Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

# Buddies, not Enemies: Fairness and Performance in Cellular Offloading

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz
Multimedia Communications Lab (KOM), Technische Universität Darmstadt, Germany
{nils.richerzhagen|bjoern.richerzhagen|michael.walter|dominik.stingl|ralf.steinmetz}@kom.tu-darmstadt.de

***Abstract*—Recently, a number of offloading approaches have been proposed to reduce the burden on the cellular infrastructure, especially during peak hours. Ranging from pure data offloading concepts using local caches to fully-fledged services operating on ad hoc networks, these approaches are mostly tuned towards their performance. Consequently, cache hit ratios, achieved throughput, or end-to-end latencies, are prominent evaluation metrics. However, when utilizing users' resources to improve the performance of a shared service, the fairness with respect to the resources contributed by individual users should be considered as well. Resource consumption, e.g., the battery lifetime, can vary significantly between individual users, depending on their contributions. This effect can have a significant impact on user acceptance and, thus, the usability of the overall system. In this paper, we examine the trade-off between overall system performance and fairness w.r.t. resource utilization for a state-of-the-art monitoring service relying on offloading. To bridge the identified gap between performance and fairness, we propose a number of protocol adjustments to increase the system's fairness. Through an extensive simulation study we show that the proposed mechanisms lead to improvements in the overall achieved fairness of up to 60%, with below 4% degradation in service quality.**

## I. Introduction

Nowadays, the amount of mobile devices – such as smartphones, tablets, and laptops – experiences vigorous growth, leading to increased mobile data traffic. Particularly, Cisco estimates that the worldwide mobile data traffic will increase tenfold until 2019 – three times faster than fixed IP traffic [1]. The load on mobile networks varies significantly throughout the day, or on specific occasions, such as large public events. In such situations, the capacity is easily exceeded by the demand, leading to unsatisfied users [2]. To reduce the load on the cellular network, offloading strategies have been proposed by the research community. Offloading utilizes idle resources of mobile consumers to improve the overall system performance. The most commonly used resource is the direct connectivity between nearby devices via technologies such as Bluetooth and Wi-Fi Direct. Direct communication is, for example, used to distribute locally relevant data to interested users [3], or to obtain insights into the local network state [4]. Depending on the target application, mobile users are often assigned different roles, e.g., relaying data via their cellular connection or storing data on behalf of other users. The assigned role determines which resources are consumed on the respective device.

As most systems are optimized with respect to the overall performance, imbalances regarding the contributions of individual users may occur. These imbalances can have a significant impact on the acceptance of the system, if, for example, the batteries of a set of users are drained faster than others. To this end, fairness measures have been proposed to examine these imbalances. However, the trade-off between achievable fairness and the overall system performance is mostly neglected in current offloading systems. This work examines the trade-off between performance and fairness using the example of the state-of-the-art adaptive monitoring service CRATER [4]. Several protocol adjustments are proposed based on an in-depth analysis of the fairness characteristics of the service. The resulting performance characteristics, as well as the achieved fairness are assessed in an extensive simulation study using the Simonstrator platform [5]. Our results reveal that fairness can be increased significantly (up to 60%) with only minor degradations (4%) in service quality.

The contributions of this paper are the following:

- The simulation based analysis of a state-of-the-art distributed monitoring system with respect to the gap between performance and fairness.
- A set of protocol adjustments and role assignment procedures to increase the fairness of the system.
- An in-depth evaluation of the resulting trade-off between performance and fairness.

The remainder of the Paper is structured as follows: At the outset, Section II provides an overview of the adaptive monitoring system CRATER. Section III discusses the concept of fairness in distributed systems. Quantitative, as well as qualitative fairness measures are reviewed. Section IV deals with the related work on fairness optimizations with respect to one or more fairness metrics in distributed systems. Section V outlines the modifications that were applied in order to improve the fairness, as well as the general performance of the modified system. The evaluation addressing the gap between fairness and performance in the adaptive monitoring system is shown in Section VI. Section VII concludes this work.

## II. Adaptive Monitoring in Dynamic Networks

CRATER embodies an adaptive monitoring solution for mobile environments that operates independent of other applications and services and assembles information on the current state of the network and the participating nodes. By adapting to environmental changes, especially with respect to user mobility and density, the system provides accurate estimates of the overall state [4]. To obtain a complete view

978-1-5090-2185-7/16/$31.00 ©2016 IEEE

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

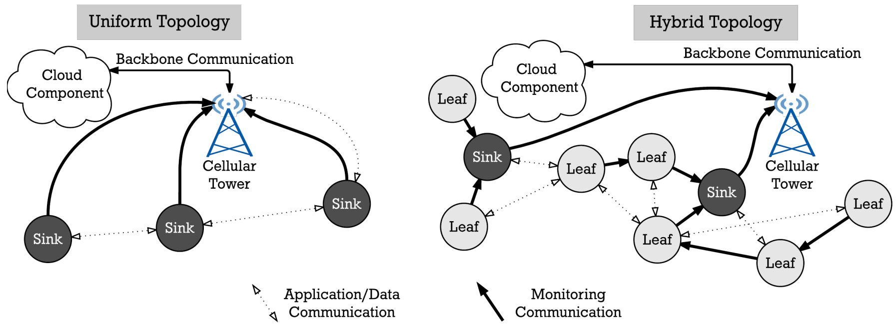


Figure 1: The two topology states of CRATER, which include sinks and leaves, the dedicated management-plane for communication of the monitoring system, and a data-plane for communication of applications and other mechanisms.

on the system, each node in CRATER periodically measures its own local state according to the set of required monitoring attributes and reports it to a central entity, referred to as cloud component. CRATER operates in two different topology states, as illustrated in Figure 1. The cloud component allocates two distinct roles to the nodes depending on the current environmental conditions: the *sink* and the *leaf* role. *Sinks* upload their monitoring information directly via their cellular connection to the cloud entity. *Leaves* do not use the cellular connection to upload their information. Instead, they affiliate to sinks using local ad hoc connectivity, e.g., Bluetooth or Wi-Fi Direct. As long as the server is not overloaded and the cellular network is able to handle the load, nodes directly upload their monitoring information, operating as sinks. In CRATER, this mode of operation is referred to as *uniform* topology. The uniform topology implements the client/server monitoring paradigm, as each node uploads only its own monitoring information. If either the central server is overloaded, or the cellular network is no longer able to handle the load, CRATER switches to a *hybrid* topology. Here, a fraction of the nodes is assigned with the sink role, while all other nodes remain leaves (cf. Figure 1). Sinks continue using their cellular link to the central server, while leaves drop their cellular connection to free additional capacities in the upload network. Consequently, leaves are obliged to choose a sink in their proximity that relays their monitoring data to the cloud component.

