Dynamic Frequency Reuse in Dense Cellular Networks

Bart Post
Eindhoven University of Technology
Eindhoven, The Netherlands

Sem Borst
Eindhoven University of Technology
Eindhoven, The Netherlands

Hans van den Berg
TNO, Den Haag
The Netherlands

***Abstract*—Network densification has emerged as a powerful paradigm to boost spectral efficiency and accommodate the continual rise in demand for wireless capacity. In dense cell deployments however, overlapping coverage areas may cause highly varying interference conditions among different cells. Moreover, denser networks experience more temporal load fluctuations due to daily and hourly changing usage patterns. The currently applied universal reuse frequency allocation is not suitable to deal with these issues, and needs to be tailored to dense cell deployments to ensure adequate performance in such scenarios.**

In this paper we present a dynamic, load aware and self-adapting frequency allocation scheme designed for dense cellular networks: the DyCRA scheme (Dynamic Cost/Reward based Allocation). The scheme makes decisions based on cost-reward trade-offs: rewards arise in the form of capacity, and costs arise in the form of interference (under spatial reuse). We quantify these costs and rewards based on SINR, and use periodic load estimates to determine if access points are in need of extra frequencies, or can spare them, and the cost/reward structure is used to determine which frequencies are allocated or released. Extensive simulation results show that the DyCRA scheme provides efficient resource allocations that adapt to changing traffic conditions and yields significant performance gains in scenarios with non-stationary traffic demands.

***Index Terms*—Self-organizing, frequency allocation, dense cellular networks, dynamic load balancing**

I. Introduction

Over the last years wireless networks have seen an immense rise in demand, taxing the networks to the limits of their capacity. Network densification is viewed as one of the key options to boost capacity by reducing cell sizes and allowing for higher spectral reuse and efficiency [1]–[3].

In dense cell deployments, physical constraints will typically make it hard to arrange the access points (APs) in an ideal hexagonal pattern. This causes the natural cell regions to be irregularly shaped and the coverage areas to significantly overlap. Small and irregular shaped cell regions result in more spatial variation of nominal traffic loads than in traditional macro cell networks. Moreover, the current standard frequency planning in cellular networks requires a (fairly) regular placement of the APs and is based on universal frequency reuse [4]. In dense cell deployments however, the overlapping coverage areas in combination with universal frequency reuse may lead to highly varying and destructive interference conditions among different cells [5].

To ensure adequate performance in dense cell deployments, the frequency allocations will have to take the load imbalances into account and direct capacity towards APs where it is most needed [6], [7], while keeping interference manageable. In addition, dense networks experience more temporal load fluctuations due to daily and hourly changing usage patterns, which furthermore requires sub-band allocation schemes to dynamically react to changing load conditions, matching demand and capacity in real time.

Despite the myriad research papers covering spectrum allocation and (fractional) frequency reuse (FFR), frequency allocation in dense cellular networks combines system characteristics which have not been studied jointly: irregular cell sizes, varying interference conditions and dynamic and imbalanced traffic demands. Varying interference and dynamic traffic is typically addressed by enhanced inter-cell interference coordination (eICIC) and/or FFR, but their focus is usually on a single macro cell that has to share its spectral resources with embedded smaller cells (e.g. pico or femto cells). The combination of irregular cell sizes and varying interference is mostly covered by interference graphs, but these types of results generally do not deal with dynamic traffic.

What furthermore adds to the complex operation of dense cellular networks, is that with the massive numbers of APs and the exact load conditions typically being hard to predict under dynamic user behaviour, manual cell planning and traffic engineering is highly impractical. This motivates the use of self-organizing schemes; the frequency allocation can be automatically adapted to match offered traffic loads with spectral capacity. Such self-organizing schemes could also be used to provide self-healing capabilities of the network: automatic reconfiguration once APs shut down due to failures.

A. Contributions and results

In this paper we present a dynamic, load aware and self-adapting frequency allocation scheme designed for dense cellular networks. The scheme, which we call the DyCRA scheme (Dynamic Cost/Reward based Allocation), has two components: (i) a trigger component which determines when to make a change in the frequency allocation, and (ii) a decision component which decides what change to make.

We use a modified version of the single load interval (SLI)-algorithm [8] as trigger. Using periodic load proxies, we classify APs into three states: (a) in need of an extra frequency, (b) no adaptation needed, (c) can spare a frequency. The load estimates for APs determine in which states the APs are

classified, depending on whether a load estimate is (a) above, (b) within, or (c) below a predefined load interval $[\rho_{\min}, \rho_{\max}]$.

The decision component of the DyCRA scheme decides which frequency to acquire/release, and can also decide not to acquire a frequency. The decisions are based on cost-reward trade-offs. Assigning a specific frequency to an AP brings a reward, namely extra capacity for that AP, but also comes at a cost: interference caused at surrounding APs. We quantify these costs and rewards based on SINR values, and only if the reward outweighs the cost by a certain margin will an AP acquire a new frequency. An AP acquires the frequency that comes with the best cost-reward trade-off. Combining the two components (i) and (ii), we obtain a dynamic resource acquisition (DRA) scheme for dense cellular networks.

Extensive simulation results show that the DyCRA scheme indeed provides a highly flexible resource allocation that adapts to changing traffic conditions and outperforms a static allocation in scenarios with non-stationary traffic demands.

B. Discussion and related work

Resource allocation has been a major theme in wireless networks research and hence there is a vast literature covering resource allocation problems in many flavours. In this section we will review the dominant research themes or directions of resource allocation in wireless networks, and also explain how our contribution goes beyond the state-of-the-art.

