A New Mobility Trace for Realistic Large-Scale Simulation of Bus-based DTNs

Michael Doering, Tobias Pögel, Wolf-Bastian Pöttner, Lars Wolf
Institute of Operating Systems and Computer Networks
Technische Universität Braunschweig
Braunschweig, Germany
[mdoering|poegel|poettner|wolf]@ibr.cs.tu-bs.de

ABSTRACT

Realistic scenarios are essential for the simulation-based evaluation of opportunistic routing protocols in vehicular networks. Synthetically generated scenarios are easy to obtain but fail to reproduce the complexity of the real world. Therefore, the generally accepted procedure is to use traces recorded in experiments. Unfortunately, this is almost impractical for large-scale scenarios. The advancing pervasion of ICT in transportation systems results in new opportunities for the research community to collect mobility traces from real systems. Using the example of one of the world's largest public transportation network, we demonstrate our approach to acquire realistic traces which are more extensive than existing traces. Moreover, we demonstrate how this data can easily be integrated into the established delay tolerant network (DTN) simulation tool 'The ONE'.

Categories and Subject Descriptors

C.2.2 [**Computer Systems Organization**]: Computer-Communication Networks—*Network Protocols*

General Terms

Experimentation, Measurement, Performance

Keywords

Mobility Traces, DTN, Routing, Simulation

1. INTRODUCTION

DTN in public transport is an active area of research. Several architectures, routing protocols and applications for public transport delay tolerant networks (PTDTNs) have been proposed, e.g. in [1, 2, 3, 4, 5]. For the evaluation of these proposals, real mobility traces are important. However, to the best of our knowledge the availability of real mobility traces for large scale simulation based-evaluation is very limited. There are traces from the Dieselnet testbed with 40 vehicles [6] and mobility traces of 1200 buses in Seattle recorded in 2001 [7]. However, a set of traces for simulation should be as large and as diverse as possible. Such a set is essential to obtain reliable and comparable results.

In this paper we demonstrate the generation of traces from publicly available data drawing on the example of Chicago and present our approach to more realistic simulation of large scale bus-based DTNs. The remainder of this paper is structured as follows. First we give an overview of related work in the field of traces, DTN simulators and trace driven evaluation. After presenting the trace acquisition and processing, we analyze its characteristics and give a comparison to the only other publicly available large scale trace [7]. Next we present an exemplary large scale scenario for 'The ONE' [8]. We conclude the paper with our findings and future work.

2. RELATED WORK

Several other projects dealt with the generation of mobility models from GPS traces. In [9] position data of taxis was used to model a scenario with about 1000 vehicles in the inner loop of Shanghai. The original data was augmented by interpolation because GPS positions are only updated about every 40 seconds. Moreover, a route-determination algorithm was applied to fill the gaps between position updates. However, the characteristics of taxis are not comparable to buses. In [10] Zhang et al. used traces of a bus-based DTN-testbed with 40 vehicles to derive a generative model of inter-contact times. Traces from the same testbed were also used in [3]. VANET simulations based on bus movement traces with a significantly larger amount of vehicles were performed in [11]. The same trace was used in [12] to evaluate a cluster-based DTN forwarding algorithm. This trace was collected in 2001 in Seattle and includes over 1200 vehicles. To the best of our knowledge there are no other publicly available large scale traces for the simulation of bus-based DTNs available besides [7]. However, we believe that a diversity of large scale traces is required for the evaluation of DTN routing algorithms and for the pre-deployment simulation of bus-based DTNs. Therefore, in this paper we present our approach for the generation of traces from continuous real-time data publicly available (for free) via an API from the Chicago Transport Authority. The data contains 1971 buses, 150 routes and 11701 bus stops.

3. TRACE CHARACTERISTICS

Our mobility traces were obtained from the Chicago Transport Authority (CTA) Bus Tracker API, which is available and documented at [13]. Automated vehicle location systems (AVL) on the buses send position updates to a central server at CTA. For a duration of 18 days in November 2009, a script was used to store timestamped vehicle IDs and positions from the Bus Tracker API in a database. From this database scenarios for 'The One' are generated by trans-

CHANTS'10, September 24, 2010, Chicago, Illinois, USA.
ACM 978-1-4503-0139-8/10/09.

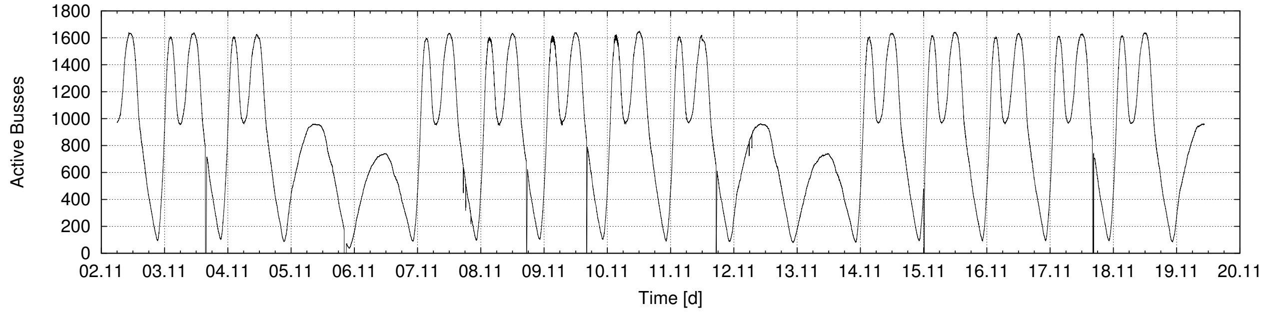

Figure 1: Amount of active buses over 18 days of the trace, with clearly visible rush hours and weekends.

