Scheduling in Densified Networks: Algorithms and Performance

Sharayu Moharir, Subhashini Krishnasamy, and Sanjay Shakkottai, *Fellow, IEEE*

***Abstract*— With increasing data demand, wireless networks are evolving to a hierarchical architecture where coverage is provided by both wide-area base stations (BS) and dense deployments of short-range access nodes (ANs) (e.g., small cells). The dense scale and mobility of users provide new challenges for scheduling: 1) high flux in mobile-to-AN associations, where mobile nodes quickly change associations with ANs (time scale of seconds) due to their small footprint and 2) multi-point connectivity, where mobile nodes are simultaneously connected to several ANs at any time. We study such a densified scenario with multi-channel wireless links (e.g., multi-channel OFDM) between nodes (BS/AN/mobile). We first show that traditional algorithms that forward each packet at most once, either to a single AN or a mobile user, do not have good delay performance. We argue that the fast association dynamics between ANs and mobile users necessitate a multi-point relaying strategy, where multiple ANs have duplicate copies of the data, and coordinate to deliver data to the mobile user. Surprisingly, despite data replication and no coordination between ANs, we show that our algorithm (a distributed scheduler—DIST) can approximately stabilize the system in large-scale instantiations of this setting, and further, performs well from a queue-length/delay perspective (shown via large deviation bounds).**

***Index Terms*— Dense Networks, downlink relay networks, wireless scheduling and routing.**

I. Introduction

THE WIRELESS industry is undergoing a sea change in cellular deployment. From a well-planned macro-cellular setting, the network is evolving to a hierarchical setting with cellular base-stations provide macro coverage (footprint of 1 km or more) and a dense deployment of access nodes (e.g., small cells [32] or femto cells [1], [2]) whose coverage range may be as little as 50 – 100 meters, provides short-range coverage. This combination – macro + dense short-range coverage – popularly referred to as *network densification*, leads to new challenges in network resource allocation.[1]

First, the access nodes' small footprints imply that mobile nodes associate and disassociate with them at a much higher rate than previously seen. A car moving at just 30 mph results in hand-offs between ANs at the time-scale of seconds. This will likely worsen with emerging technologies for 5G systems such as millimeter wave (mmWave) Broadband [28], where the radio propagation environment results in highly non-isotropic and direction-dependent short-range coverage.[2] Thus, to ensure universal coverage, operators have no recourse but to provide a very dense deployment of ANs (especially in locations with high data demand). This leads to second challenge: mobile nodes have the opportunity to associate with several possible ANs at any given time (however, this set changes rapidly over time due to mobility and coverage directionality).

In this paper, we argue that operating these dense networks in a traditional manner, where mobile nodes associate with one AN at any time, and then hands-off to a new one as the environment/location changes, can be inefficient. Instead, we study an approach where data packets are *replicated* at a collection of ANs whose footprints most-likely cover the mobile node, and these ANs deliver packets to the mobile user by making decisions in a decentralized manner using local information. We communicate directly between the base-station and the mobile node *only as a last resort* when the ANs are unable to reach the mobile node (e.g., due to uncertainty in tracking the mobile node, poor location, poor channel rates due to fading). We propose a formal model to capture this setting and analytically show the performance benefits.

Coordination in wireless communication has been studied in various contexts like Distributed/Virtual MIMO [26], [27], Network Coding [19], [21] etc. Importantly, these techniques require coordination at the packet or time-slot level. As discussed in [13], backhaul delays could be much larger

Manuscript received April 10, 2015; revised February 19, 2016; accepted May 9, 2016; approved by IEEE/ACM Transactions on Networking Editor D. Leith. Date of publication July 19, 2016; date of current version February 14, 2017. This work was supported in part by the National Science Foundation under Grants CNS-0964391, CNS-1017549, CNS-1161868, and CNS-1343383 and the U.S. Department of Transportation through the DSTOP Tier 1 University Transportation Center. An earlier version of this work was presented at the 52nd Annual Allerton Conference on Communication, Control and Computing, October 2014 [24].

S. Moharir is with IIT Bombay, Mumbai 400076, India (e-mail: sharayu.moharir@gmail.com).

S. Krishnasamy and S. Shakkottai are with The University of Texas at Austin, Austin, TX 78712 USA (e-mail: subhashini.kb@utexas.edu; shakkott@mail.utexas.edu).

Digital Object Identifier 10.1109/TNET.2016.2580614

[1] An earlier version of this work was presented at the 52nd Annual Allerton Conference on Communication, Control and Computing, October 2014 [24]. This work was partially supported by NSF grants CNS-0964391, CNS-1017549, CNS-1161868 and CNS-1343383, and the US DoT supported D-STOP Tier 1 University Transportation Center.

[2] For instance, in a mmWave Broadband system, the human body completely blocks radio propagation [20], [29]. Thus even slight movement (e.g., rotation of the human with the phone) can completely block the mmWave access node from communicating with a mobile node, thus leading to association changes that can occur within fractions of a second.

than the duration of a time-slot; further with densification, heterogeneity in backhaul delays will likely worsen. Thus, the key differentiating aspect from the above literature is that we consider the setting with delayed or sloppy coordination among the various access nodes. In our setting, access nodes do not have current knowledge of nearby nodes' instantaneous states, or indeed, even knowledge of which mobile nodes are connected to them.

A. Contributions

We study scheduling algorithms for networks with a base-station (BS) and multiple densified access nodes (AN) and multiple mobile users. We assume that the ANs are dense enough to support multi-point connectivity, i.e., each user can associate with multiple ANs at any given time. We propose an algorithm (DIST) for scheduling and evaluate its performance as detailed below.

1) **Algorithm DIST:** We propose a distributed algorithm called DIST where the BS and the ANs make their scheduling decisions independently, based only on local channel and queue-length information. Under the DIST algorithm, the BS forwards each packet to an AN that is currently connected to the intended user. If an AN cannot forward a received packet to the corresponding user because the user is no longer connected to it, unlike traditional algorithms, under the DIST algorithm, the AN forwards copies of packets to multiple ANs around it. In addition, if the ANs fail to deliver a packet to the user within a fixed number of time-slots, it is forwarded directly from the BS to the mobile user.
2) **Stability:** Under general arrival and bounded channel processes, we show that if the system scale is large enough, the DIST algorithm keeps the system stable (i.e., Markovian assumptions imply positive recurrence of the queues). Since DIST keeps packets in queues even after a transmission attempt and possibly makes multiple transmission attempts (to ANs and users), the proof involves tail analyses of coupled queues in order to even prove stability (which is not a tail behavior).
3) **Performance:** We have two performance results: *(i)* We first show that traditional algorithms like the BackPressure algorithm [33] in which the base-station forwards each packet at most once, either to a single access node or a mobile user, do not have good delay performance for mobile users, i.e., the delay rate functions are zero. *(ii)* For the proposed DIST algorithm, we show that for bounded i.i.d. arrivals and channels, the maximum queue-length rate function is strictly positive and therefore, the queue-length tails decay exponentially. Further, via simulations, we show that the DIST algorithm significantly outperforms the BackPressure algorithm in terms of the delay performance.

B. Related Work

Since the work by Tassiulas and Ephremides [33], there has been great interest in queue-length based scheduling in wireless networks (see [11] for a survey). In the many users/channels context (as in this paper), there has been recent activity to characterize stability, queue-length and delay performance, with and without relays (however without user mobility) [5]–[7], [17], [25], [31]. A key insight in these works has been the use of iterative allocations, where queues are updated to account for (partial) channel allocations even within a time-slot.

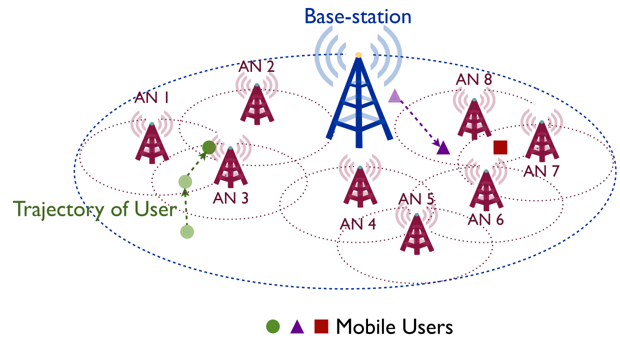

Fig. 1. A wireless network with a base-station, densely deployed ANs and mobile users. The users more in and out of the coverage area of the ANs due to mobility, but are always in the coverage area of the base-station. BS/AN image courtesy: http://intersales.com.au/GPSNetwork.aspx; see also [15].

This paper focuses on the benefits of data replication and multi-point connectivity in a mobile cellular setting (i.e., multiple access nodes maintaining active communications with a mobile user). Such access has had a long history, starting from CDMA soft-handoff (to enable make-before-break voice connections) [34], [35]. More recently, in the setting of COordinated Multi-Point (COMP) [23], there has been much work at the physical layer to develop cooperative communication strategies between a collection of base-stations and a mobile user. This is especially useful in densified settings, with increased opportunities (many base-stations/access points for coordination) and challenges (more complex interference management). These issues have been studied in various ways including simulations [8], field trials [3], [16], and information-theoretic techniques [12] (see [23] for a survey). In this paper, we focus on network level attributes – queue-lengths and delays – and show that even *local* scheduling algorithms that replicate data can significantly outperform more traditional scheduling algorithms.

Finally, as discussed in the introduction, coordination in wireless networks has a rich history and has been studied in various contexts like Network Coding, Multi-homing, virtual/distributed MIMO etc. See [36] for a discussion of challenges arising at different layers of the network protocol stack as a result of coordination in wireless communication networks. In this paper, we propose an algorithm which uses only local information, thus obviating the need for coordination between different ANs.

II. System Model

We consider a two-tiered downlink communication system with a base-station, a large number of ANs and mobile users as shown in Figure 1.

We study a multi-channel (e.g. OFDM) setting with a large number of orthogonal channels that can be used for communication simultaneously, i.e., each node (BS/AN) can transmit on multiple channels at once and each node (AN/user) can receive on multiple channels at once. This multi-channel setting, but without user mobility, was the focus in [25]. However, the fact that users are mobile and that the network is densified implies that the set of ANs that a mobile node is associated with is not time-invariant; further, a classical time-scale decomposition assumption between mobile-AN association and channel scheduling cannot be easily justified.

From a channel (average) rate perspective, our setting is one where the BS-AN, AN-AN and the AN-user links have higher data-rates than the BS-user links. Again, this is a natural setting to consider because the ANs are expected to be mounted in more suitable locations, as well as have superior hardware in terms of the number of antennas, as compared to the mobile users. Moreover, the ANs that a user is associated with are typically much closer to the user than the central BS.

