On the Efficacy of Mobility Modeling for DTN Evaluation: Analysis of Encounter Statistics and Spatio-Temporal Preferences

Gautam S. Thakur*, Udayan Kumar*, Ahmed Helmy*, Wei-Jen Hsu§
*Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL
§Cisco Systems Inc., San Jose, CA
*{gsthakur, ukumar, helmy}@cise.ufl.edu, §wehsu@cisco.com

Abstract—**In mobile networking, the main goal of mobility modeling and simulation is the ability to accurately reproduce effects of realistic mobility on the performance of networking protocols. In the areas of adhoc and delay tolerant networks (DTNs), recent work on mobility modeling focused on replicating metrics of encounter statistics and spatio-temporal preferences. No studies have been conducted, however, to show whether matching these metrics is sufficient to accurately reproduce DTN protocol performance. In this study, we address this specific problem, and attempt to show the sufficiency (or lack thereof) of existing encounter and mobility metrics in reproducing realistic effects of mobility on networking protocols.**

We first analyze the characteristics of two well-established mobility models; the random direction and the time-variant community (TVC) models, and study whether they capture encounter statistics and preference patterns observed in real-world traces. Second, we contrast the performance of epidemic routing in DTNs based on the mobility models, to that based on extensive mobility traces. We provide two main findings. First, careful parameterization of the models can indeed replicate the metrics in question (e.g., inter-encounter time distribution). Second, even carefully crafted mobility models surprisingly result in protocol performance that is dramatically different from the trace-driven performance. The difference in message delivery delays can reach 67%, while difference in reachability approaches 80%. Such findings strongly suggest the need to re-visit mobility modeling. Furthermore, they clearly show the insufficiency of existing encounter and preference metrics as a measure of mobility model goodness. Systematically establishing a new set of meaningful mobility metrics should certainly be addressed in future works.

Keywords-component; spatio-temporal performances, encounter statistics, mobility models, epidemic routing.

I. INTRODUCTION

Over the past decade, the research activity on mobility modeling has increased significantly. Such models provide an essential component for analyzing and simulating mobile networks. Much of the recent mobility modeling work focuses on communication for extreme (adhoc and delay tolerant networks [DTNs]) environments [8], where infrastructure support is minimal, and connectivity incurs long delays and frequent network partitioning. Considered a necessary evil, mobility further exacerbates the problem of severely power and memory- constrained mobile nodes. Invariably, all these issues make a significant impact on the design and development of new networking protocols and services, including vehicular [3], inter-planetary [4], and disaster relief networks [19].

In addition, the proliferation of highly capable mobile devices (e.g., laptops, smart phones, tablets) with multi-sensing capabilities greatly facilitates the capture of mobility traces [17, 20] and the direct exchange of information through encounters. Mobility traces can then be used for modeling purposes, by mimicking encounter statistics [6, 15] or mobile user location visitation preferences [13]. Much of the recent modeling work focused on encounter metrics; such as inter-encounter and hitting time distribution [6], meeting duration [6, 13], or spatio-temporal profiles [13]. These metrics are generally considered important to the operation of DTNs and adhoc networks.

Thus far, however, no studies have been conducted to show whether matching these metrics is sufficient to accurately reproduce DTN protocol performance. In this study, we thoroughly examine this specific problem, and attempt to show, for the first time, the sufficiency (or lack thereof) of existing encounter and mobility metrics in reproducing realistic effects of mobility on networking protocols.

In this paper, first we analyze spatio-temporal properties [13] and encounter statistics using two realistic wireless measurement traces. We then evaluate the same characteristic on the synthetic traces produced by two different mobility models; the random direction model [22] and the time-variant community (TVC) model [13]. Specifically, we analyze two commonly used encounter statistics; inter-meeting time and meeting duration, in addition to two spatio-temporal metrics; periodic re-appearance and location visitation preference. Finally, we perform epidemic routing [21] on the synthetic (model generated) traces and real-network traces and compares their network performance. Surprisingly, through systematic analysis, we find that even when mobility models reflect equivalent spatio-temporal and encounter statistics, they exhibit large DTN routing performance discrepancy with the real scenarios. The results of epidemic routing show mobility models deviate 79% on average reachability, 67% on the delay and 58% on the overhead as compared to realistic scenarios. Such findings strongly suggest the need to re-visit mobility modeling. Furthermore, they clearly show the insufficiency of existing encounter and preference metrics as a measure of

978-1-4244-9538-2/11/$26.00 ©2011 IEEE

mobility model *goodness*. Systematically establishing a new set of meaningful mobility metrics should certainly be addressed in future works.

