Big Wireless Measurement Campaigns: Are They Really Worth the Price?

Miguel Elias M. Campista[⋆,⋄], Marcelo D. de Amorim[⋄], and Luís Henrique M. K. Costa[⋆]

[⋆] Universidade Federal do Rio de Janeiro
GTA/COPPE-PEE/Del-Poli
Rio de Janeiro, Brasil
{miguel,luish}@gta.ufrj.br

[⋄] UPMC Sorbonne Universités
CNRS and LIP6
Paris, France
amorim@npa.lip6.fr

ABSTRACT

Performance evaluation of mobile wireless networking has to cope with two main issues: representativeness and feasibility. These two aspects are often contradictory because analyses must be thorough and, at the same time, neither costly nor time consuming. In a wireless network, the number of variables is potentially large and unexpected physical phenomena may occur at any time. Thus, balancing representativeness and feasibility often requires trace-driven approaches where the system is evaluated using real data collected from measurement campaigns. In this paper, we discuss possible criteria to allow researchers to find the traces which better fit their needs and expectations. Moreover, we suggest metrics that could be used to evaluate the value added by a new dataset with regard to existing ones. If possible, these metrics should be estimated before investing time and resources in a measurement campaign if this latter would not bring significant insights beyond those already available.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations—*Network monitoring*

General Terms

Measurement

Keywords

Trace-based performance analysis, trace similarity

1. INTRODUCTION

One major concern in performance evaluation is how to capture real system behavior to achieve representative analysis. In computer networking, researchers are engaged in this endeavor to find scenarios where their proposals can be undoubtedly checked. Experimental evaluation is preferred because it reflects real conditions and can potentially unveil anomalies, which would never be possible by other means. This approach, however, besides being expensive and time consuming, can be considered only as a snapshot of the network when experiments are conducted. Event-driven simulations are a powerful alternative, as they allow faster and cheaper investigations. Simulation also permits focusing on a given parameter, isolating all the others from the problem. These benefits can significantly simplify evaluation, but must rely on abstract simulation models and on the utilization of synthetic workloads that are, in turn, frequently based on experimental traces. Hence, there is a tradeoff between representativeness and feasibility in complex systems evaluation, reflected on what should be simulated and what should be run as a real experiment.

In wireless networking, system evaluation is even more complex because of the difficulty of repeatability [1]. Wireless environments present many physical phenomena regarding interference and signal propagation issues, which can severely deviate existing models from reality. In this case, a solution would lay somehow in-between simulations and real experimentation, aiming at combining the best of both approaches. Emulation is an attempt to fill this gap and consists of simulating part of the system and running the remainder in a real implementation [2]. Although the simulated and the real parts can be placed anywhere in a layered architecture, it is usual to see real systems evaluated under simulated physical conditions. In wireless networking, this approach is interesting because it would yield controlling the most complex part of the problem. The shortcoming, however, is still related to representativeness. It is generally unfeasible to run an emulation involving many wireless nodes. Moreover, leaving aside the most complex part of the system may not totally satisfy the representativeness requirement. Thus, emulation of wireless networks typically uses physical-layer traces collected from real measurement campaigns.

Much effort has been devoted to the acquisition of realistic traces, such as mobility patterns [3, 4, 5, 6, 7], inter-contact times among nodes [8, 9, 10], and propagation models inferred from indoor and outdoor environments [11]. Although we can point out some initiatives such as the Crawdad Project [12], the uncoordinated action is leading to a large number of real datasets that can be similar in relevant aspects. Hence, the originality of such measurements must be considered from start as well as the meaning of data. As a consequence, there is a strong need to define metrics to

HotPlanet'12, June 25, 2012, Low Wood Bay, Lake District, UK.

quantitatively measure the value-add of a trace with regard to another.

In this paper, we discuss some criteria to analyze the interest of a given trace. It is worth mentioning that measurement campaigns are fundamental and we do not aim at discouraging them. Instead, we seek to establish metrics to classify whether a trace is different from another and what those differences are. Once the differences are well understood, it becomes easier to derive conclusions from the obtained results. Furthermore, as the main characteristics of previous traces become clear, it gets easier to decide whether a new measurement campaign is worth the price. These metrics can aid researchers to justify the utilization of a trace in detriment of another or even if it is interesting to use two or more of them. In this direction, we bring up some questions that may be considered before using a dataset or starting a new measurement campaign:

- How can a dataset be analyzed?
- How can a dataset be compared with another?
- How can one analyze the impact of a dataset?

