BP-MR: Backpressure Routing for the Heterogeneous Multi-Radio Backhaul of Small Cells

Jorge Baranda   José Núñez-Martínez   Josep Mangues-Bafalluy

Centre Tecnològic de Telecomunicacions de Catalunya (CTTC)
Av. Carl Friedrich Gauss, 7
08860 Castelldefels (Barcelona), Spain
e-mail: [jorge.baranda, jose.nunez, josep.mangues]@cttc.cat

Abstract—Dense small cell (SC) deployments are expected to help handling the explosive growth of mobile data usage. However, the backhaul of these deployments will face several challenges where point-to-point (PTP) and point-to-multipoint (PMP) wireless technologies will be combined forming multipoint-to-multipoint (MP2MP) wireless mesh backhauls. In this context, routing and load balancing solutions will be of key importance to maximize the use of wireless backhaul resources. This paper presents Backpressure for Multi-Radio (BP-MR), a distributed routing and load balancing protocol specifically designed for heterogeneous MP2MP wireless mesh backhauls. The backhaul is heterogeneous in the sense that each node may embed a different number of diverse wireless interfaces. BP-MR introduces a two-stage routing process to appropriately handle Head-of-Line blocking issues that appear in such multi-radio environments. We validate these improvements with ns-3 simulations under different network conditions. As a consequence of an improved wireless link usage efficiency, results show improvements in throughput of up to 34% and in latency of up to one order of magnitude with respect to state-of-the-art approaches.

I. INTRODUCTION

In the last years, the wireless industry is experiencing an explosive growth of mobile data usage. According to [1], global mobile data traffic will increase nearly tenfold between 2014 and 2019. In such context, operators are pushing for the deployment of dense small cell (SC) networks, whose goal is to provide added capacity by bringing small base stations closer to the mobile users, hence, making a better re-use of the available spectrum. As mentioned in [2], these deployments will be driven by several challenges, such as fiber availability or installation costs. Hence, it is recognized that a) dense SC deployments will require wireless- in addition to fiber-based backhauls, particularly close to the edge, and b) there is not a single wireless backhaul technology that alone can meet the requirements of dense SC deployments.

For these reasons, we consider SC deployments in which each transport node may include multiple heterogeneous wireless radios. An example of such a setup is depicted in Figure 1, in which a mix of point-to-point (PTP) and point-to-multipoint (PMP) wireless technologies are combined forming a multipoint-to-multipoint (MP2MP)/mesh backhaul.Therefore, short- and long-range wireless backhaul radios will be combined. Specifically, a sub-6GHz backhaul link can connect several close small-cell locations because of its PMP non-line-of-sight (NLOS) nature, offering the highest level of flexibility and ease of installation, which may be


Fig. 1. Example of wireless heterogeneous backhaul deployment of SCs. The aggregation point counts with a fiber connection.

appropriate for street-level deployments. In turn, traffic may be eventually aggregated in rooftops (e.g., where sites were acquired to install macro base stations and transport radio-links). Consequently, line-of-sight (LOS) PTP links may be installed in some of the locations for reaching distant sites with potential access to fiber towards the core network, hence offloading traffic from the underlying PMP mesh backhaul (see Figure 1). Microwave radios can be used for these links because of its maturity and its competitive prices [2] and other technologies, such as millimeter wave, are also growing in importance.

But these setups bring some challenges with them. First, and due to their multi-radio nature, backhaul nodes must appropriately handle head-of-line (HoL) blocking, by which packets that cannot be transmitted through a certain interface (e.g., medium busy) block other packets queued behind them, when, in fact, these other packets could have been transmitted through another free interface. This leads to an inefficient use of wireless backhaul resources, and consequently, to the degradation of the backhaul network performance. And second, to handle such a heterogeneous wireless multi-hop backhaul, network layer intelligence in the form of flexible routing and load balancing is required at the cell site [3]. In this respect, backpressure routing is a promising solution that has been theoretically shown to be throughput optimal [4]. However, practical backpressure implementations have also been shown to experience increased queue complexity with the number of traffic flows and also high latencies especially for low traffic volumes ([5], [6], and [7]).

The main contribution of this paper is to propose Back-

pressure for Multi-Radio (BP-MR), a distributed routing and load balancing protocol specifically designed for heterogeneous multi-interface MP2MP wireless mesh backhauls. BP-MR takes routing decisions following a two-stage routing process. In the first stage, BP-MR groups data packets in queues according to their final destination to alleviate HoL problems. Therefore, BP-MR maintains a FIFO data queue for each backhaul interface. In the second stage, BP-MR actually computes the best possible next-hop according to weight calculations that take into account geographic and congestion information for all possible forwarding options in the multi-radio node that a given packet may have.

Our previous work [8], which was based on the Lyapunov drift-plus penalty theory [9], only tackled single-radio homogeneous wireless backhaul networks, i.e., the same wireless technology was used in all nodes of the network. In contrast, BP-MR works in backhaul deployments where network nodes can accommodate a different number of heterogeneous wireless backhaul radios. This requires HoL blocking management, which is handled through a two-stage routing process.

