

Reconfiguring Cell Selection in 4G/5G Networks

Qianru Li
UCLA
qianruli@cs.ucla.edu

Chunyi Peng
Purdue University
chunyi@purdue.edu

Abstract—In cellular networks, cell selection plays a critical role in providing and maintaining ubiquitous radio access. It follows standardized procedures with operator-specific policies pre-configured by tunable parameters. These parameters specify the criteria to determine whether and how to select new serving cell(s), thus impacting access quality and user experience. Recent studies reveal that today’s cell selection fails to offer good performance as it can. This is because it is configured for seamless connectivity, and thus performance is offered at “*best effort*”. In this work, we attempt to *re-configure* these parameters by taking performance into consideration. We first conduct a measurement study in one big city in the US to demonstrate that reconfiguration indeed helps improve the overall performance, without compromising connectivity. This implies that 4G/5G networks are capable of offering better performance but such potentials are under-utilized in practice. We further explore proactive reconfiguration to prevent such unnecessary performance losses. We examine technical challenges, factors and even limitations to reconfigure cell selection in a standard-compatible manner, and finally devise a simple reconfiguration algorithm based on profiling and heuristic searching to efficiently pursue promising performance gains. The evaluation over AT&T and T-Mobile in two US cities has validated its effectiveness. Performance gains outweigh losses. Reconfiguration boosts data speed in more than 30% of instances, which exceeds the ratio of losses by at least 16%; The median speed gain is at least 89.1% (up to 217 fold).

I. INTRODUCTION

Network operators have continuously and heavily upgraded their networks to provide faster mobile broadband experience to their users. They upgrade radio access technologies (say, add 5G New Radio to their 4G networks), acquire wider radio spectrums (e.g., mmWave and sub6 bands for 5G), repurpose radio resources with more efficient uses (say, 4G/3G/2G bands for 5G), and so on. They enhance raw system capabilities to provide better radio access and higher data speed. However, the enhanced capabilities do not guarantee to turn into performance gains desired and received by mobile users.

Recent measurement studies have revealed that the actual data speed a mobile phone got is much lower than what it could get at best [1], [2]; Today’s cell selection chooses poorly-performed cells to serve the mobile device in presence of good candidate cells which are able to offer much better performance. It thus under-utilizes the full potentials of enhanced network capabilities. [1] has devised a device-assisted solution to detect and correct improper cell selection at runtime. Despite effectiveness, this solution is *a remedy, not a prevention*. It seeks to reduce or recover from performance loss after such loss has already occurred.

978-1-6654-4131-5/21/\$31.00 ©2021 IEEE

In this work, we attempt to tackle this under-utilization problem in a new perspective. We explore whether and how to prevent improper cell selection at the first place, instead of fixing it afterwards. We believe that it is promising because cell selection follows standardized procedures and controlled by operator-specific policies pre-configured by tunable parameters [3], [4]. By adjusting those configuration parameters where under-utilization is originated from, we should be able to avoid or reduce the likelihood of cell selections that under-utilize network capabilities.

We conduct a showcase study over AT&T to answer *whether* and *how* questions. Different from the previous studies [1], [2], AT&T has upgraded its network into the hybrid 5G/4G mode with recent 5G rollout. Our reality check has confirmed the need and potentials of reconfiguration. For instance, by simply tuning two parameters, the perceived data speed grows from 15.3Mbps to 120.8 Mbps (§III-A); The default cell selection misses the 5G cell which is recently deployed and the 7.9-fold speed increase is resumed when preference for 5G is enforced (implicitly) through parameters tuned for performance-aware cell selection. More importantly, the gain is not uncommon. It is persistently and frequently seen in our full-region check (§III-B). We further characterize the potentials and challenges of reconfiguration (§IV-A and §IV-B). It is not easy because reconfiguration helps in some instances but may hurt others. A desired solution must statistically improve the overall gain by tuning parameters in a high-dimensional space. It is inherently complex due to three coupling effects within a cell, among co-located cells and among all the cells in a geographical area. Despite these challenges, our study shows that it is promising to reduce the likelihood of poor cell selection in advance.

We further leverage heuristics learned from our study to design RPERF (§IV). Finally, we evaluate effectiveness and efficiency of RPERF using three datasets collected in two US cities with AT&T and T-Mobile. Performance gains outweigh the losses on all the datasets. RPERF boosts data speed in more than 30% of instances (actually, 30%, 42.6% and 41.6% over three datasets), which exceeds the ratio of worsened instances by at least 16%; The speed gain is at least 89.1% (median, up to 217 fold). This is very close to the optimal results achieved by exhaustive parameter searching. But RPERF is much more efficient as it achieves near-optimal gains with execution time reduced by a factor of 3.6 to 4.8.

To our best knowledge, this is the first study to explore reconfiguration to enhance performance of cell selection.

Release. Source codes and datasets are released at [5], along with technical support from our team.

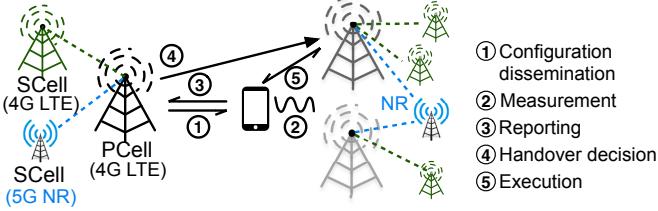


Fig. 1: Cell selection (via handover) in 4G/5G networks.

II. BACKGROUND

In 4G LTE-Advanced and beyond (say, 5G) networks, more than one cells are allowed to offer radio access to a mobile device, as illustrated in Fig. 1. The serving cell set consists of one primary cell (PCell) and several secondary cells (SCells). Both PCell and SCells are used for data transmission within one active radio connection, while the PCell is responsible for establishing and maintaining this radio connection in the control plane. By aggregating frequency carriers of PCell and SCells, this technique uses a wider frequency spectrum and thus substantially enhances data rates and network performance; It is called carrier aggregation (CA) in 4G [6] and dual connectivity in 5G [7]; 5G uses a new radio (NR) access technology, which is different from 4G. Consider 5G is currently deployed in the non-standalone (NSA) mode¹, a 5G NR cell cannot work alone as a PCell without 4G. Instead, it is added like a SCell to a 4G PCell.

Cell selection is to select one or several cells out of nearby candidate cells to serve the device. It is realized by two standardized procedures when this radio connection is established (via *initial cell selection*) or migrated (via *handover*) [3], [4]. The former procedure can be treated as a special case of the latter one (which hands over from NULL to new cells). For simplicity, we use the term of *cell selection* to represent a generalized procedure unless specified. It is performed at two phases: PCell selection and SCell selection. A PCell is first selected and then SCells (covering CA in 4G and dual connectivity in 5G NSA) are added by this PCell.

Configuring parameters for cell selection plays an essential role in giving flexibility to network operators to customize their own policies while strictly following the standard mechanism. Fig. 1 depicts a typical cell selection procedure via handover. It consists of five steps: ① configuration dissemination, ② measurement, ③ reporting, ④ decision and ⑤ execution. PCell selection starts when the current PCell broadcasts its handover-relevant parameters to all the devices; These parameters are pre-configured to define the criteria to trigger, decide and execute PCell selection at runtime. They include whether to measure neighboring cells, what cells to measure (over the same/different frequency channels), whether to report the measurement results and what to report (e.g., events A1–A6, B1–B2 [3]), how to decide the target serving cell, and so on. At runtime, measurement and reporting are triggered at the device side when the pre-configured conditions are satisfied.

¹Standalone (SA) mode allows 5G to work alone without 4G, which will be launched in the future as 5G matures [8].

