Characterizing Location Management Function Performance in 5G Core Networks

Andrea Pinto♯ Giuseppe Santaromita⋆ Claudio Fiandrino⋆ Domenico Giustiniano⋆ Flavio Esposito♯
♯Computer Science Department, Saint Louis University, USA, ⋆IMDEA Networks Institute, Madrid, Spain

***Abstract*—Despite the large attention achieved by 5G localization in standardization bodies, the integration of 5G network function modules designed for localization lacks experimental work. Assessing the performance of these modules is essential to offering location services. In this work, we present our design, implementation and evaluation of the 5G Location Management Function (LMF), the key network function in the 5G core for localization services. Our implementation complies with the 3GPP standard and OpenAirInterface, the currently most advanced framework that implements a full 5G-New Radio stack. We show that we can extend the functionality of OpenAirInterface, enabling location services. Finally, we demonstrate that the performance of our implementation satisfies the 5G Key Performance Indicators required by 3GPP for localization.**

***Index Terms*—5G localization, 5G core, network functions, localization and management function, geospatial systems.**

I. Introduction

The 5G network architecture introduces by design a high degree of virtualization of Network Functions (NFs) and services. One of these functions is the Location Management Function. Early work on localization and positioning within Long Term Evolution (LTE) wireless networks was primarily driven by emergency and law enforcement requirements. Most recently, however, Location Based Services (LBSs) relying on 5G networks are constantly growing. Industry 4.0, autonomous vehicles, emergency services, augmented reality, and Internet of Thing (IoT) for mobile health and precision agriculture are examples of commercial use cases with critical requirements for localization accuracy and reliability according to the 5G specifications [1], [2]. This uptake is driven by key 5G technological developments, such as network densification, the wider network bandwidth, and scalability support with respect to number of antennas and users. Localization is also key to help optimize specific network mechanisms [3], especially in the context of millimeter-wave networks [4].

User Equipment (UE)'s location techniques defined in the 3rd Generation Partnership Project (3GPP) are divided into three main categories: standalone, user-based, and network-based [5]. Prior releases of 3GPP mainly focused on user-based and standalone strategies, which depend on the users localizing themselves with or without network support, but mostly with Global Navigation Satellite Systems (GNSS) technology embedded in mobile devices. On the other hand, network-based techniques rely on the cellular network to perform UE's location measurements. For example, indoor localization is a typical scenario in which the benefits of network-based advantages are clear. 5G New Radio signals can improve the localization,

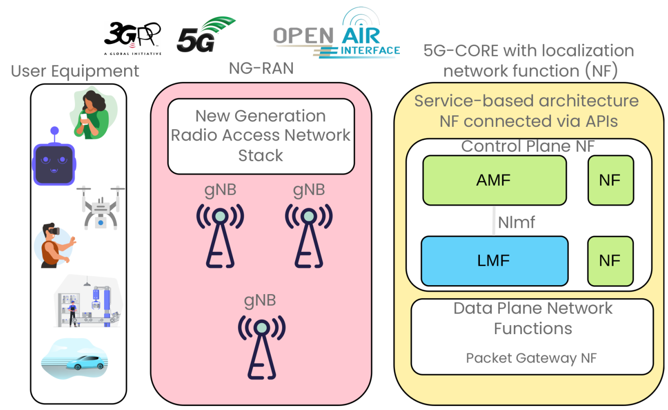

Fig. 1. 5G architecture, including NG-RAN and 5G Core, as well as the new Location Management Function (LMF). The User Equipment (UE) in the NG-RAN is connected to one or more 5G gNBs that elaborate the received localization reference signals for positioning the UE. LMF receives this information on a localization request event via an interface called Nlmf and returns the localization results to the location service client.

especially when GNSS signals alone are ineffective [6]. 5G localization techniques would also greatly help telecom operators and other stakeholders to research new ways to provide positioning services in terms of precision and scalability.

