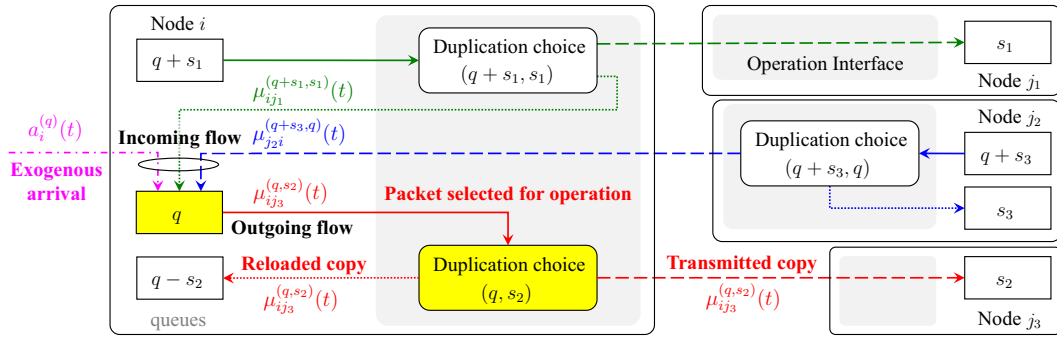


Fig. 1. Illustration of joint packet forwarding and duplication operations (solid, dashed, and dotted lines represent packets selected for operation, transmitted copies, and reloaded copies, respectively) and incoming/outgoing flow variables associated with the status q queue of node i (while only explicitly indicated for the red flow, note that all arrows with the same color are associated with the same flow variable).

in Fig. 1, for each node i and status q queue: the incoming flow of status q packets has three components: packets received from each neighbor node $j \in \delta_i^-$ after undergoing duplication $(q + s, q)$ (which creates transmitted copies of status q), i.e., $\mu_{ji}^{(q+s, q)}(t)$; packets that stay at node i after undergoing local duplication $(q + s, s)$ (which creates reloaded copies of status $q = (q + s) - s$), i.e., $\mu_{ij}^{(q+s, s)}(t)$; and exogenously arriving status q packets, i.e., $a_i^{(q)}(t)$. The outgoing flow includes all status q packets selected for operation with duplication choice (q, s) , i.e., $\mu_{ij}^{(q, s)}(t)$.

3) *Cost Performance:* The instantaneous overall resource operational cost of an admissible policy is given by

$$h(t) = h(\mu(t)) = \sum_{(i,j) \in \mathcal{E}} e_{ij} \sum_{(q,s) \in \Omega} \mu_{ij}^{(q,s)}(t) = \langle e, \mu(t) \rangle. \quad (3)$$

We employ the *expected time average cost* $\overline{\mathbb{E}\{h(t)\}}$ to characterize the cost performance, in which the expectation is taken w.r.t. the flow variables $\mu(t)$, governed by the arrival process $a(t)$ and the (possibly randomized) control policy. We denote by $h^*(\lambda)$ the minimum attainable cost under the arrival vector λ .

B. Multicast Network Stability Region

In this section, we characterize the **multicast network stability region** Λ to measure the throughput performance of the system, defined as the set of arrival vectors $\lambda = \{\lambda_i : i \in \mathcal{V}\}$ under which there exists an admissible policy satisfying the constraints in Section III-A2.

Theorem 1: An arrival vector λ is interior to the stability region Λ if and only if there exist flow variables $f = \{f_{ij}^{(q,s)} \geq 0 : \forall (q, s), (i, j) \in \mathcal{E}\}$ and probability values $\beta = \{\beta_{ij}^{(q,s)} \geq 0 : \sum_{(q', s') \in \Omega} \beta_{ij}^{(q', s')} \leq 1, \forall (q, s), (i, j) \in \mathcal{E}\}$ such that:

$$\begin{aligned} & \sum_{s \in 2^q} \sum_{j \in \delta_i^-} f_{ji}^{(q+s, q)} + \sum_{s \in 2^q} \sum_{j \in \delta_i^+} f_{ij}^{(q+s, s)} + \lambda_i^{(q)} \\ & \leq \sum_{s \in 2^q} \sum_{j \in \delta_i^+} f_{ij}^{(q, s)}, \quad \forall i \in \mathcal{V}, q \in \{0, 1\}^D, \quad (4a) \end{aligned}$$

$$f_{ij}^{(q, s)} \leq \beta_{ij}^{(q, s)} C_{ij}, \quad \forall (i, j) \in \mathcal{E}, (q, s) \in \Omega, \quad (4b)$$

$$f_{dkj}^{(q, s)} = 0, \quad \forall d_k \in \mathcal{D}, j \in \delta_{d_k}^+, q : q_k = 1, k \in \{1, \dots, D\}. \quad (4c)$$

In addition, there exists a stationary randomized policy specified by probability values β (i.e., at each time slot t , each link (i, j) selects C_{ij} status q packets for forwarding and duplication operation with duplication choice (q, s) with probability $\beta_{ij}^{(q, s)}$) and attains optimal cost $h^*(\lambda)$.

Proof: The proof follows the same steps as [11, Appendix A] with the following changes: (i) to show *Necessity*, we consider a *multicast stabilizing policy* that makes additional decisions on duplication choices (q, s) , and (ii) to show *Sufficiency*, we consider the *stationary randomized policy* presented in Theorem 1. The detailed proofs are presented in [2, Appendix B]. \square

Next, we analyze the benefit of the multicast framework exploiting in-network packet duplication to *enlarge* network stability region, compared to its unicast counterpart Λ_0 .

Proposition 1: The multicast network stability region Λ satisfies $\Lambda_0 \subseteq \Lambda \subseteq D\Lambda_0$, in which $D\Lambda_0 \triangleq \{D\lambda : \lambda \in \Lambda_0\}$. Furthermore, for any arrival vector $\lambda \in \Lambda_0$, the minimum attainable cost under the unicast framework, denoted by $h_0^*(\lambda)$, satisfies $h_0^*(\lambda) \geq h^*(\lambda)$.

Proof: Since unicast policies, which perform all duplication operations at the source node, become special solutions under the multicast framework, the attained throughput and cost performance cannot be superior to that of the optimal multicast policy, which can eliminate redundant transmissions leveraging in-network packet duplication. Hence, $\Lambda_0 \subseteq \Lambda$ and $h_0^*(\lambda) \geq h^*(\lambda)$. Next, we show $\Lambda \subseteq D\Lambda_0$. For each packet delivered along a multicast route (under a given multicast policy \mathcal{P}_m), the unicast policy \mathcal{P}_u can send each copy along the same path as in the multicast route. For each link, \mathcal{P}_u consumes at most D times the communication resources consumed by \mathcal{P}_m (when \mathcal{P}_u sends D individual copies over a link common to all unicast paths, while \mathcal{P}_m only sends 1 packet and performs duplication later). With this key observation, we can show (by contradiction) that the gain factor is bounded by D . The detailed proofs are given in [2, Appendix B]. \square

IV. PROBLEM FORMULATION

A. Queuing System

Since forwarding and duplication decisions are driven by each packet's destination set, keeping track of data packets' destination sets is essential. A key step is to construct a

queuing system with distinct queues for packets of different duplication status q . In particular, we denote by $Q_i^{(q)}(t)$ the backlog of the queue holding status q packets at node i at time t . Define $Q_i^{(0)}(t) = 0$ and $\mathbf{Q}(t) = \{Q_i^{(q)}(t) : i \in \mathcal{V}, q \in \{0, 1\}^D\}$.

