



HAL
open science

Leveraging CDR datasets for Context-Rich Performance Modeling of Large-Scale Mobile Pub/Sub Systems

Georgios Bouloukakis, Rachit Agarwal, Nikolaos Georgantas, Animesh Pathak, Valerie Issarny

► To cite this version:

Georgios Bouloukakis, Rachit Agarwal, Nikolaos Georgantas, Animesh Pathak, Valerie Issarny. Leveraging CDR datasets for Context-Rich Performance Modeling of Large-Scale Mobile Pub/Sub Systems. WiMob 2015 - 11th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, Oct 2015, Abu Dhabi, United Arab Emirates. hal-01204871

HAL Id: hal-01204871

<https://inria.hal.science/hal-01204871v1>

Submitted on 24 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Leveraging CDR datasets for Context-Rich Performance Modeling of Large-Scale Mobile Pub/Sub Systems

Georgios Bouloukakis, Rachit Agarwal, Nikolaos Georgantas, Animesh Pathak, Valérie Issarny
MiMove Team, Inria Paris - Rocquencourt, France
{firstname.lastname}@inria.fr

Abstract—Large-scale mobile environments are characterized by, among others, a large number of mobile users, intermittent connectivity and non-homogeneous arrival rate of data to the users, depending on the region’s context. Multiple application scenarios in major cities need to address the above situation for the creation of robust mobile systems. Towards this, it is fundamental to enable system designers to tune a communication infrastructure using various parameters depending on the specific context. In this paper, we take a first step towards enabling an *application platform for large-scale information management* relying on ‘mobile social crowd-sourcing’. To inform the stakeholders of expected loads and costs, we model a large-scale mobile pub/sub system as a queueing network. We introduce additional timing constraints such as (i) mobile user’s intermittent connectivity period; and (ii) data validity lifetime period (e.g. that of sensor data). Using our *MobileJINQS* simulator, we parameterize our model with realistic input loads derived from the *D4D* dataset (*CDR*) and varied lifetime periods in order to analyze the effect on response time. This work provides system designers with coarse grain design time information when setting realistic loads and time constraints.

Keywords—Publish/Subscribe, Context-Rich, Large-Scale, Call Detail Records, Queueing Networks

I. INTRODUCTION

Dependence on Internet connectivity is ever-increasing with the introduction of smart mobile devices. Recent figures predict Internet to reach 3.2 billion people by the end of year 2015 with 60% of this population residing in the developing countries [1]. Predictions also state that by the year 2020 more than 26 billion devices will be connected [2], thereby developing a necessity for creating robust mobile distributed communication systems that operate in a very large-scale environment and enable user participation. The Publish/Subscribe (pub/sub) interaction paradigm provides a loosely coupled form of interaction between *publishers* of the data and *subscribers* for the data via intermediate *brokers* [3]. Pub/sub interaction paradigm has recently observed immense growing utility in large-scale mobile distributed communication systems [4]. A large number of applications have been created and several standards have been adopted following the pub/sub interaction paradigm. These include, Web news through RSS feeds [5], [6], streaming services [7], application integration [8] and real time applications [9].

Regarding pub/sub systems operating in large-scale environments, it is fundamental to model the performance in order to tune them. For this purpose, one method is to

rely on queueing network models [10]. Based on this approach, system resources (e.g., brokers) and networks (e.g., underlying infrastructure) are represented as *queues* and the exchanged data (e.g., events) as *jobs* served at the queues. Subsequently, the common ways to evaluate the performance are: (i) using existing closed-form solutions and probability distributions [11]; and (ii) performing simulation by analyzing more complex and realistic systems [4]. Regarding the latter, the evaluation is enabled by utilizing datasets derived from real setups operating using actual data. Call Detail Records (*CDRs*) are one such dataset related to mobile networks and telephony traffic. Accordingly, *CDRs* are convenient for providing realistic workloads for mobile pub/sub systems. Additionally, *CDRs* reflect localized context and environmental conditions, since they contain spatio-temporal information about the phone interactions. The localized context relates to communication *i)* during festival period, *ii)* in certain region, *iii)* during traffic anomalies, and *iv)* during natural catastrophe, etc. Thus, depending on the application context, picking various subsets of the dataset based on spatio-temporal information leads to a *context-rich* performance modeling.

One pub/sub performance model can be customized to multiple application scenarios depending on the context. One such being the development of the application platform for large-scale transport information management that relies on *mobile social crowd-sourcing* [12], [13] where a user is able to *sense* and participate by reporting on the environment with respect to personal, social, or public context. For instance, a sensing interaction could be the data sent by a traveler reporting some incident. In terms of communication infrastructure support, the pub/sub interaction paradigm is suitable for large-scale mobile environments (characterized by a large number of mobile users, intermittent connectivity, non-homogeneous data arrival to the users, etc), since it provides loosely-coupled interaction.

