

Performance Analysis of Location Profile Routing

David R. Bild, Yue Liu, Robert P. Dick, Z. Morley Mao, and Dan S. Wallach

Abstract—We propose using the predictability of human motion to eliminate the overhead of distributed location services in human-carried MANETs, dubbing the technique *location profile routing*. This method outperforms the Geographic Hashing Location Service when nodes change locations $2\times$ more frequently than they initiate connections (e.g., start new TCP streams), as in applications like text- and instant-messaging. Prior characterizations of human mobility are used to show that location profile routing achieves a 93% delivery ratio with a $1.75\times$ first-packet latency increase relative to an oracle location service.

I. INTRODUCTION

Traditional routing protocols rely on shared global state and thus scale poorly in mobile ad hoc networks (MANETs) with frequent changes in topology. Routing overhead grows quadratically in the number of nodes for distance vector and link state protocols [1] that must distribute changes to all nodes. The natural hierarchy used to reduce the overhead traffic in networks like the Internet (e.g., CIDR) is not available. On-demand methods [2], [3], [4] delay routing table updates until needed, but only reduce overhead by constant factors—the scaling behavior is unchanged.¹ In contrast, stateless protocols that use local information to make forwarding decisions have the potential to scale.

One stateless protocol, Greedy Perimeter Stateless Routing (GPSR) [5], uses geography: messages are addressed to specific locations. Nodes already know their own locations (e.g., via GPS), allowing each intermediate step to bring the message closer to its destination. No global routing state is needed. Essentially though, this technique just shifts the complexity from routing to addressing. A forwarding node only needs its own locally-known location, but the original sender requires the current location of the recipient, a global mapping.

Distributed location services [6], [7] can maintain this identity to location mapping, but also have drawbacks. Hierarchy is imposed to manage scalability, but overhead still increases super-linearly [7]. Further, locations are sensitive information, so complicated schemes are required to protect privacy and anonymity [8]. We observe that if node locations are predictable, the mapping can be done locally as well, reducing the scaling and privacy concerns.

In fact, human locations are highly regular with $\sim 93\%$ predictability [9]. In MANETs of human-carried devices, predicative models of future locations can be pre-shared among trusted participants. These models combined with GPSR allow

low-overhead addressing and routing, with network scalability limited only by the actual traffic. We name this approach *location profile routing* (LPR) [10] and study its performance potential. We determine the number of locations that must be addressed to achieve the peak 93% packet delivery rate and derive the associated latency and traffic overheads. Finally, we determine the conditions under which LPR outperforms the Geographic Hashing Location Service.

II. DESCRIPTION OF LOCATION PROFILE ROUTING

Location profile routing (LPR) stems from the observation that humans generally have simple, repeated motions, with most time spent at a few common time-dependent locations [11] easily captured by a compact predictive model. For the many potential² applications of human-carried MANETs that can tolerate the a reduction in delivery reliability or increase in latency (we previously described a particularly compelling application—censorship-resistant personal communication [10]), LPR eliminates overhead traffic for route maintenance.

Figure 1 illustrates the main steps of LPR. Nodes continuously monitor their positions to build location profiles (step 1), which are then shared with potential future contacts directly (step 2). This sharing happens a priori when two nodes are directly connected (i.e., one hop apart), limiting bandwidth usage. A message is addressed to the location(s) predicted by the profile (step 3) and delivered via GPSR. Routing fails if none of predicted locations are correct, but delay-tolerant delivery is a possible fallback. Changes to the motion patterns are rare (e.g., when someone starts a new job or moves to a new home) and thus distributed via the network (step 4).

Location Profiles: Motion patterns can be modeled in many ways, but a simple discrete model is sufficient for our purposes. A location profile is a function P mapping a time interval (e.g., Tuesday 15:30–15:40) to a set of location–confidence tuples, with higher confidence indicating stronger belief in the node occupying that location at that time:

$$P : \text{time} \mapsto \{(\text{loc}_1, \text{conf}_1), \dots, (\text{loc}_n, \text{conf}_n)\}$$

The precise discretization level is unimportant. Both cell-tower granularity (3 km², 1 h) and WiFi AP granularity (157 m², 10 min) have similar predictabilities at 93% [9] and 92% [12].