In the ensuing text, section II illustrates related work. Details of the datasets and trace analysis are explained in section III. We devote section IV to briefly describe our two mobility models, and then in section V we explain studied characteristics and evaluate the results. Finally, we conclude our paper in section VI by giving future work directions.

II. RELATED WORK

Delay Tolerant Networks (DTNs) are essentially opportunistic networks. These types of networks do not demand permanent connectivity between source and destination; instead attempt to make best use of any scheme available that can get the message across. Mobility of the nodes is often realized for transferring the messages. Design of any communication protocols for DTNs is heavily dependent on how well the underlying mobility is understood [2, 23]. There are two elemental ways to design and test the protocols for DTNs, namely Trace-Based and Mobility Model based [1, 5].

In case of trace based design and evaluation, a mobility trace can be downloaded from a limited number of trace repositories [17, 20]. These trace are from the real world and capture real mobility patterns of the users belonging to the traces. For the trace collection environment, testing the protocol on the traces would produce most realistic results. But there are quite a few drawbacks of using real traces such as limited number of traces, not capturing all scenarios and inability to generalize the results based on a few traces. Due to these drawbacks, researchers have proposed models that capture key characteristics of human mobility and produce synthetic traces.

Due to the complexity of understanding human mobility and modeling it, models are created to reproduce few characteristics from real traces such as inter-encounter time [23], regularity [7] and community behavior [9, 13, 18]. In most cases, a synthetic trace is validated by comparing a few key characteristics against the real trace. This validation we think is not the best, as it does not test the application oriented parameters of the generated trace. In this current work, we have taken a novel approach for evaluating the mobility models; which is to compare the performance of routing protocols on synthetic trace and on real trace whose characteristics were utilized to create/validate the mobility model. This approach allows us to test and create mobility models while keeping the applicability of the generated traces.

In this work, as a case study, we consider a complex mobility model TVC [13] (along with random direction mobility model [22]). This model generates the non-homogeneous behaviors of mobile users in both space and time. The traces generated by this model show (i) skewed location visiting preferences; (ii) time dependent periodical reappearance of mobile users as seen in WLAN measurements along with other encounter statistics such as average node degree and meeting time. It uses several real traces to validate the correctness of its design.

III. DATA SET AND TRACE ANALYSIS

A. Data Set

In order to realize the efficacy of mobility models matching real protocol performance and network dynamics, we intend to compare their output against wireless measurements. In first setting, we use Wireless LAN (WLAN) longitudinal session logs of mobile users from the IBM Watson WLAN campus[25]. These traces were collected in 2006 at access point level. For the purpose of this study, our goal is to retrieve a large and unbiased sample of such users. So, we filter the repository to select 1366 active though random wireless users over a period of one month. This trace is publicly available at [17] although we customized it in a format that suits us. Initially, we investigate mobile users' preferential attachment to certain locations and their time-dependent periodic behavior. Later on we compare them with the mobility models' output.

Our second dataset comprise encounter (*radio contact between two mobile devices*) traces from IEEE Infocom 2005 iMotes experiment[26]. This data is collected using Intel's iMote, which communicate on Bluetooth protocol and log contact information of all visible Bluetooth capable devices. Such a record contains three entities - MAC address, start time, end time that correspond to each encounter between the host and foreign device. As part of the experiment, these devices were distributed in conference settings to 41 participants for a period of three-four days. We transformed the gathered data for our need to study inter-meeting time and duration of meeting among mobile users and compare output between model and reality. This dataset is available at [17].