Each sink undertakes the task of uploading its own monitoring data to the cloud component, as well as the monitoring data of all its affiliated leaves. It is important to note, that leaves do not only act as a source for monitoring information, but also forward data from nearby other leaves, if necessary. Leaves aim to choose the most stable sink to send their information to – the sink that is the least likely to recede into distance shortly. CRATER consists of mechanisms for (i) maintaining the hybrid topology and (ii) routing the information between leaves and sinks. The topology maintenance mechanism includes the *No Sink Advertising* and the *Sink Advertising* mechanisms. Depending on the current role of a node in the network, either the *No Sink Advertising*, or the *Sink Advertising* mechanism is used. The most important task for leaves is to advertise the need for affiliation to a sink if no adjacent sink is known.

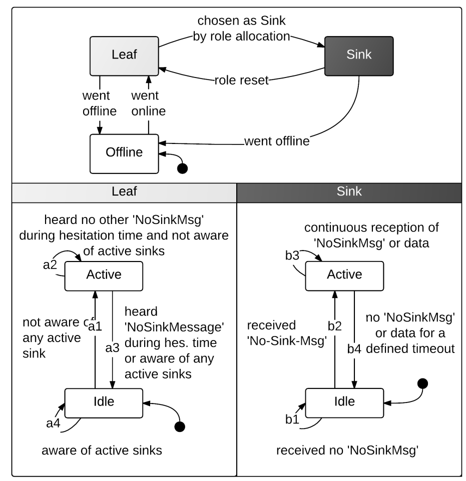


Figure 2: State chart of the main advertising components for topology maintenance in CRATER [4].

To advertise this need, leaves send *NoSinkMessages*. During affiliation, a leaf is in its *active* state, as shown in the state chart of a single node in Figure 2. From the state chart it becomes apparent, that nodes react differently to the aforementioned messages depending on their current role.

As soon as a sink receives a *NoSinkMessage*, it responds with a *SinkAdvertisingMessage*, stating its presence. Once a sink starts sending *SinkAdvertisingMessages*, it switches to its *active* state (cf. Figure 2). Leaves that hear active sinks store the result in an individual sink table. CRATER returns to its uniform topology, as soon as sufficient resources are again available to nodes. In that case, all nodes in the network operate as sinks and upload their monitoring data directly. For an in-depth description of CRATER and the different states of the components, the interested reader is referred to [4].

In the initial design the mechanism relied on a simple sink selection algorithm. Nodes are assigned as sinks until the number of sinks reaches a predefined threshold, corresponding to the capacity of the infrastructure. Sinks then preserve their role for a random time interval between 30 and 60 minutes. This rather static assignment of sinks leads to unfair resource consumption characteristics for the mobile nodes in this

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

cellular offloading scenario. To this end, more sophisticated sink selection algorithms are presented in Section V. These algorithms aim at improving the fairness as defined in the following section. Subsequently, in the course of this work the term *gateway selection* will be used instead of the term sink selection, which is more specific for [4].

## III. Fairness in Distributed Systems

In this section we introduce the characteristics of fairness measures and highlight both (i) qualitative (cf. III-A) and (ii) quantitative (cf. III-B) fairness metrics. At the end the presented fairness metrics are mapped to the metrics used for fairness estimation of eCRATER in Section III-C.

Shi et al. [6] state, that fairness is an *interdisciplinary research topic which is usually related to resource allocation* and that it is often interconnected with studies regarding, e.g., performance. Moreover, the authors state, that unfair allocation of resources in a wireless network can cause *resource starvation, wastage, or redundant allocation* at worst. The authors present three issues that must be taken into consideration for fairness within a wireless scenario: (i) a definition of fairness is required, (ii) a method for measuring whether a system is fair towards individuals must be available, and (iii) an answer to the question how a system can be made fair, must be found. Defining fairness is very difficult to accomplish in a manner such that it would achieve general consent [6]. Still, from existing definitions of fairness, the conclusion can be drawn that it is important to identify which are *equal individuals* of a system, as well as how *equal treatment* can be defined. Finally, Shi et al. claim that fairness can be considered with respect to the overall system – system fairness, but also with respect to individuals – individual fairness.

According to Jain et al. [7], fairness in a distributed system expresses the level of equality of resource distribution among the users. The authors observed that, from the literature, there is no consensus on which resources should be allocated equally. Therefore, Jain et al. believe that rating fairness of a system requires an appropriate allocation metric, as well as a formula that describes the fairness of the allocation by assigning it a quantitative value. There exist *qualitative* and *quantitative* measures for computing the fair allocation of resources in the literature which are presented in the following.

### A. Qualitative Fairness Measures

The most widely known qualitative fairness measures are *Max-Min fairness* [8], [9] and *proportional fairness* [10], [11]. A *max-min fair* algorithm uniformly distributes a shared resource between all nodes. If the resource requirements of a node are fulfilled, it no longer receives additional resource shares. The leftover resources are further distributed among the remaining nodes. This continues as long as there are resources left and resource requirements are not yet satisfied. As soon as all resources have been distributed, no resource share can be increased, without decreasing another one.

A *proportional fair* algorithm tries to maximize the sum of shares proportional to each node's requirements. Proportional fairness is achieved for a resource $x$, if the aggregate $\sum_{i=1}^{n} \frac{x_i^* - x_i}{x_i}$ becomes negative for each other feasible resource allocation $x^*$. With $n$ being the number of nodes, and $x_i$ the resource share at node $i$.

These qualitative measures are often used as a guideline to achieve fair resource allocation, however they do not assign a bound numerical value that describes the fairness level. This is a decisive disadvantage, as to the fact that the lack of such a numerical value aggravates the process of comparing different fairness levels. Thus, for comparison reasons the qualitative fairness measures are not used in this work.

### B. Quantitative Fairness Measures

Contrary to qualitative measures, quantitative measures do not lack a numerical value that describes the fairness level. Examples for quantitative measures are the ratio of minimum to maximum resource shares, Jain's index, or the variance.