Surprisingly, there is hardly any literature covering frequency re-use in dense cellular (4G) networks that does not assume universal reuse among cells. A prevalent assumption in 4G networks is universal frequency reuse (or reuse 1) among macro cells, based on regular positioning of the macro APs. This universal reuse is mostly aided by dividing the cell into three sectors and reusing the frequencies in three corresponding subsets. When smaller cells come into the picture, the focus is mostly on HetNets with a typical scenario of one macro cell with several smaller cells in its footprint. Frequency (or resource) allocation schemes then mostly concentrate on sharing the spectral resources of the macro with smaller (pico) cells, avoiding macro-pico interference as much as possible. In this direction we find eICIC (e.g. [9]–[11]) and/or FFR (e.g. [12], [13]), both of which typically work with dynamic traffic.

Frequency allocation also received significant attention from a graph colouring point of view [14], [15], where vertices of the graph represent APs, and two APs are connected by an edge if they are not allowed to transmit on the same frequency. These approaches are typically applied to stationary settings [16], and obtaining such an interference graph may require extensive field measurements [17]. When the network topology changes (e.g. installing, removing, or malfunctioning APs), the interference graph has to be adapted, possibly requiring more field measurements.

Uygungelen et al. [18] proposed a method to construct an interference graph based on SINR values of users. Their method requires a centralized processing unit and knowledge of SINR values of users, and both are perfectly possible in dense cellular networks. However, the approach of Uygungelen et al. is very conservative: an edge is created between two APs if there is at least one user whose SINR drops below a predefined threshold. Thus, a whole AP may be restricted from using the frequency set of another AP just because of the position of a single user. In addition, the proposed method does not deal with dynamic user populations. The interference graph is only recalculated when the (active) AP set changes. In our setting, we decide per frequency (and not per frequency set) if an AP is allowed to use that frequency or not.

Several works have studied dynamic resource allocation schemes in voice cellular networks. Ule and Boucherie [19] proposed a channel borrowing strategy for road covering networks, using traffic predictions for optimal channel borrowing strategies. Nanda and Goodman [20] proposed a DRA algorithm where the decision for a specific AP which frequency to acquire or release is based on a cost or reward function respectively. For resource acquisition, the resource with the lowest cost is selected, and for resource release, the resource with the highest release reward is selected. Nanda and Goodman based these rewards and costs on blocking probabilities for voice cellular networks. We on the other hand do not consider a channel based network or Erlang loss network, but a flow based network (4G flows) with a proportional fair scheduling policy at the APs. In addition, we do not work with a reuse distance as [19], [20], but rather base our selection criteria and reward/cost functions on SINR values. The fundamental difference between our approach and the one of Nanda and Goodman is that our cost and reward functions do not only determine which frequency is best to assign, but they also automatically (and dynamically) determine which frequencies can be assigned, whereas Nanda and Goodman needed the latter as input in the form of an interference graph or a reuse distance.

C. Organization of the paper

In Section II we introduce our model assumptions and some useful notation. We present our self-organizing scheme in Section III. In Section IV we present results of extensive simulation experiments, and finally in Section V we will end with some concluding remarks and suggestions for future research.

II. System description

We consider a system with L APs located in an area, and we focus on downlink communication only. Within the considered area, APs provide service to a time-varying set of users. For each AP l, the set of associated users is denoted by $\mathcal{I}_l$.

We assume that a sub-band is the smallest non-divisible slice of spectrum that can be allocated to an AP. The set of frequency sub-bands is given by $\mathcal{F}$, with $|\mathcal{F}| = F$. An AP can transmit on multiple sub-bands, but on at most $F_{\max} \leq F$ at any given time due to antenna and power restrictions. An AP uses all its allocated sub-bands to serve users by applying a proportional fair scheduling policy, as is also common in current 4G base stations and will remain so in 5G. The rate at

which users are served depends on their experienced signal to interference plus noise ratio (SINR) values on the sub-bands they are being served on. Throughout service, a user may be served on multiple sub-bands simultaneously, but it is served by only one AP. We do not account for fast fading and consider only average SINR values.

An AP l transmits on each allocated sub-band f with equal and fixed power. A user i does not receive full power of the AP but rather receives some factor of the power due to path loss factors, for example distance. The Shannon formula implies that user i can receive a maximum communication rate $R(i,l,f)$ (in bits per second) on sub-band f from AP l equal to

$$R(i,l,f) = w \log\left(1 + SINR(i,l,f)\right), \tag{1}$$

where w is the fixed bandwidth of sub-band f, and where $SINR(i,l,f)$ is the signal-to-interference-plus-noise ratio that user i experiences on sub-band f when served by AP l. We assume that if $SINR(i,l,f)$ is below a certain minimum value $SINR_{\text{min}}$, then user i which is in service at AP l cannot and/or will not be served on sub-band f.

The APs are operated through Radio-over-Fiber. That means that at the AP site there is only a simple remote radio head (RRH), and all AP intelligence is located at a centralized entity. This has the advantage that a lot of information, e.g. the current sub-band allocation, is known for the entire system, and can be used in the dynamic operation of the network.

III. DyCRA: A self-organizing scheme

In this section we present the Dynamic Cost/Reward based Allocation scheme (DyCRA scheme), which takes two types of decisions: (i) when should an AP acquire/release resources, and (ii) which resources to acquire/release. For (i) we use the SLI-algorithm [8], which exploits load proxies for the APs. For (ii), we will introduce cost and reward functions, and use these functions to decide (a) if an AP is allowed to acquire a sub-band, and if so, (b) *which* sub-band should be acquired.

