Managing Quality of Service through Intelligent Scheduling in Heterogeneous Wireless Communications Networks

David Lynch*, David Fagan†, Stepan Kucera‡, Holger Claussen§, and Michael O'Neill¶
*†¶Natural Computing Research & Applications Group
School of Business
University College Dublin
Email: * david.fagan@ucd.ie, † david.lynch.1@ucdconnect.ie, ¶ m.oneill@ucd.ie

‡§ Bell Laboratories
Nokia-Ireland
Email: ‡ stepan.kucera@nokia-bell-labs.com, § holger.claussen@nokia-bell-labs.com

***Abstract*—Small Cells are being deployed alongside pre-existing Macro Cells in order to satisfy demand during the current era of exponential growth in mobile traffic. Heterogeneous networks are economical because both cell tiers share the same scarce and expensive spectrum. However, customers at cell edges experience severe cross-tier interference in channel sharing HetNets, resulting in poor service quality. Techniques for improving fairness globally have been developed in previous works. In this paper, a novel method for service differentiation at the level of individual customers is proposed. The proposed algorithm redistributes spectrum on a millisecond timescale, so that premium customers experience minimum downlink rates exceeding a target threshold. System level simulations indicate that downlink rate targets of at least 1 [Mbps] are always satisfied under the proposed scheme. By contrast, naive scheduling achieves the 1 [Mbps] target only 83% of the time. Quality of service can be improved for premium customers without significantly impacting global fairness metrics. Flexible service differentiation will be key to effectively monetizing the next generation of 5G wireless communications networks.**

I. Introduction

Mobile traffic is exploding due to the proliferation of data-hungry User Equipments (UEs) such as smartphones and tablets [6]. Capacity can be increased by densifying the pre-existing Macro Cell (MC) tier with low-powered nodes called Small Cells (SCs) [4]. SCs are installed in an ad-hoc fashion near hotspots like transport hubs or busy shopping districts. The resulting configuration of MCs and SCs is known as a Heterogeneous Network or 'HetNet'. Cellular operators such as AT&T Inc. have embraced this multi-tiered paradigm because MCs and SCs can share spectrum, which is scarce and expensive. However, two issues constrain the capacity of channel sharing HetNets. Firstly, SCs struggle to offload UEs from much stronger MCs resulting in underutilization of the SC tier. Secondly, severe cross-tier interference occurs at SC edges, since nearby high-powered MCs transmit on the same spectrum.

The 3rd Generation Partnership Project [1] has specified mechanisms for more efficiently offloading UEs onto the SC tier, mitigating cell-edge interference, and allocating spectrum [12]. The task of designing controllers that implement these protocols has been the subject of an intensive research effort [7][8][10]. The main goal of the cited studies is to allocate resources fairly, so that all customers experience an acceptable Quality of Service (QoS). Effectively managing QoS in real time is important because channel conditions fluctuate dramatically on short timescales. Therefore, a customer may experience intermittent service outages, which may be unacceptable depending on their use case. For example, occasional lulls may not affect messaging applications, whereas smooth service is essential for multimedia streaming.

Previous work [13] has demonstrated that fairness is improved by scheduling cell-edge UEs to receive more spectrum than high-performing UEs closer to cell centres. In [13] the authors employed a technique from Natural Computing [5] called Grammar-based Genetic Programming (GP) [15] to automatically evolve schedulers for MCs and SCs. However, the state of the art does not support fine-grained service differentiation at the level of individual customers – a core goal of the ongoing 5G standardization [2][3]. In this paper, a novel method is proposed for providing QoS guarantees to customers that subscribe to a premium service plan. Schedules are first computed using evolved models in order to fairly allocate the spectrum. The proposed algorithm then allocates additional spectrum to premium customers if necessary, so that their target downlink rate is realized.

The approach is assessed by simulating a HetNet deployment serving an urban region. The experiments demonstrate that evolved schedulers improve cell-edge downlink rates compared to a baseline scheme used in practice. However, premium customers still occasionally receive less than their target rate. QoS targets are

satisfied with high probability by adjusting the optimized schedules using the proposed algorithm. The impact of prioritized scheduling on global throughput and fairness is small for moderate QoS targets. However, non-premium customers experience significantly lower downlink rates when the proportion of premium customers, or their target downlink rates, become unreasonably large. Thus, a trade-off exists between providing premium service plans for some customers, and ensuring acceptable QoS for all customers. This trade-off presents the operator with an opportunity to design tailored service plans that maximize profit.

The rest of this paper is organized as follows. In Section II, mechanisms for load balancing, interference mitigation, and scheduling in HetNets are outlined. This framework motivates the paper's main goal: achieving fairness and fine-grained service differentiation through intelligent scheduling. In Section III, a novel algorithm for realizing QoS targets is presented. The simulation environment and experiments are described in Section IV, followed by a discussion of the results in Section V. The paper concludes with directions for future work in Section VI.

II. Problem Definition

Mobile traffic is growing rapidly, but wireless spectrum is scarce and expensive. HetNets are a cost-effective means of satisfying demand because both cell tiers can share the same spectrum. However, low-powered SCs struggle to offload UEs from high-powered MCs, and severe interference at cell edges results in poor QoS for cell-edge UEs. In this section, mechanisms for load balancing and interference mitigation are outlined. The motivation for providing QoS guarantees is then discussed.

A. Load Balancing and Interference Mitigation in HetNets

Figure 1 illustrates these load balancing and interference issues in a toy deployment with three MCs m_1, m_2, m_3 and two SCs s_1, s_2. Let $\mathcal{M}$ and $\mathcal{S}$ denote the sets of MCs and SCs respectively. The signal strength that UE u receives from cell $c \in \mathcal{M} \cup \mathcal{S}$ is given by:

$$signal_{u,c} \leftarrow P_c + G_{u,c}, \tag{1}$$

where P_c is the power of cell c (in decibel milliwatts [dBm]), and $G_{u,c}$ is the signal gain from c to u (in decibels [dB]). UE u attaches to its serving cell c' according to the attachment rule:

$$c' = \arg\max_{c} \left(signal_{u,c} + \beta_c\right), \quad \forall\, c \in \mathcal{M} \cup \mathcal{S}, \tag{2}$$

where β_c is a cell selection bias that enables more efficient offloading onto the SC tier, and where $\beta_s \geq 0$ [dB], $\forall s \in \mathcal{S}$ and $\beta_m = 0$ [dB], $\forall m \in \mathcal{M}$. For example, SC s_1 in Figure 1a can utilize a positive selection bias in order to absorb the two nearby UEs from MC m_2. The blue region in Figure 1b depicts the 'expanded region' that forms around s_1 when $\beta_{s_1} > 0$. Herein, UEs u_2 and u_3 attach to s_1 despite receiving a stronger signal from m_2.