Formally, the system consists of a base-station and $M(n) = O(n)$ ANs, where n is the number of users in the system.

We say that $f = O(g)$ if $\limsup_{n\to\infty} \left|\frac{f(n)}{g(n)}\right| < \infty$, and therefore, $M(n)$ can grow sub-linearly in n.

We assume that the ANs have two RF chains, one to communicate with the BS and the other to communicate with the users and other ANs. As recommended in [1], the BS-AN communication happens at a different spectrum than the BS-user and AN-user communication. To keep the notation simple, we assume that the number of orthogonal frequency channels for BS-AN communication and AN-user communication are n each. This setting was also considered in [25]. Our results can easily be extended for other linear scalings.

A. User Mobility

We use a general notion of mobility which allows both fast moving users that move in every time-slot as well as users which move rarely. Formally, we assume that the probability that a user moves from its current position between two consecutive time-slots is $\Omega(1/poly(n))$ (at least of the order of $1/poly(n)$). This assumption allows the expected time spent at a location to be anything between one and a polynomial function of n. For example, the probability that a user moves between two consecutive time-slots can be a constant independent of n as is the case for the Levy-walk process, which is known to be a good model for human mobility in various outdoor settings including college campuses and theme parks [30]. Other popular models, for instance, (discretized versions of) the Random Waypoint Mobility model (RWM) [18] and its variants that have been shown to be more appropriate for user mobility in cellular networks [22], also satifsy this condition.

B. User-AN Connectivity

Since we consider a setting where the ANs are densely deployed, the user is very likely to be connected to multiple ANs. However, we also include the possibility that, in some time-slots, the system fails to obtain the location information of a user. This could happen for various reasons: *(i)* when a user goes out of the coverage area of the ANs, *(ii)* when the user is within the communication range of some ANs, but fails to communicate its position to those ANs, or *(iii)* when a user tracking/position learning algorithm fails. Specifically, we assume that, at the beginning of each time-slot, the location of the user is known with probability at least $1 - \epsilon(n)$, i.i.d across users.

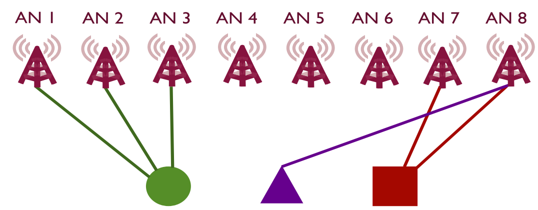

Fig. 2. Association graph between the ANs and mobile users in the network in Figure 1. Each AN is associated with all user that are currently in its coverage range, represented by an edge between the AN and the mobile user. AN image courtesy: http://intersales.com.au/GPSNetwork.aspx.

When the user location is known, it is connected to at least $C(n)$ ANs. The density of the ANs in the network can be non-uniform, and therefore, $C(n)$ imposes a lower bound on it. In this paper we work in the setting where $C(n)$ is $\Omega(\log n)$. Let I_i be the set of ANs which interfere with user i. We assume that the interference set for every user satisfies the following: For all i, $|I_i| \le n^{\beta}$, for some constant $\beta < 1$.

We also assume that the base-station can always communicate with all users albeit at lower (average) rates than the ANs. Figure 2 illustrates the user-AN association graph for the system shown in Figure 1.

1) Unpredictability of User-AN Associations: Since the users are mobile, the set of ANs a user is connected to can change between two consecutive time-slots. Let $M_u(t)$ be the set of ANs that user u is connected to in time-slot t. We assume that,

- Between two consecutive time-slots, the probability that a connected mobile user, u moves to a new location such that it is no longer associated with a previously connected AN, m is not negligible. Formally, for every t and $m \in M_u(t-1)$,
$$P(m \notin M_u(t)) \ge \mu_1(n),$$
where $\mu_1(n) = \Omega(1/poly(n))$.
- Further, the motion of the mobile user cannot be predicted with very high accuracy, i.e., for every t and $m \notin M_u(t-1)$,
$$P(m \in M_u(t)) \le 1 - \mu_2(n),$$
where $\mu_2(n) = \Omega(1/poly(n))$.

These two conditions are fairly general and are satisfied both by users moving at a very fast time-scale (every time-slot) to users that move rarely ($poly(n)$ time-slots in expectation). These conditions are also satisfied by the Levy walk process and the RWM model.

2) Concentration of Users Around an AN: In dense networks where each AN has a small footprint, it is unlikely that a large number of users will be connected to any one particular AN. Therefore, we can assume that, with high probability, not more than a constant fraction of the total number of users are connected to a particular AN at the same time. Specifically, if $U_m(t)$ is the set of users connected to AN m in time-slot t, then

$$P\left(\max_{1\leq m\leq M(n)} |U_m(t)| > n^{\nu}\right) \leq e^{-bn},$$

for a positive constant $\nu < 1-\beta$ and a constant $b > 0$. This condition is satisfied, for example, if the users are executing a lazy random walk on the network of ANs independent of other users in the system.

C. Communication Between Access Nodes

We consider the setting where each AN can communicate with $O(\text{poly}(\log n))$ other ANs located close to it. We assume that the set of ANs that a given AN m can communicate with is large enough so that even if a mobile user connected to AN m in time-slot t moves in the next two time-slots ($t+1$ and $t+2$), AN m can communicate with at least one AN in $M_u(t+2)$.

D. Interference Between Access Nodes

Although the dense deployment of ANs enables multi-point connectivity, it can cause interference at the mobile user due to simultaneous transmissions on the same channel. Let I_m be the set of ANs which interfere with an AN m, i.e., no AN in I_m can successfully transmit on the same channel as m. We assume that the interference set for every AN satisfies the following: For all m, $|I_m| \leq n^{\beta}$, for some constant $\beta < 1$. Note that this condition is quite general, and allows for interference sets that grow polynomially in the network size. As a point of reference, spatial stochastic models (where ANs are randomly scattered over the plane), and connectivity as suggested in the Gupta-Kumar model [14] have interference sets that scale only logarithmically in network size, and thus is allowed by our model.

E. Notation

We add to the notation previously used in [25] and [5]–[7] to incorporate mobility of users which leads to time-varying user-AN associations. There are n queues at the base-station and ANs (one per user). Our system evolves in discrete time $\{t = 0, 1, 2, \ldots\}$, where arrivals happen at the beginning of time-slots, and queues are updated at the end of a time-slot.

- Q_i, R_{mi} = Queue of mobile user i at the BS and at AN m respectively.
- $Q_i(t)$ = BS queue-length of mobile user i at the end of time-slot t.
- $\mathbf{Q}(t) = \{Q_i(t) : 1 \leq i \leq n\}$: BS queue-length vector (across all mobile users)
- $A_i(t), A_i^m(t)$ = Number of packet arrivals for mobile user i at the BS and AN m respectively at the beginning of time-slot t.
- $\mathbf{A}(t) = \{A_i(t) : 1 \leq i \leq n\}$: Arrival vector (across all mobile users).
- $M_i(t)$ = The set of ANs connected to mobile user i in time-slot t.
- $C(n) = \min_{i:M_i(t)\neq\phi} |M_i(t)|$ is the minimum number of ANs connected to mobile users whose locations are known.
- $U_m(t)$ = The set of mobile users connected to AN m in time-slot t.
- $X_{i,j}(t)$ = Channel rate (number of packets) for the j-th channel from the BS to mobile user i.
- $X_j^{B,m}(t)$ = Channel rate for the j-th channel from the BS to AN m.
- $X_{i,j}^m(t)$ = Channel rate for the j-th channel from AN m to mobile user i.
- $X_j^{l,m}(t)$ = Channel rate for the j-th channel from AN l to AN m.
- $H(p|q) = p\log\frac{p}{q} + (1-p)\log\frac{1-p}{1-q}$

In each time-slot, channels are allocated to service appropriate queues; this is captured via the decision variables $Y_{i,j}^{B,m}(t)$, $Y_{i,j}^m(t)$, $Y_{i,j}(t)$ and $Y_j^{l,m}(t)$ for users $1 \leq i \leq n$, channels $1 \leq j \leq n$ and ANs $1 \leq m, l \leq M(n)$ (each of the variables takes the value '1' if it corresponds to an allocation, and '0' otherwise).

Finally, $T_{i,j}(t)$ corresponds to the number of packets transmitted from user i's queue at the base-station on channel j.

III. Main Results and Discussion

A. Algorithm: DIST

DIST is a distributed algorithm where the BS and the ANs make their scheduling decisions independently. We first focus on the scheduling decisions at the BS.

The BS maintains three queues corresponding to each user for scheduling: the queue $A_i(t)$ corresponding to user i and holds all packets for user i which arrived at the BS at the beginning of time-slot i, the queue $Q_i(t)$ contains all packets that arrived for user i before time-slot $t-L-1$ and have still not been received by user i and the queue $F_i(t)$ contains all packets that arrived for user i at the beginning of time-slot $t-L$ and have still not been received by user i. By definition, at the end of time-slot t, all packets in $F_i(t)$ that have not been received by user i are moved to $Q_i(t)$. Here, we assume that at all times, the BS stores all packets which have not reached their final destination (user) even if they have been forwarded to the ANs by the BS.

The key idea is as follows. The BS first tries to forward new packets (packets in $A_i(t)$) to ANs currently connected to the users in a greedy manner, i.e., without using differential backlogs to make routing decisions. If there are unused channels after forwarding packets in $A_i(t)$, the BS focuses on packets in $F_i(t)$, i.e., packets which arrived L time-slots ago but have not yet reached the user directly to the users. It forwards these packets directly to the users. If there are unused channels after allocating packets in $A_i(t)$ and $F_i(t)$, these channels are used to forward packets in $Q_i(t)$ directly to user i.

We now formally describe the DIST algorithm at the BS. **Base-Station Algorithm:** The base-station algorithm proceeds in an iterative manner (see [6] for a detailed discussion of iterative algorithms), allocating one channel at a time. Queue-lengths are updated after each round of allocation. Channel k is allocated in iteration k.

1: *Forward New Arrivals to ANs*

$$\text{Find } \{i^*, m^*\} \in \underset{1\leq i\leq n, m\in M_i(t)}{\operatorname{argmax}} A_i^{(k-1)}(t) X_k^{B,m}(t).$$

where $A_i^{(k-1)}(t)$ is the updated (accounting for packets scheduled for transmission on channels from 1 to $k-1$) number of arrivals to user i in time-slot t and $M_i(t)$ is the set of ANs that user i is currently connected to. Packets for user i^* are scheduled for transmission from the base-station to the AN m^* on channel k.

