LGCC: Food Chain Multi-Hop Congestion Control

Peyman Teymoori, Michael Welzl

LGCC: Food Chain Multi-Hop Congestion Control

Peyman Teymoori, Michael Welzl

Department of Informatics

University of Oslo, Norway

{peymant|michawe}@ifi.uio.no

Abstract—**Technological advancements have provided wireless links with very high capacity for 5G mobile networks and WiFi 6, which will be widely deployed by 2025; however, the capacity heavily fluctuates, violating the assumption at the transport layer that the capacity is (almost) fixed. This issue has so far been targeted by mostly "hack-and-patches", which are still far from an efficient design. In this paper, we present a general and efficient, yet deployable solution to the aforementioned problem through a novel design empowered with a rich theory, allowing a significantly-improved experience in using new technologies, especially mobile cellular services.**

Index Terms—**Multi-hop Congestion Control, Logistic Growth.**

I. Introduction

In the Internet design, congestion control has been deployed at the transport layer or above, and it is supposed to operate in an "end-to-end" manner, where "end" represents the application host [1]. The commonly used transport protocol, TCP, reacts to packet drops as the main congestion signal from the network. This design greatly impedes the consideration of link-specific properties. Today, wireless links are prevalent, and the capacity that they expose to TCP are often time-varying, especially with new link layer technologies.

Notable technological advancements in communication have made it possible to exploit millimeter waves (mmWaves) [2] to benefit from the very high data rates it can provide. The mmWave spectrum has also been considered for next generation wireless networking such as cellular networks (e.g. 5G). mmWave communication is at odds with TCP's assumption of a relatively stable end-to-end path with a static capacity at the bottleneck. It promises extremely low latency, yet it will often expose a drastically fluctuating network capacity. TCP should be able to react to these fluctuations—but the reaction speed of the end-to-end control loop is inevitably in the order of end-to-end round-trip times (RTTs). Such a control loop cannot ensure low latency and good utilization in the face of capacity changes that occur at sub-RTT timescales [3].

Even with today's 4G networks, it is already common to shorten the control loop by installing Performance Enhancing Proxies (PEPs)—devices which interfere with TCP connections, e.g. by dividing them in two (acting as a receiver towards a sender and acting as a sender towards a receiver) [4]. Splitting an end-to-end connection like this (or with an application-layer proxy, which is currently under discussion in the IETF for the QUIC protocol [5]) has shown benefits, and it may even be the only feasible way to attain efficient communication over mmWave links [6], [7]. Most importantly

perhaps, connection splitting eliminates the need to operate "normal" TCP congestion control throughout the path, allowing the deployment of a more suitable congestion control mechanism on the wireless segment only.

What, then, should a mechanism for the wireless segment look like? Rather than suggesting a quick fix for one specific link layer technology, we argue that it is necessary to design a *general* congestion control mechanism which:

- can swiftly adapt to changing physical capacity,
- will stably operate when multiple instances of the mechanism are used in sequence (e.g., when transferring data from one wireless segment to another), as a chain of control loops, and
- has a low deployment barrier, e.g. by not requiring to update every router along the path.

To the best of our knowledge, this paper is the first attempt at devising a congestion control mechanism that satisfies all of these requirements.

In our previous work [8], we investigated various forms of feedback for a recursive congestion control as in the Recursive Internetwork Architecture (RINA); the main finding was that the use of overflowing queues as a feedback method can produce large delays in such scenarios. Thus, similar to earlier explicit feedback methods such as XCP or RCP [9], [10], we opt for an approach that does not rely on physical queue growth for feedback. These mechanisms were found not to be deployable in practice as they require to change all routers on an end-to-end path. Our situation is more favorable, as we aim to apply our mechanism on path segments only. By basing the feedback method on the Datacenter TCP (DCTCP) [11] style of ECN usage (interpreting ECN marks from multiple packets as a multi-bit congestion signal), we made it even easier to deploy.

Specifically, the method that we propose, called "Logistic Growth Chain Control" (LGCC):

- can be deployed by only updating end systems, and introducing some devices that support it at loop intersection points (a pronounced difference from explicit rate feedback schemes like XCP or RCP); within each control loop, commodity hardware can be used, provided that this hardware supports Random Early Detection (RED) AQM marking with Explicit Congestion Notification (ECN).
- is general enough to be adopted in different use cases, and can be easily enhanced by feeding it more information from the network. For example, it is possible to further accelerate convergence if there are more bits available in

the packet header than ECN. Because we use an *additive* signal akin to *cost* as defined in the NUM framework [12], it would also be possible for routers to mark packets before a physical queue even grows (as proposed in [13], for instance).

- is globally asymptotically stable.
- is fair, and can attain different types of fairness.
- is fast in terms of convergence due to the use of the LG function. Since we shorten control loops, flows can converge faster in case of capacity decrease/increase.
- responds quickly to capacity changes since the capacity can be measured more directly than by inferring it based on queue-related metrics such as ECN and delay. LGCC uses the concept of *carrying capacity*, the maximum available resource, denoted by K. In a 5G case, K is time varying, so we add a time parameter, i.e. $K(t)$. There are already useful signals at lower layers that can feed into $K(t)$—e.g., various MAC solutions calculate a physical rate using rate adaptation techniques, making the current physical rate a promising candidate as input to $K(t)$.

LGCC is based upon an early version of our mechanism, the biology-inspired Logistic Growth Control (LGC) [14], combined with additive ECN feedback [15]. Different from prior work, we now assume multiple consecutive control loops between each sender and receiver, one per administrative domain. This allows us to react much faster to congestion than with an end-to-end loop, but it is a more complex design approach which inevitably raises stability concerns. A stability analysis is, thus, a key contribution of this paper.

In Section II, we survey related work. Section III discusses our motivations behind this paper, and Section IV introduces the theory behind LGCC. We establish the theory of LGCC in Section VI, and in Section VIII, we analyze stability and fairness properties of LGCC. Section IX presents how LGCC is implemented, and its simulation results are presented in Section X. Finally, Section XI concludes the paper.

II. RELATED WORK

The advent of high capacity wireless links (g.e. mmWave) has raised serious challenges for traditional congestion control [16]. In a series of work [3], [7], [17]–[20], it was shown that the full potential of mmWave cannot be achieved at higher layers, especially at the transport layer.

One of the main solutions of the above problem has been to introduce some intermediate boxes helping the end-to-end congestion controller, e.g. by breaking it into several shorter loops. For example, in [6], a TCP proxy architecture was presented without any modification at the sender side. The proxy is installed in the Radio Access Network (RAN), and by collecting statistics from the 5G base station (gNB), it estimates the downlink Bandwidth-Delay Product (BDP) and performs flow window management using the advertised window field in ACKs. Another method, called X-TCP [21], was proposed to adjust the congestion window of TCP at the mobile equipment side. It estimates the current physical layer rate using the Signal to Interference plus Noise Ratio (SINR), thus keeping the buffer small.

Since the advent of hop-by-hop congestion controllers, as the extreme case of multihop, their stability has been of paramount concern because some unstable behavior was observed in some situations [22], [23]. This had led to the "fear of instability". However, as demystified by [23], the cause of instability was the unsuccessful use of non-discriminatory on-off type controls. In addition, research work such as [22]–[26] has shown that per-hop controllers accelerate the response to load changes in the path, and that certain conditions can be obtained under which the chain of control loops is stable. For example, in case of Split-TCP (where a middlebox spoofs IP addresses to act like a receiver towards the sender and act like a sender towards the receiver), the system can be stable under certain conditions [24], [27].

The assumption of having links with time-varying capacity also complicates stability conditions of the congestion controller. Analyses such as [28] show system stability with feedback delay can be achieved in the joint congestion control and scheduling for multi-hop wireless networks. In [29], a family of optimization-based distributed traffic control laws were developed to enable load balancing and/or rate adaptation in response to capacity variations in the network. Instead of proving system stability around a single equilibrium point, the authors in [30] propose a primal-dual congestion control algorithm which is proven to be trajectory stable when link capacities are time-varying; this guarantees the system is stable around a time varying reference trajectory.

Explicit feedback control protocols (e.g. XCP [9] and RCP [10]) can overcome some of the above challenges. However, they have deployment limitations in the Internet because they need drastic changes to all packet headers, routers, and end-points. The ABC mechanism [31], [32] has targeted the above problems by providing a deployable method, compared with prior explicit schemes, since it utilizes the existing Explicit Congestion Notification (ECN) [33] infrastructure. This makes the mechanism the closest competitor of the one that we present in this paper. ABC uses a Multiplicative-and-Additive-Increase/Multiplicative-Decrease (MAIMD) scheme to ensure fairness among flows. ABC adjusts the sender's window upon receiving an ACK from the receiver: it is increased by one if the ACK is marked as *accelerate*[1], and in case of a *brake* mark, it is decreased by one; this allows the sender to have a larger window size dynamic within one RTT to follow the link capacity dynamics faster. In the network, ABC routers mark packets with accelerate or brake proportional to the queue length. Additionally, ABC has a fall-back mechanism to co-exist with non-ABC traffic.

New network architectures such as the Recursive InterNetwork Architecture (RINA) [34] represent a different approach to many aspects of networking, including congestion control: here, network mechanisms are designed from scratch, without compromise.[2] In RINA, it is natural to break an end-to-end path into several shorter control loops, each with its own

---

[1]To solve the problem that "AQM schemes do not signal increases" [32].

[2]Compromises will usually be necessary when trying to interconnect the whole architecture with the Internet. This can, for example, be done by tunneling, using gateways or with direct switch-over as explained in [35] for relatively short communication paths.

Fig. 1. A simple topology used for a preliminary comparison of DCTCP, LGC, ABC, and LGCC. The links have a capacity of 500 Mbps, and Source's end-to-end RTT is 6 ms. The link capacity between Router 2 and Destination is time-varying; it drops to one third from second 2 to 3 of the simulation. Other parameters are set such that 100% utilization with minimal delay is achieved for each method.

customized congestion controller. In addition, a larger control loop can operate over several shorter ones; this is the recursive property of RINA, which can provide many benefits [36]. What we propose in this paper can also be adopted by RINA.