Each time slot is divided into two phases. In the *transmitting* phase, each node makes and executes forwarding and duplication decisions based on the observed queuing states. In the *receiving* phase, the incoming packets, including those received from neighbor nodes, reloaded copies generated by duplication, as well as exogenously arriving packets, are loaded into the queuing system, and the queuing states are updated.

1) *Queuing Dynamics*: The queuing dynamics are derived for two classes of network nodes.

(i) For an *intermediate node* $i \in \mathcal{V} \setminus \mathcal{D}$, the queuing dynamics are given by

$$Q_i^{(q)}(t+1) \leq \max \left[0, Q_i^{(q)}(t) - \mu_{i \rightarrow}^{(q)}(t) \right] + \mu_{\rightarrow i}^{(q)}(t) + a_i^{(q)}(t), \quad \forall q \in \{0, 1\}^D \quad (5)$$

where $\mu_{i \rightarrow}^{(q)}(t)$ and $\mu_{\rightarrow i}^{(q)}(t)$ are the outgoing and (controllable) incoming network flows of status q packets from/to node i , given by (as illustrated in Fig. 1)

$$\begin{aligned} \mu_{i \rightarrow}^{(q)}(t) &= \sum_{j \in \delta_i^+} \sum_{s \in 2^q} \mu_{ij}^{(q,s)}(t), \\ \mu_{\rightarrow i}^{(q)}(t) &= \sum_{j \in \delta_i^-} \sum_{s \in 2^{\bar{q}}} \mu_{ji}^{(q+s,q)}(t) + \sum_{j \in \delta_i^+} \sum_{s \in 2^{\bar{q}}} \mu_{ij}^{(q+s,s)}(t). \end{aligned} \quad (6)$$

Note that (5) holds with inequality when the number of status q packets incoming to node i is less than $\mu_{\rightarrow i}^{(q)}(t)$ (see Remark 4), and the difference is bounded by $\sum_{j \in \delta_i^-} C_{ji} + \sum_{j \in \delta_i^+} C_{ij}$.

(ii) For a *destination node* $i = d_k$ ($k = 1, \dots, D$), the queuing dynamics are given by

$$Q_i^{(q)}(t+1) \leq (I + \mu_{\rightarrow i}^{(q+b_k)}(t)) \mathbb{I}\{q_k = 0\}, \quad \forall q \in \{0, 1\}^D, \quad (7)$$

where I represents the right hand side of (5). To wit, at destination node d_k : (i) all status q queues with $q_k = 1$ are always empty, and (ii) all other queues have an additional incoming flow corresponding to the reloaded copies resulting from the automatic duplication of status $q+b_k$ packets arriving at the destination (see also Remark 3), i.e., $\mu_{\rightarrow i}^{(q+b_k)}(t)$.

Remark 4: Similar to many existing control policies, e.g., [11], [14], [15], [22], we note that flow variable $\mu(t)$ can lead to a total outgoing flow $\mu_{i \rightarrow}^{(q)}(t)$ exceeding the available packets in queue $Q_i^{(q)}(t)$, and we include the ramp function $\max[0, \cdot]$ in (5) to avoid negative queue length. If $\mu_{i \rightarrow}^{(q)}(t) > Q_i^{(q)}(t)$, the residual capability $\mu_{i \rightarrow}^{(q)}(t) - Q_i^{(q)}(t)$ is either used with *idle fill* (referred to as dummy packets) or *wasted*, which is shown not to affect the network stability region [15].

B. Problem Formulation

The goal is to develop an admissible control policy that stabilizes the queuing system, while minimizing overall operational cost. Formally, we aim to find a control policy with

decisions $\{\mu(t) : t \geq 0\}$ satisfying

$$\min_{\mu(t) \geq 0} \overline{\mathbb{E}\{h(t)\}} \quad (8a)$$

$$\text{s. t. } \overline{\mathbb{E}\{\|\mathbf{Q}(t)\|_1\}} < \infty, \text{ i.e., stabilizing } \mathbf{Q}(t) \quad (5) - (7), \quad (8b)$$

$$\sum_{(q,s) \in \Omega} \mu_{ij}^{(q,s)}(t) \leq C_{ij}, \quad \forall (i,j), \quad (8c)$$

$$\mu_{d_k j}^{(q,s)}(t) = 0, \quad \forall d_k \in \mathcal{D}, j \in \delta_{d_k}^+, q : q_k = 1, k \in \{1, \dots, D\}. \quad (8d)$$

In addition, note that by Little's Theorem [15], the average delay is linear in the queue backlog $\overline{\mathbb{E}\{\|\mathbf{Q}(t)\|_1\}}$, and thus (8b) is equivalent to guaranteeing *finite average delay*.

V. GENERALIZED DISTRIBUTED CLOUD NETWORK CONTROL ALGORITHM

In this section, we leverage the LDP theory [15] to address (8), which guides the design of the proposed GDCNC algorithm.

A. Lyapunov Drift-plus-Penalty

Consider the Lyapunov function $L(t) = \|\mathbf{Q}(t)\|_2^2/2$ and associated Lyapunov drift $\Delta(\mathbf{Q}(t)) = L(t+1) - L(t)$. The LDP approach aims to minimize an upper bound of a linear combination of the Lyapunov drift (whose derivation follows the steps in [15, Section 3.1.2], and details are shown in [2, Appendix C]) and the objective function weighted by a tunable parameter V , i.e.,

$$\Delta(\mathbf{Q}(t)) + Vh(t) \leq |\mathcal{V}|B + \langle \mathbf{a}(t), \mathbf{Q}(t) \rangle - \langle \mathbf{w}(t), \mu(t) \rangle \quad (9)$$

where B is a constant, and the *duplication utility weights* $\mathbf{w}(t) = \{w_{ij}^{(q,s)}(t)\}$ are given by

$$w_{ij}^{(q,s)}(t) = Q_i^{(q)}(t) - Q_i^{(q-s)}(t) - Q_j^{(s)}(t) - Ve_{ij}. \quad (10)$$

Equivalently, the proposed algorithm selects the flow variable $\mu(t)$ to maximize $\langle \mathbf{w}(t), \mu(t) \rangle$ at each time slot, which can be decomposed into separate problems for each link (i,j) :

$$\begin{aligned} \max_{\mu_{ij}^{(q,s)}(t) : (q,s) \in \Omega} \sum_{(q,s) \in \Omega} w_{ij}^{(q,s)}(t) \mu_{ij}^{(q,s)}(t), \\ \text{s. t. } \sum_{(q,s) \in \Omega} \mu_{ij}^{(q,s)}(t) \leq C_{ij}, \quad \mu_{ij}^{(q,s)}(t) \geq 0. \end{aligned} \quad (11)$$

The resulting *max-weight* solution is described in the next section.

B. Generalized Distributed Cloud Network Control

The developed *generalized distributed cloud network control* (GDCNC) algorithm is described in Algorithm 1, and exhibits two salient features. (i) *Decentralized*: it only requires local information exchange (i.e., queuing states of neighbor nodes) and decision making, which can be implemented in a fully distributed manner. (ii) *Sparse*: for each link (i,j) , it selects one duplication choice (q^*, s^*) at each time slot, which affects the states of status q^* and $q^* - s^*$ queues at node i , and status s^* queue at node j . Therefore, for each node i ,

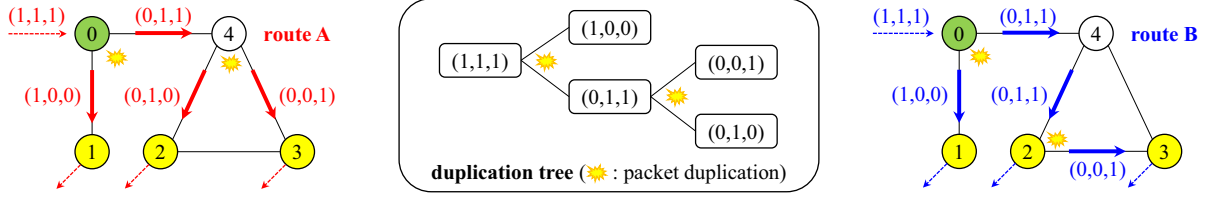


Fig. 2. Different multicast routes (i.e., routing trees) can share the same duplication choices (i.e., duplication tree). For example, the red and blue routes used to deliver the status $(1, 1, 1)$ packet from source node 0 to destination nodes $\{1, 2, 3\}$ are associated with the same duplication tree.