B. Encounter Traces

In order to pursue a study on the encounter statistics and dynamic routing in DTN, we need measurements that quantitatively depict the contact (*a.k.a. encounter*) between mobile users. An encounter occurs between them when they are in a radio communication range of each other. This is straightforward for the iMote's Bluetooth measurements that contain precise encounter information. However, the WLAN measurements are accumulated at the access point level and contain usage patterns. So, we need to convert these measurements in a way to get user encounters as well as maintaining their spatio-temporal footprints. *We consider encounter in WLAN if two users connect to same access point and share online session time*. For example, Alice and Bob are connected to access point AP-1 between 10:00 AM-02:00 PM. A counter argument can be established by saying that some WLAN devices may miss encounters beyond their coverage region of access points, but WLAN measurements have the advantage to obtain traces in much larger sizes with richer user presence. They also contain location information, which helps in spatio-temporal analysis. Mostly, a Bluetooth experiment has small set of user base for a limited time period.

IV. MOBILITY MODELS STUDIED

In this section, we discuss two mobility models used for evaluation. We use Random Direction Model [22], which does

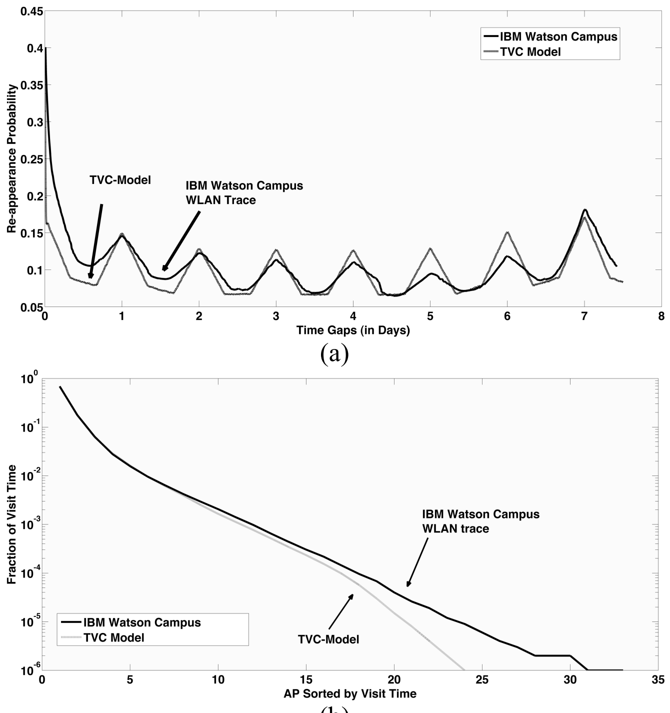

(a)

(b)

Figure 1: (a) Periodic re-appearances of mobile users in the IBM Watson Campus. TVC Model recreates similar preferences as observed in real traces. (b) TVC depicts skewed location visiting preferences as observed in IBM Watson WLAN traces.

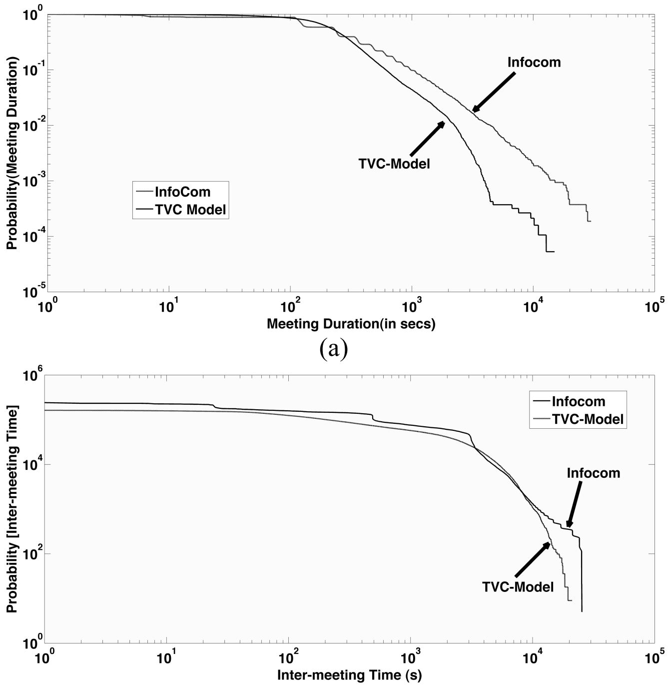

(a)

(b)

Figure 2: (a) TVC Model depicts meeting duration as measured in the real Infocom traces. (b) Inter-meeting time between mobile users are also similar for the TVC and real Infocom traces.