In this paper, to evaluate the pub/sub interaction paradigm in a context-rich large-scale setting, our contributions are as follows: *i)* we introduce a queueing network model for the end-to-end interaction within a *large-scale mobile pub/sub system* which tackles parameters such as user’s intermittent connectivity period and data validity lifetime period; *ii)* we parameterize our model with context-rich information by leveraging the *D4D* dataset [14]. The *D4D* dataset is a *CDR* provided by Orange Labs for the country of Senegal; *iii)* we extensively analyze the *D4D* dataset in order to identify the data that we are interested in and infer primary results; *iv)* we develop a

simulator called *MobileJINQS*¹ that implements our model and uses the *D4D* dataset traces as realistic input workload to the pub/sub system model over the time span of a whole year; and v) we identify ways of tuning the system parameters in order to satisfy certain design requirements, by thoroughly evaluating the behavior of the pub/sub system through simulation-based experiments. To the best of our knowledge, our work is the first which is utilizing CDRs to evaluate the performance of mobile pub/sub systems.

The rest of the paper is structured as follows: in Section II, we introduce our pub/sub model. In Section III, we summarize our analysis of the *D4D* dataset, we introduce the parameterization of our pub/sub model based on the dataset, and our simulator. Then, we present, in Section IV, our simulation experiments and their outcomes. We finally complement this work with related work in Section V, and conclude, also discussing future work, in Section VI.

II. PUB/SUB PERFORMANCE MODEL

The long-term objective in this work is to enable large-scale sensing interactions that are part of mobile social crowdsourcing applications relying on pub/sub interaction paradigm. Such an application platform guarantees that the sensing data is processed and delivered to the corresponding mobile users *on-time*, despite the intermittent connectivity of the latter and the data validity lifetime period. In the following subsections, after a brief presentation of the pub/sub interaction paradigm, we introduce our model that encompasses the above concerns.

A. Pub/Sub interaction paradigm

In the pub/sub interaction paradigm, multiple peers either take the role of publisher or a subscriber and interact via multiple intermediate brokers. Publishers produce events characterized by a specific topic to the broker. Subscribers subscribe their interest for specific topics with the broker, who maintains an up-to-date list of subscriptions. The broker matches received events with subscriptions and delivers a copy of each event to each interested subscriber (decoupled interaction). In terms of space coupling, interacting peers do not need to know each other. Events are diffused to subscribers only based on topics (one-to-many interaction). In terms of time coupling, peers do not need to be present at the same time during the interaction. Subscribers may be disconnected at the time when the events are published to the broker, who then keeps the events in a dedicated buffer for each subscriber. Thus, subscribers receive the pending events when reconnected.

B. Large-scale mobile pub/sub system

A large-scale pub/sub system relies on a network of *brokers*. It may be the case that all brokers are accessible to publishers and subscribers, or that there is a distinction between *edge brokers* (brokers at the periphery of the network) and brokers situated in the network backbone (*core brokers*). In both cases, each broker propagates the subscriptions it receives to its neighbors; in this way, the routing tables of all brokers are properly populated.




Fig. 1: End-to-end interaction between publishers and subscribers through two system access points

In a large-scale *mobile* pub/sub system, publishers and subscribers are mobile entities. Interactions then rely on volatile network access to the nearest edge broker. Publishers connect to the system to publish events. A *lifetime* limit can be assigned to these events, depending on the requirements of our application platform. Subscribers connect occasionally to the system to receive new events, and disconnect to save energy. Similar to SMS in mobile networks, moving subscribers may hand-off between brokers. These results in the old and new broker updating their subscriptions and routing tables as well as transferring between them the events stored for the subscriber but not yet delivered.

C. End-to-end interaction model of a large-scale mobile pub/sub system

In our study, we focus on end-to-end interaction between publishers and subscribers going through two access points of the system, an *input access point* and an *output access point*. The *input access point* is provided by an edge broker. Geographically close mobile publishers generate input flow to the whole system through this access point. The *output access point* is also provided by an edge broker. Geographically close mobile subscribers receive output flow from the whole system through this access point.

The *input access point* is modeled as a queueing system where the service rate μ_{in} represents matching incoming events (input flow λ_{in}) to subscriptions and routing them to corresponding subscribers and/or brokers (part A of Fig. 1). This may involve replication of some events. This may also involve dropping of some events, either due to no corresponding subscription or due to expiration.