The *min-max* [12] ratio has the advantage of not being affected by scale. However, as it only takes the outer most values into account, it does not give a thorough overview of the system's fairness. It is therefore seldom used.

*Jain's fairness index* [7] is a wide known fairness index, which is independent of scale and can be applied to any resource-allocation, or resource-sharing system. Jain's fairness index $FI$ is defined by Equation 1. While $n$ describes the number of users in the system of which each $i^{th}$ user receives an allocation $x_i$ of any kind of resource [7]. In a system of $n$ contending users the proposed index calculates the fair allocation of any kind of resource to the users.

$$FI(X) = \frac{[\sum_{i=1}^{n} x_i]^2}{n * \sum_{i=1}^{n} {x_i}^2} \tag{1}$$

*Jain's* fairness index ranges between 0 and 1. An index value near 0 implies that the fairness in the system decreases, a result of favoring or burdening only a few users. Any value between 0 and 1 asserts the percentage of the analyzed system's fairness – an index of $0.1$ implies that the system is fair to $10\%$ of the users. An index of $1$ implies fair resource-sharing for all users. However, it must be stated that Jain's fairness index does not imply to which extent users get a resource. Especially, regarding resources that have a negative impact on the users, like e.g. battery consumption or load on users, the index might result in a 100% fair share, but the total value of resources spent by each individual user may highlight potential weaknesses of a system. Thus, when considering quantitative fairness metrics, the total value of the shared resource must also be taken into account.

The *variance* $s^2$ is a measure for the distribution of data and assists in finding single outliers. It is defined as follows:

$$s^2 = \frac{1}{n} * \sum_{i=1}^{n} (x_i - \overline{x})^2 \tag{2}$$

Here, $n$ is the number of observed values, $x_1, ..., x_n$ are the actual observed values, and $\overline{x}$ is the arithmetic average of the observed values. Because the variance is not independent of

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

scale, the relative variance $\frac{s^2}{\bar{x}}$ may be used instead. Quantitative measures offer the opportunity to directly compare different levels of fairness. Therefore, to rate the level of fairness of the system under consideration, quantitative fairness measures are used. In particular, *Jain's fairness index* [7], as it is a widely studied index for computing fairness [6]. Moreover, both measures are independent of scale and dimensions. Jain's index is used to rate overall fairness of the system.

$$cdf(x) = P(X \leq x) \tag{3}$$

Cumulative distribution functions (cf. Equation 3) are used to identify the impact of outliers, i.e. unfairly treated single nodes, instead of the relative variance, as they show the valuable information about the distribution of a metric [13].

### C. Analysis of the System Fairness in eCRATER

For fairness estimations in eCRATER the following metrics are used: (i) the time ratio of nodes being in each role (Equation 4), (ii) the time ratios of nodes being active in the roles (Equation 5), and (iii) the amount of time leaves need to re-affiliate to the topology, the *refill delay* (Equation 6).

$$R_{n,\text{sink/leaf}} = \frac{T_{\text{sink/leaf}}}{T_{\text{online}}} \tag{4}$$

$$R_{n,\text{sink/leaf, active}} = \frac{T_{\text{sink/leaf, active}}}{T_{\text{sink/leaf}}} \tag{5}$$

$$T_{n,\text{refill}} = \frac{\sum_{i=1}^{m} t_{i,\text{re}}}{m} \quad \text{with} \quad t_{\text{re}} = t_{\text{fill}} - t_{\text{empty}} \tag{6}$$

The roles of individual nodes are visible in Figure 2. The refill delays for a leaf are given by the time interval from loosing connectivity to a sink (i.e. no active sink known, leafs' sink table is empty) until a leaf hears a SinkAdvertisingMessage (i.e. an active sink known, sink table filled).

## IV. Related Work

The first part of this section investigates studies that aim to evaluate the fairness of distributed systems based on fairness metrics. Other monitoring mechanisms for mobile environments are discussed in the second part with focus on the respective role allocation fairness achieved in the systems.

The fairness of routing cost among peers in a Chord-based [14] P2P overlay is analyzed by Cuevas et al. in [13]. The authors analyze routing fairness within Chord by determining the number of messages routed by each node. This analysis reveals that message routing within Chord is unbalanced – only 60% fair according to Jain's fairness index. Cuevas et al. discovered the prime cause for that is that nodes with larger zones are more likely to be selected as fingers. By modification of the finger placement the authors achieved a more balanced routing load resulting in a higher fairness level of 90 %. For traffic balancing in WSNs, Vergados et al. [15] employ an algorithm that prioritizes data flows based on a delay metric. The authors use Jain's index to prove an increase in end-to-end delay fairness. Similar to [15], Chen et al. [16] aim to improve WSN end-to-end data flow fairness as well. Within their proposed algorithm, an improved fairness is achieved by allowing nodes to individually assign bandwidth to several incoming packet streams. The assigned amount of bandwidth is proportional to the stream's weight at the forwarding node, relative to the sum of all stream-weights at this node. Each of these weights is either directly assigned to a sensor or to the aggregate of the data flows incorporated sensor weights. The authors don't take specific fairness indexes into consideration and utilize only the data rate distribution among nodes as fairness indicator. In [17], Zhu, Hung and Bensaou investigate the lifetime maximization problem and the rate allocation problem, as well as the trade-off between both in a similar environment. In their work, they state the network lifetime as the time until the first node has no more battery power available, i.e. *first node death*. A function that models the dependence of battery drain on allocated data rate is used. Maximizing this function/value yields the longest network lifetime. The other goal, a fair rate allocation, is reached by maximizing an aggregate utility function of all data flows. This can result in a max-min or proportional fair rate allocation, depending on the utility function in place. Their new approach is to combine these two conflicting problems into a single weighted function. Their distributed algorithm iteratively computes rate allocations for each node, approximating the user specified trade-off value between lifetime and fairness in each step.