TABLE I: Main configurable parameters (P=PCell, S=SCell, and RSS could be any form of RSRP or RSRQ).

Param.	Step	Criterion	P	S _{4G}	S _{5G}
θ_{A1}	②	$RSS_s > \theta_{A1}$	✓	✓	
θ_{A2}	②	$RSS_s < \theta_{A2}$	✓	✓	✓
θ_{A3}	③	$RSS_c > RSS_s + \theta_{A3}$	✓		✓
$\theta_{A5,s}, \theta_{A5,c}$	③	$RSS_s < \theta_{A5,s}, RSS_c > \theta_{A5,c}$	✓		

Afterwards, the reported measurement results are used by the current PCell to assist its handover decision, and the serving PCell switches if a handover is decided and executed. SCell selection is similar and the difference is that the criteria to use are configured by the new serving PCell and the parameter values differ from those for PCell selection, e.g., cell constraints (what cells are allowed as SCells) and reporting event thresholds. In case of no active radio connection (no serving PCell), initial cell selection is performed at the device without the reporting step. The tunable parameters and criteria are updated at step ① with the most recent PCell.

Main configurable parameters used in this work. We use AT&T and T-Mobile to study the reconfiguration problem in this work. TABLE I lists major parameters used for cell selection by both carriers, which are confirmed in our measurement study (§III-B). All the parameters and their associated criteria are regulated by 3GPP specifications [3], [4], which define a complete list of configurable options in more complex forms for global operators. Generally, the criteria are based on radio signal strength measurements (in terms of RSRP or RSRQ) of the serving cells and available candidate cells; parameters in event A1 and A2 control which cells to measure (step ②); Parameters in event A3 and A5 decide which cells to report (step ③); They work together to affect the decision of cell selection and the resulting performance.

III. MOTIVATION

In this section, we first use one real-world instance to motivate the reconfiguration problem; We then present three drive forces behind reconfiguration for better performance.

A. Example: An 8-fold Speed Increase via Reconfiguration

Figure 2 illustrates a handover instance which selects worse cells and results in much lower data speed (dropping from 120.8 Mbps to 15.3 Mbps on average). There are 8 cells involved: $P_1, P_2, Q, S_1, S_2, T_1, T_2$ and NR; Their cell information is given in Fig. 2b. These cells run over the same (marked with the same letter) or different frequency channels. Each cell is uniquely identified by its short ID and frequency channel number which corresponds to one specific frequency bandwidth regulated by 3GPP specifications [6] (for LTE) and [9] (for NR). Here, 5 channels over 4 bands are used (more observed, see §III-B); Band n5 is for 5G NR and exactly reuses band 5 for 4G (originally for 2G and 3G).

In this instance, the serving PCell hands over from P_1 to Q and then adds two SCells (S_1 and T_1); They together offer 15.3 Mbps (on average) to the device. However, this handover misses a much better choice with P_2 as PCell and S_2, T_2 and

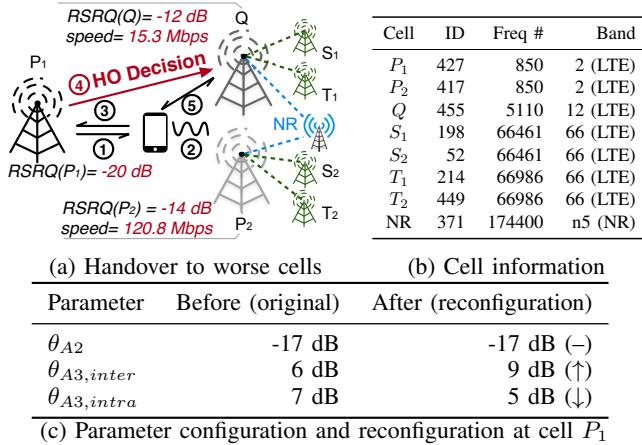


Fig. 2: An example of an 8-fold speed increase via reconfiguration (15.3 Mbps \rightarrow 120.8 Mbps, AT&T, at \star of Fig. 3)

NR as SCells, which offers 120.8 Mbps (7.9x). It is repeatedly observed at one location (\star of Fig. 3) in our study.

We next explain why the handover selects Q as the new PCell and fails to offer higher data speed it could afford. This is due to handover configurations in place. Fig. 2c lists main parameter values used by cell P_1 in the example. There are three criteria. First, there is one event A2. It specifies that inter-frequency cells are measured only when RSRQ of the PCell drops below A2 threshold θ_{A2} ; Otherwise, only intra-frequency cells are measured. Here, RSRQ of P_1 (-20 dB) is lower than θ_{A2} (-17 dB), and thus both P_2 and Q are being measured by the device. Second, there are two A3 events that specify the criteria for intra-frequency and inter-frequency measurement reporting. It is reported if the measured RSRQ of the candidate cell is offset better than PCell by θ_{A3} at least. As a result, only cell Q is reported because its RSRQ is greater than $RSRQ(P_1)$ by 8 dB, which satisfied the criterion. On the other side, P_2 is not reported because the difference in RSRQ (-20dB vs. -14 dB) is smaller than the offset $\theta_{A3,inter}$ (7 dB). Finally, only cell Q is visible to the network and gets selected eventually. In practice, several A3 events may be configured, each associated with one or multiple frequency channels of candidate cells. Last, cell P_2 (over band 2) accepts NR cells as SCells, but cell Q (over band 12) does not. It is consistently observed in our measurement study and such cell constraints are likely set by AT&T to manage her spectrum resources. At hence, the handover misses not only P_2 as a PCell but NR as a SCell (the rest two SCells running over the same frequency channels in both handover choices).

The chance of selecting better cells is eliminated by current configurations. However, current configurations are not without rational. Signal strength-centric decision has been working well for decades to provide connectivity. But it is now insufficient to reach good service, especially at locations with dense cell deployment. On the one side, those configurations do not guarantee the strongest cell is selected. Because they only target connectivity by finding cells strong enough (i.e. above some threshold). Even though the selected cell has the

highest signal strength, it does not indicate good performance compared to unselected cells. 5G NR and CA further enlarges the gap because it brings more options in terms of cell combinations which may have huge variance in capability.

We further illustrate how reconfiguration prevents such performance loss at the first place. One straightforward strategy is to include cell P_2 as a new candidate and get rid of cell Q . Fig. 2c gives one reconfiguration option, where $\theta_{A3,intra}$ decreases from 7 dB to 5 dB, $\theta_{A3,inter}$ increases from 6 dB to 9 dB and other parameter values remain the same. Therefore, cell P_2 becomes visible and finally wins when P_1 is preparing for a handover. Moreover, it is sufficient to get rid of cell Q as long as the reconfigured $\theta_{A3,intra}$ is no smaller than 8 dB; A cell over the same frequency channel (an intra-freq handover) is preferred. There are more than one effective reconfiguration options. Reconfiguration at scale is not easy as illustrated in this instance; A number of factors must be considered and we will elaborate technical challenges in §IV-A and §IV-B.

B. Three Drive Forces for Reconfiguration

We advocate reconfiguration not only for its potential performance gains, but also for its practicability and compatibility with network operations in place. It is driven by three forces.

First, reconfiguration is not new. Network operators do (often ask vendors to) configure tunable parameters to customize their operation policies while deploying and upgrading their network infrastructure. They reconfigure some or all of the parameters over time, particularly with major upgrades such as deploying a new technology (e.g., adding CA in 2016 and adding 5G NR in late 2019), acquiring new spectrum bands (e.g., adding band 12 for 4G in 2017 and band n260 for 5G in 2020), repurposing old bands (e.g., retiring band 5 for 4G and reusing it as band n5 for 5G since Nov 2019 [10]), or performing regional or national updates. This indicates that new reconfiguration strategies and algorithms for enhanced performance does not require any major physical infrastructure upgrade; They are ready to launch by leveraging the off-the-shelf interfaces and tools for reconfiguration.