The *output access point* is modeled as a queueing system where the service rate μ_{out} represents the transmission of buffered events to the intermittently connecting subscribers previously determined by routing (output flow λ_{out}), as depicted in part C of Fig. 1. Here again, some events may expire before being delivered to the destination.

In this work, we focus only on the input and output processes. We do not consider, for the moment, the rest of the pub/sub system (i.e., multiple input and output access points, routing of events among multiple brokers) since we concentrate on the limited lifetime of data and the intermittent connectivity of subscribers. Additionally, we link directly the output of the first queueing system to the input of the second one, considering the path through the pub/sub system as a simple wire. This is equivalent to having a single centralized

¹<http://xsb.inria.fr/d4d#mobilejinqs>

broker for the whole pub/sub system focusing on the timing constraints. In the next section we illustrate how we can obtain the values of λ_{in} , λ_{out} , μ_{in} and μ_{out} for a specific context using a real dataset.

III. MODEL PARAMETERIZATION AND SIMULATION

In this section, we describe the analysis of the context-rich realistic data provided to us by Orange Labs in order to parameterize and feed our pub/sub model. We then discuss the parameterization and simulation of the model based on the *D4D* dataset [14]. This data is in the context of Senegal, where the use of Internet remains limited. Indicatively, recent reports show that 31% of SMEs in Senegal have Internet access, with only 7% having a website [15]. Additionally, despite the fact that most of the people in Senegal have mobile phones, only 26% of them have mobile Internet access [16]. The above percentage is almost equal to the total Internet access, as use of fixed Internet remains very low. Hence, for a large part of the population, the only alternative for data access is SMS.

A. D4D Dataset and its Analysis

Call Detail Records (**CDRs**) are highly anonymized telephone records provided by telecom operators, such as Orange and Sonatel, for research purposes. *CDR* dataset typically consist of millions of records that are logged over a large area and time duration. *CDRs* contain spatio-temporal information about the call and SMS, more specifically: anonymized user who initiated it, base station from where it was initiated, destination base station and duration of the call. It thus also provides *context-rich* information about a country. In this work, we refer the initiation of a call or the emission an SMS as the *connection* made to a base station.

The *D4D* dataset contains *CDRs* of the users that are subscribed to the Sonatel Telecom in Senegal. This data is collected over a period of 50 weeks from 7 January 2013 until 23 December 2013. Senegal is divided into *regions*. Each region in Senegal is further subdivided into *departments*. A department has many *sub-prefectures* and each sub-prefecture has multiple *base stations*. Going forward, we address these base stations as *antennas*. The *D4D* dataset consists of 3 sub-datasets *Dataset1*, *Dataset2*, and *Dataset3*. *Dataset1* logs, the number and total duration of calls made from an antenna during each hour to another antenna. Along with calls, it also logs total number of SMSs sent from an antenna during each hour to another antenna. *Dataset2*, for each 10 mins, logs the anonymized user who made the call or sent a SMS, the associated antenna from where the call or SMS originated. *dataset3*, for each 10 mins, logs the anonymized user who made the call or sent a SMS, the associated Sub-Prefecture from where the call or SMS originated. Thus, *Dataset1* differs with *Dataset2* and *Dataset3* in terms of temporal granularity while data set *Dataset2* and *Dataset3* differ with respect to spatial granularity of the data. Further *Dataset2* has records of almost 300,000 unique subscribers of Sonatel Service while *Dataset3* has records of almost 150,000 subscribers only. Given the fine spatio-temporal granularity of the *Dataset2* and the number of subscribers recorded, we chose it for our study.

Let A be the set of antennas, T be the set of unique time intervals over which the data is collected.




Fig. 2: Visualization of $N_i \forall i \in A$ over the time period $T = [7 \text{ Jan}, 20 \text{ Jan}]$.

Definition 1. A *trace* of an antenna $i \in A$ is the set of number of connections made to the antenna i over T . Let N_i denote the trace of an antenna $i \in A$ and $N_i^t \in N_i$ be the number of connections made from antenna $i \in A$ at a given time $t \in T$.

Definition 2. Let D denote the set of unique days of the year. Let M be a matrix of size $|A| \times 365$, where $M_i^d \in M$ represents the maximum number of connections made from an antenna $i \in A$ on a given day $d \in D$. Let T' be the set of all unique time intervals on a given $d \in D$, then

$$M_i^d = \max_{t \in T'}(N_i^t)$$

Note that $T \supset T'$.

In Fig. 2, we provide a visualization of $N_i \forall i \in A$ over the time period from Jan 07, 2013 00:00:00 until Jan 20, 2013 23:50:00 (Cf. Fig. 2). An animated version of the same is available on our website². Fig. 2 also shows less utilization of mobile network during the night hours, as expected. It should be noted that antenna IDs increase from west to east (Cf. Fig. 3).