forming from WGS-84 to UTM coordinates. The position data was checked for plausibility since the AVL occasionally reports bogus positions if the GPS-reception is bad. Therefore, roughly 0.03 % of the position data was discarded. The amount of active vehicles in the trace are plotted in Figure 1. Rush hours and weekends are clearly visible.

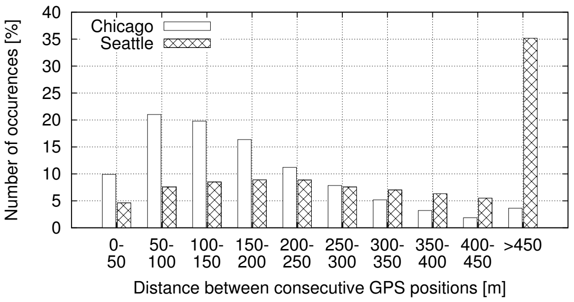

Figure 2: Distances between GPS updates

Comparing the 2001 traces of Seattle [11] with our 2009 traces of Chicago shows several similarities and differences. In order to achieve a fair comparison a Monday workday in November is selected from both traces. During this period the vehicles in Seattle sent 327880 valid position updates and those in Chicago sent 1736431, more than five times as many. The amount of active vehicles peaks to 1600 in Chicago and 800 in Seattle, as can be observed in Figure 3. Both traces show similar distinct rush hour spikes and vehicles move mainly at low speeds under 35km/h. However, the speed distribution plotted in figure 4 shows that vehicles drive more often at higher speeds in Seattle, which is due to a more rural area of operation in contrast to the denser downtown traffic in Chicago. Further, the update interval shown in Figure 5 is shorter in Chicago, leading to a larger amount of position updates. The rate is mainly between 20 and 40 seconds in Chicago but in the order of 1-2 minutes for the Seattle trace. Therefore, the Chicago trace is more suited as a basis for generating realistic mobility traces for simulations, since the higher position sample rate allows for a more precise extrapolation. For our scenario we extrapolated the positions between two updates as described in [12], but chose a one second extrapolation interval for a higher resolution.

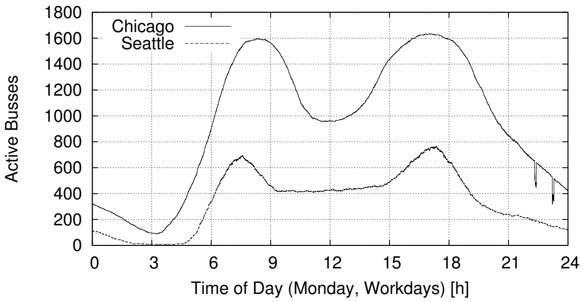

Figure 3: Number of active buses on Mondays

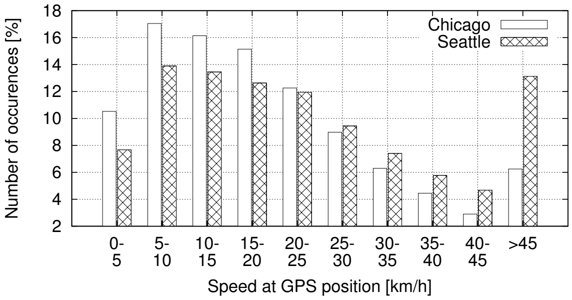

Figure 4: Distribution of Vehicle speeds

4. EXAMPLARY SIMULATION SCENARIO AND FUTURE WORK

Preparation and validation of the traces is work in progress, but the first results are very promising. A demonstration scenario for 'The ONE' and a video showing the mobility trace are available for download at `http://www.ibr.cs.tu-bs.de/bustraces`. The scenario runs for a 5 hour window from 7 to 12 a.m. on a Monday morning. It comprises 1648 buses (moving nodes) and 11701 stops (stationary nodes). In the scenario's default configuration a bundle is sent to each busstop. The bundles are generated at a gateway node in the densest area of the network, indicated by the two crossing lines in Figure 6.

