



CA++: Enhancing Carrier Aggregation Beyond 5G

Qianru Li^{†*}, Zhehui Zhang^{†*}, Yanbing Liu[‡], Zhaowei Tan[†], Chunyi Peng[‡], Songwu Lu[†]

[†]University of California, Los Angeles

[‡]Purdue University

Abstract

Carrier aggregation (CA) is an important component technology in 5G and beyond. It aggregates multiple spectrum fragments to serve a mobile device. However, the current CA suffers under both high mobility and increased spectrum space. The limitations are rooted in its sequential, cell-by-cell operations. In this work, we propose CA++, which departs from the current paradigm and explores a group-based design scheme. We thus propose new algorithms that enable concurrent channel inference by measuring one or few cells but inferring all, while minimizing measurement cost via set cover approximations. Our evaluations have confirmed the effectiveness of CA++. Our solution can also be adapted to fit in the current 5G OFDM PHY and the 3GPP framework.

CCS Concepts

• **Networks** → **Mobile networks**; **Network resources allocation**; *Network measurement*; *Network mobility*.

Keywords

Cellular Network, Carrier Aggregation, Cell Measurement, Delay Doppler Domain, CA++

ACM Reference Format:

Qianru Li, Zhehui Zhang, Yanbing Liu, Zhaowei Tan, Chunyi Peng, Songwu Lu. 2023. CA++: Enhancing Carrier Aggregation Beyond 5G. In *The 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23)*, October 2–6, 2023, Madrid, Spain. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3570361.3592500>

1 Introduction

Carrier aggregation (CA) is essential to user experience in 5G and beyond. It combines multiple component carriers (used interchangeably with “cells” in this paper), each operating on a frequency channel, to serve a mobile device. As a result, it

aggregates multiple contiguous or non-contiguous frequency spectrum resources, thereby boosting data throughput over wider spectrum. It is a valuable instrument to utilize precious radio spectrum and boost peak data rates, particularly when 5G is expanding the spectrum space into both mmWave (say, > 24GHz) and new sub-6GHz bands (say, < 6GHz) [12, 16].

In a nutshell, the current CA mandated by 3GPP uses a *cell-based* procedure. It adds a sequence of qualified cells – primary cell (PCell) first, followed by several secondary cells (SCells) – after evaluating their cell quality on a super set of candidates. The effectiveness of CA relies on two factors: (1) accurate channel quality assessment on each added cell; (2) exploration of a large super set from which PCell and SCells are selected. Ideally, the larger the super set, the better selection of cells; the more accurate the cell channel quality assessment, the more gain the CA. However, an exceedingly large super set incurs excessive cell measurements, thus prolonging the CA decision process and degrading the achievable throughput.

We have thus identified two limitations with the current CA, particularly under mobility and with spectrum expansion into the mmWave bands. First, sequential cell measurements prior to selecting cells incurs costs in proportion to the number of candidate cells, thus reducing cell connection time and degrading efficiency upon mobility. On one hand, a small set of candidate cells may lead to the exclusion of better cells. On the other hand, the absence of per-cell measurements results in inaccurate channel estimation. Second, the cell-based aggregation also suffers under mobility. The cell-by-cell addition leads to longer time for CA to complete. If the user roams at high speed, CA might not finish adding all qualified cells. Moreover, since the PCell must be added before all SCells, improper PCell selection may yield a bad set of SCells. The fundamental problem is that, CA should take a panoramic view and select cells on a group basis; otherwise, it may miss better cell combinations.

In this work, we propose CA++, which takes a new perspective on CA and overcomes both limitations. CA++ redefines cell-based operations in the legacy CA. It examines and aggregates a group of cells as a whole, rather than on an individual cell basis. It further parallelizes operations by using few cell measurements to concurrently infer all cells. CA++ works well under both low/moderate mobility (say, walking/driving) and high mobility (e.g., high-speed train). It remains effective over a wide span of radio spectrum.

The key innovations of CA++ are two new algorithms

*Co-primary authors.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ACM Mobicom '23, October 2–6, 2023, Madrid, Spain

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9990-6/23/10.

<https://doi.org/10.1145/3570361.3592500>

that enable concurrent channel quality estimations and minimize the number of cells to measure (§5). These two algorithms operate at two levels. The low-level algorithm uses a single-cell measurement from a given cell tower to infer channels for all other co-located cells. Specifically, it models the channel in the delay-Doppler (DD) domain, which makes cross-channel inference accurate under mobility. This model leverages Doppler shifts (more precisely, *fractional* Doppler shifts), which are not explicitly tailored to time. The high-level algorithm decides what cells to measure, if they are located at different cell towers. We formulate the minimization of cell measurements as the set cover problem, and solve it via greedy approximations. With both algorithms in place, we further adapt the cell-based procedure to group-based operations within the general 3GPP framework.

We implement CA++ on an SDR-based 5G testbed, and evaluate it with real-world experiments and trace-driven emulations. We assess the concurrent channel inference algorithm over a wide span of radio spectrum in two ranges: 2.45GHz - 5.55GHz for sub-6G channels and 58GHz - 62GHz for mmWave bands. The empirical validation shows that, CA++ achieves an SNR inference error of 0.37 dB - 1.03 dB at low mobility (~10 km/h) and 0.39 dB - 0.79 dB at high mobility (>200 km/h). CA++ outperforms all state-of-the-art schemes and reduces the error by more than 63%. To gauge the effectiveness of CA++ under real-world cell deployment, we run trace-driven emulations with five datasets collected from three US operators plus on high-speed trains in China [37]. With concurrent channel inference, CA++ reduces measurement time and prolongs the connection time with new CA cells, by the median value of 190 ms - 765 ms for different operators. The median throughput grows from 35.4 Mbps - 49.0 Mbps to 72.7 Mbps - 84.3 Mbps, while following the original policies. CA++ outperforms the legacy CA under a variety of policies, despite the policy impacts.

Release. Datasets and codes are available at [1].

2 Background on 5G CA

CA in 5G. CA has been introduced to better utilize the increased spectrum since 4G LTE-Advanced [8]. As shown in Figure 1, CA aggregates multiple cells and their frequency channels to serve a mobile device. In a cellular network, a cell is the basic unit to offer radio access, each running over one frequency channel (also called a component carrier). It physically resides in a cell tower, which deploys many cells running over different frequency channels [4]. Given dense cell deployment in 5G [14, 25], a location is typically covered by many (several tens of) cells and several cell towers.

CA works with a *set* of serving cells, rather than an individual cell. Specifically, the cell set in CA consists of a primary cell (PCell) and several secondary cells (SCells). The PCell is responsible for radio resource control signaling upon




Figure 1: Carrier aggregation under mobility.

mobility, while the SCells (and the PCell) are all used for data packet delivery. Note that all SCells and the PCell are co-located in the same cell tower in current 5G releases [9]¹.

CA under mobility. Figure 1 illustrates the procedure to switch to a new cell set for CA. To find the cell set, the network selects each cell separately using the standardized procedure of cell selection handover [7] (described next). It selects the PCell first, and then adds a new SCell in each round until the entire set of SCells is done.

Each cell selection procedure has four steps: ① configure, ② measure, ③ report, and ④ decide. The network first configures the device on how to measure cells and when to report the obtained measurements. The device then switches to configured cells' operating frequencies and conducts measurements. If the candidates are on multiple frequency channels, the device probes each one by one. When any reporting criterion ([3, 7]) (e.g., radio quality greater than a threshold) is met, the device reports the measurements. Once a decision is made, the network instructs the device to switch to the selected cell. The device is served by the new cell thereafter. In this paper, we consider usage scenarios under both low/moderate mobility (say, walking/driving speed) and high mobility (e.g., high-speed train).

Features of CA. From the CA operation and its associated cell selection procedure, we observe several features: (1) CA redefines the single-cell-based selection to become cell-set-based. (2) To add each new cell, the measurement and report on the new cell must be done prior to cell switching. The number of measurements/reports increases in proportion to the number of selected cells. (3) CA is characterized by *sequential* operations. It adds a new cell (PCell or SCell) one by one. (4) PCell is prioritized over SCells in that all SCells will be dropped once the PCell is not used.

3 Design Issues for CA

We describe two identified issues with the current (legacy) CA operations and illustrate them via a real-world example.

Issue 1: Sequential cell measurement. The first issue is that, current CA takes sequential operations on measuring

¹This requirement is relaxed in future 5G releases (say, with Advanced RAN Coordination [17]) and briefly discussed in §10.

the candidate cells, as well as reporting their measurements. This causes problems considering both expanded spectrum space and high mobility. Note that, the sequential measurement is well justified, since the device cannot measure several cells/channels simultaneously.

As more cells are available in 5G and beyond, measuring all cells over the wider frequency spectrum takes more time. Note that, the entire measurement duration increases in proportion to the number of candidates n_c , i.e., $T_{meas} \propto O(n_c)$. Moreover, stable radio quality measurements are typically obtained by assessing reference signals per frequency; they take tens to hundreds of milliseconds. Our experiments show that, the median measurement time per frequency channel is 294 ms with AT&T. Sequential measurement thus becomes unacceptable, given tens of candidate cells to measure in recent 5G deployment [19, 25]. Consequently, the current CA only evaluates a small set of candidate cells under mobility.