Fig. 4 provides a visualization of M . From the figure we identify an interesting spatial context of the communication pattern in major cities of Senegal. We identify the set of antennas, $A' \subset A$, in Dakar, the capital of Senegal, and in the major cities in the west and north west of Senegal (north of Gambia) to account for 65.4% of the total communication happening in Senegal over T . This set mainly included antennas with ID in $[1, 580)$. However, this set also includes some antennas where connections were not tracked over the entire duration of the

²<http://xsb.inria.fr/d4d#visualization>




Fig. 3: Representation of Antenna IDs.

dataset. We then extract the number of connections made in a given time interval at country scale.

Definition 3. Let C denote the set of number of connections made in the country in time interval $t \in T$ such that $C^t \in C$, then

$$C^t = \sum_{\forall i \in A} N_i^t$$

From Def. 3 we extract $\max_{\forall t \in T}(C^t)$, the max number of connections in the whole system (country) at any given time interval $t \in T$. Fig. 5 shows C^t . It should be noted that $\max_{\forall t \in T}(C^t) = 112,937$ which is observed on Thu, 08 Aug 2013 from 23:00:00 to 23:10:00 GMT, i.e., on the end of Ramadan, a public holiday in Senegal. This provides us with temporal context of the communication pattern in Senegal.

B. Use of the dataset to parameterize the pub/sub model

For parameterizing our model, we rely on our *D4D* dataset analysis. More particularly, we are interested in A' and $N_i \forall i \in A'$, defined in Section III-A. We map the two types of access points of our pub/sub system model to two antennas selected from the set A' . We assume that each trace in A' can equally represent the reception of a connection (call or an SMS) through the antenna. For the input access point, we map the number of connections per 10 min interval at the selected antenna to an equal number of events published over the same time interval (Table I).

Definition 4. Let λ_{in} be the input process (or input flow) at the input access point associated to the antenna $i \in A$, and N_i^t be the number of connections to the antenna i in each interval $t \in T$, as defined in Def. 1. Then λ_{in} is a non-homogeneous




Fig. 4: Visualization of maximum number of connections made from an antenna $i \in A$ on a given day $d \in D$.

time	antenna ID	number of connections = number of events
...
20:50-21:00	161	173
21:00-21:10	161	175
21:10-21:20	161	204
...

TABLE I: Trace sample as realistic load to the input access point on 07 Jan 2013

Poisson process with rate parameter $\lambda(t)$ piecewise constant in each interval $t \in T$:

$$\lambda(t) = \frac{N_i^t}{|t|}$$

For the output access point, we map the number of connections per 10 min interval at the selected antenna to an equal number of events delivered to subscribers over the same time interval (Table II), provided that there are enough events in the queue. This relates to the number of subscribers that are connected to the access point during the specific interval, i.e., that are available to receive new events present in the queue, if any. To provide this effect, we model the service process at the output access point accordingly.

Definition 5. Let μ_{out} be the service process (or service rate) at the output access point associated to the antenna $j \in A$, and N_j^t the number of connections to the antenna j in each interval $t \in T$, as defined in Def. 1. Then μ_{out} is a non-homogeneous Poisson process with rate parameter $\mu(t)$ piecewise constant in each interval $t \in T$:

$$\mu(t) = \frac{N_j^t}{|t|}$$




Fig. 5: Visualization of number of connections made in the country in time interval $t \in T$.

time	antenna ID	number of connections = number of connected subscribers
...
20:50–21:00	24	115
21:00–21:10	24	90
21:10–21:20	24	90
...

TABLE II: Trace sample as realistic load to the output access point on 07 Jan 2013

This is equivalent to having service times of events that follow an exponential distribution with mean equal to $\frac{1}{\mu(t)}$ for the interval t .

The output λ_{out} of the queueing system modeling the input access point is a process similar to the input λ_{in} , provided that the queueing system does not saturate. Then, λ_{out} is fed to the input of the queueing system modeling the output access point. We also need to take into account the possible replication and dropping of events; this results in $\lambda_{out} \neq \lambda_{in}$. Without loss of generality, we consider that in our setup there is no replication of events or dropping due to the absence of corresponding subscription; however, events may still be dropped due to expiration.

Since interaction in the pub/sub system is time-decoupled, any two antennas of the country can be selected. Nevertheless, having λ_{out} (or λ_{in}) $>$ μ_{out} , in terms of overall mean values or over specific time intervals, will result in high numbers of expired events (see IV-C). This is when the event input flow to the system is higher than the event delivery flow constrained by the limited availability of the subscribers.