Algorithm 1 GDCNC

- 1: **for** $t \geq 0$ and $(i, j) \in \mathcal{E}$ **do**
- 2: Calculate the duplication utility weight $w_{ij}^{(q,s)}(t)$ for all duplication choices (q, s) by (10).
- 3: Find the (q, s) pair with the largest weight $(q^*, s^*) = \arg \max_{(q,s) \in \Omega} w_{ij}^{(q,s)}(t)$.
- 4: Assign transmission flow:

$$\mu_{ij}^{(q,s)}(t) = C_{ij} \mathbb{I}\{w_{ij}^{(q^*, s^*)}(t) > 0, (q, s) = (q^*, s^*)\}.$$
- 5: **end for**

the number of queues with changing states in one time slot is $2|\delta_i^+| + |\delta_i^-| \sim \mathcal{O}(|\delta_i|)$, where $|\delta_i|$ denotes the degree of node i .

Remark 5: Note that when dealing with a unicast flow, where all packets have the same single destination, the only valid duplication status is $q = 1$ and the only valid duplication choice is $(q, s) = (1, 1)$, in which case GDCNC reduces to DCNC [11].

C. Performance Analysis

In this section, we analyze the delay and cost performance of GDCNC, and its complexity from both communication and computation dimensions.

1) *Delay-Cost Tradeoff:* In the following theorem, we employ the minimum attainable cost $h^*(\lambda)$ as the benchmark to evaluate the performance of GDCNC.

Theorem 2: For any arrival vector λ interior to the multicast stability region Λ , the average queue backlog and the operational cost achieved by GDCNC satisfy

$$\begin{aligned} \overline{\{\mathbb{E}\{\|Q(t)\|_1\}\}} &\leq \frac{|\mathcal{V}|B}{\epsilon} + \left[\frac{h^*(\lambda + \epsilon \mathbf{1}) - h^*(\lambda)}{\epsilon} \right] V, \\ \overline{\{\mathbb{E}\{h(t)\}\}} &\leq h^*(\lambda) + \frac{|\mathcal{V}|B}{V}, \end{aligned} \quad (12)$$

for any $\epsilon > 0$ such that $\lambda + \epsilon \mathbf{1} \in \Lambda$.

Proof: Applying *expected time average* and *telescope sum* [15] to (9), we can obtain $V\{\mathbb{E}\{h(t)\}\} \leq |\mathcal{V}|B - \epsilon\{\mathbb{E}\{\|Q(t)\|_1\}\} + Vh^*(\lambda + \epsilon \mathbf{1})$. Then, using $\{\mathbb{E}\{h(t)\}\} \geq h^*(\lambda)$ and $-\epsilon\{\mathbb{E}\{\|Q(t)\|_1\}\} \leq 0$ leads to (12). See [2, Appendix D] for a detailed derivation. \square

The above theorem is illustrated as follows. For any arrival vector interior to the stability region, GDCNC (using any fixed $V \geq 0$) can stabilize the queuing system, and thus is throughput-optimal. In addition, GDCNC achieves an

$[\mathcal{O}(V), \mathcal{O}(1/V)]$ tradeoff between delay (which is linear in queue backlog) and cost: by pushing $V \rightarrow \infty$, the attained cost can be arbitrarily close to the minimum $h^*(\lambda)$, with a tradeoff in network delay.

2) *Complexity Issues:* Next, we analyze the complexity of GDCNC.

Communication Overhead: At each time slot, GDCNC requires local exchange of queue backlog information. Instead of the entire queuing states (of size $|\{0, 1\}^D| \sim \mathcal{O}(2^D)$), we can leverage *sparsity* (see Section V-B) to reduce communication overhead, i.e., nodes only exchange information of the queues with changing states, reducing the overhead to $\mathcal{O}(|\delta_i|)$.

Computational Complexity: At each time slot, GDCNC calculates the utility weight of each duplication choice (q, s) , of computational complexity proportional to $|\Omega| \sim \mathcal{O}(3^D)$ (see Appendix B), i.e., exponential in the destination set size. This is due to the combinatorial nature of the multicast routing problem. Indeed, the state-of-the-art centralized solution to the multicast flow control problem [13] requires solving the NP-complete Steiner tree problem for route selection at each time slot.

VI. GDCNC-R WITH REDUCED COMPLEXITY

In this section, we develop GDCNC-R, a variant of GDCNC achieving sub-optimal performance with *reduced* complexity.

A. Duplication Tree

To illustrate the complexity reduction of GDCNC-R, we first define the duplication tree as a useful representation of the duplication operations performed on a given multicast packet.

Definition 2: A *duplication tree* \mathcal{T} is a binary tree with each node associated with a duplication status: the root node and the D leaf nodes represent the status of the initial multicast packet, $\mathbf{1}$, and of the copies delivered to each destination, b_k ($k \in \{1, \dots, D\}$), respectively. Each internal node q splits into two child nodes s and r , associated with duplication (q, s) and (q, r) , respectively.

As illustrated in Fig. 2, the duplication choices of a given multicast route are described by a duplication tree. Note however that different routes can share the same duplication choices, hence the same duplication tree. GDCNC achieves optimal performance by evaluating all possible duplication choices at each node, which is equivalent to consider all routes in all duplication trees. This is also the reason for its high computational complexity.

B. Proposed Approach

In contrast, the developed GDCNC-R algorithm narrows the focus on the multicast routes included in a subset of *effective* duplication trees and associated duplication choices. Given a duplication tree, the policy can select from all the multicast routes mapping into the tree for packet delivery, with the goal of reducing network congestion, and as a result, improving the delay performance. The selected duplication trees shall minimize the resulting (*throughput and cost*) *performance loss*, and the proposed approach is referred to as *destination clustering*.

Importantly, the main benefit of the multicast framework is to save redundant operations by exploiting in-network packet duplication. To maximize this gain, a duplication shall be performed only on packets with “distant” destinations in their current destination set (under the distance metrics listed at the end of this section).

To this end, we propose to construct the duplication tree as follows (in this paragraph, we refer to a node in the duplication tree as a *vertex* to distinguish from a node in the routing network). We start from the entire network node set \mathcal{V} , which includes all destinations and corresponds to the root vertex of the duplication tree. At each step, the current node set is divided into *two* disjoint clusters – in line with the *duplication* condition (see Section II-D3). The two child vertices are given by the destinations in the resulting clusters. Then, we move to the child vertices and repeat this step for the corresponding clusters, until each cluster only includes one destination node (we discard clusters not including any destinations in the procedure).

While there are many existing clustering methods, we use the widely adopted *k-means clustering* [23], described as follows. Given a set of data points with known *features* and a distance function measuring the *similarity* between two data points, the algorithm starts assigning random labels (i.e., the cluster to which the data point belongs) to the data points; then, it (i) calculates the *center* for each cluster, and (ii) updates the label of each data point to the cluster with the *nearest* center (under the given distance function) in an iterative manner until converging.

