

# Collaboration Stability: Quantifying the Success and Failure of Opportunistic Collaboration

Huber Flores<sup>1</sup>, Agustin Zuniga<sup>2</sup>, Leonardo Tonetto<sup>3</sup>, Tristan Braud<sup>4</sup>,  
Pan Hui<sup>2,5</sup>, Yong Li<sup>6</sup>, Sasu Tarkoma<sup>2</sup>, Mostafa Ammar<sup>7</sup>, Petteri Nurmi<sup>2</sup>

<sup>1</sup>Institute of Computer Science, University of Tartu, Estonia

<sup>2</sup>Department of Computer Science, University of Helsinki, Helsinki, Finland

<sup>3</sup>Department of Informatics, Technical University of Munich, Munich, Germany

<sup>4</sup>Department of Computer Science and Engineering, HKUST, Hong Kong

<sup>5</sup>Division of Integrative Systems and Design, HKUST, Hong Kong

<sup>6</sup>Department of Electronic Engineering, Tsinghua University, China

<sup>7</sup>Georgia Institute of Technology (Georgia Tech), USA

**Abstract**—We quantify the *collaboration stability of human mobility* and demonstrate its importance for networking applications. We derive a general model for collaboration stability and empirically explore the impact stability has on networking applications using a dataset of device-to-device encounters. We first derive insights on how characteristics of collaboration opportunities vary across everyday contexts and what implications this has on different multi-device scenarios. Among others, our results demonstrate that collaboration opportunities are highly dependent on the context where they take place, with diurnal patterns and spatial characteristics being particularly important. We also demonstrate the practical benefits of collaboration stability by demonstrating how it can be used to improve the selection of collaborators for sensing and computing applications.

**Index Terms**—Device-to-Device; Proximity Detection; Opportunistic Collaboration; Mobile Crowd; Edge Intelligence.

©2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

## I. INTRODUCTION

*Opportunistic collaboration* harnesses spontaneous interactions between users' devices to achieve tasks that would be infeasible or even impossible to carry out otherwise. Examples range from augmenting the capabilities of individual devices [1], [2] to sharing sensor data [3], [4] or other resources [5], [6]. In many applications collaboration is not only beneficial but also essential for accomplishing the task. For example, edge computing scenarios involving high velocity processing (e.g., real-time video) require joint efforts to offset latency requirements [1], [7] and some sensing tasks can only be accomplished when data from multiple devices are available [4], [8].

*Stability* of a system is generally defined as the property of preserving a state over time. In the context of collaboration, stability implies that the characteristics governing the execution of the shared task should remain unchanged over

time. A lack of stability can result in collaborative applications or services failing as they cannot harness sufficient pool or resources, or as the overhead of managing the participants exceeds the benefits it can bring. For example, two smartphones collaborating with each other can reach similar performance as the cloud in object detection [1], but fail to process the stream in time when there is no collaboration. Similarly, the collaborative sensing platform CoMon requires approximately 6 minutes of active collaboration for every hour spent looking for collaborators to offset the costs of collaborator discovery [3]. As we will show, this can only be achieved if the usage patterns of the application match the contexts where sufficient amounts of collaborators are available.

Despite the importance of stability for collaborative scenarios, the attributes that characterize and determine the available collaboration opportunities have thus far been understudied. Indeed, existing works have focused on intermittent and opportunistic collaborations without examining the effect different contexts or the number of simultaneous collaborators have on the collaboration opportunities. Similarly, existing works examining collaboration patterns have only considered whether a *single* additional collaborator can be found [3], [9] neglecting the potential of a larger pool of collaborators and the impact that human behavior has on collaboration opportunities. This lack of understanding makes it difficult to establish conditions under which specific types of participatory applications are likely to succeed or fail.

This paper contributes by *quantifying the collaboration stability of human mobility* and characterizing its impact on the capability to find a reliable set of collaborators for multi-device computing and networking scenarios. We derive a general model of stability that accounts for the (i) number of devices working together, (ii) the duration they need to stay together, (iii) the spatial context of the user, and (iv) the type of application that harnesses the resources from others. These facets are motivated by practical considerations as the required number of collaborators, the duration they need to remain in close proximity, and the distance that is used to define proximity depend on the characteristics of the application.

For example, edge computing scenarios require devices to be connected to the same network access point for the duration of the provisioned task [10]. Similarly, collaborative sensing requires shared interconnected sessions to last for several sensor cycles to prevent assets management overhead to offset its benefits [3], [4]. Our proposed approach accounts for these facets, offering a robust model able to quantify stability requirements for multi-device applications and to understand the conditions where different types of collaborative scenarios can succeed or where they are likely to fail.

To demonstrate the practical benefits of collaboration stability, we conduct thorough empirical experiments on a device-to-device (D2D) dataset captured by a cellular operator. First, we quantify and characterize the stability of collaboration opportunities across different everyday contexts. Secondly, we show how spatial and temporal characteristics of collaboration opportunities are highly context-dependent, suggesting the need to match the usage contexts for collaboration with those that the application is designed for. Indeed, unless these situational patterns are accounted for, collaborative applications are prone to fail as the required opportunities are not available. We also demonstrate the importance of stability for collaborative sensing applications by offering a way to substantially improve the selection of collaborators. Finally, we discuss the implications of our results for different types of applications.