This paper is organized as follows. Section 2 introduces the problem and provides examples of different datasets for wireless networking performance evaluation. Section 3 talks over a three-layer data model identified in wireless network traces. Section 4 suggests some metrics to analyze trace similarity. It also addresses completeness and repeatability issues. Finally, Section 5 concludes this work and shows future work directions.

2. PROBLEM STATEMENT

Experiments conducted with Internet traffic have revealed inherent self-similarity properties [13, 14]. The same experiments, however, are not as simple in wireless environments because characterization works are often focused on modeling physical phenomena. At the physical level, wireless communication protocols have to cope with many variants, such as mobility patterns and propagation issues. These often depend on the application scenario, human behavior, and their interactions. Thus, extracting properties from experimental traces is still a challenge in wireless networking, which explains in part the efforts of the research community in this area [15]. Statistical analyses on these traces have found distributions which can fit specific aspects of interest. The inter-contact time between mobile nodes, i.e., the contiguous time during which two nodes are not connected has shown to follow a power-law model with heavy tail [16, 17]. This finding has been used to create scenarios for wireless networking event-driven simulations and also was used as an assumption for mathematical modeling. Nevertheless, even this well-accepted model has been criticized, as in the work of Chen et al. [15]. Their main observation is that shorter inter-contact times are simpler to measure and to derive statistics from. On the other hand, the heavy tail part of the distribution may still be blurred given the difficulty of guaranteeing whether there are events larger than those measured.

All the complexity involved in wireless networks gives more and more room for new measurement campaigns using different scenarios. The goal is to confirm previous results and to obtain new insights regarding the environment under consideration. Narrowing down a little bit more our discussion, a relevant example is the number of mobility traces that have been gathered so far. Since mobility patterns are difficult to synthetically generate, mobile traces are often a target. The complexity of simulating mobility behavior comes from the interactive nature of moving, driving, or even walking. People must timely react while moving to changing conditions, such as path congestions, traffic lights, or any sorts of unexpected events (e.g., someone crossing the street out of the pedestrian line). In addition, people movement is also concerned with individual plans and behaviors.

As node mobility is considered a challenge, the well-known random waypoint, widely accepted years ago, is not in use today. It was shown that its high level of abstraction can lead to erroneous results [18]. Thus, recent investigations have contributed with traces from real scenarios. Table 1 lists a number of available people and vehicular mobility traces along with information on the campaign duration, number of nodes involved, and main purpose. Note that many of them were obtained from hundreds of participating nodes, taking days or even a few years of measurements. These numbers are not negligible and obviously require high expenses and a huge coordinated effort. Our main argument in this paper becomes clear when someone wants to choose the best mobility trace to use. As there are many, finding the one which better captures the scenario under investigation is a hard task. This characteristic may even motivate researchers to conduct their own measurement campaign because it may be easier to create another trace then choosing the most suitable one from those already available.

3. DATA MODEL

Trace analysis becomes simpler depending on the level of abstraction required. If the goal is to capture global network behavior, the analysis becomes easier because details can be neglected. This is also called a macro view of the trace, since it is concerned with tendencies and average values, e.g., the average speed of cars in a road. On the other hand, considering details makes performance evaluation more complex as it gets closer to individual behavior. This is often called a micro view of the problem.

Figure 1 illustrates the different views that can be extracted from a wireless networking trace. Note that the global view offers a greater abstraction level than the node view. The neighborhood view is somehow in-between global and node views. It aims at revealing phenomena of interest happening within the vicinity of a node. Note that in wireless networking, this level is relevant because of the physical media, which is broadcast by nature. Although all these views are classified according to the abstraction level, they are all together in the same trace. Hence, researchers are in charge of identifying the level of information required for their analyses.