not posses any spatial or temporal structure in mobility decisions, as an example from typical random mobility models. The lack of spatial and temporal structure leads to faster mixing of the mobile nodes, and sets the lower bound for delay and message delivery overhead. This, as we will show, deviates from realistic mobility traces significantly. We further consider Time Variant Community model [13] as an example of trace-based mobility models, which incorporate realistic mobility characteristics observed in real traces. Our goal is to evaluate whether such more realistic mobility models lead to more realistic evaluation of routing performances. In the following text, we briefly describe these models and construct trace driven DTN scenarios to estimate routing performance.

A. Random Direction Mobility Model

In random direction model, a mobile node makes random mobility decisions with respect to current time or location, independent of other nodes. A node randomly picks a movement direction, and takes straight-line movement towards that direction for a given distance. The node then stops for a given pause time before selecting a new direction to move. This model is more stable as compared to other random models and provides quantitatively even distribution of nodes in the simulation area. We setup this model to investigate the effect of random movements on DTN performance. We modify this model in two ways: (1) the baseline random direction model described above; (2) we add on/off behavior of mobile nodes (i.e., when a node is "off", it cannot receive/transmit packets), which corresponds to the fact that mobile devices are not always turned on.

B. Time Variant Community Mobility Model

We choose the TVC model [13] as an example of trace-based mobility models that capture realistic features of human mobility. Specifically, the TVC model allows configurations to capture (1) spatial preference and (2) temporal periodicity in human mobility. With the setting of communities, preferred locations can be designated and mobile nodes visit such locations more often. The visits are further made periodically with the setting of time periods. TVC model also includes on/off behavior of mobile nodes. It is shown that with careful community and time period setup, TVC model produces mobility characteristics that match with the real mobility traces better [13]. Since the setup of TVC model is scenario specific, in this paper we have considered two instances of TVC model setup. We synthesize mobility traces from different settings of TVC (i) with matching location-visiting preferences and periodical visits to a trace collected at a research lab (ii) with matching encounter statistics at a conference. It is our goal in this paper to evaluate whether such improved realism in mobility characteristics translates to higher similarity in terms of routing performance to real traces, when we use TVC as opposed to random models.

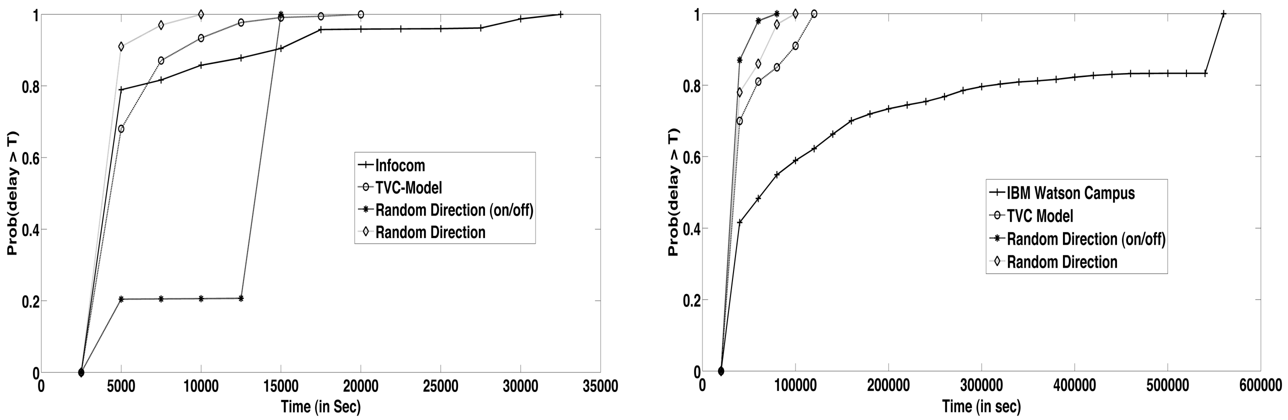

(a) Delay [Infocom]

(d) Delay [IBM Watson Campus]