## II. COLLABORATION STABILITY

Stability of collaboration can be understood as the persistence of any collaborative or cooperative encounter pattern. In the context of networking and mobile computing, collaboration stability refers to the capability of devices to execute a shared task over an extended period of time. The tasks are typically coordinated over short range communication technology and thus are dependent on the physical proximity of devices [3], [4]. Hence, in these applications, collaboration stability is inherently dependent on human mobility patterns as they determine where users spend time and for how long. According to the current scientific understanding, humans tend to spend the majority of their time at a small number of locations with the duration of their stay being inversely proportional to the personal importance of such location [11]. This implies that collaboration stability is likely to depend on the spatial and temporal characteristics of a location. For example, residential areas tend to have stable and long duration collaboration opportunities during evenings and nights, but not necessarily during the day. In contrast, encounters at a railway station or at large public spaces are likely to have shorter duration and higher variability. From a computing perspective, the requirements for stability are also highly dependent on the application. For example, in routing, having two devices collaborate for a short period of time is sufficient for ensuring good performance, whereas real-time object detection requires a larger number of devices to collaborate until all data has been processed and the task completes.

We define collaboration stability formally for an application  $A$  as the function  $f_A(d, t, s)$  where  $d$  is the number of collaborators (devices),  $t$  is the temporal context for collaboration, and  $s$  is the spatial context. The exact forms of these variables depend on the available data, e.g.,  $t$  could represent hours or minutes, whereas  $s$  could represent a certain radius or a specific point of interest. A collaborative application can then succeed if the overall stability, given by the function  $f$ , of a set of collaborators is sufficiently high in all the relevant spatial and temporal contexts. The function  $f$  can take different forms. For example, in our experiments we consider first a binary indicator function that is subject to constraints, i.e., thresholds on maximum distance and minimum time, and then combine it with a form of entropy for assessing stability. The proposed model extends existing approaches, which specify temporal constraints or execution bounds for tasks [3], [12], by accounting for the number of collaborators and the contexts where collaboration takes place.

## III. APPLICATIONS FOR COLLABORATION STABILITY

Collaboration stability can be highly beneficial to a wide range of multi-computing applications. Below we briefly discuss representative examples of applications that benefit from the capability to quantify collaboration stability.

**Contact tracing** monitors encounters between individuals with the goal of discovering potential infections in a population. The idea is to identify individuals that have been in close proximity to an infected individual and to estimate their risk of infection from the proximity characteristics of the encounters. In this case, the presence of a stable collaboration thus poses a risk for the individuals whenever one of the persons is infected. Quantifying collaboration stability is essential for understanding the overall risks of infection, and it can also be used to support new functionality. For example, collaboration stability provides a way to quantify which spaces meet a given regulation policy, e.g., restaurants can operate at 25% capacity.

**Fog computing** is a paradigm that assumes computing resources are located in the environment so that collected sensor data can be analyzed close to its source. Smart devices working in collaboration have the potential to perform large amounts of computations over big data streams [1] and thus can be used to dynamically establish a fog-like structure. Collaboration stability is essential for understanding the extent to which smart devices can contribute to fog provisioning. Indeed, as human mobility is governed by context, it is difficult to estimate the duration that the resources of a device can be used in a fixed location to support a fog-like structure [8]. By understanding collaboration stability of human mobility, it is possible to obtain quantifiable estimates that can be used to aggregate tasks and schedule them optimally to avoid disruptions as the device composition changes. This type of fog constellation can also complement existing edge computing deployments by elastically augmenting the resources based on the availability of individuals.

**Autonomous vehicles** can benefit from stable collaboration between the vehicles. For example, the network can be sliced so that vehicles moving together share networking resources. As an example, delivery drones that use the same distribution routes could share the load of executing services, storing and caching data. This can potentially enable vehicles with seamless connectivity to automate tasks such as data transfer, traffic control, smart parking and collision avoidance. Understanding collaboration stability also offers better capacity for planning and enables more resilient approaches to be integrated into existing communication network architectures to provision services to users. For instance, estimating the amount of autonomous vehicles required to deliver a (distributed and replicated) service to a mobile crowd of users.

**Crowd computing** collects and analyses data by aggregating contributions and resources from a set of users. This can improve task performance and augment the spatial coverage of the collected data. Assigning tasks to individuals is challenging and the optimal assignment is dependent on the user’s context. For example, users may be more willing to take on complex tasks when they are static for a long time rather than when they are moving. Collaboration stability can improve task scheduling by identifying better suited contexts for execution and by finding the best-suited candidates from a pool of individuals to carry out the task.