Sequential measurements also impede prompt and correct selection of CA cells under high mobility. Note that measurements of the candidate cells cannot complete simultaneously. Results from those early-measured cells may become obsolete as later measurements are yet to come. Consequently, the network may not select the proper cells at runtime.

The impact of mobility on higher-frequency channels (say, those mmWave ones) is more severe. Radio channels remain statistically coherent within channel coherence time [34]: $T_c \propto \frac{1}{f_m} = \frac{c}{v \cdot f}$, where f_m is the maximum Doppler shift, decided by the speed of light c , velocity v , and frequency f . From 2.4 GHz to 39 GHz, higher-frequency channels vary 16× faster at the same velocity, thus making measurement less reliable. Faster measurements are needed to not miss good cells on higher frequency channels.

Issue 2: Sequential, cell-by-cell aggregation. The second issue is cell-by-cell aggregation, which suffers from two downsides under mobility. First, cell-by-cell addition leads to longer time for CA to complete. If the user roams at high speed, CA may only select a subset of candidates and might not complete adding all qualified cells. Second, given the order of adding the PCell first and adding SCells later, improper PCell selection may lead to a bad set of SCells. The fundamental problem is that, CA cannot select cells on a group basis by considering the PCell and SCells together.

An Illustrative Example. We use a real-world instance with AT&T 5G networks to illustrate the above issues.

In Trace 1, the phone is looking for its new cell set. To start, it is configured to search PCell candidates on four 4G channels. Each channel is denoted by F with a unique channel number specified by 3GPP [8, 10]. The phone performs sequential measurements for about 1.3 seconds until finding the first good cell on channel F9820 (centered on 2355 MHz). Note that the “good” cell is determined by the policy below,

```
20:19:31.945 Configure measurement on frequency
channel F5110, F850, F9820, F66486.
20:19:32.054 Measure cell #478 on F5110, RSRQ=-15dB
20:19:32.405 Measure cell #462 on F66486, RSRQ=-16dB.
20:19:33.045 Measure cell #61 on F9820, RSRQ=-9dB.
20:19:33.365 Measure cell #244 on F66936, RSRQ=-8dB.
20:19:33.685 Report cell #61 on F9820.
20:19:33.738 Select cell #61 on F9820 as PCell.
20:19:33.738-20:19:37.222 Search 4G SCell(s) on
frequency channel F850,F5110,F66486,F66661,F66936;
Finally add SCell #370 on F66661.
20:19:33.886-20:19:37.222 Search 5G SCell(s) on
mmWave channels F2253331,F2251665,F2254997, and
sub-6G channel F174270; Finally add SCell #561 on
F174270 (5G sub-6G).
```

Trace 1: An illustrative example observed in AT&T.

Measuring order	(Serving PCell: RSRQ = -18dB)	RSRQ (dB)	Throughput (Mbps)	No change	Inter: 3 -> 1 Intra: 5 -> 3	Inter: 3 -> 5 Intra: 5 -> 7
	Cell (F5110,#478)	-15	46.6-65.5	✗	✓	✗
	Cell (F9820,#61)	-9	15.7-54.3	✓	✓	✓
	Cell (F66936,#244)	-8	79.3-95.1	✓	✓	✓
	Cell (F850,#306)	-12	148.0-216.7	✓	✓	✗

✗ Not meet criteria
✓ Meet criteria

Selected by legacy CA (Cell F9820, #61)
Best candidate cell (Cell F850, #306)

Table 1: Legacy CA selection v.s. best candidate.

which specifies the criterion to trigger a measurement report via event A3 [3, 7] (more policies in Table 5):

$$\begin{cases} RSRQ_{candidate} > RSRQ_{serving} + 3dB, & \text{for intra-freq cells} \\ RSRQ_{candidate} > RSRQ_{serving} + 5dB, & \text{for inter-freq cells} \end{cases}$$

An intra-freq (inter-freq) candidate cell runs over the same (different) channel of the serving one (here, F5110). The first two measurements over F5110 and F66486 do not meet the reporting policy ($RSRQ_{serving} = -18dB$, Table 1). Upon receiving the measurement report of cell #61 over F9820 ($RSRQ_{candidate} = -9dB$), the network selects it as the new PCell. Afterwards, it takes 3.5 seconds to measure channels on 4G, 5G sub-6G and 5G mmWave bands; Finally, another 4G cell and a 5G sub-6G cell are added as SCells. This cell set achieves data speed of about 15.7 – 54.3Mbps. However, our analysis shows that there exists a better PCell option on channel F850 (1955 MHz). Along with other SCells including 5G mmWave cells, this new cell set achieves data speed of 148.0 – 216.7 Mbps (using cell #306 over F850 as the new PCell, Table 1). Due to sequential measurements (Issue 1), this cell is missed, as the network immediately selects the first good cell to avoid connection loss. Moreover, SCells are added by the new PCell, which does not equally consider all available cells, but a subset of cells (Issue 2). As a result, the network does not assess all available cells and makes a worse choice.

We notice that, CA performance is also affected by the operator policies, specifying which cells are preferred, how many are to be aggregated at most, and what are the minimal RSRP/RSRQ thresholds to be considered for CA, to name a

few. Next, we use the same example to illustrate policy impacts on CA and show that addressing the identified two issues can boost CA performance with various policies.

Table 1 summarizes the results with these two policy changes. For simplicity, we list statistical throughput per PCell over extensive speedtest experiments at the same location (methodology in §8.1). We tune the policy parameters to impact whether a candidate PCell is reported (here, A3 thresholds). Legacy CA follows the same order to perform *sequential* measurements while the best CA makes the decision by getting the measurements of *all* qualified cells, which are determined by the reporting policies.

First, we reduce A3 thresholds (here, by 2dB) and thus include more candidates. In this case, legacy CA would pick another cell over F5110 as the PCell and reach a slightly better cell set (46.6–65.6Mbps). However, it still misses the best cellset which uses cell #306 over F850 as the PCell because of sequential measurements and aggregation (Issue 1 and Issue 2). Second, we increase A3 thresholds (by 2dB), which shrinks the candidate pool. In this case, the legacy CA retains the same selection. However, the best CA choice changes under this policy because cell #306 over F850 is not qualified for reporting; It chooses another cellset using a PCell over F66936. The performance headroom is reduced due to this policy change. From the above example, we can see that tuning policies may improve or reduce CA performance at certain locations. Indeed, policy tuning may offer another dimension for design. However, it cannot fully address the performance issues caused by sequential CA operations.

In this work, we focus on the technical solution to sequential CA operations, under the current policy constraints. We seek to devise new algorithms that will work with existing policies. We consider the impacts of policies in our evaluation, using the original policies, as well as policy variations (§8.4).

4 Overview of CA++

Goals and challenges. CA++ seeks to achieve two goals by addressing three technical challenges.

The first goal is concurrent channel estimation for all cells in the super set upon mobility. Ideally, the super set includes every candidate cell the CA seeks to evaluate and select from. Note that, these cells can be deployed on multiple cell towers in the spatial proximity of the mobile device.

To achieve the above goal, we address two issues:

- Given multiple, colocated cells on a single cell tower, *how to measure or infer channel quality of all the cells quickly?*
- Given the super set of cells on multiple towers, *how to minimize the measurements cost to cover all the cells?*

The second goal is concurrent cell aggregation. We seek to consider all to-be-aggregated cells together, thus departing from the current paradigm of cell-by-cell exploration and addition. To this end, we must address the third issue:




Figure 2: Overview of our solution CA++.

- *How to enable group-based, rather than cell-by-cell, aggregation within the 3GPP framework?*

Solution. CA++ uses three components (shown in Figure 2), which work in concert to address all three issues:

- **New algorithm for concurrent channel inference via once-for-all measurement.** Our algorithm accurately models the shared multipaths between the device and cell tower, and infers colocated cells from a single measurement. It works well under mobility and with wide spectrum span.
 - **New algorithm for minimal cell measurement via set cover problem abstraction.** CA++ embraces cells on different towers, and minimizes the incurred measurement cost. It abstracts this issue into a classic set cover problem and provides efficient approximations.
 - **Group-based management on aggregated cells.** Given the two algorithms, CA++ further adapts to the current CA procedure and enables group-based cell aggregation within the 3GPP CA framework. It covers the entire process of measurement, feedback/report, and decision making.
- Figure 2 illustrates how different modules work together in CA++. Concurrent channel inference runs two algorithms, which infer co-located cells per tower from one measurement and cover all towers. Both algorithms work with the group-based measurement module within the 3GPP framework. The obtained inference results are sent via group-based feedback, which subsequently triggers group-based decision for CA.
- The three components are next elaborated in §5 and §6.

5 Algorithms for Concurrent Channel Inference: Measure Few, Yet Infer All

We design two new algorithms for concurrent channel inference with minimal cell-level measurements. The first algorithm exploits the channel representation in the DD domain to infer all cells colocated on the same tower using a single cell measurement (§5.1–§5.3). The second algorithm describes our set cover abstraction on cell measurements to different towers (§5.4). Its approximation scheme tends to minimize cell-level measurements upon mobility.