Simulation results with ns-3 [10] confirm the increased efficiency in wireless resource usage of BP-MR by showing throughput and latency improvements under a wide range of network conditions with respect to state-of-the-art routing approaches targeting similar goals, such as Greedy Perimeter Stateless Routing (GPSR) [11] and Backpressure for Sparse Deployments (BS) [8]. In summary, BP-MR adapts to the specific setup of SCs and their associated transport network node equipped with multiple heterogeneous (PTP and PMP) wireless backhaul radios, while scaling with the number of interfaces. In our simulations, up to five wireless backhaul radios were evaluated. Additionally, BP-MR alleviates performance limitations due to HoL, as the two-stage routing process used to take forwarding decisions significantly improves the use of wireless links built with multi-radio nodes. In particular, BP-MR obtains gains of up to 34% in throughput and one order of magnitude in latency with respect to GPSR and BS.

The remainder of this paper is organized as follows. The related work in Section II is followed by the operational details of BP-MR in Section III. After that, Section IV presents the performance of BP-MR under different network scenarios and conditions before concluding the paper with Section V.

II. RELATED WORK

Tassiulas and Ephremides developed the roots of dynamic backpressure routing for multi-hop wireless networks in [4]. In essence, it is a centralized policy to route traffic in a multi-hop network that attains throughput optimality by minimizing the Lyapunov drift in the network, that is, minimizing the sum of the queue backlogs in the network from one time slot to the following one. Although this protocol showed optimality in terms of throughput, it suffers from high queue complexity (per-flow queuing system) and increased end-to-end latencies derived from the last-packet problem, whereby packets belonging to a flow may get excessively delayed in queues due to the lack of subsequent packet arrivals of the same flow. This happens because decisions are based on sending the packet to the neighbor generating the highest queue backlog differential for a given flow. Therefore, if there is


Fig. 2. Last-Packet problem: Flow with higher queue backlog (flow B) is served first because of its highest queue backlog differential, increasing packet latency of flows with lower queue backlog (flow A).

little traffic, no queue backlog differential is generated and packets are not forwarded, as depicted in Figure 2. Several modifications have been proposed in the literature to reduce the effect of these problems ([5], [6], and [7]). Nevertheless, these alternatives, which we will refer to as legacy backpressure, require per-flow or per-destination queue maintenance information, hence higher queue management complexity than our proposed solution.

Closer to our approach, Neely extended the concepts of Tassiulas et al. and defined the Lyapunov-drift-plus penalty approach to optimize the routing problem in wireless multihop networks [9]. Based on this work, the authors of [12] presented a practical backpressure implementation in the context of wireless sensor networks, which presents very different requirements from those of a wireless backhaul. More specifically, this implementation only considers many-to-one traffic communications and uses a last-in first-out (LIFO) queue per node to reduce the end-to-end latency.

In [8], we have presented Backpressure for Sparse Deployments (BS), a simple yet efficient distributed backpressure routing scheme based on the Lyapunov-drift-plus penalty weight of [9]. This weight combines queue backlog information (backpressure component) with geographic information, and the relative importance of each component can be adjusted through a knob dynamically adjusted to find the best trade-off. It is relevant to note the importance of the geographic component in our approach because of the scalability and low overhead properties that geographic routing presents. Such properties are desirable in a high-scale wireless mesh backhaul. In this sense, GPSR [13] is a reference geographic routing protocol which proposes two modes of operation to forward packets: greedy and perimeter mode. The greedy mode is based on forwarding packets to the node that minimizes the Euclidean distance to the destination. When a local minimum is found, GPSR enters in recovery mode, which performs routing operations based on the right-hand rule.
As in [12], BS uses a single queue, in this case following a first-in first-out (FIFO) scheme, to handle any-to-any traffic communications under arbitrary mesh backhaul deployments of network nodes equipped with a single PMP backhaul interface to communicate with its surrounding neighbors. In brief, the aim of the proposed Lyapunov-drift-plus penalty scheme is to route packets trying to find the appropriate balance between the shortest and the less congested path, while avoiding routing loops in the presence of network voids.

With respect to legacy backpressure algorithms, the BS

approach is more scalable, given its single queue for all flows and destinations, hence presenting lower queue management complexity. Moreover, it is free from the last-packet problem, since packets from all flows share the queue, which more easily generates queue backlog differentials, and the geographic component makes it act as shortest path when there is low load. Therefore, it increases the fairness amongst packets belonging to different traffic flows. However, the performance of BS decreases when nodes are equipped with multiple interfaces (as will be shown in Section IV). In such scenario, BS suffers from the well-known head-of-line (HoL) blocking effect, hence leading to loss of possible transmission opportunities.