2: *Direct Forwarding to Users*

If channel k is not used by the base-station to forward new arrivals to the ANs, search for the queue index

$$i^* \in \underset{1\leq i\leq n}{\operatorname{argmax}} F_i^{(k-1)}(t-1) X_{i,k}(t),$$

breaking ties in the favor of the smaller user index, where $F_i^{(k-1)}(t)$ is the updated (accounting for packets scheduled for transmission on channels from 1 to $k-1$). Allocate channel k to transmit $X_{i^*,k}(t)$ from the queue for user i^* at the base-station directly to user i^*.

If channel k is not used by the base-station to forward new arrivals to the ANs, and $F_{i^*}^{(k-1)}(t) X_{i^*,k}(t)$, search for the queue index

$$i^* \in \underset{1\leq i\leq n}{\operatorname{argmax}} Q_i^{(k-1)}(t-1) X_{i,k}(t),$$

breaking ties in the favor of the smaller user index, where $F_i^{(k-1)}(t-1)$ is the updated (accounting for packets scheduled for transmission on channels from 1 to $k-1$). Allocate channel k to transmit $X_{i^*,k}(t)$ from the queue for user i^* at the base-station directly to user i^*.

3: *Update Queue-lengths*

Update all queue-lengths before allocating the next channel.

Remarks: The salient features of DIST (at the Base-station level) are:

i. Step 1 – Local Information + Greedy: *Unlike the BackPressure algorithm, the DIST algorithm does not use differential backlogs to make its routing decisions and therefore does not try to balance the load at the ANs. Instead, the algorithm tries to push packets to the ANs in a greedy manner whenever it sees high channel rates.*

ii. Step 2 – Direct Forwarding over Free Channels: *Unused channels (i.e., unused by BS-to-AN transmissions) are used by the base-station to route packets which are queued at the base-station directly to the users. These packets may have previously been successfully received by one or more ANs, but which failed to forward it to the intended user within L time-slots. The base-station transmits these packets directly to the users.*

Access Node Algorithm: We now describe how each AN carries out the task of channel allocation.

For each AN m, we define two sets:

- $V_m :=$ the set of ANs that AN m can communicate with.
- $D_m(t) := \{u : m \in M_u(t-1) \setminus M_u(t)\}$ be the set of users which were connected to AN m in the previous time-slot, but are not connected to AN m in this time-slot.

Remarks: Before we formally describe the algorithm, the key features of DIST (at the AN level) are:

i. Local Information: *Each AN makes its decisions using local queue-length and channel information (channel rates to from m to ANs in V_m and users connected to AN m).*

ii. Forwarding Strategy: *For users that are connected to the AN, the AN forwards packets directly to the users. Packets for users that were connected to the AN in the previous time-slot, but are no longer connected to the AN in the current time-slot (users in the set $D_m(t)$), are forwarded to neighboring ANs (ANs in the set V_m).*

iii. Channel Randomization: *Each AN chooses the channel it transmits on uniformly at random from the set of channels which have the highest channel rate. This can lead to collisions, but, since we work in the large scale multi-channel setting, the expected number of collisions are a vanishing fraction of the supportable load.*

iv. Purging Old Packets: *The ANs do not cache packets for more than a fixed number of time-slots (say L). Any packet which arrived at the base-station at the beginning of time-slot t is deleted by the ANs at the end of time-slot $t + L$.*

Formally, each AN m implements the following steps:

1: Initialize $J = \{1, 2, \ldots n\}$, $k = 1$, $B_u^{l(0)}(t) = A_u^m(t)$ for $u \in D_m(t)$ and $l \in V_m$, and $A_i^{m(0)} = A_u^m(t)$ for $u \in U_m(t)$.

2: *Forward Packets to Connected Users*

If $\max_{i\in U_m(t)} A_i^{m(k-1)} = 0$, and goto step 4. Else,

$$\{i^*, j^*\} \in \underset{i\in U_m(t), j\in J}{\operatorname{argmax}} A_i^{m(k-1)} X_{i,j}^m(t),$$

breaking ties uniformly at random. Allocate channel j^* to serve the queue for user i^* and update $J = J \setminus j^*$.

3: $A_{i^*}^{m(k)} = (A_{i^*}^{m(k-1)} - X_{i^*,j^*}^m(t))^+$, $k = k + 1$, and goto Step 2.

4: *Forward Packets to Neighboring ANs*

$$\{l^*, u^*, j^*\} \in \underset{l\in V_m, u\in D_m(t), j\in J}{\operatorname{argmax}} B_u^{l(k-1)} X_j^{m,l}(t),$$

breaking ties uniformly at random. Allocate channel j^* to forward packets for user $u^* \in D_m(t)$ to AN l^* and update $J = J \setminus j^*$.

5: $B_{u^*}^{l^*(k-1)} = (B_{u^*}^{l^*(k-1)} - X_{j^*}^{m,l^*}(t))^+$, $k = k + 1$, and goto Step 4.

B. Stability

Our first result characterizes a necessary condition for stability.

Theorem 1: If the load $\lambda > 1$, no algorithm can stabilize the queues (i.e., render the queue to be positive recurrent).

Next, we focus on the DIST algorithm. The DIST algorithm allows the base-station to retransmit packets which have already been received by one or more ANs. Retransmission can lead to the instability of queues in the system, but we show that under some reasonable assumptions on the channel and arrival processes, the DIST algorithm stabilizes the queues in the system. These assumptions are analogous to those in [25] (see also [6], [10]), with the natural additions to account for user mobility.

Assumption 1 (Bounded Channel Processes):

- *The channel processes are i.i.d. across time-slots (and independent of the arrival process).*
- $X^{B,m}_{i,j}(t) \le C_{max} < \infty$.
- $X^{m}_{i,j}(t) \le C_{max} < \infty$.
- $X^{m,l} \le C_{max} < \infty$.
- $X_{i,j}(t) \le C^{d}_{max} < C_{max} < \infty$.
- *For every j and user i in time-slot t,*
$$P(X_{i,j}(t) = C^{d}_{max}) \ge q_1^{(C^d_{max})} > 0.$$
- *For every j, t and user i connected to AN m in time-slot t,*
$$P(X^{m}_{i,j}(t) = C_{max}) \ge q_2^{(C_{max})} > 0.$$
- *For every j, t and every AN l which can communicate with AN m,*
$$P(X^{m,l}_{j}(t) = C_{max}) \ge q_3^{(C_{max})} > 0.$$

Assumption 2 (Arrival Process):

- *We assume that $\boldsymbol{A}(t)$ (arrival vector per time-slot) is an aperiodic, irreducible, finite state Discrete Time Markov Chain.*
- $A_i(t) \le \kappa(n)$ *such that* $\kappa(n)\epsilon(n) = o(1)$ *and* $\kappa(n)n^{\nu}n^{\beta} = o(n^{\alpha})$ *for some* $\alpha < 1$.
- *We define the load* $\lambda = \frac{1}{n}E\left[\sum_{i=1}^{n}\left\lceil\frac{A_i(0)}{C_{max}}\right\rceil\right]$. *Then,*
$$P\left(\sum_{i=1}^{n}\left\lceil\frac{A_i(0)}{C_{max}}\right\rceil = n(\lambda+\delta)\right) = o\left(\frac{1}{n}\right),$$
for any $\delta > 0$.

Recall that $\epsilon(n)$ is the probability that a user cannot be located by the ANs in a time-slot; thus requiring the BS to use a (lower rate) channel to directly transmit packets to the mobile. Clearly, as $\epsilon(n)$ increases, $\kappa(n)$ has to decrease to maintain stability of the system. The assumption $\kappa(n)\epsilon(n) = o(1)$ quantitatively captures this effect. For example if $\epsilon(n) = 1/\sqrt{n}$, users can have up to $o(\sqrt{n})$ arrivals in a time-slot.

Recall that w.h.p., the number of users connected to an AN in a time-slot is less than n^{ν} and the size of the interference set for each AN is at most n^{β}. Therefore, $\kappa(n)$ has to be small enough to ensure that using n channels, it is possible for each AN to forward all incoming packets to the corresponding users or other ANs without coordinating with other ANs, yet the number of collisions in each time-slot is a vanishing fraction of the total load on the system.

Assumption 3 (Base-station to AN Channel Process): Consider the event F_1 that for channel j, $X^{B,m}_{j} < C_{max}$ for all ANs user i is connected to in time-slot t. This is equivalent to saying that in time-slot t, channel j cannot be used at rate C_{max} by the base-station to forward packets for user i to the ANs. Then,
$$P(F_1) = o\left(\frac{1}{n^2}\right).$$

Assumption 4 (AN to Users Channel Process): For an AN m and user i connected to AN m, consider the event F_2 that there exist at least $n\frac{q_2^{(C_{max})}}{2}$ channels such that $X^{m}_{i,j}(t) = C_{max}$ for each channel. Then,
$$P(F_2^c) = o\left(\frac{1}{n^3 M(n)}\right).$$

Assumption 5 (AN to AN Channel Process): For an AN m which can communicate with AN l, consider the event F_3 that there exist at least $n\frac{q_3^{(C_{max})}}{2}$ channels such that $X^{m,l}_{j}(t) = C_{max}$ for each channel. Then,
$$P(F_3^c) = o\left(\frac{1}{n^3 \log n}\right).$$

Assumption 6 (Base-station to Users Channel Process):

- *Let I be a set of users such that that $|I| \ge kn^{1-\alpha}$, for constants $\alpha < 1$ and $k < 1$. Consider the event F_4 that for a channel j and for every user $i \in I$, $X_{i,j}(t) < 1$, $\forall i$. Then,*
$$P(F_4) = o\left(\frac{1}{n^3}\right).$$
- *Let I be a set of user such that that $|I| = kn^{1-\alpha}$, for constants $\alpha < 1$ and $k < 1$, and let J be a set of channels such that $|J| = \frac{2kn^{1-\alpha}}{q_1^{(C^d_{max})}}$, where*
Consider the event F_5 that for every relay in I there exist $kn^{1-\alpha}$ channels in J such that $X_{i,j}(t) = 1$. Then,
$$P(F_5^c) = o\left(\frac{1}{n^3}\right).$$

For instance, these assumptions (1.3 – 1.6) are satisfied by i.i.d. Bernoulli(q) channels, or more generally, by correlated (across users) channels that have a spatial correlation decay property (modeled via the α-mixing condition [4]).

Theorem 2: Under Assumption 1, for a given load $\lambda < 1$, there exists $n_0(\lambda)$ such that for all $n > n_0(\lambda)$, the Markov Chain corresponding to the queue-lengths at the base-station and access nodes is positive recurrent.