III. MOTIVATION

In this section, we motivate the need for a deployable multi-hop congestion controller with preliminary simulation and evaluation; this shows why the controller needs to be fast, in both rate increase and decrease, and also in terms of convergence rate, with the direct knowledge of link capacity and its changes as well as employing ECN to increase deployability.

In previous sections, we discussed about limitations of end-to-end congestion control and why shortening control loops can respond faster to congestion and capacity change. ABC is a very recent state-of-the-art mechanism, which, according to its authors, operates at the Pareto frontier when compared with other congestion controllers such as Cubic, Vegas, XCP, BBR, and PCC [32]. Although ABC is empowered by ABC routers on the path to signal a rate increase/decrease via ECN, it still operates in an end-to-end manner. This means that ECN marks first arrive at the destination, and then they are echoed back to the sender. Considering the topology illustrated in Fig. 1, in Fig. 2 we show a preliminary simulation[3] comparison of LGCC with ECN-based mechanisms: LGC (our previous end-to-end approach [14]), DCTCP [11], and ABC. We observe that DCTCP produces the largest queue because it cannot incorporate link capacity, its changes, and fast increase; LGC can incorporate link capacity, but it cannot track capacity changes, and starts fluctuating when the capacity drops, leading to a larger queue; ABC only incorporates link capacity at ABC routers, but still the ABC sender can track changes faster. LGCC can, however, outperform all of the others in terms of rate fluctuation, convergence speed, and keeping the queue small because of a multi-hop design where Router 1 and Router 2 act as connection splitters, shortening the control loop, and incorporating the link capacity and its changes directly in their rate dynamic.

In terms of convergence rate, we need a fast method for flows to reach their target rate. Although ABC flows can follow the underlying link capacity fast with the help of ABC routers, they do not get a fair share of the capacity as quickly as they fill the capacity. The main problem is that when a new flow joins, to attain fairness among flows, ABC's MAIMD

adds/subtracts less than 1 packet per RTT to/from the cwnd of the flows, and the increase/decrease becomes smaller as flows converge.

We evaluate ABC's convergence rate through a simple analysis: assume n flows with equal RTTs and congestion window w_i, competing for a bottleneck capacity. ABC's congestion window of flow i increases by $1 + 1/w_i$ per accelerate and decreases $-1 + 1/w_i$ per brake signal. Therefore, during an RTT where the percentage of accelerates is f, the window increases by $f(w_i + 1)$ and decreases by $(1 - f)(w_i + 1)$. Adding these two changes to w_i to obtain the new window size, and considering the bottleneck capacity is completely filled, we have

$$n \sum_{i=1}^{n} \left(w_i + w_i f^* - w_i(1 - f^*) + 1 \right) = \text{BDP} . \quad (1)$$

Solving for f^* yields

$$f^* = \frac{\text{BDP} - n}{2 \sum_{i=1}^{n} w_i} = \frac{1}{2} \frac{\text{BDP} - n}{\text{BDP}} < \frac{1}{2} .$$

As an example, if $n = 2$ and BDP $= 200$, then $f^* = 0.495$. If $w_1 = 198$ and $w_2 = 2$, in the next RTT, $w_1 = w_1 + w_1(2f - 1) + 1 = 197.02$, and $w_2 = w_2 + w_2(2f - 1) + 1 = 2.98$. Since $(2f - 1) < 0$, the only term that increases w_2 is 1. This means that in each RTT, the w_2 increase is less than 1, and it becomes smaller as w_1 and w_2 get closer. For example, after 50 RTTs, $w_2 = 40$ and $w_1 = 160$, and after 400 RTTs, they are approximately equal, which can be a very long time. Simulation results also confirm this. In case of LGCC, after 25 RTTs we have $w_2 = 40$ and $w_1 = 160$, and after 60 RTTs, $w_1 \approx w_2$. It should be noted that RTTs are shorter in LGCC because of shorter control loops. Fig. 3 shows a convergence comparison of ABC and LGCC. The first flow fills the capacity, and then the second flow joins. We can observe that it takes a long time for the ABC flows to converge. However, this is shorter than half a second for LGCC. In the following sections, we elaborate on the design of LGCC and show how it incorporates link capacity, capacity changes, and ECN in its rate dynamic.

IV. LOGISTIC GROWTH (LG)

The base of LGCC is the Logistic Growth (LG) function. LG is usually used to model the growth of a species population over time. This function is defined by the following differential equation

$$\dot{N} = \frac{\gamma N(t)(K - N(t))}{K} \quad (2)$$

where N is the population size, K the "carrying capacity" determining the maximum value of N, and γ is growth rate. By the dot notation, we mean time differentiation, i.e. $\dot{N} = \frac{dN}{dt}$. We summarize the notation in Table I. The idea of growth constrained by a capacity limit in LG has inspired the use of this function in network congestion control [14], [37], [38].

The LG function (and its dynamics (2)) is usually normalized by defining $x(t) = N(t)/K$, which yields

$$\dot{x} = \gamma x(t)\bigl(1 - x(t)\bigr) \quad (3)$$

---

[3]It is performed in OMNeT++; we used the code we had implemented in [14] for DCTCP and LGC. We re-implemented ABC, and validated it vs. the results in [32].

(a) DCTCP send rate

(b) LGC send rate

(c) ABC send rate

(d) LGCC send rate

(e) DCTCP bottleneck link

(f) LGC bottleneck link

(g) ABC bottleneck link
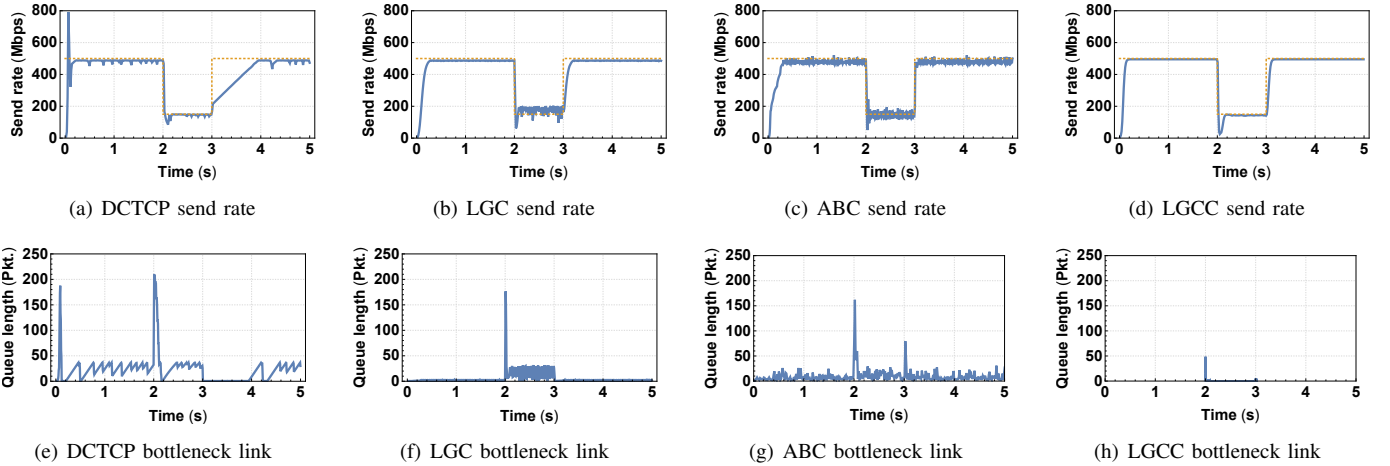
(h) LGCC bottleneck link

Fig. 2. Performance comparison of 1 flow of DCTCP, LGC, ABC, and LGCC. The dotted line in the top row illustrates the link capacity, and the blue lines show the senders' rate using each of the aforementioned methods.


(a) ABC

(b) LGCC

Fig. 3. Convergence comparison of two flows

TABLE I
NOTATIONS

Symbol	Description
K	Carrying capacity
R	The set of sources
L	The set of links
Q	The set of queues
A	The set of aggregating points/nodes
x_r	Send rate of source r
γ	Growth rate
λ_l	Lagrange multiplier of link l
p_l	Marking probability of link l
$U_r(.)$	Utility function of source r
R_l	The set of the sources crossing link l
L_r	The set of the links source r crosses
O_r	The number of loops of source r to reach its destination
$L_{r,i}$	The set of the links source r crosses in loop i
q_r	The total cost source r is charged by links L_r
q_r	The total marking probability of source r traversing links L_r
$p_{r,i}$	The marking probability received by flow r in loop i
c_l	Capacity of link l
b_l	Buffer size of link l
y_l	The sum of rates of sources crossing link l
$.^*$	Optimum value of . that optimizes some function
$\varrho_r(x_r)$	Step size function of source r
$\sigma_l(p_l)$	Step size function of link l
β	Step size parameter
γ	Step size parameter
ϕ	Logarithm base parameter

with the solution

$$x(t) = \frac{1}{1 + \left(\frac{1}{x(0)} - 1\right)e^{-\gamma t}} \tag{4}$$

where $x(0)$ denotes the initial population. Clearly,

$\lim_{t \to \infty} x(t) = 1$.

A. Lotka-Volterra Model

The LG function specifies how one species population grows over time. The Lotka-Volterra model consists of a set of equations examining the effect of inter-species competition where two or more species compete for some limited resource. The growth rate of a species is then not only bounded by the carrying capacity, but also by the growth rate of a competing species. In Section VI, we will use this relationship to model the interconnected loops of a path like individual species, using this relationship to interconnect them. If a generalized Lotka-Volterra model for n interacting species has a non-trivial equilibrium, $x_i^* > 0$, then there can be found a Lyapunov function that is zero at the equilibrium and negative at other points [39]. This implies that the model is globally stable.