#### IV. DATASET AND EXPERIMENTAL SETUP

##### A. Dataset

We conduct our experiments using a dataset that contains real world mobility traces collected from a cellular operator in Shanghai over a one-week period. Human mobility has been shown to have strong daily and weekly patterns and hence covering a full week captures the most important behaviours in human mobility [13]. The use of a data collected through a cellular operator ensures that the data has extensive coverage and is representative of both the actual human mobility and the encounter patterns one would expect in a real-world deployment of a multi-device application. The data contains session information from users connecting to base stations resulting from calls, messaging, and data transfer activity linked with mobile applications. Each sample includes device identifier, the start and end time of a session, the amount of data transferred during the session (in bytes), the identifier of the base station that handles the connection, and the GPS coordinates of the base station. In total, the dataset contains data from 137 495 devices and 10 363 base stations. For privacy reasons the identifiers of devices and base stations are anonymized. Our previous work has shown the data to closely mirror the characteristics and overall distribution of human proximity detected by short-range communication technology [8].

##### B. Processing

We filter the dataset to consider only measurements where devices are in a proximal communication range through a base station. Devices are co-located and available for collaboration

when they are simultaneously connected to the same base station. We identify such devices by analysing the session information in the data, and we estimate the encounter duration of these devices by examining the overlap in session times. Prior to the analysis, we discard sessions that have very short or very long duration to ensure the results are representative of common encounter patterns rather than biased by infrequently occurring patterns. The pruning is done by estimating the distribution of encounter duration and discarding all sessions below the 10th percentile or above the 90th percentile. As our goal is to support long and stable collaborators, we analyze collaborations that can be harnessed during a time period of at least five minutes, which has been previously defined as the minimum time for a sensing task to be beneficial in a collaboration between two devices [3]. We also discard base stations and users whose total number of measurements is below the 10th percentile. We then perform temporal alignment of all sessions per device using hourly bin intervals and match the sessions of devices that share the same interval. We focus our analysis on the 20 km<sup>2</sup> area with the highest density of users and the 1000 users with the highest amount of samples. After these steps, the final dataset contains approximately 1 000 000 encounters.

We focus on six areas that differ in terms of user concentration and land use (Residential, Pubs, Park, Shopping, Train Station and Financial). The areas were selected by overlaying publicly available satellite imagery (through Google Maps) to the selected area. Land use was verified using OpenStreetMap and public information made available by the City of Shanghai. We chose these areas as they are representative examples of locations that people encounter as part of everyday routines on a daily basis [14]. For example, a financial area exposes working hours’ patterns, a train station describes users’ transportation mobility, and residential their habitual housing patterns. Each location has an area of 500 meters × 500 meters, which is sufficient to cover the area and to avoid overlap with other locations. The size of the area is fixed to ensure locations can be compared in a meaningful way. Lastly, to ensure the characterization of the selected locations is correct, we repeat the same steps for additional locations of the same type that are available in the dataset.

##### C. Experiments

We use the aforementioned data to carry out two experiments. First, we perform experiments to quantify collaboration stability and analyze the factors that influence it. Second, we conduct experiments on collaborator selection using collaborative sensing as a representative example of an application that would benefit from the quantification of collaboration stability. The experiments we conduct are briefly described below.

**Quantification:** The first experiment analyses collaboration stability present in human mobility by monitoring continuous proximity between devices. We quantify the stability of the different areas described above. Since the mobility of users varies between weekdays and weekend, we separately estimate collaboration stability for both cases. We achieve this by

measuring the number of devices available as a group and estimating the churn of the group over time using a retention rate metric [15]. A group depicts devices that are in proximal communication range for a period of time. We then compare the stability of different areas using daily patterns of users. Lastly, we evaluate whether stability values can be used as unique properties to predict the type of an area.

**Collaborator Selection:** The second experiment focuses on demonstrating how collaboration stability can improve the selection of collaborators. We compare two methods. *Familiarity* refers to the frequency of encounters between devices and is based on the collaborator selection algorithm of the CoMon collaborative monitoring platform [3]. Two devices that encounter each other several times are very familiar whereas two devices that encounter a few times are less familiar. *Regularity* uses the collaborator selection algorithm of the COSINE framework proposed in our previous work [8]. Regularity is based on collaboration stability and recommends collaborators based on Markov trajectory entropy by finding those collaborators that most consistently stay within each other’s proximity for the longest possible time. Both methods use a threshold parameter  $\alpha$  to determine collaborators. In the case of familiarity, the threshold is applied on the frequency of encounters, whereas in the case of regularity the threshold is applied on the entropy of the encounters.

## V. RESULTS

We next quantify collaboration stability and show it is highly dependent on context with both the type of location and the time-of-day affecting it. We also show how collaboration stability can help improve the selection of collaborators.

**Stability of contexts:** We first investigate how the number of devices that are available for collaboration varies over contexts. We consider the six locations described in the previous section, and calculate the number of devices separately for weekdays and weekends. Figure 1 shows the results. The values along the x-axis correspond to the hours of day, the y-axis indicates the duration of proximity, and the values along the z-axis correspond to the number of devices that are available. As expected, during weekdays the amount of devices is higher than on weekends for all of the areas, reflecting higher device use and user activity patterns. Nevertheless, it is still possible to find collaboration groups during the weekend, but their stability is generally lower. Such groups generally can only support tasks with short duration as otherwise overhead from group management might offset benefits from collaboration.