III. BP-MR: Backpressure Routing for Multi Radio

This section presents the operation of BP-MR to perform per-packet forwarding decisions when the transport node associated with a given SC is equipped with multiple heterogeneous technologies. The process is divided in two stages. In the first one, ingress packets are distributed in the different per-interface queues maintained by BP-MR. In the second one, the next-hop is determined for the packets present at the head of the queues using the Lyapunov drift-plus-penalty approach. And so, decisions are independently taken for each packet depending on network conditions and without building and maintaining end-to-end routes, which entails a lower control overhead compared to conventional mesh routing protocols (e.g., AODV, and OLSR), and similar one compared to the schemes described in Section II.

A. Stage 1: Per-interface queue management system

The negative effects of HoL blocking are decreased when the different backhaul interfaces at a SC can be used concurrently. To achieve this, BP-MR maintains a FIFO data queue per each backhaul interface. Although the queue management complexity increases compared to the single-queue scheme of [8], it is more scalable than the legacy backpressure solutions that maintain per-flow queues, as presented in Section II.

The distribution of ingress packets to each of the available queues plays a key role to alleviate the effects of the aforementioned HoL blocking problem. When a packet is received through any backhaul interface, BP-MR delivers the packet to the upper layers of the local node (e.g., the 3GPP protocol stack of the SC) if it is the destination, or places it into a given queue associated with a backhaul interface. According to the proposed SC deployment, each backhaul interface is associated with a set of neighbors. In particular, there is one neighbor for PTP backhaul interfaces and potentially several neighbors for PMP backhaul interfaces.

The packet is placed in the queue attached to the interface whose neighbors are geographically closer to the destination in terms of euclidean distance. The aim of this criteria is to limit the number of hops needed to reach the intended destination. Note that this is a default routing policy that will be later reconsidered in next stage in the case that the SC faces network congestion. To break ties, in the case there is more than one interface whose neighbors are closer to the intended destination, BP-MR enqueues the packet in the less loaded interface, i.e., the one that minimizes the ratio between the queue backlog and the link rate of the radio interface. As there are several alternatives to connect to a certain destination, BP-MR selects the one that is likely to forward the packet faster, trying to reduce the time a packet spends in the queue before being transmitted. We assume that GPS (or any other means to build virtual coordinates) is available at each SC to facilitate geographic information. Each SC/node periodically sends HELLO packets through each of its interfaces, which carry its queue backlog at the moment such HELLO packet is sent and geographic location to its neighbors. In this way, neighbors have an updated view of congestion and geolocation information that is used to take per-packet forwarding decisions.

Aggregating packets headed towards similar geolocated regions in the same queue helps alleviating the HoL blocking effect, given that similar next-hop resolutions are expected for these packets. The aim of this process is to continuously feed each interface with packets that are likely to experience similar next-hop resolutions due to the geographical proximity of their respective destinations. In this way, there is a more efficient use of per-link transmission opportunities and the impact of additional latency on queued packets produced by a packet in the head of the queue waiting to be served by the suitable interface is reduced. Note that this stage decides a possible distribution of packets among the available backhaul interfaces, but the final routing decision is taken in stage 2.

B. Stage 2: Next-hop computation

The actual packet forwarding decision is performed on a per-packet basis. At each transmission opportunity in a backhaul interface, the packet head of the queue associated to the backhaul interface is scheduled to be forwarded. The next-hop is determined by computing a weight based on information exchanged with neighboring nodes through HELLO packets. As mentioned before, this weight relies on the theoretical framework behind Lyapunov drift-plus penalty described in [9]. In essence, this weight is a function of two main routing components, namely backpressure and geographic routing. The relative importance of each of the components when taking forwarding decisions is dynamically adjusted as a function of network conditions. This is done by introducing a $V(t)$ parameter, whose value tries to find the best trade-off between following the shortest path to the destination and avoiding congested spots.

The key element when taking forwarding decisions is the calculation of drift-plus-penalty weights. The exploitation of multiple heterogeneous interfaces requires a complete redefinition of such weight and the associated procedures to calculate it (e.g., HELLO message content) compared to previous work. The design of this weight is aligned with the per-interface queuing system previously presented. The Lyapunov drift-plus penalty value determines the next hop for a given packet, hence the actual interface through which it is transmitted. This weight (denoted by $w_{i_k j_l}$) combines four components and is calculated for all the neighbors directly connected to each of the interfaces in the local node i according to the following expression:

$$w_{i_k j_l}(t) = (\Delta Q_{i_k j_l}(t) - V_{i_k j}(t) c_{ij}^d(t)) R_{i_k j}, \qquad (1)$$

where i is the local node (i.e., that taking the forwarding

decision), i_k is the interface of the local node i connecting with neighbor j and j_l is an output interface l of neighbor j.

Therefore, this expression is used to calculate the metric that determines all forwarding decisions. In fact, the packet will be forwarded through the i_k interface to neighbor j presenting the highest non-negative value of $w_{i_k j_l}(t)$. In this way, the best possible trade-off between congestion and shortest distance is pursued. If all weights are negative, the packet is kept at node i, meaning that all nodes in the area are congested and it is better to wait for a better transmission instant and not waste wireless resources unnecessarily.