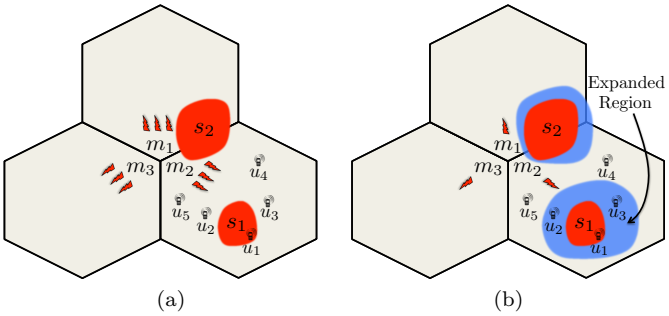

Fig. 1: Load balancing and interference mitigation in a toy HetNet. On the right, SC s_1 broadcasts a positive cell selection bias in order to reduce congestion on MC m_2. Severe interference in the resulting expanded region of s_1 is mitigated by periodically muting the interfering MCs.

The severe interference experienced by u_2 and u_3 in the expanded region of s_1 can be mitigated by periodically muting m_2 and other nearby MCs. Note that cells transmit packets to their attached UEs during 1 [ms] intervals called subframes, and a full frame consists of 8 subframes. When a MC mutes we say that it executes an Almost Blank Subframe (ABS), since only minimal reference signals are transmitted therein. Interference mitigation via ABSs is referred to as enhanced Inter-cell Interference Coordination (eICIC) [1]. For example, let MCs $m \in \mathcal{M}$ mute in the first four out of eight subframes in every frame, so that $ABS_{m,f}^{pattern} = [0, 0, 0, 0, 1, 1, 1, 1]$. Hence, SC attached UEs u_1, u_2 and u_3 experience reduced interference from MCs, and a higher downlink rate from SC s_1. Note that SCs do not implement ABSs so they transmit in every subframe: $ABS_{s,f}^{pattern} = [1, 1, 1, 1, 1, 1, 1, 1]\,,\ \forall s \in \mathcal{S}$.

B. Downlink Rates

Downlink rates are given by Shannon's formula. The signal to interference and noise ratio (*SINR*) received by UE u in subframe f of some frame $\mathcal{F}_t$ (where $t \in \mathbb{N}$ counts the number of elapsed frames) is given by:

$$SINR_{u,f} = \frac{signal_{u,c'} \times ABS_{c',f}^{pattern}}{noise + \sum\limits_{c \in \mathcal{M} \cup \mathcal{S} \setminus c'} signal_{u,c} \times ABS_{c,f}^{pattern}}, \tag{3}$$

where the signal strength is converted to Watts [W], and $noise = 4 \times 10^{-16}$ [W] is background electromagnetic noise. Let $Q_{u,f}$ denote the 'channel quality' experienced by u in f:

$$Q_{u,f} = \log_2\left(1 + SINR_{u,f}\right). \tag{4}$$

Hence, Shannon's formula yields the instantaneous downlink rate:

$$R_{u,f} = Q_{u,f} \times sch_{u,f} \times (20 \times 10^6), \tag{5}$$

where $0.0 \leq sch_{u,f} \leq 1.0$ is the proportion of the spectrum (20×10^6 [Hz] in our experiments) that u is scheduled to receive during f. In every subframe, exactly 100% of the spectrum is utilized by the UEs $u \in \mathcal{A}_c$, where $\mathcal{A}_c$ denotes the set of UEs that attach to cell c. That is,

$$\sum_{u \in \mathcal{A}_c} sch_{u,f} = 1.0, \quad \forall f \in \mathcal{F}_t.$$

Finally, the (average) downlink rate for u over frame $\mathcal{F}_t$ is given by:

$$\overline{R}_u = |\mathcal{F}_t|^{-1} \sum_{f \in \mathcal{F}_t} R_{u,f}. \quad (6)$$

C. Fairness and Service Differentiation via Scheduling

Perceived QoS does not depend linearly on received downlink rate. Extremely low rates ($\ll$1 [Mbps]) will result in many dropped packets, causing service outages and excessive loading times. However, applications may run smoothly at either 25 [Mbps] or 1 [Mbps]. Therefore, extra resources should be directed towards the lowest-rate UEs, by sacrificing those with higher rates.

Shannon's formula in Equation 5 suggests that downlink rates can be managed by controlling the channel quality experienced by UEs (i.e. $Q_{u,f}$), or by computing optimized schedules (i.e. $sch_{u,f}$). Previous work demonstrated that channel conditions are improved at SC edges by optimizing the powers and biases of SCs, and the muting patterns of MCs [8][13]. The proposed technique is compatible with these existing methods for optimizing the HetNet configuration. However, in this paper we implement a static HetNet configuration, in order to isolate the effects of the proposed scheduling method. SC powers are set to 35 [dBm], SC biases are fixed at 10 [dB], and MCs mute during the first four out of eight subframes in every frame.

The next section describes how GP is employed to evolve schedulers for allocating spectrum fairly. The main contribution of the paper is then presented: an algorithm for fine-grained service differentiation at the level of individual customers.

III. Methods

In this section, the shortcomings of a baseline scheduler that is used in real HetNets are identified, motivating GP as a means for automatically constructing intelligent schedulers. An algorithm is then developed for fine-tuning the schedules produced by evolved models, so that QoS targets are satisfied for premium customers.