The number of devices that are available for collaboration also depends on the time of the day. For example, when sorting the locations based on the number of devices that are available on weekdays (Residential, Pubs, Park, Shopping, Train Station, Financial) and on weekends (Residential, Pubs, Train Station, Shopping, Park, Financial), the order of the locations changes. While some locations retain their position, the number of available devices changes depending on the hour of the day. As an example, for pubs device availability is highest during the

evening on weekends and during the afternoon on weekdays. Similarly, for train station, on weekdays there are many devices available in the morning, but during the weekend device availability changes. This indicates that collaboration stability in a location is dynamic and changes during the day. Collaborative applications need to account for these changes to ensure they can benefit from collaboration and be successful.

**Quantifying stability:** We next compare and quantify collaboration stability in the different areas. We focus on the Financial, Residential and Shopping areas as the amount of available devices at these locations varies more than at the other location. To account for oscillations in device availability, we model stability based on mobility patterns that are indicative of users’ routines [14]. We consider three common daily routines to demonstrate the dynamics of stability in different contexts: the first routine distinguishes between Daytime, and Nighttime; the second between Rest, Work and Leisure; and the last considers seven activities that are part of a generic weekday schedule. Table I shows the results.

We observe that collaboration stability changes significantly depending on the granularity of the daily routines. For instance, the first routine (Daytime) shows that 117 devices can be harnessed as a group during 47 minutes in a Financial context (8:00-20:59). However, a Financial location mostly has activity during working hours (9:00-16:59) as shown by the second routine (Work). Indeed, during these times the amount of devices that can be harnessed changes to 161 and the duration increases to up to 53 minutes. This also correlates with the categories in the third routine (Work 08:00-11:59, Lunch 12:00-14:59 and Work 15:00-17:59) that depict activities with high congregations of people. In contrast, if we consider the other times in the second routine (Rest and Leisure), the amount of available devices is reduced significantly to 16 and 21, respectively and the duration of collaboration decays considerably. For rest the duration decreases to 18 minutes and for leisure to 27 minutes. Residential provides the most consistent collaboration stability with average values of 281 devices to be harnessed with an average duration of 55 minutes, followed by Shopping with a stability of 185 devices with an average duration time of 49 minutes, and Financial with a stability of 66 devices that can be harnessed during 33 minutes on average.

**Stability predicts context:** Based on our findings, we proceed to determine whether stability is a factor that can be used to predict the context, and specifically the type of location. To achieve this, we analyze further other locations that are similar to the ones initially selected for the experiment. In other words, per each location described previously, we augmented our analysis with additional four location from the same category, e.g., four additional parks, four additional train stations and so on. We then develop prediction models using SVM (Support Vector Machine), RR (Random Forest) and LR (Logistic Regression) to evaluate whether collaboration stability can be used to infer the type of a location. The classification performance varies from 51% (random forest) to 70% (SVM).