The 3GPP, since Release 16, has been proposing 5G Core (5GC) modules to handle functionalities of suitable position measurements between the 5GC and the UE, enabling several network location-based applications [1]. Considering Fig. 1, the Location Management Function (LMF) is the key enabler for localization purposes proposed by 3GPP. This module is part of the 5GC and can run various localization algorithms; it gathers location data from the Next Generation Radio Access Network (NG-RAN) for each Next Generation Node Base station (gNB) involved in the process, and returns the UE's estimated location.

Our LMF implementation is the first fully integrated localization solution within OpenAirInterface (OAI), an open-source 5GC reference implementation that adheres to the most recent 3GPP standards [1]. We expect our implementation to be a key enabler for a plethora of case studies in the 5G localization context. In particular, the key contributions of this paper are:

(i) we investigate whether the current 5G NFs for positioning can fulfill the required 5G Key Performance Indicators (KPIs), in terms of latency and system scalability, following the definitions of the 3GPP standards;

[1]Code release at: https://github.com/pintauroo/5G_LMF.git.

(ii) we present the first positioning implementation that meets the requirements of 3GPP standards, and it is integrated with OAI, one of the most popular open-source 5GC software suites;
(iii) we analyze and test our system tuning several policies to verify the fulfillment of 3GPP requirements in various scenarios, especially under intensive workloads, and in a real-world scenario using a trace-driven evaluation;
(iv) our investigation characterizes the latency of the location network function services in all its sub-parts.

The remainder of this paper is organized as follows. In Section II we present some related work. In Section III we detail the design of the LMF module and its implementation within OAI. In Section IV we provide our performance evaluation under multiple scenarios, including a trace-driven evaluation with real-world traffic. Finally, in Section V we conclude the work and outline some future research directions.

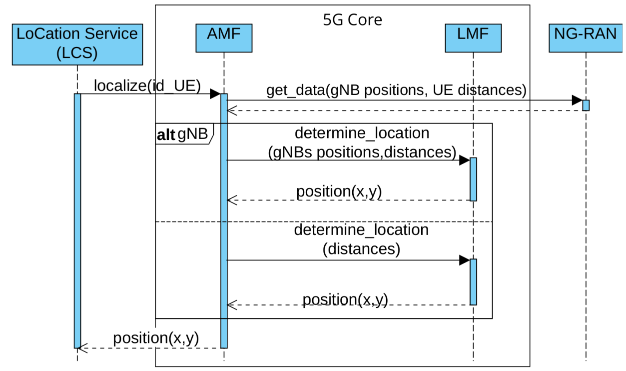

Fig. 2. Our localization service procedure designed. The sequence diagram shows the direct communications between the LCS and the 5GC modules, such as AMF and LMF, via HTTP. This diagram also describes the AMF signals data retrieval process from the NR-RAN that happens via NGAP protocol. The architecture is compliant with the 3GPP standards and, in particular, with TS 38305: Location Service Support by NG-RAN [10].

II. Related Work

Existing work on positioning has mainly focused on the study of suitable methods using 3GPP Reference Signals (RSs) measured in the NG-RAN, addressing positioning accuracy for a single UE and its applications [6]–[8]. No publicly available work featuring the integration of localization functionalities within the 5GC as the LMF NF has reached our practical level. Given the novelty of the study, it was not possible to compare our study to other's work in terms of performance or implementation.

Among the few solutions that considered networking aspects with multiple UEs, the authors in [1] described 5G use cases and KPI presented by the 3GPP where accurate positioning is required. The authors specifically mentioned that the KPI requirements in 5G scenarios could be met by utilizing various signal processing techniques and technologies, which typically have limited communication, processing, and memory resources. In [6], the authors described how cellular networks can be exploited for localization (in the context of contact tracing), especially considering 5G and beyond's NFs. Bartoletti et al. [5] contain a high-level analysis of cutting-edge applications in 5G and beyond, focusing on people-centric and network-centric location-based analytics, and a proposal to send positioning data from LMF to a new module called LDAF to extract on-demand analytics that serves third-party applications or optimize the network performance. Examples are analytics for mobility clustering, network optimization, and network diagnosis. The authors in [9] proposed the positioning uncertainty quantification as an analytics function in the 5GC. They developed a virtualized system based on a simulated LMF module and a machine learning-based approach for predicting and updating the level of localization uncertainty in a monitored environment.