C. Simulation of the pub/sub model

We have developed a simulator that implements our pub/sub model and depending on the specific context, it receives as input the Antenna set for parameterizing the model. Our simulator, *MobileJINQS*, is an open-source library for building simulations encompassing constraints of mobile systems. *MobileJINQS* is an extension of *JINQS*, a Java simulation library for multiclass queueing networks [17]. *JINQS* provides a suite of primitives that allow developers to rapidly build simulations for a wide range of stochastic queueing network models [10].

MobileJINQS retains the generic model specification power of *JINQS*, while it provides additional features of interest to mobile or other systems such as: (i) lifetime limitation for each customer entering a queue, (ii) intermittently available (on-off) queue server or server with variable service rate over time to represent mobile users' behavior, and (iii) input flow with variable customer arrival rate over time to represent real input dataflow traces. In particular, a system designer is able to assign lifetimes, on-off intervals, variable service rates and arrival rates following well-known probability distributions.

By relying on the *D4D* dataset trace, we set up and execute a set of experiments with *MobileJINQS*. Our experiments and their results are presented in the following section.

IV. SIMULATION RESULTS

In this section, we provide results of *MobileJINQS* simulations for our pub/sub system with varied incoming loads, service delays and lifetime periods. We use the *D4D* dataset to derive realistic traces for incoming loads and service delays. We perform *event transactions* by sending events, according to the rate of the selected incoming load, between the input and the output access point. For each event, we assign appropriate lifetime periods to tune the system. We demonstrate that varying incoming loads and service delays has a significant effect on response time. In the case of varying lifetime periods, the tradeoff involved between the rate of successful event transactions and response time is also evaluated. The results of our experiments enable thorough analysis of the behavior of the system and can assist a system designer in the configuration of the system parameters in order to satisfy a set of design requirements.

A. Selecting representative input load

Based on Section III-B, our large-scale pub/sub system consists of 300,000 unique peers, 1,666 input and output access points, and 1,204,451,385 published events. Nevertheless, as defined in III-A we select traces from the set A' . Then, our pub/sub system consists of 579 input and output access points and 787,635,426 published events. To perform our simulations with varied traces, we classify the load of each selected trace into three main categories: (i) *low load antenna*; (ii) *medium load antenna* and (iii) *high load antenna*. Fig. 6 depicts four antennas used for the experiments of this section. Antenna 9 has a low load trace with an overall average rate of 0.04 connections over the 50-week period. Antennas 24 and 14 have medium load traces with overall average rates 0.075 and 0.082, respectively. Finally, Antenna 161 has a high load trace with overall average rate 0.129.




Fig. 6: Low, Medium and High load of antennas used for our experiments




Fig. 7: End-to-end transaction Response Times from low load Antenna 9 to high load Antenna 161

B. Response time for infinite lifetime

Fig. 1 represents the model used for the simulation. To get end-to-end response time between input and output access points, we tune the system with parameters: λ_{in} , μ_{in} and μ_{out} . In this experiment, to avoid very high response times we choose $\lambda_{in} < \mu_{out}$.

At the input access point, we map the load of Antenna 9 to the input flow λ_{in} . The service delay μ_{in} , which represents matching incoming events to subscriptions and routing them to corresponding brokers (in our case, to the single output access point), is set to 1 event per second. This means that we consider rapid event relaying to avoid event expirations at the input access point. At the output access point, the input flow λ_{out} is equal to λ_{in} , provided that the input access point does not saturate. For the service delay μ_{out} , we use the load of antenna 161. At a 10 min interval each connected subscriber is receiving one single event (i.e., subscribers' connectivity rate = 1). Finally, in this experiment we consider infinite lifetime period for each event.

Transaction response times are shown in Fig. 7a over a period of 50 weeks where users' connectivity rate equals to 1. We depict average values of response times over 2-week periods. The lowest response time value over the whole period is 40 mins. Between the 5th and 15th week, we get even higher response times (73 mins). This implies that the rate at which subscribers connect to receive events is not sufficient with respect to the rate of incoming events.

Consequently, getting lower response times depends on




Fig. 8: End-to-end transaction Success Rates and Response Times from low load Antenna 9 to high load Antenna 161 with varying lifetime periods

the subscribers' behavior. Fig. 7b shows response times if we change the subscribers' behavior by multiplying their connectivity rate with a factor of 10. This means that they connect 10 times more for receiving events. By this way, the highest and lowest response times become 3.5 mins and 1.8 mins respectively.