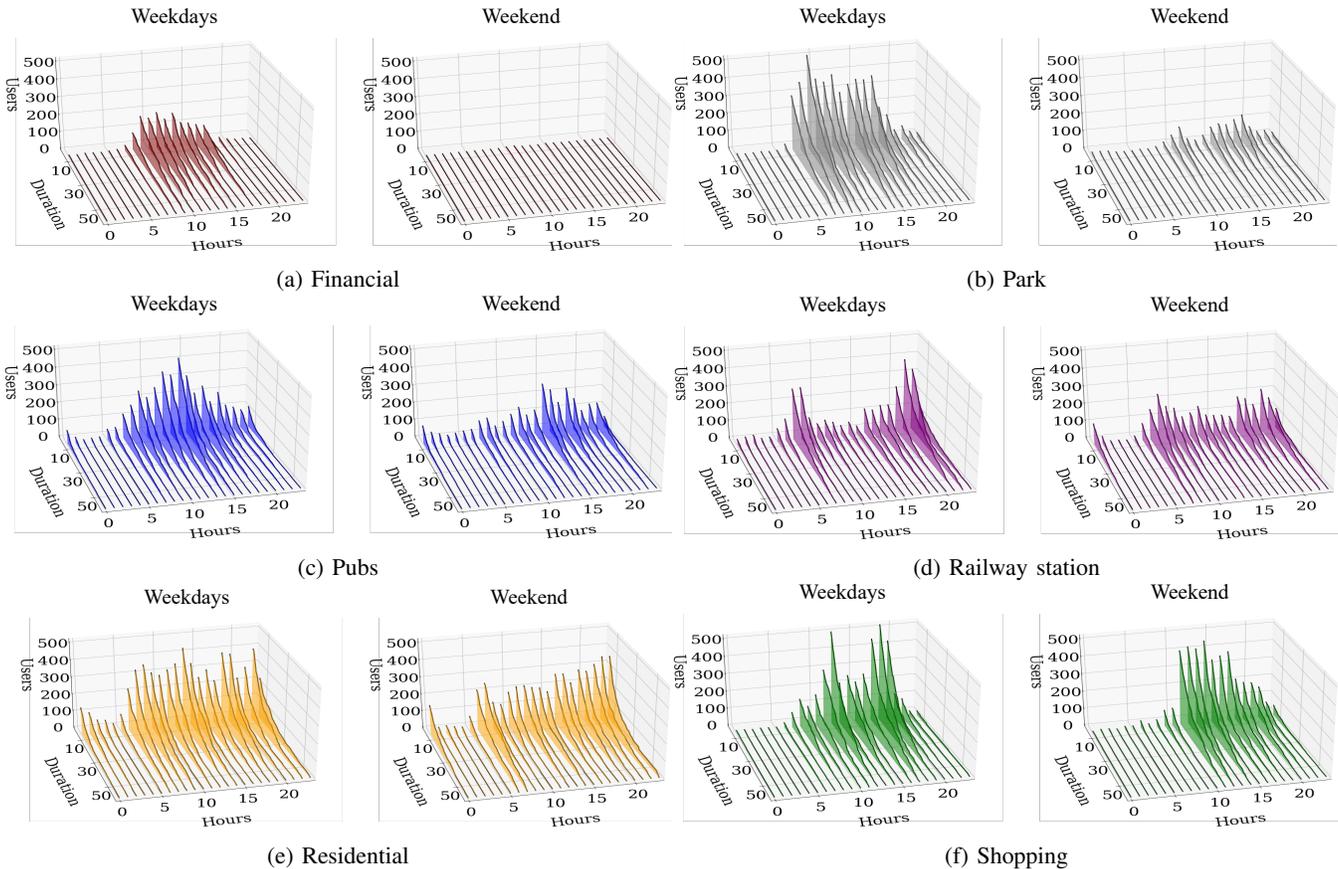


Fig. 1: Decline of available users for D2D collaborations over time in different contexts.

TABLE I: Stability for different contexts, duration time and number of available devices based on different daily routines.

| Routine      | Category<br>(time window)   | Time<br>range | Financial |           | Residential |            | Shopping  |            |
|--------------|-----------------------------|---------------|-----------|-----------|-------------|------------|-----------|------------|
|              |                             |               | Duration  | Devices   | Duration    | Devices    | Duration  | Devices    |
| 1            | Daytime                     | 08:00-20:59   | 47        | 117       | 56          | 339        | 57        | 299        |
|              | Nighttime                   | 21:00-07:59   | 14        | 5         | 53          | 192        | 36        | 28         |
| 2            | Rest (Early Morning)        | 01:00-08:59   | 18        | 16        | 51          | 170        | 32        | 35         |
|              | Work (Morning-Afternoon)    | 09:00-16:59   | 53        | 161       | 56          | 342        | 57        | 267        |
|              | Leisure (Afternoon-Evening) | 17:00-00:59   | 27        | 21        | 57          | 322        | 55        | 250        |
| 3            | Rest (Early Morning)        | 01:00-05:59   | 5         | 2         | 46          | 82         | 17        | 7          |
|              | Rush hours (Morning)        | 06:00-07:59   | 34        | 18        | 59          | 297        | 59        | 56         |
|              | Work (Morning)              | 08:00-11:59   | 55        | 166       | 54          | 336        | 57        | 190        |
|              | Lunch break                 | 12:00-14:59   | 55        | 161       | 57          | 393        | 56        | 344        |
|              | Work (Afternoon)            | 15:00-17:59   | 49        | 112       | 57          | 279        | 57        | 323        |
|              | Rush hours (Evening)        | 18:00-20:59   | 27        | 14        | 55          | 350        | 57        | 377        |
|              | Leisure (Evening)           | 21:00-00:59   | 15        | 2         | 58          | 279        | 48        | 41         |
| <b>Total</b> |                             |               | <b>33</b> | <b>66</b> | <b>55</b>   | <b>282</b> | <b>49</b> | <b>185</b> |

Collaboration stability thus not only depends on the location, but is a characteristic of the location. Applications that are mostly used at specific contexts thus need to ensure their expected collaboration patterns align with the patterns that characterize the context as otherwise the prerequisites for collaboration are not met.

**Selection of collaborators:** We next focus on the practical benefits of collaboration stability and demonstrate how it can be used to improve the selection of collaborators in multi-device sensing and computing tasks. As described in

Section IV-C, we compare two methods, a familiarity-based collaborator selector and a regularity-based method that uses entropy to quantify stability and to identify collaborators.

Figure 2a plots the cumulative distribution function (CDF) of collaboration opportunities estimated using familiarity. The distribution closely mirrors earlier results [3], suggesting that there are collaboration opportunities in the environments captured by the dataset. As familiarity increases, so does the duration of collaboration opportunities. The likelihood of finding devices to collaborate is high when familiarity requirements

are low ( $\alpha = 0$ ) but decreases sharply as requirements increase (from  $\alpha = 25$  to  $\alpha = 450$ ). Figure 2b further characterizes collaboration opportunities by looking at how the number of available devices and the duration of collaboration varies when familiarity between devices changes. We can observe that the available collaboration opportunities depend on the application requirements. For example, when familiarity requirements are low ( $\alpha = 25$ ), up to 1000 devices can be harnessed during 20 minutes with a likelihood of 0.4. However, when familiarity requirements are high ( $\alpha = 450$ ), only 2 devices can be harnessed during 20 minutes with the same probability. Thus applications that assume high familiarity can only benefit from short collaborations whereas applications allowing interactions with less familiar devices can benefit from more opportunities, even if they are shorter in duration.

