A New Paradigm of Communication-Aware Collaborative Positioning for FutureG Wireless Systems

Yu-Tai Lin and Karthikeyan Sundaresan

Georgia Institute of Technology

USA

ABSTRACT

Future generation wireless systems are rapidly evolving to support positioning as a native feature using their communication infrastructure. Notwithstanding the dependence on a densely deployed infrastructure (for accurate positioning). We highlight the significant degradation due to the overhead of such an infrastructure-based positioning (IP) approach can bring communication performance warranting timely attention. We propose a fundamentally different paradigm of communication-aware collaborative positioning *CO2P*, whereby the burden of positioning is offloaded to client devices in an intelligent, communication-aware and collaborative (peer-peer) manner that reduces overhead and improves spatial reuse to preserve communication performance, without compromising on positioning accuracy. Through technically-sound algorithms *CO2P* addresses the underlying tradeoff between communication performance and positioning accuracy to deliver an efficient coexistence, providing a two-fold increase in both throughput and accuracy over conventional IP approaches. *CO2P* is also orchestrated as a practical, distributed, adoption-friendly solution that is realized using WiFi's positioning protocol.

CCS CONCEPTS

• **Networks → Mobile networks**; **Mobile ad hoc networks**; **Network design and planning algorithms**.

KEYWORDS

5G, next-generation wireless system, WiFi, wireless positioning

1 INTRODUCTION

Indoor positioning (a.k.a localization) for interactive, immersive, and robot-human applications, forms a key service targeted by future generation wireless systems (both WiFi and cellular). These systems are rapidly evolving to support positioning as a native feature within their communication paradigm (802.11mc ranging protocol [2] in WiFi, and upgraded positioning protocol [6] in 5G new radio - NR), potentially becoming one of the first, key applications offered in the dis-aggregated RAN framework (Open RAN [3]). Leveraging the existing communication infrastructure of access points (APs) and base stations (BSs) to localize and position client devices (a.k.a. infra-based positioning), has the potential to bring high-performance positioning and tracking to indoor environments, and is expected to be the standard approach. However, not much thought has been given to its impact on the communication performance itself and hence their coexistence.

Limitations of Infrastructure Positioning: Existing positioning standards (both WiFi and cellular), require a client to take measurements from multiple BSs (to determine its location) making it challenging to provide wide coverage. Notwithstanding the dependence on a densely deployed infrastructure (for accurate positioning) and the associated cost, the measurement overhead of an infra-based positioning approach (Fig. 1a) can bring significant degradation to communication performance. While allocating a dedicated spectrum for positioning can address this degradation, it would come at the expense of the limited spectrum allocated for communication, and hence not viable. From the realistic network simulation results shown in Fig. 3a, it is clear that while measuring multiple ranges to several BSs helps increase the positioning accuracy, it comes at the expense of the network's throughput capacity that further degrades with the number of clients, resulting in a drop of over 40% (for 3m accuracy). When the communication interface is further leveraged to track clients, the periodic measurements (higher positioning refresh rate) needed to keep track of client mobility will easily overwhelm the network's communication performance (Fig. 3b). This brings us to the fundamental tradeoff between communication (throughput capacity) and positioning (accuracy) performance that needs to be addressed for a successful roll-out of these positioning solutions. Given the growing importance of 5G positioning, this work aims to take an important, timely step in enabling the seamless coexistence of positioning and communication.

CO2P– **A New Paradigm:** Our proposed approach explores a fundamentally different paradigm of *COmmunication-aware COllaborative Positioning*, *CO2P* (as shown in Fig. 1b), whereby the burden of positioning is offloaded to client devices in an intelligent, communication-aware and collaborative (peer-peer) manner that reduces overhead by an order of magnitude and fosters improved spatial reuse to preserve communication performance without compromising on positioning accuracy. While collaborative positioning has been explored in isolation for stationary sensor
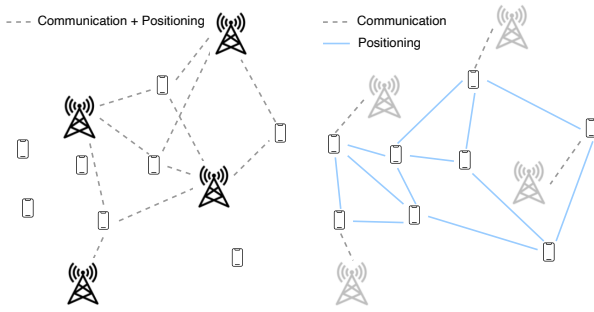
Figure 1: (a) Infrastructure Positioning (b) Collaborative Positioning

networks [23], bringing it to communication networks changes the nature of the problem and faces two key challenges:
- How to characterize the trade-off between communication and positioning performance, and adapt the paradigm of collaborative peer-peer positioning to design a seamless and efficient coexistence framework.
- How to realize coexistence in practical wireless systems through the design of scalable, distributed, and efficient medium access mechanisms that enable a seamless operation of collaborative positioning along with communication data traffic.

CO2P **Design:** Towards addressing these challenges, *CO2P*'s design incorporates two key components:

(i) *A centralized framework for characterizing the communication-positioning tradeoff with collaborative positioning (CP):* CP involves the estimation of distances (a.k.a. ranging) between different pairs of clients, captured through Euclidian distance matrices (EDMs) and jointly localizes the *topology* of clients as a whole by applying multi-dimensional scaling (MDS) on complete EDMs [11]. In contrast to stationary sensor networks, 5G networks cater to communication traffic to/from mobile clients. This poses two critical challenges while more measurements between peer clients increase the completeness of EDM and hence the localization accuracy of all clients (say N) collectively: (i) the associated overhead ($O(N^2)$) significantly reduces the communication capacity, and (ii) client mobility renders measurements collected beyond a certain duration (coherence time related to mobility) stale and detrimental to localization accuracy. The coexistence with communication renders existing solutions ineffective, leading to a novel, unique variant of the CP problem that we refer to as latency/overhead-bounded CP.

Given the bounded overhead, *CO2P*'s centralized approach introduces three key overhead-reducing features: *progressive spectrum ranging*, *rigidity-aware ranging*, and *passive ranging*. In progressive spectrum ranging, *CO2P* leverages the observation that ranging on a larger channel bandwidth (e.g. 80 vs 20 MHz) yields a higher accuracy of range estimation with a smaller interference footprint (power spread across a wider bandwidth), thereby also leading to higher spatial reuse. It maximizes the number of non-interfering peer-peer range measurements that can be scheduled on an 80 MHz channel first. Then it estimates the location of the clients and determines if longer links need to be established in the topology through additional ranges on a 40 MHz channel (or 20 MHz if needed) if increased accuracy is desired. Through rigidity-aware ranging, *CO2P* intelligently selects the set of links for ranging that maximize spatial reuse while enabling a tri-lateration ordering of the clients – the latter brings rigidity to the topology,

allowing *CO2P* to iteratively improve its estimate of the missing ranges through a process of sequential multi-lateration and location perturbation. Finally, through passive ranging, *CO2P* devises a novel, collaborative time-of-flight estimation approach that enables passive clients with over-heard transmissions to compute their ranges to the transmitting clients without expending any overhead. *CO2P*'s algorithms allow us to characterize the potential of CP in effectively addressing the communication-positioning tradeoff, obtaining a localization accuracy that is better than IP, at a fraction of the overhead ($O(1)$), and impact on communication performance.