Notice also that a packet initially enqueued at interface i_k of the local node in stage 1 may be transmitted through another interface $i_{k'}$ depending on weight calculations. In this case, the packet is dequeued from the initial queue and sent for transmission to the suitable interface (if it is free/ready), or it is enqueued in the head of the queue associated to the suitable interface so that it is transmitted as soon as the interface is free. Therefore, forwarding decisions are taken just before packet transmission (i.e., packet at the head of the queue). Consequently, they are taken based on the most up-to-date information about surrounding network congestion conditions, according to the Lyapunov-drift-plus penalty weight.

Next, a more detailed description of each of the four components of the weight defined in Equation 1 is provided.
Backpressure Routing: This component of the weight/metric is in charge of minimizing the Lyapunov-drift to attain an even load distribution amongst the nodes in the wireless mesh backhaul. $\Delta Q_{i_k j_l}(t)$ is the queue backlog difference between interface i_k of node i and the output interface j_l of neighboring node j. Q_{i_k} represents the backlog of packets waiting to be served at interface i_k of node i. The queue backlog information is included in the HELLO packets that each node sends to its neighboring nodes. Q_{j_l} denotes the backlog of packets in neighboring node j at output interface j_l. The number of backhaul output interfaces of neighboring node j varies depending on the number of backhaul interfaces installed in node j. As illustrated in Figure 3, a PMP NLOS interface installed in a neighboring node j is always considered as potential output interface for a packet being transmitted from node i to node j. On the other hand, a PTP interface installed in node j is considered as output interface for a packet being transmitted from node i to neighboring node j if this PTP interface is not connecting local node i with neighboring node j. This represents an important difference with respect to our previous work [8], which only considers backhaul nodes equipped with a single WiFi-based backhaul interface.
 Geographic routing: Derived in [8], the penalty function ($c_{ij}^d(t)$) generally rewards decisions that push packets closer to the destination, and penalizes decisions that move the packet away from its destination. But when facing a network void, i.e., no neighbor is closer to destination, this penalty function rewards decisions selecting non-traversed nodes farther from the destination. In this way, it avoids packets getting trapped in data queues (which would increase the experienced latency) and penalizes decisions generating 1-hop loops. A 1-hop loop happens when a packet in a SC is routed back to the SC from which the packet was just received. Note that the penalty function between node i and node j is calculated without using per-interface information. It only leverages geolocation infor-


Fig. 3. Node i considers all the PTP interfaces of neighboring j as possible output interfaces except the one connecting node i and j. The PMP interface is always considered as output interface due to its multipoint capabilities.

mation of the local and neighboring nodes. This information from neighboring nodes is provided in HELLO packets.
The link rate: $\mathbf{R_{i_k,j}}$ is the link rate between interface i_k in local node i and neighboring node j. This term gives priority to links with a higher capacity. Note that this term is calculated on a per-interface basis, since each interface i_k in local node i may have a different data rate.
The V parameter: $V(t)$ is a non-negative function in charge of finding the appropriate trade-off between distributing the load among neighbors (backpressure component) and approaching the destination (geographic component). In BP-MR, we propose to calculate the V value on a per-interface basis. This represents an important difference with respect to our previous work [8], where the V value was the same for all the neighboring nodes as they were connected through the same interface. In [8], all the nodes were contending for accessing the same transmission medium, whereas in BP-MR, the PMP and the PTP interfaces do not share the same medium, as they work in different frequency bands. For this reason, the V parameter in BP-MR is calculated in local node i on a per-interface basis. The per-interface value $V_{i_k j}$ is upper bounded by the queue size limit denoted by Q_{MAX}. As showed by Equation (2), $V_{i_k j}$ depends on the queue backlog $Q_{i_k j}$ in the considered interface i_k of local node i connecting with neighbor j:

$$V_{i_k j}(t) = Q_{MAX} - Q_{i_k j} \qquad (2)$$

An important characteristic of the proposed per-interface V value is that it is calculated solely with local information. Since the neighboring node j can have several interfaces, it is not known a priori which one will be forwarding the packet. In [8], the V parameter was computed by using the maximum queue backlog of the local node i and its different neighbors, hence not using only local information. As neighboring nodes only had a single interface, it was known a priori which interface of the neighboring node was going to be used to forward the packet. Because of the nature of backpressure for distributing traffic in order to minimize the Lyapunov drift in the network, the local queue backlog information is a good approximation of surrounding congestion for this interface.