V. LOGISTIC GROWTH CONTROL (LGC)

LGC [14] was presented as a congestion controller using the Lotka-Volterra competition model [40]. In this controller, the rate of source r is updated by

$$\dot{x}_r = x_r \gamma_r \left(1 - x_r - \hat{p}_r\right) \tag{5}$$

or in the discrete form,

$$x_r[n+1] = x_r[n]\gamma_r\left(1 - x_r[n] - \hat{p}_r[n]\right) + x_r[n] \tag{6}$$

where γ_r is a constant value and $\hat{p}_r$ is the (end-to-end) measured marking probability at source r. This yields a normalized rate which sources can multiply by their link capacity to obtain a send rate in bits per second. Since LGC needs the link capacity, it was proposed for data centers in which link capacities are known. If there is only one bottleneck on the path, it is proven that $\hat{p}_r$ converges to $\frac{S-1}{S}$ where S is the number of competing flows, and hence, $x_r^* = \frac{1}{S}$.

To illustrate how LGC works, we numerically evaluate (6): we assume that there are 5 flows competing for a bottleneck

(a) Rate   (b) Marking probability

Fig. 4. 5 LGC competing flows with $\gamma = 1$

capacity, and excess traffic is marked proportional to the capacity. Formally speaking,

$$p_r = \left[p_r + \left(\sum_{i=1}^{5} x_i - 1 \right) \right]_0^1 \qquad (7)$$

where $[.]_0^1 = \max(\min(.,1),0)$. Fig. 4 shows how the 5 flows compete when they join one by one. The first flow starts at time step 0, and then after every 20 time steps, one new flow joins. We see that the normalized rate of flows converges to 0.2, and their observed marking probability converges to 0.8. At step 150, flows leave one by one every 20 steps until there is only one flow left. We can also see that the convergence speed of pure logistic growth decays as more flows enter the system—flow 5 is quite slow at obtaining its bandwidth share. We have fixed this by making γ adaptive in [14], and we apply the same method in this paper, in Section IX-A.

A. Deflated LGC

In [15], we obtained the utility function of LGC, and showed how LGC can be adopted by the NUM theory. This also helps change the equilibrium marking probability of LGC. In other words, we can deflate its marking probability, which can become very high when the number of competing flows is large. The send rate dynamic is

$$\dot{x}_r = x_r \gamma_r \left(-\log_{\phi_1}(x_r) + \log_{\phi_2}(1 - \hat{p}_r) \right). \qquad (8)$$

where ϕ_1 and ϕ_2 are the base parameters that can be configured to control the marking probability. If $\phi_1 = \phi_2$, then (8) reduces to (5). In case there is one bottleneck, the equilibrium marking probability can be obtained, which is given by

$$p_r^* = 1 - e^{\frac{\log(\phi_2)\log(\frac{1}{S})}{\log(\phi_1)}} \qquad (9)$$

where S is the number of competing flows. In this case, if $\phi_1 > \phi_2$, the equilibrium marking probability becomes smaller. For example, if we set $\phi_1 = 10$ and $\phi_2 = 2$, then $p_r^* = 0.1897$ (this was 0.8 in Fig. 4(b)). however, the congestion controller behavior does not change—it exhibits the exact behavior shown in Fig. 4(a).
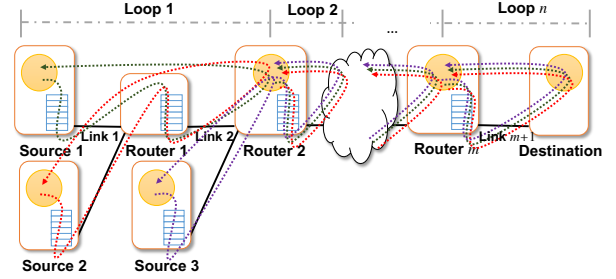
VI. LGC Chain (LGCC) Model

A. Network Model

In this section, we introduce how we construct a chain of LGCs and discuss its deployment considerations. Our main

(a) A sample network topology

(b) Corresponding loops

Fig. 5. An example of how loops are formed.

aim is to have a deployable solution with minimal required changes to current commodity hardware. ECN is a widely-deployed mechanism in commodity hardware with broad utility. We utilize this feature and assume that routers in the network support packet marking via RED. However, some of the nodes on the path towards the destination need to support LGCC, which we call *LGCC router*; these nodes could be domain border routers. In other words, LGCC does not need to be implemented as a hop-by-hop solution although that is also possible.

Fig. 5(a) illustrates an example network topology of a source and a destination with m routers on the path. This can be mapped into the diagram in Fig. 5(b) in which only some of the routers deploy LGCC, meaning that an end-to-end congestion controller between the source and the destination is broken in several smaller congestion control loops. The other routers are assumed to have RED deployed, which ECN-marks packets if the output queue is non-empty. Therefore, there is one LGC loop from the source to Router 2, another one from Router 2 to some other router, and finally, there is one from Router m to the destination. There is one LGC loop per source from each sources 1 and 2 to Router 2, one loop from source 3 to Router 2, another one per source from Router 2 to some other router (continued in that way until Router m), and finally, there is one per source from Router m to the destination. This example already highlights the possibility of aggregating traffic to reduce unnecessary competition between flows, but we deem traffic aggregation out of scope for the present paper.

The number of LGCC routers and where they operate depends on the network design. For example, each loop can be formed over one domain, or if the medium is wireless, the loop can operate over only that medium specifically. We assume that it is a design issue and up to the network manager where to deploy LGCC routers; however, we will provide recommendations about an optimal deployment in Section **??**.

Fig. 6 shows a high level internal structure of an LGCC

Fig. 6. Loop management architecture in LGCC routers

router. It has several "Receivers" to control the congestion control loops ending at it, and several "Senders" to start new ones towards their destination/next LGCC router. The role of "Schedulers" is to prioritize taking packets from receivers and pass them to the right next sender. If packets are accumulated at receivers waiting for schedulers to pass them, a virtual queue management scheme can be deployed to ECN-mark packets. These marks are sent back towards the previous senders via acks from the receivers.

B. Congestion Controller Model

Here, we introduce LGCC. In particular, we form a chain of (5) where the carrying capacity is not normalized to 1. Instead, it is determined by the send rate of next loop. In addition, we use notation $\hat{p}_{r,i}$ to reflect the marking probability of only loop i that is measured by its sender. Assume a flow broken into n loops from a source to a destination. The rate dynamics at the sender of each loop are given by

$$\dot{x}_{r,1} = x_{r,1}(t)\gamma_{r,1}\left(1 - \frac{x_{r,1}(t)}{k_{r,1}(t)} - \hat{p}_{r,1}(t)\right) , \qquad (10a)$$

$$\dot{x}_{r,2} = x_{r,2}(t)\gamma_{r,2}\left(1 - \frac{x_{r,2}(t)}{k_{r,2}(t)} - \hat{p}_{r,2}(t)\right) , \qquad (10b)$$

$$\vdots$$

$$\dot{x}_{r,n} = x_{r,n}(t)\gamma_{r,n}\left(1 - \frac{x_{r,n}(t)}{k_{r,n}(t)} - \hat{p}_{r,n}(t)\right) , \qquad (10c)$$

where

$$k_{r,1}(t) = \min\left(x_{r,2}(t), c_l(t)\right), \quad \forall l \in L_{r,1} , \qquad (11a)$$

$$k_{r,2}(t) = \min\left(x_{r,3}(t), c_l(t)\right), \quad \forall l \in L_{r,2} , \qquad (11b)$$

$$\vdots$$

$$k_{r,n}(t) = \min\left(x_{r,n+1}(t), c_l(t)\right), \quad \forall l \in L_{r,n} , \qquad (11c)$$

and $x_{r,n+1}(t)$ and c_l denote the destination's processing rate and link capacity, respectively. LGC uses the notion of carrying capacity, and so does LGCC. Therefore, K should be determined in each loop. This can be the minimum of the send rate of the next loop and the link capacity in the current loop, which is obtained by (11). Since we assume that the formation of each loop is under one authority/domain, link capacities

may be known in that domain. In case a link capacity is changing over time, e.g. in a wireless medium, information about the current capacity and modulation is usually available at the output link of the LGCC router. This could be used as the *current* link/carrying capacity.

If packets of two or more receivers are forwarded to the same sender in the next loop, i.e. the loops are aggregated into one flow in the next loop, we can enable schedulers to prioritize packets of certain receivers while passing to their next sender. Here, we assign weights to receivers and employ a weighted fair queuing discipline in the scheduler. In this case, the rate dynamic of loop i is given by

$$\dot{x}_{r,i} = x_{r,i}(t)\gamma_{r,i}\left(1 - \frac{x_{r,i}(t)}{k_{r,i}(t)} - \frac{\hat{p}_{r,i}(t)}{\omega_{r,i}}\right) \qquad (12)$$

where $\omega_{r,i}$ denotes the weight assigned to the send rate of flow r in loop i.

C. Benefits of the LGCC Design

The notion of carrying capacity is an important feature of the LGCC design. It helps the controller to have an accurate estimate of the upper capacity limit, which is a direct signal sent back to previous controllers, relieving them from having to wait for indirect congestion/cost signals such as queue marks to adjust their rates. In other words, each loop uses the send rate of the next loop as its own K, which leads to faster rate adjustment to available capacity at next loops.

LGC's increase/decrease behavior is non-linear. As opposed to controllers with, e.g. additive increase, how fast to increase/decrease depends on the difference between $x_{r,i}$ and $(1-\hat{p}_{r,i})k_{r,i}$. This means that in case of a sudden increase in the medium capacity (assume a wireless link that is recovered from a rate drop due to an intermittent noise), capturing the available capacity can be fast.

D. Utility Function

For LGCC, we apply the same method that we used to obtain the LGC utility function in [15]. In equilibrium, we know that $q_{r,i}^* = U'(x_{r,i})$ where $q_{r,i}^* = -\log_{\phi_2}(1 - p_{r,i}^*)$. Solving for p yields $p_{r,i}^* = 1 - \phi_2^{q_{r,i}^*}$. We also have $1 - \frac{x_{r,2}^*(t)}{k_{r,2}(t)} - p_{r,2}^*(t) = 0$. These give

$$U(x_{r,i}) = x_{r,i}\left(\frac{1}{\log \phi_1} - \log_{\phi_1}\left(\frac{x_{r,i}}{k_{r,i}}\right)\right) , \qquad (13)$$

which is a strictly concave, twice differentiable function in the range $(0, k_{r,i}]$. These properties facilitate its stability and convergence.