In all of the cases mentioned above, LMF was given as an existing and already working module. Our LMF implementation and study would have improved their study quality, promising a real emulation test of the system with the entire 5GC stack, including details of our analyzed KPIs.

III. Location Management Function: Design and Implementation within the 5G Core

5G Core Network: Overview and Context. The 5G Core network architecture is designed to be natively deployed and operated on top of virtualized infrastructures using NFs. The main components of the 5G architecture, shown in Fig. 1, are the UE, i.e., any device used directly by an end-user or a communicating application process, the gNB, the NG-RAN, and the 5GC. The NG-RAN is a distributed network of gNBs that orchestrates multiple functionalities, such as radio spectrum usage or a handle range of connectivity options. The 5GC is a Service-based Architecture (SBA), that uses NF to offer one or more services to other NF via Application Programming Interface (API). The overall architecture is integrated with modern virtualization and softwarization process, separating the NFs from the hardware. Consequently, the core is divided into control and data planes. While the control plane gives directives on the 5G network handling, the data plane contains a NF that acts as a gateway to the public Internet.

Considering such 5G core architecture, in the rest of this section, we illustrate the LMF support in the 5GC and the design of our prototype that follows the 3GPP principles. We characterize and analyze our LMF implementation following these specifications. Specifically, we measure latency and system scalability requirements under different system loads, frequency between two consecutive localization requests (localization inter-packet gap (IPG)), and Quality of Service (QoS) to control traffic and ensure the performance of critical applications.

Generally, LMF can be used to determine the position of a UE, to coordinate and schedule resources, or to calculate or verify a final location, estimate velocity, and compute the achieved location accuracy [11].

A. Location Management Function: Design and Integration with System 5G Core

The LoCation Service (LCS) Client is an entity whose purpose is to obtain location information to run location services [12]. Providing LCS to a UE is a collaborative endeavor, including multiple 5GC NFs and APIs. Different applications require different Location policies based on the UE's supported technologies and requirements. Specifically, these policies comprise different parameters like the QoS, the required coordinate shape universal geographic area description (GAD), and IPG. The QoS is used to characterize the location request based on accuracy and response time values. The error primarily determines the accuracy in the measurements of the signals received by the NG-RAN and the localization algorithm selected depending on the required accuracy and latency.

Localization service procedure. Fig. 2 shows a generic LCS requesting localization services to 5GC to start the determine-location stimulus via an HTTP GET request on the *localization-ue/UE_id* endpoint. According to the specification, the UE ID should be handled by Access and Mobility Function (AMF) [12]. The localization request on this endpoint might be external or internal to the 5GC; an authorized stakeholder that aims at determining the UE location over an endpoint (that we define as *localize_ue*) passes the UE id[2]. Then AMF interacts with LMF and sends the two input values such as the gNB positions and the relative UE distances via the *determine_location* endpoint. Moreover, AMF has to retrieve the input values from the NG-RAN to feed the localization algorithm running in LMF. Note that the connection happens via the NGAP protocol, but we simulate this mechanism as described below. Once LMF receives UE data from one or more gNB, the selected LMF localization algorithm computes the user position, and the result is sent back to the LCS via AMF.

Network-centric localization. Our focus is on network-centric localization, which enables UE localization purely based on signal processing in the 5GC. Note that the NG-RAN processes the UE signals first, and then the 5GC elaborates upon these results. Based on these signals and the location policy set by the user or required by the application, LMF should guarantee different services levels, and for each of these, dispose of different algorithms able to calculate the UE position using the measurements received from the NG-RAN.

Although we design our LMF module to be flexible and incorporate different localization algorithms, in this work, we adopt and integrate a well-known trilateration algorithm implemented in C++ [13]. The algorithm, used in Microsoft localization competitions, implements the Newton method to minimize the localization cost function and uses the Armadillo library to operate with matrices at each step of the gradient descent algorithm. The applicability of the LMF to 5G networks can be in contexts wherein the NG-RAN computes the distances using Multi-RTT (round trip time), which is mapped to estimated distances. Two RSs are used to compute multi-RTT in the 3GPP standards, that is, the downlink positioning reference signal (DL-PRS) and Uplink sounding reference signals (UL-SRS). However, our LMF implementation is flexible, and other algorithms can be integrated in the future. As a result, since the UE can move and different gNB can participate in the localization process, AMF must be concerned with identifying which are the reference gNBs by retrieving this information from the NG-RAN. If the UE moves, AMF has to send the new gNBs positions to LMF to determine the location. Otherwise, AMF only needs to send the updated distances.