IV. EVALUATION

A. Methodology

The initial performance evaluation of the BP-MR protocol is done via the ns-3 simulator [10]. To carry backhaul traffic,

Fig. 4. Example of an evaluated 25 nodes regular deployment. In addition to the PMP NLOS links in each node, four nodes count with an additional PTP LOS link connecting to the aggregation point (i.e., SC labeled as SC13), which counts with a total of 5 interfaces.

every node is equipped with a sub-6 GHz PMP NLOS link. Additionally, several nodes are equipped with a PTP LOS link, which interconnects them with the aggregation point, as depicted in Figure 4. This topology has been simulated using an IEEE $802.11a$ interface configured to the same channel at a link rate of 54Mbps as example of PMP NLOS sub-6GHz link, and a PTP net device with a bit rate of 311Mbps was set up to represent a microwave PTP LOS link. The distance between neighboring nodes is of 100 meters. The set of neighbors of a given SC are the nodes within the range of 100 meters and those that are directly connected with the PTP link (if the SC has a microwave link). We use a simple channel model featuring the K-*hop* interference model (with K equal to two) for the PMP links. On the other hand, we assume that PTP links do not interfere with each other due to their orthogonal channel assignment. In particular, the grid mesh backhaul topologies under evaluation features 90% of PMP NLOS links, whereas the remaining 10% are PTP LOS links. In this way, SCs endow either 1 PMP link, or 1 PMP and 1 PTP link, or 1 PMP and several PTP links (up to four PTP links in our simulations). An example of evaluated deployment is depicted in Figure 4. Notice that the central node, which acts as the aggregation point, endows several PTP links.

The aim of this section is to evaluate the capabilities of BP-MR under a diverse set of network conditions. In particular, our evaluation covers different wireless mesh backhaul topologies considering different congestion levels and spatial traffic load variations. In subsection IV-B, we evaluate ten different topologies with twenty-five nodes connected in a regular manner and deployed in a square region of 400m by 400m, where each topology differs in the sets of nodes equipped with PTP LOS links. In subsection IV-C, we evaluate other ten different topologies composed by twenty nodes distributed in the same region, but connected in a non-regular manner, that is, including the presence of different sets of network voids. In this case, we consider only a single setup of PTP links. Regarding variations in the traffic load, we want to highlight that UDP constant bit rate (CBR) flows are considered to evaluate the overall network behavior when approaching saturation. Although using CBR flows, randomness in the spatial traffic distribution is created due to the number of different source-destination pairs evaluated.

To assess the performance of BP-MR, we compare it with the following routing protocols, because out of all previous work, they are the closest ones to our scheme in some of their design decisions.

Greedy Perimeter Stateless Routing (GPSR): a state of the art (SoA) geographic routing protocol [13] that shares with BP-MR the reduced routing control overhead and its suitability for handling network dynamism. The aim of comparing BP-MR with GPSR is to assess the gains offered by BP-MR when the geographic routing protocol also features load balancing capabilities, unlike GPSR. It is worth mentioning that, for the sake of fairness, we adapted the ns-3 model of GPSR provided by the authors of [14] to support the multi-radio nodes. In particular, we extended this implementation so that HELLO packets originated by GPSR are transmitted over all the wireless backhaul interfaces endowed in a SC.

Backpressure for Sparse Deployments (BS) [8]: a backpressure routing protocol designed to deal with network dynamism that merely features a single-queue per SC. We consider BS as the main benchmark to assess the improvements experienced by BP-MR in a multi-interface setup when tackling the HoL blocking problems presented in Section II.

Note that the comparison with routing protocols such as [5], [6], and [7] is omitted in our evaluation. This is because these protocols suffer from scalability issues due to their per-flow or per-destination queue maintenance, which causes them to be unpractical for their use in high-scale wireless mesh backhauls.

Regarding the control overhead of each backhaul routing protocol, BP-MR and BS send HELLO broadcast messages of 110 bytes every 100ms, whereas GPSR sends HELLO messages of 135 bytes every 100ms. In turn, GPSR includes an additional header in the data traffic, adding 50 bytes to the configured packet size (1488 bytes). In terms of queuing architecture at the SC, BS maintains a single FIFO queue of a maximum of 200xL packets, where L denotes the number of interfaces, whereas BP-MR and GPSR maintain a separated FIFO queue of a maximum of 200 packets for each available interface. Table I summarizes the general simulation parameters.

TABLE I. GENERAL SIMULATION PARAMETERS

Topology	Regular and non-regular node deployment in a region of 400m^2
SC Backhaul Size	[20,25] nodes
Backhaul interfaces per SC (L)	1 PMP NLOS and [0,1,4] PTP LOS
PMP link rate	54 Mbps
PTP link rate	311Mbps
Backhaul Traffic	UDP CBR model
Backhaul Routing Protocols	BP-MR, GPSR, and BS
Packet Size	1488 bytes
HELLO rate	10 packets/s
SC buffer capacity	200xL packets
Queuing discipline	FIFO drop-tail

The performance of each protocol is characterized by measuring the throughput and the latency in every simulation. For each of the metrics under study, we use the average value and boxplots to represent their statistical distribution. In particular, the box stretches from the 25^{th} to the 75^{th} percentiles, and the whiskers represent the 5^{th} and the 95^{th} percentiles. Note that for some cases, the boxplots collapse in a single point because the obtained distribution concentrates on the average value.