Some effective choices for the inter-node distance are listed as follows. In the wired case: To optimize the throughput performance, we can use the *reciprocal transmission capacity* as the inter-node distance; while *unit transmission cost* is a proper metric to deal with the cost performance. For nodes that are not directly connected, we assume that there is a link of zero capacity and infinity cost between them. Given the pairwise similarities, we can apply spectrum clustering techniques [24] to define features for the data points, followed by *k-means* clustering. In the wireless scenario, a straightforward distance metric is the *Euclidean distance*, which leads to using the *geographical location* of each network node as its feature, and whose effectiveness is validated by numerical experiments (see Fig. 3a and 3b).

C. Complexity Analysis

As shown in Appendix B, each duplication tree includes $2D - 1$ nodes, with $D - 1$ internal nodes (including the root node) and D leaf nodes. There are 3 possible duplication choices associated with each internal node q (with child nodes s and r), i.e., (q, q) , (q, s) and (q, r) , while only one duplication choice (b_k, b_k) is associated with each leaf node b_k . Therefore, every duplication tree includes $4D - 3 \sim \mathcal{O}(D)$ possible duplication choices.

GDCNC-R uses K duplication trees, with K chosen to strike a good balance between performance optimality and computational complexity. Specifically, GDCNC-R has $\mathcal{O}(KD)$ complexity, i.e., polynomial in the number of duplication trees and destination set size.

VII. EXTENSIONS

In this section, we present extensions to the GDCNC algorithms, including: (i) a cloud network control policy for multicast AgI service delivery, (ii) a modified algorithm for MEC scenarios (with wireless links), and (iii) a variant of GDCNC with *Enhanced* delay performance, EGDCNC.

A. Multicast AgI Service Delivery

In line with [11], [13], the additional packet processing decisions involved in the distributed cloud network setting can be handled as follows.

1) *Cloud Network Model*: Consider a cloud network composed of compute-enabled nodes that can process data packets via corresponding service functions, with the available processing resources and associated costs defined as: (i) *processing capacity* C_i , i.e., the computation resource units (e.g., computing cycles per time slot) at node i , and (ii) *unit processing cost* e_i , i.e., the cost to run one unit of computation resource in one time slot at node i .

2) *AgI Service Model*: Consider an AgI service modeled as SFC, i.e., a chain of $M - 1$ functions, through which source packets must be processed to produce consumable results, resulting in the end-to-end data stream divided into M *stages*. Each processing step can take place at different network locations hosting the required service functions, and each function ($m \in \{1, \dots, M - 1\}$) in the service chain is described by two parameters: (i) *scaling factor* $\xi^{(m)}$, i.e., the number of output packets per input packet, and (ii) *workload* $r^{(m)}$, i.e., the amount of computation resource to process one input packet.

3) *Processing Decision*: We create different queues to hold packets of different processing stages and duplication status. Let $Q_i^{(m,q)}(t)$ denote the backlog of (stage m , status q) packets, and $\mu_i^{(m,q,s)}(t)$ the scheduled *processing flow*, i.e., the amount of (stage m , status q) packets selected for (processing and duplication) operation with duplication choice (q, s) at node i at time t .

Following the procedure in Section V-A on a properly constructed cloud layered graph [13], the processing decisions for each node i at time t are given by

- (i) Calculate the duplication utility weight $w_i^{(m,q,s)}(t)$ for each (m, q, s) tuple:

$$w_i^{(m,q,s)}(t) = [Q_i^{(m,q)}(t) - \xi^{(m)} Q_i^{(m+1,s)}(t) - Q_i^{(m,q-s)}(t)] / r^{(m)} - V e_i. \quad (13)$$

- (ii) Find the (m, q, s) tuple with the largest weight: $(m^*, q^*, s^*) = \arg \max_{(m,q,s)} w_i^{(m,q,s)}(t)$.
- (iii) Processing flow $\mu_i^{(m,q,s)}(t) = (C_i / r^{(m)}) \mathbb{I}\{w_i^{(m^*,q^*,s^*)}(t) > 0, (m, q, s) = (m^*, q^*, s^*)\}$.

B. MEC Scenario

In line with [22], the packet transmission decisions can be modified to handle wireless distributed cloud network settings, e.g., MEC [25], as follows.

1) *Wireless Transmission Model*: Consider a MEC network composed of two types of nodes, i.e., user equipments (UEs) and edge servers (ESs), collected in \mathcal{V}_{UE} and \mathcal{V}_{ES} , respectively. Each ES, equipped with massive antennas, is assigned a separate frequency band of width B_0 , and with the aid of beamforming techniques, it can transmit/receive data to/from multiple UEs simultaneously without interference (assuming that the UEs are spatially well separated) [26]; while each UE is assumed to associate with only one ES at each time slot. For each wireless link (i, j) , the transmission power $p_{ij}(t)$ is assumed to be constant during a time slot (and the maximum power budget of node i is denoted by P_i), incurring cost $(\tilde{e}_i \tau) p_{ij}(t)$, with \tilde{e}_i denoting the *unit energy cost* at node i and τ the time slot length. Besides, we assume that the channel gain $g_{ij}(t)$ is i.i.d. over time and known by estimation, and the noise power is denoted by σ_{ij}^2 . The ESs are connected by wired links, as described in Section II-B.

2) *Wireless Transmission Decision*: In addition to flow assignment $\mu(t) = \{\mu_{ij}^{(q,s)}(t) \geq 0 : (i, j), (q, s)\}$, the control policy makes decisions on power allocation, i.e., $\mathbf{p}(t) = \{p_{ij}(t) \geq 0 : (i, j)\}$, and UE-ES association, i.e., $\chi(t) = \{\chi_{ij}(t) \in \{0, 1\} : (i, j)\}$, where $\chi_{ij}(t)$ indicates if node i is associated with node j ($\chi_{ij}(t) = 1$) or not ($\chi_{ij}(t) = 0$).

The LDP bound (9) remains valid, and the resulting problem (11) is given by

$$\max \sum_{(i,j) \in \mathcal{E}} \chi_{ij}(t) \Psi_{ij}(t),$$

$$\Psi_{ij}(t) \triangleq \sum_{(q,s) \in \Omega} w_{ij}^{(q,s)}(t) \mu_{ij}^{(q,s)}(t) - V(\tilde{e}_i \tau) p_{ij}(t) \quad (14a)$$

$$\text{s. t. } \sum_{j \in \delta_i^+} \chi_{ij}(t) \leq 1 \text{ for } \forall i \in \mathcal{V}_{\text{UE}};$$

$$\chi_{ij}(t) = 1 \text{ for } \forall i \in \mathcal{V}_{\text{ES}}, j \in \delta_i^+, \quad (14b)$$

$$\sum_{(q,s) \in \Omega} \mu_{ij}^{(q,s)}(t) \leq \tau R_{ij}(t),$$

$$R_{ij}(t) \triangleq B_0 \log_2(1 + g_{ij}(t) p_{ij}(t) / \sigma_{ij}^2), \forall (i, j) \quad (14c)$$

$$\sum_{j \in \delta_i^+} \chi_{ij}(t) p_{ij}(t) \leq P_i, \forall i \in \mathcal{V} \quad (14d)$$

where $w_{ij}^{(q,s)}(t) = Q_i^{(q)}(t) - Q_i^{(q-s)}(t) - Q_j^{(s)}(t)$, and (14b) represents the constraint that each UE can only select one ES to communicate with at each time slot. Following the

procedure in [22], we can solve (14) to derive the wireless transmission decisions for each link (i, j) at time t as

$$\text{Power allocation: } p_{ij}^*(t) = \min[p_{ij}(t; 0), P_i], i \in \mathcal{V}_{\text{UE}};$$