(ii) *A distributed MAC solution for efficient coexistence of collaborative positioning and communication traffic:* Translating *CO2P*'s centralized framework to practical wireless networks such as WiFi and cellular faces numerous challenges. Given the lack of scalability with a BS/AP scheduling the peer-peer measurements, distributed client-based algorithms are needed to enable neighbor discovery and channel access for peer-peer measurements, in a manner that realizes efficient coexistence with communication traffic. Further, the extremely distributed nature of peer-peer transmissions, increases the incidence of hidden terminals, impacting the utility of the collected measurements.

CO2P proposes a novel *positioning access control* that brings together the two types of traffic (communication and positioning measurements) under a common umbrella of utility-based distributed medium access control, whereby optimal policies to control throughput (communication), accuracy (positioning), and fairness between clients and across traffic types are established. Towards practical deployment, it further instantiates the distributed solution within WiFi's MAC protocol using the latter's differentiated access contention mechanism, while adapting WiFi's ranging protocol along with intelligent channel allocation (for ranging) to enable efficient neighbor discovery both within and across cells, as well as accounting for the relative priority of different measurements in facilitating passive ranging.

We implement both the centralized and distributed solutions in *CO2P* in NS3 [4], while leveraging WiFi's 802.11mc FTM measurements as the basic ranging protocol with the FTM extension module [29]. Our extensive evaluations highlight that *CO2P* delivers an effective joint communication-positioning system that is capable of accurately (under 2m) tracking a large network of sixty mobile (1 m/s) clients, while significantly limiting the impact on communication performance (0.2x) on the same spectrum. This is in contrast to IP approaches that suffer over a 10x drop in communication performance to deliver a 3m accuracy.

Our contributions in this work are multi-fold:
- We propose and showcase an alternate paradigm of *communication-aware collaborative positioning* (*CO2P*) for efficient coexistence of communication and positioning traffic in FutureG systems. To the best of our knowledge, this is a first-of-its kind proposal towards an important, timely problem.
- We design a centralized framework with the associated algorithms needed to address the unique overhead-bounded CP problem.
- We propose a distributed MAC called *positioning access control* and apply it within WiFi's existing MAC and positioning framework for practical deployment and adoption.

We will open-source our algorithms in NS3 to facilitate future research into the novel paradigm of *CO2P*. We believe *CO2P*'s
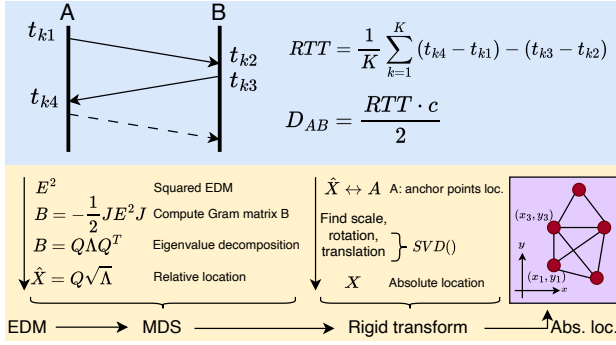
Figure 2: (a) Two-way ranging (b) EDM-based CP

Figure 3: IP: Trade-off (a) Overhead (b) Mobility

core design, albeit instantiated for WiFi, equally applies to cellular networks as well, with the growing popularity of peer-peer communication through ProSe [5] (Proximity Service) in LTE/NR.

2 BACKGROUND AND RELATED WORK

Infra-based positioning (IP): A client takes positioning measurements to multiple BSs (3 or more for 2D positioning), which are then aggregated along with knowledge of the BSs' own locations, to estimate the location of the client, as shown in Fig. 1a. The measurements typically include ranges, obtained indirectly through time-of-flight (round-trip-time RTT in cellular [6], fine time measurements FTM in WiFi [2], two-way ranging TWR in UWB [19]). In particular, ranges are estimated through a two-way packet exchange between the entities, whereby recording the local timestamps of transmission-reception of the packets, the relative time offset can be eliminated in the ToF estimation (as shown in Fig. 2a). For MIMO-based BSs and clients, measurements can also include angular (angle-of-arrival AoA, angle-of-departure AoD) information. With access to higher bandwidths (through carrier or channel aggregation) in 5G systems, ranging and hence positioning resolution is moving closer to sub-meter and is becoming attractive for several applications. While the body of work in IP is quite large, tackling both WiFi [18, 24, 26, 28] and cellular [8, 12, 15, 20] systems, we have yet to understand the impact of such positioning on communication performance as the current and future wireless standards move towards their integration under a common interface.

Collaborative positioning (CP): The notion of collaborative positioning has been popular in sensor networks [17, 22, 27]. As shown in Fig. 2a, clients measure ranges to each other, and the collection of such ranges is captured through an Euclidian distance matrix (EDM), which can be decomposed (using multi-dimensional scaling, MDS, Fig. 2b) to estimate the locations of the clients that best capture the measured ranges [9, 11, 14]. It relies on the construct of geometric rigid bodies [16]. More ranges (edges) from a client lead to better estimations. Theoretically, the minimum number of measurements required to resolve its position in r dimensions ($r = 2$ for 2D) is $r + 1$. With the practical challenge of obtaining a large number of ranges, EDM completion (filling in missing entries) approaches have become important [11]. Such approaches lead to relative localization of the clients jointly as a topology, i.e. invariant upto a translation and rotation; they are converted to absolute positions through a rigid transform [10]
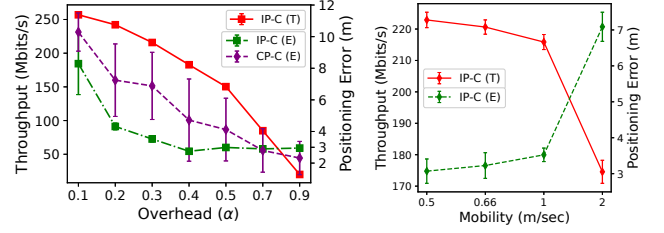
using 3 or more clients (in the entire topology), whose positions are either known a priori or obtained through IP. While the notion of CP has existed before, it has understandably focused on relative localization algorithms, more so for stationary clients with sufficient measurements, without considering the impact of their associated overhead on communication performance in operational networks. In fact, when one limits the measurement overhead to enable coexistence with communication, maintaining the topology's rigidity becomes difficult, leading to poor accuracy in tracking mobile clients or even a large number of stationary clients. This in turn leads to a novel, overhead-bounded variant of the CP problem, that is much harder to address.

We intend to address the underlying tradeoff between positioning and communication through our novel *CO2P*, *communication-aware* collaborative positioning framework. While we consider the prevalent case of sub-6 GHz networks with single antenna devices, several aspects of our work will apply to mmWave networks as well (albeit, with adaptations to incorporate directionality information).

3 POTENTIAL FOR COLLABORATIVE POSITIONING

3.1 Motivation

We conduct a comprehensive system simulation in NS3 leveraging WiFi's FTM ranging protocol for IP (centralized - no contention) in a network with 3 BSs and 60 clients. We control the relative channel time allocation (α, i.e. measurement overhead) to positioning traffic within each allocation epoch and study the tradeoff between positioning accuracy and communication throughput.

From Fig. 3a, it can be seen that with less measurement in channel access, IP incurs a high error (8-10m), while larger overhead leads to better accuracy (3m error) at the expense of a significant impact on communication performance (40%). As clients' mobility increases (Fig. 3b), the number of required measurements within a given time increases considerably, leading to a large impact on communication.