3.1 Global view

The global view is connected to statistical analysis of tendencies and average node behaviors. In a high-level view, we can observe the average speed of cars in a road or even in a city. In addition, we can conclude that the inter-contact time between nodes follows a power-law model with heavy tail. Thus, statistical modeling fits well when performance evaluation is willing to obtain the global picture of the scenario under consideration. This approach can make the analysis

Table 1: Summary of different mobility traces for people and vehicular wireless networks. It is worth mentioning that the number of nodes or access points is not always clearly documented in the traces shown below. When they are not, we extract them from the available traces, either counting the number of different IDs or MAC addresses.

Trace	Measurement duration	Number of nodes	Purpose
People Network			
USC Wireless LAN [9]	93 days	25,481 users + 79 access points	Users mobility and traffic characterization in an infrastructure wireless network
UCSD [10]	77 days	200 access points + 275 user devices	Connectivity and location traces of wireless users
Dartmouth [19]	1,177 days	13,888 users + 624 access points	Users mobility and traffic characterization in an infrastructure wireless network
Haggle [20]	4 days	20 static nodes + 4,519 wireless devices	Users mobility and intercontact behavior characterization
NCSU [4]	602 days	91 users	User mobility characterization
RollerNet [6]	3 hours	1,112 users	User mobility characterization
Vehicular Network			
TAPASCologne [7]	2 hours	117,484 cars	Vehicle mobility characterization
Monarch [3]	14 days	1,319 buses + 8 access points	Bus mobility characterization
UMassDieselNet [8]	20 days	35 buses + 6 relay nodes + 4,221 access points	Mobility, connectivity, contact duration, and data transfer traces between buses in a DTN and between buses and access points
DieselNet - AP connectivity [21]	20 days	35 buses nodes + 972 access points	Connectivity traces between hybrid networks composed of buses and access points
EPFL [5]	25 days	536 taxis	Taxi mobility characterization

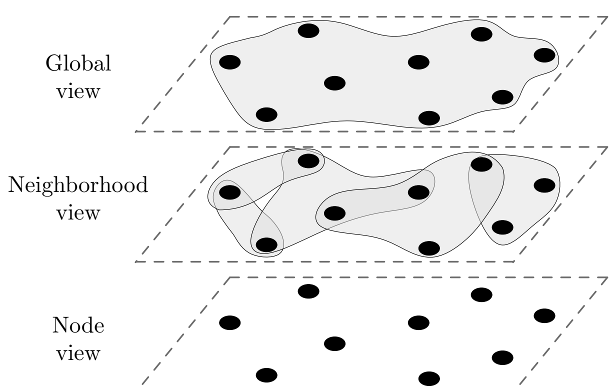

Figure 1: Different abstraction levels used in a wireless trace.

easier because all the details may not be considered. If one is interested in evaluating the behavior of a specific node, the global view can fail because it hides all the peculiarities and anomalies. Rare events can appear in a real trace and can be masked in a statistical model. Considering the time variable, a long-term analysis can fit better an average behavior than a short-term analysis. The problem is concerned with finding the observation time frame needed for an analysis to get closer to the expected behavior. This problem is addressed by collecting measurements from longer periods of time, without further investigations of how long should it last.

3.2 Neighborhood view

In wireless networks, neighborhood knowledge must be extracted from traces so as to better characterize each node surrounding environment. It is well-known that the node performance in the wireless medium is influenced by its neighbor's behavior. A challenge, in this case, is to delimit such neighborhood so as not to go farther than the area of influence. As a consequence, the node position must be considered, making the spatial notion an additional issue.

Similarly to the global view, the problem is also time-dependent. The characterization of the scenario is influenced by the neighborhood of a node, which changes with time. For instance, measuring a node degree at a time instant can give a clue about the density of its location and what kind of relationship is established among the nodes in that area. Nevertheless, it is not straightforward to affirm the consequences of such density. On the one hand, a small number of neighbors can lead to lack of connectivity, especially in mobile networks, whereas a high number of neighbors can lead to saturation and low data throughput. On the other hand, a large number of neighbors can lead to connectivity improvements, whereas less nodes can reduce the probability of achieving network saturation.

Note that coupling density to traffic behavior during wireless network characterization requires a deeper trace inspection. Since information regarding node vicinity is required, the global view cannot meet these needs. The neighborhood view, however, may require additional data, such as nodes' mobility and traffic measurements. One can have real mobility but may not know anything about network connectivity and traffic patterns. In this case, both can be simulated using existing modules. This approach can impact the final results, because it can be far from the reality. Observe that, at the neighborhood level, multiple traces can be required to fully characterize a wireless network.