$$p_{ij}^*(t) = p_{ij}(t; \nu^*), i \in \mathcal{V}_{\text{ES}}, \quad (15a)$$

$$\text{Flow assignment: } \mu_{ij}^{*(q,s)}(t) = \tau R_{ij}(t) \mathbb{I}\{(q, s) = (q^*, s^*)\},$$

$$(q^*, s^*) = \arg \max_{(q,s)} w_{ij}^{(q,s)}(t), \quad (15b)$$

$$\text{UE-ES association: } \chi_{ij}^*(t) = \mathbb{I}\{j = j^*, \Psi_{ij}^*(t) > 0\}, i \in \mathcal{V}_{\text{UE}};$$

$$\chi_{ij}^*(t) = 1, i \in \mathcal{V}_{\text{ES}} \quad (15c)$$

where $p_{ij}(t; \nu) \triangleq \max[B_0 w_{ij}^{(q^*, s^*)}(t) / (\tilde{e}_i V + \nu) / \ln 2 - \sigma_{ij}^2 / g_{ij}(t), 0]$, $\nu^* = \max\{0, \nu_0\}$ with ν_0 satisfying $\sum_{j \in \delta_i^+} p_{ij}(t; \nu_0) = P_i$, and $j^* = \arg \max_j \Psi_{ij}^*(t)$.

C. EGDCNC with Enhanced Delay

In line with [17], [11], a *biased* queue that incorporates network topology information can be designed to enhance the delay performance of multicast flow control.

In the unicast setting, [17] defines the bias term for each node i , $H_U(i, d)$, as the *minimum hop-distance* between node i and destination d , and combines it with the physical queue as $\tilde{Q}_i(t) = Q_i(t) + \eta H_U(i, d)$. The bias term creates an intrinsic pressure difference that pushes packets along the shortest path to the destination. The parameter η can be found by grid search to optimize the combined effect of hop-distance and queue backlog on the total network delay.

In the multicast setting, we propose to modify the biased queue as

$$\tilde{Q}_i^{(q)}(t) = \|q\|_1 Q_i^{(q)}(t) + \eta H_M(i, q),$$

$$\text{with } H_M(i, q) \triangleq \sum_{k=1}^D q_k H_U(i, d_k). \quad (16)$$

Differently from the unicast case, (i) the bias term $H_M(i, q)$ is now the *sum of minimum hop-distances* to all current destinations, and (ii) the backlog term $Q_i^{(q)}(t)$ is now weighted by $\|q\|_1$ in line with its contribution to the average delay (i.e., the delay a status q packet experiences will impact the delay of the $\|q\|_1$ copies resulting from its subsequent duplications; the detailed proof is given in [2, Appendix E]).

EGDCNC works just like GDCNC, but using $\tilde{Q}(t)$ in place of $Q(t)$ to make forwarding and duplication decisions. Following the procedure in [17], we can show that EGDCNC (using any fixed $\eta \geq 0$) does not lose throughput and cost optimality as compared to GDCNC.

VIII. NUMERICAL RESULTS

A. Network Setup

We consider a MEC network within a square area of $200 \text{ m} \times 200 \text{ m}$, including 9 UEs and 4 ESs. In a Cartesian coordinate system with the origin located at the square center, the ESs are placed at $(\pm 50 \text{ m}, \pm 50 \text{ m})$. Each user is moving according to Gaussian random walk (reflected when hitting the boundary), and the displacement in each time slot distributes in $\mathcal{N}(\mathbf{0}, 10^{-4} \mathbf{I}) \text{ m}$. The length of each time slot is $\tau = 1 \text{ ms}$.

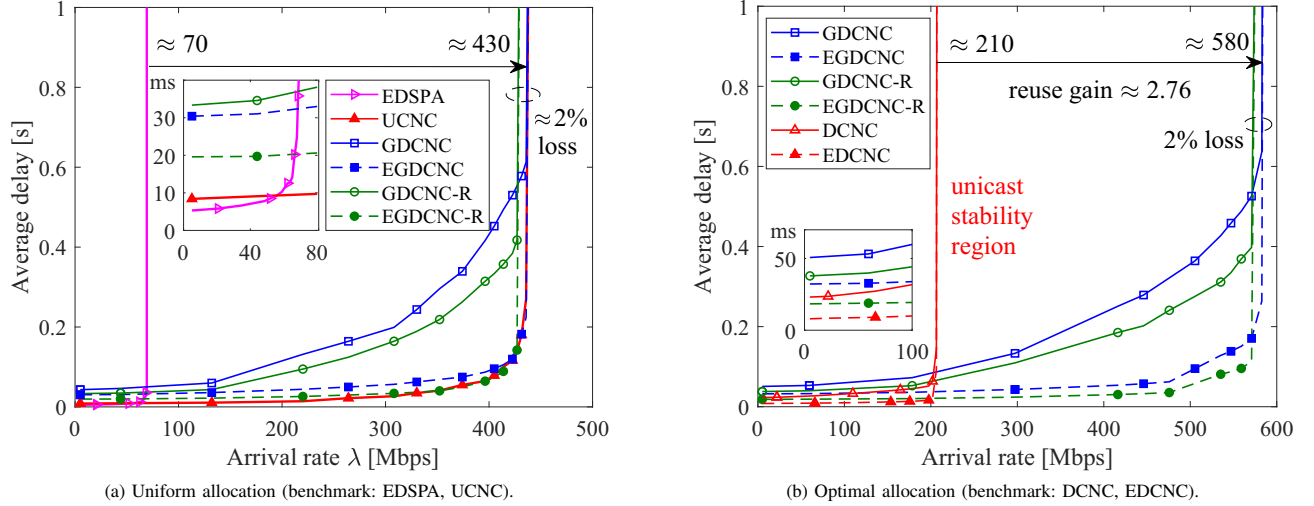


Fig. 3. Network stability regions attained by the benchmark and proposed algorithms.

TABLE I
NETWORK RESOURCES AND OPERATIONAL COSTS OF THE STUDIED SYSTEM

Processing		Transmission	
UE	$C_i = 1$ GHz, $e_i = 2$ /GHz	UE-ES	$i \in \mathcal{V}_{\text{UE}}: P_i = 200$ mW, $\tilde{e}_i = .01$ /J; $i \in \mathcal{V}_{\text{ES}}: P_i = 1$ W, $\tilde{e}_i = .005$ /J
ES	$C_i = 5$ GHz, $e_i = 1$ /GHz	ES-ES	$C_{ij} = 1$ Gbps, $e_{ij} = 1$ /Gbps for (i, j) if $\text{distance}(i, j) = 100$ m

We employ the transmission model described in Section VII-B, where ESs are connected by wired links, while UEs and ESs communicate via wireless links with the following parameters: bandwidth $B_0 = 100$ MHz, path-loss = $32.4 + 20 \log_{10}(f_c) + 31.9 \log_{10}(\text{distance})$ dB with $f_c = 30$ GHz (urban microcell [27]), communication range = 150 m, standard deviation of shadow fading $\sigma_{\text{SF}} = 8.2$ dB, and $\sigma_{ij}^2 = N_0 B_0$ with noise spectral density $N_0 = -174$ dBm/Hz. The processing/transmission capacities and associated costs are shown in Table I.