A. Baseline Scheduling

Consider again the toy HetNet displayed in Figure 1. Here, five UEs download packets during a single 8 [ms] frame $\mathcal{F}_t$. Downlink rates can be computed for $u_1, u_2, \ldots, u_5$ given the channel gains $G_{u,c}$ displayed in Table I and knowledge of the HetNet configuration (recall Section II-C). The steps are as follows.

	u_1	u_2	u_3	u_4	u_5
m_1	−100	−98	−110	−117	−99
m_2	−89	−88	−82	−92	−72
m_3	−120	−92	−85	−96	−101
s_1	−82	−80	−65	−105	−100
s_2	−85	−112	−89	−98	−123

TABLE I: Channel gains $G_{u,c}$ [dB] for the toy HetNet depicted in Figure 1 in some frame $\mathcal{F}_t$. For example, the signal gain from cell m_1 to UE u_1 is −100 [dB] during $\mathcal{F}_t$.

Firstly, Equations 1 and 2 imply that u_1, u_2, u_3 attach to s_1, whereas u_4, u_5 attach to m_2. Hence, Equations 3 and 4 give the channel qualities $Q_{u,f}$ experienced by the UEs in every subframe $f \in \mathcal{F}_t$. Under baseline (BL) scheduling, a cell splits the spectrum evenly among its attached UEs. Therefore, s_1 assigns a proportion $sch_{u,f} = 1/3$ of the spectrum to $u_1, u_2, u_3 \in \mathcal{A}_{s_1}$ in every subframe f. Similarly, m_2 assigns a proportion $sch_{u,f} = 1/2$ of the spectrum to $u_4, u_3 \in \mathcal{A}_{m_2}$ in every subframe f. Given the values of $Q_{u,f}$ and $sch_{u,f}$, Shannon's formula in Equation 5 yields the instantaneous downlink rates $R_{u,f}$. Finally, average downlink rates over frame $\mathcal{F}_t$ are given by Equation 6.

Method	$\overline{R}_{u_1}$	$\overline{R}_{u_2}$	$\overline{R}_{u_3}$
BL	5.1	7.3	24.2

TABLE II: Downlink rates (in [Mbps]) received by UEs u_1, u_2, u_3 attached to SC s_1 in the toy HetNet.

Table II displays the downlink rates received by the SC attached UEs. Notice how BL scheduling has resulted in a highly unfair outcome. Here, the downlink rate for u_3 is over four times greater than that for u_1. The next section describes how evolved models improve fairness by granting more spectrum to the poorly performing UEs.

B. Evolving Schedulers with Genetic Programming

Previous work has established that Genetic Programming (GP) is a powerful framework for evolving schedulers in HetNets [13][14]. In this paper, GP is augmented with a novel algorithm that enables fine-grained service differentiation at the level of individual customers. Evolved models (i.e. a symbolic expressions) map channel quality reports to schedules on a millisecond timescale. Algorithm 1 and Figure 2 describe how a schedule is produced for SC s_1 of the toy HetNet.

The leftmost panel in Figure 2 displays the channel qualities $Q_{u,f}$ that are experienced by $u_1, u_2, u_3 \in \mathcal{A}_{s_1}$ (the UEs attached to s_1) during frame $\mathcal{F}_t$. Note that UEs send reports of $Q_{u,f}$ to their serving cell after every frame. In step 1, the evolved model maps features computed over the channel qualities to the real values $\left(model^{\text{GP}}_{u,f}\right)$. Table IV defines the features and arithmetic operators. For instance, $T8_{u,f}$ is the average of all reported channel qualities over frame $\mathcal{F}_t$. In step 2, the real valued outputs from the scheduler are scaled, so that exactly 100% of the spectrum

is utilized in every subframe. For example, consider the following simple scheduler:

$$model^{\text{GP}}_{u,f} = (T12_{u,f} * T12_{u,f})\%(T12_{u,f} * T1_{u,f}) = \frac{f * f}{\sqrt{1 + (f * Q_{u,f})^2}}. \quad (7)$$

The rightmost panel of Figure 2 displays the schedule generated by the model in Equation 7.

Algorithm 1 Generate schedule using evolved model.

procedure COMPUTE PRIORITIZED SCHEDULE($Q_{u,f}$, $\forall (u,f) \in \mathcal{A}_c \times \mathcal{F}_t$)
  $sch^{\text{GP}}_{u,f} = 0.0, \quad \forall (u,f) \in \mathcal{A}_c \times \mathcal{F}_t$; #initialise schedule
  #Step 1 – execute model:
  for $u \in \mathcal{A}_c$ **do** #iterate over UEs
    for $f \in \mathcal{F}_t$ **do** #iterate over subframes
      Compute features in Table IV given $Q_{u,f}$;
      Hence, execute the scheduler yielding $model_{u,f}$;
  #Step 2 – scaling:
  for $u \in \mathcal{A}_c$ **do**
    for $f \in \mathcal{F}_t$ **do**
      $sch^{\text{GP}}_{u,f} = \frac{model^{\text{GP}}_{u,f}}{\sum_{u \in \mathcal{A}_c} model^{\text{GP}}_{u,f}}$;
  return $sch^{\text{GP}}_{u,f}, \quad \forall (u,f) \in \mathcal{A}_c \times \mathcal{F}_t$;

	u_1	u_2	u_3
1	0.4	0.5	1.3
2	0.4	0.5	1.3
3	0.4	0.5	1.3
4	0.4	0.5	1.3
5	1.2	1.8	6.3
6	1.2	1.8	6.3
7	1.2	1.8	6.3
8	1.2	1.8	6.3

$Q_{u,f}$

Step 1 ⟼

	u_1	u_2	u_3
1	19.09	15.52	6.13
2	16.48	13.46	5.35
3	13.85	11.38	4.58
4	11.18	9.28	3.80
5	3.26	2.20	0.63
6	2.41	1.64	0.48
7	1.54	1.07	0.31
8	0.64	0.49	0.16

$model_{u,f}$

Step 2 ⟼

	u_1	u_2	u_3
1	0.47	0.38	0.15
2	0.47	0.37	0.16
3	0.46	0.38	0.16
4	0.46	0.38	0.16
5	0.53	0.37	0.10
6	0.52	0.37	0.11
7	0.52	0.37	0.11
8	0.50	0.38	0.12

$sch^{\text{GP}}_{u,f}$

Fig. 2: In step 1, channel quality reports from the UEs attached to SC s_1 are mapped to real values. In step 2, outputs are scaled so that 100% of the spectrum is utilized per subframe. Fairness is achieved by allocating more spectrum to UEs with the lowest channel qualities.