C. Response time vs. delivery success rate with varying lifetime periods

Considering that a system designer has no control over the subscribers' behavior, another way to get lower response times is to limit the lifetime of events, certainly at the cost of some events not getting delivered due to expiration. At the same time, this guarantees the freshness of delivered data, e.g., in the case of a traveler reporting some incident in the transportation network.

In this simulation experiment, we keep the same settings as in Section IV-B and we assign a lifetime period to each input event. Lifetime periods follow exponential distributions with mean values of 1, 10, 20 and 30 mins. End-to-end response times and event delivery success rates are shown in Fig. 8. As expected, with higher levels of lifetime periods, we observe higher success rates, along with much higher response times over the 50 weeks period. More specifically, by increasing the lifetime period, response times get higher rapidly, but success rates increase slightly after a certain lifetime value. When the lifetime period equals 1 min, we get the lowest response time, equal to 9.5 seconds, with a message success rate of 62 %.

In a subsequent set of experiments, we simulate different pairs of antenna loads at the input/output access points. We use the antennas shown in Fig. 6 and perform experiments for 3 different types of end-to-end interactions: input access point with medium load to output access point with (i) low load; (ii) medium load; and (iii) high load antennas.

Fig. 9 shows the mean value of response time and success rate for the 50 weeks period. For the first experiment, success rate is very low (0.04 %) even when using 30 min of lifetime period. On the other hand, response time is much higher comparing to the other pairs. For the second experiment, success rates are between 10 and 20 % with response times between 2 and 7.5 mins. Finally, for the third experiment, we




Fig. 9: Mean end-to-end transaction Response Times and Success Rates between Medium load Antenna and Low, Medium, High load Antennas with varying lifetime periods

get 35 - 50 % success rate with reasonable response times of 0.3 - 2.2 mins.

V. RELATED WORK

The Pub/Sub interaction paradigm is increasingly being used in the design of large-scale distributed systems. Consequently, investigating evaluation techniques of such systems has become crucial. *SCRIBE* [18] builds a multicast topic-based infrastructure and relies on several algorithms to optimize the routes from the publisher to each subscriber based on latency. To parameterize and evaluate *SCRIBE*, Zipf distribution for the topic sizes is used.

Corona [19] proposes a pub/sub architecture that is compatible with the existing pull-based architecture of the Web. In order to allow user subscriptions that span multiple events, the authors of [20] propose a solution based on event processing. To evaluate both the systems' scalability, authors use real-life RSS traces (100,000 RSS feeds and RSS V2.0 feeds from 415 websites, correspondingly). Similarly, to achieve event processing for large-scale, heterogeneous, and open environments, in [21] authors propose an approach to associate events and subscriptions with tags to describe their semantic themes. To evaluate the approach, authors use a large amount of events (14,743), where semantic heterogeneity holds, that is created by sensors deployed in the *SmartSantander* smart city project and the *Linked Energy Intelligence (LEI)* dataspace.

To extend pub/sub middleware systems operating in the mobile wireless environments, the approach in [22] guarantees that all the published messages are successfully delivered to all interested subscribers in their publishing order regardless of the current location of the mobile subscribers. Here, similarly to our approach, intermittent connectivity period has been taken into account. However, authors evaluate this approach using a *small-scale* setup (six instances of the JMS broker).

To allocate resources (i.e., minimum amount of resources needed, effective way to allocate and cost of hosting them) for a large-scale pub/sub system it is critical to get insights from the workload it drives and maximize the overall quality of service given to the subscribers. In [23] and [4], authors analyze the traces from a real deployment of *Spotify* and *Twitter* traces collected via public APIs. The analysis provides several interesting observations which can benefit pub/sub

system designers. The *Spotify* traces consists of about 1.1 million topics and 4.9 million subscribers forming about 12 million topic-subscriber pairs. The traces were gathered for 10 days (from 9th Jan 2013 to 19th Jan 2013) from *Spotify's* datacenter in Stockholm. Twitter traces provided around 8 million active users, 30 million subscribers, and around 683.5 million topic-subscriber pairs. This data was gathered for 10 days (from 30th Oct 2013 to 9th Nov 2013). In [11], authors provide a model to predict quality of service in terms of delivery probability and timeliness when deployed over unreliable, best-effort public networks. Probabilistic and real-world event traces used to validate the algorithms' accuracy and effectiveness. Regarding real-world event traces, a real-time stock market quote service has been used, where each publisher publishes real-time quotes of a stock to subscribers that are interested in that stock. NASDAQ stock quote event trace is obtained from Google Finance between the 4th and 5th December of 2009. The trace consists of 258,853 events from 2,792 stocks on the first day, and 272,974 events from 2,832 stocks on the second day.