In IP, each client performs ranging to a limited number of (typically 3) BSs that leads to saturation of accuracy (Fig. 3a). In contrast, a conventional CP approach (centralized, based on Sec. 2) has the potential to (i) further increase accuracy by using numerous neighbors as virtual anchors [13], and (ii) increasing spatial reuse in the network through local peer-peer measurements. However, as Fig. 3a reveals, this potential is realized only at the cost of high overhead because of the challenges we discuss below.
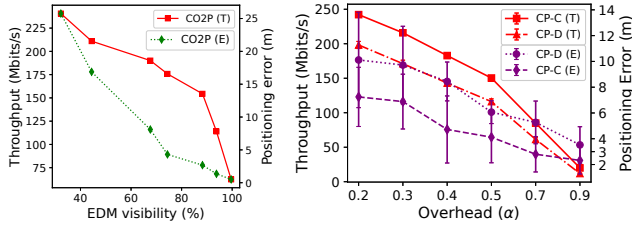
Figure 4: (a) CP: Impact of ranging overhead N^2 for EDM, (b) Limitations of existing MAC for CP (centralized vs. distributed)

3.2 Challenges

Conventional CP as an approach faces challenges at two levels: (i) the larger the number of measurements between different client pairs, the better the accuracy of the localized topology, but the larger the impact on communication (Fig. 4a). Indeed, EDM-based solutions expect a complete EDM, i.e. range estimates of every pair of clients ($O(N^2)$). However, in an *overhead-bounded* setting (key differentiator in a communication-aware CP), intelligent algorithms are required to identify the appropriate set of ranges. The measurement will maximally increase the collective positioning accuracy while incurring the least overhead and communication impact, lest it might not be feasible to position any of the clients. (ii) Further, the nature of measurements has an impact on communication performance as well – the larger the bandwidth used for ranging, the better the accuracy of ranging (seen from our prior experimental work [7]) and the better the spatial reuse through a smaller signal footprint, but the lesser the spectrum available for communication traffic.

CP in a practical, distributed network, peer-peer transmissions increase the number of hidden terminals, making it harder to discover neighbors and their ranges and leads to a significant degradation in both communication throughput and positioning error compared to its centralized counterpart, as seen in Fig. 4b (distributed CP employs 802.11 DCF for contention). Further, clients across cells could be operating on different channels preventing the discovery of ranges important to the topology. The key challenge lies in ensuring fair and efficient channel access between ranging measurements and communication traffic, as well as realizing such a solution that alleviates hidden terminal impact in practice.

We will tackle the challenges of CP through a centralized CO2P framework and associated algorithms that address the tradeoff between communication and positioning effectively. Then, we will address the distributed realization of CO2P within the context of WiFi's FTM ranging protocol to yield a practical solution.

4 *CO2P*: DESIGN

4.1 Overview

CO2P's key contributions lie in devising intelligent algorithms that employ a small set of carefully selected peer-peer range measurements ($O(1)$ overhead) to help fill in the remaining missing ranges. This leads to high positioning accuracy, while also offering a significantly better communication performance. As shown in Fig. 5 and detailed in Section 4.2, *CO2P*'s design incorporates 3 innovative components that are common to both its centralized and distributed manifestations: (i) *Progressive spectral ranging*: First, *CO2P* maximizes the number of non-interfering
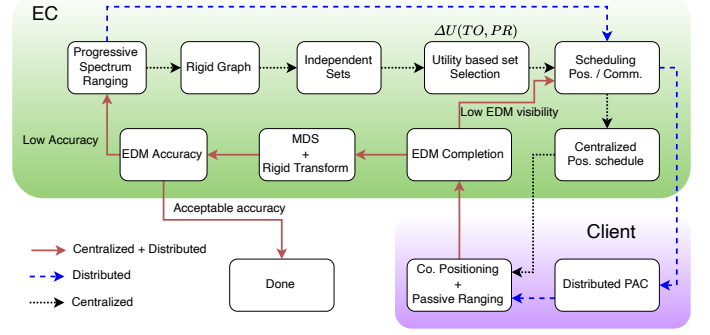

Figure 5: Overview of *CO2P* Operations

peer-peer measurements which can be scheduled simultaneously by employing the largest bandwidth channel (e.g. 80 MHz) for high-accuracy ranging. Then, it estimates the location of the clients and determines if longer links need to be established in the topology through additional ranging on a channel with progressively lower bandwidth (e.g. 40 MHz followed by 20 MHz) in subsequent iterations. This allows for maximum spatial reuse of the measurements, resulting in the least overhead (for communication) with maximum utility (for positioning). (ii) *Rigidity-aware ranging:* Second, in each iteration, it increases the accuracy of the filled-in ranges by prioritizing ranges that establish a tri-lateration ordering of the clients – the latter brings rigidity to the topology, allowing CO2P to iteratively improve its estimate of the missing ranges through a process of sequential multi-lateration and location perturbation. (iii) *Passive ranging:* Third, as the ranging measurements are conducted in each iteration, any client which overhears a measurement between another pair of clients infers its own range to each of the over-heard nodes, by recording and intelligently processing its own ToF of the over-heard measurements. This contributes to a significantly large number of additional ranges without expending any overhead and allows *CO2P* to further decrease its overhead on active measurements.

In addition, *CO2P* designs a distributed protocol in Section 4.3 that enables practical deployment in distributed wireless networks with ranging capabilities (e.g. WiFi FTM [1]). Here, the challenge lies in executing its ranging mechanisms in a completely distributed manner, while also coexisting in a fair and efficient manner with the communication transmissions. *CO2P* proposes a novel positioning access control (PAC) framework that brings together the two types of traffic (communication and positioning measurements) under a common umbrella of utility-based distributed medium access control, and establishes an optimal policy (executable at each client locally) to control throughput (communication), accuracy (positioning) and fairness between clients and across traffic types. This policy is then realized within WiFi's MAC using its prevalent QoS-based contention framework (802.11e), whereby different access categories and their respective contention parameters are adapted to distributively prioritize different ranging measurements based on their contribution to the overall positioning accuracy and spatial reuse. Further, WiFi's FTM ranging protocol is adapted to (i) enable and foster maximum passive ranging during active measurements while also leading to fewer hidden terminals, and (ii) leverage multiple channels, particularly for edge-clients (that overhear multiple BSs) to be able to discover peer clients

---

Algorithm 1: Bounded-overhead CP

input : Client set R; Position error threshold ξ; EDM visibility threshold θ; Period time T
output: Client positions P

1 **for** BW; $BW \in \{80\ MHz,\ 40\ MHz,\ 20\ MHz\}$ **do**
2   Construct Rigid graph G with tri-lateration ordering;
3   Construct independent sets S by estimating the corresponding conflict graph in G;
4   Adopt smallest α;
5   **while** $\alpha \leq 1$ **do**
6   Compute FTM access duration A_p^t;
7   **while** *FTM measurement in* A_p^t **do**
8   Select independent sets based on utility ΔU and start parallel transmissions;
9   **end**
10   Collecting active + passive FTM measurement from all nodes in R and construct EDM;
11   **if** *EDM visibility* $< \theta$ **then**
12   Increase α;
13   continue;
14   **end**
15   Get complete EDM by sequential multi-lateration + location perturbation;
16   Apply MDS on complete EDM to get relative location;
17   Apply rigid transform on relative location to get absolution location of clients P;
18   **if** *Client positions error* $\leq \xi$ **then**
19   **return** *Client positions* P;
20   **end**
21   **end**
22 **end**
23 **return** *Client positions* P;

Figure 6: Bounded-overhead CP Algorithm

across cells operating on different channels, thereby still enabling measurements between them.

4.2 *CO2P*: Centralized Framework

An enterprise controller (EC) is responsible for collecting the outcome of the ranging measurements from its clients through its communication interface, and computing the client positions.