Various implementations are possible, but for completeness we summarize the Prediction-by-Partial-Match (PPM) approach of Burbey and Martin [12]. PPM is a variable-order Markov model over a sequence S of observed time-interval–location pairs, $S = \{T_1L_1T_2L_2\dots T_nL_n\}$. This defines a probability distribution over the next location conditioned on the prior k elements of context. In our case, prior locations are not known, so our definition of P corresponds to the first-order

¹This work was supported by NSF under award TC-0964545.

D. R. Bild, Y. Liu, R. P. Dick, and Z. Morley Mao are with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109. E-mail: drbild,liuyue,dickrp,zmao@umich.edu

D. S. Wallach is with the Department of Computer Science, Rice University, Houston, TX 77005. E-mail: dwallach@cs.rice.edu

²We assume that sender and receiver locations are not correlated; that could change the scaling behavior.

²Ad hoc networks are not yet widely used by the general public.

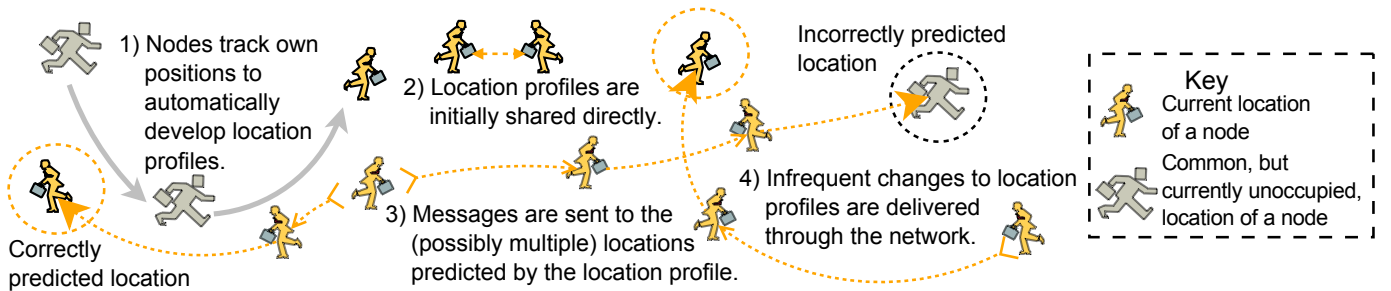


Fig. 1: Illustration of the main components in location profile routing [10].

variant ($k = 1$, i.e., context is the current time). We briefly discuss zero- (no context) and third-order (context includes the previous location) variants. This scheme captures most of the predictability (90% [12] vs. the 93% reported maximum [9]).

Profile Distribution: Location profiles are disseminated a-priori and out-of-band, similar to telephone numbers or email addresses. For our envisioned applications—communication between friends—the profiles can be exchanged face-to-face. In other cases, a centralized service, similar to a telephone directory, might be needed. Regardless, the salient point is that the profiles are known a-priori and thus can be exchanged outside of the network. Although changes could be disseminated out-of-band as well, in-network propagation is feasible because updates are infrequent and sent only to select participants (e.g., friends). Opportunistically updating when devices are in close proximity further bounds the overhead.

Addressing Policy: The addressing policy translates the location–confidence tuples output by the profile into a message delivery strategy specifying when and where packets will be sent. Only one of the locations can be correct, so the order and method in which they are tried influences the network throughput and latency trade-off. Their spatial correlations influence the minimum cost routing strategy (e.g., Steiner tree) to reach all locations. The primary focus of this paper is analyzing these performance characteristics and trade-offs.

Fallback Method: LPR fails outright when nodes are in unpredictable locations, i.e., at least 7% of the time [9]. Although this may be tolerable for many applications in which messages can be redelivered later, it is non-ideal. As this is not our focus, we omit details here, but possible strategies include delay tolerant delivery (in-network buffering of the message at a common location until the node’s return) or rendezvous delivery (messages are sent to a rendezvous location which the node, when not in a predictable location, apprises of current forwarding instructions). Such schemes allow for reliable delivery with average overheads still drastically lower than traditional routing approaches.