Approaches that use gateway selection and clustering algorithms to achieve a defined non-functional requirement are presented in the following. DAMON [18] uses an agent-sink topology for distributed data collection in multi-hop networks. Using agents to control the flooding of the network, leads to less overhead and reduces the possibility of collisions. However, the sinks in DAMON are static and pre-defined by the network operator, comparable to a non-changing gateway selection in this work. Differentiating into detailed local and a sparse global network view, as Nanda and Kotz [19] do, is helpful to gain a more precise monitoring result. Nonetheless, the used hierarchy, consisting of static mesh nodes that operate as sinks and mobile nodes. Assuming static nodes that perform the sink task is inappropriate for the case of mobile environments. The small scale evaluation scenario of Mesh-Mon with less than 25 nodes is not appropriate for heterogeneous network scenarios as presented in this work. HMAN by Battat and Kheddouci [20] establishes a three-tiered topology based on weights of nodes. Those weights incorporate factors, such as energy consumption, the distance to other nodes, and the leftover storage capacity. However, the weights used in the approach are pre-defined, which results in a limited applicability. Furthermore, due to the weighted functions, it is very likely that a small quantity of nodes remain in their hierarchical roles over time, which has a significant impact on the fairness as it is similar to a static role allocation. BlockTree [21] is a fully decentralized monitoring approach for MANETs. It establishes a hierarchy, built by location-aware nodes. However, requiring detailed information on the nodes' location, which are provided by additional services, such as GPS, rendering such approaches useless for

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

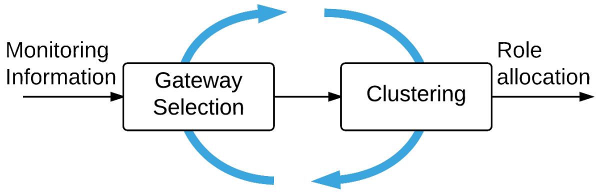


Figure 3: The process of gateway selection and clustering in eCRATER. The sequence is not fixed, as for some algorithms the clustering process may precede.

indoor scenarios or when the localization is not as accurate as needed. Summarizing, all approaches lack the important bridge between resource allocation fairness and system performance, which is examined in this work. Allocating different roles to nodes involves varying resource consumption of the individuals, resulting in an unfair system.

## V. eCRATER: Gateway Selection in an Adaptive Monitoring System using Cellular Offloading

The initial gateway selection in CRATER is based on a simple random gateway selection algorithm (cf. Section II). Due to the long time period sinks stay in their respective role, the sink selection process is not able to adjust the allocation based on the current network state. Similar to DAMON [18] and Mesh-Mon [19], this results in a complete static, to in the best case, very static gateway selection where roles are seldom reassigned. However, as leaves and sinks within CRATER perform different tasks during the course of time, it is evident that unfairness may occur. To circumvent or at least reduce such unfairness prior to its occurrence, we introduce several modifications to the system with more sophisticated gateway selection and clustering algorithms. The respective modifications are detailed in the following by introducing enhanced CRATER (*e*CRATER). The gateway selection and clustering problem can be solved with exact algorithms. However, for allocating $n$ nodes in a network in $k$ clusters one would have to solve the NP-complete problem shown in Equation 7.

$$S_n^{(k)} = S_{n,k} = \frac{1}{k!}\sum_{j=0}^{k}(-1)^{k-j}\binom{k}{j}j^n\,. \tag{7}$$

This results in over $2.75 \cdot 10^{93}$ combinations for a fairly small network of 100 nodes and 10 clusters ($S_{100}^{(10)}$). Solving that problem in a few seconds of time to return a valid solution for the mobile network is not feasible for larger networks. Hence, heuristic approaches are used in this work to solve the given problems within a reasonable time window. The process of role allocation can be split into (i) *gateway selection* and (ii) *clustering*. Both internal components may be exchanged, i.e. the gateway selection is performed first, followed by clustering the nodes to the gateways, or vice-versa (cf. Figure 3). However, eCRATER uses only the resulting gateways. The chosen gateways are then allocated the sink role by the cloud component. The monitoring information used by the role allocation algorithms is obtained via global

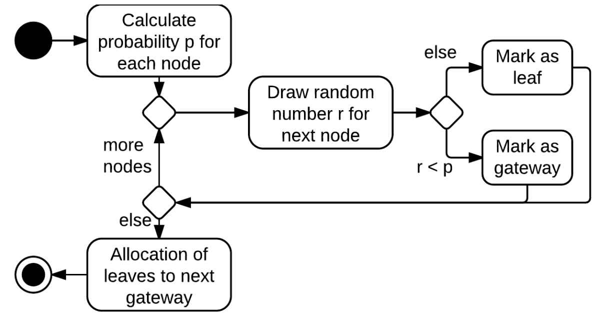


Figure 4: Process diagram of SEC. **S**tochastic gateway selection with a number of gateways following an **e**xpectation value. Final **c**lustering of non gateway nodes to the gateways.

knowledge. This abstraction is necessary in first place as the focus lies on examining the potential performance and fairness gains through more sophisticated gateway selection.

In the following, we propose the architecture of eCRATER. From existing literature [22]–[24] it becomes apparent that the selection of gateways can be deterministic or stochastic. The cluster assignment process can be performed successively for each cluster or for all at once. The combination of both results in four different schemes: (i) SEC, (ii) D1C, (iii) CD, and (iv) CS. SEC and D1C describe the allocation process with initial gateway selection and subsequent clustering. SEC stands for **s**tochastic gateway selection with the number of gateways following an **e**xpectation value and successive **c**lustering of non gateway nodes to the chosen gateways. D1C describes a **d**eterministic gateway selection followed by a successive selection of **one** gateway and the associated **c**luster. CD and CS describe the **c**lustering followed by choosing a number of gateways either **d**eterministic, or **s**tochastic.

The process of SEC, which is shown in Figure 4, is partially based on ALEACH [22]. At the beginning each node is assigned a probability $p_i$ that is calculated via a given valuation method. A random number $r_i$ is drawn in the next step for each node. If the value of the random number lies below the assigned probability ($r_i < p_i$), the node is marked as gateway. Otherwise, the node is marked as leaf. Finally, the clusters are calculated with a best-fit algorithm. The best-fit algorithm used in this work assigns a leaf to the gateway that has the shortest euclidean distance. However, other algorithms such as shortest path or highest RSSI may also be used. The number of clusters in SEC is determined by the distribution of the probabilities. Due to the stochastic process the gateways are distributed randomly in the network. Because of that it may happen that leaves are not covered by a sink. In the SEC approach the well known stochastic selection approaches LEACH [23], ALEACH [22], and DEEC [25] are used.