4.2.1 Bounded-overhead Collaborative Positioning. With the scheduling of data and positioning traffic being executed in epochs, the overhead of positioning can be bounded by the EC by fixing the fraction (α) of an epoch devoted to range measurements. *CO2P*'s focus is on maximizing the number of high resolution ranges that can be obtained within the small measurement period α to maximize positioning accuracy. If the estimated accuracy is not satisfactory at the end of the measurement period, its duration is increased to accommodate additional ranges.

CO2P's selection of links for ranging proceeds iteratively. The interference between peer-peer measurements is captured through a conflict graph[1], and the independent sets of which provide the set of active measurements that can be scheduled in parallel. By starting with the largest bandwidth for ranging, *CO2P* not only maximizes the accuracy of collected ranges, but the limited interference footprint of transmissions (power is distributed across a wider bandwidth) leads to a smaller node degree in conflict graph, thereby allowing for increased spatial reuse. In each iteration, *CO2P* composes its schedule as follows: it selects a previously un-scheduled independent set (s^*) that maximizes the aggregate utility for overall positioning accuracy. It accomplishes this by evaluating the incremental utility provided by each link (i, j) in the set, determined based on the contribution of the link (positioning throughput, ΔU contributed by each of its clients' transmissions; discussed in Section. 4.3.1) to a tri-lateration ordering (TO)[2] of

[1]Ranging links are vertices, and an edge indicates interference between two links, precluding them from operating in tandem.
[2]Every node has 3 neighbors in a manner that enables sequential multi-lateration of all the nodes.

clients in the topology, as well as the additional passive ranges (PR) enabled by it, i.e. $s^* = \arg\max_s \sum_{(i,j) \in s} (\Delta U_i + \Delta U_j)$.

Once the schedule is executed, the active and passive ranges are estimated. However, to enable joint positioning of all the clients, a complete and accurate EDM is needed to apply MDS as outlined in Section 2. To fill in the missing ranges as accurately as possible, *CO2P* leverages its tri-lateration ordering as follows. With the collected ranges, it enables a sequential multi-lateration of all the nodes to estimate their potential locations. Using these estimated locations, another estimated EDM is generated where the RMSE with respect to the actual measured ranges is, then, used to perturb the location of the nodes by gradient descent. The process is repeated till the RMSE cannot be further improved. If the RMSE is satisfactory, the estimated complete EDM along with 3 of the clients' absolute locations which obtained through IP, is used to obtain the absolute location of all the clients (using the approach discussed in Section 2). As additional measurements are collected in subsequent iterations, the location estimates of the clients become more accurate. Every few iterations, *CO2P* samples a few of the clients and conducts IP to them directly from the BSs and compares the error in the location estimates between IP and its CP. If the error difference is within a threshold ξ, the measurement period can be pre-empted (giving the remaining reserved duration to communication traffic). However, if additional measurements are needed due to low EDM visibility, then α is increased, and measurements are conducted in the same bandwidth (if un-scheduled links remain) before moving to progressively lower bandwidths to increase the topology's connectivity. Fig. 6 summarizes the overall procedure.

4.2.2 Passive Ranging. The lesser the missing EDM ranges, the higher the positioning accuracy. Hence, when a set of measurements are transmitted between clients, *CO2P* leverages the broadcast nature of the wireless channel, coupled with the information carried in the measurement packets and the recording of the time of reception of these packets by other neighboring clients, to *passively* infer additional ranges between the overhearing and transmitting clients through intelligent processing.

Consider a canonical topology with 4 clients, K, L, M, and N. Let t_i^t and $t_{j,i}^r$ denote the time at which the client i transmitted the measurement request and the time at which client j received the request from i locally. Then, based on TWR process, the range between i and j can be estimated as follows.

$$(t_{i,j}^r - t_i^t) - (t_j^t - t_{j,i}^r) = \frac{2d_{i,j}}{C} \tag{1}$$

While TWR is typically between a pair of clients, *CO2P* adapts it for peer-peer ranging, by allowing multiple clients to respond sequentially to a TWR transmission. If L initiates the first transmission, followed by responses from K, M and N, then L's distance to K, M and N can be estimated as,

$$(t_{L,X}^r - t_L^t) - (t_X^t - t_{X,L}^r) = \frac{2d_{L,X}}{C}, \ \forall X \in \{K, M, N\} \tag{2}$$

In addition to the active ranges enabled by these measurements directly, *CO2P* infers two additional types of ranges between neighboring clients learnt passively from overheard transmissions, namely (i) passive T (transmitting clients): while two transmitting clients (e.g. M and N) are involved in TWR measurements, their transmissions are overhead at each other ($t_{M,N}^r$ at M, $t_{N,M}^r$ at N),

which in turn can be used to infer $d_{M,N}$ from,

$$(t_{M,N}^r - t_M^t) - (t_N^t - t_{N,M}^r) = \frac{2d_{M,N}}{C} \qquad (3)$$

and (ii) passive R (non-transmitting clients): a client O can overhear the transmissions between K, L, M, and N, and infer its ranges to each of these four transmitting clients, while not participating in a transmission itself (Fig. 7a). Recording the time of reception of each of the four transmissions, and taking the time difference of arrival between links (e.g. O-L and O-K, O-L and O-M, and O-L and O-N) with respect to a reference link (e.g. O-L), O can compensate for the lack of its own transmission. This results in the following equations,

$$(t_{O,X}^r - t_{O,L}^r) - (t_X^t - t_{X,L}^r) = \frac{d_{O,X} - d_{O,L} + d_{L,X}}{C}, \ \forall X \in \{K, M, N\} \qquad (4)$$

Estimating Passive Ranges: Using the three TDOA equations above, *CO2P* estimates the ranges for any overhearing client O to the existing clients, K, L, M, and N as follows. Given the range estimates in Equations 2 and 3, *CO2P* computes a solution for the location of the clients K, M, and N (with L as reference). Any solution tuple can be selected while multiple solution tuples are possible for K, M, and N. Given the location of these 4 clients, and employing the three hyperbolic equations governing the ranges of an overhearing node O, the latter's location can be estimated, from which its ranges to the other nodes are then obtained. To establish that the passive ranges inferred by *CO2P* are accurate, we have the following lemma.

LEMMA 4.1. *Given the relative distances of three clients (K, M, N) to a reference node L, three time difference of arrival equations with respect to L are sufficient for an overhearing node O to infer its own ranges to the other nodes (K, L, M, and N).*

PROOF. (Sketch) It is easy to observe that while multiple location tuples are possible for (K, M, N), and each tuple will result in a corresponding location for O, the relative distance (between K, L, M, and N) and TDoA constraints (between links OK-OL, OM-OL and ON-OL) will restrict the relative distance of O from other clients to be the same, irrespective of the solution tuple chosen. Thus, while accurate location estimates cannot be resolved locally, the ranges are still inferred accurately. □

Overhead Scaling Benefits:

THEOREM 4.2. *CO2P incurs a constant measurement overhead ($O(1)$) that does not scale with the number of clients, unlike a conventional CP approach, whose overhead in turn scales as $O(N^2)$ compared to $O(N)$ in an IP approach.*

PROOF. Consider N clients in a single contention region. In IP, each client will need to conduct TWR with at least 3 BSs, resulting in a total of $6N$ transmissions. In conventional CP, every pair of clients needs to perform TWR, resulting in $N(N-1)/2$ transmissions. In contrast, with *CO2P*, only 4 clients are sufficient to perform a single transmission each, allowing all the other clients to infer their relative ranges to each other through its passive ranging feature. In a larger network with multiple contention regions, spatial reuse across the regions will allow this overhead to still scale as $O(1)$. □

Implications: This result highlights the importance of *CO2P*'s passive ranging not just from the perspective of providing constant overhead irrespective of growing client density in the network,
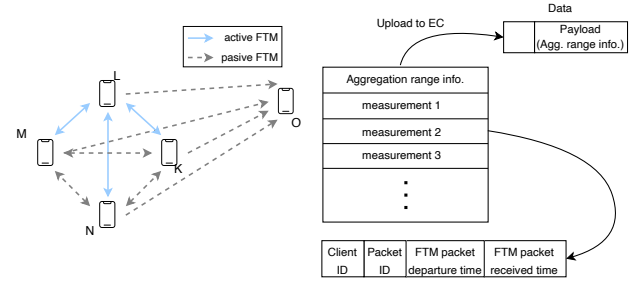

Figure 7: (a)Passive ranging illustration: Canonical Topology (b) Passive ranging information structure

but also in establishing the feasibility of deploying CP as a viable paradigm for positioning in FutureG systems, whose absence would lead to a high quadratic overhead ($O(N^2)$).