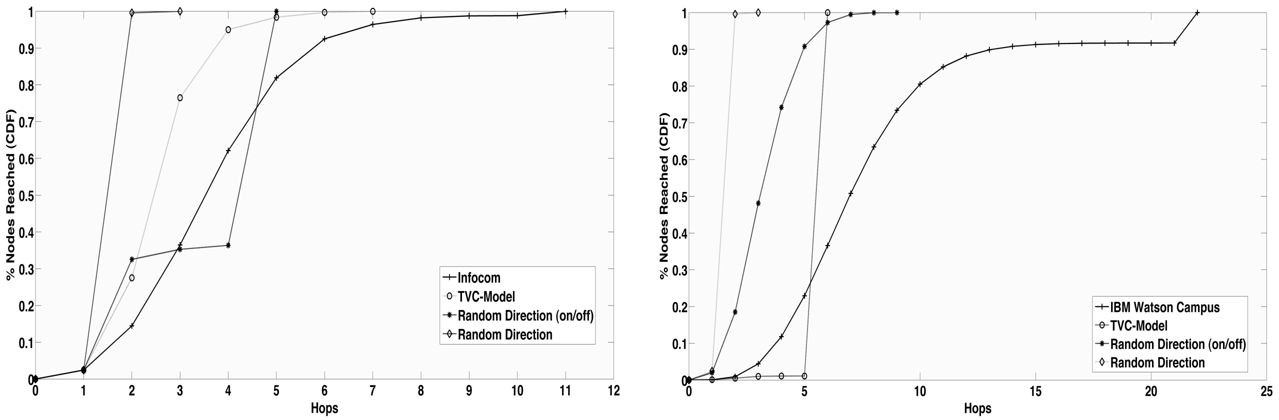

(b) Reachability [Infocom]

(e) Reachability [IBM Watson Campus]

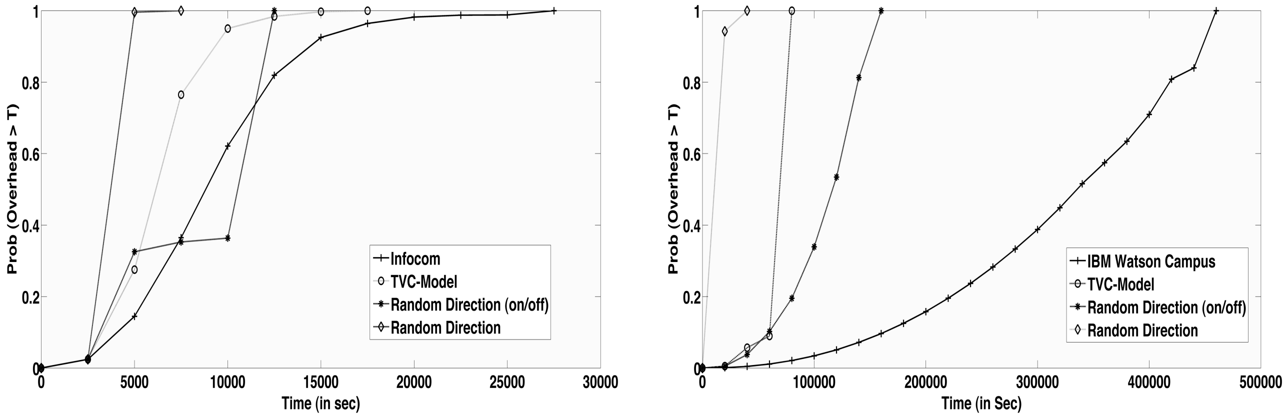

(c) Overhead [Infocom]

(f) Overhead [IBM Watson Campus]

Figure 3: (a-c) Show Epidemic routing results for the Infocom settings. (d-f) Show Epidemic routing results for the IBM Watson settings. As seen, largely deviate in their network performance for delay, overhead and message delivery compared to real measurement results.

V. EVALUATING MOBLITY CHARACTERISTICS

All mobility models are designed to demonstrate specific characteristics of mobility. In general settings, their capabilities are restricted to a few selected features observed in realistic mobility because of the structural and dynamical complexity involved in implementing them. The very fundamental question is, which are those metrics that define major aspects of mobility. The metrics proposed in [6, 13] and thereafter focus on: 1) Preferential attachment and spatio-temporal preferences; 2) Encounter statistics; and 3) Routing performances. Since humans invariably drive mobility, some work also incorporate human behavioral space [10, 14].