5.1 Limitations of Time-frequency Domain

Channels model in the time-frequency (TF) domain involves Doppler shift tailored to time and it is hard to accurately infer the channel quality under mobility. The wireless channel in

the TF domain is represented as $h(t, f) = \sum_{i=1}^P a_i e^{j2\pi(t v_i - f \tau_i)}$, where P is the total number of propagation paths, and the i -th path has path attenuation a_i , delay τ_i , and Doppler shift v_i upon mobility. It is clear from the above expression that the channel at time t is coupled with Doppler shift v on each path. We cannot separate those paths and extract each Doppler shift. Consequently, the time-frequency channel cannot be accurately estimated under mobility.

Prior efforts simplify $h(t, f)$ as $\hat{h}(t, f) = \sum_{i=1}^P a_i e^{-j2\pi \frac{d_i}{\lambda} + j\Phi_i}$, where d_i is the distance, λ is the wavelength, and Φ_i is the phase shift. Both OptML [13] and R2F2 [35] assume a static setting on channel estimation; the phase shift is assumed to be invariant across different frequency channels. This assumption does not hold under high mobility, where Φ_i would vary along with the cell's frequency.

5.2 Inference in Delay-Doppler Domain

We switch to the DD domain because the channel representation explicitly models the Doppler shifts under mobility. Moreover, the multi-path parameters can be readily mapped from one frequency to another. The DD channel directly models those multiple propagation paths [32]:

$$h(v, \tau) = \sum_{i=1}^P a_i \delta(v - v_i) \delta(\tau - \tau_i), \quad (1)$$

where P is the total number of propagation paths. The i -th path is characterized by its path attenuation a_i , delay τ_i and Doppler shift v_i . $\delta(\cdot)$ is the Dirac delta function.

In the DD domain, the multi-path parameters $\{(a_i, v_i, \tau_i)\}_{i=1}^P$ are not coupled with time. They can be converted across different frequencies. Path attenuation and delay are invariant of frequency; The Doppler shift is derived from the frequency and an invariant coefficient, i.e., $\frac{v}{c}f$, where v, c, f are the moving speed, the speed of light and frequency, respectively.

Infer channels in the delay-Doppler domain. Here is the conceptual procedure: We measure the channel response h_A from a cell on frequency f_A , retrieve the multipath profile, and map it to the profile of another cell on frequency f_B .

In practice, discrete signal symbols are transmitted over the wireless channel. Therefore, the channel response is a matrix spread on a grid with Doppler span $N\Delta v$ and delay span $M\Delta\tau$. The DD grid is associated to another $N \times M$ grid in TF domain, with quantization steps: $\Delta v = \frac{1}{N\Delta t}$ and $\Delta\tau = \frac{1}{M\Delta f}$, where Δt is the symbol duration and Δf is the sub-carrier spacing (Figure 3). Applying the channel response matrix to original symbols $x_{DD}[k, l], k = 0, \dots, N-1, l = 0, \dots, M-1$ generates $y_{DD}[k, l]$ at the receiver:

$$y_{DD}[k, l] = \frac{1}{NM} \sum_{k'=0}^N \sum_{l'=0}^M \mathbf{H}[k', l'] x_{DD}[k' - k, l' - l] \quad (2)$$

The channel response matrix can be derived from the multi-path profile [21, 32]:




Figure 3: Concurrent channel inference in delay-Doppler domain (DD embedded into OFDM).

$$\mathbf{H}[k, l] = h_w(k\Delta v, l\Delta\tau) = \sum_{i=1}^P a_i e^{-j2\pi v_i \tau_i} \mathcal{G}(k\Delta v, v_i) \mathcal{F}(l\Delta\tau, \tau_i)$$

$$\mathcal{G}(v, v_i) \triangleq \sum_{k'=0}^{N-1} e^{-j2\pi(v-v_i)k'T}, \quad \mathcal{F}(\tau, \tau_i) \triangleq \sum_{l'=0}^{M-1} e^{j2\pi(\tau-\tau_i)l'\Delta f}.$$

From the measured channel response, we can retrieve the multi-path parameters of the cell on f_A , say $\{(a_i^A, v_i^A, \tau_i^A)\}_{i=1}^P$. For the channel for cell on f_B , we can derive its multi-path parameters, with $v_i^B = v_i^A \frac{f_B}{f_A}$ and a_i^A, τ_i^A being invariant.

A new challenge: wider radio spectrum. The recent algorithm REM [26] also exploits the DD domain. However, its accuracy declines when being applied over a wider spectrum (say, 5G and beyond). This is stemmed from its assumption that each path's delay and Doppler are exactly *integer* multiples of the quantization steps on the DD grid. Equivalently, each path is located on the discrete DD grid with integer coordinates $\kappa_i = \frac{v_i}{\Delta v}, l_i = \frac{\tau_i}{\Delta\tau}$ (Figure 3). It is the key to its path decoupling via singular value decomposition (SVD).

However, the above assumption for REM does not hold over the mmWave bands. Instead, we must consider *fractional* Doppler for more accurate inference. The Doppler quantization step ($\Delta v = \frac{1}{N\Delta t}$) runs large, thus being coarse-grained at high frequencies. Note that the quantization step of Doppler $\Delta v = \frac{1}{N\Delta t}$ depends on the grid duration $N\Delta t$. In a typical 5G setting with $f_c = 38$ GHz, $N\Delta t = 5$ ms², the resolution with integer Doppler $\Delta v = 200$ Hz is pretty low, equivalent to 5.7 km/h in terms of moving speed difference.

5.3 Path Decoupling over Wide Spectrum

We next show how to accurately retrieve the multipath profile over a wide spectrum. To make it work precisely with mmWave channels, the key is path decoupling under *fractional* Doppler shifts. CA++ uses the delay factor to decouple those paths in the delay-Doppler domain.

Decouple multiple paths by delay. We decouple multiple propagation paths based on a critical observation: Paths are separated along the delay index.

Specifically, each path has a distinct, integer delay coordinate (i.e. $l_i = \frac{\tau_i}{\Delta\tau}$). This holds true in reality for two reasons.

²Based on the duration of reference signal in 5G (i.e. SSB burst set) [11].

(a) The paths in a single cell are sparse compared to the range for delays. Given limited reflectors, the number of propagation paths is much smaller than the range of delay coordinates (typically over 200), i.e. $P \ll M$. (b) The delay quantization step is fine-grained ($\Delta\tau$). A mmWave cell with 100 MHz channel-width has a delay step of 10 ns. It corresponds to 3 m of difference in path length. Such a fine resolution makes path decoupling feasible in most scenarios. Measurement studies [30, 36] have shown that outdoor paths are separated by over 100ns under high mobility. With indoor setting, propagation paths which get closer can be merged and characterized by one set of parameters. Our indoor experiments (§8.2) confirmed that the algorithm works well and outperforms the prior work.

Given the insight, we next derive the mathematical form for decoupling. Denote the channel response matrix with M vectors, i.e., $\mathbf{H} = [\vec{\mathbf{h}}_1, \vec{\mathbf{h}}_2, \dots, \vec{\mathbf{h}}_M]$. The i -th path with delay coordinate l_i is only associated with its vector $\vec{\mathbf{h}}_{l_i}$. We thus decouple the multiple paths by taking the corresponding vectors. Theorem 5.1 describes this one-to-one mapping (see proof sketch in §A in [29]).

THEOREM 5.1 (PATH DECOUPLING). *Given any two paths with distinct delays (i.e., $l_i \neq l_j, \forall i \neq j$). For any vector $\vec{\mathbf{h}}_{l_i}$, the following holds true: (i) If there exists a path i such that $l = l_i$, the k -th ($0 \leq k \leq N - 1$) element of the vector is given by: $\vec{\mathbf{h}}_{l_i,k} = Ma_i e^{-j2\pi v_i \tau_i} \mathcal{G}(k\Delta v, v_i)$; (ii) Otherwise, we have $\vec{\mathbf{h}}_{l_i} = \mathbf{0}$.*

Extract parameters for each path. With separated paths, we extract each path's parameters from the corresponding vector in the channel response matrix. From Theorem 5.1, the i -th propagation path with delay coordinate l_i is associated with $\vec{\mathbf{h}}_{l_i}$. Therefore, we have N relations between each element in the vector and the path parameters $\{a_i, \kappa_i, l_i\}$:

$$\vec{\mathbf{h}}_{l_i,k} = Ma_i e^{-j2\pi v_i \tau_i} \frac{k - e^{-j2\pi(k-\kappa_i)}}{k - e^{-j\frac{2\pi}{N}(k-\kappa_i)}}, k = 0, \dots, N - 1. \quad (3)$$

We solve these equations to retrieve the i -th path's parameters from vector $\vec{\mathbf{h}}_{l_i}$. Note that we already know the index of path delay (say l_i) through path decoupling. For the Doppler shift, we divide $\vec{\mathbf{h}}_{l_i,0}$ by $\vec{\mathbf{h}}_{l_i,\frac{N}{2}}$ and obtain:

$$\kappa_i = \frac{N}{\pi} \left(x\pi \pm \arg \cot \left| \frac{\vec{\mathbf{h}}_{l_i,0}}{\vec{\mathbf{h}}_{l_i,\frac{N}{2}}} \right| \right), \quad (4)$$

where x is an integer. Given the range $0 < \kappa_i < N$, (4) has two possible values in that range. We further locate the correct one by checking it satisfies $\left| \sin \frac{\pi}{N} \cot \frac{\pi \kappa_i^*}{N} - \cos \frac{\pi}{N} \right| = \left| \frac{\vec{\mathbf{h}}_{l_i,0}}{\vec{\mathbf{h}}_{l_i,1}} \right|$.