Second, reconfiguration to enhance data speed is largely missing despite feasibility. The above instance is not rare. It is commonly observed in our reality check in Los Angeles (C1), one of the largest cities in the US, where AT&T have full 4G coverage and early 5G rollout. Our results are consistent with recent measurement studies in a small college town [1], [2].

Methodology and dataset. We follow the methodology in [1] but use two new phone models including Pixel 4a and Pixel 5 which support 5G in AT&T. This measurement study is mainly conducted in a $1.3 \text{ km} \times 1.6 \text{ km}$ commercial region in April - May 2021. To reduce possible biases with selected locations and routes, we run driving experiments along all accessible roads to fully cover the test region. We sample out of the region because some driving routes across the surrounding areas impact the initial cell set. In each experiment, we run an elephant flow (speedtest via file downloading) or a mice flow (ping every second) at our test phones. The former is to collect cell selection instances and data speed samples while both are

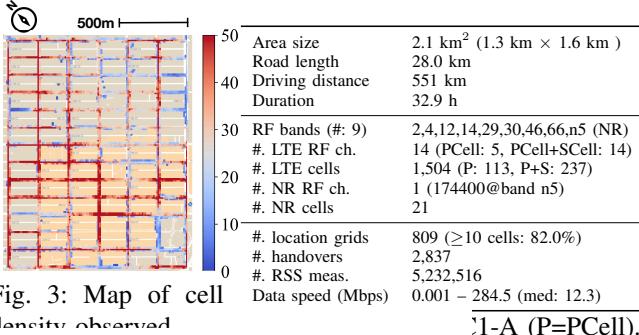


Fig. 3: Map of cell density observed

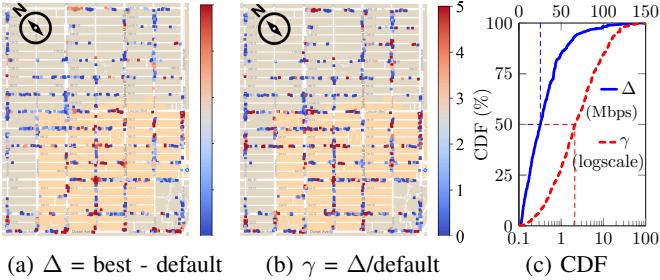


Fig. 4: Distribution of missed data speeds in our study.

used to measure radio quality and cell deployment. TABLE II gives basic information of our dataset C1-A. Fig. 3 shows the map, along with cell density observed. There are abundant candidate cells at any location (>40 at hotspots). This is thanks to continuous and heavy investment on network infrastructure upgrades (e.g., acquiring more bands, deploying denser cells and rolling out 5G).

Reality check. The purpose is to check whether default cell selection (configurations) could lead to high data speed as it can. We evaluate the performance of the selected serving cells, based on the comparison with the bound of affordable performance at the same location. The best serving cell is learned from its performance profile. We ran extensive experiments to collect sufficient performance data in the selected region and build profiles for each cell set. At a location, the best cell set has the highest median data speed among all available ones. We collect all cell selection instances and analyze whether they have chosen the best serving cells. We define a serving cell set is α -optimal if the ratio between its median speed and the median of the best cell set is no less than α ($0 \leq \alpha \leq 1$). The selection of a sub-optimal cell set implies that the existing configurations are improper for not preferring or even ignoring the better candidates. In the whole region, we observe that only 28.3% of handovers lead to 90%-optimal cells. Fig. 4 shows the speed gaps of non-optimal selections. Note that cell selections do not happen everywhere and we only show the gaps at locations of cell selections. We use the absolute and relative gaps between the median data speed by the best and selected serving cell sets. At more than 50% of instances, the data speed gap is larger than 25.4 Mbps or 214%. In the worst case, the gap goes up to 148 Mbps; It is likely larger as 5G grows. This implies that current parameters are not well tuned towards higher data speed.

Last, reconfiguration is not a remedy, but a prevention. [1]

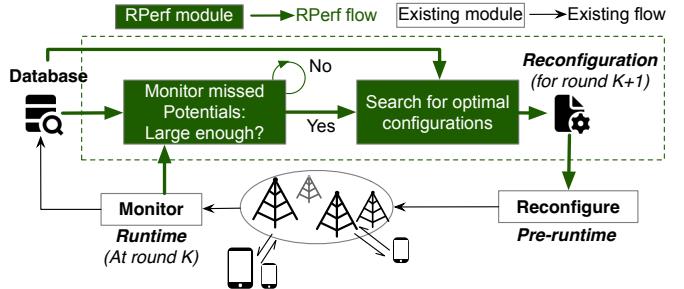


Fig. 5: Overview of the RPERF design.

has devised a device-assisted solution to detect and correct improper cell selection at runtime to boost data speed. Despite effectiveness, the solution is a remedy. It seeks to reduce or recover from performance loss after such loss has already occurred. It requires heavy profiling and training in advance and raises runtime monitoring and learning which can not complete right away (takes at least several seconds). Naturally, it fails to help when traffic flows are short or traffic patterns vary over time. A more desirable solution is to prevent under-utilization instead of mitigation after it happens. Just by tuning some parameters, we can change the result of cell selection towards the target with better performance. More importantly, it is aligned with the needs of both users and operators. The trending technology is to make 5G more intelligent and maximize the efficiency of network resources. This calls for reconfiguration beyond connectivity by taking user performance into account. Reconfiguration should be proactive, not passive. It should timely and intelligently monitor performance of cell selection and tune parameters to reduce the likelihood of poor ones, rather than take actions upon bulk user complaints.

IV. THE RPERF DESIGN

We propose RPERF to reconfigure parameters used for cell selection to enhance data performance afterwards. The overall design is depicted in Figure 5. RPERF is built on top of two existing network functions: monitoring at runtime and reconfiguration at pre-runtime; It adds two modules to trigger and execute performance-driven reconfiguration. Reconfiguration for the next round is triggered when the potentials of better performance missed by configuration at the current round has become large enough (see the triggering condition in §IV-D); It is then executed by efficiently searching parameters that achieve better performance in all the impacted cell selection instances (§IV-C). Before we dig into RPERF, we first present its technical challenges (§IV-A) and design heuristics (§IV-B).

A. Reconfiguration is not Easy

Intuitively, RPERF is to change the result of cell selection towards the target with better performance, just by tuning some parameters. However, it is not easy as illustrated before.

First, the mechanism of cell selection is based on radio signal strength, not designed for performance. To be compatible with minimal changes to the existing network infrastructure and operations, RPERF cannot directly change the outcome of cell selection but tune radio-centric criteria to *implicitly* impact

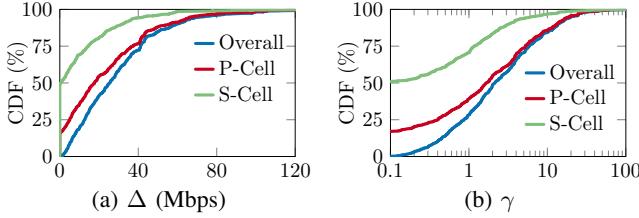


Fig. 6: CDF of the speed gaps caused by PCell/SCell selection.

the outcome. The alternative solution of explicitly changing the cells to select is discussed in §VII. Specifically, RPERF must tune these threshold parameters to change the criteria so as to get rid of bad candidate and reach optimal (or close) cells. To be qualified for being considered as a target cell (step ④), a neighboring cell must be first measured by the device (step ②) and has its radio signal strength higher than the reporting threshold (step ③). To tell good from bad, we need to build up profiles of performance and signal strength based on historical measurements. Note that performance of any unselected cell at a specific time never exists and thus is unobservable. Facilitated with the knowledge of good/poor candidates, the serving cell is able to figure out the reconfiguration towards enhanced performance and better resource utilization.