3.3 Node view

Modeling the behavior of a nodes instead of the behavior of the entire network makes the analysis more complex because it requires individualization. It is easier to compute the maximum speed achieved in a given scenario than it would be to compute the maximum speed achieved by a specific node. The former does not need individualization whereas the latter requires associating the node to its respective speeds. Hence, node view compels a deeper inspection within data traces.

Researchers must judge in advance how deep they should go within their analysis to satisfactorily characterize a wireless network. Note that, as deeper we go, the more realistic and complex our analyses become. Hence, similarly to our initial problem, here again it is important to evaluate the impact of a deeper investigation in the final result, as we should do before starting a new measurement campaign.

4. TRACE ANALYSIS

In this section, we discuss some metrics that can be considered before starting a new measurement campaign or when comparing two or more traces.

4.1 Similarity

Finding trace similarity is a key issue for saving time and investments in new measurement campaigns, for identifying traces with different characteristics, and also for avoiding the storage of huge amounts of data without clear meaning. Hence, two different aspects can be considered before storing or considering traces. On the one hand, *self-similarity* is a property that looks for repeated data sequences inside the same trace so as to reduce file sizes [22]. Nevertheless, extracting information from masses of data is not straightforward. In this direction, compression methods are undertaken to reduce trace sizes by identifying and removing repeated patterns. On the other hand, *inter-trace similarity* is a property that aims at comparing two different traces so as to identify if they are complementary or not. The challenge in both cases, however, is similar and is concerned with finding an appropriate metric to evaluate representativeness.

A trivial approach to reduce trace size is by only keeping summarized information. The risk of this approach is to not collect enough information to identify behaviors. Avoiding this problem usually leads to storing as much information as possible. Reducing file sizes then requires the identification of repeated patterns inside the same trace. Determining self-similarity is, however, not simple because there are no guarantees of finding exactly the same data section within a single trace. Then, reducing traces by finding self-similarity can eliminate some data, even if some information is lost. Methods for finding similar sequences of events in traces are based on a parameter that can be relatively compared. For example, considering the execution of the same algorithm, the resulting data sections within the trace are concerned with the same sequence of procedures. Then, the duration of the first and of the second run of the same algorithm can be used as an indicative of self-similarity. If their duration are similar, based on a previously established threshold, they can be considered the same and one of them can be eliminated. This method uses a "distance" metric between two data sections. Other metrics for computing distances can consider a vector of values, instead of only the total duration, extracted from each section [22].

Another metric to verify self-similarity is the Hurst exponent, which computes the autocorrelation of the time series and the rate it decreases as the distance between two values grows. First, the dependence of the rescaled series is estimated by dividing the series in n sections and computing the ratio between the section range ($R(n)$) and its standard deviation ($S(n)$). If $R(n)/S(n)$ increases for an increasing number of sections, it means that range and standard deviation are inversely proportional with an increasing n, and, therefore, the distribution presents self-similarity. Plotting the expected value of $R(n)/S(n)$ for an increasing n in power-law fashion, the Hurst exponent is the linear coefficient of the curve found.

The inter-trace similarity problem is even more difficult because traces may not have the same format as they have in a single file. Therefore, the first challenge is finding parameters in common to identify similarities. Once those parameters are identified, there are statistical tests that allow computing a metric to estimate the similarity of a pair of traces. The Kolmogorov-Smirnov is a nonparametric test to compare continuous one-dimensional probability distributions. The comparison can be between an empirical distribution and a reference one (e.g., exponential, Pareto, and log-normal) or between two empirical distributions [23]. To this end, the test measures the distance D between the two probability functions $F_{1,n}(x)$ and $F_{2,n}(x)$ for n observations. This test can reject or fail to reject the hypothesis of similarity. If the distance D is large enough, the hypothesis is rejected and the distributions are not considered similar. Otherwise, if the test fails to reject the hypothesis, the distributions are considered similar for a given confidence level $1-\alpha$. An alternative test to the Kolmogorov-Smirnov is the Cramér-von Mises criterion, which also compares two distribution functions by computing a distance metric to estimate their goodness of fit.