Replicating the calculations from Section III-A yields the downlink rates received by u_1, u_2, and u_3 under the fair schedule. Table III shows that fairness is improved relative to the baseline, since the three UEs that attach to s_1 now experience similar downlink rates.

Method	$\overline{R}_{u_1}$	$\overline{R}_{u_2}$	$\overline{R}_{u_3}$
BL	5.1	7.3	24.2
GP	7.70	8.19	8.56

TABLE III: Downlink rates (in [Mbps]) received by $u \in \mathcal{A}_{s_1}$ under baseline (BL) and fair (GP) scheduling.

High performing models are evolved via grammar-based GP [9][15]. Table V displays the Backus-Naur-Form grammar which defines the search space of symbolic expressions

Feature	Interpretation	Operator	Interpretation
$T1_{u,f}$	$Q_{u,f}$	$(x+y)$	$x+y$
$T2_{u,f}$ - $T4_{u,f}$	avg, min, max $\{Q_{u,*}\}$	$(x-y)$	$x-y$
$T5_{u,f}$ - $T7_{u,f}$	avg, min, max $\{Q_{*,f}\}$	$(x \times y)$	$x \times y$
$T8_{u,f}$ - $T10_{u,f}$	avg, min, max $\{Q_{*,*}\}$	$(x\%y)$	$\frac{x}{\sqrt{1+y^2}}$
$T11_{u,f}$	$\text{ID}_u \in [1,2,\ldots,\lvert\mathcal{A}_c\rvert]$	$plog(x)$	$\log(1+\lvert x\rvert)$
$T12_{u,f}$	$f \in [1,2,\ldots,8]$	$sine(x)$	$\sin(x)$
		$psqrt(x)$	$\sqrt{\lvert x\rvert}$

TABLE IV: Interpretation of the grammar elements. A cell's attached UEs are sorted in ascending order based on received channel quality and assigned identifiers (ID_u).

⟨e⟩ ::= ⟨r⟩ | ⟨r⟩ | ⟨r⟩ | ⟨T⟩
⟨r⟩ ::= ⟨A$_1$⟩(⟨e⟩) | (⟨e⟩⟨A$_2$⟩⟨e⟩)
⟨A$_1$⟩ ::= $plog$ | $sine$ | $psqrt$
⟨A$_2$⟩ ::= + | − | × | %
⟨T⟩ ::= $T1_{u,f}$ | $T2_{u,f}$ | ,..., | $T12_{u,f}$ | ⟨n⟩ | ⟨n⟩
⟨n⟩ ::= −1.0 | −0.9 | ,..., | 0.9 | 1.0

TABLE V: Backus-Naur-Form grammar.

that is explored by GP. The standard hyperparameter settings recommended by [9] are adopted, since GP reliably yields highly fit models [14]. The function and terminal sets are defined in Table IV. A GP run proceeds as follows:

1) Initialize a population of 1000 randomly generated derivation trees.
2) Assign a fitness to each individual (i.e. derivation tree) by evaluating the corresponding model (i.e. phenotype) in simulation.
3) Select parents from the current population using tournament selection with tournament size 5.
4) Each pair of selected parents undergoes subtree crossover with a probability of 0.7. Subtrees rooted at randomly selected non-terminals in each parent individual are swapped.
5) Mutate all of the resulting children using subtree mutation. A randomly selected non-terminal in the derivation tree is replaced by a new random subtree.
6) Replace the worst 99% of the current population with the children resulting from subtree crossover and mutation. The remaining 1% are the fittest 'elites' found so far, and they enter into the next generation unchanged.
7) Return the fittest individual (on validation data) found after 250 generations.

In item 2, models are evaluated using 30 channel gain matrices (similar to that in Table I) corresponding to 30 different frames. In this paper, channel gain data are generated in simulation, but UEs report channel gains in real HetNets. Hence, the fitness of a model is given by the sum logarithm of downlink rates [8]:

$$\textit{fitness} = \sum_u \log\left(\overline{R}_u\right), \quad \forall u \text{ receiving data}, \quad (8)$$

where *fitness* is high when spectrum is fairly allocated among UEs. Of the 30 channel gain matrices, 25 serve

as training cases and the remaining 5 are validation cases used for model selection. GP automatically discovers complex models that are tailored to the deployment scenario [13]. Automation is valuable in this domain because manually designing schedulers for different deployment scenarios is a labor intensive task.

C. Prioritized Scheduling

Algorithm 2 specifies a procedure for fine-tuning the schedule generated by an evolved model, so that premium customers receive their target downlink rate.

Algorithm 2 Modify schedule: $\overline{R}_u^{\text{PR}} \geq \overline{R}_u^{\text{target}}, \quad \forall u \in \mathcal{P}.$