VII. CONGESTION CONTROL ALGORITHMS

The LGCC dynamics (10) can be converted into a Network Utility Maximization (NUM) problem. This enables us to analyze LGCC better and find deployable solutions, especially on commodity routers. In [15], we obtained the utility function of LGC, and showed how LGC can be adopted by the NUM theory. Here, we use the utility function behind LGCC

dynamics, and then, we divide the optimization logic between sources and routers. Since sources and LGCC routers use the same module to send packets, by *sender dynamics* we mean the controller logic in both of them, and by *router dynamics* we mean routers that do not break connections. Our assumption about these routers is that they can either implement our specific packet marking scheme or at least, they support RED, which is a widely-deployed mechanism.

Fig. 7. 5 loops chained towards the destination

A. Primal Algorithm

In this algorithm, the optimization logic is deployed at the sender side. The network (routers) only return a cost to senders, which is a direct mapping of the current aggregate traffic of a link to a non-negative real value. It is given by

$$\dot{x}_{r,i} = x_{r,i}\gamma_{r,i}\left(U'(x_{r,i}) + \log_{\phi_2}(1 - \hat{p}_{r,i})\right) \quad (14)$$

where $U(.)$ is the utility function of senders [15]. Substituting the derivative of (13) in (14) yields

$$\dot{x}_{r,i} = x_{r,i}\gamma_{r,i}\left(-\log_{\phi_1}\left(\frac{x_{r,i}}{k_{r,i}}\right) + \log_{\phi_2}\left(1 - \hat{p}_{r,i}\right)\right). \quad (15)$$

Since our assumption is that routers support RED, the network cost can be sent back to senders via ECN marks. This is denoted in (14) by $\hat{p}_{r,i}$. In the next subsection, we introduce an ECN marking algorithm for this purpose.

B. Dual Algorithm

In the dual algorithm, the optimization logic is deployed in routers, and senders directly map the cost they get from the network into a send rate. First, we discuss a general router which could run this algorithm, and then discuss how routers with RED AQM can employ it.

We use the fact that at equilibrium, $q_{r,i}^* = U'(x_{r,i}^*)$ where $q_{r,i}^*$ is the cost sender i gets. As we already showed in [15], in case of employing RED, $q_{r,i} = -\log_{\phi_2}(1 - \hat{p}_{r,i})$. Combined, these two yield

$$x_{r,i} = k_{r,i} \ e^{\left(\log(1-\hat{p}_{r,i})\log(\phi_1)/\log(\phi_2)\right)} \quad (16)$$

meaning that the sender can directly calculate its optimal send rate using $\hat{p}_{r,i}$. However, routers need to iterate on their cost until they converge to the optimal value. The first dual algorithm, called **Dual1** is given by

$$\dot{p}_l = \beta_l\left[y_l - C\right]_{p_l}^+ \quad (17)$$

where p_l and y_l denote the marking probability of link l (at the router's output queue to the link) and the aggregate traffic sent over link l in the last interval. β_l denotes the step size. At output queues, $C = c_l$ where c_l is the link capacity, and at receivers' virtual input queues, $C = x_{r,i+1}$, i.e. the send rate of the next loop.

As we showed in our previous work, we can use RED with a special configuration (see Appendix E of [15]) as the dual algorithm. This enables us to have a more deployable

solution. The cost/marking probability dynamics at routers' output queues (links) is given by

$$p_l[n] = \left[\frac{b_l[n]}{\max_{\text{th}}}\right]_0^1 \quad (18)$$

where b_l, $\max_{\text{th}}$, and n denote the queue length, the maximum threshold parameter of RED over which all the packets are marked, and the interval number. Other RED parameters are $\min_{\text{th}} = 1$, $w_q = 1$, and $\max_p = 1$. This algorithm implies that the current backlogged size can be divided by the maximum threshold to have a marking probability for the current interval.

One of the benefits of the dual algorithms of LGCC is that, due to the notion of carrying capacity, senders can sometimes *jump* to a send rate if, for example, the carrying capacity suddenly changes. In case the sender's loop is not shared with other loops, jumping to higher rates does not affect the network. Otherwise, if some links of a loop might be shared with other loops, the primal algorithm can be used instead of (16); the use of a primal-dual algorithm can also handle this case, which will be explained in the next subsection.

C. Primal-Dual Algorithm

This type of algorithm divides the optimization logic between senders and routers. The combinations (15)–(17) and (15)–(18) form two sets of primal-dual algorithms.

D. Numerical Evaluation

In order to have a notion on how much (10) is effective and can accelerate convergence of a chain of loops due to the usage of the carrying capacity, we perform a numerical evaluation here on the topology shown in Fig. 7. We consider two scenarios: 1) a chain of 5 loops where inside each loop, there is no other competing loop (i.e. Source 2 to Source 5 do not send any packets), and 2) a chain of 5 loops where inside each loop, there is another competing loop, which will be aggregated into one loop in the next loop (i.e. all sources send packets, which will be aggregated at the next router). Therefore, in the last loop, the two flows compete for the bottleneck capacity, and each one should get half of the capacity; in the previous loop, the two flows that are aggregated in the next loop using one of the flows should also get half of the send rate of that flow.

If we assume a normalized capacity of 1 in the last loop, then last-loop flows get 0.5, each flow in the previous loop gets 0.25, and the values are 0.125, 0.0625, and 0.03125, respectively for the flows in the previous loops. Apparently, in both scenarios, the rate of loop i, $x_{r,i}$, is the carrying capacity

of loop $i-1$, but $p_{r,i}^*$ is equal to 0 and 0.5 in scenarios 1 and 2, respectively. In the figures, we only report the rate of one flow in each loop labeled with "Loop i". At step 200, we cut the capacity of the last loop to 0.5 to observe how fast the chain of congestion controllers can converge.

We evaluated the Primal-Dual and Dual algorithms under the above scenarios. In addition, to see how a chain of controllers that do not employ the notion of carrying capacity behaves, we used a controller with a logarithmic utility function, i.e. $U(x) = \log(x)$. Hence, the rate dynamics in loop i are given by

$$\dot{x}_{r,i} = \gamma_{r,i} \left(\frac{1}{\hat{p}_{r,i}} \right) . \tag{19}$$

The Dual algorithm (17) was used for the algorithms, and the controllers were also tuned for fastest convergence. Fig. 8 and Fig. 9 illustrate the results of scenario 1 and scenario 2, respectively, and in Fig. 10 shows the results of the logarithmic utility function. From the figures we observe that LGCC can converge faster when the capacity drops to 0.5 in loop 5 at step 200, and also at step 400, when it changes back to 1. It might seem that a Dual method is superior to the other one. As mentioned before, if links in a loop is not shared by other flows, a Dual method can be used to *jump* to the rate routers imply; however, it is not true in a general topology.

VIII. ANALYTICAL EVALUATION

In this section, we analytically evaluate properties of LGCC such as equilibrium, stability, and fairness.

A. Equilibrium

Depending on where the bottleneck in the network is, LGCC can converge to different values. For this analysis, we assume that the bottleneck is always in some loop n, and then, the min operator in (11) always kicks in, and for the sake of simplicity, we assume that schedulers in Fig. 6 employs a fair queuing discipline. Hence, assuming $x_{r,i+1}(t) \leq c_l(t)$, $\forall l \in L_{r,i}, \forall t$:

$$x_{r,i}^*(t) = x_{r,i+1}^*(1 - p_{r,i}^*) \tag{20}$$

and when the marking probability is deflated

$$x_{r,i}^* = x_{r,i+1}^* \, e^{\left(\log(1 - p_{r,i}^*) \log(\phi_1)/\log(\phi_2) \right)} \tag{21}$$

Employing a fair queuing discipline in the scheduler implies that $x_{r,i}^* = x_{s,i}^*$ for competing flows r and s in loop i. Using the fact that $\sum_r x^* r, i = x_{r,i+1}^*$, we get

$$e^{\left(\log(1 - p_{r,i}^*) \log(\phi_1)/\log(\phi_2) \right)} = \frac{1}{S} \tag{22}$$

where S is the number of competing flows. Solving for $p_{r,i}^*$ yields

$$p_{r,i}^* = 1 - e^{\left(\log(1/S) \log(\phi_2)/\log(\phi_1) \right)} . \tag{23}$$

B. Stability Analysis

Here, we discuss the stability of LGCC. The following theorem will prove its global stability.

Theorem VIII.1. *The primal-dual algorithm* (15)–(17) *is globally, asymptotically stable.*

Proof. See Appendix A-A. □

Although Theorem VIII.1 proves that the system of equations (15)–(17) is globally stable, we also analyze its local stability; this provides us a relationship between the step size parameters γ and β, and we can use this as a guideline for setting those parameters.

Local stability refers to the analysis of the system behavior close to an equilibrium. We use the dual algorithm (17) as the cost generator at routers. This indeed includes a step size parameter, β, which can affect the system stability.

We use an aggregation tree topology that is illustrated in Fig. 11 with variable depth (the number of loops of the chains – denoted by δ) and breadth (the number of competing flows in each loop – denoted by ζ). The circles represent LGCC routers at which incoming flows compete. For the sake of simplicity, we assume that the flows in each loop only compete at the input queue of the next loop, and they are all aggregated into one flow in the next loop.

Theorem VIII.2. *The system of equations* (15)–(17) *is locally stable if and only if*

$$\beta_d \leq \frac{1}{4K} \gamma_d \, \zeta^{(\delta - d - 2)} \quad for \; 1 \leq d \leq \delta . \tag{24}$$

Proof. See Appendix A-B. □

The above theorem implies that the stability in each loop i only depends on the number of competing flows in that loop; this means that the value of γ_i and β_i determine the stability, and as we get closer to the bottleneck link, i.e. $i = 1$ in this case, γ should be smaller than β.