Finally, we derive path attenuation $a_i = \frac{1}{M^2 N^2} \sqrt{\sum_{k=0}^{N-1} |\vec{\mathbf{h}}_{l_i,k}|^2}$.

The above algorithmic operations apply to both static and mobile scenarios. In static cases, the Doppler shift κ_i becomes zero. We derive $\kappa_i = 0$ by checking $\vec{\mathbf{h}}_{l_i,k} = 0$, for $k \neq 0$.

Algorithm 1. Here is the full algorithm for channel inference. Initially, we obtain the channel response matrix $\mathbf{H}_A$ of cell C_A on frequency f_A through measurement. Our goal is to infer the channel matrix $\mathbf{H}_B$ of its co-located cell C_B on frequency f_B and estimate its radio quality. Based on Theorem 5.1 (path decoupling), each vector with non-zero values in the original channel matrix reveals one propagation path (Line 2). We separate paths by taking the non-zero vectors and deriving the parameters (Lines 3-7). Next, we project the shared parameters onto C_B on frequency f_B . Path attenuation and delay (a_i, τ_i) remain invariant with frequency. The Doppler shift $v_{i,B}$ is in proportion to the frequency; it is equal to $v_{i,A} \frac{f_B}{f_A}$ (Line 6). With parameters for all paths extracted and mapped to frequency f_B , we reconstruct the channel matrix $\mathbf{H}_B$ (Line 8). Finally, we infer the radio quality (here, SNR, RSRP and RSRQ) for C_B (Line 9).

Algorithm 1 ONETOWER: Once-for-all measurement

Input: Cell C_A 's channel matrix $\mathbf{H}_A$, frequency f_A , delay-Doppler grid $N_A, M_A, \Delta v_A, \Delta \tau_A$; Co-located Cell C_B 's frequency f_B , grid setting $N_B, M_B, \Delta v_B, \Delta \tau_B$

Output: C_B 's radio quality

- 1: $i = 1, \mathbf{P} = \emptyset$;
 - 2: **for** each column vector with non-zero values in $\mathbf{H}_A$ **do**
 - 3: $l_i \leftarrow$ the index of the vector; $\tau_i \leftarrow l_i \Delta \tau_A$;
 - 4: Derive κ_i based on (4); $v_{i,A} \leftarrow \kappa_i \Delta v_A$;
 - 5: Derive a_i from $\vec{\mathbf{h}}_{l_i,k}$;
 - 6: $v_{i,B} \leftarrow v_{i,A} \frac{f_B}{f_A}$;
 - 7: $\mathbf{P} \leftarrow \mathbf{P} \cup \{(a_i, v_{i,B}, \tau_i)\}, i \leftarrow i + 1$;
 - 8: Compute $\mathbf{H}_B$ based on (3) and path parameters in $\mathbf{P}$, and the grid setting $N_B, M_B, \Delta v_B, \Delta \tau_B$;
 - 9: Compute SNR/RSRP/RSRQ of C_B ;
-

Complexity. The algorithm has polynomial computation complexity of $O(NMP)$. The overhead stems from channel matrix reconstruction for the cell to infer, which dominates the cost of retrieving shared multipaths ($O(P)$). Our algorithm is faster than those solutions in the time-frequency domain [35], which are based on non-convex optimizations. It also outperforms the recent DD-domain algorithm [26] by reducing the processing cost by a factor of $\frac{\max(N,M)}{P}$.

Note that the frequency gap between sub-6GHz and mmWave channels may result in unpredictable discrepancy in physical propagation conditions. Therefore, channel inference is operated separately on two frequency ranges (here, 2.5 – 6GHz for sub-6G, and 58 – 62 GHz for mmWave). This is consistent with the real-world deployment, where cells in different frequency ranges reside on different towers.

5.4 Inference for All Cell Towers

Algorithm 1 (§5.3) accelerates channel inference among cells on the same tower. Given that multiple towers are deployed

in the neighborhood of the mobile device, we thus need to decide on what to measure and what to infer.

We seek to minimize the number of frequency channels to measure, thus speeding up CA++. Note that when the device measures a frequency channel, it acquires measurements for all cells on the same frequency channel. We illustrate the design issue with another real-world instance. At one location, we have observed two 5G cell towers, CT-A and CT-B. CT-A carries 4 cells on the mmWave frequency channels, denoted as F1 - F4. CT-B has 8 mmWave cells: 4 of them on the same 4 frequencies as CT-A; the other 4 are denoted as F5 - F8. By measuring only *one* frequency channel shared by both towers, say F1, the device can measure the cells over F1 on both towers and infer all other cells using Algorithm 1.

Set Cover Abstraction. We abstract the above issue as an equivalent set cover problem, where we aim to minimize the frequency channels to *physically* measure. The measured and inferred frequencies shall cover all cells on every tower in the neighborhood of the mobile device. The formulation can be formally described as:

PROPOSITION 5.1. *To find the minimum number of frequency channels to measure is equivalent to the following set cover problem. A universal set $S = \{CT_1, CT_2, \dots, CT_Q\}$ represents Q neighbor cell towers. Cell tower i carries k_i frequency channels. There are c unique frequency channels; each frequency f_j can be represented by a non-empty subset $S_j \subset S$ including towers that carry the frequency. We have $\max_{1 \leq i \leq Q} k_i \leq c \leq \sum_{i=1}^Q k_i$. Moreover, CT_i occurs k_i times among all subsets. Hence, minimizing the number of frequencies to measure is equivalent to a set cover problem, i.e., to find the minimal index sets $I \subset \{1, 2, \dots, c\}$ s.t. $\bigcup_{i \in I} S_i = S$.*

Algorithm 2. We approximate the optimal solution to the NP-hard set cover problem and develop Algorithm 2 over greedy optimizations. The algorithm incurs low overhead, with time complexity of $O(\bar{k}Q^2)$. For its effectiveness, Theorem 5.2 offers the upper bound on the number of measured frequency channels (the full proof in the Appendix of [29]).

Algorithm 2 ALLTOWERS: Minimizing cost of measure

Input: A universal set $S = \{CT_1, CT_2, \dots, CT_Q\}$, subsets $\{S_1, S_2, \dots, S_c\}$

Output: Index set I representing selected subsets

- 1: $I \leftarrow \emptyset, X \leftarrow S$
- 2: **while** $X \neq \emptyset$ **do**
- 3: Let i be the index maximizing $|X \cap S_i|$
- 4: $I \leftarrow I \cup \{i\}, X \leftarrow X \setminus S_i$

THEOREM 5.2. *Assume Q cell towers, cell tower CT_i carries k_i channels, and c distinct frequencies ($\max_i k_i \leq c \leq \sum_i k_i$) in total. The number of channels ALLTOWERS measures $\leq$*

$$\min \left\{ \left\lceil \frac{\log(Q - \max_{1 \leq j \leq c} |S_j|)}{\log \frac{c}{c - k_m}} \right\rceil + 2, Q \right\}, \text{ where } k_m = \min_{1 \leq i \leq Q} k_i.$$




Figure 4: Frequencies to measure (Legend for $Q, \bar{k}$).

Figure 4 compares the upper bound of measured frequencies by ALLTOWERS with the legacy per-frequency cell measurement. We use three parameters to characterize the deployment of cell towers and frequencies: $Q, \bar{k}, c$ as the number of cell towers, the *average* number of frequency channels carried by a tower, and the total number of frequency channels nearby. The legacy scheme measures c frequency channels to obtain the global view, invariant with $Q, \bar{k}$. We consider practical settings of $Q = 3, 6$ and $\bar{k} = 4, 6, 8$ based on empirical data. ALLTOWERS reduces the measurement time by a factor of $2.3\times - 11.5\times$. The efficiency depends on whether the towers have more similarities (small c) or more heterogeneities (large c) in deployed frequency channels. For example, measuring a frequency channel shared by all towers (highest similarity) would cover all cells. In an extreme case of highest heterogeneity, towers do not share frequency channels. The device has to measure one frequency per tower. The measurement is still accelerated by a factor k , as $\frac{c}{Q} = \bar{k}$.

6 Group-Based CA Management

We next present our group-based operations within the 3GPP CA framework. To make full use of our algorithms in §5, we further adapt the cell-based procedures of measurement, feedback and decision making to group-based operations.

6.1 Group-Based Measurement

We first enable group-based measurements within 5G PHY. As shown in Algorithm 1, a single cell measurement is used to infer the entire group of cells residing on the same tower. We next address two issues: (1) How to multiplex reference signals in the standardized OFDM PHY? (2) How to cope with heterogeneity among co-located cells inside a group?