1: **procedure** COMPUTE PRIORITIZED SCHEDULE
2: $sch_{u,f}^{\text{PR}} = sch_{u,f}^{\text{GP}}, \quad \forall (u,f) \in \mathcal{A}_c \times \mathcal{F}_t;$
3: **for** $f \in \mathcal{F}_t$ **do**
4: $R_{u,f}^{\text{PR}} = Q_{u,f} \times sch_{u,f}^{\text{PR}} \times (20 \times 10^6), \quad \forall (u,f) \in \mathcal{A}_c \times \mathcal{F}_t;$
5: $\overline{R}_u^{\text{PR}} = \frac{1}{|\mathcal{F}|} \sum_{f \in \mathcal{F}} R_{u,f}^{\text{PR}}, \quad \forall u \in \mathcal{A}_c;$ a
6: $\mathcal{P}' = \left\{ u \in \mathcal{A}_c \cap \mathcal{P} \middle| \overline{R}_u^{\text{PR}} < \overline{R}_u^{\text{target}} \right\};$
7: $\delta_{u,f} = \frac{|\mathcal{F}| \times \left(\overline{R}_u^{\text{target}} - \overline{R}_u^{\text{PR}} \right)}{Q_{u,f} \times (20 \times 10^6)}, \quad \forall u \in \mathcal{P}';$
8: $\Delta_{u,f} = sch_{u,f}^{\text{PR}}, \quad \forall u \in \mathcal{A}_c \setminus \mathcal{P}';$
9: $\delta_f = \sum_{u \in \mathcal{P}'} \delta_{u,f};$ #spectrum needed by premium UEs
10: $\Delta_f = \sum_{u \in \mathcal{A}_c \setminus \mathcal{P}'} \Delta_{u,f};$ #total spectrum available
11: **if** $\delta_f < \Delta_f$ **then**
12: $sch_{u,f}^{\text{PR}} = sch_{u,f}^{\text{PR}} + \delta_{u,f}, \quad \forall u \in \mathcal{P}';$
13: $sch_{u,f}^{\text{PR}} = sch_{u,f}^{\text{PR}} - \delta_{u,f} \left(\frac{\Delta_{u,f}}{\Delta_f} \right), \quad \forall u \in \mathcal{A}_c \setminus \mathcal{P}';$
14: **else if** $\delta_f \geq \Delta_f$ **then**
15: $sch_{u,f}^{\text{PR}} = sch_{u,f}^{\text{PR}} + \Delta_f \left(\frac{\delta_{u,f}}{\delta_f} \right), \quad \forall u \in \mathcal{P}';$
16: $sch_{u,f}^{\text{PR}} = sch_{u,f}^{\text{PR}} - \Delta_{u,f}, \quad \forall u \in \mathcal{A}_c \setminus \mathcal{P}';$
return $sch_{u,f}^{\text{PR}}, \quad \forall (u,f) \in \mathcal{A}_c \times \mathcal{F}_t;$

Intuitively, Algorithm 2 redistributes the spectrum from non-premium to premium UEs, so that the latter receive their target rate over a frame. Let $\mathcal{P}$ denote the subset of UEs that subscribe to the premium service plan, and let $\overline{R}_u^{\text{target}}$ denote their target rate. Recall that $\mathcal{A}_c$ denotes the set of UEs that are attached to cell c during frame $\mathcal{F}_t$. On line 2, the prioritized (PR) schedule is initialised to the schedule generated by the evolved model. The algorithm then iterates over subframes (line 3). Downlink rates that would be received under the current prioritized schedule are computed on lines 4 and 5. The subset of premium UEs ($\mathcal{P}'$) that require more spectrum in order to reach their target rate is then determined on line 6. The amount of extra spectrum ($\delta_{u,f}$) needed by each $u \in \mathcal{P}'$ in subframe f in order to reach their target rate is computed on line 7. Similarly on line 8, the amount of spectrum that non-premium UEs can sacrifice ($\Delta_{u,f}$) is computed. Lines 11–13 describe how spectrum is redistributed to premium UEs if they require less spectrum than is available in f. Similarly, lines 14–16 deal with the case when premium UEs request more spectrum than is available in f. Spectrum is diverted from non-premium UEs based on their ability to sacrifice it (line 13), and granted to premium UEs in proportion to their need (line 15). These steps preserve the constraint that 100% of the spectrum is utilized per subframe.

For illustration, let u_1 from the toy HetNet subscribe to a premium service plan offering minimum downlink rates of 10 [Mbps]. That is, $u_1 \in \mathcal{P}$ and $\overline{R}_{u_1}^{\text{target}} = 10$. Figure 3 displays the prioritized (PR) schedule generated by Algorithm 2 after it modifies the fair (GP) schedule from Figure 2. Notice that all of the spectrum is awarded to u_1 during the first four subframes. Hence, the average downlink rate for u_1 is increased by slightly sacrificing non-premium UEs u_2 and u_3 over the frame.

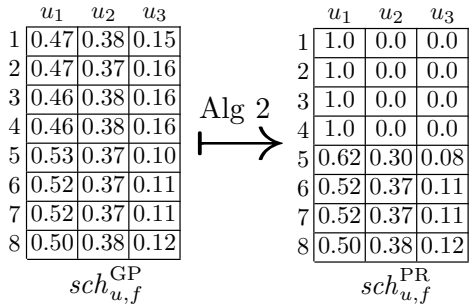

Fig. 3: The prioritized schedule for s_1 ensures that premium UE $u_1 \in \mathcal{P}$ receives an average downlink rate of at least $\overline{R}_u^{\text{PR}} = 10.0$ [Mbps]. Extra spectrum is granted to u_1 at the expense of non-premium UEs u_2 and u_3.

Finally, Table VI displays the average downlink rates under the baseline (BL: Section III-A), fair (GP: Section III-B), and prioritized (PR: Section III-C) schedules. The last row indicates that the target downlink rate for u_1 is realized when s_1 executes the prioritized schedule. Furthermore, the remaining spectrum remains fairly distributed between the non-premium UEs since $\overline{R}_{u_2}^{\text{PR}} \approx \overline{R}_{u_3}^{\text{PR}}$.

Method	$\overline{R}_{u_1}$	$\overline{R}_{u_2}$	$\overline{R}_{u_3}$
BL	5.1	7.3	24.2
GP	7.70	8.19	8.56
PR	10.0	6.08	6.32

TABLE VI: Downlink rates $\overline{R}_u$ [Mbps] for $u \in \mathcal{A}_{s_1}$ under various scheduling regimes.

IV. EXPERIMENTS

A HetNet with 21 MCs and 21 SCs serving Dublin City centre was simulated. Realistic channel gains were computed by modelling obstacles in the environment such as buildings, roadways, and parks. Table VII displays

the parameters of the simulation and path loss model. Figure 4a displays the terrain. The path loss model was described in [11]. Figure 4b shows how channel quality varies throughout the network. Seven trisector MC towers were spaced out on a regular grid. SCs were distributed randomly in order to reflect their ad-hoc installation where hotspots arise. The blue hue at cell edges indicates regions of low channel quality where inter-cell interference is significant.