4.3 *CO2P*: Distributed Protocol

We now detail how *CO2P*'s ranging mechanisms can be efficiently and systematically incorporated in a practical WiFi network, leveraging the latter's FTM ranging protocol.

4.3.1 Positioning Access Control. CO2P proposes a novel positioning access control framework that brings together the two types of traffic (communication and positioning measurements) under a common umbrella of utility-based distributed medium access control and establishes an optimal policy (executable at each client locally) to control throughput (communication), accuracy (positioning) and fairness between clients and across traffic types. Noting that medium access control for communication and positioning have different objectives. Hence, *CO2P* leverages network utility based optimization theory [21] to first design optimal, stochastic access mechanisms jointly for communication and positioning, whose probabilistic nature allows for subsequent easier, distributed realization in practice. The overall utility based optimization problem has the following objective.

$$\text{Maximize} \sum_{c \in C} (1 - \alpha) U_C(\bar{R}_c) + \sum_{p \in \mathcal{P}} \alpha U_P(\bar{R}_p) \qquad (5)$$

where C and $\mathcal{P}$ denote the set of clients participating in communication (BSs, clients) and positioning (clients) respectively in our collaborative framework. $U_C()$ and $U_P()$ represent the utility functions that capture communication and positioning performance respectively as a function of the client's average throughput, while $\alpha \in [0, 1]$ is a constant, tunable parameter that relatively weights the objectives to prioritize communication over positioning flexibly during medium access. In communication, the average throughput of a client at time t depends on its transmission rate history and can be related to its medium access probability (q_c) and contention loss probability (p_c that is measured locally) through an exponentially-weighted moving average as,

$$\bar{R}_c(t) = (1 - \eta) \cdot \bar{R}_c(t - 1) + \eta \cdot R_c(t) \cdot q_c (1 - p_c) \qquad (6)$$

In positioning, we are interested in positioning accuracy rather than the number of bits transferred. However, given the challenge of directly translating accuracy to throughput, *CO2P* captures an equivalent *positioning throughput* in terms of the number and quality of peer measurements needed by a client to increase the accuracy of its ranges, contribution to passive ranges and rigidity of the measurement topology, all of which directly impact the positioning accuracy. This can be captured as,

$$\bar{R}_p(t) = \left(\bar{R}_p(t-1) + R_p(t) \cdot q_p(1-p_p)\right) \cdot \eta'$$

$$\text{where, } \bar{R}_p(t-1) = (1-\delta_p^b)M_p^t + \sum_{j \in N(p)} \log(1 + (1-\delta_j^b)M_{pj}^r) \quad (7)$$

$$\text{and } R_p(t) = (1-\delta_p^b) + |N(p)| \cdot \log\left(1 + \frac{(1-\delta_p^b)}{1+(1-\delta_p^b)M_p^t}\right)$$

where the average rate $\bar{R}_p(t)$ of client p at any instant is related to the total number of measurements obtained with its neighbors ($j \in N(p)$) thus far which contributes to the overall positioning accuracy. This includes both its own transmissions (M_p^t) as well as receptions from neighbors ($M_{pj}^r, \forall j \in N(p)$), where the former has a higher impact (linear vs. logarithmic for receptions) in being able to contribute to both itself and its neighbors through passive ranging. In contrast, the instantaneous rate focuses on the contribution of the single measurement towards itself and its incremental value ($\log(1 + (1-\delta_p^b)(M_p^t + 1)) - \log(1 + (1-\delta_p^b)M_p^t)$) to neighbors. The rate also incorporates the accuracy of the measurement estimates (given by $(1-\delta_p^b)$), where the measurement uncertainty (noise) δ_p^b depends on the bandwidth (b) used for positioning ($\delta \propto 1/BW$). This captures the coupling between communication and positioning – the larger the bandwidth used for positioning, the better the measurement accuracy, but the lesser the potential number of neighbors and hence topology rigidity, as well as reduced bandwidth for communication. The final term $\eta' = \frac{\eta R_c^{max}}{R_p^{max}}$ is used to normalize and translate the transmissions that capture positioning throughput (# measurements) to the communication throughput (# bits) so as to allow for joint access control between the two.

Noting the constraint that the sum of access probabilities of all clients within a contention region is bounded by one, we can now use a standard Langrangian-based dual optimization to perturb the system about its optimal point. This allows us to derive optimal, distributed access policies for both modalities. We will consider concave utility functions ($U() = \log()$ for proportional fairness employed in today's wireless systems [21, 25]) to capture the notion of diminishing returns with increasing throughput (bps and measurements respectively) to enable fairness in access. We can now derive the following policies for adapting the access probabilities of entities in the respective modalities.

$$q_c(t) \leftarrow q_c(t-1) + (1-\alpha) - p_c\left\{(1-\alpha) + \left(\frac{\beta}{\eta}\right)\left(\frac{\bar{R}_c(t)}{R_c(t)}\right)\right\}, \ c \in C \quad (8)$$

$$q_p(t) \leftarrow q_p(t-1) + \alpha - p_p\left\{\alpha + \left(\frac{\beta}{\eta'}\right)\left(\frac{\bar{R}_p(t)}{R_p(t)}\right)\right\}, \ p \in \mathcal{P} \quad (9)$$

We omit the proofs, which are standard, in the interest of space. The adaptation policies indicate how the access probability can be adapted by each client (or BS) independently (based on its communication or positioning role) during both success (incrementing q by α or $1-\alpha$) and failure (decrement related to $1/\Delta U(t) = \bar{R}(t)/R(t)$) of medium access. This allows the joint communication-positioning system (described in Equation 5) as a whole to converge to a stable and optimal operating point.

4.3.2 PAC in Practice with WiFi FTM and 802.11e. While the above stochastic policies can be easily realized at clients for CP in a synchronous (slotted) access system such as cellular networks, asynchronous access networks such as WiFi do not follow a slotted

system. Guided by the optimal policies, *CO2P* realizes their essence within the contention access framework provided by WiFi.

In particular, *CO2P* adopts WiFi's prevalent QOS-based contention framework (802.11e) for differentiated traffic access and instruments it to achieve our desired policies. Note that 802.11e employs three key parameters to control channel access: (i) TXOP - transmission opportunity that controls the maximum time for a transmission, (ii) AIFS - access interframe space, which is the time that every client needs to defer between new transmissions (equivalent of DIFS in regular MAC). AIFS is adapted based on AIFSN values that are varied across traffic types as follows.

$$AIFS_i = t_s \cdot AIFSN_i + SIFS \quad (10)$$

where t_s is the slot duration and SIFS is the short interframe space; and (iii) $CW_{\min}$ and $CW_{\max}$ – minimum and maximum contention window parameters used in contention resolution.