In this section, we showcase the presence of such characteristics in realistic measurements. Then we construct mobility models to demonstrate them and compare the results. Finally, we run epidemic routing on real and synthetic measurements and analyze the performance.

A. *Evaluation Metrics*

1) *Analysis of Spatio Temporal Preferences*

The non-homogenous behavior of mobile users in space and time is captured by: (i) Skewed location visiting preferences (ii) Periodical reappearances. Studies carried out in [7, 11, 12, 13, 16] tell us that mobile user exhibit preferential attachment and periodic reappearances to few locations in DTNs. We assume, understanding these distributions aid to better message dissemination, prediction of information transmission and the message delivery in opportunistic setting.

2) *Analysis of Encounter Statistics*

In dynamic infrastructure-less mobile networks (like DTNs etc.), the routing is performed by data carrying mobile nodes. The exchange of information takes place when two nodes encounter (a.k.a meet) each other. Intuitively, we can improve routing mechanism given we understand the statistics of these encounter patterns. So, we analyze two encounter statistics: (i) Intermeeting Time, which is the time gap that separates two consecutive mobile encounters. (ii) Meeting Duration, which is the single uninterrupted meeting duration surrounded by intermeeting times. Thus, our statistics alternate between each other. These statistics as mentioned in [15] can have important implications on the performance of opportunistic forwarding algorithms in challenged networks.

3) *DTN Performance Metrics*

We look into delay, reachability and message overhead to estimate epidemic routing performance. First, we convert the model's synthetic usage traces into encounter traces. Then we use an implicit time scale of trace duration to model network dynamics and encounters as a medium to exchange any information. Finally, we run epidemic routing on the traces to compare the performance of Random Direction, TVC and real measurements. Formally, *Reachability* is number of hops it took for source node's messages to reach all the recipients; *Delay* is the total time taken by source node's messages to reach all the recipients; *Overhead* is the average message count percentage incurred during the simulation.

B. *Results*

Here, we investigate how synthetic measurements produced by the mobility models deviate from realistic consideration. Initially, we construct TVC model to imitate spatio-temporal preferences observed in the realistic mobility patterns of IBM Watson WLAN campus. Then we re-configure the model to demonstrate identical encounter statistics of Infocom iMotes experiment. Since random direction is not designed to model these characteristics, we skip its evaluation and instead use its traces to focus on the routing performance only. Then we take combined synthetic traces earlier generated by the TVC (for IBM Watson WLAN and Infocom) and random direction model to evaluate the epidemic routing performance against the respective set of real traces. For interested readers, details of mobility setup and accompanying files are made available at [20].

1) *Spatio-temporal Analysis for IBM Watson Campus*

We construct the TVC model to generate a month long synthetic trace for 1366 nodes. In Fig-1 (a) and (b), we see that TVC model demonstrates realistically close location visiting and periodic reappearance properties. For brevity, periodic reappearances are plotted for seven days only. The reappearance of spikes demonstrates users visit the same location(s) with higher probability in a periodic fashion. A normalized curve of location preferences show nodes visit very few locations although spending significant amount of their online time. These two characteristics when combined results in better predicting the mobility and on/off patterns of mobile nodes. Furthermore, it can also help to identify hotspots and to measure an approximate delay in message reception. Next, we analyze the state space of encounters among mobile nodes.

2) *Encounter Statistic Analysis for Infocom*

The TVC model has ability to generate measurements to analyze encounter statistics. So, we configure it for Infocom setting to generate individual mobility traces for the same number of 41 nodes and for an equivalent duration of four days. The simulation area is modeled like a conference setting with flexibility to visit hotel rooms and outside locations. We later on process the generated traces and plot them along real measurements. The CDF plots in Fig.-2 show that model significantly matches real encounter statistics. We see intermeeting time follows Powerlaw distribution up to a characteristic time period after which it decays exponentially. This made us to believe that TVC can also be used to model encounter patterns for unknown scenarios. We conclude, TVC model is statistically accurate on these metrics and closely follow observed realities. We hope that it shows similar protocol performance as well, which we see next. In our case, we assume these metrics if captured via a model are vital in achieving identical performance in routing.