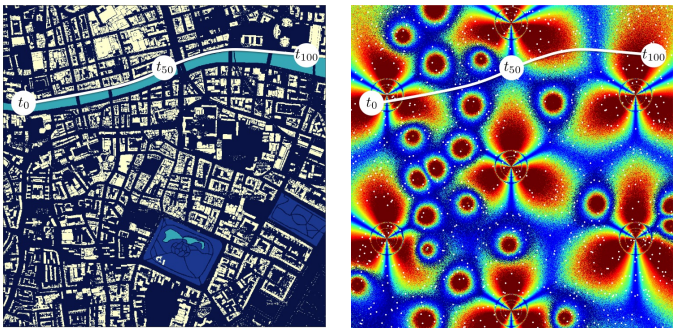

(a) Terrain. (b) Channel Quality Heatmap.

Fig. 4: A HetNet with 21 MCs and 21 SCs serving a 3.24 km^2 region of Dublin City is displayed. UEs are represented by white dots. Starting and end points for a single UE traversing the HetNet are displayed.

Parameters	Value
Scenario	
Indoor/outdoor map	Dublin
MC BS placement	7 eNodeB with 3 sectors each (hexagonal grid)
SC BS placement	Uniformly randomly distributed
Inter-MC BS distance	800 [m]
Scenario resolution	2 [m]
LPC power weight set	$\{0, 1\}$ (binary on/off power control)
Noise density	-174 [dBm/Hz]
Channel	
Carrier frequency	2 [GHz]
Bandwidth	20 [MHz]
NLOS path-loss	$G_{\mathrm{Pn}} = -21.5 - 39\log_{10}(d)$ (MC)
	$G_{\mathrm{Pn}} = -30.5 - 36.7\log_{10}(d)$ (SC)
LOS path-loss	$G_{\mathrm{Pl}} = -34.02 - 22\log_{10}(d)$
Shadow fading (SF)	6 dB std dev.
SF correlation	$R = \mathrm{e}^{-1/20d}$, 50% inter-site
Environment loss	$G_{\mathrm{E},n} = -20$ [dB] if indoor, 0 [dB] if outdoor
Antenna	
Height	25 [m] (MC), 10 [m] (SC)
Maximum gain	$G_{\max} = 15.5$ [dBi] (MC), 7.06 [dBi] (SC)
H. halfpow. beamwidth	$\alpha = 65^\circ$
V. halfpow. beamwidth	$\beta = 11.5^\circ$ (MC)
Front-to-back ratio	$\kappa = 30$ [dB] (MC)
Downtilt	$\delta_1 = 8.47^\circ$ (MC)
Elements & spacing	4 element dipole, $d_{\mathrm{elem.}}$=0.6λ (SC)
Phase difference	$\delta_{\mathrm{phase}} = 95^\circ$ (SC)
Element amplitude	$a_{\mathrm{elem.}} = [0.9691, 1.0768, 1.0768, 0.8614]$ (SC)

TABLE VII: Simulation Parameters.

The best model from 30 runs of GP was executed during 'snapshots' that simulated several minutes of activity in the HetNet. Exactly 1000 UEs were dropped randomly onto the map at the beginning of a snapshot. Each UE then moved at uniform velocity towards a randomly chosen end point. Downlink rates for all 1000 UEs were computed in 100 frames that were sampled uniformly throughout the snapshot. A total of 30 snapshots were simulated in order to resolve the following research questions:

1) To what extent can Algorithm 2 satisfy QoS targets for premium UEs?
2) How rapidly does the probability of achieving QoS targets decrease as the proportion of premium UEs and their target downlink rate increase?
3) How is overall network throughput and fairness impacted by redistributing spectrum from non-premium to premium UEs?

V. Results and Discussion

The experimental results are presented and discussed in this section. We first confirm that evolved models improve fairness relative to the baseline method. A case study then motivates the proposed prioritized scheduling algorithm. Properties of Algorithm 2 are analyzed, including its ability to achieve QoS targets. Finally, the impact on global fairness incurred by redistributing spectrum to premium customers is assessed.

A. Managing QoS through Intelligent Scheduling

Figure 5 shows that downlink rates are increased for the worst performing UEs when evolved models are used for scheduling. Better fairness improves the QoS experienced by vulnerable customers at cell-edges, since they receive a higher downlink rate.

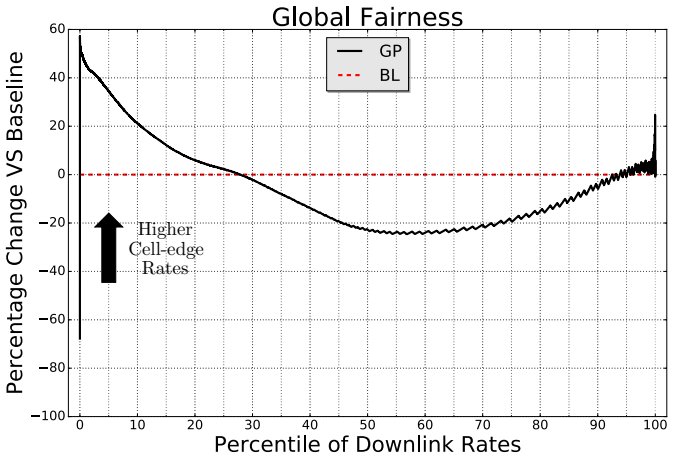

Fig. 5: Downlink rates for the worst performing UEs (0$^{\mathrm{th}}$–5$^{\mathrm{th}}$ percentile) are increased by between 35% and 58% versus baseline (BL), when evolved models (GP) are used for scheduling. Fairness is achieved by slightly sacrificing the best performing UEs (roughly, 25$^{\mathrm{th}}$–100$^{\mathrm{th}}$ percentile).