Other statistical tests that can be applied over probability distributions are the Kuiper's test, which considers periodical repetition of a given behavior. This is an improvement over Kolmogorov-Smirnov test, which would lose such property because it considers uniform distributions. Since the periodical behavior is spread all over the distribution, the Kolmogorov-Smirnov does not capture periodical patterns.

4.2 Completeness

Experimental measurements are always conducted in a time frame. If this time is short, we can lose some events of interest after finishing the measurement or even during it. If we increase the time used for collecting traces, it is still possible to lose an event during the measurement or losing an event that lasts longer than the measurement time. Considering power law distributions with heavy tail, a long event can occur after the measurement has finished. For example, two nodes can lose contact during the measurement campaign and can reestablish contact only after the measurement has finished. Even if the contact is reestablished during the measurement, this event can be lost and not recorded into the trace file. The contact recovery event would be lost and would be erroneously considered for statistical analysis or be simply discarded. In both situations, the distribution found for inter-contact time would diverge from the real one.

Events lost during measurements or after them are called *censored data*. Computing the amount of censored data is possible by running a survivability analysis, which estimates the probability of losing an event. The survivability analysis uses an estimator, which defines a survivorship function ($S(t)$). For example, the Kaplan-Meier Estimator (K-M Estimator) [24] considers events not censored to compute the survivorship function of the censored events. Hence, algorithms to remove censored data can be used to approximate the data measurement from the real one [15].

In wireless environments, events can also be lost when sniffing the network by simply failing on overhearing all the traffic injected. Hence, using multiple nodes to overhear the wireless network reduces the probability of losing data. To assemble all data in a single trace, all the collected data are merged based on the recognition of equivalent patterns [25].

4.3 Repeatability

The repeatability problem is a recurrent issue from experimental analysis which does not have stationary behavior. In wireless networking, finding the same scenario and running the same test is a difficult, or even impossible task. Hence, repeating the same experiment using the same conditions previously set, i.e. addressing repeatability, is barely viable. Moreover, comparing different proposals must be carefully done so as not to lead to biased results. Usually, promoting fair comparisons is obtained by running long measurements so as to reduce the probability of rare events that could mask final results [1].

A trivial way of achieving repeatability is by using trace-driven approaches. Although there are testbeds for wireless networking which claim to guarantee repeatability, e.g., ORBIT [26], even them cannot address all possible situations one could face in a real scenario because they are limited to hundreds of nodes. Again, we face all the challenges discussed so far in order to find the more representative trace to achieve concluding and also fair results.

5. CONCLUSIONS

In this work, we draw attention to the numerous measurement campaigns that have been conducted on similar problems. As simulating complex networks and running experimentations are costly and time consuming, many works are using trace-driven approaches. Nevertheless, it lacks criteria to choose the traces that can better fit the performance evaluation required or to decide if another measurement campaign is really needed. In this work, we have also observed that the effort to extract information from a trace depends of the level of detail required and that there are metrics that can be used in order to estimate the similarity of a trace compared with others or even with itself.

As future work, we plan to quantitatively evaluate available traces used in similar problems to verify their impact. To accomplish this, we aim at establishing a framework to compute similarity between different traces starting from a high-level view to a low-level view.

6. ACKNOWLEDGMENTS

The authors would like to thank the Brazilian agencies CAPES, CNPq, Faperj, FINEP, and FUNTTEL and French agency ANR for their support. In addition, the authors would also like to thank the Crawdad Project [12] for maintaining a huge repository of wireless traces for the research community. Finally, we would like to thank LINCS (Laboratory of Information, Network and Communication Sciences) for hosting the authors.

7. REFERENCES

[1] A. C. Viana, S. Maag, and F. Zaidi, “One step forward: Linking wireless self-organizing network validation techniques with formal testing approaches,” *ACM Computing Surveys*, vol. 43, no. 2, pp. 7:1–7:36, February 2011.

[2] E. Lochin, T. Pérennou, and L. Dairaine, “When should I use network emulation?,” *Annals of Telecommunications (Online)*, pp. 1–9, July 2011.

[3] J. G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A. K. Saha, and D. B. Johnson, “Design and evaluation of a metropolitan area multitier wireless ad hoc network architecture,” in *IEEE Workshop on Mobile Computing Systems & Applications (WMCSA)*, October 2003, pp. 32–43.