A. The SLI-algorithm

The Single Load Interval (SLI) algorithm [8] determines when an AP should (attempt to) acquire or release sub-bands. In short, the concept is as follows: choose an interval $[\rho_{\min}, \rho_{\max}]$. For each AP l, determine a load estimate $\rho_l(t_k)$ at decision time t_k and attempt acquisition or release a sub-band if the load estimate is above $\rho_{\max}$ or below $\rho_{\min}$ respectively.

The idea is to keep the AP loads within the control interval $[\rho_{\min}, \rho_{\max}]$. The upper bound of the interval protects against overloaded APs. The lower bound of the interval frees capacity when possible. The decision times are specified by deterministic intervals, and occur with a frequency of ν_{load}. The $\sigma_l(t_k)$ is defined as the percentage of time that AP l was busy in the time interval $[t_{k-1}, t_k]$. Then, the load estimate $\rho_l(t_k)$ of AP l is determined by a moving average: $\rho_l(t_k) = (1-\varepsilon)\rho_l(t_{k-1}) + \varepsilon(\sigma_l(t_k))$, where $\varepsilon > 0$ determines the magnitude of the updates and is typically small.

The eventual choice of the parameters ν_{load} and ε should ensure that the load estimates are not too sensitive to the temporal load variations, but that systematic changes in the underlying load parameters are detected sufficiently fast.

Remark III.1. *The SLI-algorithm is designed to deal with scenarios where there is just enough capacity to deal with all traffic, as it will try to operate the APs such that their loads are in the interval $[\rho_{min}, \rho_{max}]$. In scenarios with an abundance of capacity, the SLI-algorithm will leave capacity unused, while in scenarios with a significant shortage of capacity it is impossible to sustain the traffic demands anyhow.*

B. DyCRA: Sub-band acquisition

Governed by the SLI-algorithm, APs will attempt sub-band acquisition based on an acquisition cost and an acquisition reward. If the reward outweighs the cost, the AP will acquire a sub-band. The acquisition reward should reflect how valuable a specific sub-band is to an AP. The acquisition costs should reflect the impact it would have on the rest of the system if that sub-band were acquired.

We base our construction of the cost and reward functions on user experienced SINR values. We assume that the users in service at a specific AP only report SINR measurements of the sub-bands that the AP has in use, as we cannot expect users to listen to all potential sub-bands and report all SINR measurements of these sub-bands. However, we do assume that the users receive pilot signals from the APs on a separate sub-band, and these pilot signals are such that they are representative for the signal propagations that users can expect from the APs. If the pilot signal is received with sufficient strength, a user can identify from which AP it originated and deduce the received signal strength from that AP.

1) Acquisition value: Suppose that the SLI-algorithm indicates that AP l should attempt to acquire a sub-band, and consider sub-band f. We will now define the acquisition reward of sub-band f for AP l. Using the pilot signals, suppose we know for all users $i \in \mathcal{I}_l$ the SINR, received from AP l on sub-band f: $SINR(i,l,f)$. Under the current sub-band allocation, user i receives average service rate $R(i,l)$ (in bits per second). If AP l acquired sub-band f, the least profit (in terms of bits per second) that user i would gain is

$$R^+(i,l,f) = w \log_2(1 + SINR(i,l,f))/I_l, \tag{2}$$

where $I_l = |\mathcal{I}_l|$. This follows from the Shannon rate limit, assuming that the sub-band is shared equally among all users. The rate increase (2) is a lower bound, since once the scheduler is going to include the new sub-band in its scheduling policy, users may benefit more. Define the benefit factor $\zeta_i(l,f)$ for user i of AP l acquiring sub-band f as

$$\zeta_i(l,f) = \frac{R^+(i,l,f) + R(i,l)}{R(i,l)}. \tag{3}$$

Then we define the acquisition reward $v(l,f)$ for AP l of sub-band f by

$$v(l,f) = \frac{1}{I_l} \sum_{i \in \mathcal{I}_l} \zeta_i(l,f). \tag{4}$$

The acquisition reward is the ratio between the (minimum) rate a user would experience once AP l acquires sub-band f, and its current rate, averaged over all users. In other words, it is an average of ratios of improvement. When $v(l,f)=1$, we expect no significant improvement, but we typically expect $v(l,f)>1$.

2) Acquisition costs: We will now focus on the cost for AP l of acquiring sub-band f. The cost reflects the impact on the system if AP l acquires sub-band f. Let $P(i,l)$ be the power user i receives from AP l on the pilot signal. Assuming that this indeed is representative for the signal propagation user i experiences from AP l, we can write the SINR as

$$SINR(i,l,f)=\frac{P(i,l)}{\eta+\sum_{l'\in\mathcal{L}(f)}P(i,l')}, \tag{5}$$

where $\mathcal{L}(f)$ is the set of APs that have acquired sub-band f (i.e. sub-band f has been allocated to those APs). Note that AP $l\notin\mathcal{L}(f)$, since we assumed f is a candidate sub-band for acquisition at AP l.