When regularity is used to select collaborators, the duration that the collaborators are availability is more stable, as shown in Figure 2c. This is also supported by Figure 2d, which quantifies the amount of devices available for collaboration. We observe that  $\alpha = 3$  provides the optimal results for selecting collaborators with the longest duration of co-proximity. By using this strategy, collaborators are available for up to 22 minutes. In contrast, while the number of available devices remains almost the same for other  $\alpha$  values, the duration that collaborators are available decreases significantly to between 17 and 19 minutes. In practice these results mean that collaborator selection using devices that have similar duration of co-proximity (i.e., higher collaboration stability) can improve system stability and result in better selection of collaborators than what the familiarity-based selection offers. While regularity provides better collaborator selection than familiarity, it only offers selection of a single collaborator and does not offer insights about how to select *groups of devices* to collaborate with. Generally the amount of devices that have the same spatial and temporal characteristics is unknown and hence it is also essential to quantify the collaboration stability at different contexts to have an understanding of how the availability of devices varies and how it can be harnessed optimally.

**Implications:** The results demonstrate that the availability of devices oscillates over time and locations. Understanding changes in the number of available devices, and hence on stability, is critical for deciding how to optimally schedule collaborative tasks and applications may succeed or fail in similar situations at different times. For example, applications that are mostly used during weekends and that rely on long-term collaboration are likely to fail whereas applications that adapt the expected benefits of the tasks over time are more likely to succeed. Applications that are mostly used in specific locations need to make sure they are used at times-of-day where the expected collaboration patterns are high. For example, applications that are used mostly within a financial district can succeed during the day on weekdays, but would be likely to fail during weekends as there would not be sufficiently many devices to form a stable collaboration group. The suitability of

tasks for different locations also varies, e.g., train stations do not support tasks requiring stable groups with a long duration. As such, train stations are better suited for collaborative sensing rather than for collaborative computing tasks. Our results shed light on how these collaboration patterns vary over time, and how the changes in stability need to be accounted to optimally harness collaboration.

The analysis also highlights how characteristics of *individual collaborations* largely have remained the same over time – which is to be expected due to characteristics of human mobility [16]. However, it is increasingly possible to find *several* devices to collaborate with, even if the stability of these groups varies depending on task duration and group size requirements. Existing collaborator identification methods have been designed largely to find (individual) collaborators, rather than look at how to harness the potential of devices that are available in a given location or context. Harnessing collaboration capabilities of locations requires novel methods that can tailor the selection to the context where they operate.

## VI. DISCUSSION

**Applications and Stakeholders:** Collaboration stability offers applications a possibility to adapt and improve execution performance by avoiding over-provisioning of resources in capacity planning. For instance, device-to-device discovery and monitoring can be optimized so that discovery routines are triggered based on context. This can help applications such as digital contact tracing and device-to-device networking. Similarly, stability can be also exploited to manage and share the resources of devices in a more sophisticated way. For instance, the quality of microphones, cameras and environmental sensors can vary considerably across devices and collaborator selection can be extended to consider also the quality of the sensors. Collaboration stability can also encourage application providers to equip applications with functionality that requires resource intensive processing, especially if the expected collaboration patterns match with the needs of the processing. For instance, edge intelligence and distributed machine learning models can be adjusted depending on stability with more powerful models preserved for situations where more devices are available for longer periods.

**Selection of collaborators in a group:** Collaboration groups can consist of a massive number of devices, but the optimal formation depends on the characteristics of the task to be executed. Devices need mechanisms to identify the best possible collaborators for a task, taking into consideration the context of the device and the collaborators (e.g., encounter frequency, group size, remaining battery). In some applications it may be possible to harness all available devices, such as in edge provisioning, but even in that case devices need to agree on who is coordinating the execution of tasks.

**Harnessing stability:** We demonstrated that collaboration stability enables forming device groups for a specific duration. Thus, groups can be harnessed as a whole to execute tasks that are common. While we demonstrate the potential of groups

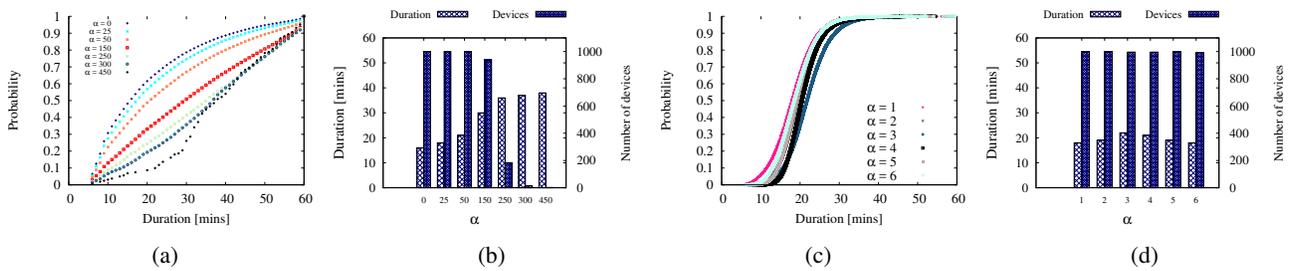


Fig. 2: Selection of collaborators using CoMon, a) CDF, b) Available collaborators; Selection of collaborators using COSINE, c) CDF, d) Available collaborators.

by quantifying them in terms of the number of devices and the time they stay together in different user contexts, the way device groups are harnessed depends on the architecture in which devices are arranged. For instance, in a collaborative architecture, the execution of a task is scheduled among multiple devices, such that one device at a time executes the task, but all devices leverage the result of that execution. In contrast, in a centralized architecture, a device can contribute to the execution of a task even if it does not benefit its own constrained resources, e.g., through federated learning. In addition, devices can be arranged into hierarchical formations inside the group to reduce communication effort of end points, disseminate information with ease and assign optimal roles to devices.