Thus, for n large enough, the DIST algorithm stabilizes the system for all loads $\lambda < 1$. Further, this is tight in the sense that beyond $\lambda = 1$, we cannot stabilize the queues by any means. This result is interesting because user mobility and collisions at the second hop (AN-user links) lead to retransmissions of those packets by the base-station and yet in the large-scale setting, the DIST algorithm keeps the system stable.

The proof leverages the fact that as the system scale increases, even if a user moves, at least 1 AN that the user is currently connected to has a copy of all the packets which arrived at the base-station less than L time-slots before the currently time-slot. Therefore, even if the user changes its position, it can receive packets from the ANs it is currently connected to. Moreover, as the number of channels increases, there are sufficient degrees of freedom in the system to ensure that the number of collisions is a vanishing fraction of the supportable load. Therefore, for a given load, as the system scale increases, there is sufficient additional capacity in the system to retransmit packets which are lost due to collisions and directly forward packets from the base-station to those users whose location information is not known. Therefore, we conclude that, in large scale systems, the benefits of multi-point connectivity can be achieved without the overhead of coordination.

C. Performance

1) Single Transmission (ST) Algorithms: We first characterize the performance of a class which we refer to as Single Transmission (ST) algorithms. An algorithm belongs to this class if it satisfies the following two conditions:

i. Each packet is transmitted successfully by the base-station at most once i.e. once the intended receiver (AN/user) of a packet receives it successfully, the base-station deletes that packet from its queue.
ii. Each AN forwards a received packet only to the corresponding user.

This class of algorithms includes the BackPressure algorithm [33] which is known to be throughput optimal for multi-hop systems. Iterative versions of the BackPressure algorithm and the MaxWeight algorithm were proposed for multi-channel systems in [25] and were shown to have good buffer-usage or delay performance for system in which users are not mobile. These algorithms too belong to the ST class of algorithms. The next theorem characterizes the performance of algorithms belonging to the ST class for mobile users.

Theorem 3: For a mobile user in a system implementing an ST algorithm, the delay for a packet that is routed to an AN by the base-station is such that

$$d := \limsup_{n\to\infty} \frac{-1}{n} \log P(\mathit{Delay} > r) = 0,$$

for any $r < \infty$.

We thus conclude that traditional algorithms like BackPressure/MaxWeight [33] do not have good delay performance for mobile users.

2) DIST: We study the buffer overflow probability for the largest queue at the base-station:

$$c := \liminf_{n\to\infty} \frac{1}{r+1} \frac{-1}{n} \log P\bigg(\max_{1\le i\le n} Q_i(0) > r\bigg).$$

This value of r is a bound on the rate of decay of the longest queue (large deviations rate function). Note that the queues at the access nodes delete packets within a small number of time-slots; thus stability or performance of these access node queues is not the focus here.

If an algorithm results in a positive value of r, then we have that (neglecting constants outside the exponent)

$$P\bigg(\max_{1\le i\le n} Q_i(0) > r\bigg) \approx e^{-c(r+1)n}.$$

Therefore, the probability that the system has any backlogged packets goes to zero very quickly which means that all packets that enter the system are served almost immediately, thus leading to low delay.

We analyze the performance of DIST for a restricted set of arrival and channel processes.

Assumption 7 (Multi-level Bounded Arrivals and Channels):

- $A_i(t) = k$ *w.p.* p_k *for* $0 \le k \le K$ *and 0 otherwise*
- $X_i^{B,m}(t) = c$ *w.p.* $q_1^{(c)}$ *for* $0 \le c \le C_{max}$ *and 0 otherwise.*
- $X_{i,j}^m(t) = c$ *w.p.* $q_2^{(c)}$ *for* $0 \le c \le C_{max}$ *if user* i *is connected to AN* m *and 0 otherwise.*
- $X_j^{m,l}(t) = c$ *w.p.* $q_3^{(c)}$ *for* $0 \le c \le C_{max}$ *if AM* m *can communicate with AN* l *and 0 otherwise.*
- $X_{i,j}(t) = 1$ *w.p.* q_4 *and 0 otherwise.*
- $\epsilon(n) = o(1)$.

The arrival and channel processes are i.i.d. across users, ANs and time-slots. In addition we assume that $C(n) \ge 2\log n$.

Theorem 4: Under Assumption 7, for the DIST algorithm, for any integer $r \ge 0$,

$$c = \liminf_{n\to\infty} \frac{-1}{n} \log P\bigg(\max_{1\le i\le n} Q_i(0) > r\bigg) > 0.$$

From this theorem we conclude that under Assumption 7, using multi-point connectivity, good buffer-usage performance can be achieved without the overhead of multi-point coordination.

Like the proof of Theorem 2, this proof too leverages the fact that as the system scale increases, multi-point connectivity and the large number of channels ensure that the number of collisions is small and direct retransmission of packets from the base-station to users ensures that no packets stays in the system for too long.

IV. Proof Outlines

In this section, we provide proof outlines. Refer to the Appendix for detailed proofs.

A. Stability of DIST (Theorem 2)

Stability of multihop systems has been studied in literature in numerous settings, [25] being closest to the setting in the paper. In [25], stability of a static multihop system (no user mobility) for an iterative version of the MaxWeight algorithm was proved in a sequential manner by first showing the stability of base-station queues followed by showing that the relay queues are also stable. The reason why such a decoupling is possible in [25] is that the MaxWeight algorithm is an ST (Single Transmission, see Section III-C) algorithm and therefore, once a packet is forwarded by the base-station to a relay/user, it is deleted from the queue at the base-station. The queue process at the base-station is therefore independent of the packet transmissions at the second hop (relay-user links). However, for the DIST algorithm, every packet in the system which has not reached its final destination (user) is queued the base-station even if it has been forwarded to the ANs. This couples the queue processes at the base-station with the channel allocation at the second hop (AN-user links).

Therefore, unlike [25], where stability was proved in a sequential manner, we have to analyze the entire system at once which requires a different proof structure. Moreover in our setting, since all packets which have not reached their destination (user) are queued at the base-station, it suffices to show that the base-station queues are stable in order to show stability of the system.

Apart from this key difference, the analysis of DIST has three other new aspects.

1) *Dealing with missing user location information:* Unlike settings considered previously, we deal with users whose location is sometimes unknown. We show that there is sufficient unused capacity for DIST to directly forward packets to such users (see also (3) below).
2) *Decentralized nature of DIST:* The ANs forward packets received to connected users and scheduling decisions are made in a distributed manner. This can lead to two bad events: (i) there are packets which no AN forwards to a user, and (ii) due to collisions, packets are not received successfully by the users. We show that for the DIST algorithm, the number of such bad events in the second hop (AN-user links) is $o(n)$ with probability $\geq 1 - o(e^{-n})$.
3) *Splitting packets at the BS into new and old packets:* The base-station forwards new arrivals to the ANs and old packets (packets that arrived more than L time-slots before the current time-slot) directly to the users. We show that all new arrivals for users whose location is known are forwarded to the ANs by the BS in a given time-slot with probability ($\geq 1 - o(1/n)$). We then show that there is sufficient additional capacity in the system (channels unused by the BS-AN links) to ensure that all packets that arrived L slots before the current time-slot t, and which could not be forwarded by the ANs either due to collisions due to the decentralized nature of DIST or because the location of those users was not known can be sent directly from the base-station to the users in time-slot t.

Using these properties of the DIST algorithm, we show that on average, the base-station queues can serve more packets than they receive in a time-slot (accounting for both new arrivals and old packets that re-enter the base-station queues because the ANs fail to forward them to the users). We then use the standard Foster's Lyapunov technique for Markov Chains (with a quadratic Lyapunov function) to show stability. In other words, for a given load λ, there exists a constant n_0 such that the DIST algorithm ensures that the base-station queues in a system with $n > n_0$ channels are positive recurrent.

B. Performance Analysis of ST Algorithms (Theorem 3)

1) We consider a packet p for a mobile user u which is sent to AN m by the base-station in time-slot t. Let F be the event that the user never connects to m in time-slots $t+1$ to $t+r$. By the assumptions made in Section II,
$$P(F) \geq \min\{\mu_1, \mu_2\}.\mu_2^{r-1}$$
Since $\mu_1, \mu_2 = \Omega(1/poly(n))$, we have that $P(F) = \Omega(1/poly(n))$.
2) Conditioned on F, the packet cannot reach the user u before the end of time-slot $t+r$. Therefore we conclude that,
$$d = \limsup_{n\to\infty} \left(-\frac{1}{n} \log P(\text{delay} > r) \right) = 0.$$

C. Performance Analysis of DIST (Theorem 4)

We use Markov Chain coupling results in [7] to prove this theorem. Unlike in [7] where the coupling results were introduced, or in [25] where multihop static networks were studied using similar coupling arguments, the analysis of the DIST algorithm for the setting in this paper is more challenging because of user mobility, missing user location information, collisions at the second hop (AN-user links) and most importantly, due to coupling between the base-station queues and transmissions on the AN-user links as discussed in the proof outline for Theorem 2. Instead of sequentially looking at base-station queues followed by relay queues as in [25], we characterize the buffer-usage at the base-station queues since all packets which have not reached their final destination are queued at the base-station.

Our key step is the construction of a new random variable $Y^{(n)}(t)$ which dominates the maximum queue-length process at the base-station in our system. $Y^{(n)}(t)$ is a Markov Chain which increases by at most a constant with a very small probability (e^{-rn}, for some constant $r > 0$) and decreases by '1' with constant probability. This construction enables us to use coupling results from [7] to lead to the desired result.

The first step in the construction is to show that for the base-station queues, the probability that the maximum queue-length increases in a slot is small ($\leq e^{-nr}$). We do this via the following steps:

1) *Number of Collisions is Small:* From the proof of Theorem 2 (stability of DIST), we know that the number of bad events (collisions and packets not being forwarded by ANs) in the second hop (AN-user links) is $o(n)$ with probability $\geq 1 - o(e^{-n})$.
2) *Most new arrivals reach users in L time-slots:* We show that all new arrivals to the base-station for users whose location is known are forwarded to the ANs in a given time-slot with high probability ($\geq 1 - e^{-nr}$) and by (1), we know that all but $o(n)$ of them are forwarded to the users or ANs connected to that user in the next two time-slots.
3) *Direct Transmission of Old Packets:* We show that there is sufficient additional capacity in the system (channels not used by the BS-AN links for new packets) so that all packets which could not be forwarded by the ANs to the users by L time-slots after their arrival into the system can be sent directly from the base-station to the users in time-slot t with high probability ($\geq 1 - e^{-nr}$).