Second, reconfiguration is not optimized for individuals, but statistically for all the instances. It can not be tailored for every handover instance. Instead, it is applied to all the impacted handovers within an area. As a result, reconfiguration probably helps some cases while hurting some others. For example, reconfiguring A3 thresholds indeed increases the chance of selecting P_2 , not Q , in the above instance. But it may also degrade performance in case Q (along with its SCells) performs better than P_2 at a different location but is not selected due to reconfiguration. Therefore, the goal of reconfiguration is all about improving the *overall* performance. Whether to trigger reconfiguration, depends on whether performance gains in all the cell selection instances outweigh losses, if the losses occur. We next show the “net” gain in a what-if study (§IV-B), where gains in more cases outweigh losses in fewer cases.

Last, reconfiguration has a huge, high-dimensional space. In principle, it must tune all the parameters of all the cells *together* due to three coupling effects. (1) Parameters used by one single cell must work together to determine the steps of cell selection leading to the final target and thus cannot be tuned independently. Parameter values may be associated with different radio frequency channels (e.g. θ_{A3}); the configuration space per cell expands with the growing frequency channels to use. For each cell, there are tens or even hundreds parameters to tune. It also matches with a previous global-scale handover configuration measurement study [11]. (2) Parameters for all the cells involved in a cell selection must be tuned at the same time. Starting from the original serving cell, the device might go through one or multiple handovers until reaching the final target (without further switch). Note the target cell of a single handover may just be a transient state. Once the serving cell switches, the configured parameters and the associated criteria

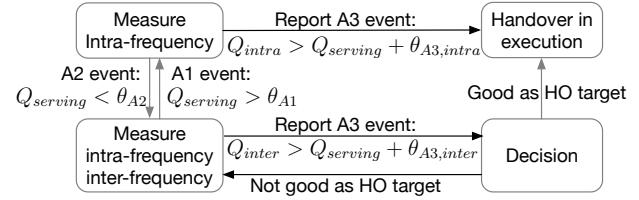


Fig. 7: Model of the handover process learned from real traces.

change accordingly. The outcome of cell selection depends on the parameters of all the involved cells. (3) All the parameters of all the cells impact each other with coverage locality. Parameters tuned to optimize performance at some locations may hurt data performance of cell selection at other locations. It is challenging to efficiently search for new parameter values across the high-dimensional space. Therefore, RPERF uses heuristics to reduce complexity of reconfiguration (§IV-B).

B. Heuristics for RPERF

We find that reconfiguration can be simplified with several heuristics learned from real-world traces.

First, PCell selection takes the major blame of missed performance. We examine the speed gaps caused by the selection of PCell or SCell(s) in our reality check. Except 28.3% handover instances that achieve 90%-optimal, 51.6% instances are caused by an improper PCell and 20.1% instances can be fixed by using different SCell(s). We further examine the distribution of speed gaps contributed by PCell/SCell selection in Fig. 6. The gaps caused by PCell selection goes very close to the overall one. It is aligned with two more observations: The largest speed gaps exist between the serving cell sets with different PCells. With P-Cell fixed, the missing performance by SCell selection goes much smaller. Therefore, we should prioritize reconfiguring PCell selection.

Second, we find that not all the configuration parameters are equally important to cell selection. In fact, a subset of tunable parameters play a decisive role. Unfortunately, such information is decided by the operator and not released to public. We follow the approach in [12] to infer the handover model. This is based on the mechanism defined by 3GPP standards [3], [4] and inference from real traces. We use MobileInsight [12] to collect 5G/4G signaling messages exchanged for each handover instance and learn a model of the handover process used by AT&T (Fig. 7). The model is represented by a state machine, demonstrating which configuration parameters to use for cell selection and how. It is highly accurate and predicts the outcomes of cell selection at precision of 99.0%. Therefore, when trying different parameter values without actual deployment, this model is able to predict new targets with high confidence.

In general, we find that the following parameters are used by AT&T to influence cell selection in our test area (city): (1) θ_{A2} and θ_{A1} set the criterion to enable (disable) measurement on inter-frequency channels. There is no inter-freq measurement until the serving cell’s signal strength runs below θ_{A2} , while intra-freq measurement is always on. In our measurement

study, we observe that θ_{A1} always equals to θ_{A2} as it functions in the opposite way. (2) θ_{A3} specifies the condition to trigger reporting for the measured cells: only when the neighboring cell's signal strength is stronger than the serving one by θ_{A3} at least. θ_{A3} has two parameter values for the intra-freq and inter-freq cells. We notice that some handovers (< 20%) are caused by unpredictable A5 reporting. We leave them because we have no access to all information available to the network operator. It can be resolved when reconfiguration is performed by the operator with all the information available on the network side (discussed in §VII). We focus on proof-of-concept reconfiguration in RPERF. As a result, our study focus on three parameters per cell: A2 threshold (θ_{A2}), A3 offsets for intra-frequency ($\theta_{A3,intra}$) and inter-frequency neighbors ($\theta_{A3,inter}$). In our dataset, all parameters are based on RSRQ.

Last but not least, we find that it is promising to reduce the configuration dimensions as many factors can be decoupled without impacting the performance. We conduct a what-if study to examine the need and feasibility of performance-driven reconfiguration. We enumerate all possible values for those critical configurations to estimate the highest possible reward for all handover instances. We simplify the whole-space search with two tricks. First, parameters are tuned within a rational range (not too far away from actual values observed in our dataset), namely, $\theta_{A2} \in [-17, -8]$ dB, $\theta_{A3,intra}, \theta_{A3,inter} \in [0, 10]$ dB. This is reasonable to make reconfiguration practical at both the network and device sides. Second, we reconfigure each frequency channel separately, not per cell. Carriers would like to simplify configuration with area-specific policies instead of cell-specific ones. As our test area is small enough, we consider distinct configurations per frequency channel to use. Our study further shows that such simplifications are reasonable and the impacts are negligible.

Our what-if study is performed with three steps:

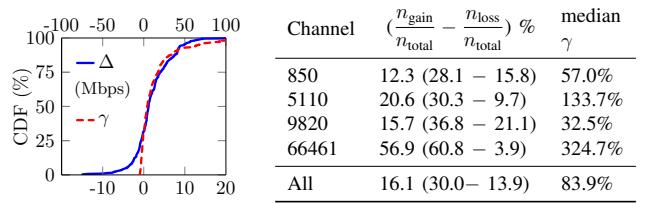
- 1) For each parameter setting, predict the new target of each handover instance based on the model (Fig. 7) and profiles.
- 2) Estimate the overall reward considering gains and losses over all the cases. In this step, we define the reward as the possibility of gains minus the possibility of loss:

$$R = \frac{n_{gain}}{n_{total}} - \frac{n_{loss}}{n_{total}},$$

where n_{gain} and n_{loss} refer to the number of cell selection instances with gains and losses after reconfiguration and n_{total} is the total number of all instances.

- 3) Repeat step 1 & 2 to enumerate all the possible settings and select the values associated with the highest reward.

We note that the network operator could assess the benefit with different reward functions, considering the aggregated gains and losses together. For example, a conservative operator might first minimize the number of cases with performance losses, and then maximize the gains atop of that. In the what-if study, we have to use the profile extracted from our dataset, because the device cannot measure performance of multiple cells at the same time.