The D1C procedures use combined metrics such as weighted algorithms, as shown in Figure 5. Starting with a list of potential gateway candidates the procedure calculates weights for each candidate. The candidate with the highest weight in that round is chosen as gateway. Once the gateway is chosen, a cluster is created around that gateway with the

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

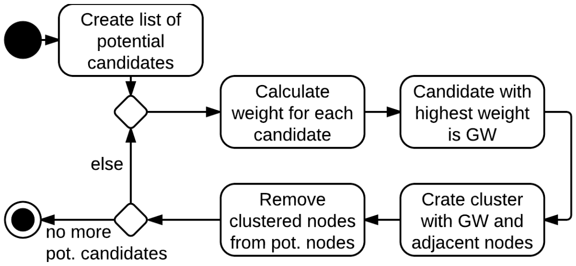


Figure 5: Process diagram of D1C. **D**eterministic gateway selection with a selection of **one** gateway and the **c**luster.

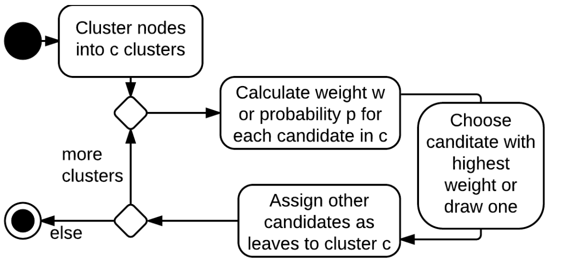


Figure 6: Process diagram of CD and CS. **C**lustering nodes according to given cluster algorithm. Either assign candidates as gateways in a **d**eterministic or in a **s**tochastic approach.

adjacent nodes. In the next step all clustered nodes and the gateway are removed from the list of potential nodes. If there are any candidates left in the list, the procedure restarts with calculating weights per candidate. WCA [24] and FWCABP [26] are used in D1C in comparison with the static behavior.

Starting with clustering, followed by either deterministic or stochastic gateway selection, the approaches CD and CS are designed (Figure 6). First, the nodes are clustered according to a given cluster algorithm. In the current version the density based DBScan [27] and the partitioning k-Means [28] clustering solutions are used together with a grid-density approach, which is based on the work by de Berg in [29]. Afterwards, one gateway is calculated for each cluster, either by calculating a weight (DS), or by calculating a probability (CS) for each node in a cluster. In the case of DS, the node with the highest weight is chosen as gateway, for CS one gateway is drawn. The other nodes in that cluster are leaves accordingly.

## VI. Performance Evaluation

The two main goals targeted by the evaluation are the following: (i) examining the impact of the proposed role allocation schemes SEC and D1C, including state-of-the-art gateway selection algorithms, on the system performance and fairness of an adaptive approach, which uses cellular offloading techniques, and (ii) comparing different clustering algorithms and their respective impact on fairness and system performance. The first part of the evaluation shows that improved fairness does not imply high degradation in system performance. That the choice of the used clustering algorithm is crucial for adaptive systems, such as eCrater, that rely on relations between the nodes in the network, is shown in the

Table I: Scenario and simulation setup

| | |
|---|---|
| Simulated Area | $1500\,m \times 1500\,m$ |
| Max. Wi-Fi Range | $88\,m$ |
| Max. Server Conn. | min. 37% |
| Role Calc. Interval | 5 min |
| Mov. Speed $[\frac{m}{s}]$ | $1.5 - 2.5$ |
| Density $[\frac{\text{nodes}}{km^2}]$ | $44.4 - 177.7$ |
| SEC GW | static, random, LEACH, ALEACH, DEEC |
| D1C GW | static, random, WCA, FWCABP |
| Clustering | DBScan, k-Means, Grid with (BECLeach, WCA) |

second part. In the following, we detail (i) the modeling of the scenario and the used evaluation parameters, (ii) the impact of the role allocation schemes SEC and D1C with state-of-the-art gateway selection algorithms such as ALEACH [22], DEEC [25], and WCA [24] (cf. Section V), and the (iii) comparison of different clustering algorithms including DBScan, k-Means, and the grid-density based approach.

### *A. Modeling of the Scenario and Evaluation Setup*

As the robustness of the system is already shown in [4] we focus on the proposed role allocation in this setup. The evaluation is performed in the Simonstrator platform [5]. The Simonstrator comprises the IEEE 802.11g standard from the ns-3 simulator [30] to model the Wi-Fi ad hoc communication between the nodes. Accordingly, to reduce node island formations and inter-connection times, that support gateway selection mechanisms, we rely on the Gauss-Markov Movement model [31]. Two hours of operations are simulated while the measurement starts after 20 minutes, resulting in 100 minutes of measurement time. The first 20 minutes are used to reach a steady state in the scenario. The simulations are repeated with five different seeds. Bar charts show the average with 95% confidence interval. For comprehensiveness reasons, the distributions of the results are shown in box plots and cdf's. The median is represented by a solid line inside the box, while the lower and upper quartile are represented by the boxes. Whiskers show the upper (lower) data point within 1.5 of the interquartile range. Outliers are represented by crosses.

The simulation setup is summarized in Table I. The different settings for the two main parts of the evaluation are shown. The movement speed is uniformly distributed in the given interval. To enforce the adaptivity of eCrater, a churn model is used that achieves a fluctuating density by joining and leaving nodes during the simulation. In doing so, the network density fluctuates from sparse to dense triggering the system to adapt between the uniform and hybrid topology.

### *B. Impact of Gateway Selection Algorithms*

Both the deterministic (D1C), as well as the stochastic (SEC) gateway selection with a followed best-fit clustering, are compared with simple static and random gateway selection algorithms. The static gateway selection approach is used, as it describes the baseline used in the initial version of Crater [4], as well as other monitoring solutions, such as DAMON [18] and Mesh-Mon [19]. This comparison of