[2]Our implementation simplifies the 3GPP mechanism, which defines that the LCS has to interact with an LCS server called Gateway Mobile Location Centre (GLMC), in charge of handling interactions with AMF.

B. LMF Implementation

Location Management Function in the 5G Core. The 5GC is designed by the new 3GPP defined SBA as a set of Virtual Network Functions (VNFs). Each NF can have very different performance requirements. The message bus between the SBA parts is called Service-based Interface (SBI). The SBI employs REpresentational State Transfer (RESTful) API principles over HTTP/2. The most popular implementation of the 5GC is called OAI[3]. The main OAI objective is to develop a fully 3GPP compatible 5GC stack as open-source software. Specifically, the OAI SBI are implemented using the C++ framework called Pistache[4].

We provide a new NF called LMF, which is implemented in the OAI 5G Core in accordance with their Dockerized environment. LMF has a single SBI with AMF called *Nlmf*. 3GPP describes all the available endpoints with information about the exchange data in an OpenAPI format [11]. Our implementation focuses on the *determine-location* endpoint, designated to retrieve the *LocationData* within the Determine Location Response message. This endpoint uses an HTTP POST request accordingly to Fig. 2, and the data exchanged in the HTTP body are encoded in a JSON format.

Localization policies. A localization request needs to comply with the current KPIs, e.g., latency, accuracy, and inter-packet gap (IPG). Localization services accept a variety of UE policy requirements based on the UE's speed, location, or application requirements. As an example, an emergency localization request policy might require low response latency, high localization accuracy, and low IPG to have high-frequency updates. However, only a limited number of UEs may be served to fulfill these application requirements. Latency is also affected by the QoS provided to that specific user. In our implementation, we define two types of QoS: "Stringent QoS" and "Slack QoS". The stringent QoS causes no additional delays other than those caused by system processing and signaling; The slack QoS allows an extra latency of 10 ms. Finally, the Inter Packet Gap measures the interval between the UE requests and, in our setting, it ranges between 100 and

[3]OpenAirInterface: https://openairinterface.org

[4]Pistache framework: https://github.com/pistacheio/pistache/

500 ms. Small intervals may be needed, e.g., for cyber-physical system applications such as haptic wearable devices [14], and larger intervals for power-constrained UEs, such as those used in precision agriculture [15].

IV. Prototype evaluation

This section presents the evaluation results that demonstrate the effectiveness of the LMF network function. Our prototype has been tested using real-world traffic collected from mobile production networks.

Evaluation metrics and methodology. The KPI for 5G location-based applications have been chosen according to standard practice [1] and we focus on two of these: 1) the *latency* describes the time elapsed between the starting time of the process that queries the determination of the position-related data and the availability of such data at the positioning system interface; 2) and the *system scalability* specifies the number of devices for which the positioning system can calculate position-related data in a given time and at a given update rate.

To evaluate the latency, we follow the 3GPP [16] guidelines. In particular, we evaluate the "higher layer latency" involving NG-RAN, LCS, AMF, and LMF nodes and signaling delays between nodes. The processing times in accordance with the 3GPP standards are T^{proc}_{AMF}, T^{proc}_{LMF}, and T^{proc}_{gNB}. Thus, the overall latency is determined by the AMF request serving time overhead, the LMF request serving time overhead with the algorithm execution time, and the time required to retrieve input data for the localization algorithm from the NG-RAN. The signaling times between each component are negligible given that the (Docker) containers run on a single server, and thus omitted to avoid biasing the measurements with the performance of the lightweight virtualization system used e.g., Docker, Kubernetes, or native Linux Containers LXC. We did instead consider the interaction between AMF and NG-RAN.