(a) Dist. of gain and loss. (b) Overall reward of optimal reconfig.

Fig. 8: Performance gain/loss of ‘optimal’ reconfigurations.

Reconfiguration helps more. Reconfiguration could bring promising benefits as expected. We notice that it is double-edged but the gain outweighs the loss. Fig. 8a shows the distribution of gains and losses in all the impacted handover instances. Performance are enhanced in 30.0% of all cell selection instances. Meanwhile, performance is degraded in 13.9 % cases. Considering all gains and losses, the median increase is by 83.9% or 5.9 Mbps as absolute change. There are no changes in the rest 54.3% of instances because reconfiguration would not impact the outcome of every cell selection.

On each individual channel, the percentage of cases w/ performance gain also dominates those w/ loss. In our study, there are five frequency channels used for PCells. Table 8b demonstrates the optimal reward for all channels, after filtering one channel with insufficient samples (here, 5330). On all the four observed channels, the percentage of improved cases exceeds the hurt cases by at least 12.3%. The median speed increase is at least 32.5% on all channels, and reached up to 327.4%. Note that we could obtain all such benefits by simply tuning some parameters within reasonable ranges. These results indicate great improvement to be achieved with proper reconfiguration per frequency channel.

The searching space reduces. Most importantly, we find that it is feasible to decouple those inter-dependent parameters by applying restrictions to reconfiguration. Instead of tuning all parameters together, we can reduce the search space with three key observations.

1) The possibility of immediate switches after one handover can be largely reduced as long as the signal strength is stable. Without impacting the final handover decision (impacting the reward), we can adopt the following tactics: (1) consider only non-negative values for A3 offsets, and (2) use θ_{A2} as the lower bound of the candidate's signal strength when inter-freq handover is considered. Note that A2 and A3 are used to tune the measurement and reporting steps. All rules above will greatly reduce the possibility of continuous handovers at the same location, rather than eliminate it.

2) Changing A2 threshold would only cause marginal difference to the reward. Figure 9a demonstrates the trend of the maximal rewards by tuning θ_{A2} . Across the value range, the absolute difference between the maximum and minimum rewards is no more than 10%. In addition, the maximum reward falls into [-17, -15] dB on all the tested channels. It justifies that the gain of reconfiguration is not compromised when A2 threshold is restricted to a small range for complexity reduction. In the following reconfiguration search, we should

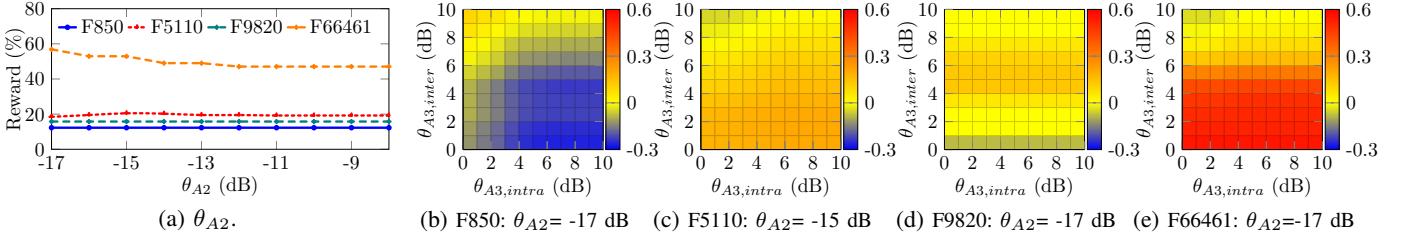


Fig. 9: The impact of θ_{A2} , $\theta_{A3,intra}$ and $\theta_{A3,inter}$ in our what-if study.

Algorithm 1 Fast search for optimal configurations

```

1: function FASTSEARCH
2:   Set  $\theta_{A2}, \theta_{A3,intra}, \theta_{A3,inter}$  to initial values
3:    $r_{max} = -Inf$ 
4:   for  $\theta'_{A2}$  in [-17, -15] dB do
5:      $\theta'_{A3,intra}, \theta'_{A3,inter}, r' \leftarrow \text{MAXREWARD}(\theta'_{A2})$ 
6:     if  $r' > r_{max}$  then
7:        $r_{max} \leftarrow r', \theta_{A2} \leftarrow \theta'_{A2}$ 
8:        $\theta_{A3,intra} \leftarrow \theta'_{A3,intra}, \theta_{A3,inter} \leftarrow \theta'_{A3,inter}$ 
9:     end if
10:   end for
11:   return  $\theta_{A2}, \theta_{A3,intra}, \theta_{A3,inter}$ 
12: end function
13: function MAXREWARD( $\theta_{A2}$ )
14:   Set  $\theta_{A3,intra}, \theta_{A3,inter}$  to initial values
15:    $\theta'_{A3,inter} \leftarrow \arg \max R(\theta_{A2}, \theta_{A3,intra}, \theta_{A3,inter})$ 
16:    $\theta'_{A3,intra} \leftarrow \arg \max R(\theta_{A2}, \theta_{A3,intra}, \theta'_{A3,inter})$ 
17:   return  $\theta_{A3,intra}, \theta_{A3,inter}, R(\theta_{A2}, \theta_{A3,intra}, \theta_{A3,inter})$ 
18: end function

```

prioritize a small range of [-17, -15] dB.

3) A3 offsets play a more critical role in impacting the reconfiguration reward. These configurations directly determine the qualification as target cells. We examine the reward of all combinations of $\theta_{A3,intra}$ and $\theta_{A3,inter}$ from 0 to 10 dB. The range covers the dominant values used by AT&T: $\theta_{A3,intra} = 3$ dB and $\theta_{A3,inter} = 5$ dB for all channels. Fig. 9 depicts how the reward changes with regards of A3 offsets, given a fixed A2 threshold. All channels have strong correlation with inter-freq configuration. Comparatively, intra-freq A3 offset has less impact on the reward. Take channel 66461 as an example (Figure 9e). With parameter $\theta_{A3,inter}$ fixed, tuning $\theta_{A3,intra}$ only changes reward slightly (absolute change less than 4%). In the opposite way, tuning $\theta_{A3,inter}$ covers the reward from -3.9% to 56.9%. Another important observation is that, for all value of $\theta_{A3,intra}$, the best choice of $\theta_{A3,inter}$ is nearly constant.

C. Fast search

We now incorporate these heuristics to RPERF. The core is to efficiently search for new values of configurations. The brute-force approach is unrealistic given high-dimensional configuration space and tremendous handover instances. In RPERF, we design fast search by prioritizing sub-space search and then using linear search for acceleration (Algorithm 1).

Subspace prioritization. To reduce the searching complexity, our first take is to decouple those inter-dependent parameters by applying restrictions to reconfigurations. In particular, we focus on A2 thresholds in a range of $\theta_{A2} \in [-17, -15]$ dB and only positive values for A3 offsets. Meanwhile, we also prioritize candidates with $\text{RSRQ} \geq \theta_{A2}$ in the stage of final decision. Our study shows that 93.5% of original handovers would not precede another handover at the same location, given the above subspace. This is good enough to enable distributed reconfiguration on each channel, which makes the design scalable. Note that our design cannot make 100% handovers stable; Otherwise more stringent condition is expected to select the target, which may impede on-time handover and even hurt connectivity.