The above steps ensure that the maximum queue-length does not increase (with exponentially high probability). When combined with long-term stability arguments as in Lemma 8, [6], it can be shown that within a finite number of time-steps, the maximum BS queue-length (in the dominating system) decreases by a fixed amount with constant probability. Finally, by explicitly analyzing the dominating system, we obtain the desired bounds on the decay rate of the maximum base-station queue-length.

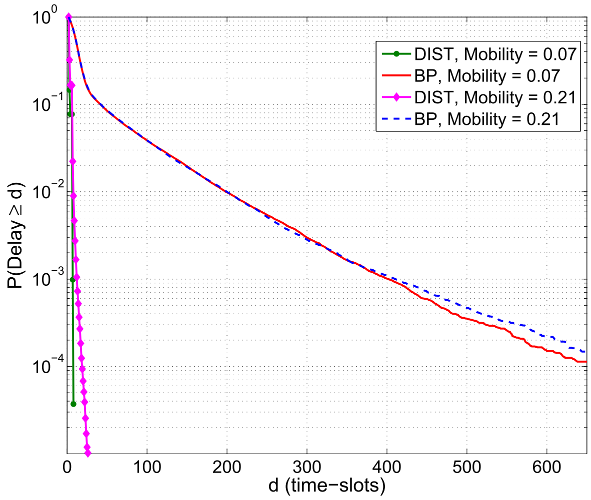

Fig. 3. BackPressure v/s DIST: Delay Performance.

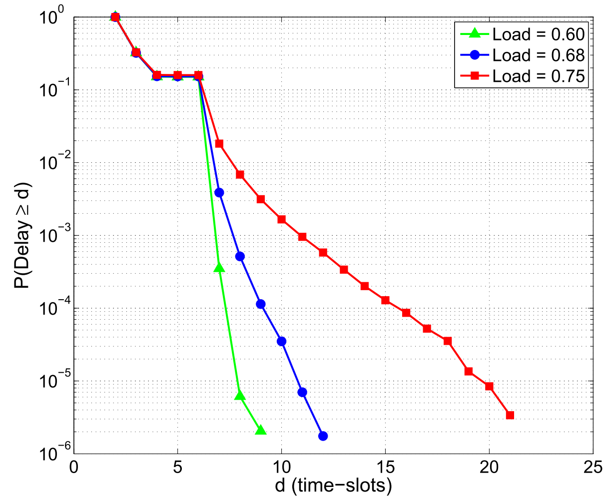

Fig. 4. DIST: Delay Performance for Different Loads.

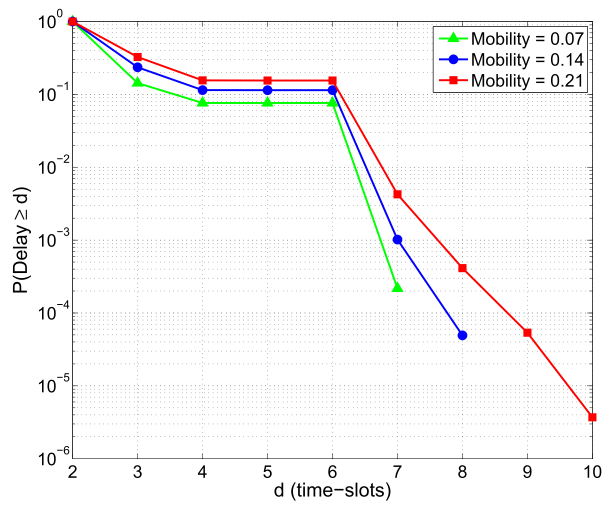

Fig. 5. DIST: Delay Performance for Different User Mobility.

V. Simulation Results

We compare the delay performances of the DIST algorithm and the BackPressure algorithm and also study the effect of various system parameters on the delay performance of DIST via simulations. We choose the BackPressure algorithm as a benchmark because it is known to be throughput-optimal for multihop networks.

We consider a system consisting of a base-station, 50 users, 50 channels and 50 ANs and run it for 10^4 time-slots. We assume that the ANs lie on a line. The mobility of users is a lazy random walk on this line. Unless specified, each user is assumed to be connected to the three nearest ANs. Each AN can communicate with the four nearest (two on either side) ANs and interferes with the six nearest (three of either side) ANs. The BS-AN channels and the AN-User channels take the value 0, 1 and 2 with probabilities 0.2, 0.3 and 0.5 resp and the BS-User channels take the value 1 and with probabilities 0.8 and 0 otherwise, i.i.d. across users, ANs and time-slots. We assume that $\epsilon(n) = 0.01$. The following plots show delays for arrivals across time-slots and users (the arrival process is symmetric).

Figure 3 compares the performance of DIST with $L = 5$ and the BackPressure algorithm. In this plot, the parameter Mobility is defined to be the probability that a user moves between two consecutive time-slots. The load on the system in this plot is 0.7. There is a significant difference in the performance of the two algorithms.

Figure 4 summarizes the performance of DIST with $L = 5$ for three different loads. Since $L = 5$, if a packet does not reach its destination (intended user) in 5 time-slots, the base-station tries to forward it to the mobile user directly. Such packets reach the user with a delay of at least $L + 1$. From Figure 4, we see that most packets reach their destination in $L+1$ (=6) time-slots. As expected, packet delays increase with load, but, compared to the performance of the BackPressure algorithm in 3, the delay performance of DIST is significantly better even for higher loads.

Figure 5 compares the performance of DIST with $L = 5$ for different values Mobility (the probability that a user moves between two consecutive time-slots) and load $= 0.7$. The delay performance worsens as the mobility of users in the system increases, but, remains comparable even when the probability a user moves between two consecutive time-slots triples. We conclude that the delay performance of the DIST algorithm is quite robust to user mobility.

Figure 6 compares the performance of DIST for load $= 0.7$ and Mobility $= 0.1$ for different values of the parameter L which is the number of time-slots the BS waits to let the ANs try and delivery packets to the intended users. If a packet does not reach the intended user via an AN within L time-slots after its arrival, the BS directly sends it to the user. As can be seen in Figure 6, most of the packets reach their final destination within $L + 1$ slots after their arrival.

VI. Discussion

The DIST algorithm outperforms traditional algorithms like the BackPressure algorithm in terms of end to end delay for two reasons: the first reason is that each AN which receives a packet for a particular user is not constrained to forward the

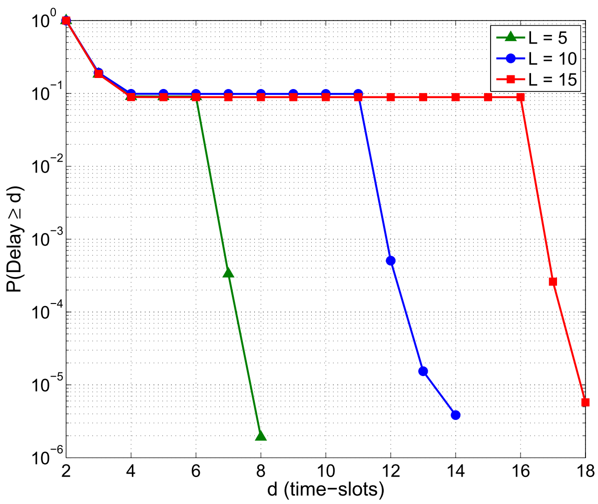

Fig. 6. DIST: Delay Performance for values of the parameter L (number of time-slots the BS waits to let the ANs try and delivery packets to the intended users before directly forwarding it to the user).

paper only once and only to the intended user and secondly, if the ANs fail to deliver a packet to the user in L time-slot, the BS forwards the packet directly to the user as a last resort. This improved delay performance is achieved as a result of increased energy consumption. For instance, under the DIST algorithm, if the BS forwards a packet for user i to AN j in time-slot t, unless user i stays connected to AN j in time-slot $t+1$, the AN forwards the packet to all its neighboring ANs. Compared to this, in the traditional BackPressure type algorithms, each packet is transmitted at most twice, once by the BS and at most once by at most one AN.

VII. Proofs

A. Stability

Theorem 5: Under Assumption 1, if the load $\lambda > 1$, then the system is unstable under any scheduling algorithm.

Proof: If $\lambda > 1$, then the mean number packet arrivals to the system in a given time-slot is more than the maximum number of packets that can be served by the base-station in a given time-slot ($= nC_{max}$). Hence the system is unstable under any scheduling algorithm. □

We now prove that if the load on the system $\lambda < 1$, the DIST algorithm stabilizes the system.

Lemma 1: The arrival process at the ANs is bounded by $(L+1)\kappa(n)$, i.e., $\max_{1\le i\le n, 1\le m\le M(n)} A_i^m(t) \le (L+1)\kappa(n)$ for all t.

Proof: In time-slot t, the BS only sends those packets which arrived at the BS at the beginning of time-slot t to the ANs. In addition, at the end of time-slot $t-1$, the ANs delete all packets in their queues which arrived at the BS before time-slot $t-L-1$. Since the arrival process at the base-station is bounded by $\kappa(n)$ by Assumption 1, we have that the arrival process at the ANs is also bounded by $(L+1)\kappa(n)$. □

Lemma 2: Let $U(t)$ be the set of users whose location is known at the beginning of time-slots t and $t+1$. Consider the set of all packets for users $u \in U(t)$ which are received and queued at least one AN at the beginning of time-slot t. A packet for user u in this set is said to be "lost" in time-slot t if at the beginning of time-slot $t+1$, it is neither received by the user u nor by at least one AN connected to user u in time-slot $t+1$. Let $\sigma = (1-\lambda)/7$ and E_1 be the event that in time-slot t, at most σn of the packets are "lost". For the DIST-AN algorithm,

$$P(E_1^c) \le o\left(\frac{1}{n^2}\right).$$

Proof: Let $q_2 := q_2^{(C_{max})}$. Consider the event E_2 that for each (AN, user) pair, in a time-slot, the number of channels that have channel rate C_{max} is at least $\frac{nq_2}{2}$. By Assumption 4,

$$P(E_2^c) \le o\left(\frac{1}{n^2}\right).$$

The rest of this proof is conditioned on E_2. Let $q_3 := q_3^{(C_{max})}$. Consider the event E_3 that from all ANs to all other ANs within their communication radius, in a time-slot, the number of channels that have channel rate C_{max} is at least $\frac{nq_3}{2}$. By Assumption 5,

$$P(E_3^c) \le o\left(\frac{1}{n^2}\right).$$

The rest of this proof is conditioned on E_3. Let E_4 be the event that not more than n^ν users are connected to any 1 AN for a constant $\nu < 1$, as discussed in Section II. Then, we have that,

$$P(E_4^c) \le e^{-bn},$$

for some $b > 0$. The rest of this proof is conditioned on E_4. By Lemma 1, the arrival process at the ANs is bounded by $(L+1)\kappa(n)$. Therefore, every AN has at most $(L+1)n^\nu\kappa(n)$ packets to send in a time-slot and can use at most $k_1(L+1)n^\nu\kappa(n)\log n$ channels to do so since by our assumption in Section II, each AN can communicate with at most $k_1 \log n$ other ANs. Therefore, each packet is sent to either $k_1 \log n$ other ANs or one user. So, for each AN, in any round of channel allocation, there are at least $\frac{nq_2}{2} - k_1(L+1)n^\nu\kappa(n)\log n$ channels that have channel rate C_{max} for every user and $\frac{nq_3}{2} - k_1(L+1)n^\nu\kappa(n)\log n$ channels that have channel rate C_{max} for every AN within its communication radius.