First let us consider the impact for a user i in service at AP $\hat{l}\neq l$, if AP l acquires sub-band f, where $\hat{l}\in\mathcal{L}(f)$. Before acquisition, user i experiences $SINR(i,\hat{l},f)$. After acquisition of sub-band f by AP l, user i will experience an SINR of

$$SINR_l^+(i,\hat{l},f)=\frac{P(i,\hat{l})}{P(i,l)+\eta+\sum_{l'\in\mathcal{L}(f)\setminus\{l\}}P(i,l')} \tag{6}$$

$$=\left(\frac{1}{SINR(i,\hat{l},f)}+\frac{P(i,l)}{P(i,\hat{l})}\right)^{-1}. \tag{7}$$

Let the impact ratio for a user i in service at AP $\hat{l}$ in terms of rate, after AP l acquires sub-band f, be given by

$$\psi_i(\hat{l},l,f)=\frac{\log_2\left(1+SINR_l^+(i,\hat{l},f)\right)}{\log_2\left(1+SINR(i,\hat{l},f)\right)}. \tag{8}$$

We define the *impact of acquisition* $\gamma_{\hat{l}}(l,f)$ of sub-band f by AP l on AP $\hat{l}$ as

$$\gamma_{\hat{l}}(l,f)=\begin{cases}\frac{1}{I_{\hat{l}}}\sum_{i\in\mathcal{I}_{\hat{l}}}\psi_i(\hat{l},l,f) & \text{if } \hat{l}\in\mathcal{L}(f)\\ 1 & \text{otherwise}\end{cases} \tag{9}$$

The impact of acquisition (9) will typically have a value in $[0,1]$, and the closer to one, the less impact the acquisition of sub-band f by AP l has on AP $\hat{l}$. The above-described impact represents an impact per AP, but we need the impact on the entire system. In line with the previous definitions, the cost of acquisition $c(l,f)$ (the impact on the system) will be the average of the impacts per AP:

$$c(l,f)=\frac{1}{L-1}\sum_{\hat{l}\in\mathcal{L}\setminus\{l\}}\gamma_{\hat{l}}(l,f). \tag{10}$$

3) Acquisition rule: Considering the acquisition cost and reward as defined above we set a threshold Δ_{Acq}, and AP l can acquire sub-band f if and only if

$$v(l,f)\cdot c(l,f)\geq\Delta_{\text{Acq}}, \tag{11}$$

i.e. when the reward outweighs the cost by a certain threshold. If there are multiple candidate sub-bands that satisfy (11), the AP acquires the sub-band with the highest value for $v(l,f)\cdot c(l,f)$, breaking ties at random.

C. DyCRA: Sub-band release

When the SLI-algorithm signals that an AP has to release a sub-band, we wish to release the sub-band which brings relatively low capacity to the AP. Hence, AP l releases the sub-band f which minimizes $v(l,f)$ among all sub-bands allocated to AP l, where we compute $v(l,f)$ as if AP l had not acquired sub-band f yet.

D. Choosing the acquisition threshold

The threshold Δ_{Acq} has a similar influence as the interference radius in an interference graph model: setting the threshold too high results in no re-use of sub-bands across the entire system.

Lemma III.1. *Assume that the rates $R^+(i,l,f)$ and $R(i,l)$ are lower and upper bounded by R_{min} and R_{max} respectively. Then there exists a threshold Δ^*_{Acq} such that*

$$v(l,f)\cdot c(l,f)\leq\Delta^*_{Acq}, \tag{12}$$

for all APs $l\in\mathcal{L}$, and al sub-bands $f\in\mathcal{F}$.

Proof. By (7) we have $SINR_l^+(i,\hat{l},f)\leq SINR(i,\hat{l},f)$. It then follows that $c(l,f)\leq 1$, which means we only still have to upper bound $v(l,f)$. To do so, first consider $(R^+(i,l,f)+R(i,l))/R(i,l)$, which is maximized for $R^+(i,l,f)=R_{\max}$ and $R(i,l)=R_{\min}$. That implies

$$v(l,f)=\frac{1}{I_l}\sum_{i\in\mathcal{I}_l}\frac{R^+(i,l,f)+R(i,l)}{R(i,l)} \tag{13}$$

$$\leq\frac{1}{I_l}\sum_{i\in\mathcal{I}_l}\frac{R_{\max}+R_{\min}}{R_{\min}}=1+\frac{R_{\max}}{R_{\min}}. \tag{14}$$

Hence, $v(l,f)\cdot c(l,f)\leq\Delta^*_{\text{Acq}}=1+R_{\max}/R_{\min}$. □

Lemma III.1 indicates that if the threshold is too high, i.e. $\Delta_{\text{Acq}}\geq\Delta^*_{\text{Acq}}$, sub-bands will never be reused, not even over larger distances. On the other hand, if the threshold is too low (say 0 in the extreme case) then sub-bands will be re-used too often, resulting in strong interference conditions. In Section IV we will present simulation results for different values of Δ_{Acq}.

IV. Numerical results

In this section we present various results of numerical experiments we conducted to gain insight in the performance of the DyCRA scheme. We considered an area of $1000m\times500m$ with 10 APs where users appear uniformly at random according to a two-dimensional Poisson process with rate $\nu=10$ users per second. The file sizes of users are independent

and exponentially distributed with a mean of $\mu = 2.5\,MB$. These chosen arrival and service distributions are not essential for the DyCRA scheme to operate, but are primarily used for convenience in the simulations. The decision times of the SLI-algorithm (see Section III-A) occur with frequency $\nu_{\text{load}} = 0.1/s$, i.e. once every 10 seconds. For the AP positions we used two configurations that were generated uniformly at random, see Figures 1 and 2 (interference graphs are plotted as explained in Section IV-A).