Consider two AgI services composed of 2 functions with parameters given by (the subscript of each parameter denotes the associated service ϕ , and workload $r_\phi^{(m)}$ is in GHz/Mbps):

$$\begin{aligned} \xi_1^{(1)} &= 1, r_1^{(1)} = \frac{1}{300}, \xi_1^{(2)} = 2, r_1^{(2)} = \frac{1}{400}; \\ \xi_2^{(1)} &= \frac{1}{3}, r_2^{(1)} = \frac{1}{200}, \xi_2^{(2)} = \frac{1}{2}, r_2^{(2)} = \frac{1}{100}. \end{aligned}$$

Each service has 1 source node and $D = 3$ destination nodes, which are randomly selected from the UEs, and the arrival process is modeled as i.i.d. Poisson with λ Mbps.

B. Uniform Resource Allocation

First, we compare the proposed GDCNC algorithms with two state-of-the-art cloud network control policies: UCNC (a throughput-optimal source routing algorithm) [13] and EDSPA (a widely used multicast routing technique) [20]. Since resource allocation exceeds the scope of the benchmark algorithms, the following policy is employed for all algorithms for fair comparisons: each node/link uses maximum processing/transmission capacity, and each ES allocates equal transmission power for each UE. In the GDCNC algorithms,

we select $V = 0$ to optimize the delay performance; in GDCNC-R, we select $K = 1$ duplication tree by k -means clustering using geographic distance; in UCNC, route selection is based on delayed information resulting from hop-by-hop transmission of the queuing states from all nodes to the source.

Fig. 3a depicts the average delay attained by the algorithms under different arrival rates. First, we focus on the throughput performance. We observe an identical critical point (≈ 430 Mbps) for GDCNC, EGDCNC, and UCNC, at which point the average delay blows up, indicative of the stability region boundary, validating Proposition 2, i.e., the throughput-optimality of GDCNC/EGDCNC (because UCNC is throughput-optimal [13]). GDCNC-R/EGDCNC-R achieve sub-optimal throughput by only $\approx 2\%$, illustrating the marginal performance loss from precluded duplication trees. Finally, EDSPA only achieves a maximum rate ≈ 70 Mbps, which is far from the stability region boundary, due to the lack of route diversity.

When looking at the delay performance, we observe that the enhanced variants EGDCNC/EGDCNC-R effectively reduce the delay of the initial solutions GDCNC/GDCNC-R, by adding the shortest path bias to improve the delay performance. In low-congestion regimes, the end-to-end delay mainly depends on the hop-distance of the route; as a result, EDSPA, which delivers the packet along the *shortest* path, and UCNC, which requires the entire path to be *acyclic* for hop-distance reduction, outperform EGDCNC/EGDCNC-R. As arrival rate increases, queuing delay becomes the dominant component, and gaps between EGDCNC/EGDCNC-R and UCNC vanish, demonstrating the good performance of the algorithms in high-congestion regimes.

Finally, we study the complexity of the algorithms. We

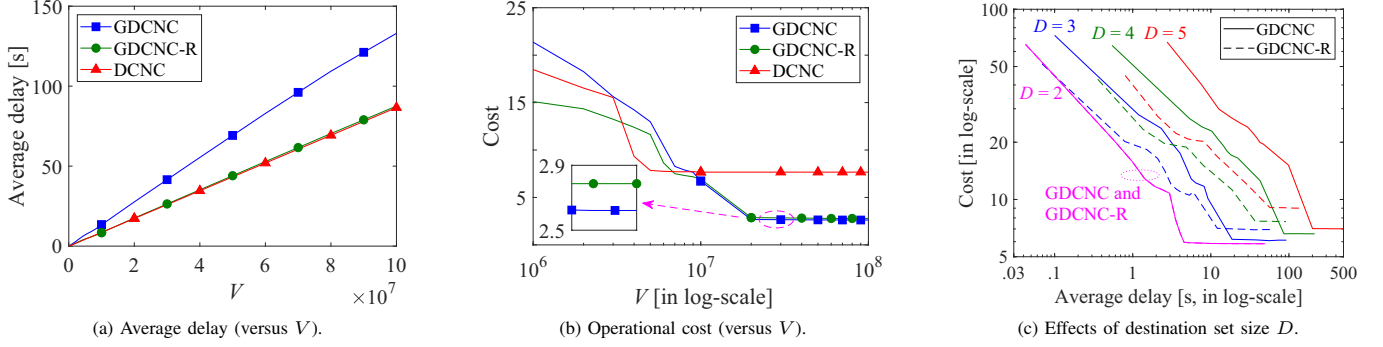


Fig. 4. Delay-cost tradeoffs attained by the three algorithms under different V parameters and destination set sizes D .

first sort the algorithms by *communication overhead*, i.e., required information for decision making. Noting that network topology changes much slower than queuing states, we have: EDSPA (requiring network topology) \ll GDCNC (requiring local queuing states) \approx EGDCNC (requiring local queuing states and network topology) \ll UCNC (requiring global queuing states). We then present the running time of the algorithms as a measure for *computational complexity*: EDSPA (1.3 s) $<$ GDCNC \approx EGDCNC (2.7 s) \ll UCNC (5.8 min), among which: EDSPA runs fastest because it uses fixed routes; GDCNC can operate efficiently to complete simple algebraic operations (10); and UCNC requires solving the NP-complete Steiner tree problem (in a 39-node layered graph) at each time slot, which incurs high computational complexity that can increase when applied to larger networks.

To sum up: in low-congestion regimes, EDSPA is a good choice due to its low complexity and superior delay performance; in high-congestion regimes, UCNC works better for small networks that impose low overhead for information collecting and decision making; and EGDCNC is competitive in all regimes and especially suitable for large-scale distributed cloud networks.

C. Optimal Resource Allocation

Next, we demonstrate the cost performance of the GDCNC algorithms, employing DCNC [11] as the benchmark, which is throughput- and cost-optimal for unicast flow control.

1) *Network Stability Region*: Fig. 3b shows the network stability regions attained by the algorithms. We find that: the designed power allocation policy (15) boosts the stability region (≈ 580 Mbps), compared to *uniform* power allocation (≈ 430 Mbps, as shown in Fig. 3a). Another observation is: compared to DCNC (which attains the *unicast* stability region ≈ 210 Mbps), in the studied setting (with $D = 3$ destinations), the multicast framework enlarges the stability region by a factor of 2.76, which is bounded by D , validating Proposition 1. Finally, the performance loss of GDCNC-R is marginal ($\approx 2\%$) compared to GDCNC, which, together with the results shown in Fig. 3a, validates its competitive throughput performance.

2) *Delay and Cost Performance*: Next, we present the tunable delay and cost performance of the algorithms, under

$\lambda = 150$ Mbps and $\eta = 0$ (we fix η and focus on the effects of V).

As shown in Fig. 4a, the average delay attained by each algorithm increases linearly with V . Note however that DCNC can achieve a better delay than GDCNC (and almost the same as GDCNC-R) for arrival rates within the unicast stability region (low congestion regime in the enlarged multicast stability region), because it can select separate paths for each copy to optimize the individual delays. This is expected, as the throughput-optimal design of GDCNC, which *jointly* selects the copies' paths to reduce network traffic and enlarge the stability region, can lead to higher delay in low-congestion regimes.

Fig. 4b shows the reduction in operational cost with growing V , validating Theorem 2. By pushing $V \rightarrow \infty$, the curves converge to the corresponding optimal costs, given by: DCNC (7.7) $>$ GDCNC-R (2.8) \approx GDCNC (2.65). For $V < 10^7$, GDCNC-R attains an even lower cost than GDCNC, even though it has sub-optimal *asymptotic* performance (e.g., $V > 10^7$). However, note that a large V also leads to an excessive delay, making it a sub-optimal choice in practical systems. For example, when increasing V from 10^6 to 10^7 , the cost attained by GDCNC reduces from 21.4 to 6.7, while the delay grows from 1.3 to 13.4 seconds (these results are for comparison purpose, and can be improved by EGDCNC, as shown in Fig. 3a and 3b).