**Future Directions:** The experiments considered one week of data from a single mega-city, Shanghai, and characterized collaboration stability within the city. Larger time-spans or other areas with differing population density patterns may have different results and examining stability at these locations can offer further insights of collaboration stability in a richer spectrum and diversity of contexts, e.g., considering cultural and ethnicity differences. Nevertheless, mobility characteristics contain many universal characteristics [17] and we would expect many of the findings to be similar across different locations or time-spans. While we considered different types of locations and temporal contexts, we did not track specific devices during our study. As a result, our results do not account for influence of transportation or other mobility variations. Group mobility is very important as scenarios in which a group of devices traveling together can be stable enough to execute collaborative tasks on the move, e.g., mobile edge computing.

## VII. SUMMARY AND CONCLUSIONS

We contributed by quantifying and modeling the collaboration stability of devices for computing scenarios involving collaboration between multiple devices. We found that different stable groups of devices sharing same spatial and temporal properties can be formed in different contexts, and the stability of each group is quantifiable in terms of duration and the amount of devices. The results indicate that stability is very sensitive to the context, and can change significantly depending on the time of the day. For instance, for a financial context, the amount of available collaborators can increase up

to 166 devices that can be harnessed during a 53 minute period, but it also can decay drastically to just 2 devices available during a 5 minute period. Naturally, the time window that is considered to estimate stability influences the stability results. We also found that in some contexts, such as in residential areas, stability is almost constant during the day. This suggests that by understanding the stability of devices in a context, it is then possible to schedule the execution of a task based on its computational complexity and processing requirements. We also demonstrated how collaboration stability can support the selection of collaborators and reduce group management overhead in collaborative computing and sensing applications.

## VIII. ACKNOWLEDGMENTS

This research is supported by the European Regional Funds through the IT Academy Programme and the Academy of Finland grants 317875, 319669, 325570 and 339614. The publication only reflects the authors' views.

## REFERENCES

- [1] E. Lagerspetz, X. Li, J. Hamberg, H. Flores, P. Nurmi, N. Davis, and S. Helal, "Pervasive data science on the edge," *IEEE Pervasive Computing*, 2019. [Online]. Available: <https://doi.org/10.1109/MPRV.2019.2925600>
- [2] K. Habak, M. Ammar, K. A. Harras, and E. Zegura, "Femto clouds: Leveraging mobile devices to provide cloud service at the edge," in *IEEE CLOUD '15*. IEEE, 2015, pp. 9–16. [Online]. Available: <https://doi.org/10.1109/CLOUD.2015.12>
- [3] Y. Lee, Y. Ju, C. Min, S. Kang, I. Hwang, and J. Song, "Comon: Cooperative ambience monitoring platform with continuity and benefit awareness," in *ACM MobiSys 12'*, UK, June 25–29 2012. [Online]. Available: <https://doi.org/10.1145/2307636.2307641>
- [4] J. Leppänen, M. Pelkonen, H. Guo, S. Hemminki, P. Nurmi, and S. Tarkoma, "Collaborative and energy-efficient speech monitoring on smart devices," *Computer*, vol. 49, no. 12, pp. 22–30, Dec 2016. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/MC.2016.380>
- [5] D. Chatzopoulos, C. Bermejo, E. ul Haq, Y. Li, and P. Hui, "D2d task offloading: A dataset-based q&a," *IEEE Communications Magazine*, vol. 57, no. 2, pp. 102–107, 2018. [Online]. Available: <https://doi.org/10.1109/MCOM.2018.1700873>
- [6] A. Amiri Sani, K. Boos, M. H. Yun, and L. Zhong, "Rio: a system solution for sharing i/o between mobile systems," in *ACM MobiSys 14'*. ACM, 2014, pp. 259–272. [Online]. Available: <https://doi.org/10.1145/2594368.2602434>
- [7] M. Satyanarayanan and N. Davies, "Augmenting cognition through edge computing," *Computer*, vol. 52, no. 7, pp. 37–46, 2019. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/MC.2019.2911878>