Access between Positioning and Communication Traffic: While the proposed distributed access policies in Section 4.3.1 will work in a completely decentralized network, *CO2P* can leverage WiFi's inherent BS/AP-driven access policy to devise a more efficient access mechanism between the two traffic types. The ratio of channel access between *CO2P*'s communication and positioning (FTM) traffic ($\frac{1-\alpha}{\alpha}$) is controlled by regulating their respective access duration (A_i^t) in each epoch (T) as follows.

$$\frac{A_C^t}{A_P^t} = \frac{1-\alpha}{\alpha}, \text{ and } A_C^t + A_P^t = T \quad (11)$$

The enterprise controller (EC) configures a trigger interval at its APs to correspond to the periodicity of the FTM measurement phase, namely the epoch (T). The trigger control packets (e.g. CTS-to-Self in WiFi) [1] notify (i) the time allocated for FTM measurements between clients (A_P^t), which immediately follows the trigger packet, and (ii) the bandwidth for use by clients for their FTM measurements to enable progressive ranging (Section 4.2.1). Notified by these triggers, clients contend for peer-peer FTM transmissions only during their allocated duration (A_P^t), while the AP waits to schedule its data transmissions on the same channel in the following A_C^t duration of the epoch. Isolating and coordinating channel access between the two traffic types using trigger packets, also helps ensure efficient channel access within each of the respective traffic classes, especially given the significant difference in the size of the data and FTM packets. Further, the EC can periodically evaluate its positioning accuracy as in the centralized approach to adapt its α as necessary.

Decentralized Access between Positioning Traffic: While access between communication/data transmissions can follow conventional DCF contention parameters, *CO2P* needs to regulate decentralized access between peer positioning clients, while accounting for FTM-specific features. Inspired by our distributed access policy from Section 4.2.1, *CO2P* leverages the incremental utility based on positioning throughput ($\Delta U_p(t) = \frac{R_p(t)}{\bar{R}_p(t)}$ from Equation 7), as the priority function used by each client, either to initiate an FTM measurement or to respond to one, to determine its appropriate AIFSN and hence the deferred period during channel access. For a client p, its AIFSN is adapted as follows.

$$AIFSN_p = f\left(\frac{1}{\Delta U_p(t)}\right) \quad (12)$$

where $f(\cdot)$ maps the inverse incremental utility to an integer value in the window [2,7] using a logarithmic function. Thus,

---

Algorithm 2: Distributed PAC operations

input : FTM overhead α; Period time T

1 **while** *Transmissions* **do**
2   Calculate A_C^t, A_P^t by given α, T;
3   **if** *FTM transmission* **then**
4   APs notify FTM measurements to all clients with access duration (A_P^t);
5   Client access channel with AIFSN by computing Eq. (12);
6   If a client gets channel access, broadcast FTM transmission;
7   Other clients receive broadcast FTM and record the timestamps they received the packet;
8   **else if** *Data transmission* **then**
9   APs notify Data transmission to all clients with access duration (A_C^t);
10   Clients upload FTM measurements to EC;
11   Clients start data transmission;
12   **end**
13   Move on to next epoch;
14 **end**

Figure 8: PAC operations

the higher the incremental utility, the smaller the AIFSN number and hence the AIFS defer period, providing prioritized access. As can be seen from Equation 12, a higher priority is given to a client with fewer FTM transmissions, unless a sufficient number of transmissions from neighbors have already been received by it. Further, adapting AIFS without adapting their contention windows, retains the properties of the underlying backoff-based contention mechanism in WiFi, while still allowing for prioritization of relative clients' FTM measurements.

In contrast to FTM ranging that was designed for IP between two specific devices (a client and an AP), *CO2P* enables *concurrent ranging* that is more efficient for CP between peer clients. Here, a client i initiates an FTM packet as a broadcast transmission, which is followed by other neighbors sending their own broadcasts (governed by channel access described above). Every client listens to each of these broadcast FTM transmissions and records their received timestamps (Fig. 7b). The aggregated information allows for the estimation of ranges between multiple client pairs with a much lower overhead (3 transmissions per contention region) and lends itself nicely for *CO2P*'s passive ranging, where the emphasis is simply on transmissions from clients, without differentiating between destination-specific FTM requests and responses. Further, lesser transmissions also contribute to less hidden terminals. Note that, every client will contend for access for a maximum of S FTM transmissions (estimated to be 5 from our experiments.), to provide sufficient measurements to its neighbors for an accurate range estimate at each of them. The information collected from each packet (Fig. 7b) at a client is aggregated and conveyed through a data packet to the central controller to compute the desired ranges. An overview is provided in Fig. 8.

4.3.3 Ranging across Multiple Channels. In a practical enterprise WiFi network with multiple APs, it is possible for neighboring APs to be operating on a different channel to reduce inter-cell interference. This would limit the ability of clients (operating on different channels) at the edge of two cells to discover neighbors and establish ranges across the cells, thereby affecting the connectivity of the peer-peer topology, which in turn is essential to enable CP.

Given the ability of current WiFi (802.11ac/ax) networks to enable primary and secondary operating channels for each AP and its clients, *CO2P* leverages this capability to operate clients on multiple channels (albeit not simultaneously), including those used by its neighboring cell for positioning, to address this challenge. *CO2P* employs the following mechanism involving two steps: (i) *Identifying edge clients*: each client i measures the RSSI from two

strongest APs (one being its own associated AP j, and another neighboring AP k) and compares their relative RSSI. If the net difference is within a threshold, indicating that the client received comparable signal strength from both APs and hence is likely to be an edge client. While RSSI is known to be a sub-optimal indicator of relative client locations, we can use a vector of such RSSIs collected across multiple channels to jointly establish the edge nature of a client with higher accuracy. (ii) *Stochastic ranging on multiple channels:* A client that is an edge to two APs, A_1 and A_2 operating on channels C_1 and C_2 respectively, will probabilistically select one of these two channels to contend on for its FTM transmissions in each epoch, while data transmissions will continue to operate on its AP's primary channel. This allows edge clients from neighboring cells to stochastically discover neighbors through rendezvous on each others' channel. As time progresses, an edge client can choose to weight its probability of employing a particular channel for ranging depending on its history of measurements across channels. Note that, while the allocation of channels to APs for data traffic can be better optimized by taking into account the impact of positioning as well, this is beyond the scope of this work.

5 EVALUATION

CO2P is fully implemented and comprehensively evaluated in NS3. We start with our simulation set-up (Sec. 5.1), followed by a characterization of *CO2P*'s overall performance in efficiently addressing the communication-positioning tradeoff (Sec. 5.2), and a deeper analysis into *CO2P*'s design components that contribute to its performance (Sec. 5.3).

5.1 Simulation Setup

We consider an 802.11ax WiFi-based enterprise network consisting of three APs, each separated about 15 m from one another in a triangle deployment. This canonical network is sufficient to capture the desired contention dependencies between different transmissions such that the resulting inferences from our study are equally applicable to a larger network with more cells. Unless otherwise mentioned, the APs and clients operate on an 80 MHz channel, with WiFi's conventional rate (MCS) adaptation on the link. A saturating traffic of 100 Mbps is considered for each of the clients. For FTM, measurements are conducted in bursts and each burst contains 5 FTM measurements to provide a better estimate of the range. The error model for FTM measurements varies based on the bandwidth employed (lesser mean and variance with higher bandwidth) and is informed by our experiments with commercial devices [7]. Each simulation lasts for several seconds spanning numerous epochs, where an epoch duration lasts a few hundred milliseconds. For example, within an epoch duration of 100 ms, an $\alpha = 0.2$ would correspond to 20 and 80 ms of channel access for FTM and communication traffic respectively.