3) *Effects of Destination Set Size*: Finally, we present in Fig. 4c the delay-cost tradeoffs attained by the GDCNC algorithms, under different destination set sizes varying from $D = 2$ to 5, and $\lambda = 300$ Mbps.

We make the following observations. A larger destination set size D results in (i) increasing delay and cost, since more network resources are consumed to handle copies for additional destinations, which also results in longer time waiting for available network resources, (ii) growing running time, due to the greater number of duplication choices, e.g., GDCNC (3.2 s) \approx GDCNC-R (2.9 s) when $D = 3$, and GDCNC (8.9 s) $>$ GDCNC-R (5.8 s) when $D = 5$, (iii) a widening gap between the average delay performance of GDCNC and GDCNC-R (in order to attain a given cost, e.g., 20), which validates the benefit of GDCNC-R reducing the queuing system and improving the delay performance (which is observed also in Fig. 3a and 3b).

Again we emphasize that the selection of parameter V targeting optimal cost performance can lead to excessive delay (tens of seconds when $D = 3$), which is a sub-optimal choice in practical systems. With appropriate V values, GDCNC-R can attain suitable delay-cost performance pairs $(0.8\text{ s}, 20)$, $(1.3\text{ s}, 23)$, $(2\text{ s}, 27)$ under destination set sizes $D = 3, 4, 5$, respectively.

To sum up, although we cannot provide an analytical bound on the performance loss of GDCNC-R, numerical results validate that it can remain competitive in throughput performance (with negligible performance loss), while striking an even better delay-cost tradeoff than GDCNC.

IX. CONCLUSIONS

We addressed the problem of decentralized control of mixed-cast AgI services in distributed cloud networks. Under the multicast framework, we characterized the enlarged network stability region and analyzed the benefit of exploiting in-network packet duplication. By extending LDP control to a novel queuing system that allows accommodating the duplication operation and making flow control decisions driven by data packets' current destinations, we designed the first decentralized, throughput- and cost-optimal packet processing, routing, and duplication policy for generalized (multicast) flow control, GDCNC, as well as practical variants targeting reduced complexity, enhanced delay, and extended scenarios. Via numerical experiments, we validated the performance gain attained via effective in-network packet duplication, as well as the benefits of joint processing, routing, and duplication optimization for the efficient delivery of multicast AgI services over distributed cloud networks.

APPENDIX A

Condition c) Duplication: For any policy \mathcal{P}_1 that can create (and operate on) multiple copies at each time slot (via replication), consider a policy \mathcal{P}_2 that makes the same replication and routing decisions as \mathcal{P}_1 , but creating only two copies at each time slot (via duplication) and hence spending multiple slots to complete a replication operation that creates more than two copies under \mathcal{P}_1 . It follows that \mathcal{P}_2 achieves the same network flow (see [2, Appendix A] for detailed proof), and associated throughput and operational cost, as \mathcal{P}_1 .

Condition d) Non-consolidation: For any policy \mathcal{P}_1 that includes consolidation (co-located packets of identical content are combined into one), a policy \mathcal{P}_2 prevented from performing this operation is presented as follows. Suppose that under \mathcal{P}_1 , two copies of status $(q_1 + s_1)$ and $(q_2 + s_2)$ are created via duplication at node j , sent along different routes while undergoing additional duplications to deliver copies to destinations in s_1 and s_2 , respectively, and then rejoined at node i with updated status q_1 and q_2 , where they get combined. Then, in \mathcal{P}_2 , we can create two copies with status $(q_1 + q_2 + s_1)$ and s_2 at node j and send them along the same routes. After duplications for destinations in s_1 and s_2 , only the first packet of updated status $(q_1 + q_2)$ will arrive to node i . Such operation leads to the same end state without incurring additional resource consumption, thus yielding the same throughput and cost performance.

APPENDIX B

GDCNC: The complexity is proportional to the number of duplication choices (q, s) , given by $\sum_{k=1}^D C(k; D)(2^k - 1) = 3^D - 2^D \sim \mathcal{O}(3^D)$, where $C(k; D)$ (i.e., the combinatorial number of *choosing k from D*) denotes the number of D -dimensional vectors with k entries equal to 1, and $2^k - 1$ is the cardinality of $2^q \setminus \{0\}$, i.e., the number of possible vectors s for a given q .

GDCNC-R: In this part, we show: each duplication tree includes $D - 1$ internal nodes (and evidently, D leaf nodes). Note that (i) each tree involves $D - 1$ duplication operations, and each operation involves 3 tree nodes, (ii) when counting the tree nodes based on duplication operations, every node occurs twice (one time as the parent, and the other time as a child), other than the root node and the leaf nodes (which are counted once). Therefore, the number of internal nodes (including the root node) is $[3 \times (D - 1) - 1 \times (D + 1)]/2 + 1 = D - 1$.


REFERENCES

- [1] Y. Cai, J. Llorca, A. M. Tulino, and A. F. Molisch, "Optimal multicast service chain control: Packet processing, routing, and duplication," in *Proc. IEEE Int. Conf. Commun.*, Montreal, Canada, Jun. 2021, pp. 1–7.
- [2] Y. Cai, J. Llorca, A. M. Tulino, and A. F. Molisch, "Decentralized control of distributed cloud networks with generalized network flows," arXiv:2204.09030. [Online]. Available: <https://arxiv.org/abs/2204.09030>, Apr. 2022.
- [3] Y. Cai, J. Llorca, A. M. Tulino, and A. F. Molisch, "Compute- and data-intensive networks: The key to the Metaverse," in *2022 1st International Conference on 6G Networking (6GNet)*, Paris, France, Jul. 2022, pp. 1–8.
- [4] M. Weldon, The future X network: A Bell Labs perspective. Boca Raton, FL, USA: CRC Press, 2016.
- [5] M. Huang, W. Liang, Y. Ma, and S. Guo, "Maximizing throughput of delay-sensitive NFV-enabled request admissions via virtualized network function placement," *IEEE Trans. Cloud Comput.*, vol. 9, no. 4, pp. 1535–1548, Oct.-Dec. 2021.
- [6] Y. Yue, B. Cheng, M. Wang *et al.*, "Throughput optimization and delay guarantee VNF placement for mapping SFC requests in NFV-enabled networks," *IEEE Trans. Netw. Service Manag.*, vol. 18, no. 4, pp. 4247–4262, Dec. 2021.
- [7] M. Barcelo, J. Llorca, A. M. Tulino, and N. Raman, "The cloud service distribution problem in distributed cloud networks," in *Proc. IEEE Int. Conf. Commun.*, London, UK, May 2015, pp. 344–350.
- [8] M. Barcelo, A. Correa, J. Llorca, A. M. Tulino, J. L. Vicario, and A. Morell, "IoT-cloud service optimization in next generation smart environments," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 12, pp. 4077–4090, Oct. 2016.
- [9] H. Feng, J. Llorca, A. M. Tulino, D. Raz, and A. F. Molisch, "Approximation algorithms for the NFV service distribution problem," in *Proc. IEEE INFOCOM*, May 2017, pp. 1–9.
- [10] Y. Kim, J. Kwak, H.-W. Lee, and S. Chong, "Dynamic computation and network chaining in integrated SDN/NFV cloud infrastructure," *IEEE Trans. Cloud Comput.*, pp. 1–1, 2021.
- [11] H. Feng, J. Llorca, A. M. Tulino, and A. F. Molisch, "Optimal dynamic cloud network control," *IEEE/ACM Trans. Netw.*, vol. 26, no. 5, pp. 2118–2131, Oct. 2018.
- [12] H. Feng, J. Llorca, A. M. Tulino, and A. F. Molisch, "Optimal control of wireless computing networks," *IEEE Trans. Wireless Commun.*, vol. 17, no. 12, pp. 8283–8298, Dec. 2018.
- [13] J. Zhang, A. Sinha, J. Llorca, A. M. Tulino, and E. Modiano, "Optimal control of distributed computing networks with mixed-cast traffic flows," *IEEE/ACM Trans. Netw.*, vol. 29, no. 4, pp. 1760–1773, Aug. 2021.
- [14] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
- [15] M. J. Neely, *Stochastic network optimization with application to communication and queueing systems*. San Rafael, CA, USA: Morgan & Claypool, 2010.