Embedding OTFS signals over OFDM. To perform measurement in DD domain, we need to use OTFS, a DD modulation scheme [20], for reference signals. We achieve this by embedding OTFS signals over the current OFDM-based PHY. 5G NR uses two types of reference signals for cell measurements: SSB blocks and CSI-RS blocks [11]. Signals from different cells are allocated on orthogonal OFDM grids (i.e., different time slots or subcarriers). CA++ places OTFS symbols at the same locations as SSB burst and dynamic CSI-RS grids. In Figure 3, a 4x2 grid for Cell 1 is scheduled for SSB. CA++ transforms OTFS signals into that 4x2 OFDM grid. We thus obtain OTFS measurements via the OTFS overlay on top of existing OFDM grids. The transformation between OTFS and OFDM is realized with ISFFT/SFFT. We further transform

Type	Group-based event
Select group	(G-P) Neighbor eligible as PCell, $Q_n > \theta_{G1}$
	(G-S) Neighbor eligible as SCell, $Q_n > \theta_{G2}$
SCell update	(U) SCell worse than threshold, $Q_s < \theta_U$

Table 2: Group-based criteria.

the measurements into standard-compliant SNR/RSRP/RSRQ metrics on radio signal level and quality [11]. RSRP and RSRQ are used for cell selection in the current practice [14].

Note that, the OTFS scheme is used for reference signals only, whereas OFDM is still used for other data types (e.g., packet delivery). We thus reduce the prohibitive cost of completely revamping the OFDM based modulation at PHY.

Tackling heterogeneity inside a group. In our group-based measurement, cells in the same group might use heterogeneous OFDM PHY configurations (i.e., numerology and subgrid size). We cope with such heterogeneity by extracting multi-path profiles in the delay-Doppler domain that are *agnostic* to the PHY configurations.

The first diversity is numerologies (i.e., symbol duration Δt and sub-carrier Δf^3). Consequently, the transformation to the DD domain also carries diverse numerology ($\Delta\tau, \Delta\nu$). Figure 3 illustrates an example of inferring Cell-2 based on Cell-1 while they use distinct numerologies. We first estimate Cell-2's multi-path parameters based on parameters retrieved from Cell-1. We then map those parameters to the coordinates on the OTFS grid with Cell-2's numerology.

The second diversity is that, cells in the same group might use different subgrid sizes due to different scheduling of reference signals. With fractional Doppler covered, our inference algorithm (see Algorithm 1) can tolerate flexible choices of grid size (i.e., N, M). As shown in Figure 3, Cell-2 has a small N and thus a large Doppler step; The propagation path can still be represented with fractional Doppler coordinates.

6.2 Group-Based Feedback

We further devise a group-based feedback scheme within the 3GPP framework. The 5G specifications organize measurement reports by cells' frequencies; this is designated for the legacy single-cell operation. Since CA involves cells operated on different frequency channels, the current practice would incur isolated reports; thus it avoids evaluation and addition of PCell and SCells together. We propose group-based feedback to stay aligned with group-based cell aggregation and reduce the signaling overhead.

Feedback to change PCell and SCell(s) together. We devise group-based, hierarchical criteria of sending reports (Table 2). Our criteria are structured with the minimum requirement to become PCell and SCell, denoted as conditions **G-P**(primary) and **G-S**(secondary), respectively. A group of cells on the same tower is reported if and only if one or more

cells meet the PCell criterion; the report also includes all cells which meet SCell criterion.

Feedback to update SCell(s) only. We further support reports to update SCell(s) w/o changing the current PCell, denoted as criterion **U**. It is triggered when any SCell becomes weaker than a threshold. The report also includes candidates whose channel quality is above a threshold, from the same tower as the current PCell. Unlike the legacy CA, CA++ could generate such reports without physically measuring any candidate; all colocated cells on the same tower can be inferred with Algorithm 1.

6.3 Group-Based Decision

CA++ have also cleared roadblocks to make group-based decisions. With accurate radio quality reports on the cells, the network can assess each group as a whole and select the PCell and SCells concurrently. Here CA++ adopts a straightforward selection strategy to maximize the aggregated channel-width. In reality, operators have the freedom to use any selection logic they prefer. However, it is an independent topic beyond the scope of this work. Instead, CA++ focuses on the algorithmic aspects for concurrent selection of the PCell and SCells based on accurate channel inference under mobility and over wide spectrum (sub-6GHz and mmWave bands). It works in parallel with any decision logic.

7 Implementation

We have implemented CA++ with USRP, acting at both sides of a cell tower and a mobile device. The prototype is built on an open-source 5G/4G testbed Flora [24]. It incurs minor changes to the device and infrastructure after integration.

Network side. Our implementation at the network follows the standardized cell selection for CA. The key change is to perform each operation on the entire cell set, rather than a single cell. It involves changes in a few steps. First, we implemented Algorithm 2 to decide which cells to measure and which to infer among all cells. Second, we added the group-based criteria to trigger reports. The new form is a set of standardized criteria for single cells (e.g., A1-A6 in [7]). Finally, we implemented group-based selection. The current version used the basic scheme that prefers the largest aggregated channel-width. For PHY transmission, we change only reference signals (using OTFS) over Flora. Specifically, we revise a subgrid of OFDM symbols over 12 subcarriers and 14 reference symbols for OTFS-based measurements.

Mobile device. We made two updates at the device. First, we implemented Algorithm 1. In order to mitigate false-positive paths caused by interferences and noises, we eliminate the weak paths that yield attenuation 30 dB lower than the strongest one. Second, we report measurements on the cell-set basis in concert with the group-based operations.

³5G currently supports 4 numerologies, with 15kHz, 30kHz, 60kHz, 120kHz as subcarrier spacing Δf (extended to 960kHz in recent releases [7]).




Figure 5: Experimental settings.

8 Evaluation

Using the methodology (§8.1), we assess each component, i.e., inference accuracy (§8.2), effectiveness of group-based management (§8.3), followed by its overall performance (§8.4).

8.1 Methodology

Our evaluation uses both testbed-based experiments and trace-driven emulations. We gauge the accuracy and efficiency of concurrent channel inference with the SDR testbed under real-world channels. We further quantify the throughput boost by CA++ and its efficiency using real-world traces. **Testbed and experiment settings.** Figures 5a and 5b show our testbed setup. The testbed supports both sub-6G and mmWave bands. The cell tower uses USRP X300 for sub-6G, and operates with the 60GHz HMC6350 TX frontend [15] for mmWave. USRP N210 acts as the sub-6G client device; USRP X300, together with HMC6350 RX frontend and horn antennas, serves as the mmWave client. The USRPs run baseband processing and feed signals to sub-6G and mmWave frontends. They are connected to the server running Ubuntu 18.04 (Intel i7-9700k CPU with 32G RAM) for baseband signal generation. We use the testbed to conduct indoor (at 14 locations in a 10.2 m × 11.2 m lab, Figure 5c) and outdoor experiments (Figure 5d), in both static settings and low-mobility (≤ 10 km/h) scenarios.

Trace-driven emulation. Since CA++ is not deployed by mobile carriers, we use large-scale datasets collected from operational networks to perform trace-driven emulations. Effectively, we conduct a “what-if” study with two steps. First, we run walking/driving tests with file downloading tasks on phones. Cell deployment, throughput and signal strengths are all recorded. Meanwhile, we extract logs on CA operations using MobileInsight [27]. Second, for each CA instance, we assess *all* available cell combinations and remake the choice on a group basis by exploiting historical data at the given location. Since we could not measure RSRP/RSRQ and throughput for any *unselected* cell set at runtime, we use the median values measured in the historical runs. This approximates the cell conditions and performance in case the device connects to them. Table 3 shows the five datasets used for emulations: (1-4) **A-C1, V-C1, T-C1 and A-C2**: We run tests with three major operators (A, V and T short for AT&T, Verizon, and T-Mobile, respectively) in one US Midwest city (C1), and with A in one large city (C2). A and V support 5G over both sub-6G and mmWave bands in the test areas, whereas T supports only sub-6G. Since we need sufficient data on cell

Dataset	A-C1	V-C1	T-C1	A-C2	HST [37]
Date	Apr 2021 - March 2022				Nov 2018
Region	1.65×1.85 km ²				1,300-km
Speed (km/h)	driving: 10-40 (mostly); walking: <5				300 - 350
RAT	5G + 4G (T supports 5G over sub-6G only)				4G
Max# CA CH	6	5	5	7	3
# CA groups	5,681	2,037	492	3,031	534
Max CA CW	430 MHz	140 MHz	135 MHz	445 MHz	50 MHz

Summary of 5G cells					
# sub-6G cells	62	21	55	38	N/A
# mmWave cells	372	100	0	39	
# sub-6G CH	3	1	3	2	N/A
# mmWave CH	16	4	0	9	
sub-6G freq.	826–2116	885	626–2608	826–2116	N/A
mmWave freq.	38.6–39.5	27.9–28.3	0	38.6–39.5	
sub-6G CW	5	10	15–100	5	N/A
mmWave CW	100	100	N/A	100	

Summary of 4G cells					
# cells	1,719	1,228	878	1,490	1,910
# Channel	20	40	12	18	8
Freq. (MHz)	709–5824	701–5825	701–2539	709–5824	1740–2155
CW (MHz)	5/15/10/20 MHz				

Table 3: Datasets. CH: Channel. CW: Channel-Width.

Sub-6G freq. are in MHz. MmWave freq. are in GHz. deployment and CA usage, we run extensive experiments to scan the test regions (over 5,700 km and 400 hrs in total). We keep downloading files (500MB each from Google Cloud) to measure the downlink throughput. (5) **4G-HST**: We use a public dataset [37] collected on High-Speed-Train (HST) commuting between Shanghai and Beijing, China. 5G was not deployed on HST yet.