One large-scale dataset, other than those mentioned above, is the *CDR*. Analysis of such information leads to various inferences. A survey about results obtained from the analysis of *CDRs* can be found in [24]. Among the results and inferences obtained, the most common one are the mobility and social networks related results. *CDRs* are also used to evaluate large-scale *mobile crowdsensing applications* [25] regarding the overall energy consumption in data transfer. Regarding pub/sub systems, in [26] a socially-aware broker allocation algorithm is proposed to trade-off between efficiency and overhead regarding message forwarding. However, the algorithm is evaluated using uniform distribution of 100 nodes as an initial condition in an area of $1km \times 1km$ and uses mobility models rather than very large-scale actual mobility data. As *CDRs* are also interpreted as mobility data sets, we utilize *CDRs* to evaluate our model. To the best of our knowledge, *CDRs* have not been utilized thus far to evaluate any aspects in pub/sub systems, especially quality of service metrics for large-scale systems with mobile characteristics.

VI. CONCLUSIONS

In this work, we study the behavior of the underlying communication infrastructure under context-rich realistic workload. We rely on the *Pub/Sub* interaction style to support loosely coupled interactions with additional constraints, in large-scale environments. Parameters such as intermittent connectivity of mobile users and freshness of delivered events are modeled using a queueing network for end-to-end interaction. The *MobileJINQS* simulator implements our model and leverages incoming loads and service delays derived from the *D4D* dataset. We demonstrate that varying incoming loads and service delays have a significant effect on response time. Furthermore, by introducing varied lifetime periods in the published events, we evaluate the trade-off between the rate of successful event delivery and response time. Results indicate that system designers should apply data validity lifetime periods depending on the context to achieve high performance on their applications.

Our future work includes the refinement of our model regarding the distributed pub/sub system (part B of Fig. 1)

by introducing multiple brokers, as well as applying limits to the buffers of each broker. Moreover, includes the study of the response time and success rate for the various combinations of antennas in more fine-grained scales (e.g., check what their evolution is over one day). Also, we intend to study and experiment with appropriate interaction styles (based on message-passing, events, data sharing) on top of 3G/2G/SMS data connections (or even future 4G [27]), further depending on the requirements of our application platform. For this, we plan to rely on our ongoing work on *Extensible Service Bus*, (XSB) [28], a framework for handling interconnection between interaction protocols applying the above-mentioned interaction styles. We are particularly interested in interaction adaptation depending on the network conditions (e.g., switching to SMS-based protocol when the 3G/2G network is unavailable). Furthermore, to deal with large-scale sensing interactions, we intend to build upon our recent work on *MobIoT* [29], a service-oriented middleware enabling efficient sensing over ultra-large population of mobile sensors.

ACKNOWLEDGEMENT

This work has been partially supported by the European Unions Horizon 2020 Research and Innovation Programme H2020 / 2014-2020 under grant agreement number 644178 (project CHOReVOLUTION³). Moreover, the authors would like to thank the D4D organizers for providing the dataset.