Therefore,

$$P(\text{AN } m \text{ uses channel } j \text{ to transmit packet } p|E_2,E_3,E_4) \le \frac{1}{\frac{nq^{(c)}}{2} - k_1(d+1)n^\nu\kappa(n)\log n},$$

where $q^{(c)} = \min\{q_2, q_3\}$.

The event E_1^c implies that more than σn packets are lost. Let $\mathcal{L}$ be a subset consisting of σn packets, i.e., $\mathcal{L} = \{p_1, p_2, , p_{\sigma n}\}$. Let "$\mathcal{L}$ lost" imply that all packets in $\mathcal{L}$ are lost. We have that,

$$P(E_1^c|E_2,E_3,E_4) \le \binom{n(d+1)C_{max}}{n\sigma} P(\mathcal{L} \text{ lost})$$

Let $p \leftarrow p^*$ mean that packet p is lost due to a collision with packet p^* at the destination of p. The size of the interference set for a user or an AN is at most n^β. Therefore, each packet sent to a user or an AN can collide with at most $k_1(L+1)n^\nu n^\beta \kappa(n) \log n$ other packets.

$$\begin{aligned} P(\mathcal{L} \text{ lost}) &\leq k_1(L+1)n^\nu n^\beta \kappa(n) \log n \\ &\quad P(\mathcal{L} \setminus \{p_1, p_1^*\} \text{ lost} | p_1 \leftarrow p_1^*) P(p_1 \leftarrow p_1^*) \\ &\leq \left(\frac{k_1(d+1)n^\nu n^\beta \kappa(n) \log n}{\frac{nq^{(c)}}{2} - k_1(d+1)\kappa(n)n^\nu \log n} \right) \\ &\quad \times P(\mathcal{L} \setminus \{p_1, p_1^*\} \text{ lost} | p_1 \leftarrow p_1^*). \end{aligned}$$

Repeating this arguement for all the packets in $\mathcal{L} \setminus \{p_1, p_1^*\}$, we have that,

$$P(\mathcal{L} \text{ lost}) \leq \left(\frac{k_1(d+1)n^\nu n^\beta \kappa(n) \log n}{\frac{nq^{(c)}}{2} - k_1(d+1)\kappa(n)n^\nu \log n} \right)^{\sigma n/2}$$

Therefore,

$$\begin{aligned} P(E_1^c | E_2, E_3, E_4) &\leq \binom{n(d+1)C_{max}}{n\sigma} \\ &\quad \times \left(\frac{k_1(d+1)n^\nu n^\beta \kappa(n) \log n}{\frac{nq^{(c)}}{2} - k_1(d+1)\kappa(n)n^\nu \log n} \right)^{\sigma n/2} \\ &= o\left(\frac{1}{n^2}\right). \end{aligned}$$

by Assumption 1, and the assumptions made in Section II. □

Lemma 3: If the load $\lambda < 1$, all arrivals to the base-station at the beginning to time-slot t for users whose location information is available are sent to the ANs by the BS in $(\lambda + \sigma)n$ rounds of channel allocation with probability $= 1 - o(1/n)$, where σ is as defined in Lemma 2.

Proof: Let E_6 be the event that not more than $n(\lambda + \sigma)C_{max}$ packets arrive at the base-station at the beginning of time-slot t. By Assumption 1, we have that,

$$P(E_6^c) = o\left(\frac{1}{n}\right).$$

The rest of this proof is conditioned on the event E_6. Let E_7 be the event that in the first $(\lambda + \sigma)n$ rounds of channel allocation, all channels can be used by the base-station to send packets from the longest BS queue (after updating after previous channel allocations) to the ANs, at rate C_{max}. Under Assumption 3, the probability that a particular channel k cannot be used to forward packets from the longest queue to the ANs connected to the corresponding user is $o(1/n^2)$. Taking a union bound over all channels, we have that,

$$P(E_7^c) = o\left(\frac{1}{n}\right).$$

Conditioned on E_7, all packets which arrived at the base-station at the beginning of time-slot t will be served in the first $(\lambda + \sigma)n$ rounds of channel allocation. Therefore, all arrivals to the base-station at the beginning to time-slot t are sent to the ANs in the first $(\lambda + \sigma)n$ rounds of channel allocation, with probability $= 1 - o\left(\frac{1}{n}\right)$. □

Lemma 4: For any t,

1) $\max_{1 \leq u \leq n} F_u(t) \leq \kappa(n)$
2) *In a time-slot t.*

$$P\left(\sum_{i=1}^{n} F_i(t) \leq 4n\sigma\right) \geq 1 - o\left(\frac{1}{n}\right),$$

where σ is as defined in Lemma 2.

Proof: The first part of the lemma follows from the fact that the arrival process for a user u at the BS is bounded by $\kappa(n)$ and therefore, the number of packets for a user u that arrived before time-slot $t - L$, but could not be received by the corresponding users by time-slot t is $\leq \kappa(n)$.

This proof is conditioned on the event E_7 defined in Lemma 3. We compute the probability of the event E_8 that in a time-slot, not more than $\frac{n\sigma}{\kappa(n)}$ users are not connected to the ANs or their locality information is not available. The probability that a user's information is not known is $\epsilon(n)$, where $\epsilon(n)\kappa(n) = o(1)$. By the Chernoff bound, we have that,

$$P(E_8^c) = o\left(\frac{1}{n}\right).$$

The rest of this proof is conditioned of E_8 for time-slots $t - L$ and $t - L + 1$. The number of packets that arrived at the BS at beginning of time-slot $t - L - 1$ for users that were not connected to the ANs in either time-slot $t - L$, or $t - L + 1$, or both is $2n\sigma$.

Now consider the packets that arrived at the BS at beginning of time-slot $t - L - 1$ for users that were connected to the ANs in both time-slots $t - L$ and $t - L + 1$. Conditioned on the event E_1 defined in Lemma 2, the number of packets "lost" in time-slots $t - L$ is less than $n\sigma$. If a packets is not lost in time-slot $t - L$, it either reaches the intended user by the end of time-slot $t - L$ or an AN connected to the user in time-slot $t - L + 1$.

Consider the set of packets for users which are received by at least 1 AN connected to the corresponding user in time-slot $t - L + 1$. Since at most σn packets are lost in time-slot $t - L + 1$, all but at most σn of such packets reach the corresponding users by the end of time-slot $t - L + 1$.

From this, we conclude that, for users that are connected to ANs in time-slot $t - L$ and $t - L + 1$, all but $2\sigma n$ of the arrivals at the BS in time-slot $t - L - 1$ are received by the users by the end of time-slot $t - L + 1$. Therefore, conditioned on E_1 and E_8 for time-slots $t - L$ and $t - L + 1$, we have that,

$$P\left(|F(t)|_1 \leq 4n\sigma\right) \leq o\left(\frac{1}{n}\right).$$

□

Lemma 5: Let $S_i(t) = \sum_{j=1}^{n} X_{i,j}(t)Y_{i,j}(t)$ be the service allocated to queue i at the base-station by the DIST algorithm in time-slot t. Dropping the time index for simplicity, let E_9 be the event that

$$\cap_i \{F_i \leq S_i\} \cap \{S_{i^*} \geq F_{i^*} + 1\},$$

where $i^* \in \arg\max_i Q_i(t-1)$. The event E_9 implies that all the new arrivals and the feedback arrivals to the base-station queues at the beginning of slot t are served in slot t and at least one of the longest queues is served by at least 1 additional channel. Then,

$$P(E_9^c) = o\left(\frac{1}{n}\right).$$

Proof: The proof is conditioned on the event E_7. Let $q_{min}^{(d)} = \min_{i,j} P(X_{i,j}(t) = C_{max}^d)$. Pick any δ in

$$\left(0, \frac{q_{min}^{(d)}\sigma}{2\kappa(n)(2-q_{min}^{(d)})}\right).$$

Let F_r be the set of queues which received r new packets at the beginning of slot t, and $F_r^{(k)}$ be the set of queues which have r new packets after k arounds of channel allocation. We know that $|F_r| = 0$ for $r > \kappa(n)$. Let $r = \kappa(n)$.

Case I: $|F_r| = |F_r^{(0)}| \geq \delta n$.

Define $w_0 = |F_r^{(0)}| - \delta n$. By Assumption 1, we have that after the first w_0 rounds of service, $|F_r^{(w_0)}| \leq \delta n$ w.p. $\geq 1 - \delta n o(1/n^3)$.

Consider the next $v_0 = \dfrac{2\delta n}{q_{min}^{(d)}}$ rounds of allocation, By Assumption 1, we have that $|F_r^{(v_0+w_0)}| = 0$ w.p. $\geq 1 - o(1/n^3)$.

Case II: $|F_r| = |F_r^{(0)}| \leq \delta n$.

Consider the first $v_0 = \dfrac{2\delta n}{q_{min}^{(d)}}$ rounds of allocation, By Assumption 1, we have that $|F_m^{(v_0)}| = 0$ w.p. $\geq 1 - o(1/n^3)$.

Conditioned on E_7, there are $(1-\lambda-\sigma)n$ channels unused by the BS-AN links. The proof now follows by repeatedly applying the above procedure for $r = \kappa(n), \kappa(n)-1, ...1$. As a result, all new feedback arrival packets are served at the end of the next $5\sigma n$ rounds of allocation with probability

$$\geq 1 - \kappa(n)\left(\delta n o\left(\frac{1}{n^3}\right) + o\left(\frac{1}{n^3}\right)\right).$$

In the remaining σn rounds of allocation, by Assumption 1, at least one channel serves the longest BS queue with probability $= o(1/n^3)$. Therefore, $P(E_9^c) = o\left(\dfrac{1}{n}\right)$. □

Theorem 6: Under Assumption 1, for a given load $\lambda < 1$, the system of queue-lengths at the base-station is stabilized by the DIST algorithm for $n > n_0$ which is a function of λ. Here, by stability we mean that the Markov chain corresponding to the queue-lengths at the base-station is positive recurrent.