Linear search. We then iteratively search for the optimal parameters in the space to explore. Given distinct impacts of configuration parameters (Fig. 9), we take two strategies: (1) enumerate all values of A2 threshold, given a small range, and (2) tune $\theta_{A3,inter}$ before $\theta_{A3,intra}$, as the reward is largely decided by $\theta_{A3,inter}$. It helps us to benchmark the highest reward “level”. Then, tuning $\theta_{A3,intra}$ further optimizes the reward on that level. The time complexity is $O(NK(|\theta_{A3,inter}| + |\theta_{A3,intra}|))$, in which N is the total number of handover instances, K is the number of frequency channels and $|\theta_{A3}|$ is the size of the range.

We argue that such heuristics-based fast search may sacrifice the reward optimality but it is acceptable and practical. First, the extra reward from the sub-space search to the whole space search is marginal. This is likely because network operators do not reconfigure parameters for performance and thus the reward is significant with such reconfiguration. At hence, there is no much need to push to the limit once the potential of reconfiguration is almost fulfilled. Second, current parameter values are not set randomly. They came from many-year experience and professional field trials. The engineers and technicians do radio planning and (re)configure these parameters for radio connectivity. Abundant cell deployment and increasing capabilities result in good radio \neq good performance, which opens room for performance-driven reconfiguration. However, good values must comply with good radio coverage; It is often harmless to narrow down reconfiguration to a small subspace (validated by empirical studies).

D. Triggering Reconfiguration

Reconfiguration is triggered based on the possible performance reward of all handover instances. Generally, the

network evaluates it using statistical measures periodically (e.g. every day or two) to avoid frequent changes. Which measure to monitor is up to the operator’s decision. For example, our design uses the ratio of handover instances whose targets are *not* 90%-optimal. Reconfiguration is invoked when the ratio goes above 30%. The operator can define measures out of their needs, as long as they are consistently used for reconfiguration triggering and optimization. Given selected measures, a triggering condition is then created to indicate when the gap goes beyond tolerance.

V. EVALUATION

RPERF is evaluated by trace-driven emulation. Unfortunately, we cannot test RPERF on real systems or large-scale testbeds, since we do not have internal access to change configurations. In order to emulate results in practice, we use real data from collected traces to approximate the actual network conditions, including handover instances, cell signal strength and performance.

In addition to the previous C1-A dataset (Los Angeles, AT&T), we use two more datasets for evaluation: C1-T and C2-A. C1-T is collected at the same region as C1-A to test RPERF with T-Mobile. We ran similar driving experiments over T-Mobile for 17.8 hours and 305 km in Aug 2021. C2-A is a public dataset over AT&T in West Lafayette, IN (C2) [1]; It contains performance and radio information in a region of 2.5 km² with 876 grids. Note that C2-A does not include 5G measurement as a result of no coverage. Then, we run RPERF on three datasets and evaluate the overall improvement (§V-A) and efficiency (§V-B). We also compare the results on different datasets and show insights on 5G (§V-C).

A. Overall Improvement

AT&T. We evaluate RPERF by analyzing the performance gain and loss. In C2-A, we see that AT&T use the same parameters as C1-A: $\theta_{A2}, \theta_{A3,intra}, \theta_{A3,inter}$, which are the decisive factors to cell selection. We make three observations.

First, RPERF would greatly enhance the performance by improving a large proportion of users. We use the metric R defined in §IV-B, the “net” gain which takes both improved cases and worsened cases into account. Larger values indicates that the number of improved cases is much more than the number of worsened ones. As shown in TABLE III, the “net” gain is 16.0% in C1-A (30% of cases with gains v.s. 14.0% of cases with losses); It is higher in C2-A, which reaches 27.8% (42.6% v.s. 14.8%). We examine the results per frequency channel and see the percentage of better instances is always higher than the worse one in both datasets.

Second, considering all cell selection instances impacted by reconfiguration, the majority still get a big surge in performance. We use the median of absolute speed increase (i.e., Δ in Mbps) and relative increase (i.e., γ in %) in performance over all instances with performance change (not just limited to gain). Larger numbers indicate higher increase overall. In C1-A, the median speed increase is 5.9 Mbps, or 81.3% as relative value (TABLE III). In C2-A, the speed increases by

9.6 Mbps or 56.2%. Therefore, despite worsened cases, RPERF still benefits users with a decent overall increase.

Last, the gain outperforms the loss in terms of the increased data speed. Fig. 10 shows the absolute difference (Δ) and the relative difference (γ) in three datasets. In AT&T, the median speed grows by 13.6 Mbps (200.0%) in those improved instances, while the median drop is 7.7 Mbps (45.9%) in those worsened cases in dataset C1-A; We see that the gain declines a little bit in C2-A: the median gain is 14.0 Mbps (89.1%) while the median drop is 13.3 Mbps (32.7%) in those worse instances. This is because the gains over some frequency channels (1125, 9820, 9840) are fewer than those at other channels and these channels are observed in C2-A; According to the median speeds, the absolute gain outperforms the loss on most frequency channels: 3 out of 4 in C1-A and 6 out of 11 in C2-A.

T-Mobile. RPERF is applicable to other carriers. We first find that the handover model used by T-Mobile is almost the same as the one by AT&T, except for A5 event used to select inter-frequency cells. This model imitates the network’s operations with high confidence, given the prediction accuracy of 98.7%. Accordingly, there are three tunable parameters critical to cell selection: $\theta_{A2}, \theta_{A3,intra}$ and $\theta_{A5,c}$. The first two are used to monitor serving cell and intra-freq candidates, which are the same as AT&T; $\theta_{A5,c}$ is used to search inter-freq candidates, and $\theta_{A5,s}$ is omitted because it always overlaps with θ_{A2} . T-Mobile uses RSRP, rather than RSRQ by AT&T. TABLE III and Fig. 10 show that RPERF works well over T-Mobile. 41.6% of cell selection instances are improved with median gain of 10.7 Mbps (99.8%), while 24.6% of instances have median performance drop of 10.4 Mbps (by 41.9%). Compared to C1-A, RPERF benefits more instances while hurts more instances at the same time, which leads to the similar “net” gain of 17.0%.

B. Efficiency

Next, we evaluate the efficiency of RPERF. As mentioned in §IV, RPERF utilizes several heuristics and adopts fast search to pursue considerable performance gain while minimizing the complexity of reconfiguration. To assess the efficiency of RPERF, we introduce reconfiguration via brute-force search as the baseline. Generally, brute-force search enumerates all combinations of parameters values and ends up with the optimal reconfiguration on the searching space. We compare RPERF with the baseline, by checking the difference in performance improvement and the execution time needed to figure out corresponding reconfiguration.

RPERF turns out to approach the optimal performance while reducing the cost by a factor 3.6 to 4.8. We compare the “net” increase in TABLE III, which is the goal for RPERF and the brute-force approach to optimize. The absolute value of gap between the reward achieved by RPERF and the optimal reconfiguration is negligible: 0.1% in C1-A, 0.6% in C2-A, and even no loss for C1-T at all. Moreover, RPERF could achieve the optimal reward on 3 out of 4 channels in C1-A and 7 out of 11 channels in C2-A. On other channels, the amount

TABLE III: Gain and loss after applying reconfiguration. (Configurations: AT&T - RSRQ, T-Mobile - RSRP.)