B. Regular deployments

In the following, we focus on a regular deployment of twenty five nodes forming a 5x5 grid mesh backhaul network. Each of these nodes counts with a PMP wireless links. In addition to this, we consider to include PTP wireless links in five different SCs/nodes. The goal here is to assess how BP-MR behaves regardless of the different position of the PTP LOS wireless links, which may not be feasible at all SC locations [2]. We compare BP-MR with GPSR and BS also considering spatial traffic variations and different workloads in our ns-3 simulations.

In particular, we consider ten different sets of PTP microwave link locations in the 5x5 grid, resulting in ten different topologies. Figure 4 shows an example of one of the ten evaluated topologies. Spatial traffic variations are generated randomly selecting a different set of source-destination pairs. Each of the ten different considered topologies is evaluated with forty different sets of twenty traffic flows of the same intensity. The evaluated intensities are 0.5Mbps, 1Mbps, 1.5Mbps, and 2Mbps, so a total of 10Mbps, 20Mbps, 30Mbps, and 40Mbps of input traffic load is injected in the backhaul network. Each experiment considering the same set of flows and intensity was repeated for each routing protocol under consideration. Thus, the performance of each protocol is evaluated over 1600 different simulations, resulting in a total of 4800 simulations per 5x5 backhaul grid, considering the three routing protocols.

Figure 5 and Figure 6 show the average throughput and the average latency distribution exhibited by BP-MR, BS, and GPSR in the considered ten different 5x5 grid backhaul networks. With the lowest offered load evaluated (i.e., 10Mbps), the three protocols present a similar performance. BP-MR shows substantial gains with respect to GPSR, and BS under more demanding traffic conditions (i.e., twenty flows of 1.5Mbps and 2.0Mbps).

When the input rate is of 20Mbps, Figure 6 shows that GPSR starts suffering from congestion. Average delay increases because packets are queued at nodes as GPSR always routes the traffic through the same path following a greedy forwarding approach. At an input rate of 30 Mbps, simulation traces show that GPSR starts experiencing queue overflows and is not able to deliver all the offered load, as Figure 5 shows. In contrast, BS experiences congestion problems according to the latency values presented in Figure 6. With BS, packets remain more time than needed at data queues as a consequence of having only a single FIFO queue in SCs equipped with multiple backhaul radios. However, the offered workload is mostly served. In contrast, BP-MR is able to serve all 30 Mbps while not suffering from latency degradation using efficiently the multiple interfaces and leveraging alternative paths to forward packets, hence relieving the network congestion.

At an input rate of 40Mbps, BS experiences a noticeable degradation in its performance, both in terms of achieved throughput and latency. The packet accumulation at node queues is stressed due to the capacity misuse observed with BS. The HoL blocking effect produces the loss of transmission opportunities in multi-interface nodes leading also to queue overflows. Under such conditions, the value of the V parameter decreases its value to nearly zero in most nodes. Thus, the backpressure routing component, in charge of minimizing the


Fig. 7. Hop cumulative distribution of the experiments run in the topology showed in Figure 4 at an input rate of 40Mbps.

Lyapunov drift, gains influence when taking routing decisions at the expense of the geographic routing component. Therefore, traffic will be distributed through all the network aiming at load balancing, but not necessarily ensuring that packets head to the destination until they reach a less congested spot. Hence, BS cannot manage the offered load efficiently, negatively impacting the latency and the throughput performance of BS.

On the other hand, BP-MR is able to handle almost all the input traffic, showing improvements of up to 34% in throughput with respect to BS and GPSR, while maintaining a latency one order of magnitude lower than GPSR and BS (notice the logarithmic scale). The packet distribution process in per-interface queues alleviates the negative effects of the HoL blocking problem. Then, the per-interface weighting metric computation is able to deliver appropriate forwarding decisions to make an efficient use of the multiple available paths, but not increasing noticeably the stretch of such paths, as depicted in Figure 7.

Figure 7 shows the cumulative distribution of the number of hops traversed by data packets over the backhaul topology depicted in Figure 4 with an input rate of 40Mbps. In a regular deployment, like the ones considered in this subsection, the greedy forwarding approach followed by GPSR is equivalent to the shortest path. In this sense, we can see how BP-MR uses slightly longer paths than GPSR, but achieves more throughput (see Figure 5) and reduced latencies (see Figure 6). This confirms the efficiency of BP-MR load balancing capabilities to find alternative paths to route traffic with respect to GPSR. On the other side, Figure 7 confirms the behavior described previously for BS. We can observe that only a slight percentage of data packets (around a 5%) experience a high number of hops to reach their destinations. This increase in terms of number of hops is because of congestion. Under congestion, a slight percentage of packets suffer from the fact that backpressure prioritizes routing decisions based on the minimization of the Lyapunov drift (minimization of queue backlog differentials) causing an excessive number of hops to reach the destination. However, in general (for most of the traffic) these results reveal that the main component causing higher latencies in BS with respect to BP-MR is in fact the inefficient use of wireless backhaul interfaces due to the HoL blocking effect.