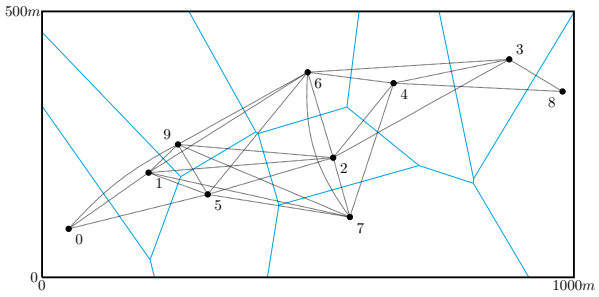

Fig. 1. Configuration 1, with an interference graph based on $r = 400m$.

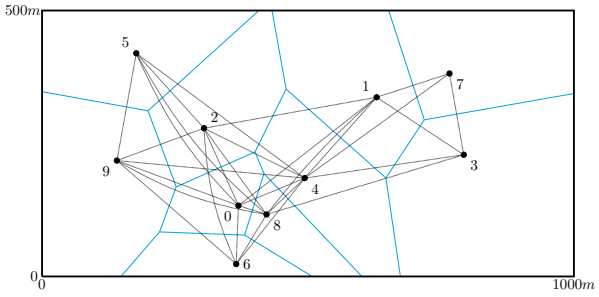

Fig. 2. Configuration 2, with an interference graph based on $r = 400m$.

Users are assigned to the AP that provides them with the strongest signal. In free space that means that the areas in which users are assigned to the same AP are Voronoi cells, with the AP acting as cell center. These Voronoi cells are also drawn in Figures 1 and 2, in order to give insight in the approximate offered traffic that an AP experiences.

Each AP transmits with equal power of $24\,dBm$ (on each sub-band that it transmits on). The signal propagation and path loss follows the 3GGP urban micro model defined in 3GPP 36.814 v9.0.0, where the path loss (in dB) from AP l to user i is given by $PL(i,l) = 140.7 + 36.7\log_{10}(d(i,l))/1000)$, and $d(i,l)$ is the distance in meters between user i and AP l. Furthermore, each sub-band has a bandwidth of $180\,kHz$ (similar to the bandwidth of an LTE resource block), and we assume a thermal noise of $-174\,dBm/Hz$.

The number of users that can be in service at an AP simultaneously is limited by 100 users. If there are 100 users in service and a new user initiates a connection, then that user will be denied service, and leave the system directly without receiving service. One way to measure the performance of a system is to consider the number - or fraction - of service denials. A lower fraction of service denials implies a more efficient resource allocation.

Based on the two configurations shown in Figures 1 and 2, we consider three different scenarios. Scenario 1 and 2 have APs positioned as in Configuration 1. Scenario 3 has APs positioned as in Configuration 2. Scenario 1 has all other system parameters as described above, but Scenarios 2 and 3 have additional user arrivals in the form of a non-stationary hotspot. The hotspot is a $200m \times 100m$ area, and moves over time. It starts with its south-west corner at $(200, 100)$, then it moves to $(400, 100)$, $(600, 100)$, back to $(400, 100)$ and finally returns to $(200, 100)$, after which this pattern repeats. The hotspot has a relative arrival rate of 10 times the normal arrival rate.

A. Benchmark systems

We compare the DyCRA scheme against two benchmark systems which both operate under a static sub-band allocation. The first benchmark system allocates to each AP all sub-bands, and we refer to this system as the full re-use (FR) system.

The second benchmark system applies a sub-band allocation based on an interference (or conflict) graph, and we will refer to this as the interference graph (IG) system. The IG benchmark is constructed in three steps, described in the following three paragraphs.

First, given an interference radius r, we construct an interference graph on the APs by connecting two APs with an edge if the distance between them is less than r meters. When two APs are connected by an edge, they can not use the same sub-band. Two examples of interference graphs have been plotted in Figures 1 and 2.

In the second step, we use a discretization of the area into unit squares to estimate the offered traffic at each AP as follows. For each unit square, we determine the expected offered traffic in bits per unit time. In a uniform setting such as Scenario 1, this is $\nu\mu/(1000 \cdot 500)$. For Scenarios 2 and 3, the HotSpot zone has to be averaged in time over the unit squares it is covering. The offered traffic from the unit square is then assigned to the AP from which the received signal, received at the center of the unit square, is strongest. To obtain a notion of demand at the AP, the assigned offered traffic is divided by the rate at which the AP can serve the center of the unit square, where the rate is in turn obtained by applying the Shannon rate formula (1). The resulting demand for an AP can then be interpreted as the number of sub-bands that the AP needs to sustain the offered traffic, i.e. to serve all traffic offered to the AP per unit time in expectation. For the interference conditions that influence the rates following from (1) we use the following approximation. For an AP l, we assume that no other AP within radius r is using the same sub-band, but all other APs are transmitting on the same sub-band. This assumption gives an upper bound on the interference that we can expect in an interference graph model in free space.

Third and last, we apply a graph colouring heuristic that finds a sub-band allocation that respects the interference graph, and allocates to each AP the number of sub-bands as described

in the previous step. The output of the heuristic is thus a sub-band allocation, which implies a set of sub-bands $\mathcal{F}$ needed to realize this allocation.

To obtain fair comparisons, we first initialize the IG systems as described above, and provide all other systems with the same set (or rather the same number) of sub-bands $\mathcal{F}$.