8.2 Concurrent Channel Inference

We use the SDR testbed to assess concurrent channel inference under a variety of experimental settings on frequency bands, mobility, locations in both indoor scenarios (in a lab space of Figure 5c with its floor plan in Figure ??) and outdoor scenarios (at the top level of a parking structure of Figure 5d). We compare CA++ with prior schemes in the delay-Doppler domain (REM [26]) and in the time-frequency domain (OptML [13] and R2F2 [35]). Both OptML and R2F2 require to configure the number of paths for high accuracy. We use their optimal configurations in our tests.

Inference accuracy under mobility. We first assess the inference accuracy under both low and high mobility. We run low mobility experiments using a mobile cart that carries the client devices and moves at the speed of about 10 km/h. We emulate high mobility in outdoor scenarios. We extract the propagation model from the low-mobility traces and replay these traces with the mobility settings collected from the HST dataset (i.e., 150 km/h – 300 km/h). The propagation paths are extracted by fitting the channel responses with the optimal Doppler shifts $\{(v_i)\}_{i=1}^P$. To replay the propagation paths under high mobility, we scale the Doppler shifts according to the moving speed v , i.e., $v_i = \frac{v}{c}f$. We use the scaled Doppler shift to evaluate CA++ under high mobility. This scaling follows the practice in 3GPP channel models (e.g., [5]).




Figure 6: SNR inference errors under low mobility.




Figure 7: SNR inference errors under high mobility.

We consider two frequency ranges: (1) sub-6GHz (i.e., measuring over 2.45 GHz to infer four unlicensed sub-6G channels at 2.55 GHz/3.65 GHz/5.45 GHz/5.55 GHz) and (2) mmWave (i.e., measuring over 58 GHz to infer 62 GHz). We assess the inference accuracy in terms of SNR/RSRP/RSRQ errors. The results are similar (see next); we only present SNR results in Figures 6 and 7, due to space limit.

CA++ outperforms all the state-of-arts under low mobility. For inference within sub-6G, it yields a median error of about 1 dB ([0.91 dB, 1.03 dB]) when measuring 2.45 GHz. In contrast, REM incurs a median error of about 3 dB ([2.66 dB, 3.21 dB]), while OptML and R2F2 perform much worse with the median errors > 6 dB. CA++ outperforms over mmWave as well. Its error even goes down to 0.37 dB, much lower than 2.03 dB, 1.92 dB, and 4.60 dB by REM, OptML and R2F2. All schemes perform better over mmWave, because the measurement-inference frequency gap becomes relatively smaller (4GHz over 58–62 GHz versus > 1 GHz over 2.45–5.55GHz except in the 2.55 GHz case). CA++ performs much better, because its delay granularity increases with the center frequency. We use the median error unless specified.

CA++ outperforms other approaches even more under high mobility. Compared to low mobility, the error is smaller over sub-6G (1dB $\rightarrow$ 0.4 dB) and slightly larger over mmWave (0.37 dB $\rightarrow$ 0.64 dB). REM, OptML and R2F2 all perform worse over mmWave under high mobility. For sub-6G bands, the inference errors grow to 3.94 dB–4.16 dB (REM), 6.3 dB–10.2 dB (OptML) and 8.53 dB–11.0 dB (R2F2). CA++ performs better under fast mobility, because the Doppler shifts for diverse paths become more significant. All three prior schemes cannot provide accurate estimation, as they fail to capture or precisely model time-varying Doppler.

We note that, high accuracy gain tends to diminish over mmWave (ultra-high frequencies). It confirms the fact that radio signal over ultra-high frequencies fades much faster than over low frequencies. Nevertheless, CA++ achieves high inference accuracy (all below 1 dB). Since OptML and R2F2 perform much worse than CA++ and REM, we only show




Figure 8: SNR/RSRP/RSRQ errors under low mobility.




Figure 9: Impacts of measured frequencies.

the empirical comparisons with REM next.

SNR versus RSRP/RSRQ. Figure 8 plots the inference errors under low mobility (outdoor). Results under high mobility are similar and omitted. In Figure 8a, we combine all samples when inferring each of four sub-6GHz channels using the measurements over 2.45 GHz. RSRP measurements are within $[-140\text{dBm}, -45\text{dBm}]$, and RSRQs are mostly in $[-20\text{dB}, -4.5\text{dB}]$. The reported RSRP (RSRQ) values are quantized with a resolution of 1 dBm (0.5dB) [6, 7]. CA++ consistently outperforms REM, regardless of the radio quality metric. Compared with REM, CA++ reduces the median error from 3 dB/dBm to below 1 dB/dBm over sub-6G. Its RSRP/RSRQ errors even go to zero (partly due to quantization) over mmWave.

Effectiveness over various frequencies. CA++ remains effective regardless of measured frequencies. Figure 9 shows the results when measuring 5.55 GHz (sub-6G) and 62 GHz (mmWave) under low mobility (outdoor). In fact, we run tests with all five sub-6GHz channels, and omit the results using three other sub-6GHz channels because they are consistent. CA++ outperforms REM by about 2 dB, reducing the SNR errors from 2.16 dB/2.89 dB to 0.5 dB/0.86 dB (min/max). This also applies to mmWave (from 62 GHz to 58 GHz).

Robustness to various settings. CA++ is robust to various numerology settings. We test with various subcarrier spacing Δ_f : 15kHz - 30kHz, 15kHz - 60kHz, 15kHz - 120kHz, and 60kHz - 120kHz under high mobility (the plot is omitted). CA++ contains the median error within 1dB for cross-numerology channel inference.

We also gauge the inference accuracy under weaker signals and stronger noises, where SNR degrades from 30dB to 14dB, as we move the client away from the cell tower (from 3 m to about 6 m). In case of measuring 2.45 GHz to infer other sub-6GHz channels, CA++ increases its inference error by < 0.5 dB, whereas REM increases by > 1 dB (Figures are omitted). This is because CA++ is more robust in inferring path attenuation under noises. We also test CA++ in indoor experiments under low-mobility (walking) or static settings

	Min	Max	Mean
R2F2	2,464	59,294	3,548
OptML	211	285	238
REM	2.0	131.1	51.7
CA++	0.6	44.1	8.2

Table 4: Inference time (ms).

(where the cells are placed at the blue icon and the clients are placed at the rest locations marked in Figure 5c). CA++'s inference accuracy is comparable in these indoor experiments (the plots are omitted).

Efficiency. We compare the efficiency of all four schemes by measuring the time needed for channel inference. CA++ is much more efficient than all the other approaches, at least 3x faster in all the experiments. Table 4 shows the time needed in one setting (sub-6G inference, low mobility, outdoor). CA++ takes 8.2 ms on average, compared with 51.7 ms by REM, 238 ms by OptML, and 3548 ms by R2F2. CA++ and REM outperform those OFDM-based algorithms by at least an order of magnitude, as they do not rely on optimization with many iterations. Compared with REM, CA++ reduces its execution time by decoupling sparse propagation paths and performing concurrent inference. It is even faster than measuring one 5G frequency (>40/80 ms).

8.3 Group-based CA Management



We use two micro-benchmarks to assess its effectiveness.

Measurement acceleration. We first gauge how much the group-based operation speeds up measurements. We compare CA++ with the legacy mechanism, using the average time to detect a cell *eligible* as PCell/SCell. We define cells' eligibility based on whether the cell is ever used at the location. For the legacy scheme, we use results from real traces. For CA++, we estimate the time to measure *all* eligible cells in the same location using Algorithm 2.

Figure 10 shows the measurement accelerations by CA++. Under low/moderate mobility, CA++ decreases the median measurement time from 307 ms to 18 ms for AT&T in C1. CA++ achieves 399 ms → 35 ms for AT&T in C2, 417 ms → 114 ms for Verizon, and 132 ms → 42 ms for T-Mobile. The acceleration is thus by 15.4× (4.7×), 7.0× (1.9×), 5.0× (2.0×), and 3.0× (1.5×) at the median (p90) rate. On the high-speed rail under 4G, legacy CA and CA++ spend 405 ms and 89 ms, with a 4.4× speed-up in 50% cases.

Comparing different datasets, we note that CA++ further accelerates measurements as more frequency channels are deployed. It yields faster measurements for AT&T 5G than Verizon 5G, T-Mobile 5G, and 4G on HST in China. This is because AT&T 5G has deployed more frequencies than others (Table 3). A single measurement covers more colocated cells on different channels. Our design thus yields more benefits with expanding spectrum resources in 5G and beyond.

Measurement acceleration results in faster selection of

**Figure 10: Measurement acceleration by CA++.****Figure 11: Benefits of group-based operations in CA++.**

the serving cell set, thus increasing the effective connection time. Figure 11a shows that, the median connection time improvement (CA++ over legacy) is 190.5 ms – 765.0 ms for 5G operators, and 369.3 ms on 4G high-speed trains. It thus benefits all usage scenarios. In static/low-mobility settings, mobile users spend more time on stable connections. Under extreme mobility (say, high-speed trains), the network offers seamless service by lowering the risk of lost connectivity.