C. Non-regular deployments

In the following, we evaluate how BP-MR adapts to randomly generated non-regular topologies while exploiting
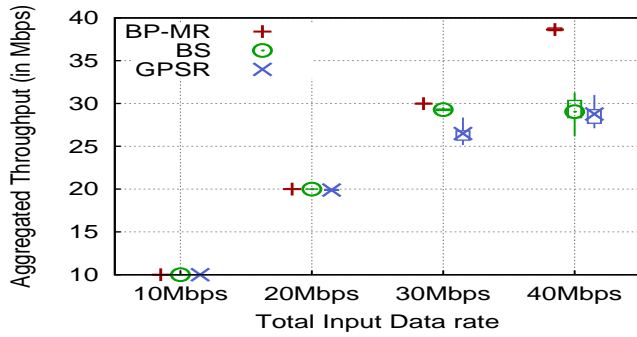
Fig. 5.   Throughput performance comparison in the regular deployment.


Fig. 6.   Latency performance comparison in the regular deployment.

Fig. 9.   Throughput performance comparison in the non-regular deployment.
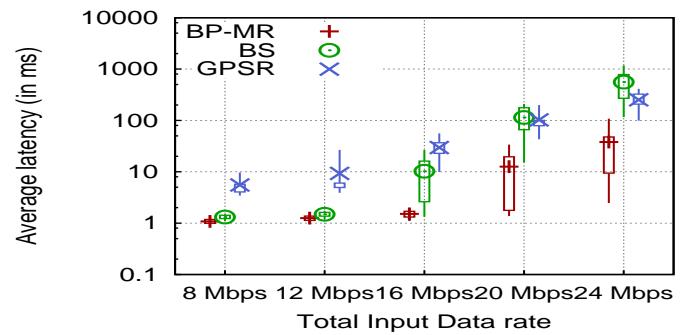

Fig. 10.   Latency performance comparison in the non-regular deployment.


Fig. 8.   Example of an evaluated 20 nodes non-regular wireless mesh backhaul deployment.

the added capacity brought by nodes equipped with multiple backhaul radios. We simulate ten different non-regular and random SC backhaul topologies composed of twenty nodes similar to that in Figure 8. The resulting topology does not include any isolated node and all nodes include a PMP wireless link and up to five nodes include a PTP wireless link. With respect to the previous subsection, we keep fixed the set of nodes with the PTP LOS link. However, in some of the evaluated scenarios, not all the five nodes counting with the PTP link are available. The offered load consists of a different number of 2Mbps flows from the set {4,6,8,10,12}, resulting in a maximum offered load of 24Mbps. For each offered load and backhaul topology, we conducted forty simulations. The different set of flows are generated randomly to simulate

spatial traffic variations and the same set is considered in each experiment for all the routing protocols. Thus, we conduct over 2000 ns-3 simulations for BP-MR, BS, and GPSR, resulting in a total of 6000 simulations.

Figure 9 and Figure 10 show the average throughput and the average latency distribution exhibited by BP-MR, BS, and GPSR for the considered non-regular backhaul deployments. As explained in the previous subsection, when data packets are not facing a network void, GPSR always routes a traffic flow following the same path using a greedy forwarding approach without considering the possible congestion in the network, whereas BP-MR and BS are able to use multiple paths, hence load balancing traffic under congestion conditions. When facing a network void, the behavior of GPSR is different from that of BP-MR and BS. GPSR follows the right-hand rule to circumvent a network void, which ensures neither the optimal path in terms of number of hops nor a low congested one [13]. This, jointly with null load balancing capabilities of GPSR explains the observed high latency values of GPSR in Figure 10 even for the 8Mbps and 12Mbps cases compared to BS and BP-MR. BS and also BP-MR circumvent voids by relying in the penalty function of the geographic routing component of Equation 1 explained in Section III without incurring into excessive end-to-end latencies.

As the offered load increases, BS starts experiencing a significant latency degradation due to the HoL blocking problems caused by packets intended towards different directions queued in the same FIFO data queue. This confirms that BS experiences an inefficient use of backhaul interfaces, generating excessive queuing latencies of packets. Such inefficiency
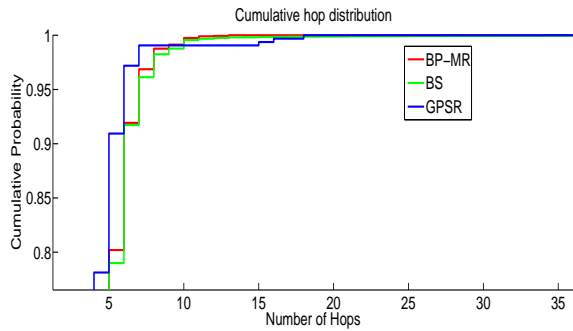
Fig. 11. Hop cumulative distribution of the experiments run in the topology showed in Figure 8 at an input rate of 20Mbps.

also affects the attained throughput with BS, showing worse performance than both BP-MR and GPSR when considering the highest offered load (i.e., 24Mbps). In contrast, BP-MR performs better under all tested conditions both in terms of throughput and latency. This confirms that stage 1 and stage 2 of BP-MR enable a better use of the available network resources also in non-regular deployments, showing improvements of up to one order of magnitude in terms of latency under high loads. This is especially achieved by minimizing the time data packets wait at FIFO queues to be transmitted, which also helps reducing possible queue overflows as the offered load increases.