We compare both centralized and distributed versions of *CO2P* for CP with the corresponding versions of IP, and study their performance as a function of increasing (i) overhead (allocated α), (ii) client density that exacerbates the need for more measurements, and (iii) client mobility that restricts the overhead for FTM to deliver timely, yet accurate location estimates. While 802.11's DCF MAC is used for communication and positioning traffic in IP, *CO2P*
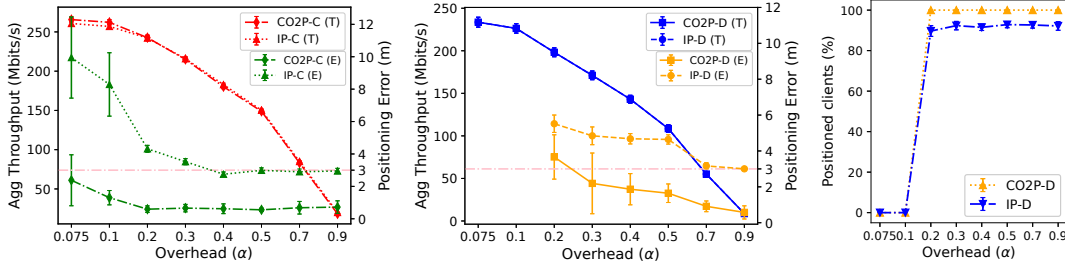
Figure 9: *CO2P* vs. IP left to right: (a) Communication vs. Positioning (Centralized) (b) Communication vs. Positioning (Distributed) (c) Positioned clients (Distributed)
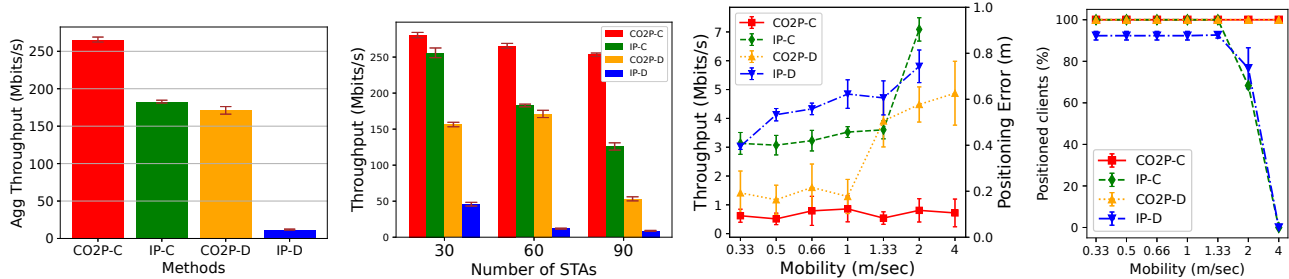

Figure 10: Left to right: (a) Throughput performance (PosErr ≤ 3m) (b) Throughput across node density (PosErr ≤ 3m) (c) Mobility: positioning error (d) Mobility: positioned clients

employs DCF MAC only for the communication traffic, while leveraging its proposed distributed PAC for FTM traffic. Aggregate communication throughput and average positioning error are the metrics of interest, and results are averaged over tens of topologies.

5.2 Overall Performance

Impact of limiting overhead: Figs. 9a and 9b present both the communication throughput and positioning error as a function of increasing the measurement budget for positioning in the network with 60 clients. While an impact on throughput with increasing overhead is expected, it is interesting to note that while *CO2P* requires very little overhead (0.075) to reduce its positioning error to under 3m, IP incurs substantial overhead (0.4) to improve its accuracy to 3m (Fig. 9a). This is further exacerbated in a practical distributed network, where additional overhead is needed in both schemes (0.3 in *CO2P* and 0.7 in IP) to achieve the same accuracy of 3m as shown in Fig. 9b. However, IP suffers more due to its inability to obtain sufficient measurements for all its clients as shown in Fig 9c, resulting from increased hidden terminal impact – with every client reaching out to two of the other, potentially farther APs (other than its own) for their FTM measurements, this increases the scope of collisions that are concentrated at the APs.

Fig. 10a better captures the resulting impact on communication throughput for a fixed/desired positioning error of 3m, where *CO2P* delivers a gain of over 60% in the centralized case, and multiple folds in the distributed case. *By leveraging both its overhead-reducing ranging mechanisms along with an efficient distributed MAC, CO2P is able to significantly alleviate the impact of hidden terminals, while still obtaining sufficient measurements to deliver a seamless coexistence of communication and positioning traffic.*

Impact of client density: Fig. 10b captures the impact on throughput for desired positioning error of 3m, as a function of increasing client density. Clients are also mobile with a speed up to

1m/s. Increasing density increases the need for more positioning measurements, affecting the throughput performance. This can be observed with both *CO2P* and IP. Compared to *CO2P*, IP suffers a higher degradation in communication performance, delivering a gain of over 100% for *CO2P*. While *CO2P* incurs a slightly larger impact for 90 clients, note that IP is able to position only 89% of the clients (100% for *CO2P*).

Impact of client mobility: We consider client mobility ranging from almost no mobility (0.3 m/s) human clients to an autonomous agent (4 m/s ≃ 9 mph), in a network with 60 clients and an overhead of 0.3. Fig. 10c and 10d highlight that *CO2P* is able to deliver under 2m accuracy even with its distributed scheme, for speeds at or under 1 m/s, which can be further extended to higher speeds (for autonomous agents) with a slightly increased overhead of 0.4. In contrast, IP is unable to position all clients at high speeds and reach 3m accuracy even for almost static clients. This showcases *CO2P's ability to accurately track numerous high speed agents in the enterprise without impacting communication performance appreciably.*

5.3 Understanding the Benefits

To better understand the merits of *CO2P*'s design, we now isolate the impact of its key features.

Impact of passive ranging: Fig. 11a captures the number of measurements incurred in the distributed schemes of IP and *CO2P*. To achieve a desired 3m positioning accuracy in a network with 60 clients, *CO2P* requires an order of magnitude fewer measurements. This can be related to passive ranging that results in $O(1)$ overhead that does not scale appreciably with growing density, compared to IP's $O(N)$ overhead. This is a key feature that allows *CO2P* to realize CP in larger network densities, which otherwise would typically require $O(N^2)$ overhead.

Impact of distributed positioning access control: Fig. 11b captures the benefits of *CO2P*'s decentralized positioning access
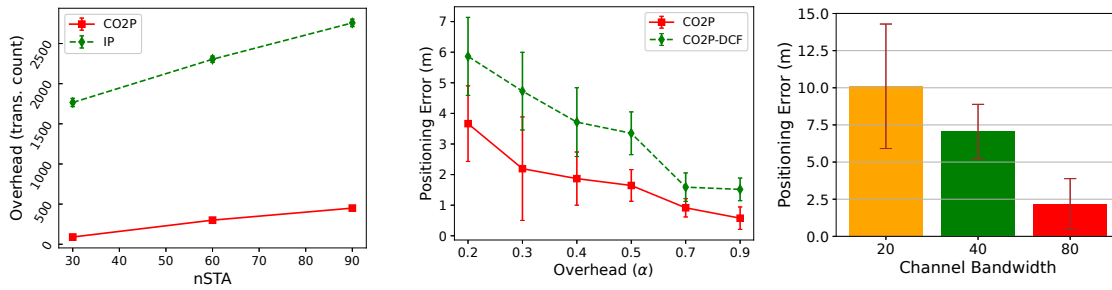
Figure 11: Left to right: (a) FTM overhead for reaching positioning error threshold (3m) (b) CO2P MAC vs. DCF (c) CO2P Progressive Ranging Benefit

control scheme, by comparing it to other versions of itself that differ in only the MAC employed (i.e. using DCF MAC for FTM) in a network with 60 mobile clients (1 m/s speed) and overhead of 0.2 − 0.9. It can be clearly seen that *CO2P*'s ability to introduce differentiated access within FTM traffic, allows it to maximize the utility of the measurements collected within a given overhead, thereby delivering significantly better positioning accuracies.