Dataset	Freq.	Fast search							Optimal (brute-force)* \$(\frac{n_{\text{gain}} - n_{\text{loss}}}{n_{\text{total}}})\%"\$
		\$(\frac{n_{\text{gain}} - n_{\text{loss}}}{n_{\text{total}}})\%"\$	median \$\gamma\$ (%)	median \$\delta\$ (Mbps)	overhead \$\downarrow\$	\$\theta_{A2}\$	\$\theta_{A3,\text{intra}}\$	\$\theta_{A3,\text{inter}}\$ (A) / \$\theta_{A5,\text{inter}}\$ (T)	
C1-A	850	12.3 (28.1 - 15.8)	57.0	4.5	3.6\$\times\$	[-17,-15]	0	10	12.3 (28.1 - 15.8)
	5110	20.5 \$\downarrow\$ (30.5 - 10.0)	131.7	8.3	3.3\$\times\$	-15	10	2	20.6 (30.3 - 9.7)
	9820	15.7 (36.8 - 21.1)	32.5	2.8	6.2\$\times\$	[-17,-15]	10	5	15.7 (36.8 - 21.1)
	66461	56.9 (60.8 - 3.9)	184.7	8.3	3.7\$\times\$	-17	3	0	56.9 (60.8 - 3.9)
	Overall	16.0 \$\downarrow\$ (30.0 - 14.0)	81.3	5.9	3.6\$\times\$	N/A	N/A	N/A	16.1 (30.0 - 13.9)
C2-A	850	0.7 (12.2 - 11.5)	19.9	5.8	3.9\$\times\$	[-17,-15]	2	10	0.7 (12.2 - 11.5)
	1125	8.3 (50.0 - 41.7)	24.2	2.8	5.6\$\times\$	[-16,-15]	10	[2,3]	8.3 (50.0 - 41.7)
	1150	16.2 \$\downarrow\$ (49.4 - 33.2)	30.1	4.5	3.9\$\times\$	-15	10	3	21.8 (49.4 - 27.6)
	2425	37.5 \$\downarrow\$ (51.4 - 13.9)	136.3	17.8	3.8\$\times\$	-15	10	3	38.9 (51.4 - 12.5)
	5145	19.5 (34.1 - 14.6)	51.3	5.3	4.4\$\times\$	-15	10	1	19.5 (34.1 - 14.6)
	9820	1.8 \$\downarrow\$ (37.5 - 35.7)	12.6	3.1	4.3\$\times\$	[-16,-15]	10	6	3.6 (39.3 - 35.7)
	9840	19.1 (38.1 - 19.0)	17.2	3.0	4.9\$\times\$	[-16,-15]	10	3	19.1 (38.1 - 19.0)
	66486	62.9 \$\downarrow\$ (72.6 - 9.7)	67.7	12.2	3.8\$\times\$	-15	10	1	63.6 (73.1 - 9.5)
	66661	57.9 (75.4 - 17.5)	82.2	9.3	4.0\$\times\$	-15	10	1	57.9 (75.4 - 17.5)
	66911	84.6 (84.6 - 0.0)	169.7	22.6	4.0\$\times\$	[-16,-15]	10	5	84.6 (84.6 - 0.0)
C1-T	66936	24.4 (51.1 - 26.7)	69.4	5.9	4.4\$\times\$	-15	10	1	24.4 (51.1 - 26.7)
	Overall	27.8 \$\downarrow\$ (42.6 - 14.8)	56.2	9.6	4.0\$\times\$	N/A	N/A	N/A	28.4 (42.8 - 14.4)
C1-T	1125	24.2 (42.5 - 18.3)	40.8	2.7	5.4\$\times\$	-101	6	[-117,-114]	24.2 (42.5 - 18.3)
	5035	26.7 (26.7 - 0)	1511	6.4	5.9\$\times\$	[-106,-100]	[0,10]	[-104,-100]	26.7 (26.7 - 0)
	66786	15.9 (41.5 - 25.6)	31.2	3.7	4.7\$\times\$	[-118,-117]	9	[-105,-100]	15.9 (41.5 - 25.6)
	68886	50.0 (60.0 - 10.0)	96.2	12.0	5.1\$\times\$	[-107,-100]	10	[-115,-114]	50.0 (60.0 - 10.0)
	Overall	17.0 (41.6 - 24.6)	31.8	3.6	4.8\$\times\$	N/A	N/A	N/A	17.0 (41.6 - 24.6)

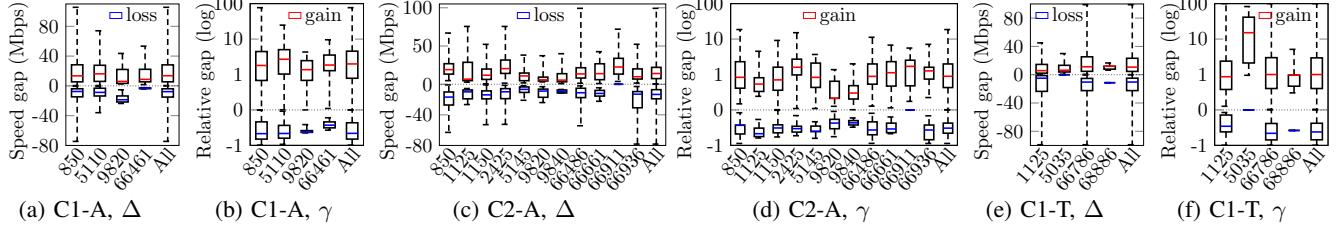


Fig. 10: Performance gain/loss with RPERF.

of unachieved reward is within 0.1% and 5.6%, respectively. In the meanwhile, RPERF would speed up the reconfiguration by a factor of 3.6 and 4.8. The overhead gets reduced greatly, almost without comprising the performance gain.

C. Comparing Results on Difference Datasets

While RPERF has achieved considerable performance improvement for both datasets, we still observe some differences and obtain insights from the comparison.

Reconfiguration towards 5G. As we know, 5G is only deployed in C1 but not C2. To be more specific, AT&T 5G cell are used as SCells only if the PCell is from frequency channel 850. Such restriction is probably enforced by radio resource planning. Therefore, in order to increase the usage of 5G service, the operator should tune parameters in favor of cells on channel 850. Parameter values chosen by RPERF have already implied a similar intention. On channel 850, RPERF sets $\theta_{A3,\text{intra}} = 0$ dB and $\theta_{A3,\text{inter}} = 10$ dB. Such new parameters indicate high preference of intra-freq handovers over inter-freq handovers. It would help secure the user on channel 850, instead of migrating to a different channel. As a result, RPERF has successfully avoided performance loss on 28.1% cases with only 15.8% cases getting worse. On the contrary, we notice the gain on the same channel in C2-A is negligible. It proves that channel 850 has been enhanced

by aggregation with 5G cells, which makes cells on channel 850 more likely outperform others. It also shows the urgency of reconfiguration: As the operators are rolling out mmwave cells, the potential gap between good and bad cells will be further enlarged. Therefore, cell selection is encouraged to consider performance and RPERF could be easily patched onto the infrastructure to prevent under-utilization.

Reconfiguration on the level of frequency channel. We also notice that the overall gain in C2 is larger than C1. This is mainly because 4 bad channels on band 66 (the last 4 channels in TABLE III) are densely deployed and frequently selected as PCell. Cells on those channels do not accept any SCell, which results in much more narrow channel width compared to others. Therefore, tuning parameters on band 66 towards inter-freq handovers could greatly save loss.

This finding sheds light on reconfiguration on the level of frequency channel, instead of per cell. It reveals the tendency of operators to manage radio resource for each frequency channel, instead of individual cells. For example, in our dataset, cells on the same frequency share the same channel width. In addition, operators may just support limited combinations of frequencies for carrier aggregation. Such behavior of radio planning make cells on the same frequency share common capabilities. Therefore, reconfiguration on a channel is aligned with such behavior of radio planning. It will get promising gain

for making good use of the commonality within one channel and discrepancy among channels.