REFERENCES

- [1] "ICT Facts and Figures," <http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2015.pdf>, May 2015.
- [2] O. Mazhelis, H. Warma, S. Leminen, P. Ahokangas, P. Pussinen, M. Rajahonka, R. Siuruainen, H. Okkonen, A. Shveykovskiy, and J. Myllykoski, "Internet of Things Market, Value Networks, and Business Models: State of the Art Report," December 2013.
- [3] P. T. Eugster, P. A. Felber, R. Guerraoui, and A. M. Kermarrec, "The many faces of publish/subscribe," *ACM Computing Surveys (CSUR)*, vol. 35, no. 2, pp. 114–131, 2003.
- [4] V. Setty, G. Kreitz, G. Urdaneta, R. Vitenberg, and M. Van Steen, "Maximizing the number of satisfied subscribers in Pub/Sub systems under capacity constraints," in *INFOCOM, 2014 Proceedings IEEE*. IEEE, 2014, pp. 2580–2588.
- [5] H. Liu, V. Ramasubramanian, and E. G. Sirer, "Client behavior and feed characteristics of RSS, a publish-subscribe system for web microwebs," in *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*. USENIX Association, 2005, pp. 3–3.
- [6] M. Petrovic, H. Liu, and H. A. Jacobsen, "G-ToPSS: fast filtering of graph-based metadata," in *Proceedings of the 14th international conference on World Wide Web*. ACM, 2005, pp. 539–547.
- [7] V. Setty, G. Kreitz, R. Vitenberg, M. Van Steen, G. Urdaneta, and S. Gimåker, "The hidden pub/sub of Spotify:(industry article)," in *Proceedings of the 7th ACM international conference on Distributed event-based systems*. ACM, 2013, pp. 231–240.
- [8] J. Reumann, "GooPS: Pub/Sub at Google," *Lecture & Personal Communications at EuroSys & CANOE Summer School*, 2009.
- [9] "PubNub," <http://cdn.pubnub.com/pubnub-datasheet-2013.pdf>, 2013.
- [10] E. D. Lazowska, J. Zahorjan, G. S. Graham, and K. C. Sevcik, *Quantitative system performance: computer system analysis using queueing network models*. Prentice-Hall, Inc., 1984.
- [11] T. Pongthawornkamol, K. Nahrstedt, and G. Wang, "Probabilistic QoS modeling for reliability/timeliness prediction in distributed content-based publish/subscribe systems over best-effort networks," in *Proceedings of the 7th international conference on Autonomic computing*. ACM, 2010, pp. 185–194.
- [12] G. Bouloukakakis, N. Georgantas, R. Agarwal, A. Pathak, and V. Issarny, "Towards Mobile Social Sensing for Transport Information Management," in *D4D Challenge Senegal Poster Session, NetMob*, 2015.
- [13] S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell, "Bikenet: A mobile sensing system for cyclist experience mapping," *ACM Transactions on Sensor Networks (TOSN)*, vol. 6, no. 1, p. 6, 2009.
- [14] Y. A. De Montjoye, Z. Smoreda, R. Trinquart, C. Ziemlicki, and V. D. Blondel, "D4D-Senegal: The Second Mobile Phone Data for Development Challenge," *arXiv preprint arXiv:1407.4885*, 2014.
- [15] F. Afrik, "La Chambre de Commerce de Dakar et Orange Business Services pour booster l'entrepreneuriat au Sénégal," <https://goo.gl/FTd6J8>, April 2014.
- [16] ITmag, "Exclusivité : Tigo nouveau leader de l'internet mobile au Sénégal," <http://www.itmag.sn/news/exclusivite-tigo-nouveau-leader-de-linternet-mobile-au-senegal>, June 2014.
- [17] T. Field, "JINQS: An extensible library for simulating multiclass queueing networks, v1. 0 user guide," 2006.
- [18] M. Castro, P. Druschel, A. M. Kermarrec, and A. Rowstron, "SCRIBE: A large-scale and decentralized application-level multicast infrastructure," *Selected Areas in Communications, IEEE Journal on*, vol. 20, no. 8, pp. 1489–1499, 2002.
- [19] V. Ramasubramanian, R. Peterson, and E. Sirer, "Corona: A High Performance Publish-Subscribe System for the World Wide Web," in *NSDI*, vol. 6, 2006, pp. 2–2.
- [20] A. Demers, J. Gehrke, M. Hong, M. Riedewald, and W. White, "Towards expressive publish/subscribe systems," in *Advances in Database Technology-EDBT 2006*. Springer, 2006, pp. 627–644.
- [21] S. Hasan and E. Curry, "Thematic event processing," 2014.
- [22] A. A. Gaddah, "A pro-active mobility management scheme for publish/subscribe middleware systems," Ph.D. dissertation, Citeseer, 2008.
- [23] V. J. Setty, "Publish/Subscribe for Large-Scale Social Interaction: Design, Analysis and Resource Provisioning," 2015.
- [24] V. D. Blondel, A. Decuyper, and G. Krings, "A survey of results on mobile phone datasets analysis," *arXiv preprint arXiv:1502.03406*, 2015.
- [25] H. Xiong, D. Zhang, L. Wang, J. P. Gibson, and J. Zhu, "Enabling energy-efficient mobile crowdsensing with anonymous participants," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 6, no. 3, p. 39, 2015.
- [26] Y. Zhao and J. Wu, "Socially-aware publish/subscribe system for human networks," in *Wireless Communications and Networking Conference (WCNC), 2010 IEEE*. IEEE, 2010, pp. 1–6.
- [27] TeleGeography, "Senegals mobile operators to all deploy 4G pilot trials before year-end, Reuters," <http://goo.gl/SPxjUO>, December 2013.
- [28] N. Georgantas, G. Bouloukakakis, S. Beauche, and V. Issarny, "Service-oriented Distributed Applications in the Future Internet: The Case for Interaction Paradigm Interoperability," in *European Conf. on Service-Oriented and Cloud Computing*, 2013.
- [29] S. Hachem, A. Pathak, and V. Issarny, "Service-oriented middleware for large-scale mobile participatory sensing," *Pervasive and Mobile Computing*, vol. 10, pp. 66–82, 2014.

³<http://www.chorevolution.eu>