[4] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, “SLAW: A new mobility model for human walks,” in *IEEE Infocom*, April 2009, pp. 855–863.

[5] M. Piorkowski, N. Sarafijanovoc-Djukic, and M. Grossglauser, “A parsimonious model of mobile partitioned networks with clustering,” in *International Conference on COMmunication Systems and NETworkS (COMSNETS)*, January 2009, pp. 1–10.

[6] P. U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. D. de Amorim, and J. Whitbeck, “The accordion phenomenon: Analysis, characterization, and impact on DTN routing,” in *IEEE Infocom*, April 2009, pp. 1116–1124.

[7] S. Uppoor and M. Fiore, “Large-scale urban vehicular mobility for networking research,” in *IEEE Vehicular Networking Conference (VNC)*, November 2011, pp. 62–69.

[8] N. Banerjee, M. D. Corner, D. Towsley, and B. N. Levine, “Relays, base stations, and meshes: enhancing mobile networks with infrastructure,” in *ACM Mobicom*, September 2008, pp. 81–91.

[9] A. Helmy, F. Bai, and W. Hsu, “Mobilib: Community-wide library of mobility and wireless networks measurements,” Available at

http://nile.cise.ufl.edu/MobiLib/USC_trace, July 2008.

[10] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," *Mobile Computing Communications Review*, vol. 9, no. 2, pp. 40–55, April 2005.

[11] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *ACM Sigcomm*, August-September 2004, pp. 121–132.

[12] Dartmouth College, "Crawdad Project," Available at http://uk.crawdad.org/index.php, April 2012.

[13] M. E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, February 1997.

[14] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, February 1994.

[15] Y.-C. Chen, P. Sreedevi, K.-T. Chen, and L.-J. Chen, "Analysis of opportunistic networks based on realistic network traces," Institute of Information Science, Academia Sinica, Tech. Rep., August 2006.

[16] L.-J. Chen, C.-H. Yu, T. Sun, Y.-C. Chen, and H. hua Chu, "A hybrid routing approach for opportunistic networks," in *ACM SIGCOMM workshop on Challenged Networks (CHANTS)*, September 2006, pp. 213–220.

[17] Y. Wang, S. Jain, M. Martonosi, and K. Fall, "Erasure-coding based routing for opportunistic networks," in *ACM SIGCOMM workshop on Delay-tolerant networking*, August 2005, pp. 229–236.

[18] J. Yoon, M. Liu, and B. Noble, "Random waypoint considered harmful," in *IEEE Infocom*, April 2003, pp. 1312–1321.

[19] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," *Elsevier Computer Networks*, vol. 52, no. 14, pp. 2690–2712, October 2008.

[20] A. Chaintreau, P. Hui, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Impact of human mobility on opportunistic forwarding algorithms," *IEEE Transactions on Mobile Computing*, vol. 6, no. 6, pp. 606–620, June 2007.

[21] A. Balasubramanianm, B. N. Levine, and A. Venkataramani, "Enhancing interactive web applications in hybrid networks," in *ACM Mobicom*, September 2008, pp. 70–80.

[22] K. Mohror and K. L. Karavanic, "Evaluating similarity-based trace reduction techniques for scalable performance analysis," in *ACM Conference on High Performance Computing Networking, Storage and Analysis*, November 2009, pp. 55:1–55:12.

[23] P.-U. Tournoux, J. Leguay, F. Benbadis, J. Whitbeck, V. Conan, and M. D. de Amorim, "Density-aware routing in highly dynamic DTNs: The RollerNet case," *IEEE Transactions on Mobile Computing*, vol. 10, no. 12, pp. 1755–1768, December 2011.

[24] E. L. Kaplan and P. Meier, "Nonparametric estimation from incomplete observations," *Journal of the American Statistical Association*, vol. 53, no. 282, pp. 457–481, June 1958.

[25] T. Claveirole and M. D. de Amorim, "WiPal: Efficient offline merging of IEEE 802.11 traces," *ACM Mobile Computing and Communications Review*, vol. 13, no. 4, pp. 39–46, October 2009.

[26] S. Ganu, H. Kremo, R. Howard, and I. Seskar, "Addressing repeatability in wireless experiments using orbit testbed," in *International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (Tridentcom)*, February 2005, pp. 153–160.