Signaling efficiency. We examine the efficiency of group-based feedback in terms of signaling overhead. We compare the number of reports needed to include all eligible cells under legacy CA and CA++. Figure 11b confirms that CA++ reduces the signaling overhead. Under low/moderate mobility and with AT&T 5G in C1(or C2), the legacy CA incurs the median value of 7 (or 2) reports, while CA++ reduces the median to 3 (or 1). For Verizon 5G, T-Mobile 5G and China Unicom 4G (on HST), CA++ reduces the number of reports in 13.2%, 66.2%, and 31.0% of cases, respectively. This is because fewer frequency channels are deployed on each tower by these operators. The median reduction factor is no smaller than two for all operators. CA++ becomes more efficient with denser cell deployment over more frequency channels.

8.4 Overall Improvement by CA++

We use trace-driven emulations to assess the overall improvement by CA++. Since the device never knows runtime performance of those un-selected cell sets, we use historical data to profile performance of the cell sets available. We perform a “what-if” study to compare the CA options enabled by CA++ and selected by the legacy (current) practice. Note that this assessment is only feasible for our 5G datasets; The public 4G-HST dataset lacks sufficient speedtests. We first assess CA++ with the existing policies and then evaluate it with policy changes (Table 5). CA++ turns out to boost both the aggregated channel-width (Figure 13) and throughput (Figure 14) in all the cases, following the original policies in place. We use A-C1 to showcase the improvement per location in Figure 12 (all test locations along main roads). We see that the gains of CA++ are impacted by policies but largely hold with most policy changes (Figure 15).




Figure 12: Overall throughput improvement (Mbps) observed in A-C1 in the tested region.




Figure 13: CDFs of the aggregated channel-width (CW).




Figure 14: CDFs of median throughput in each region.

Aggregated channel-width. CA++ is assumed to only aggregate cells that are observed to work collectively in real traces. This premise complies with device capacity and network policies. CA++ greatly enlarges the aggregated channel-width. For AT&T, 74.7% (46.5%) of cases in C1 (C2) have the aggregated channel width over 400MHz, compared with 27.3% (6.0%) before. City C1 sees more enhancements than C2, because more 5G frequency channels and cells, particularly mmWave ones, are deployed in C1 (Table 3). We observe that Verizon and T-Mobile add no more than one 5G cell in most cases, and their aggregated channel widths are up to 140 MHz and 135 MHz, respectively. CA++ is effective in making more spectrum resources in use. With Verizon, CA++ can reach the largest channel-width (140 MHz) in 66.7% of cases, while the legacy one reaches 6.2% only. With T-Mobile, CA++ increases the percentage with >100 MHz from 46.9% to 89.4%. The gain is smaller because the legacy practice in T-Mobile does better by using wider spectrum at more locations.

Throughput boost. In our what-if study, CA++ makes new decision on the serving cell set by prioritizing mmWave cells (sub-6G cells for T-Mobile) for wider spectrum whenever they are available. Figure 14 compares the results for all three operators. With AT&T in C1(C2), CA++ increases the median throughput from 35.4 Mbps (29.1 Mbps) to 83.7 Mbps (54.0 Mbps) out of all tested locations. The median throughput rises from 43.3 Mbps to 84.3 Mbps (Verizon), and from 49.0 Mbps to 72.7 Mbps (T-Mobile). Throughput gains are higher in AT&T and Verizon because their legacy CA performs worse over 5G (particularly over mmWave). T-Mobile 5G sidesteps mmWave, and supports sub-6Ghz bands [2].

Impact by policies. We further assess the impacts of

Policy	Cell type	Policy	Change		
			A	V	T
A2	PCell	Monitor inter-freq. when $RSS_s < \Theta_{A2}$	-2(Q)	-10(P)	-10(P)
A3		Report cell if $RSS_c > RSS_s + \Theta_{A3}$	-2(Q)	-2(Q)	-5(P)
A5		Report cell is $RSS_s < \Theta_{A5}^1, RSS_c > \Theta_{A5}^2$	+2(Q)	+2(Q)	+5(P)
			-	-10(P)	-10(P)
A1	4G SCell	Report cell if $RSS_c > \Theta_{A1}$	-	+10(P)	+30(P)
B1	5G SCell	Report cell if $RSS_c > \Theta_{B1}$	+2(Q)	+10(P)	+30(P)
Agg-5G	5G SCell	Allow aggregation with 5G cells	+10(P)	+10(P)	+20(P)
		Allow aggregation with 5G cells	Allow 5G → No 5G		

Table 5: Policy changes (Q: RSRQ, P:RSRP).

policies on CA++, by changing six policies listed in Table 5. All the policies are used in operational mobile networks. The first five policies are regulated in [7] and retrieved from the signaling messages sent to mobile phones. The last one is observed from the long-term historical data. In particular, A2, A3 and A5 are used for PCell selection: A2 specifies when to measure inter-freq channels (otherwise only intra-freq channels). A3 and A5, in different forms, determine when to report a candidate cell. Used for SCell selection, A1 and B1 include the condition for 4G and 5G, respectively. The last policy Agg-5G mandates whether the serving PCell can aggregate with 5G cells or not. To evaluate the policy impacts on CA++, we apply one policy change at a time, while keeping others unchanged. Table 5 lists detailed changes for all three US operators. For threshold-based policies, we apply a relative difference to all cells, since their original values may vary. For the policy about aggregation with 5G cells, we disable such aggregation in any case.

Figure 15 plots the throughput gains by CA++ under different policy changes. We make three observations: First, CA++ outperforms the legacy CA with most tested policy changes. However, there is one exception: CA++ has lower/comparable throughput (no more than 5.2 Mbps) when aggregation with 5G cell is disabled. This is because this policy greatly limits the power of CA. CA++ even performs slightly worse because it simply adopts a coarse-grained approach that chooses the group with the largest aggregate channel-width. It works well with 5G cells (in particular over mmWave bands) but does not yield any gain if the aggregation of 5G cells is disabled. We believe that CA++ can be further enhanced by explicitly considering the metric of achieved throughput. Second, CA++ achieves more improvements when changed policies allow for a larger candidate pool. Reduced thresholds (e.g., increasing thresholds in A3/A5/A1/B1) would bring more candidate cells, thus making more headroom for improvement. Given the increasing number of cells being used by CA in upcoming 5G/6G releases, the potentials of CA++ will grow. Third, policy impacts on the legacy CA are relatively smaller (absolute up-and-down within 19.4 Mbps). This illustrates the important of addressing sequential operations in the legacy CA, which plays a more decisive role than policies in determining what cells are considered for aggregation. The gains of CA++ are




Figure 15: Legacy v.s. CA++ under policy change (+/- means to increase/reduce the policy parameter.)

impacted by policies which changes all cells available for aggregation. It is important to design better policies to work with CA or CA++ but it is an orthogonal problem beyond the scope of this paper.

9 Related Work

Several recent studies seek to improve CA using the standardized mechanisms in 4G/5G [14, 18, 23, 25, 28]. [14, 25] focus on better cell selections for CA with either network or device-centric solutions. Other prior efforts [18, 23, 28] use resource scheduling or joint carrier selection to maximize utilization. CA++ takes a new perspective. We unveil the fundamental limitations in the sequential, cell-by-cell procedure of the current CA. We explore a group-based paradigm for CA to accommodate the expanding spectrum space (say, mmWave bands) and high mobility scenarios.

Our solution approach also differs. The main innovation is two new algorithms to both enable concurrent channel inference and minimize the cell-level measurement cost, thus measuring few yet inferring all. To this end, we draw insights from recent studies on OTFS based modulation in the delay-Doppler domain [21, 22]. However, we go beyond its usage on wireless communications, and exploit the compact, frequency-independent multi-path representations to make accurate channel inference. Furthermore, our algorithm explicitly characterizes fractional Doppler shift, to embrace the expanded spectrum space (into mmWave bands) in 5G and beyond. In contrast, prior algorithms [13, 26, 35] cannot handle the wide frequency spectrum well, since none considered the fractional Doppler effect. Finally, we focus on CA over multiple cells and renovate the group-based management for CA, while [26, 38] studied the single-cell setting and resolved the policy conflicts by modifying the related configurations.

10 Discussion

We discuss several issues related to CA++.

Dual connectivity and multi-connectivity (DC/MC). 5G NR is rolling out DC [31] and MC [33], which aggregate cells from two or more towers. CA++ is conceptually applicable to DC/MC as well. The component on minimal cell-level measurement can stay as it is. We can further adapt CA++ to select cells from multiple cell towers.

Decision making for CA. The decision process to determine which CA group to select remains a separate topic;

it can serve as a plug and play component in CA++. In the future, we may explore other decision criteria with performance-related factors (e.g., latency/reliability/availability etc.).

Phased array. In this paper, we evaluated CA++ with horn antennas, which are different from the phased arrays used in real deployments. Due to the physical constraints of our testbed, we are unable to assess CA++ with the phased arrays which have imperfect beams, sidelobes, and more multipath effects. In principle, co-located cells adopt different antenna phases and angles, which might impact concurrent channel inference. We leave the extension as our future work.

Other RAN technologies. Besides CA, there are many important RAN innovations such as CoMP and vRAN. The cloudification of RAN facilitates the deployment of CA++: different cell towers can exchange deployment information and measurements via the high-speed backhaul network, and thus Algorithm 2 can quickly respond to real-time updates.