Figure 6 displays the downlink rates received by a single customer traveling through the HetNet. The path traced out by this customer along the River Liffey in Dublin City was indicated in Figure 4. The target downlink rate of 2 [Mbps] is rarely achieved under fair scheduling (GP)

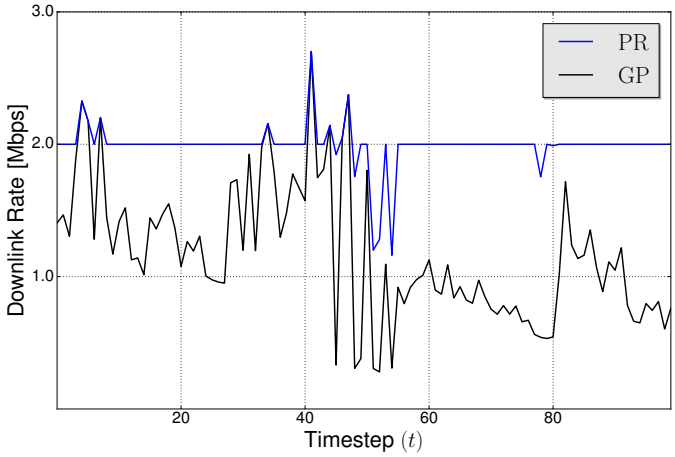

Fig. 6: Downlink rates received by a particular UE during a snapshot, under fair (GP) and prioritized (PR) scheduling. The target downlink rate is 2 [Mbps].

alone. However, the target rate is typically (though not always) achieved after the optimized schedules are modified using Algorithm 2. Notice how prioritized scheduling smooths out the large fluctuations that occur during the handover between two MCs from timesteps $t = 40$ to $t = 60$. Figure 7 shows that downlink rates are increased by over 450% during some frames. Figure 7 also indicates that the target rate is occasionally achieved by evolved models, and thus there is sometimes no need to modify the schedule. Algorithm 2 automatically detects when a cell's schedule should be modified, and it precisely redistributes the spectrum so that target rates are satisfied for premium customers.

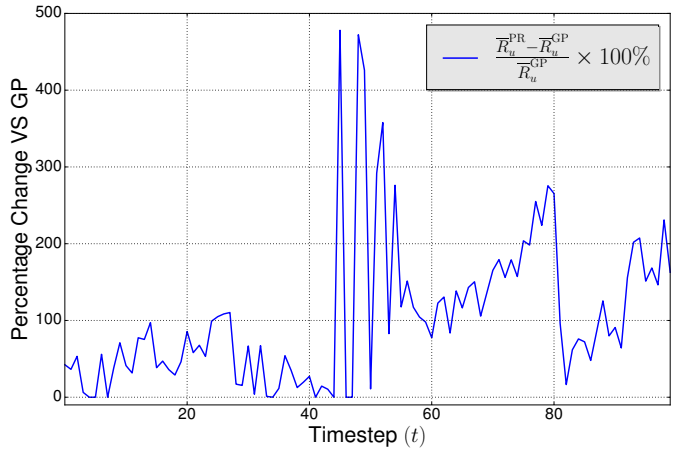

Fig. 7: Percentage increase in downlink rates versus fair scheduling for a single customer traveling through the HetNet.

B. Statistical Analysis

The preceding case study demonstrated the benefits of prioritized scheduling for a single customer. In this section, a statistical analysis of Algorithm 2 is presented. Downlink rates were computed for 1000 UEs in the 100 frames of 30 different snapshots (recall Section IV).

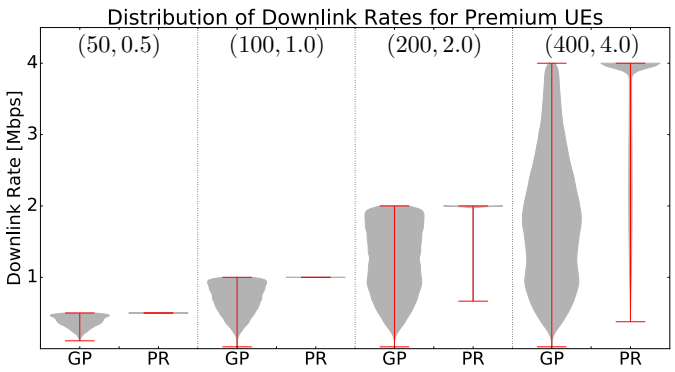

Fig. 8: Violin plots visualizing how downlink rates are distributed for premium UEs under fair (GP) and prioritized (PR) scheduling. The values in parentheses indicate the number of premium UEs (out of 1000) and their target rate (in [Mbps]) respectively.

Figure 8 displays the distribution of downlink rates for premium customers under fair (columns labeled "GP") and prioritized (columns labeled "PR") scheduling. The subset of premium customers with downlink rates below the target under fair scheduling are first identified. Their downlink rates are then computed after executing Algorithm 2. Fair scheduling typically results in downlink rates for premium customers that are below their target rate. In contrast, the distribution closely straddles the target value after schedules have been modified using Algorithm 2. In fact, the rightmost column indicates that target rates as high as 4 [Mbps] are usually achieved.

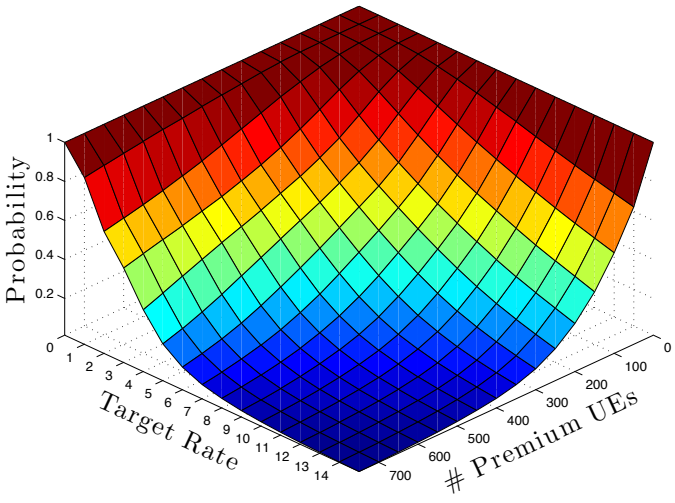

Fig. 9: Probability that a premium customer will receive their target downlink rate for different settings of the target rate and proportion of premium customers.

Of course, it is not always possible to satisfy QoS targets

since non-premium UEs can only liberate a finite amount of spectrum. Figure 9 visualizes how the probability that target rates are achieved decreases, as the proportion of premium customers and their target rate increases. QoS targets can be satisfied with high probability for a wide range of settings. For example, target rates are delivered to premium customers with a probability exceeding 0.95 if up to 150 premium customers have a target rate of 3.0 [Mbps].