Impact of progressive ranging: Fig. 11c captures the importance of *CO2P*'s progressive ranging. Starting with a higher bandwidth of 80 MHz for ranging (coupled with its passive ranging), helps it identify a sufficient number of lower error ranges to determine accurate positioning for all clients without moving to lower bandwidths for identifying longer links in the topology. Ranges obtained on lower bandwidth links are vulnerable to higher range errors, affecting the positioning accuracy of all the clients.

6 CONCLUSIONS

This work proposed a fundamentally different paradigm of communication-aware collaborative positioning *CO2P*, whereby the burden of positioning in future wireless systems is offloaded to client devices in an intelligent, communication-aware and collaborative manner that reduces overhead and improves spatial reuse to preserve communication performance, without compromising on positioning accuracy. Through technically-sound algorithms and a practical, distributed solution that is realized using WiFi's positioning protocol, *CO2P* addresses the underlying tradeoff between communication performance and positioning accuracy that impacts today's infrastructure positioning solutions and delivers an efficient coexistence.

7 ACKNOWLEDGEMENT

REFERENCES

[1] 2016. IEEE Standard for Information technology—Telecommunications and information exchange between systems Local and metropolitan area networks—Specific requirements. *IEEE Std 802.11-2016* (2016), 1–3534.
[2] 2016. Wi-Fi ranging with RTT. https://developer.android.com/guide/topics/connectivity/wifi-rtt.
[3] 2021. Operator Defined Open and Intelligent Radio Access Networks. https://www.o-ran.org/resources.
[4] 2022. Network Simulator (ns-3 v3.36). https://www.nsnam.org/.
[5] 3GPP. 2017. *Proximity-based services (ProSe); Stage 2.* Technical Specification (TS) 23.303. 3rd Generation Partnership Project (3GPP). https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=840 Version 17.0.0.
[6] 3GPP. 2018. Study on NR positioning support. *Release 16 Technical Report* (2018).
[7] Shivang Aggarwal, Ramanujan K Sheshadri, Karthikeyan Sundaresan, and Dimitrios Koutsonikolas. 2022. Is Wifi 802.11mc Fine Time Measurement Ready for Prime-Time Localization? *(ACM WiNTECH)*.
[8] Ala'a Al-Habashna and Gabriel Wainer. 2020. RSSI-based Indoor Localization with LTE-A Ultra-Dense Networks *(IEEE SPECTS)*.
[9] Abdo Y Alfakih, Amir Khandani, and Henry Wolkowicz. 1999. Solving Euclidean distance matrix completion problems via semidefinite programming. *Computational optimization and applications* 12, 1-3 (1999), 13–30.
[10] K. S. Arun, T. S. Huang, and S. D. Blostein. 1987. Least-Squares Fitting of Two 3-D Point Sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1987), 698–700.
[11] Jose A Costa, Neal Patwari, and Alfred O Hero III. 2006. Distributed weighted-multidimensional scaling for node localization in sensor networks. *ACM Transactions on Sensor Networks (TOSN)* 2, 1 (2006), 39–64.
[12] José A. del Peral-Rosado, Ronald Raulefs, José A. López-Salcedo, and Gonzalo Seco-Granados. 2018. Survey of Cellular Mobile Radio Localization Methods: From 1G to 5G. *IEEE Communications Surveys Tutorials* 20, 2 (2018), 1124–1148. https://doi.org/10.1109/COMST.2017.2785181
[13] Carmelo Di Franco, Enrico Bini, Mauro Marinoni, and Giorgio C. Buttazzo. 2017. Multidimensional scaling localization with anchors *(IEEE ICARSC)*.
[14] Ivan Dokmanic, Reza Parhizkar, Juri Ranieri, and Martin Vetterli. 2015. Euclidean distance matrices: essential theory, algorithms, and applications. *IEEE Signal Processing Magazine* 32, 6 (2015), 12–30.
[15] S. Dwivedi. 2021. Positioning in 5G networks. *Arxiv 2102.03361* (2021).
[16] Tolga Eren, OK Goldenberg, Walter Whiteley, Yang Richard Yang, A Stephen Morse, Brian DO Anderson, and Peter N Belhumeur. 2004. Rigidity, computation, and randomization in network localization *(IEEE INFOCOM)*.
[17] Sanvidha CK Herath and Pubudu N Pathirana. 2013. Optimal sensor arrangements in angle of arrival (AoA) and range based localization with linear sensor arrays. *Sensors* 13, 9 (2013), 12277–12294.
[18] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi *(ACM SIGCOMM)*.
[19] Daniel Neuhold, Christian Bettstetter, and Andreas F Molisch. 2019. HiPR: High-precision UWB ranging for sensor networks *(ACM MSWIM)*.
[20] Lei Ni, Yuxin Wang, Haoyang Tang, Zhao Yin, and Yanming Shen. 2017. Accurate localization using LTE signaling data *(IEEE CIT)*.
[21] D.P. Palomar and Mung Chiang. 2006. A tutorial on decomposition methods for network utility maximization. *IEEE Journal on Selected Areas in Communications* 24, 8 (2006), 1439–1451. https://doi.org/10.1109/JSAC.2006.879350
[22] Neal Patwari, Joshua N Ash, Spyros Kyperountas, Alfred O Hero, Randolph L Moses, and Neiyer S Correal. 2005. Locating the nodes: cooperative localization in wireless sensor networks. *IEEE Signal processing magazine* 22, 4 (2005), 54–69.
[23] Neal Patwari, Alfred O Hero, Matt Perkins, Neiyer S Correal, and Robert J O'dea. 2003. Relative location estimation in wireless sensor networks. *IEEE Transactions on signal processing* 51, 8 (2003), 2137–2148.
[24] Elahe Soltanaghaei, Avinash Kalyanaraman, and Kamin Whitehouse. 2018. Multipath Triangulation: Decimeter-level WiFi Localization and Orientation with a Single Unaided Receiver *(ACM MobiSys)*.
[25] K. Sundaresan and Sampath Rangarajan. 2009. Efficient resource management in OFDMA Femto cells *(ACM MobiHoc)*.
[26] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Decimeter-level Localization with a Single WiFi Access Point *(USENIX NSDI)*.
[27] Henk Wymeersch, Jaime Lien, and Moe Z. Win. 2009. Cooperative Localization in Wireless Networks. *Proc. IEEE* 97, 2 (2009), 427–450. https://doi.org/10.1109/JPROC.2008.2008853
[28] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. 2019. mD-Track: Leveraging Multi-Dimensionality for Passive Indoor Wi-Fi Tracking *(ACM MobiCom)*.
[29] Anatolij Zubow, Christos Laskos, and Falko Dressler. 2022. FTM-NS3: WiFi Fine Time Measurements for NS3 *(IEEE WONS)*.