Remark IV.1. *Note that we only use the interference graph to create a good sub-band allocation as a benchmark. Once the allocation has been found, we "forget" the interference graph and all service rates of users are again determined by the SINR model described in Section II. In particular, when $r = 0$, the IG and FR systems are the same.*

Remark IV.2. *In simulations, we need to solve a proportional fair scheduling problem in order to know at what rates users are served, at what time the next user will finish service, and hence to determine the end of the time interval in which these rates are valid. The simulation time is primarily influenced by the time to solve these scheduling problems.*

B. Performance

We consider two performance indicators: service denials (in percentages, as explained above) and experienced throughput in $Mbit/s$. The user-experienced throughput is defined as the size (in bits) of the file that the user downloaded, divided by the time it took the system to deliver the file to the user.

All results are based on a sequence consisting of 50 000 users, where the sequence was randomly generated, but the same sequence was presented to each system to gain useful comparisons. In Table I we present the percentage of service denials for scenarios with $r = 300m$ and $r = 400m$. Due to long simulation times, we do not present results for Scenario 3 with $r = 300$.

TABLE I
PERCENTAGE OF SERVICE DENIALS.

System	$r = 300m$		$r = 400m$		
	S1	S2	S1	S2	S3
$\Delta_{\text{Acq}} = 0.9$	1.61	3.82	1.55	3.80	5.75
$\Delta_{\text{Acq}} = 1.0$	1.64	3.94	1.60	3.71	4.75
$\Delta_{\text{Acq}} = 1.25$	1.76	3.92	1.66	3.71	9.13
$\Delta_{\text{Acq}} = 1.5$	1.76	3.92	1.66	3.71	9.83
IG	0	4.01	0	4.30	5.03
FR	2.99	7.30	0.85	6.34	13.39

Remark IV.3. *The difference between $r = 300$ and $r = 400$ is in the initialization of the IG benchmark, which also determines the eventual number of available sub-bands in the system. In Scenario 2 with $r = 300$, the number of available sub-bands was 194, while with $r = 400$ it was 212, and in Scenario 3 with $r = 400$ it was 141.*

Remark IV.4. *We do not present results for $\Delta_{Acq} < 0.9$, as we observed that this was (in many cases) equivalent to the FR benchmark. We also do not present results for $\Delta_{Acq} > 1.5$, since in most cases (that we considered) no frequencies were re-used by the DyCRA scheme (in line with Lemma III.1 and existence of Δ^*_{Acq}).*

First notice that the choice for Δ_{Acq} may depend on the AP positions. In Scenario 2, all considered values result in better service denial performance than the IG benchmark, while in Scenario 3 only $\Delta_{\text{Acq}} = 1.0$ performs better. We observe that with dynamically changing demand conditions such as in Scenario 2 or 3, the DyCRA scheme can serve more users than both the benchmarks (assuming $\Delta_{\text{Acq}} = 1.0$ in Scenario 3). In statistically stable demand conditions, like Scenario 1, we do not expect the DyCRA scheme to outperform (with respect to service denials) the IG benchmark for the following reason. The DyCRA scheme, through the SLI-algorithm, releases frequencies at APs when the APs experience low load conditions. In other words, the DyCRA scheme operates (or at least tries to operate) the APs at load levels that fall in the interval $[\rho_{\min}, \rho_{\max}]$. Hence, in times when the loads happen to be low for the moment, the scheme releases frequencies, meaning that the users that are in service at the AP at that moment do not reap the full benefits of the momentarily low load conditions. The IG scheme on the other hand has a fixed number of sub-bands allocated to each AP at all times, such that users that are in service at APs with temporarily low load conditions get a lot of resources. The DyCRA scheme is designed to relieve APs that are facing high demands, and therefore outperforms the IG system in scenarios where APs are suddenly faced with a demand they cannot meet.

Looking at Table I it looks like the DyCRA scheme is only doing a marginally better job than the IG system. However, not only is the DyCRA scheme realizing lower service denial percentages, it also realizes a better user-experienced throughput, as can be seen in Figures 3 to 5, where we plotted the cumulative fraction of users as a function of their received throughput, for several values of Δ_{Acq} (cyan curves) and the two benchmarks IG and FR (red and black curves respectively).

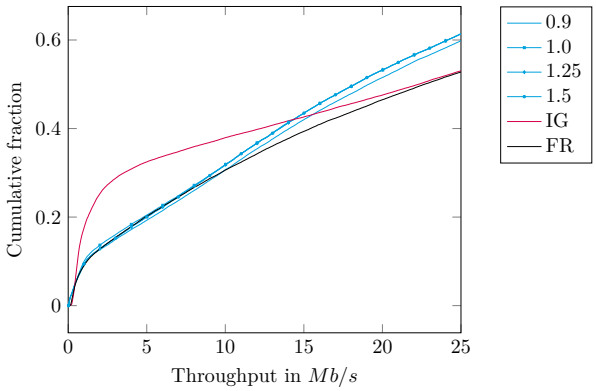

Fig. 3. User-experienced throughput for Scenario 2, with $r = 300m$.

In Figures 3 and 4, if for a specific throughput value T one curve is lower than the other, it means that in the first case more users experienced a higher throughput than T than the users of the system belonging to the second curve.

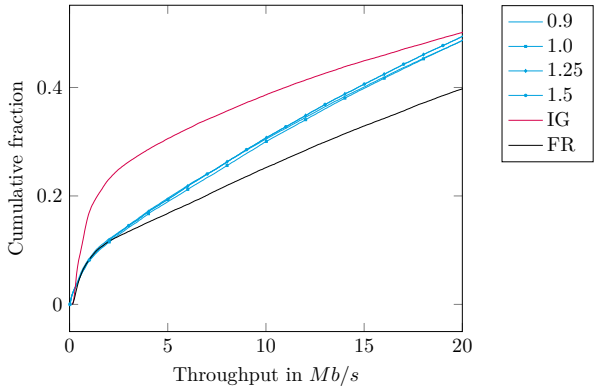

Fig. 4. User-experienced throughput for Scenario 2, with $r = 400m$.