The interaction with the NG-RAN is simulated, with appropriately added delays in compliance with the 3GPP guidelines [17]. Specifically, we assume that the T^{proc}_{gNB} is fixed at 1 ms plus a 3 ms delay for the T^{sign}_{gNB} signaling procedure. In total, the interaction with the NG-RAN is expected to take a minimum of 4 ms, given the above mentioned times. In our implementation, the T^{proc}_{LMF} is composed of two parts: the LMF server processing and the localization algorithm execution time. The localization algorithm is a policy, hence the execution time might change with different algorithms. Its performance in our implementation range between 500 μs and 1 ms. Thus the latency time λ we consider is defined as:

$$\lambda = T^{proc}_{AMF} + T^{proc}_{gNB} + T^{sign}_{gNB} + T^{proc}_{LMF}$$

As a result of these observations, we conclude that the minimum delay for a single localization request is around 5 ms. We describe our evaluation findings in the following section.

A. Performance Evaluation

In this section, we discuss the performance evaluation findings relative to our LMF implementation in terms of latency and scalability. To emulate the LCS requests for a specific UE system workload, we used JMeter[5], a well-known performance evaluation tool.

The goal of this performance analysis is to find the largest possible number of users that our system can handle optimally. As a first step, we characterize the user behavior that we replicate in our testbed by setting constraints based on the localization policy described in Section III. In our experiments, we have set all user requests with an IPG of 300 ms. Then, we vary the number of users for each test case to process different system loads and determine the usable, nominal, and *knee capacity*. Such capacity describes the point after which the system degradation accelerates, throughput decreases, and latency increases rapidly. As a *usable capacity*, we denote instead the maximum throughput achievable without exceeding a pre-specified response-time limit. Finally, the *nominal capacity* denotes the maximum achievable throughput in requests per second under ideal workload conditions, but the response time is too high.

The number of users in each emulation ranges from 100 to 1000, with an interval of 100 for each experiment, and to sample the population, we assumed that for each emulation, each user has to send 1000 requests Each emulation is executed 30 times to find a representative population statistic over each user number.

We monitor the system statistics during each simulation, considering the most impacting ones as CPU, memory, and network utilization to determine the system scalability over the given constraints. The results of these experiments are reported in Fig. 3. The x-axis represents the number of users in each of the three subfigures, while the y-axis represents CPU utilization in Fig. 3(a), throughput in terms of requests per second in Fig. 3(b), and latency in milliseconds in Fig. 3(c). Observing the upper plot, we note that as the number of users grows, the CPU utilization of both AMF and LMF grows almost equally. Note that a 100% CPU level corresponds to a full single core utilization. The middle and lower plot show the throughput of our system and its latency, respectively.

The minimum latency experienced starts at 5 ms for poor loads. However, each emulation is affected by an additional server's processing time overhead due to the high loads increasing the users number. Thus we will show much load our system can handle for each emulation.

Observing Fig. 3, our implementation meets the 3GPP latency requirements, however, as the load increases, the available resources limit the throughput and the latency increases. Under the given conditions, the knee capacity is met with 400 users load since the latency is low and the throughput gradually increases. The usable capacity is around 500 user load, where latency increases and throughput almost reaches its peak of 1500 requests per second (req/sec). Then, depending on the latency requirements, we can set the nominal capacity given our constraints at 1000 users, assuming that throughput is at its peak and latency is always under 1 second. Potentially our

[5]JMeter: https://jmeter.apache.org/

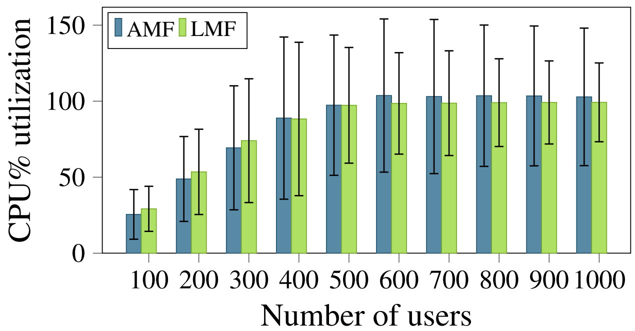

(a) CPU% utilization.