III. PERFORMANCE ANALYSIS

We use prior empirical studies of human motion patterns to develop analytical models suitable for studying the performance of LPR. Barabási et al. studied six-month location traces of 100,000 European cellphone users [11], [9] at cell-tower granularity, reporting a maximum predictability of 93%. The size and duration of the traces make this best source to date.

To confirm that locations are as predictable at WiFi granularity, we turn to Burbey and Martin’s study [12] of traces from 275 WiFi users at UCSD [13]. They found similar predictability, 92%, confirming that cellular granularity is not limiting.

A. How Predictable are Common Locations?

A location profile returns multiple locations in order of likelihood, so delivery cost and latency depends on how many, K , must be targeted to reach the user. Intuitively, most time is spent in two locations—home and work—so a zero-order model (i.e., not conditioned on current time) might be sufficient. The pmf is $\tilde{\pi}(k) = p_k \prod_{i=1}^{k-1} (1 - p_i)$, where p_i is the probability that the target is in location i . The p_i ’s are roughly distributed³ as $p_i \propto i^{-1}$ with proportionality constant $c \approx 0.48$ [11]. K is equivalent to a beta-geometric distribution, $K \sim \text{Geom}(L)$ with $L \sim \text{Beta}(c, 1 - c)$, and has CDF

$$\tilde{\Pi}(k) = 1 - \frac{1}{k B(k, 1 - c)}. \quad (1)$$

The match⁴ to measured data [9] is shown in Figure 2. The first moment diverges, but two locations suffice only 60% of the time and ten achieve only 80% delivery. Conditioning the model on time of day is necessary.

The first-order model (with 10 min intervals) is 90% accurate for the first location on the UCSD dataset [12], nearing the 93% upper bound and suggesting marginal gains for additional guesses. A third-order model is surprisingly only slightly better at 92%. The larger cellular dataset (with 1 h intervals) is more pessimistic. The accuracy $R(t)$ of the first-order model here is given by

$$R(t) = c_1 \sin\left(\frac{2\pi}{24}t + \frac{2\pi}{8}\right) + c_2 \sin\left(\frac{2\pi}{12}t - \frac{2\pi}{24}\right) + c_3, \quad (2)$$

where $c_1 = 0.148$, $c_2 = 0.077$, $c_3 = 0.657$ and $t \in [0, 167]$ is the hour of the week, i.e., $t = 0$ is Monday 00:00–0:59 and $t = 167$ is Sunday 23:00–23:59. As shown in Figure 3, this form captures one-day and half-day periodicities. On weekends, the variability is lower and the intervals of highest predictability occur later in the day. The accuracy on weekdays ranges from 55% to 90%, averaging $\bar{R} \approx 65\%$.

³A true Zipfian distribution requires a bounded domain $i \in [1, N]$ with $c = \frac{1}{H_N}$ for the p_i ’s to total one. The following results are for the reported empirical form, not a true distribution.

⁴ $L \sim \text{Beta}(0.60, 0.72)$ yields a tighter fit, but lacks an explanatory origin. It *might* result from a mixture of different upper bounds N in the Zipfian model of the p_i ’s—individuals have different numbers of common locations.




Fig. 2: The probability that a user currently occupies one of his k most-common locations is well-modeled by Equation 1.




Fig. 3: The time-dependent regularity $R(t)$, i.e., the probability the user is in the most common location associated with that time interval.




Fig. 4: Success rate of a first-order profile versus the number of locations attempted. Rates during maximum (night) and minimum (day) predictability are shown too.

Assuming the power law form, $p_i \propto i^{-1}$, holds during each time interval⁵, equations 1 and 2 can be combined as

$$\tilde{\Pi}_1(k) = 1 - \int_0^{168} \frac{D(t)}{k B(k, 1 - R(t))} dt, \quad (3)$$

where $D(t)$ is the traffic density at time t , to yield the average probability that packet addressed to the k -most common locations reaches the target, shown in Figure 4. We assume a uniform density, $D(t) = \frac{1}{168}$, but other known traffic patterns can be substituted. $k = 5$ achieves 85% success and 93% requires only $k = 12$. More locations are required during the day and fewer at night. The exact number of locations to attempt is application-specific, depending on the trade-off between desired delivery rate and cost, i.e., increased latency and traffic overhead.