Proof: Consider the Lyapunov function $V(t)$ where $V(\mathbf{Q}(t), \mathbf{A}(t)) = ||Q(t)||^2$. We drop the time index for convenience.

$$\begin{aligned} E[V(t+1) - V(t)|\mathbf{Q}(t)] &= ||Q(t+1)||^2 - ||Q(t)||^2 \\ &= ||Q+F-S+U||^2 - ||Q||^2 \\ &= ||Q||^2 + ||(F-S)||^2 \\ &\quad + 2\langle Q, F-S\rangle + ||U||^2 \\ &\quad + 2\langle U, (Q+F-S)\rangle - ||Q||^2 \\ &\leq n^3 C_{max}^2 + 2\langle Q, (F-S)\rangle. \end{aligned}$$

We use the fact that $U = (Q+F-S)^+ - (Q+F-S)$, therefore $\langle U, (Q+F-S)\rangle = -||U||^2 \leq 0$.

For the DIST algorithm and the event E_9 defined above, $P(E_9^c) = o(1/n)$. By the definition of event E_9, we have that

$$E[\langle Q, F-S\rangle|\mathbf{Q}(t), E_9] \leq -Q_{max}.$$

Also,

$$E[\langle Q, F-S\rangle|\mathbf{Q}(t), E_9^c] \leq Q_{max} C_{max} n.$$

Therefore,

$$\begin{aligned} E[V(t+1) - V(t)|\mathbf{Q}(t)] &\leq n^3 C_{max}^2 + 2\langle Q, (F-S)\rangle. \\ &\leq n^3 C_{max}^2 - 2Q_{max}P(E_9) \\ &\quad + 2Q_{max}C_{max}nP(E_9^c) \\ &\leq n^3 C_{max}^2 - Q_{max}P(E_9), \end{aligned}$$

for n large enough. For $Q_{max} > \frac{n^3 C_{max}^2 - 1/2}{P(E_9)}$, the drift is $\leq -\frac{1}{2}$. Therefore, by Foster's theorem, the queues are stabilized by the DIST algorithm. □

B. Performance

Theorem 7: For a mobile user (as discussed in Section II), for any SUT algorithm, for a packet that is routed to an AN by the base-station,

$$d := \limsup_{n\to\infty} \frac{-1}{n} \log P(\mathit{Delay} > r) = 0,$$

for any $r < \infty$.

Proof: Let a packet for a mobile user u be sent to AN m in time-slot t. Let E be the event that the user is not connected to AN m in the next r time-slots.

$$P(E) \geq (\min\{\mu_1, \mu_2\})^r.$$

Conditioned on E, the packets cannot reach the user before time-slot $t+r$. Hence the result follows. □

Lemma 6: Recall that $U(t)$ is the set of users whose location is known at the beginning of time-slot t. Consider the set of all packets for users $u \in U(t)$ which are received and queued at least one AN at the beginning of time-slot t. As defined in Lemma 2, a packet for user u in this set is said to be "lost" in time-slot t if by the end of time-slot t, it is neither received by the user u nor at least one AN connected to user u in time-slot $t+1$. Let $\gamma > 0$ be a constant. Let G_1 be the event that in a time-slot, at most γn are "lost". Under Assumption 7, for the DIST-AN algorithm,

$$\begin{aligned} P(G_1^c) \leq o(e^{-n}) + nM(n)\exp\left(-nH\left(\frac{q_2^{(C)}}{2}\middle| q_2^{(C)}\right)\right) \\ +(M(n))^2\exp\left(-nH\left(\frac{q_3^{(C)}}{2}\middle| q_3^{(C)}\right)\right) + e^{-bn}. \end{aligned}$$

Proof: Consider the event G_2 that for any (AN, user) pair, in a time-slot, the number of channels that have channel rate C is at least $\dfrac{nq_2^{(C)}}{2}$. Since channels are i.i.d. across users and ANs,

$$P(G_2^c) \leq nM(n)\exp\left(-nH\left(\frac{q_2^{(C)}}{2}|q_2^{(C)}\right)\right).$$

The rest of this proof is conditioned on G_2 for all (AN, user) pairs. Consider the event G_3 that for any AN in the communication radius of an AN, in a time-slot, the number of channels that have channel rate C is at least $\frac{nq_3^{(C)}}{2}$. Since channels are i.i.d. across users and ANs,

$$P(G_3^c) \leq (M(n))^2 \exp\left(-nH\left(\frac{q_3^{(C)}}{2}\middle| q_3^{(C)}\right)\right).$$

The rest of this proof is conditioned on G_3 for all such (AN, AN) pairs. Let G_4 be the event that not more than n^ν users are connected to any 1 AN where ν is as defined in Section II. By the assumptions in Section II,

$$P(G_4^c) \leq e^{-bn},$$

for some $b > 0$. The rest of this proof is conditioned on G_4.

The arrival process at the ANs is bounded by K and each AN can communicate with at most $O(\log n)$ other ANs. Therefore, every AN has at most $K_1K(L+1)n^\nu \log n$ packets to send in a time-slot and can use at most $K_1K(L+1)n^\nu \log n$ channels to do so. So, for each AN, in any round of channel allocation, there are at least $\frac{nq_2^{(C)}}{2} - K_1K(L+1)n^\nu \log n$ channels that have channel rate C for every connected user and neighboring AN.

Let G_5 be the event that at most γn packets are lost. Then from Lemma 2, we have that for n large enough,

$$\begin{aligned} P(G_5^c|G_2, G_3, G_4) &\leq \binom{n(d+1)C_{max}}{n\gamma} \\ &\quad \times \left(\frac{K_1Kn^\beta n^\nu \log n}{\frac{\eta n q^{(C)}}{2} - K_1Kn^\nu \log n}\right)^{\gamma n/2} \\ &\leq 2^{nH(\gamma|(d+1)C_{max})} \\ &\quad \times \left(\frac{K_1Kn^\beta n^\nu \log n}{\frac{\eta n q^{(C)}}{2} - K_1Kn^\nu \log n}\right)^{\gamma n/2} \\ &\leq 2^{nH(\gamma|(d+1)C_{max})} C_1 e^{-(1-\beta-\nu)\gamma n \log n} \\ &= o(e^{-n}). \end{aligned}$$

Therefore,

$$\begin{aligned} P(G_1^c) \leq o(e^{-n}) &+ nM(n) \exp\left(-nH\left(\frac{q_2^{(C)}}{2} | q_2^{(C)}\right)\right) \\ &+ (M(n))^2 \exp\left(-nH\left(\frac{q_3^{(C)}}{2} | q_3^{(C)}\right)\right) + e^{-bn}. \end{aligned}$$

□

Lemma 7: Fix

$$\gamma \in \frac{1}{7}\left(0, \min\left(\frac{p_0}{K}, \frac{1 - \sum_{k=1}^K p_k \lceil \frac{k}{C} \rceil}{\lceil \frac{K}{C} \rceil C}\right)\right).$$

and let M be the set difference between the probability simplex in K dimensions and an γ ball around the probability vector p. Let

$$\tau := \inf_M \sum_{k=0}^K z_k log \frac{z_k}{p_k}.$$

Let G_6 be the event that all arrivals to the base-station at the beginning to time-slot t are sent to the ANs by the in $(\lambda + 2\gamma)n$ rounds of channel allocation. Then, for a positive constant $\rho < 1$,

$$P(G_6^c) \leq e^{-n\tau(1-\rho)} + exp(-n\gamma^2/2).$$

Proof: By Sanov's theorem, we have that the load on the system in time-slot $t \leq (\lambda + \gamma)$ with probability $\geq e^{-n\tau(1-\rho)}$ for any $\rho < 1$. We condition the rest of the proof of this event. By Assumption 7, the probability that channel k cannot be used at rate C to serve the longest queue at the base-station updates after $k-1$ rounds of allocation is $\geq (1-q_1^{(C)})^{2\log n} = o(1)$. Define $\epsilon'(n) = (1 - q_1^{(C)})^{2\log n}$

$$P(G_6^c) \leq \exp(-nH(\gamma|\epsilon'(n))),$$

Using Pinsker's inequality [9], for n large enough, we have that, $P(G_6^c) \leq \exp\left(-n\frac{\gamma^2}{2}\right) + e^{-n\tau(1-\rho)}$. □

Lemma 8: For any t,

1) $\max_{1 \leq u \leq n} F_u(t) \leq K$
2)

$$\begin{aligned} P(|F(t)|_1 > 4n\gamma) &\\ \leq 2 \exp&\left(-n\frac{\gamma^2}{2K^2}\right) \\ &-o(e^{-n}) + 2e^{-bn} \\ &-2nM(n) \exp\left(-nH\left(\frac{q_2^{(C)}}{2}\middle| q_2^{(C)}\right)\right) \\ &-2(M(n))^2 \exp\left(-nH\left(\frac{q_3^{(C)}}{2}\middle| q_3^{(C)}\right)\right) \\ &-2 \exp\left(-n\frac{\gamma^2}{2}\right) \\ &-2e^{-n\tau(1-\rho)}. \end{aligned}$$

where γ is as chosen in Lemma 7.

Proof: The first part of the lemma follows from the fact that the arrival process for a user u at the ANs is bounded by K and therefore, the number of packets for a user u that cannot be served by the ANs in a given time-slot is $\leq K$.

This proof is conditioned on the event G_1 defined in Lemma 6 and G_6 defined in Lemma 7 for time-slots $t - L$ and $t - L + 1$ and follows on the same lines as that of the proof of Lemma 4.

Conditioned of G_1 and G_6, we need to compute the probability of the event that not more than $\frac{\gamma}{K}n$ users each are not connected to the ANs or their locality information is not available in time-slots $t-1$ and t.