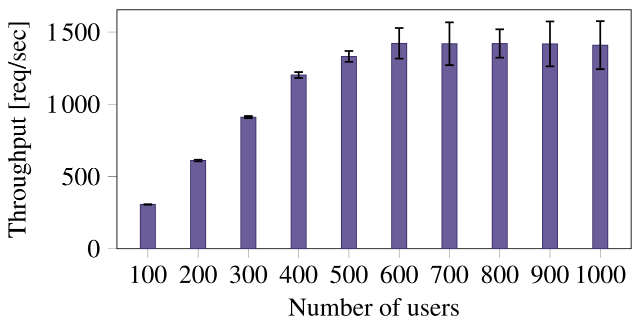

(b) Throughput in terms of served localization requests per second.

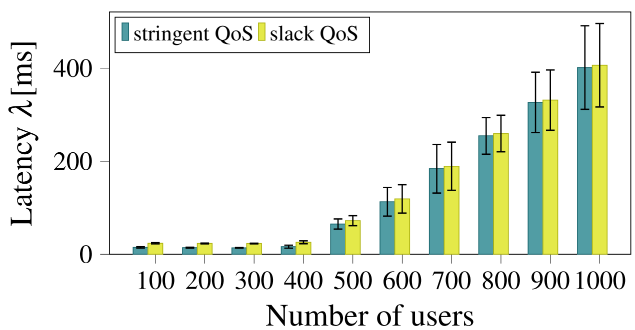

(c) Mean latency in milliseconds.

Fig. 3. 5GC AMF and LMF module performance evaluation. The x-axis depicts the workload given by the number of users for each emulation. By varying this parameter, we can compare the CPU utilization, the throughput in terms of served requests per second, and the effective mean latency with their respective confidence intervals. It is worth noting that the results were obtained by repeating the experiment 30 times for each user workload.

representation is strictly bounded to our simulation scenario, but it gives us an idea of the system performance.

In the next subsection, we show the results of an analysis in a realistic environment where the load is defined by the presence of users in real production networks and the localization requests come with different QoS requirements.

B. Trace-Driven Evaluation

For the trace-driven analysis, we exploit a dataset of LTE traffic allocations from multiple BSs located in different areas of Madrid, Spain [18]. The dataset is collected using a Software-Defined Radio (SDR)-based LTE sniffer tool, i.e., FALCON [19]. FALCON allows decoding the TTI-level (unencrypted) traffic allocations that LTE eNB send to the UEs over

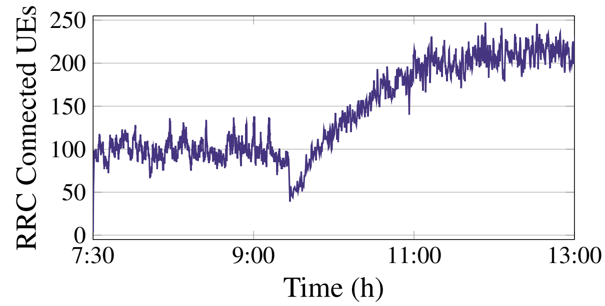

Fig. 4. Number of RRC connected UEs from a production eNB.

the Physical Downlink Control CHannel (PDCCH) channel. Specifically, we gather the temporary user ID (known in 3GPP jargon as Cell Radio Network Temporary Identifier (C-RNTI), the ID of the frame containing the traffic allocation for the C-RNTI. We filter out background traffic by removing Radio Network Temporary Identifiers (RNTIs) that transmit traffic in less than 5 active Transmission Time Interval (TTI) over the entire activity period [20]. We also discard RNTIs that are reserved for random access (i.e., RA-RNTI with ID 1-960), paging and system notification (i.e., P-RNTI with ID 65534), and broadcast system information (i.e., SI-RNTI with ID 65535). We then use the methodology described in [18] to estimate the individual user lifetime. We determine the number of users that are currently connected in active to the Radio Resource Control (RRC) at any point in time. In our analysis, we use this information to assess a realistic number of localization requests that the LMF should handle. Fig. 4 shows an example of active users in one eNB of our dataset for a 5 and a half hour period [21].