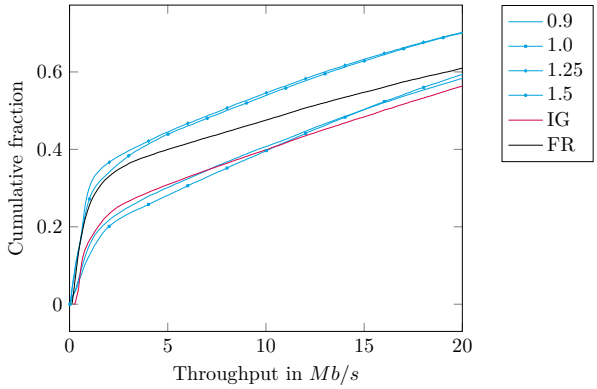

Fig. 5. User-experienced throughput for Scenario 3, with $r = 400m$.

Hence, a lower curve implies a better performance. We observe that the DyCRA scheme outperforms the IG system in the lower throughput region, which is typically where we expect users from highly loaded APs. This confirms that the DyCRA scheme indeed moves frequencies from APs with lower load towards APs that face high traffic demands.

C. Stability

Now that we have established the effectiveness of the DyCRA scheme when it comes to matching frequencies with demand, we will consider "stability". We do this by looking at the frequency allocations provided by the DyCRA scheme, and looking at the load proxies for APs.

To be able to meet capacity with demand is a desirable property, but it is undesirable to completely change the frequency allocation all the time. In Figures 6 and 7 we plotted the number of changes in the frequency allocation at each decision epoch of the DyCRA scheme, for Scenarios 1 and 2, and for different values of Δ_{Acq}. The number of changes is given by the sum of the number of acquisitions and the number of releases, over all APs in the system. For example, when one AP acquired a frequency, and one AP released a frequency (and all other APs do not change their frequency set) the number of changes is two.

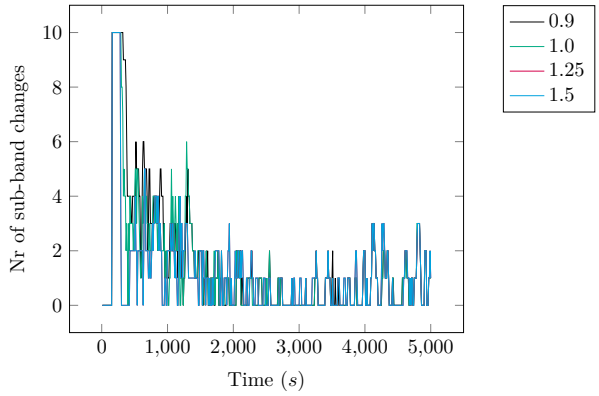

Fig. 6. Number of sub-band acquisitions and releases per update moment for Scenario 1, with IG based on $r = 400m$.

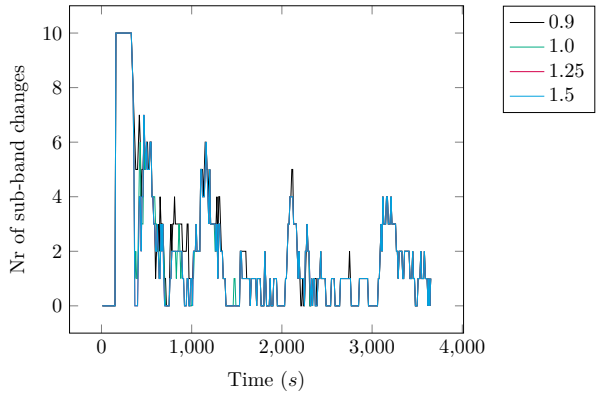

Fig. 7. Number of sub-band acquisitions and releases per update moment for Scenario 2, with IG based on $r = 400m$.

Figure 6 shows that the DyCRA scheme provides quite stable frequency allocations in a setting where the demands are stable: after some warm up-period it mostly performs none to two changes. In Figure 7, with changing demand characteristics, we see that the DyCRA scheme has more peaks in the number of changes, following to the (structural) changes in demand.

In Figure 8 we plotted the load proxies $\rho_l(t_k)$ for APs 1, 5 and 7, whose cells overlap with the HotSpot at various times.

We can clearly observe the peaks in loads due to the moving HotSpot, and we can also observe that the DyCRA scheme tries to operate the APs in the (estimated) load regime $[0.5, 0.8]$. Figure 8 shows that the DyCRA scheme actively reacts to the changing demand conditions and effectively pushes the AP loads into the desired range.

V. Conclusion

In this paper we presented a dynamic, load aware and self-adapting frequency allocation scheme, the DyCRA scheme, specifically designed for dense cellular networks. The DyCRA scheme relies on load measurements at APs and SINR measurements from users, to make favourable changes in the

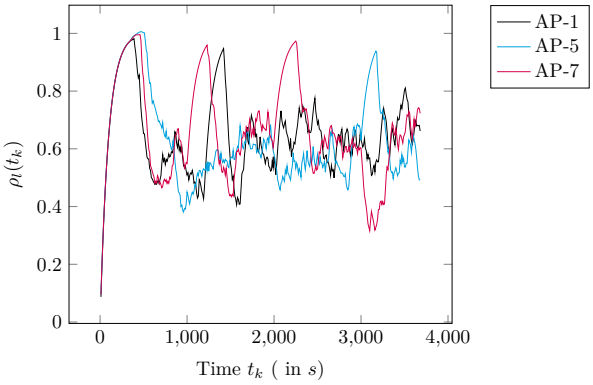

Fig. 8. Load proxies for AP 1 in Scenario 2 with $r = 300$.

frequency allocations, moving surplus capacity towards APs that face a (too) high demand. We introduced cost and reward functions for assigning a frequency to an AP, and only assign a new frequency to an AP if the reward outweighs the cost by a tunable parameter Δ_{Acq}.