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

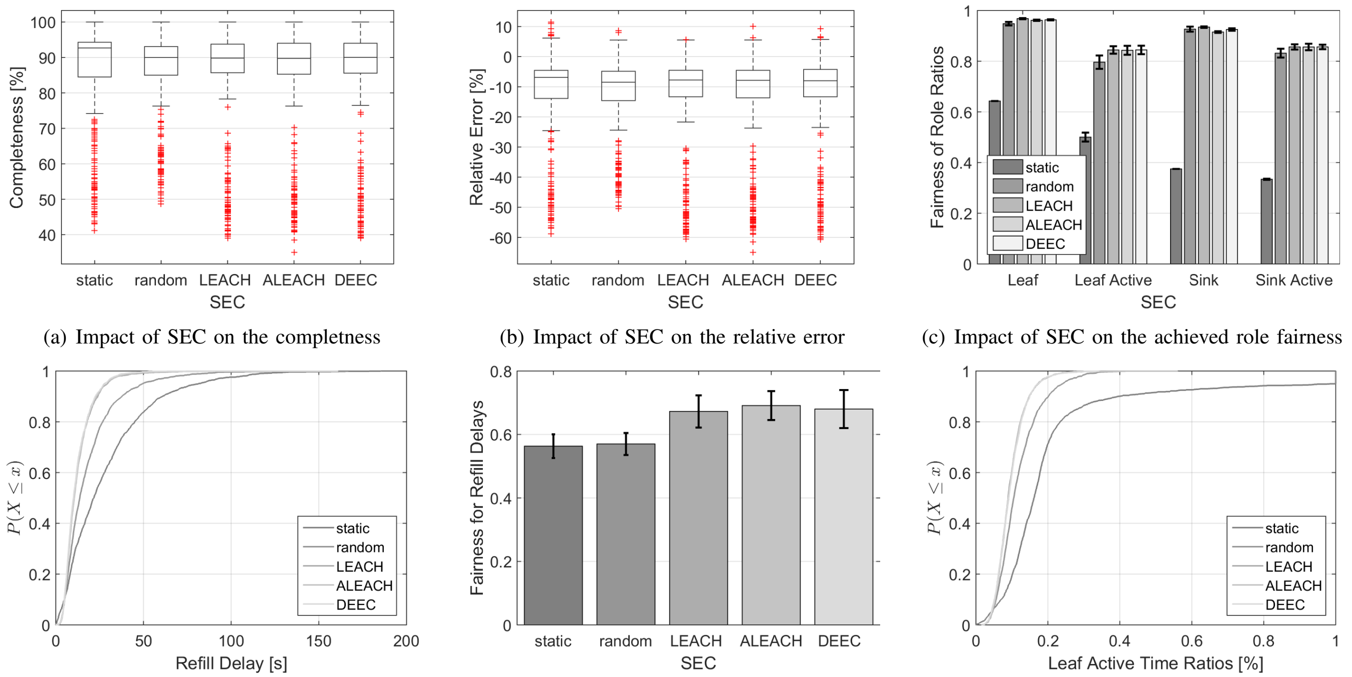


(a) Impact of SEC on the completness

(b) Impact of SEC on the relative error

(c) Impact of SEC on the achieved role fairness

(d) Impact of SEC on the responsiveness

(e) Impact of SEC on the refill delay fairness

(f) Impact of SEC on the active ratio of leaves

Figure 7: Evaluation of the impact of SEC role allocation in comparison with a static and random gateway selection.

eCRATER with other fundamental works is important to value the results achieved with both D1C and SEC in this work. Furthermore, we chose to use a random gateway selection approach, which should generally achieve good results regarding the fairness of the gateway selection process. Comparing the fairness of the approaches is essential in mobile networks as discussed in Section IV. For the evaluation of the system's performance, we plot the relative error as well as the completeness of the monitoring data. Both are based on the duplicate sensitive monitoring attribute *node count*.

eCRATER is able to deliver at least 90% of the monitoring data on the median, using the stochastic role allocation approach (SEC). But, as visible in Figure 7(a), the static selection approach is able to deliver 4% more monitoring data on the median. However, the increased completeness and slightly better accuracy for the static selection (cf. Figure 7(b)) comes at the expense of unfair role allocation. This unfairness is visible in Figure 7(c), where the respective fairness indexes for role ratios $FI(R_{\text{sink/leaf}})$ and role active ratios $FI(R_{\text{sink/leaf, active}})$ are shown. The figure shows that the static gateway selection delivers unfair results with respect to the role allocation of gateways for over 60% of the nodes. By using the random strategy, i.e. without considering e.g. the strategic placement of nodes, the fairness is significantly increased, as the probability for role changes increases. LEACH, ALEACH and DEEC are able to deliver accordingly high fairness for gateways. Additionally, the active ratios for both leafs and sinks are spread more evenly resulting in increased fairness for the role active ratios compared to the static and random gateway selection approaches. The reactiveness of eCRATER is increased as leaves in the network reconnect faster to the topology. Figure 7(d) shows that for 90% of the leafs the refill delay is reduced significantly by up to 60% (in total from 60s to 25s) using the stochastic role allocation approaches of SEC in comparison to the static allocation. Even compared to the random gateway selection an improvement of over 30% is achieved (in total from 37s to 25s) using SEC approaches. Moreover, not only the responsiveness of the system is improved. At the same time the fairness for the refill delays $FI(T_{\text{refill}})$ is improved by more than 10% compared to the static and random selection approaches (cf. Figure 7(e)). Furthermore, eCRATER is not only improving the responsiveness of the system. It also guarantees better coverage with the SEC approaches, the times leaves are active is reduced. As visible in Figure 7(f), leaves are actively trying to reconnect to the topology in 90% of the cases for 15% of their time being leaf using SEC approaches. A significant reduction of over 62% compared to the static selection (from 40% to 15%), and even 25% compared to the random selection (from 20% to 15%). This improvement of the fairness, the coverage, and the reactiveness do come for a slightly higher relative error and reduced completeness.

Similar to the SEC approaches eCRATER is able to deliver good performance with D1C approaches. The completeness (cf. Figure 8(a)) is slightly affected, but still in median 90% of the monitoring data are delivered. The choice of the D1C approach has an impact on the fairness (cf. Figure 8(b)). WCA is able to keep up with a random sink selection, whereas FWCABP looses 10% fairness, i.e. eCRATER is fair to 83% of the nodes when choosing the gateways with FWCABP com-

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

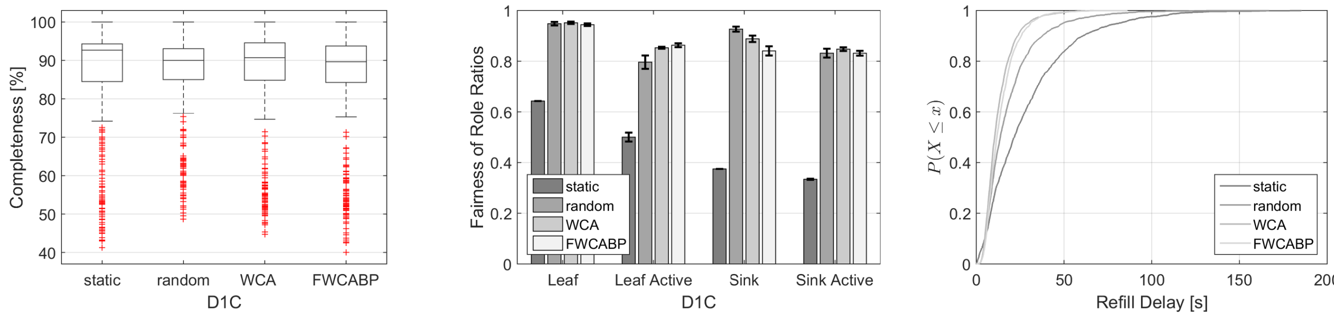


(a) Impact of D1C on the completness (b) Impact of D1C on the acheived fairness (c) Impact of D1C on the responsiveness

Figure 8: Evaluation of the impact of D1C role allocation in comparison with a static and random gateway selection.

pared to 93% fairness using the random selection. However, the system is able to reduce the refill delays compared to the random selection as shown in Figure 8(c).