- [8] H. Flores, A. Zuniga, F. Faghihi, X. Li, S. Hemminki, S. Tarkoma, P. Hui, and P. Nurmi, "Cosine: Collaborator selector for cooperative multi-device sensing and computing," in *IEEE PerCom '20*. IEEE, 2020. [Online]. Available: <http://doi.org/10.1109/PerCom45495.2020.9127364>
- [9] J. Scott, J. Crowcroft, P. Hui, and C. Diot, "Haggle: A networking architecture designed around mobile users," in *WONS '06*, 2006, pp. 78–86.
- [10] T. X. Tran, A. Hajisami, P. Pandey, and D. Pompili, "Collaborative mobile edge computing in 5g networks: New paradigms, scenarios, and challenges," *IEEE Communications Magazine*, vol. 55, no. 4, pp. 54–61, 2017. [Online]. Available: <https://doi.org/10.1109/MCOM.2017.1600863>
- [11] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, no. 7196, pp. 779–782, 2008. [Online]. Available: <https://doi.org/10.1038/nature06958>
- [12] A. Hekmati, P. Teymouri, T. D. Todd, D. Zhao, and G. Karakostas, "Optimal mobile computation offloading with hard deadline constraints," *IEEE Transactions on Mobile Computing*, vol. 19, no. 9, pp. 2160–2173, 2019. [Online]. Available: <https://doi.org/10.1109/TMC.2019.2920819>
- [13] A. Sevtsuk and C. Ratti, "Does urban mobility have a daily routine? learning from the aggregate data of mobile networks," *Journal of Urban Technology*, vol. 17, no. 1, pp. 41–60, 2010. [Online]. Available: <https://doi.org/10.1080/10630731003597322>
- [14] F. Xia, J. Wang, X. Kong, Z. Wang, J. Li, and C. Liu, "Exploring human mobility patterns in urban scenarios: A trajectory data perspective," *IEEE Communications Magazine*, vol. 56, no. 3, pp. 142–149, 2018. [Online]. Available: <https://doi.org/10.1109/MCOM.2018.1700242>
- [15] A. Zuniga, H. Flores, E. Lagerspetz, P. Nurmi, S. Tarkoma, P. Hui, and J. Manner, "Tortoise or hare? quantifying the effects of performance on mobile app retention," in *The World Wide Web Conference*. ACM, 2019, pp. 2517–2528. [Online]. Available: <https://doi.org/10.1145/3308558.3313428>
- [16] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human mobility," *Science*, vol. 327, no. 5968, pp. 1018–1021, 2010. [Online]. Available: <https://doi.org/10.1126/science.1177170>
- [17] M. Schläpfer, M. Szell, H. Salat, C. Ratti, and G. B. West, "The hidden universality of movement in cities," *arXiv preprint arXiv:2002.06070*, 2020.

**Huber Flores** is an Associate Professor at the University of Tartu. He completed his Ph.D. in Computer Science at the University of Tartu, Estonia in 2015. His research interests include distributed, mobile and pervasive computing systems. Contact him at [huber.flores@ut.ee](mailto:huber.flores@ut.ee)

**Agustin Zuniga** is a Doctoral candidate at the University of Helsinki. He completed his MSc. in Computer Science at the University of Helsinki in 2018. His current research focuses on pervasive data science, machine learning and Internet of Things. Contact him at [agustin.zuniga@helsinki.fi](mailto:agustin.zuniga@helsinki.fi)

**Leonardo Tonetto** is a Doctoral candidate at the Technical University of Munich. His research interests include human mobility sensing and modeling, and machine learning. Contact him at [tonetto@in.tum.de](mailto:tonetto@in.tum.de)

**Tristan Braud** is an Assistant Professor at HKUST. He obtained his Ph.D. in Computer Science from University Grenoble Alpes, France in 2016. His research interests include, augmented reality, pervasive computing and device to device communications. Contact him at [tbraud@cse.ust.hk](mailto:tbraud@cse.ust.hk).

**Pan Hui** is a Professor of Computational Media and Arts at the Hong Kong University of Science and Technology and the Nokia Chair in Data Science at the University of Helsinki. He obtained his PhD from University of Cambridge. His research interests include opportunistic networks, augmented reality and mobile computing. Contact him at [panhui@ust.hk](mailto:panhui@ust.hk).

**Yong Li** is an Associate Professor with Tsinghua University. He received the Ph.D. degree in electronic engineering from Tsinghua University in 2012. His research interests include machine learning and big data mining. Contact him at [liyong07@tsinghua.edu.cn](mailto:liyong07@tsinghua.edu.cn).

**Sasu Tarkoma** is a Full Professor at the Department of Computer Science, University of Helsinki. He completed his Ph.D. in Computer Science at the University of Helsinki in 2006. His research interests include mobile computing, Internet technologies, and AI. Contact him at [sasu.tarkoma@helsinki.fi](mailto:sasu.tarkoma@helsinki.fi)

**Mostafa Ammar** is a Regents' Professor at the Georgia Institute of Technology. He received his Ph.D. from the University of Waterloo, Ontario, Canada. His research interests are in network architectures, protocols and services. Contact him at [ammar@cc.gatech.edu](mailto:ammar@cc.gatech.edu)

**Petteri Nurmi** is an Associate Professor at the University of Helsinki. He completed his Ph.D. in Computer Science at the University of Helsinki in 2009. His research interests include distributed systems, pervasive data science, and sensing systems. Contact him at [petteri.nurmi@helsinki.fi](mailto:petteri.nurmi@helsinki.fi)