The simulations were performed on a server with two Intel Xeon E5520 (2,26GHz, 4 cores each) and 32GB RAM. Processing time mainly depends on the routing algorithm and is in the order of several days. 'The ONE' is able to handle this large-scale scenario, but we noticed that some of

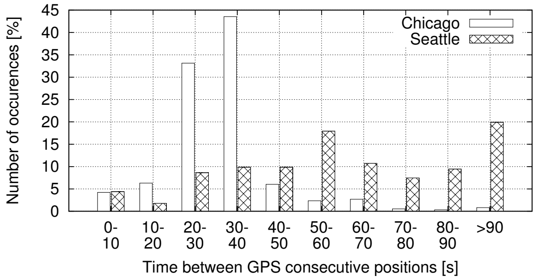

Figure 5: Distribution of GPS update rates

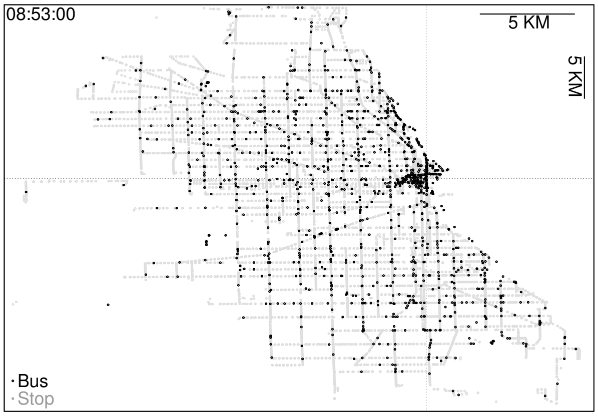

Figure 6: Exemplary snapshot of bus and bus stop positions to illustrate the vehicle distribution. (aspect ratio is optimized for printing)

the routing algorithms included in 'The ONE' do not scale well in the simulator. Especially MaxProp turned out to be very resource hungry. Besides the obvious optimization of the routing implementation we also assume that the scalability could be improved by optimizing the simulator for parallel processing since we noticed that only one core of our multiprocessor simulation server was busy.

In future work we will perform an in-depth analysis of the mobility trace and the buses' contacts. We also plan to contribute a set of reference scenarios based on specific use cases and our mobility traces. Another important area of research is the influence of the vehicles uneven distribution on routing performance. We will also investigate the optimal placement of gateway nodes.

5. ACKNOWLEDGMENTS

We thank the Chicago Transit Authority for developing the Bus Tracker API and for making the data publicly available. We also would like to thank the developers of 'The One'. This work was supported by the European Regional Development Fund under project number W2-80028895.

6. REFERENCES

[1] A. Balasubramanian, Y. Zhou, W. B. Croft, B. N. Levine, and A. Venkataramani, "Web search from a bus," in *CHANTS '07: Proceedings of the second ACM workshop on Challenged networks*. New York, NY, USA: ACM, 2007, pp. 59–66.

[2] A. Balasubramanian, B. N. Levine, and A. Venkataramani, "Enabling Interactive Web Applications in Hybrid Networks," in *Proc. ACM Mobicom*, September 2008, pp. 70–80. [Online]. Available: http://prisms.cs.umass.edu/brian/pubs/balasubramanian.mobicom.2008.pdf

[3] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, "MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks," in *In Proc. IEEE INFOCOM*, 2006.

[4] D. K. Jörg Ott, "Bundling the web: Http over dtn," in *WNEPT 2006 Workshop on Networking in Public Transport, QShine Conference*, 2006.

[5] S. Lahde, M. Doering, W.-B. Pöttner, G. Lammert, and L. Wolf, "A practical analysis of communication characteristics for mobile and distributed pollution measurements on the road," *Wireless Communications and Mobile Computing*, vol. 7, no. 10, pp. 1209–1218, Dec. 2007.

[6] Umass, "UMass Trace Repository." [Online]. Available: http://traces.cs.umass.edu/index.php/Network/Network

[7] J. G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A. K. Saha, and D. B. Johnson, "CRAWDAD data set rice ad hoc city (v. 2003-09-11)," Downloaded from http://crawdad.cs.dartmouth.edu/rice/ad_hoc_city, Sep. 2003.

[8] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE Simulator for DTN Protocol Evaluation," in *SIMUTools '09: Proceedings of the 2nd International Conference on Simulation Tools and Techniques*. New York, NY, USA: ICST, 2009.

[9] H.-Y. Huang, P.-E. Luo, M. Li, D. Li, X. Li, W. Shu, and M.-Y. Wu, "Performance Evaluation of SUVnet With Real-Time Traffic Data," *Vehicular Technology, IEEE Transactions on*, vol. 56, no. 6, pp. 3381–3396, Nov. 2007.

[10] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang, "Study of a bus-based disruption-tolerant network: mobility modeling and impact on routing," in *MobiCom '07: Proceedings of the 13th annual ACM international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2007, pp. 195–206.

[11] 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 *Mobile Computing Systems and Applications, 2003. Proceedings. Fifth IEEE Workshop on*, 2003, pp. 32–43.

[12] S. Ahmed and S. S. Kanhere, "Cluster-based Forwarding in Delay Tolerant Public Transport Networks," in *LCN '07: Proceedings of the 32nd IEEE Conference on Local Computer Networks*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 625–634.

[13] Chicago Transit Authority, "Bus Tracker API." [Online]. Available: http://www.transitchicago.com/developers/bustracker.aspx