To summarize the impact of either stochastic (SEC) or deterministic (D1C) gateway selection, compared to state-of-the-art approaches used in the related work, eCRATER is able to deliver significant improvement regarding the fairness of role allocations in the network as a system relying on cellular offloading. However, by enhancing the fairness in the network the completeness and the accuracy are affected. The gap between fairness and system performance exists, though with small loss in the achieved performance the system is able to make substantial improvements in the achieved fairness.

### C. Impact of Clustering

Due to lack of space the evaluation of the CS approach is not shown in this work. Accordingly, the results for the deterministic clustering (CD) are shown in the following. For estimation of the impact of clustering, one density based (DBScan [27]), one partitioning (k-Means [28]), and one grid-density based approach similar to [29] are chosen. The choice of the clustering scheme strongly affects the achieved completeness and accuracy of eCRATER. Choosing the density based clustering approach DBScan (*DBS* in figures) reduces the achieved completeness (cf. Figure 9(a)) significantly compared to k-Means (*kM* in figures) and the grid approach. The reason for that is that DBScan may choose nodes in a dense area as one cluster. However, clustering all nodes in one limited region into one cluster signifies higher load on that sink and a greater potential for collisions resulting in data loss. The relative error is in median less than -10%, whereas k-Means is able to preserve the good results achieved with the SEC and D1C approaches. The grid clustering approach ranges between DBScan and k-Means in both completeness and the relative error as visible in Figures 9(a) and 9(b). Figure 9(c) shows that not only the used gateway selection approach (as shown in Section VI-B) has an impact on the achieved fairness. Also the used clustering approaches have an impact on the fairness and as well on the performance achieved by the system. Here used with BECLeach [32] and WCA [24] gateway selection.

The choice of the clustering algorithm used in eCRATER and similar systems relying on cellular offloading techniques is crucial not only for performance, but also for fairness reasons. Beside the generic clustering schemes, the mechanism itself is important as it might depend on the outcome of the clustering and the gateway selection or, as in the case of eCRATER, only on the gateways. Similar to SEC and D1C, CD demonstrates that high fairness can be achieved without degrading the performance of systems using cellular offloading too much. However, clustering followed by the gateway selection shows stronger dependency on the used clustering algorithm.

## VII. Conclusions

In this paper, we analyzed the trade-off between performance and fairness with respect to an individual user's resource consumption for mechanisms relying on cellular offloading. As offloading systems are mostly analyzed with respect to their performance, the question arises whether performance and fairness in such systems are mutually exclusive goals. Based on an in-depth discussion of fairness measures for distributed systems, the fairness characteristics of systems relying on cellular offloading are assessed on the example of the adaptive monitoring solution CRATER [4]. By proposing essential extensions in form of state-of-the-art gateway selection and clustering approaches into the role allocation process of the system, proposing eCRATER, the achieved fairness within the system rises significantly. Instead of relying on static or semi-static allocation schemes, eCRATER uses stochastic and deterministic gateway selection algorithms for role allocation. While the achieved fairness is increased by up to 60%, only a minor decrease (below 4%) in the system's performance is observed. The results reveal that the order of gateway selection and clustering in the role allocation process is important and application specific.

In its current state, the fairness is estimated based on the responsiveness of the system and the duration of a client's contribution to the system. We are currently investigating the accuracy of system independent estimates for fairness, e.g. based on the traffic or the battery consumption. This allows the assessment of the fairness characteristics in a setting, where these indicators are also influenced by other active components on the mobile device. Considering eCRATER, we

Nils Richerzhagen, Björn Richerzhagen, Michael Walter, Dominik Stingl, Ralf Steinmetz. *Buddies, not Enemies: Fairness and Performance in Cellular Offloading* In: Proceedings 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), p: 1-9, IEEE, June 2016, ISBN 978-1-5090-2185-7/16.

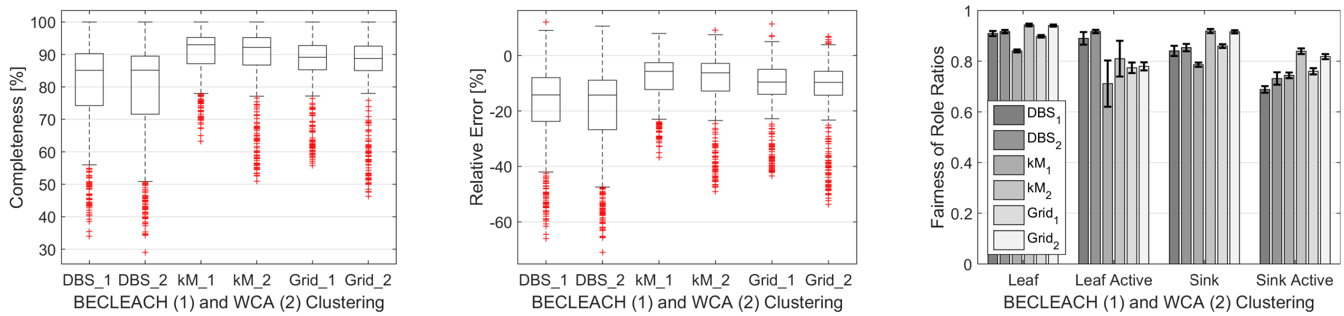


(a) Impact of CD clustering on the completness

(b) Impact of CD clustering on the relative error

(c) Impact of CD clustering on the fairness

Figure 9: Impact of DBScan (*DBS*), k-Means (*kM*), and grid clustering approaches in the CD clustering scheme.

expect that better performance can be achieved by triggering the role allocation process based on specific events, rather than periodically. Additionally, the role allocation scheme can be adapted or exchanged based on application requirements or environmental conditions, providing best performance under various conditions. Instead of obtaining the required data for the role allocation algorithms using global knowledge, the impact of the latency and the potential error introduced by using eCRATER for obtaining the data is highly interesting as this reveals multiple dependencies between the monitoring and the mechanism relying on the monitoring.