B. What Additional Latency and Traffic is Induced by LPR?

Some packets must be sent to multiple locations to have an adequate packet delivery rate, increasing latency and traffic by constant factors. Note that the costs increase only for the first packet in a stream. Subsequent packets are sent directly to the now-known current location. The true average overhead depends on the percentage of first packets, which is low for applications like text-messaging and email and higher for interactive applications like voice chat. We report overheads for first packets only, which readers should scale by the first packet percentage of their applications.

We assume that receiver common locations and sender locations are uniformly distributed in the network.⁶ Thus, we can report overheads relative to the average latency (round-trip time) and traffic cost (round-trip hop count) for a single delivery attempt, e.g., a $2\times$ increase.

Parallel delivery to all k common locations does not increase latency, but increases traffic by $k\times$. Serial delivery—attempting each location only if the previous failed, using ACKs and a timeout to detect failure—reduces the traffic overhead. The pmf of the factor increase T , plotted in Figure 6, is

$$\Pr[T = t] = \tilde{\pi}_1(t), \quad (4)$$

⁵The number of common locations is inversely correlated with $R(t)$ [9] (Fig. 3B), suggesting that it does.

⁶Spatially-correlated locations can reduce overheads (see Section III-C).

where $\tilde{\pi}_1$ is the pmf associated with Equation 3. Latency increases similarly, as shown in Figure 5.

A combined approach—addressing a subset of the locations in parallel—can fine-tune the trade-off. For example, four different groupings can be used when trying three locations ($\sim 81\%$ success rate).



All locations within a group (a box in the diagram) are tried concurrently and groups are tried serially from left to right, as needed. Formally, a grouping G is a partition of the common locations, $G = \{g_1, g_2, \dots\}$, with the property that for $i < j$, all locations in group g_i are more probable than those in g_j . Let $\kappa(g)$ denote the index of the most common location in g , e.g., $\kappa(g_1) = 1$. Then, the probability that group g is tried, i.e., that all previous groups failed, is $\Phi_1(g) = 1 - \tilde{\Pi}_1(\kappa(g) - 1)$. Thus, the average latency increase \bar{L} for a grouping G is

$$\bar{L}(G) = \sum_{g \in G} \Phi_1(g) \quad (5)$$

and the average traffic overhead \bar{T} is

$$\bar{T}(G) = \sum_{g \in G} |g| \Phi_1(g). \quad (6)$$

Figure 7 shows the Pareto fronts for several average success rates, i.e., the maximum number of locations attempted. At the knees, $\bar{L} \approx 1.25\times$ and $\bar{T} \approx 3\text{--}4\times$. These curves are network averages. At runtime when a specific location profile is known, the precise trade-offs for that instance can be computed.

C. Under What Conditions Does LPR Outperform Location Services?

LPR trades the cost of updating a location service as devices move for multiple transmissions at the first packet. We use a simple analytical model to derive the network conditions under which LPR outperforms the Geographic Hashing Location Service (GHLS) [7], a scalable distributed location service. Let f be the network-wide location update rate (which increases with node movement), r be the network-wide first-packet rate, s be the average number of hops between a node and its GHLS location server, p be the average number of hops between a source and destination, and \bar{T} , as previously defined, be




Fig. 5: PMF of the latency increase for the first packet in a stream induced by trying multiple locations in turn. Concurrent attempts do not impact latency.




Fig. 6: PMF of the traffic overhead for the first packet in a stream induced by trying multiple locations in turn. Concurrent attempts have a fixed overhead.




Fig. 7: Pareto front of the first packet latency-traffic trade-off of a combined parallel-series strategy for several average success rates.

the average number of destinations attempted by LPR. The location update, location query, and first-packet delivery costs (i.e., transmission counts) for GHLS are⁷ fs , $2rs$, and $2rp$. LPR has only the first-packet delivery cost, $2\bar{T}rp$. After rearranging the total costs in terms of $\frac{f}{r}$ and $\frac{p}{s}$, we see that LPR has lower overhead when

$$\frac{f}{r} > \frac{p}{s}(2\bar{T} - 2) - 2. \quad (7)$$

When $s = p$ (source and destination are uniformly distributed over the entire field) and $\bar{T} \approx 3$ (from Figure 7), this simplifies to $\frac{f}{r} > 2$; LPR outperforms GHLS when the location update rate is more than twice the first-packet rate. This bound further decreases when sources and destinations are spatially concentrated, i.e., $p < s$.