3) *Analysis of Epidemic Routing*

Our implementation of epidemic routing input time varying mobile encounter sessions. Essentially, they serve a basis for intermittently connected dynamic network topology setting where each encounter is viewed as an opportunity to receive and forward messages. We plot the routing performance in Fig-3 to validate real and synthetic traces in all cases of delay, reachability and overhead. Alongside, a quantitative report is shown in Table-1. Surprisingly, despite models claim to exhibit vital mobility characteristics, they dramatically deviate in network routing performance benchmarks. We observe epidemic results on Infocom experiment trace takes an average of 11 hops to deliver message to all other nodes; while it takes only seven in case of TVC traces and even less in case of Random Models. Meanwhile for the delay, there are at least two folds of difference between real measurement and synthetic trace. The TVC and Random models take much less time in delivering messages compared to the observed ones in the real scenario. Finally, message overhead for synthetic traces also diverge from realistic values.

C. *Results Discussion*

Mobility models are designed with a particular scenario in mind. However, in this study we would like to question the efficacy those metrics that are widely adopted or expected to

Table 1: Summary of performance measurement for epidemic routing.

Metric	**Infocom iMote**					**IBM Watson Campus**				
	O	***T***	***R1***	***R2***	***% Dev***	***O***	***T***	***R1***	***R2***	***% Dev***
Reachability	11	7	4	3	57%	22	4	7	3	78%
Delay	89	24	34	0.17	79%	66	37	2.4	0.18	80%
Overhead	0.006	0.004	0.003	0.001	55%	0.0012	0.0002	0.0036	0.002	61%

O = Real Traces; T = TVC Model; R1 = Random Direction (on/off); R2 = Random Direction. %Dev = deviation of TVC (T) to real trace (O)

be vital in closing the performance gap between modeling and reality. Now, there is an immense need to identify them and a perception should be made to use them for correct estimation. Our results show current metrics are inadequate; because models completely miss out on the performance criteria. We believe it is important for the researchers to search for fundamental characteristics that drive the dynamics in challenged networks. Not only we should maintain the current characteristic but also look out for structural semblance and topological realisms between simulation and similarity. A good research direction would be look into measures that affect globally in a similar way routing decision are made.

VI. CONCLUSION AND FUTURE WORK

In this paper, we scrutinize mobility models on routing performance benchmarks. We testify that despite models capture realistic human behavioral patterns; their routing performance deviate from reality. To make this point, we use trace-based mobility models to exhibit identical spatio-temporal preferences and encounter statistics as seen in two realistic measurements. Later on we used the same synthetic traces to run epidemic routing and measure performance. By doing so, we find that mobility models performance is not analogous to reality. The average number of hops and delay it takes to reach all nodes is respectively 68% and 79% less compared to the real trace overheads. The synthetic mobility traces' indeed carry no structural similarity. These dramatic deviations from realism indicate serious flaws in the existing models and their inadequacy as testbed tools for any kind of performance evaluation purposes. In this paper, we limit our work in verifying epidemic routing against two well-known mobility models. In future, we are looking to test other routing protocols and models that parameterize. In [24], we elaborated on the presence of similarity among mobile users and the detection of collective behavior via community detection in wireless networks. We showed the gap between reality and current mobility models in demonstrating collective behavior. In our on-going work, we are developing a multi-dimensional mobility framework that helps scientists to develop mobility metrics and verify current models against realistic settings and provide guidelines to develop new models. We are looking into a global perspective of clustering and mobility coefficient and maintaining structural properties and performance by revisiting mobility modeling, which is vital for the evaluation and design of next-generation behavior-aware protocols.

ACKNOWLEDGMENT

This work is funded by National Science Foundation (NSF) grants - NeTS 0832043 and Cisco Systems Inc.