C. Fairness Analysis

Fairness properties of a congestion controller determine how fairly the capacity is shared among a number of competing flows and how much each flow can expect to receive. Since LGCC chains controllers, we focus on two types of fairness: intra-loop and inter-loop.

1) Intra-Loop Fairness: Here, we focus on only one loop. We consider a number of flows competing for links capacity in an arbitrary topology. First, we define a new type of fairness.

Definition 1 (Logarithmic Proportional Fairness)**.** *Every rate change from the equilibrium proportional to the logarithm of the equilibrium is non-positive. Formally speaking,*

$$\sum_r -\log(x_r^*)(x_r - x_r^*) \leq 0 . \tag{25}$$

This is obtained from the fact that if the feasible set is convex and $U(.)$ has a globally maximized at $\mathbf{x}^*$, then for any feasible allocation vector $\mathbf{x}$, we have $\nabla(U)|_{\mathbf{x}^*}(\mathbf{x} - \mathbf{x}^*) \leq 0$ meaning that the gradient of U, evaluated at $\mathbf{x}^*$, times the change is not positive. The above definition also implies that
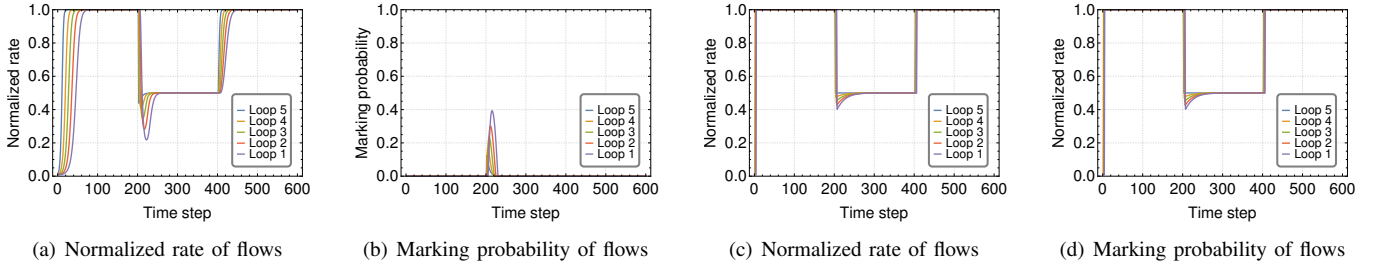
(a) Normalized rate of flows (b) Marking probability of flows (c) Normalized rate of flows (d) Marking probability of flows

Fig. 8. Chaining 5 loops of LGCC. (a) and (b): Primal-Dual, and (c) and (d): Dual.

(a) Normalized rate of flows (b) Marking probability of flows (c) Normalized rate of flows (d) Marking probability of flows
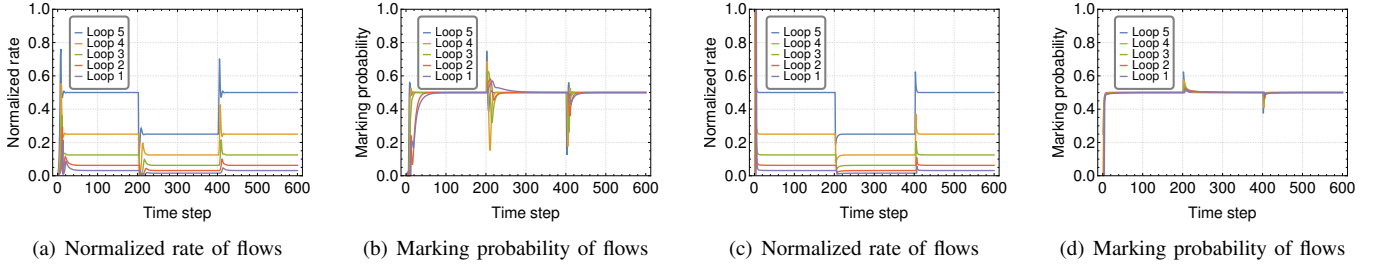
Fig. 9. Chaining 5 loops of LGCC where there are two competing flows in each loop, aggregated into one the next loop. (a) and (b): Primal-Dual, and (c) and (d): Dual.

(a) Normalized rate of flows (b) Marking probability of flows (c) Normalized rate of flows (d) Marking probability of flows
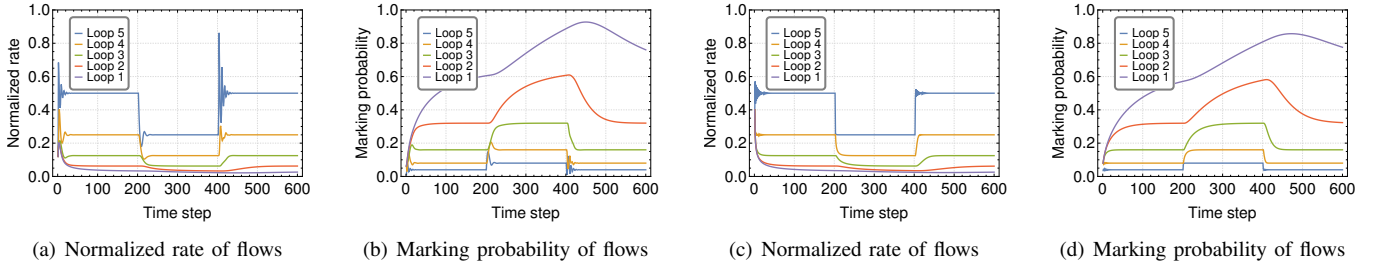
Fig. 10. Chaining 5 loops of the logarithmic controller where there are two competing flows in each loop, aggregated into one the next loop. (a) and (b): Primal-Dual, and (c) and (d): Dual.


Fig. 11. Aggregation tree

if there are n sequential links with the normalized capacity of 1, 1 n-hop flow with rate x_1 crossing all the links, and n 1-hop flows crossing each link with rate x_2, then x_1^* is the multiplication of the equilibrium rate of other 1-hop flows in its path. In other words,

$$x_1^* = \prod_{i=1}^{n} x_2^* .$$

2) Inter-Loop Fairness: To focus on only inter-loop fairness, here, we assume that all routers support LGCC. This implies that loops are formed in a hop-by-hop fashion. In this case, schedulers determine the fairness type. We consider a number of flows in a loop arriving at an LGCC router and analyze two types of chaining: 1) the flows are aggregated into one flow in the next loop, and 2) a separate flow is created in the next loop for each flow. We employ a weighted fair queuing discipline in LGCC routers. Through the following theorems, we discuss how the type of chaining and weight assignment determines fairness.

Theorem VIII.3. *In case of aggregated chaining, if weights are set equal to the number of aggregated flows, LGCC attains max-min fairness.*

Proof. See Appendix B-A. □

Theorem VIII.4. *Let $p_{r,i}^u$ denote the upstream marking probability of flow r, i.e. the marking probability at the output queues of LGCC routers 1 to i. In case of per-flow chaining, if weights are set inversely proportional to $-\log(1 - p_{r,i}^u)$ at the scheduler in LGCC router $i + 1$ for all routers i on the*

Algorithm 1 LGCC congestion control for flow r in loop i

1: **procedure** INIT
2:     $x \leftarrow x_{\text{init}}$
3:     $\hat{p} \leftarrow 0$
4:     $\gamma \leftarrow \gamma_{\text{init}}$
5: **end procedure**
6: **procedure** ONINTERVALENDS
7:     $m \leftarrow$ ECN-marked percentage of ACKs
8:     $\hat{p} \leftarrow (1 - \alpha)\,\hat{p} + \alpha\,m$
9:     $k \leftarrow \min\left(x_{r,i+1}, c_l\right) \quad \forall l \in L_{r,1}$
10:     **if** Primal-Dual **then**
11:         $\hat{x} \leftarrow k\,e^{\left(\log(1-\hat{p})\,\log(\phi_1)/\log(\phi_2)\right)}$
12:         **if** $|x - \hat{x}| \leq g_1$ **then**
13:             $\gamma \leftarrow \gamma_{\text{conv}}$
14:         **else if** $|x - \hat{x}| \geq g_2$ **then**
15:             $\gamma \leftarrow \gamma_{\text{init}}$
16:         **else**
17:             $\gamma \leftarrow |x - \hat{x}|\frac{\gamma_{\text{init}} - \gamma_{\text{conv}}}{g_2 - g_1} + \gamma_{\text{conv}}$
18:         **end if**
19:         $x \leftarrow x\gamma\left(\log_{\phi_2}\left(1 - \hat{p}\right) - \log_{\phi_1}\left(\frac{x}{k}\right)\right) + x$
20:         $x \leftarrow \max(x, x_{\text{init}})$
21:     **else if** Dual **then**
22:         $x \leftarrow k\,e^{\left(\log(1-\hat{p})\,\log(\phi_1)/\log(\phi_2)\right)}$
23:     **end if**
24: **end procedure**

---

path, then LGCC attains proportional fairness.

Proof. See Appendix B-B. $\qquad\square$

Since in LGCC we assume that each loop operates with its own marking probability, accessing $p_{r,i}^u$ at downstream routers would need a special bit in the LGCC header to convey the marks received so far at the output queues of LGCC routers.

IX. LGCC ALGORITHM AND IMPLEMENTATION

Algorithm 1 shows how each sender adjusts its rate. It can be either a Primal-Dual or a Dual algorithm. At startup, send rate x is set to the initial value x_{init}; this can be interpreted as the initial window size of TCP. However, the unit of x is bit per second. We estimate the loop marking probability, $\hat{p}$, using an exponential smoothing average with parameter α. As a rate-based method, we define intervals to sample ECN-marks, and at the end of each interval, $\hat{p}$ is updates. The carrying capacity is also obtained: it can be the minimum link capacity in the loop, or the last send rate of the next loop, which is echoed back to the previous loop via the receiver window size field; the value of this field is interpreted as a rate in bit per second. Then, depending the algorithm, and in case of Primal-Dual, (15) is run. Before the rate calculation, from line 11 to 18, the adaptive step size algorithm is run. The goal is to be more aggressive if the calculated rate from a Dual algorithm is *far* from the current send rate; this will be discussed in the next subsection. Then, the send rate is set, ensuring it is not lower than the initial send rate. In case of Dual, the send rate is directly calculated from the measured marking probability,

$\hat{p}$. Details of the parameter setting are presented in the next section.