Thus, the number of active users is determined by the time series shown in Fig. 4 that we simulate using Locust framework[6], an open-source load testing tool. We consider all the users from this time series requesting the localization service. We reproduce five distinct users request patterns differentiated by IPG and QoS. The IPG for localization requests ranges between 100 ms and 500 ms to simulate different localization frequencies. A QoS mechanism influences user latency by prioritizing users who require higher positioning service levels. We equally distribute generated users connected to the base station between two levels of latency: stringent-quality and slack-quality. Then for this emulation scenario, we evaluate the system's characteristics in responding to the localization requests given by the time series.

Observing the latency distribution in Fig. 5, or the likelihood of a latency under the lambda value, we can notice that the IPG does not affect the UE latency experience; on the contrary it is mostly affected by the QoS. For the stringent case, the latency falls around 20 ms with a very high probability, which is within the standard's requirements. The latency performance gap between all the different IGP is minimal, around 10 ms, and negligible after. Whereas the bottom CDF shows the slack

[6]Locust: https://locust.io/

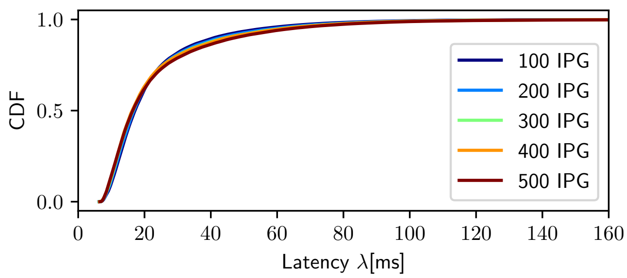

(a) Stringent QoS localization requests.

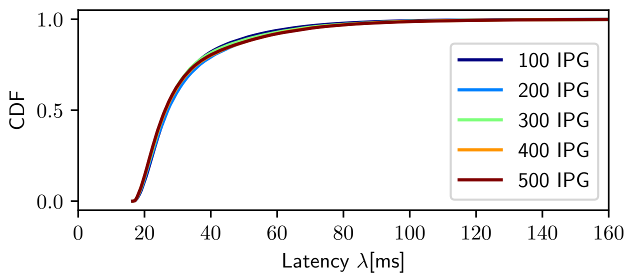

(b) Slack QoS localization requests latency.

Fig. 5. CDF of latency λ of different IPG values ranging from 100 to 500 ms under real users time-series traffic shape (connected UEs from a production base station). Privileged and unprivileged QoS respectively illustrated on top and bottom. The IPG value does not affect the user latency experience but only the number of requests sent over time. The unprivileged user QoS has for each IPG value a 10 ms gap with the privileged users.

QoS with around 20 ms latency at 50% probability. It deserves to be noted that 3GPP requires 1 second latency for most of the localization use cases and a minimum of 15 ms for some demanding localization use cases [17]. We can observe that our system under the user time series meets the 3GPP standards.

V. Conclusion

This paper presented the first design and implementation of the LMF, the key network function in the 5G core for localization services. We integrated it in the 5G Core and released as open source. We described the implementation details of LMF, which is compliant to the widely used OAI 5GC, and characterized its performance and impact as an NFV integrated in the 5G core. We designed LMF in alignment with the latest 3GPP standard requirements. For the evaluation, we performed an extensive system characterization under a broad range of settings and scenarios. We found that under realistic workloads obtained from traffic in mobile production networks, 50% of the localization requests are served in less than 15 ms, which aligns with the standard for stringent localization requirements [17]. Stress tests with heavier workloads to assess the system scalability have shown that localization requests are served in less than 500 ms. We believe that our effort of the LMF fills the gap between prior research of single user localization algorithms and the pressing need to research location-based applications and analytics with a large number of UEs.