## Acknowledgment

This work has been funded by the German Research Foundation (DFG) as part of projects B1/C2/C3 within the Collaborative Research Centre (CRC) 1053 – MAKI.

## References

[1] Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2014–2019," Cisco, White Paper, 2015. [Online]. Available: http://www.cisco.com/.../white_paper_c11-481360.html

[2] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang, "A First Look at Cellular Network Performance During Crowded Events," in *ACM SIGMETRICS*, 2013.

[3] B. Richerzhagen, D. Stingl, R. Hans, C. Gross, and R. Steinmetz, "Bypassing the Cloud: Peer-assisted Event Dissemination for Augmented Reality Games," in *IEEE P2P*, 2014.

[4] N. Richerzhagen, D. Stingl, B. Richerzhagen, A. Mauthe, and R. Steinmetz, "Adaptive Monitoring for Mobile Networks in Challenging Environments," in *IEEE ICCCN*, 2015.

[5] B. Richerzhagen, D. Stingl, J. Rückert, and R. Steinmetz, "Simonstrator: Simulation and Prototyping Platform for Distributed Mobile Applications," in *SIMUTOOLS*, 2015.

[6] H. Shi, R. Prasad, E. Onur, and I. Niemegeers, "Fairness in Wireless Networks:Issues, Measures and Challenges," *Communications Surveys Tutorials*, vol. 16, no. 1, pp. 5–24, 2014.

[7] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. John Wiley & Sons, 1991.

[8] D. Bertsekas and R. Gallager, *Data Networks*. Prentice-Hall Inc., 1987.

[9] B. Radunović and J.-Y. L. Boudec, "A Unified Framework for Max-min and Min-max Fairness with Applications," *Transactions on Networking*, vol. 15, no. 5, 2007.

[10] F. Kelly, "Charging and Rate Control for Elastic Traffic," *European Transactions on Telecommunications*, vol. 8, no. 1, pp. 33–37, 1997.

[11] J. Mo and J. Walrand, "Fair End-to-end Window-based Congestion Control," *Transactions on Networking*, vol. 8, no. 5, pp. 556–567, 2000.

[12] T. Ozugur, M. Naghshineh, P. Kermani, C. M. Olsen, B. Rezvani, and J. A. Copeland, "Balanced Media Access Methods for Wireless Networks," in *ACM/IEEE MobiCom*, 1998.

[13] R. Cuevas, M. Uruenña, and A. Banchs, "Routing Fairness in Chord: Analysis and Enhancement," in *IEEE INFOCOM*, 2009.

[14] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications," in *ACM SIGCOMM*, 2001.

[15] D. D. Vergados, D. J. Vergados, A. Sgora, D. Vouyioukas, and I. Anagnostopoulos, "Enhancing Fairness in Wireless Multi-hop Networks," in *ACM MOBIMEDIA*, 2007.

[16] S. Chen and Z. Zhang, "Localized Algorithm for Aggregate Fairness in Wireless Sensor Networks," in *ACM/IEEE MobiCom*, 2006.

[17] J. Zhu, K.-L. Hung, and B. Bensaou, "Tradeoff Between Network Lifetime and Fair Rate Allocation in Wireless Sensor Networks with Multi-path Routing," in *ACM MSWiM*, 2006.

[18] K. N. Ramachandran, E. M. Belding-Royer, and K. C. AImeroth, "DAMON: A Distributed Architecture for Monitoring Multi-hop Mobile Networks," in *IEEE SECON*, 2004.

[19] Soumendra Nanda and David Kotz, "Mesh-Mon: A Multi-radio Mesh Monitoring and Management System," *Computer Communications*, vol. 31, no. 8, 2008.

[20] N. Battat and H. Kheddouci, "HMAN: Hierarchical Monitoring for Ad Hoc Network," in *IFIP EUC*, 2011.

[21] D. Stingl, C. Gross, L. Nobach, R. Steinmetz, and D. Hausheer, "BlockTree: Location-aware Decentralized Monitoring in Mobile Ad Hoc Networks," in *IEEE LCN*, 2013.

[22] M. S. Ali, T. Dey, and R. Biswas, "ALEACH: Advanced LEACH Routing Protocol for Wireless Microsensor Networks," in *IEEE ICECE*, 2008.

[23] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient Communication Protocol for Wireless Microsensor Networks," in *IEEE HICSS*, 2000.

[24] M. Chatterjee, S. Das, and D. Turgut, "WCA: A Weighted Clustering Algorithm for Mobile Ad Hoc Networks," *Cluster Computing*, vol. 5, no. 2, 2002.

[25] L. Qing, Q. Zhu, and M. Wang, "Design of a Distributed Energy-efficient Clustering Algorithm for Heterogeneous Wireless Sensor Networks," *Computer Communications*, vol. 29, no. 12, 2006.

[26] A. H. Hussein, A. O. Abu Salem, and S. Yousef, "A Flexible Weighted Clustering Algorithm Based On Battery Power for Mobile Ad Hoc Networks," in *IEEE ISIE*, 2008.

[27] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Kdd*, vol. 96-34, 1996.

[28] A. K. Jain, "Data Clustering: 50 Years Beyond K-means ," *Pattern Recognition Letters*, vol. 31, no. 8, 2010.

[29] M. de Berg, M. van Kreveld, M. Overmars, and M. Schwarzkopf, *Computational Geometry*. Springer, 2000, ch. Quadtrees, pp. 1–17.

[30] T. R. Henderson, S. Roy, S. Floyd, and G. F. Riley, "ns-3 Project Goals," in *ACM WNS2*, 2006.

[31] T. Camp, J. Boleng, and V. Davies, "A Survey of Mobility Models for ad hoc Network Research," *Wireless Communications and Mobile Computing*, vol. 2, no. 5, 2002.

[32] H. Deng, C. Yang, and Y. Sun, "A Novel Algorithm for Optimized Cluster Head Selection," *Science Journal of Electrical & Electronic Engineering*, 2013.