- [16] L. Bui, R. Srikant, and A. Stolyar, "Novel architectures and algorithms for delay reduction in back-pressure scheduling and routing," in *Proc. IEEE INFOCOM*, Rio de Janeiro, Brazil, Apr. 2009, pp. 2936–2940.
- [17] L. Ying, S. Shakkottai, A. Reddy, and S. Liu, "On combining shortest-path and back-pressure routing over multihop wireless networks," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 841–854, Jun. 2011.
- [18] S. Islam, N. Muslim, and J. W. Atwood, "A survey on multicasting in software-defined networking," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 1, pp. 355–387, Firstquarter 2018.
- [19] J. Kleinberg and Éva Tardos, *Algorithm design*. Delhi, India: Pearson Education India, 2006.
- [20] M. T. Ananta, J.-R. Jiang, and M. A. Muslim, "Multicasting with the extended Dijkstra's shortest path algorithm for software defined networking," *Int. J. Appl. Eng. Res.*, vol. 9, no. 23, pp. 21 017–21 030, 2014.
- [21] A. Sinha and E. Modiano, "Optimal control for generalized network flow problems," *IEEE/ACM Trans. Netw.*, vol. 26, no. 1, pp. 506–519, Feb. 2018.
- [22] Y. Cai, J. Llorca, A. M. Tulino, and A. F. Molisch, "Mobile edge computing network control: Tradeoff between delay and cost," in *Proc. IEEE Global. Telecomm. Conf.*, Taipei, Taiwan, Dec. 2020, pp. 1–6.
- [23] E. Alpaydin, *Introduction to machine learning*. Cambridge, MA, USA: MIT press, 2020.
- [24] M. Meilă and W. Pentney, "Clustering by weighted cuts in directed graphs," in *Proceedings of the 2007 SIAM international conference on data mining*, SIAM, 2007, pp. 135–144.
- [25] M. Chen and Y. Hao, "Task offloading for mobile edge computing in software defined ultra-dense network," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 3, pp. 587–597, Mar. 2018.
- [26] A. Adhikary, J. Nam, J.-Y. Ahn, and G. Caire, "Joint spatial division and multiplexing – The large-scale array regime," *IEEE Trans. Inf. Theory*, vol. 59, no. 10, pp. 6441–6463, Oct. 2013.
- [27] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," *3rd Generation Partnership Project (3GPP), Tech. Rep.*, vol. 38, 2018.



Yang Cai (Graduate Student Member, IEEE) received his B.E. and M.S. degrees from the Department of Electronic Engineering, Tsinghua University, in 2015 and 2018, respectively. Now he is a Ph.D. candidate in the Department of Electrical and Computer Engineering at University of Southern California, supported by the Annenberg Graduate Fellowship award. He is currently working on modeling, analysis, evaluation, and control of next-generation computing networks and services.



Jaime Llorca (Member, IEEE) is a Technology Consultant and Research Professor with the NYU Tandon School of Engineering, NY. He previously held a Senior Research Scientist position with the Network Algorithms Group at Nokia Bell Labs, NJ, a Research Scientist position with the End-to-End Networking Group at Alcatel-Lucent Bell Labs, NJ, and a post-doctoral position with the Center for Networking of Infrastructure Sensors, MD. He received M.S. and Ph.D. degrees in Electrical and Computer Engineering from the University of Maryland, College Park, MD. His research interests are in the field of network algorithms, optimization, machine learning, and distributed control, with applications to next generation communication networks, distributed/edge cloud, and content distribution. He has made fundamental contributions to the mathematics of content delivery and distributed cloud networks, including pioneering cooperative caching, network coding, and cloud network control algorithms. He has authored more than 100 peer-reviewed publications and 20 patents. He currently serves as Associate Editor for the IEEE/ACM Transactions on Networking. He is a recipient of the 2007 IEEE ISSNIP Best Paper Award, the 2016 IEEE ICC Best Paper Award, and the 2015 Jimmy H.C. Lin Award for Innovation.



Antonia M. Tulino (Fellow, IEEE) is a full professor at Università degli Studi di Napoli Federico II. She held research positions with the Center for Wireless Communications in Oulu, Princeton University, a Università degli Studi del Sannio, Italy and Bell Labs, NJ. Since 2019, she is Research Professor at Dep. of Electrical and Computer Engineering NYU Tandon School of Engineering, and director of the 5G Academy jointly organized by the Università degli Studi di Napoli, Federico II and several leading company in digital transformation. Her research interests lay in the area of communication networks approached with the complementary tools provided by signal processing, information theory, and random matrix theory. From 2011 to 2013, she has been a member of the Editorial Board of the IEEE Transactions on Information Theory and in 2013, she was elevated to IEEE Fellow. From 2019 to 2021, she has been chair of the Information Theory society Fellows Committee.

She has received several paper awards, including the 2009 Stephen O. Rice Prize in the Field of Communications Theory. She was recipient of the UC3M-Santander Chair of Excellence from 2018 to 2019 and selected by the National Academy of Engineering for the Frontiers of Engineering program in 2013.



Andreas F. Molisch (Fellow, IEEE) received his degrees (Dipl.Ing. 1990, PhD 1994, Habilitation 1999) from the Technical University Vienna, Austria. He spent the next 10 years in industry, at FTW, AT&T (Bell) Laboratories, and Mitsubishi Electric Research Labs (where he rose to Chief Wireless Standards Architect). In 2009 he joined the University of Southern California (USC) in Los Angeles, CA, as Professor, and founded the Wireless Devices and Systems (WiDeS) group. In 2017, he was appointed to the Solomon Golomb – Andrew

and Erna Viterbi Chair.

His research interests revolve around wireless propagation channels, wireless systems design, and their interaction. Recently, his main interests have been wireless channel measurement and modeling for 5G and beyond 5G systems, joint communication-caching-computation, hybrid beamforming, UWB/TOA based localization, and novel modulation/multiple access methods. Overall, he has published 5 books (among them the textbook "Wireless Communications", third edition in 2022), 21 book chapters, 300 journal papers, and 370 conference papers. He is also the inventor of 70 granted (and more than 10 pending) patents, and co-author of some 70 standards contributions. His work has been cited more than 62,000 times, his h-index is > 100, and he is a Clarivate Highly Cited Researcher.

Dr. Molisch has been an Editor of a number of journals and special issues, General Chair, Technical Program Committee Chair, or Symposium Chair of multiple international conferences, as well as Chairperson of various international standardization groups. He is a Fellow of the National Academy of Inventors, Fellow of the AAAS, Fellow of the IEEE, Fellow of the IET, an IEEE Distinguished Lecturer, and a member of the Austrian Academy of Sciences. He has received numerous awards, among them the IET Achievement Medal, the Technical Achievement Awards of IEEE Vehicular Technology Society (Evans Avant-Garde Award) and the IEEE Communications Society (Edwin Howard Armstrong Award), and the Technical Field Award of the IEEE for Communications, the Eric Sumner Award.