In both of the algorithms, we assume that routers run (18), which can be realized through a special configuration of RED. At the end of the loop, the receiver module of the next LGCC router, echoes back ECN marks.

The following subsections elaborate on some aspects of this algorithm that go beyond the previously discussed basic LGCC dynamics: adapting γ, obtaining the right capacity value to use (in line 9, we assume that the minimum of all link capacities is known), and how often to trigger "OnIntervalEnds".

A. Adaptive Step Size

To improve the performance of LGCC via faster convergence, we use the same technique that we used in LGC [14] to adjust the step size of the Primal-Dual algorithm, i.e. γ. In LGCC, $\hat{p}$ and accordingly, (16) provides an indication of what the target rate can be, and how *far* the current send rate is from that target. In the Primal-Dual algorithm, when the current rate is far from the target, we use a larger value of γ. When the rates are close, then we use a small value for γ to have a smoother behavior. The algorithm (line 11 to 18 in Algorithm 1) determine if the difference of x and $\hat{x}$ are larger than the threshold g_2. If so, γ is set to a large value, i.e. γ_{init}; if the difference is smaller than the threshold g_1, then γ_{conv} is used; otherwise, it is set linearly based on the proportion of the difference.

B. Carrying Capacity

As shown in (11), LGCC senders need to know the carrying capacity of their loop to obtain their rate; this could be either the minimum link capacity on their path in the loop, or the current send rate of the next sender in the next loop of the chain. In [14], we showed that if an LGC sender cannot get the carrying capacity, it still can operate (as we see in Fig. 2(b)). However, the queue length fluctuates more. When each loop is formed under one domain or over a specific medium, e.g. a wireless link, obtaining the carrying capacity is sometimes straightforward. In wireless links with time-varying capacity, the current send rate (depending on the noise level and modulation) is the carrying capacity. This is how the current transmission rate is obtained in some work such as [21]. ABC also uses a method based on the output queue drain rate to estimate the current link capacity. Hence, we assume that if there is a link with time-carrying capacity, an LGCC router is run on that node to start a new loop over that link; in case of wired links, LGCC routers can be run on interior/exterior routers of each authority domain, e.g. ISPs.

C. Rate Update Time Intervals

The rate is regularly updated by each sender per interval. The LGCC dynamics, (10), imply that the fair rate only depends on $\hat{p}$ in each loop, and as long as all sources *see* the same $\hat{p}$, they converge to the same rate. This means that the steady state behavior of the controller is not sensitive to the update interval: updating it more or less often will naturally

make the control converge faster or slower, but the point of convergence is not affected (but updating it extremely fast could lead to oscillations). There are several options for the update frequency such as using the smallest RTT value seen over a certain time interval (an estimate of the RTT without queuing delay). If a flow's smallest RTT is still a long time interval, especially for a large-RTT flow, we can use shorter intervals. Based on our observation in simulations (performed in [14], the performance is robust if a flow can receive at least 4 ACKs per interval. Hence, another option is to use small, equal-sized intervals that can fit receiving a few ACKs.

X. SIMULATION AND EXPERIMENT RESULTS

A. Simulation Setup

In this section, we present simulation results of LGCC over various topologies and scenarios, compared with ABC as the state-of-the-art. Simulations were performed in the INET framework[4] of OMNeT++.

B. Comparison with ABC

As the first topology, we used a dumbbell topology with 5 sources, sending packet to 5 different destinations. Fig. 12 shows the result of the first scenario in which, the flows start together, and from second 2 to 3, the capacity drops to 150 Mbps. The links have a capacity of 500 Mbps, and the link propagation delay is 1 ms. The queue length of the bottleneck router is also shown. Both ABC and LGCC have a special router at the bottleneck link. From the figures we see that LGCC operates faster and smoother in terms of convergence, which also results in a smaller queue at the bottleneck.

In the next scenario over the previous topology, flows join one-by-one every two seconds, and they last for 12 seconds. Fig. 13 illustrates how the two protocols operate; again, LGCC converges faster and behaves more smoothly than ABC with a smaller queue length.

We ran the above two scenarios 10 times, each with a different random seek and a random start of flows in shorter that one link RTT. Fig. 14 shows how LGCC and ABC operate. In terms of short-term fairness, calculated over the average values of time windows of 24 ms. This also confirms that LGCC operates more smoothly, and flows converge fast. The queue length is accordingly smaller because of faster reactions to any changes in the number of flows.

In the next scenario, we ran ABC and LGCC over three real LTE link capacity traces of the AT&T, TMobile, and Verizon operators. The topology includes a destination, connected through an LTE link whose capacity changes over time (the capacity follows the pre-measured trace). Fig. 15 illustrates the difference between the rate of the flow with the actual measured rate in the trace at any time. In case of 2 flows, each flow should get half of the current capacity, and the difference is calculated with this. Referring the the figure, we observe that LGCC flows can track capacity changes faster, leading to a smaller queue at the LTE link. As the number of flows increases, the difference between ABC and LGCC become

[4]The code is available at https://cutt.ly/2yV7eQK

(a) ABC send rate   (b) ABC bottleneck link

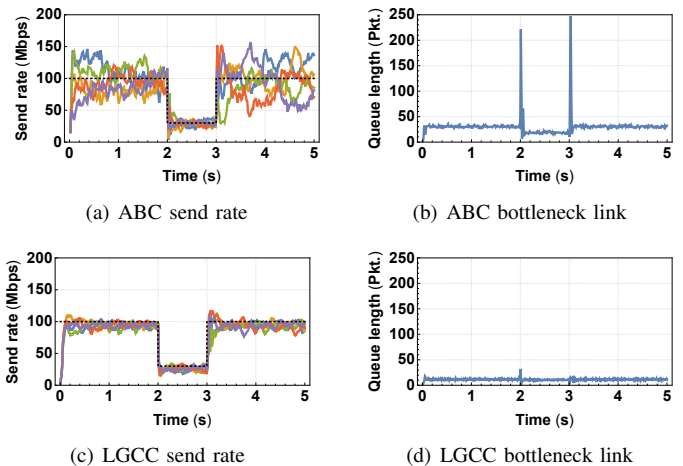(c) LGCC send rate   (d) LGCC bottleneck link

Fig. 12. Performance comparison of ABC and LGCC in a dumbbell topology: 5 flows start at the same time, and the capacity of the bottleneck link drops to 150 Mbps in $[2, 3]$s.

(a) ABC send rate   (b) ABC bottleneck link
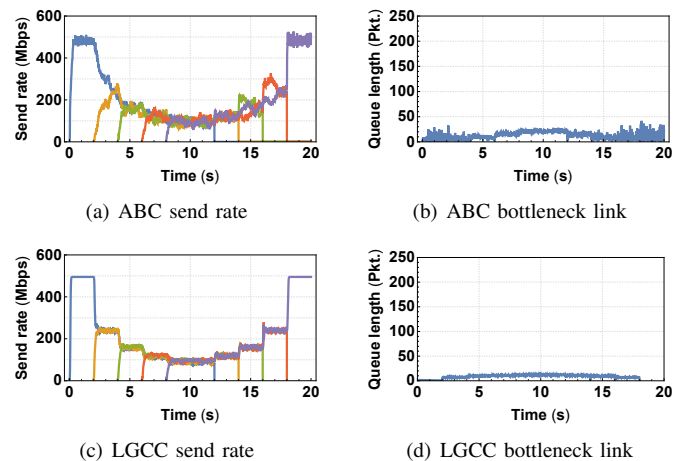
(c) LGCC send rate   (d) LGCC bottleneck link

Fig. 13. Performance comparison of ABC and LGCC in a dumbbell topology: 5 flows join one-by-one at every two seconds, and leave after 12 seconds of transmission.
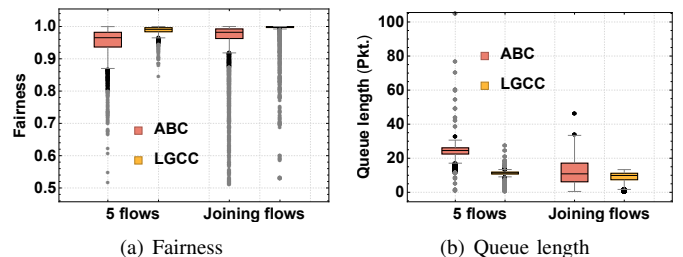

(a) Fairness   (b) Queue length

Fig. 14. Box-Whisker plot of short-term Jain's fairness index and queue length of ABC and LGCC in the two different scenarios.

smaller because each flow should receive a smaller share of the capacity. The link utilization of the protocols are 92% for ABC, and 94% in case of LGCC.

C. Fairness Evaluation

Here, we evaluate fairness properties of LGCC using simulation. Fig. 16

(a) Queing delay (b) Rate difference

Fig. 15. Box-Whisker plot of rate difference and queuing delay of ABC and LGCC.

(a) Send rate at senders (b) Marking prob. at bottleneck

Fig. 16. Marking probability of 5 flows when the carrying capacity changes.

(a) Send rate at the first and the second loop (b) Marking prob. at first LGCC router

Fig. 17. Send rate and marking probability of 2 aggregated flows when the carrying capacity changes.

D. Flow Aggregation

One of the potentials of LGCC is flow aggregation. Although a detailed evaluation of this is out of the scope of this paper, however, we demonstrate, using a simple topology, how it looks in practice. Fig. 17 shows the results of two flows aggregated at an LGCC router into one flow on a dumbbell topology. The link capacity is 500 Mbps, but the capacity at the second link drops to 150 Mbps in $[2, 3]$s. Since there is only one flow operating on the second link, there is no competition there. The LGCC router only applies a fair queuing discipline on the flows to be aggregated. So, the flows receive half of the capacity, and in case of a drop, they flow the capacity change fast. The scheduling does not change the equilibrium marking probability *seen* by sources; the spikes only happen when on-the-fly packets arrive at the bottleneck link whose capacity is just dropped.