C. Impact of Prioritized Scheduling

Global fairness is damaged by diverting spectrum away from non-premium UEs in order to satisfy QoS targets for premium UEs. It follows that a trade-off exists between overall fairness, and the number of premium UEs and their target rates. Figure 10 displays the percentage change in downlink rates compared to the baseline method for fair and prioritized scheduling. As discussed in Section V-A, evolved models improve fairness by increasing downlink rates for the worst performing cell-edge UEs (e.g. 0^{th}–5^{th} percentile). The blue curve implies that target rates of 0.75 [Mbps] can be delivered to 50 (out of 1000) premium UEs without significantly degrading cell-edge throughput. The cyan curve suggests that target rates up to 1.5 [Mbps] can be tolerated. However, cell-edge throughput is severely damaged if the number of premium UEs and their target rate increases above 100 and 1.5 [Mbps] respectively.

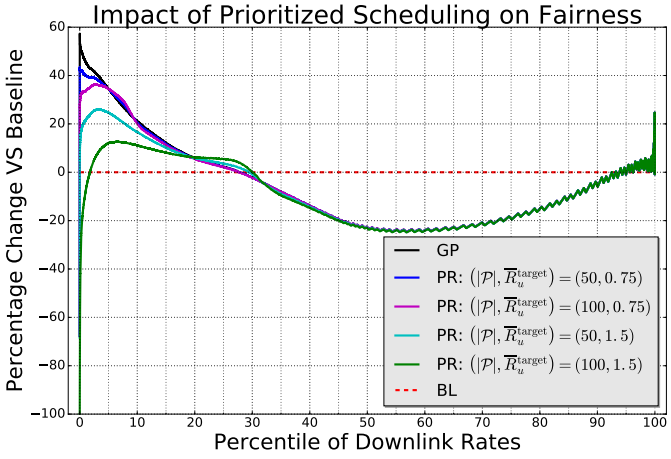

Fig. 10: Evolved models (GP) improve fairness compared to baseline scheduling (BL). There is a trade-off between improving QoS for premium UEs and sustaining global fairness.

VI. Conclusions

Maintaining high customer satisfaction is vital in the fiercely competitive wireless telecommunications industry. The technique for service differentiation that was outlined in this paper provides stable QoS for customers that avail of a premium service plan. However, the HetNet's ability to support such a scheme depends on the number of participants, and their QoS targets. Therefore, network operators should price plans appropriately, so that a small subset of all customers take part. In practice, customers could be dynamically assigned to the premium class, depending on the application they are running. For example, a customer could be given preferential treatment temporarily when streaming high definition video. Prioritized scheduling has minimal impact on network-wide performance under reasonable conditions. In real HetNets, scheduling contributes to better QoS and lower latency, but congestion control and traffic partitioning is also necessary. Future work could explore the potential for natural computing techniques in 5G HetNets.

Acknowledgements

This research is based upon works supported by the Science Foundation Ireland under grant 13/IA/1850.

References

[1] 3GPP, Release 10. Technical report, December 2014.

[2] R. Agrawal, A. Bedekar, R. J. La, and V. Subramanian. Class and channel condition based weighted proportional fair scheduler. *Teletraffic Science and Engineering*, 4:553–567, 2001.

[3] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty. Next generation/dynamic spectrum access/cognitive radio wireless networks: A survey. *Computer networks*, 50(13):2127–2159, 2006.

[4] N. Bhushan, J. Li, D. Malladi, R. Gilmore, D. Brenner, A. Damnjanovic, R. Sukhavasi, C. Patel, and S. Geirhofer. Network densification: the dominant theme for wireless evolution into 5G. *IEEE Communications Magazine*, 52(2):82–89, 2014.

[5] A. Brabazon, M. O'Neill, and S. McGarraghy. *Natural computing algorithms*. Springer, 2015.

[6] Cisco. Cisco visual networking index: Global mobile data traffic forecast update, 2014 – 2019. Technical report, 2015.

[7] A. Damnjanovic, J. Montojo, Y. Wei, T. Ji, T. Luo, M. Vajapeyam, T. Yoo, O. Song, and D. Malladi. A survey on 3GPP heterogeneous networks. *IEEE Wireless Communications*, 18(3), 2011.

[8] S. Deb, P. Monogioudis, J. Miernik, and J. P. Seymour. Algorithms for enhanced inter-cell interference coordination (eICIC) in LTE HetNets. *IEEE/ACM Transactions on Networking (TON)*, 22(1):137–150, 2014.

[9] M. Fenton, J. McDermott, D. Fagan, S. Forstenlechner, E. Hemberg, and M. O'Neill. PonyGE2: Grammatical evolution in python. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1194–1201. ACM, 2017.

[10] Y. Liu, C. S. Chen, C. W. Sung, and C. Singh. A Game Theoretic Distributed Algorithm for FeICIC Optimization in LTE-A HetNets. *IEEE/ACM Transactions on Networking*, 2017.

[11] D. Lopez-Perez, H. Claussen, and L. Ho. Improved frequency reuse schemes with horizontal sector offset for LTE. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2013 IEEE 24th International Symposium on*, pages 2159–2164. IEEE, 2013.

[12] D. Lopez-Perez, I. Guvenc, G. De la Roche, M. Kountouris, T. Q. Quek, and J. Zhang. Enhanced intercell interference coordination challenges in heterogeneous networks. *IEEE Wireless Communications*, 18(3):22–30, 2011.

[13] D. Lynch, M. Fenton, D. Fagan, S. Kucera, H. Claussen, and M. O'Neill. Automated Self-Optimization in 5G Heterogeneous Wireless Communications Networks. 2018.

[14] D. Lynch, M. Fenton, S. Kucera, H. Claussen, and M. O'Neill. Scheduling in Heterogeneous Networks using Grammar-based Genetic Programming. In *Genetic Programming*, pages 83–98. Springer, 2016. LNCS 9594.

[15] R. McKay, N. X. Hoai, P. A. Whigham, S. Yin, and M. O'Neill. Grammar-based Genetic programming – A Survey. *Genetic Programming and Evolvable Machines*, 11(3-4):365–396, 2010.