Easy deployment. CA++ needs upgrades at both client and network side in operational logic and algorithms only. We do not change protocols or network infrastructure. CA++ is thus a practical solution for the deployed systems.

11 Conclusion

CA is a promising technology in 5G and beyond. It combines fragmented radio spectrum to boost performance. However, current CA design suffers with expanding spectrum space (say, new sub-6GHz and mmWave bands) and high mobility. The fundamental problem is that, the legacy CA uses a sequential, cell-by-cell procedure. It cannot evaluate the radio quality of all cells concurrently and make prompt selections.

We thus devise CA++ to embrace both increased spectrum and high mobility. The principle is to enable group-based feedback and selection. The enabler is two novel algorithms. One characterizes channels in the delay-Doppler domain and admits and calibrates fractional Doppler shifts, thus enabling accurate and concurrent channel inference for all nearby cells. The second algorithm minimizes the cell measurement cost via a new set cover abstraction and its greedy approximation. Our overall solution also fits in the legacy OFDM based 5G PHY and the general 3GPP framework.

Acknowledgements. We greatly appreciate our shepherd, and all anonymous reviewers for their intellectual insights and constructive feedback. The work has been partially supported by NSF grants CNS- 1750953, CNS-1910150, CNS-2008026 and CNS-2112471.

References

- [1] 2023. CA++ Release. <https://github.com/mssn/ca-plus>.
- [2] Nov 23, 2021. How and why T-Mobile sidestepped mmWave 5G. <https://www.lightreading.com/5g/how-and-why-t-mobile-sidestepped-mmwave-5g/d/d-id/773678>.
- [3] 3GPP. 2015. TS36.331: Radio Resource Control (RRC). <http://www.3gpp.org/ftp/Specs/html-info/36331.htm>
- [4] 3GPP. 2018. NG-RAN; Architecture description.
- [5] 3GPP. 2018. TR 38.900: Study on channel model for frequency spectrum above 6 GHz.
- [6] 3GPP. 2021. TS38.133: NR: Requirements for support of radio resource management. V16.7.0.
- [7] 3GPP. 2021. TS38.331: NR: Radio Resource Control (RRC). V16.7.0.
- [8] 3GPP. 2022. TS36.101: E-UTRA; User Equipment (UE) Radio Transmission and Reception. V16.12.0.
- [9] 3GPP. 2022. TS36.300: E-UTRA and E-UTRAN; Overall description; Stage 2 (Release 17). https://www.etsi.org/deliver/etsi_ts/136300_136399/136300/17.01.00_60/ts_136300v170100p.pdf
- [10] 3GPP. 2022. TS38.104: NR; Base Station (BS) Radio Transmission and Reception. V16.10.0.
- [11] 3GPP. 2022. TS38.211: NR; Physical channels and modulation. V16.8.0.
- [12] 5G America. 2019. Global 5G: Implications of a Transformational Technology. <https://www.5gamerica.org/wp-content/uploads/2019/09/2019-5G-Americas-Rysavy-Implications-of-a-Transformational-Technology-White-Paper.pdf>.
- [13] Bakshi, Arjun and Mao, Yifan and Srinivasan, Kannan and Parthasarathy, Srinivasan. 2019. Fast and Efficient Cross Band Channel Prediction Using Machine Learning. In *ACM MobiCom*. ACM, 37.
- [14] Haotian Deng, Qianru Li, Jingqi Huang, and Chunyi Peng. 2020. Icell-speed: Increasing cellular data speed with device-assisted cell selection. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–13.
- [15] Analog Devices. 2022. Getting Started with the EK1HMC6350 Evaluation Kit and Software. <https://www.analog.com/media/en/technical-documentation/user-guides/EK1HMC6350-UG-1031.pdf>.
- [16] Ericsson. 2021. Carrier aggregation in 5G. <https://www.ericsson.com/en/ran/carrier-aggregation>.
- [17] Ericsson. 2021. What, Why and How: the Power of 5G Carrier Aggregation. <https://www.ericsson.com/en/blog/2021/6/what-why-how-5g-carrier-aggregation>.
- [18] Weidong Gao, Lin Ma, and Gang Chuai. 2017. Joint optimization of component carrier selection and resource allocation in 5G carrier aggregation system. *Physical Communication* 25 (2017), 293–297.
- [19] Xiaohu Ge, Song Tu, Guoqiang Mao, Cheng-Xiang Wang, and Tao Han. 2016. 5G ultra-dense cellular networks. *IEEE Wireless Communications* 23, 1 (2016), 72–79.
- [20] Ronny Hadani and Anton Monk. 2018. OTFS: A new generation of modulation addressing the challenges of 5G. *arXiv preprint arXiv:1802.02623* (2018).
- [21] Ronny Hadani, Shlomo Rakib, Shachar Kons, Michael Tsatsanis, Anton Monk, Christian Ibars, Jim Delfeld, Yoav Hebron, Andrea J Goldsmith, Andreas F Molisch, et al. 2018. Orthogonal Time Frequency Space Modulation. *arXiv preprint arXiv:1808.00519* (2018).
- [22] Ronny Hadani, Shlomo Rakib, Michail Tsatsanis, Anton Monk, Andrea J Goldsmith, Andreas F Molisch, and R Calderbank. 2017. Orthogonal time frequency space modulation. In *2017 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.
- [23] Roghayeh Joda, Medhat Elsayed, Hatem Abou-Zeid, Ramy Atawia, Akram Bin Sediq, Gary Boudreau, and Melike Erol-Kantarci. 2021. QoS-aware joint component carrier selection and resource allocation for carrier aggregation in 5G. In *ICC 2021-IEEE International Conference on Communications*. IEEE, 1–6.
- [24] UCLA WiNG Lab. [n. d.]. Flora: Flexible Mobile Network Platform. <http://metro.cs.ucla.edu/index.html>
- [25] Qianru Li and Chunyi Peng. 2021. Reconfiguring Cell Selection in 4G/5G Networks. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.
- [26] Yuanjie Li, Qianru Li, Zhehui Zhang, Ghufan Baig, Lili Qiu, and Songwu Lu. 2020. 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*. 344–358.
- [27] Yuanjie Li, Chunyi Peng, Zengwen Yuan, Jiayao Li, Haotian Deng, and Tao Wang. 2016. Mobileinsight: Extracting and Analyzing Cellular Network Information on Smartphones. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (New York City, New York).
- [28] Hong-Sheng Liao, Po-Yu Chen, and Wen-Tsuen Chen. 2014. An efficient downlink radio resource allocation with carrier aggregation in LTE-advanced networks. *IEEE Transactions on Mobile Computing* 13, 10 (2014), 2229–2239.
- [29] Authors of MobiCom'23 submission #49. [n. d.]. Supplementary Materials to MobiCom'23 Submission #49. <https://drive.google.com/file/d/1YzbhaQXNKv2HHd4ZKu1icHssPRLz-OHi/view?usp=sharing>
- [30] Jae-Joon Park, Juyul Lee, Kyung-Won Kim, and Myung-Don Kim. 2020. 28-GHz high-speed train measurements and propagation characteristics analysis. In *2020 14th European Conference on antennas and propagation (EuCAP)*. IEEE, 1–5.
- [31] Jens Bergqvist Patrik Rugeland. [n. d.]. Key insights: Early measurements for improved carrier aggregation and dual connectivity setup. <https://www.ericsson.com/en/blog/2020/10/carrier-aggregation-dual-connectivity>
- [32] Patchava Raviteja, Khoa T Phan, Yi Hong, and Emanuele Viterbo. 2018. Interference cancellation and iterative detection for orthogonal time frequency space modulation. *IEEE Transactions on Wireless Communications* 17, 10 (2018), 6501–6515.
- [33] R. Mumtaz S. A. Busari and J. Gonzalez. [n. d.]. Multi-connectivity in 5G new radio standards. <https://www.standardsuniversity.org/e-magazine/december-2020/multi-connectivity-in-5g-new-radio-standards/>
- [34] David Tse and Pramod Viswanath. 2005. *Fundamentals of wireless communication*. Cambridge university press.
- [35] Vasisht, Deepak and Kumar, Swarun and Rahul, Hariharan and Katabi, Dina. 2016. Eliminating Channel Feedback in Next-Generation Cellular Networks. In *Proceedings of the 2016 ACM SIGCOMM Conference*. ACM, 398–411.
- [36] Cheng-Xiang Wang, Ammar Ghazal, Bo Ai, Yu Liu, and Pingzhi Fan. 2015. Channel measurements and models for high-speed train communication systems: A survey. *IEEE communications surveys & tutorials* 18, 2 (2015), 974–987.
- [37] Jing Wang, Yufan Zheng, Yunzhe Ni, Chenren Xu, Feng Qian, Wangyang Li, Wantong Jiang, Yihua Cheng, Zhuo Cheng, Yuanjie Li, et al. 2019. An active-passive measurement study of tcp performance over lte on high-speed rails. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–16.
- [38] Zengwen Yuan, Qianru Li, Yuanjie Li, Songwu Lu, Chunyi Peng, and George Varghese. 2018. Resolving policy conflicts in multi-carrier cellular access. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 147–162.