VI. RELATED WORK

We introduce the most relevant work only. Our work is inspired by recent studies [1], [2]. [2] is the first measurement study to reveal and characterize the missed performance problem due to “improper” cell selection; and [1] is a follow-up of [2] which proposes a device-assisted fix. Our work differs as we attempt to *prevent* such improper cell selection, rather than fix them afterwards. We use our measurement study in one big US city and their datasets to demonstrate the promising benefits of performance-driven reconfiguration. Several handover studies (e.g., [11], [13], [14]) examine how handover is performed in operational cellular networks (4G/3G networks) and disclose that performance may become worse after a handover. [15], [16] optimizes handover policies to reduce failures in extreme mobility. Our study is center on tuning parameters, to increase data performance which is missed by current configuration.

VII. DISCUSSION AND CONCLUSION

We present RPERF to prevent improper cell selection which fail to select the cells with good performance in today’s 4G/5G networks. Instead of fixing them afterwards, we investigate a new perspective to make the thing right at the start. We demonstrate the need, feasibility and potential benefits of parameter reconfiguration despite limitations.

RPERF aims to optimize data performance impacted by cell selection but itself is far away from “optimal” due to its inherent limitations and remaining issues.

Limited traces. We design and evaluate RPERF based on real-world traces. However, our traces are limited as they are collected from mobile phones, not from the network side. Network operators have a much larger sample set and more complete ground-truth regarding their cell selection configurations and operations. In addition, the handover model is extracted from their operation directly and thus the overall reward is much less biased. Our effort is to leverage what we can to demonstrate the potentials of reconfiguration and call for attentions and actions from network operators. If the standards could enforce consideration of throughput onto the cell selection, we will have not only a more complete view of network performance, but also new perspectives to solve this problem in policies and mechanisms other than configurations.

Spatial granularity for reconfiguration. In this work, reconfiguration is performed on two selected regions of $2 - 2.5 \text{ km}^2$. As the regions are small, we do not split them into smaller sub-regions for area-specific reconfiguration. However, given broad (nationwide) coverage, an operator has to split the entire area into smaller regions for separate reconfiguration. A practical solution needs a proper spatial granularity to trade off the performance gain and practicality. It can be aligned with network deployment (the network is divided to serve different geographical areas). An alternative solution is to start with reconfiguration in small regions and merge adjacent

regions with close parameters into larger ones. After merge, reconfiguration is performed in big regions. It is feasible as regions in close proximity are likely to share common features in radio resource management, data usage etc. As the operator keeps updating cell deployment and radio planning, previously merged regions may gradually lose their commonality. Therefore, the operator should adjust reconfiguration regions regularly or after a big system update. To validate this solution, we will enable measurements in much wider areas. We will also release our tools to conduct measurement and analysis of cell selections at places of user interest.

Run-time dynamics. Run-time dynamics like radio quality fluctuation, scheduling and cell load could be impact factors. We aim to reduce the impact of transient factors and focus on for the overall reward affected by *persistently* worse cell selections. This is first validated by [2] and further confirmed by our latest experiments with 5G/4G. Accordingly, reconfiguration proposed in RPERF is to prevent such persistent performance loss by promoting cell selections towards better cells. We admit that reconfiguration learned from the historical data may not work for cell selections in the next second. But it seeks for the overall reward of statistical significance which eliminates that impact of runtime dynamics. An alternative solution is to make decision based on run-time situations, which provides a different angle for cell selection. It could ensure quality cell selection individually. This will complement the proposed reconfiguration and warrants future work.

5G-related issues. In our study, AT&T and T-Mobile deploys their 5G networks by adopting dynamic spectrum sharing (DSS) technology which runs two generations of cellular networks (4G and 5G) over the same frequency channel. As a result, the achieved speed in 5G is quite comparable to the legacy 4G, despite of small speed growth. This is why we observe similar gains in both cities while no 5G is deployed in C2. However, with more advanced 5G technologies including mmWave and Standalone 5G, we believe that 5G can be much faster, which will raise a pressing need for reconfiguration to reconfiguration to reduce poorly-performed cell selection and promote good ones.

Miscellaneous. There are unexplored design options in triggering and executing reconfiguration. Instead of heuristic-based search, advanced ML techniques like neural networks can be exploited with a much larger dataset. Reinforcement learning seems to fit by iteratively tune parameters. When to trigger can be performed with periodic checking, runtime monitor over down-sampled traces, or hybrid. RPERF is far away from a perfect solution. Instead, it is more like a proof-of-concept demo which demonstrates that reconfiguration is simple, ready-to-launch with immediate benefits. More practical solutions will follow once network operators take actions.

Acknowledgments. We appreciate our shepherd Prof. Feng Qian and anonymous reviewers for their constructive comments. This work was partially supported by NSF grants: CNS-1750953, CNS-2027650 and IIS-2112471.

REFERENCES

- [1] H. Deng, Q. Li, J. Huang, and C. Peng, “iCellSpeed: Increasing Cellular Data Speed with Device-Assisted Cell Selection,” in *MobiCom’20*, 2020.
- [2] H. Deng, K. Ling, J. Guo, and C. Peng, “Unveiling the Missed 4.5G Performance In the Wild,” in *HotMobile’20*, March 2020.
- [3] 3GPP, “TS36.331: E-UTRA; Radio Resource Control (RRC),” 2020, v15.9.0.
- [4] ———, “TS38.331: NR; Radio Resource Control (RRC),” 2021, v16.4.1.
- [5] “Rperf release (icnp’21): Datasets and source codes,” <https://github.com/cathyli93/icnp21-reconfig-dataset>, 2021.
- [6] 3GPP, “TS36.101: E-UTRA; User Equipment (UE) radio transmission and reception,” 2021, v16.9.0.
- [7] ———, “TS37.340: NR; Multi-connectivity; Overall description; Stage-2,” 2021, v16.5.0.
- [8] GSMA, “5g implementation guidelines: Sa option 2,” <https://www.gsma.com/futurenetworks/wp-content/uploads/2020/06/5G-SA-Option-2-ImplementationGuideline-v1.3.pdf>, June 2020.
- [9] 3GPP, “TS38.104: NR; Base Station (BS) radio transmission and reception,” 2021, v16.7.0.
- [10] OpenSignal, “Analyzing at&t’s spectrum usage to understand its 5g rollout plans,” <https://www.opensignal.com/2020/03/09/analysing-atts-spectrum-usage-to-understand-its-5g-rollout-plans>, 2020.
- [11] H. Deng, C. Peng, A. Fida, J. Meng, and C. Hu, “Mobility Support in Cellular Networks: A Measurement Study on Its Configurations and Implications,” in *ACM Internet Measurement Conference*, ser. IMC’18, 2018.
- [12] Y. Li, C. Peng, Z. Yuan, J. Li, H. Deng, and T. Wang, “MobileInsight: Extracting and Analyzing Cellular Network Information on Smartphones,” in *ACM International Conference on Mobile Computing and Networking*, ser. MobiCom’16, 2016.
- [13] Y. Li, H. Deng, J. Li, C. Peng, and S. Lu, “Instability in distributed mobility management: Revisiting configuration management in 3g/4g mobile networks,” in *ACM SIGMETRICS*, 2016.
- [14] S. Xu, A. Nikravesh, and Z. M. Mao, “Leveraging context-triggered measurements to characterize lte handover performance,” in *PAM*, 2019.
- [15] Y. Li, Q. Li, Z. Zhang, G. Baig, L. Qiu, and S. Lu, “Beyond 5g: Reliable extreme mobility management,” in *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, 2020, pp. 344–358.
- [16] Y. Li, E. Datta, J. Ding, N. Shroff, and X. Liu, “Bandit policies for reliable cellular network handovers in extreme mobility,” *arXiv:2010.15237*, 2020.