Acknowledgment

A. Pinto and F. Esposito have been supported by NSF Awards: # 1908574 and # 2201536. The work by IMDEA Networks was supported by the European Union's Horizon 2020 research and innovation program under Grant no. 871249, by the Juan de la Cierva grant from the Spanish Ministry of Science and Innovation (IJC2019-039885-I) and by the Regional Government of Madrid and the EU through the European Regional Development Fund (ERDF) as part of the response from the EU to the COVID-19 pandemic in the context of REACT-CONTACT-CM-23479 research project, awarded under the Agreement formalized between the Regional Government of Madrid and IMDEA Networks.

References

[1] S. Bartoletti *et al.*, "5G localization and context-awareness," in *5G Italy White Book: From Research to Market*. National, Inter-Univ. Consortium for Telecommunications, 2018, pp. 167–187.

[2] ——, "Positioning and sensing for vehicular safety applications in 5G and beyond," *IEEE Communications Magazine*, vol. 59, no. 11, pp. 15–21, 2021.

[3] D. Macagnano *et al.*, "MAC performances for localization and tracking in wireless sensor networks," in *Workshop on Positioning, Navigation and Communication*, 2007, pp. 297–302.

[4] C. Fiandrino *et al.*, "Scaling millimeter-wave networks to dense deployments and dynamic environments," *Proceedings of the IEEE*, vol. 107, no. 4, pp. 732–745, 2019.

[5] S. Bartoletti *et al.*, "Location-based analytics in 5G and beyond," *IEEE Communications Magazine*, vol. 59, no. 7, pp. 38–43, 2021.

[6] D. Giustiniano *et al.*, "5G and beyond for contact tracing," *IEEE Communications Magazine*, vol. 59, no. 9, pp. 36–41, 2021.

[7] H. Ryden *et al.*, "Cellular network positioning performance improvements by richer device reporting," in *Proc. of IEEE VTC Spring*, 2018, pp. 1–5.

[8] S. Dwivedi *et al.*, "Positioning in 5G networks," *IEEE Communications Magazine*, vol. 59, no. 11, pp. 38–44, 2021.

[9] S. Bartoletti *et al.*, "Uncertainty quantification of 5G positioning as a location data analytics function," in *Proc. of EuCNC/6G Summit*, 2022, pp. 255–260.

[10] 3GPP, "TS 38305, NG radio access network (NG-RAN); stage 2 functional specification of user equipment (UE) positioning in NG-RAN," in *3rd Generation Partnership Project, Release 17.1*, 2022.

[11] ——, "TS 29572, 5G system; location management services," in *3rd Generation Partnership Project, Release 17*, 2022.

[12] ——, "TS 23273, 5G system (5GS) location services (LCS); stage 2," in *3rd Generation Partnership Project, Release 17.5*, 2022.

[13] M. Rea *et al.*, "Time-of-flight wireless indoor navigation system for industrial environment," in *Proc. of ACM WiNTECH*, 2019, pp. 37–44.

[14] A. Sangiorgi *et al.*, "Demo: Propelling haptic devices into the mobile world to advance K-12 STEM learning," in *Proc. of ACM MobiHoc*, July 2019.

[15] W. Muhammad *et al.*, "Polly: A tool for rapid data integration and analysis in support of agricultural research and education," *Internet of Things*, vol. 9, p. 100141, 2020.

[16] 3GPP, "TS 38857, study on NR positioning enhancements," in *3rd Generation Partnership Project, Release 17.5*, 2021.

[17] ——, "TS 22.261, service requirements for the 5G system," in *3rd Generation Partnership Project, Release 18*, 2020.

[18] G. Attanasio *et al.*, "Characterizing RNTI allocation and management in mobile networks," in *Proc. of ACM/IEEE MSWiM*, 2021, p. 189–197.

[19] R. Falkenberg *et al.*, "FALCON: An accurate real-time monitor for client-based mobile network data analytics," in *Proc. of IEEE GLOBECOM*, Dec 2019, pp. 1–7.

[20] N. Bui *et al.*, "Data-driven evaluation of anticipatory networking in LTE networks," *IEEE Transactions on Mobile Computing*, vol. 17, no. 10, pp. 2252–2265, 2018.

[21] C. Fiandrino *et al.*, "Traffic-driven sounding reference signal resource allocation in (beyond) 5G networks," in *Proc. of IEEE SECON*, 2021, pp. 1–9.