Extensive simulations demonstrated the effectiveness of the DyCRA scheme to dynamically match capacity with demand, allowing service to more users that would otherwise have been denied service. The DyCRA scheme realized both less service denials and better throughputs for users, compared to a benchmark allocation that is optimized to deal with the average traffic demands. We furthermore demonstrated that even though the DyCRA scheme dynamically adapts the frequency allocation, it does so with only a small number of changes, leading to a quite stable frequency allocation.

To the best of our knowledge, it is the first self-adapting frequency allocation scheme specifically designed for all challenges that arise in dense cellular networks. We wish to stress the fact that the DyCRA scheme realizes good performance without the need of prior optimization, field measurements, or the knowledge of an interference graph. The DyCRA scheme could be improved by further tuning the algorithm parameters and studying more advanced scenarios. Also, for specific properties of frequency allocations it may be necessary to develop other cost/reward trade-offs, opening up interesting directions for future research.

VI. Acknowledgements

We thank Phil Whiting for his insightful remarks and useful suggestions over the course of this project.

References

[1] 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 Comm. Mag.*, vol. 52, no. 2, pp. 82–89, 2014.

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

[3] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What will 5G be?" *IEEE J. Sel. Areas Comm.*, vol. 32, no. 6, pp. 1065–1082, 2014.

[4] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[5] I. Hwang, B. Song, and S. S. Soliman, "A holistic view on hyper-dense heterogeneous and small cell networks," *IEEE Comm. Mag.*, vol. 51, no. 6, pp. 20–27, 2013.

[6] Q. Ye, B. Rong, Y. Chen, M. Al-Shalash, C. Caramanis, and J. G. Andrews, "User association for load balancing in heterogeneous cellular networks," *IEEE Trans. Wireless Comm.*, vol. 12, no. 6, pp. 2706–2716, 2013.

[7] J. G. Andrews, S. Singh, Q. Ye, X. Lin, and H. S. Dhillon, "An overview of load balancing in HetNets: Old myths and open problems," *IEEE Wir. Comm. Mag.*, vol. 21, no. 2, pp. 18–25, 2014.

[8] B. Post, S. C. Borst, and A. M. J. Koonen, "Dynamic resource allocation in radio-over-fiber enabled dense cellular networks," in *Proc. WiOpt*. IEEE, 2018, pp. 1–8.

[9] A. Tall, Z. Altman, and E. Altman, "Self organizing strategies for enhanced ICIC (eICIC)," in *Proc. WiOpt'12*. IEEE, 2014, pp. 318–325.

[10] M. Simsek, M. Bennis, and A. Czylwik, "Dynamic inter-cell interference coordination in HetNets: A reinforcement learning approach," in *GLOBECOM*. IEEE, 2012, pp. 5446–5450.

[11] S. Park and S. Bahk, "Dynamic inter-cell interference avoidance in self-organizing femtocell networks," in *Proc. ICC*. IEEE, 2011, pp. 1–5.

[12] S. H. Ali and V. C. M. Leung, "Dynamic frequency allocation in fractional frequency reused OFDMA networks," *IEEE Trans. Wireless Comm.*, vol. 8, no. 8, 2009.

[13] A. L. Stolyar and H. Viswanathan, "Self-organizing dynamic fractional frequency reuse in OFDMA systems," in *Proc. INFOCOM*. IEEE, 2008, pp. 691–699.

[14] M. M. Halldórsson and G. Kortsarz, *Multicoloring: Problems and Techniques*. Springer Berlin Heidelberg, 2004, pp. 25–41.

[15] M. M. Halldórsson and T. Tonoyan, "How well can graphs represent wireless interference?" in *Proc. 47th ACM Symp. Th. of Comp.* ACM, 2015, pp. 635–644.

[16] X. Zhou, Z. Zhang, G. Wang, X. Yu, B. Y. Zhao, and H. Zheng, "Practical conflict graphs for dynamic spectrum distribution," in *SIGMETRICS Perf. Eval. Rev.*, vol. 41. ACM, 2013, pp. 5–16.

[17] J. Yang, S. C. Draper, and R. Nowak, "Learning the interference graph of a wireless network," *IEEE Trans. Sig. Inf. Proc. Netw.*, vol. 3, no. 3, pp. 631–646, 2017.

[18] S. Uygungelen, G. Auer, and Z. Bharucha, "Graph-based dynamic frequency reuse in femtocell networks," in *VTC*. IEEE, 2011, pp. 1–6.

[19] A. Ule and R. J. Boucherie, "Adaptive dynamic capacity borrowing, in road-covering mobile networks," *Res. Alloc. Next Gen. Wireless Netw.*, vol. 5, p. 67, 2006.

[20] S. Nanda and D. J. Goodman, "Dynamic resource acquisition: distributed carrier allocation for TDMA cellular systems," in *GLOBECOM*, vol. 2. IEEE, 1991, pp. 883–889.