Let G_7 be the event that more than γn users' location information is not known in a given time-slot. The probability of this event is $\epsilon(n)$, i.i.d. across users and time-slots. Therefore,

$$P(G_7|G_1, G_6) = 2 \exp(-nH(\gamma|\epsilon(n))).$$

Using Pinsker's inequality [9], for n large enough, we have that,

$$\begin{aligned} P(|F(t)|_1 \leq 4n\gamma) \geq 1 - \exp&\left(-n\frac{\gamma^2}{2K^2}\right) \\ &-P(G_1^c) - P(G_6^c). \end{aligned}$$

Substituting the value of $P(G_1^c)$ and $P(G_6^c)$, the result follows. □

Lemma 9: Let $\xi(t) := \max_{1\leq i\leq n} Q_i(t)$ be the maximum queue-length at the end of time-slot t. Fix a constant

$$\delta \in \left(0, \frac{q_4\gamma}{K(2-q_4)}\right).$$

Then,

$$\begin{aligned} P(\xi(t) > \xi(t-1)) \leq \exp\left(-n\frac{\gamma^2}{2K^2}\right) \\ +P(G_3^c) + P(G_4^c) \\ +Kn(1-q_4)^{n\delta} \\ +Kn\delta\exp\left(\frac{2n\delta}{q_4}H\left(\frac{q_4}{2}\middle|q_4\right)\right). \end{aligned}$$

Proof: Conditioned on the properties of the feedback arrivals from Lemma 8 and ([7, Lemma 4]), we have that all feedback arrivals are forwarded directly to the users by the end of n rounds of channel allocation with probability $\geq Kn(1-q_4)^{n\delta} + Kn\delta\exp\left(\frac{2n\delta}{q_4}H\left(\frac{q_4}{2}\middle|q_4\right)\right)$. This completes the proof. □

Lemma 10: Let $\xi(t) := \max_{1\leq i\leq n} Q_i(t)$ be the maximum queue-length at the end of time-slot t. There exists a constant k_0, such that

$$P(\xi(t) < \xi(t-k_0)) \geq 1/2.$$

Proof: This result follows from Lemma 9 and [7, Lemma 8]. □

Theorem 8: Under Assumption 7 and γ, τ, ρ as defined in Lemma 7 and δ defined in Lemma 9, for the DIST algorithm, for any integer $r \geq 0$,

$$\begin{aligned} c^{(DIST)} &= \liminf_{n\to\infty} \frac{-1}{n} \log P\left(\max_{1\leq i\leq n, 1\leq r\leq Rn} Q_i(0) > v\right) \\ &= \frac{r+1}{K}\min\left(\frac{\gamma^2}{2K^2}\log\gamma, \tau(1-\rho),\right. \\ &\quad H\left(\frac{q_2^{(C)}}{2}\middle|q_2^{(C)}\right), b, n\delta\log\frac{1}{1-q_4}, \\ &\quad \left. H\left(\frac{q_3^{(C)}}{2}\middle|q_3^{(C)}\right), H\left(\frac{q_4}{2}\middle|q_4\right)\right). \end{aligned}$$

Proof: The proof follows from Lemma 9, Lemma 10 and [7, Theorem 5]. □

References

[1] J. G. Andrews, "Seven ways that hetnets are a cellular paradigm shift," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 136–144, Mar. 2013.

[2] J. G. Andrews, H. Claussen, M. Dohler, S. Rangan, and M. C. Reed, "Femtocells: Past, present, and future," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 3, pp. 497–508, Apr. 2012.

[3] A. Barbieri *et al.*, "Coordinated downlink multi-point communications in heterogeneous cellular networks," in *Proc. IEEE Inf. Theory Appl. Workshop (ITA)*, Feb. 2012, pp. 7–16.

[4] P. Billingsley, *Probability and Measure*. New York, NY, USA: Wiley, 1995.

[5] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant, "Scheduling in multi-channel wireless networks: Rate function optimality in the small-buffer regime," in *Proc. 11th SIGMETRICS/Perform. Conf.*, 2009, pp. 121–132.

[6] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant, "Low-complexity scheduling algorithms for multi-channel downlink wireless networks," in *Proc. IEEE INFOCOM*, Mar. 2010, pp. 1–9.

[7] S. Bodas, S. Shakkottai, L. Ying, and R. Srikant, "Scheduling for small delay in multi-rate multi-channel wireless networks," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 1251–1259.

[8] W. C. Cheung, T. Q. S. Quek, and M. Kountouris, "Throughput optimization, spectrum allocation, and access control in two-tier femtocell networks," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 3, pp. 561–574, Apr. 2012.

[9] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley, 2012.

[10] A. Eryilmaz, R. Srikant, and J. R. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Netw.*, vol. 13, no. 2, pp. 411–424, Apr. 2005.

[11] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource allocation and cross-layer control in wireless networks," *Found. Trends Netw.*, vol. 1, no. 1, pp. 1–144, 2006.

[12] D. Gesbert *et al.*, "Multi-cell MIMO cooperative networks: A new look at interference," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1380–1408, Dec. 2010.

[13] A. Gopalan, C. Caramanis, and S. Shakkottai, "On the value of coordination and delayed queue information in multicellular scheduling," *IEEE Trans. Autom. Control*, vol. 58, no. 6, pp. 1443–1456, Jun. 2013.

[14] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.

[15] *Dense Networks*, accessed on Oct. 2014. [Online]. Available: http://www.ericsson.com/news/120223_it_comes_back_to_backhaul_244159020_c

[16] R. Irmer *et al.*, "Coordinated multipoint: Concepts, performance, and field trial results," *IEEE Commun. Mag.*, vol. 49, no. 2, pp. 102–111, Feb. 2011.

[17] B. Ji, G. R. Gupta, X. Lin, and N. B. Shroff, "Performance of low-complexity greedy scheduling policies in multi-channel wireless networks: Optimal throughput and near-optimal delay," in *Proc. IEEE INFOCOM*, Apr. 2013, pp. 2589–2597.

[18] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*. New York, NY, USA: Springer, 1996, pp. 153–181.

[19] S. Katti, S. Gollakota, and D. Katabi, "Embracing wireless interference: Analog network coding," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, pp. 397–408, Aug. 2007.

[20] F. Khan, Z. Pi, and S. Rajagopal, "Millimeter-wave mobile broadband with large scale spatial processing for 5G mobile communication," in *Proc. IEEE 50th Annu. Conf. Commun., Control Comput. (Allerton)*, Oct. 2012, pp. 1517–1523.

[21] J. N. Laneman and G. W. Wornell, "Distributed space-time-coded protocols for exploiting cooperative diversity in wireless networks," *IEEE Trans. Inf. Theory*, vol. 49, no. 10, pp. 2415–2425, Oct. 2003.

[22] X. Lin, R. K. Ganti, P. J. Fleming, and J. G. Andrews, "Towards understanding the fundamentals of mobility in cellular networks," *IEEE Trans. Wireless Commun.*, vol. 12, no. 4, pp. 1686–1698, Apr. 2013.

[23] A. Lozano, R. W. Heath, Jr., and J. G. Andrews. (2012). "Fundamental limits of cooperation." [Online]. Available: https://arxiv.org/abs/1204.0011

[24] S. Moharir, S. Krishnasamy, and S. Shakkottai, "Scheduling in densified networks: Algorithms and performance," in *Proc. 52nd Annu. Conf. Commun., Control Comput. (Allerton)*, Sep./Oct. 2014, pp. 683–690.

[25] S. Moharir and S. Shakkottai, "Maxweight vs. backpressure: Routing and scheduling in multi-channel relay networks," in *Proc. IEEE INFOCOM*, Turin, Italy, Apr. 2013, pp. 1537–1545.

[26] A. Nosratinia, T. E. Hunter, and A. Hedayat, "Cooperative communication in wireless networks," *IEEE Commun. Mag.*, vol. 42, no. 10, pp. 74–80, Oct. 2004.

[27] A. Ozgur, O. Leveque, and D. N. C. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," *IEEE Trans. Inf. Theory*, vol. 53, no. 10, pp. 3549–3572, Oct. 2007.

[28] Z. Pi and F. Khan, "An introduction to millimeter-wave mobile broadband systems," *IEEE Commun. Mag.*, vol. 49, no. 6, pp. 101–107, Jun. 2011.

[29] S. Rajagopal, S. Abu-Surra, Z. Pi, and F. Khan, "Antenna array design for multi-Gbps mmWave mobile broadband communication," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Dec. 2011, pp. 1–6.

[30] I. Rhee *et al.*, "On the levy-walk nature of human mobility," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 630–643, Jun. 2011.

[31] M. Sharma and X. Lin, “OFDM downlink scheduling for delay-optimality: Many-channel many-source asymptotics with general arrival processes,” in *Proc. ITA*, Feb. 2011, pp. 1–10.
[32] *LTE Small Cells*, accessed on Oct. 2014. [Online]. Available: http://en.wikipedia.org/wiki/Small_cell
[33] 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.
[34] A. J. Viterbi, A. M. Viterbi, K. S. Gilhousen, and E. Zehavi, “Soft handoff extends CDMA cell coverage and increases reverse link capacity,” *IEEE J. Sel. Areas Commun.*, vol. 12, no. 8, pp. 1281–1288, Oct. 1994.
[35] D. Wong and T. J. Lim, “Soft handoffs in CDMA mobile systems,” *IEEE Pers. Commun.*, vol. 4, no. 6, pp. 6–17, Dec. 1997.
[36] W. Zhuang and M. Ismail, “Cooperation in wireless communication networks,” *IEEE Wireless Commun.*, vol. 19, no. 2, pp. 10–20, Apr. 2012.

Sharayu Moharir received the B.Tech. degree in electrical engineering from IIT Bombay, Mumbai, India, in 2009, the M.Tech. degree in communication and signal processing, and the Ph.D. degree in electrical and computer engineering from The University of Texas at Austin. She is an Associate Professor with the Department of Electrical Engineering, IIT Bombay.

Her research interests include algorithms and performance analysis for wireless networks and content delivery networks.

Subhashini Krishnasamy received the B.Tech. degree in electronics and communication engineering from Jawaharlal Nehru Technological University at Hyderabad, Hyderabad, India, in 2009, and the M.E. degree in telecommunications from the Indian Institute of Science, Bengaluru, India, in 2011. She is currently pursuing the Ph.D. degree with The University of Texas at Austin, Austin, TX, USA. She was with Texas Instruments as a Systems Engineer from 2011 to 2012. She was with Alcatel Lucent Bell Labs as an Intern in 2014. Her current research interests include modeling and algorithm design for networks, especially of large-scale.

Sanjay Shakkottai (M’02–SM’11–F’14) received the Ph.D. degree from the Electronics and Communication Department, University of Illinois at Urbana–Champaign, in 2002. He was with The University of Texas at Austin, where he is currently a Professor with the Department of Electrical and Computer Engineering. His current research interests include network architectures, algorithms and performance analysis for wireless networks, and learning and inference over social networks. He received the NSF CAREER Award in 2004.