Figure 11 shows the cumulative hop distribution experienced by data packets over the backhaul topology depicted in Figure 8 with an input rate of 20 Mbps. In this figure, we can see the effects of the inefficiency of the right-hand rule for the GPSR protocol. This is especially noticeable for a slight percentage of packets (around a 2%), which need more than 15 hops to arrive from a source to their intended destination. Therefore, the reason for the significant improvements in terms of latency with respect to GPSR are the BP-MR load balancing capabilities. We can also observe that the hop distribution obtained by BS is similar to the one presented by BP-MR. However, according to Figure 10, the average latency value presented by BS is similar to the one achieved by GPSR. Thus, high latencies experienced by BS are caused mostly by the HoL blocking effect, which contributes to the inefficient use of the backhaul interfaces.

V. CONCLUSIONS

This paper proposes BP-MR, a distributed routing and load balancing protocol for heterogeneous wireless mesh backhauls deployments where each SC site can be equipped with a different number and type of wireless backhaul radios. In BP-MR, each SC maintains a data queue per interface and carries out the routing process in two stages. In the first stage, BP-MR distributes ingress packets among the per-interface queues present in each SC with the goal of reducing the HoL blocking effect in a multi-radio SC. In the second stage, BP-MR uses the Lyapunov drift-plus penalty weight for each interface to compute the actual outgoing interface and the next-hop for each packet at the head of the queues. The two-stage routing process of BP-MR enables traffic load balancing while showing low packet latencies without incurring into an excessive number of hops and alleviating HoL blocking problems of single-queue routing approaches.

To corroborate this, we evaluated BP-MR under several mesh backhaul network scenarios, including different topologies and diverse traffic demands. Ns-3 simulation results reveal that BP-MR obtains remarkable throughput and latency improvements compared to BS, a backpressure routing scheme maintaining a single-queue per node and GPSR, a state of the art geographic routing protocol maintaining per-interface queues. In particular, we showed throughput and latency improvements of up to 34% and one order of magnitude, respectively.

REFERENCES

[1] Cisco Systems, "Cisco visual networking index: Global mobile data traffic forecast update 2014-2019, February 2015."

[2] Senza Fili Consulting, "Small-cell backhaul: Industry trends and market overview," *White Paper*, 2013.

[3] Aviat Networks, "The case for layer 3 intelligence at the mobile cell site," *White Paper*, 2015.

[4] L. Tassiulas and A. Ephremides, "Stability properties of constrained queuing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, Nov. 1992.

[5] L. Ying, S. Shakkottai, A. Reddy, and L. Shihuan, "On combining shortest-path and back-pressure routing over multihop wireless networks," *IEEE/ACM Trans. on Networking*, vol. 19, no. 3, pp. 841–854, 2011.

[6] L. Bui, R. Srikant, and A. Stolyar, "A novel architecture for reduction of delay and queueing structure complexity in the back-pressure algorithm," *IEEE/ACM Trans. on Networking*, vol. 19, no. 6, pp. 1597–1609, 2011.

[7] K. Choumas, T. Korakis, I. Koutsopoulos, and L. Tassiulas, "Implementation and end-to-end throughput evaluation of an ieee 802.11 compliant version of the enhanced-backpressure algorithm," in *Proc. of the 8th EAI TridentCom, Thessaloniki, Greece*, 2012.

[8] J. Núñez-Martínez, J. Baranda, and J. Mangues-Bafalluy, "A self-organized backpressure routing scheme for dynamic small cell deployments," *Ad Hoc Networks, Elsevier*, vol. 25, Part A, pp. 130–140, 2015.

[9] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Synthesis Lectures on Communication Networks, in Morgan & Claypool Publishers, 2010.

[10] "The ns-3 network simulator, Available at: http://www.nsam.org."

[11] B. Karp and H.-T. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proc. of the 6th Annual Int. Conf. on Mobile Computing and Networking (MobiCom)*. Boston, MA, USA, 2000: ACM, 2000, pp. 243–254.

[12] S. Moeller, A. Sridharan, B. Krishnamachari, and O. Gnawali, "Routing without routes: the backpressure collection protocol," in *Proc. of the 9th ACM/IEEE Int. Conf. Information Processing in Sensor Networks (IPSN)*. Stockholm, Sweden: ACM, 2010, pp. 279–290.

[13] D. Chen and P. Varshney, "A survey of void handling techniques for geographic routing in wireless networks," *Communications Surveys Tutorials, IEEE*, vol. 9, no. 1, pp. 50–67, 2007.

[14] A. Fonseca, A. Camoes, and T. Vazao, "Geographical routing implementation in ns3," in *Proceedings of the 5th Int. Conf. Simulation Tools and Techniques (ICST)*. Brussels, Belgium: ACM, 2012, pp. 353–358.