REFERENCES

[1] Bai F. and Helmy, " A Survey of Mobility Modeling and Analysis in Wireless Adhoc Networks ", Book Chapter in the book on "*Wireless Ad Hoc and Sensor Networks", Kluwer Academic Publishers, 2008*

[2] Bai, F., Sadagopan, N. and Helmy, A. IMPORTANT: A framework to systematically analyze the Impact of Mobility on Performance of RouTing protocols for Adhoc NeTworks. *IEEE Infocom*, *2003*

[3] Zhang, X., Kurose, J., Levine, B. N., Towsley, D. and Zhang, H. Study of a Bus-Based Disruption Tolerant Network: Mobility Modeling and Impact on Routing. *ACM Mobicom, 2007, 195-206*

[4] Burleigh, S., Hooke, A., Torgerson, L., Fall, K., Cerf, V., Durst, B., Scott, K. and Weiss, H. Delay-tolerant networking: an approach to interplanetary Internet. *IEEE Communications Magazine, 41, 6 (2003), 128-136*

[5] Camp, T., Boleng, J., Davies, V. A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing, 2002*

[6] Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R. and Scott, J. Impact of Human Mobility on Opportunistic Forwarding Algorithms. *IEEE Transactions on Mobile Computing, 2007, 606-620*

[7] Eagle, N. and Pentland, A. Eigenbehaviors: Identifying Structure in Routine. *Proc. Roy. Soc. A*, *2006*

[8] Fall, K. A delay-tolerant network architecture for challenged Internets. *ACM SIGCOMM '03*.

[9] Girvan, M. and Newman, M. E. J. Community Structure in Social and Biological Networks. *PNAS, 2002*

[10] Hsu, W., Dutta, D. and Helmy, A. CSI: A Paradigm for Behavior-oriented Delivery Services in Mobile Human Networks. *IEEE ToN, 2008*

[11] Hsu, W., Dutta, D. and Helmy, A. Mining behavioral groups in large wireless LANs. *ACM MobiCom 2007, 338-341*

[12] Hsu, W. and Helmy, A. On nodal encounter patterns in wireless LAN traces. *IEEE WiNMee*, *2006*

[13] Hsu, W., Spyropoulos, T., Psounis, K. and Helmy, A. Modeling spatial and temporal dependencies of user mobility in wireless mobile networks. *IEEE/ACM Trans.Netw., 17, 5 (2009), 1564-1577*

[14] Hui, P., Crowcroft, J. and Yoneki, E. Bubble rap: social-based forwarding in delay tolerant networks. *MobiHoc*, *2008, 241-250*

[15] Karagiannis, T., Le B., J. and V. Milan. Power law and exponential decay of inter contact times between mobile devices, *ACM MobiCom 2007*

[16] Kim, M. and Kotz, D. Periodic properties of user mobility and access-point popularity. *Personal Ubiquitous Comput., 11, 6 (2007), 465-479*

[17] Kotz, D. and Henderson, T. CRAWDAD: A Community Resource for Archiving Wireless Data at Dartmouth. *IEEE Pervasive Computing, 2005*

[18] Musolesi, M. and Mascolo, C. A community based mobility model for ad hoc network research. *REALMAN*, *2006, 31-38*

[19] Nelson, S. C., Harris, I.,Albert F. and Kravets, R. Event-driven, role-based mobility in disaster recovery networks. *In Proceedings of the 2nd ACM workshop on Challenged networks (ACM Mobicom CHANTS '07)*

[20] NOMADS Lab. *MobiLib: Community-wide Library of Mobility and Wireless Networks Measurements*

[21] Vahdat, A. and Becker, D. Epidemic Routing for Partially Connected Ad Hoc Networks. *CS-200006. Duke University, 2000*

[22] Royer, E. M., Y, P. M. M. and Y, L. E. M. An analysis of the optimum node density for ad hoc mobile networks. *IEEE ICC*, 2001*, 857-861*

[23] Spyropoulos, T., Psounis, K. and Raghavendra, C. S. Performance analysis of mobility-assisted routing. *ACM MobiHoc, 2006, 49-60*

[24] Thakur S. Gautam, Ahmed Helmy, and Wei-Jen Hsu. 2010. Similarity analysis and modeling in mobile societies: the missing link. *In Proceedings of the 5th ACM workshop on Challenged networks (ACM CHANTS '10)*

[25] M. Balazinska and P. Castro. CRAWDAD data set ibm/watson (v. 2003-02-19). http://crawdad.cs.dartmouth.edu/ibm/watson, *Feb. 2003*

[26] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau, "CRAWDAD data set cambridge/haggle *(v. 2009-05-29)," Downloaded from http://crawdad.cs.dartmouth.edu/cambridge/haggle, May 2009*