XI. CONCLUSION

In this paper, we presented a novel multi-hop congestion controller using the well-known Logistic Growth (LG) function in a food chain style. The idea of carrying capacity in LG parallels with the bottleneck link capacity, which in case of link layer capacity measurements, can provide an excellent

feedback to sources to adapt their send rate, especially for links with time-varying capacity. This combined with employing ECN enables us to deploy LGCC over path segments, and react faster to available capacity changes.

We analytically evaluated LGCC for stability and fairness; we proved that the system is globally stable, and provided conditions for the step size parameter setting. Through simulation study, we evaluated the performance of LGCC in terms of latency (queuing), utilization, and short-term fairness. We also used real-life traces of some LTE link capacities, and compared LGCC with ABC. Results confirm that LGCC can improve performance, and in addition, its ability to provide flow aggregation can provide benefits for congestion control.

As future work, we will perform a deeper analysis of flow aggregation, devise a flow aggregation protocol, and provide a scalable solution. Performing experiments on real use cases is of our interest.

REFERENCES

[1] P. Teymoori, M. Welzl, S. Gjessing, E. Grasa, R. Riggio, K. Rausch, and D. Siracusa, "Congestion control in the recursive internetworking architecture (rina)," in *2016 IEEE International Conference on Communications (ICC)*. IEEE, 2016, pp. 1–7.

[2] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-wave cellular wireless networks: Potentials and challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, 2014.

[3] M. Zhang, M. Mezzavilla, R. Ford, S. Rangan, S. Panwar, E. Mellios, D. Kong, A. Nix, and M. Zorzi, "Transport layer performance in 5g mmwave cellular," in *Computer Communications Workshops (INFOCOM WKSHPS), 2016 IEEE Conference on*. IEEE, 2016, pp. 730–735.

[4] D. A. Hayes, D. Ros, and Ö. Alay, "On the importance of tcp splitting proxies for future 5g mmwave communications," in *2019 IEEE 44th LCN Symposium on Emerging Topics in Networking (LCN Symposium)*. IEEE, 2019, pp. 108–116.

[5] M. Kühlewind, Z. Sarker, T. Fossati, and L. Pardue, "Use Cases and Requirements for QUIC as a Substrate," Internet Engineering Task Force, Internet-Draft draft-kuehlewind-masque-quic-substrate-00, Mar 2020, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/html/draft-kuehlewind-masque-quic-substrate-00

[6] M. Polese, M. Mezzavilla, M. Zhang, J. Zhu, S. Rangan, S. Panwar, and M. Zorzi, "milliproxy: A tcp proxy architecture for 5g mmwave cellular systems," in *Signals, Systems, and Computers, 2017 51st Asilomar Conference on*. IEEE, 2017, pp. 951–957.

[7] R. Ford, M. Zhang, M. Mezzavilla, S. Dutta, S. Rangan, and M. Zorzi, "Achieving ultra-low latency in 5g millimeter wave cellular networks," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 196–203, 2017.

[8] D. A. Hayes, P. Teymoori, and M. Welzl, "Feedback in recursive congestion control," in *European Workshop on Performance Engineering*. Springer, 2016, pp. 109–125.

[9] D. Katabi, M. Handley, and C. Rohrs, "Congestion control for high bandwidth-delay product networks," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 89–102, 2002.

[10] N. Dukkipati, "Rate control protocol (rcp): Congestion control to make flows complete quickly," Ph.D. dissertation, Stanford, CA, USA, 2008.

[11] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan, "Data Center TCP (DCTCP)," in *Proc. ACM SIGCOMM*, New Delhi, India, 2010, pp. 63–74.

[12] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, 1998. [Online]. Available: http://dx.doi.org/10.1057/palgrave.jors.2600523

[13] M. Alizadeh, A. Kabbani, T. Edsall, B. Prabhakar, A. Vahdat, and M. Yasuda, "Less is more: Trading a little bandwidth for ultra-low latency in the data center," in *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. San Jose, CA: USENIX, 2012, pp. 253–266. [Online]. Available: https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/alizadeh

[14] P. Teymoori, D. Hayes, M. Welzl, and S. Gjessing, "Even lower latency, even better fairness: Logistic growth congestion control in datacenters," in *2016 IEEE 41st Conference on Local Computer Networks (LCN)*, Nov 2016, pp. 10–18.

[15] P. Teymoori, D. A. Hayes, M. Welzl, and S. Gjessing, "Estimating an additive path cost with explicit congestion notification," University of Oslo, Tech. Rep. ISBN 978-82-7368-452-3, 2019.

[16] X. Wang, L. Kong, F. Kong, F. Qiu, M. Xia, S. Arnon, and G. Chen, "Millimeter wave communication: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1616–1653, 2018.

[17] P. Jimenez Mateo, "Analysis of tcp performance in 5g mm-wave mobile networks," Ph.D. dissertation, Universidad Carlos III de Madrid, Spain, 2017.

[18] M. Zhang, M. Mezzavilla, J. Zhu, S. Rangan, and S. Panwar, "Tcp dynamics over mmwave links," in *Signal Processing Advances in Wireless Communications (SPAWC), 2017 IEEE 18th International Workshop on*. IEEE, 2017, pp. 1–6.

[19] M. Pieska and A. Kassler, "Tcp performance over 5g mmwave linksâĂŤtradeoff between capacity and latency," in *Wireless and Mobile Computing, Networking and Communications (WiMob),*. IEEE, 2017, pp. 385–394.

[20] M. Polese, R. Jana, and M. Zorzi, "Tcp and mp-tcp in 5g mmwave networks," *IEEE Internet Computing*, vol. 21, no. 5, pp. 12–19, 2017.

[21] T. Azzino, M. Drago, M. Polese, A. Zanella, and M. Zorzi, "X-tcp: a cross layer approach for tcp uplink flows in mmwave networks," in *Ad Hoc Networking Workshop (Med-Hoc-Net), 2017 16th Annual Mediterranean*. IEEE, 2017, pp. 1–6.

[22] V. Kulkarni, S. Bohacek, and M. Safonov, "Stability issues in hop-by-hop rate based congestion control," in *Proc. Annual Allerton Conference on Communication Control and Computing*, vol. 36. University of Illinois, 1998, pp. 79–88.

[23] P. P. Mishra and H. Kanakia, "A hop by hop rate-based congestion control scheme," in *ACM SIGCOMM CCR 22(4)*, 1992.

[24] S. Bohacek, "Stability of hop-by-hop congestion control," in *Proc. IEEE Decision and Control*, vol. 1, 2000, pp. 67–72.

[25] P. P. Mishra, H. Kanakia, and S. K. Tripathi, "On hop-by-hop rate-based congestion control," *IEEE/ACM TON*, vol. 4, no. 2, pp. 224–239, 1996.

[26] Y. Yi and S. Shakkottai, "Hop-by-hop congestion control over a wireless multi-hop network," *IEEE/ACM ToN*, vol. 15, no. 1, pp. 133–144, 2007.

[27] F. Paganini, J. Doyle, and S. Low, "Scalable laws for stable network congestion control," in *Proc. IEEE Decision and Control*, 2001.

[28] F. Qiu and Y. Xue, "Robust joint congestion control and scheduling for time-varying multi-hop wireless networks with feedback delay," *IEEE Transactions on Wireless Communications*, vol. 13, no. 9, pp. 5211–5222, 2014.

[29] W. Su, C. M. Lagoa, and H. Che, "Optimization-based, qos-aware distributed traffic control laws for networks with time-varying link capacities," *Automatica*, vol. 72, pp. 158–165, 2016.

[30] G. Zhang, Y. Wu, and Y. Liu, "Stability and sensitivity for congestion control in wireless mesh networks with time varying link capacities," *Ad Hoc Networks*, vol. 5, no. 6, pp. 769–785, 2007.

[31] P. Goyal, M. Alizadeh, and H. Balakrishnan, "Rethinking congestion control for cellular networks," in *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, 2017, pp. 29–35.

[32] P. Goyal, A. Agarwal, R. Netravali, M. Alizadeh, and H. Balakrishnan, "ABC: A simple explicit congestion controller for wireless networks," in *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. Santa Clara, CA: USENIX Association, Feb 2020, pp. 353–372. [Online]. Available: https://www.usenix.org/conference/nsdi20/presentation/goyal

[33] K. Ramakrishnan, S. Floyd, and D. Black, "RFC3168: The Addition of Explicit Congestion Notification (ECN) to IP," USA, 2001.

[34] J. Day, *Patterns in Network Architecture: A Return to Fundamentals*. Prentice Hall, 2007.

[35] K. Ciko and M. Welzl, "First contact:can switching to rina save the internet?" in *2019 22nd Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, Feb 2019, pp. 37–42.

[36] M. Welzl, P. Teymoori, S. Gjessing, and S. Islam, "Follow the model: How recursive networking can solve the internet's congestion control problems," in *IEEE ICNC 2020, Big Island*, Feb 2020.

[37] M. Welzl, *Scalable performance signalling and congestion avoidance*. Springer Science & Business Media, 2003.

[38] G. Hasegawa and M. Murata, "TCP symbiosis: Congestion control mechanisms of tcp based on lotka-volterra competition model," in *Proceedings from the 2006 Workshop on Interdisciplinary Systems Approach in Performance Evaluation and Design of Computer & Communications Sytems*, ser. Interperf '06. New York, NY, USA: ACM, 2006. [Online]. Available: http://doi.acm.org/10.1145/1190326.1190338

[39] B. S. Goh, "Global stability in many-species systems," *The American Naturalist*, vol. 111, no. 977, pp. 135–143, 1977.

[40] J. Vano, J. Wildenberg, M. Anderson, J. Noel, and J. Sprott, "Chaos in low-dimensional lotka–volterra models of competition," *Nonlinearity*, vol. 19, no. 10, p. 2391, 2006.

[41] H. K. Khalil, *Nonlinear Systems*. Prentice Hall, 2001.

[42] B. Radunovic and J. Le Boudec, "A unified framework for max-min and min-max fairness with applications," *IEEE/ACM Transactions on Networking*, vol. 15, no. 5, pp. 1073–1083, 2007.

APPENDIX A
STABILITY ANALYSIS

A. Proof of Theorem VIII.1

Proof. We define a Lyapunov function based on the Primal-Dual algorithm in the form of

$$\mathcal{V}(x,p) = \sum_{r \in R} \sum_{i=1}^{O_r} \int_{x_r^*}^{x_r} \frac{1}{x_{r,i}\gamma_{r,i}} \big(\varphi - x_r^*\big) d\varphi$$
$$+ \sum_{l \in Q} \int_{p_l^*}^{p_l} \frac{1}{\sigma_l(\psi)} \big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-\psi)\big) d\psi \quad (26)$$

which is inspired from the Lyapunov function in [15] where the marking probability is mapped to cost using the function $-\log_\phi(1-p_l^*)$. Fixing the cost dynamics to the Dual1 algorithm and taking the time derivative of $\mathcal{V}(x,p)$ yields

$$\frac{d\mathcal{V}}{dt} = \sum_{r \in R} \sum_{i=1}^{O_r} \frac{\dot{x}_{r,i}}{x_{r,i}\gamma_{r,i}}(x_{r,i} - x_{r,i}^*)$$
$$+ \sum_{l \in L} \big[y_l - c_l\big]_{p_l}^{+} \big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big)$$
$$+ \sum_{l \in A} \big[y_l - x_{r,i+1}\big]_{p_l}^{+} \big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big)$$
$$\leq$$
$$\sum_{r \in R} \sum_{i=1}^{O_r} \big(U^{'}(x_{r,i}) + \log_{\phi_2}(1-p_{r,i})\big)(x_{r,i} - x_{r,i}^*)$$
$$+ \sum_{l \in L} (y_l - c_l)\big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big)$$
$$+ \sum_{l \in A} (y_l - x_{r,i+1})\big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big)$$

$$= \sum_{r \in R} \big(\log_{\phi_2}(1-p_{r,i}) - \log_{\phi_2}(1-p_r^*)\big)(x_{r,i} - x_{r,i}^*) \quad (27a)$$
$$+ \sum_{l \in L} \big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big)(y_l - y_l^*) \quad (27b)$$
$$+ \sum_{l \in A} \big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big)(y_l - y_l^*) \quad (27c)$$
$$+ \sum_{l \in L} (y_l^* - c_l)\big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big) \quad (27d)$$
$$+ \sum_{l \in A} (y_l^* - x_{r,i+1})\big(\log_{\phi_2}(1-p_l^*) - \log_{\phi_2}(1-p_l)\big) \quad (27e)$$
$$+ \sum_{r \in R} \big(U^{'}(x_{r,i}) + \log_{\phi_2}(1-p_{r,i}^*)\big)(x_{r,i} - x_{r,i}^*) \quad (27f)$$
$$\leq 0 .$$

The sum of (27a), (27b), and (27c) is equal to zero. (27d) and (27e) are less than or equal to zero: if, for example, $y_l^* = c_l$, then (27d) is zero, and if $y_l^* < c_l$, then p_l^*, which makes $\log_\phi(1-p_l^*)$ zero, and since $-\log_\phi(1-p_l) \geq 0$, this term is non-positive. The same argument holds for (27e). In order for (27f) to be non-positive, if $x_r \leq x_r^*$, then we should have $U_r^{'}(x_r) \geq -\log_\phi(1-p_r^*)$, or vice versa. This means that, for example, in case of $x_r \leq x_r^*$, $x_{r,i+1}$ should be large enough to let the first term become non-negative. Since we connect each loop to the next one (i.e. the carrying capacity in each loop is only dependent on the rate of the next loop, not the

previous one), loops converge from the last one, and for each loop, $x_{r,i+1}$ will increase until it meets the above condition and stabilizes the system. In case of $x_r \geq x_r^*$, the first term should be non-positive. Since the value of $x_{r,i+1}$ does not depend on $x_{r,i}$, this means that $x_{r,i+1}$ should be small enough for the term to become non-positive. Again, as the next loop converges first, and it does not depend on the previous one, this condition is met. This means that the Lyapunov function is negative, and it is zero at the equilibrium, making the system globally, asymptotically stable. $\square$

B. Proof of Theorem VIII.2

The following theorem establishes the necessary and sufficient conditions for local stability.

Theorem A.1 ([41]). *The equilibrium* $\mathbf{x}^*$ *of the system of equations* $\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x})$ *is locally stable if and only if all the eigenvalues of the Jacobian of the system of equations evaluated at* $\mathbf{x}^*$ *have real parts less than zero.*

Now we prove Theorem VIII.2:

Proof. First, we write the system equations and then, we obtain the equilibrium. We use (10) to obtain system dynamics. Consider some LGCC router i at depth d of the aggregation tree whose flows are aggregated into flow $x_{k,j}$ in loop j at depth $d+1$. The incoming flows to router i are governed by

$$\dot{x}_{1,i} = x_{1,i}(t)\gamma_d \left(1 - \frac{x_{1,i}(t)}{x_{k,j}(t)} - p_i(t)\right) , \quad (28a)$$

$$\dot{x}_{2,i} = x_{2,i}(t)\gamma_d \left(1 - \frac{x_{2,i}(t)}{x_{k,j}(t)} - p_i(t)\right) , \quad (28b)$$

$$\vdots$$

$$\dot{x}_{\zeta,i} = x_{b,i}(t)\gamma_d \left(1 - \frac{x_{b,i}(t)}{x_{k,j}(t)} - p_i(t)\right) , \quad (28c)$$

where

$$\dot{p}_i = \beta_d \Big[\big(\sum_{y=1}^{\zeta} x_{y,i}\big) - x_{k,j}\Big] . \quad (29)$$

In the loop at depth δ, $x_{k,j} = K$. The equilibrium of the above system is given by

$$p_i^* = \frac{\zeta - 1}{\zeta} , \quad (30)$$
$$x_{.,i}^* = (1 - p_i^*)x_{k,j}^* , \quad (31)$$

where $x_{k,j}^* = K$ for $d = \delta$. Now we obtain the Jacobian of the above system. In each loop, there are $\zeta + 1$ variables: ζ equations representing flows' rate, and one equation adjusting the cost. This makes a Jacobian matrix of dimension $\frac{(1+\zeta)(-1+\zeta^\delta)}{-1+\zeta}$ (this is obtained by considering that the aggregation tree is a perfect m-ary tree, where the number of nodes determines the number of rate dynamics, and the number of non-leaf nodes also determines the number of LGCC routers). Although obtaining eigenvalues of a matrix is not trivial, since the above matrix has a special form, we can calculate its eigenvalues for

an arbitrary ζ and δ. Each element of the Jacobian matrix is given by $\frac{\partial \dot{\chi}}{\partial v}$ where χ and v are any two system variables. For example, for $0 \leq b \leq \zeta$

$$\frac{\partial \dot{x}_{b,i}}{\partial x_{b,i}} = \gamma_d \left(1 - \frac{x_{b,i}}{x_{k,j}} - p_i \right) - \gamma_d \frac{x_{b,i}}{x_{k,j}} \quad , \tag{32}$$

$$\frac{\partial \dot{x}_{b,i}}{\partial p_i} = \gamma_d x_{b,i} \quad , \tag{33}$$

$$\frac{\partial \dot{p}_i}{\partial x_{b,i}} = \gamma_d \quad , \tag{34}$$

and the partial derivatives are zero for other variables. This matrix has a number of eigenvalues; some of them are negative, and some of them are in the form of

$$\frac{1}{4} \left(-\gamma_d + \sqrt{-4K\zeta^{(d+2-\delta)}\gamma_d \beta_d + \gamma_d^2} \right) \quad .$$

Constraining this eigenvalue to be negative and solving for a relationship between γ and δ yields (24). $\qquad\square$

APPENDIX B
FAIRNESS ANALYSIS

A. Proof of Theorem VIII.3

First, we provide a definition and theorem from [42] on max-min fairness.

Definition 2 (Bottleneck Link). *Link l is a bottleneck for source r if and only if link l is saturated, and source r has the maximum rate among all sources crossing link l.*

Theorem B.1. *A feasible rate allocation vector to all sources is max-min fair if and only if every source has a bottleneck link.*

Now we prove Theorem VIII.3:

Proof. At every LGCC router, there is a schedule per-output port. Taking into account that all the incoming flows that are forwarded towards the same output port are aggregated, and they are fairly scheduled, if this is LGCC router 1, then it is a bottleneck for source r; if the bottleneck is at some router i, then our weight assignment guarantees that again this is a bottleneck for source r, and hence, every source has a bottleneck, confirming that the policy attains max-min fairness. $\qquad\square$

B. Proof of Theorem VIII.4

Lemma B.2. *A proportionally-fair rate allocation vector assigns rates inversely proportional to the end-to-end cost the network charges each source r.*

Proof. With a logarithmic utility function, which attains proportional fairness, $x_r^* = \frac{1}{\sum_{l \in L_r} q_l}$. $\qquad\square$

Now we prove Theorem VIII.3:

Proof. According to the employed scheduling policy, as flow r crosses multiple routers, each router's cost is accumulated. Consider the last router with a non-zero cost: if it constrains x_r to $\frac{1}{\sum_{l \in L_r} q_l}$, then every flow r gets a share of the bandwidth inversely proportional to its cost. Since the sum of costs is less constrained at previous routers, it is the role of the last router to enforce fairness; previous loops of flow r are backpressured due to this constraint. This is similar to limiting the carrying capacity at the last router, not to increase more than a proportionally fair share. In [15], we proved that the function $-\log(1 - p_{r,i}^u)$ provides the cost using ECN. Therefore, setting weights inversely proportional to this value guarantees proportional fairness. $\qquad\square$