D. Reducing Overhead Via Spanning Trees

The preceding overhead and latency analysis assumed linear routing, i.e., one transmission from the source per attempted destination. A branching route (e.g., the Euclidean Steiner tree containing the source and destinations) would reduce this overhead, particularly when destinations are spatially-clustered relative to the source. Unfortunately, this works only for dense networks in which nodes are guaranteed to exist at the branching (e.g., Steiner) points. Many real-world networks are too irregular, and the linear approach should be used.

In dense networks, the branching approach is feasible. One desires a routing tree with low total weight to minimize traffic but also with short source-to-destination path lengths to minimize latency. Although seemingly conflicting, both goals are achievable. Taking n as the size of the network, trees with weights within $o(n)$ of the $O(n)$ -length minimal Steiner tree and source-to-destination path lengths within $o(\log n)$ of the $O(\sqrt{n})$ straight-line distances exist [14]. We refer the reader to Aldous and Kendall for details and construction [14].

IV. CONCLUSION

We have argued that predictable motion patterns can replace expensive distributed location services in human-carried MANETs, reducing overhead transmissions. The promising

potential of LPR highlights the advantages of considering deployment- and application-specific behaviors in lower levels of the network stack. Our analysis focused on MANETs, but the approach applies equally to delay-tolerant networks. In light of growing commercial interest in device-to-device communication for smartphones (e.g., Wi-Fi Direct), we hope this work spurs further interest in adopting human behavior to improve ad hoc network performance.

REFERENCES

- [1] P. Jacquet, P. Muhlethaler, T. Clausen, A. Laouiti, A. Qayyum, and L. Viennot, "Optimized link state routing protocol for ad hoc networks," in *Proc. Int. Multi-Topic Conf.*, Dec. 2001, pp. 62–68.
- [2] M. Abolhasan, T. Wysocki, and E. Dutkiewicz, "A review of routing protocols for mobile ad hoc networks," *Ad Hoc Networks*, vol. 2, no. 1, pp. 1–22, Jan. 2004.
- [3] C. E. Perkins and E. M. Royer, "Ad-hoc on-demand distance vector routing," in *Proc. Wkshp. on Mobile Computing Systems and Applications*, Feb. 1999, pp. 90–100.
- [4] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," *Mobile Computing*, vol. 353, pp. 153–181, 1996.
- [5] B. Karp and H. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proc. Int. Conf. Mobile Computing and Networking*, Aug. 2000, pp. 243–254.
- [6] J. Li, J. Jannotti, D. S. J. D. Couto, D. R. Karger, and R. Morris, "A scalable location service for geographic ad hoc routing," in *Proc. Int. Conf. Mobile Computing and Networking*, Aug. 2000.
- [7] S. M. Das, H. Pucha, and Y. C. Hu, "Performance comparison of scalable location services for geographic ad hoc routing," in *Proc. Int. Conf. Computer Communications*, Mar. 2005, pp. 1228–1239.
- [8] A. Beresford and F. Stajano, "Location privacy in pervasive computing," *IEEE Pervasive Computing*, vol. 2, pp. 46–55, Jan. 2003.
- [9] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human motion," *Science*, vol. 327, pp. 1018–2021, Feb. 2010.
- [10] D. R. Bild, Y. Liu, R. P. Dick, Z. M. Mao, and D. Wallach, "Using predictable mobility patterns to support scalable and secure MANETs of handheld devices," in *Proc. Int. Wkshp. on Mobility in the Evolving Internet Architecture*, Jun. 2011, pp. 13–18.
- [11] M. C. González, C. A. Hidalgo, and A.-L. Barabási, "Understanding individual human mobility patterns," *Nature*, vol. 453, pp. 778–782, Jun. 2008.
- [12] I. Burbey and T. L. Martin, "Predicting future locations using prediction-by-partial-match," in *Proc. Int. Wkshp. Mobile Entity Localization and Tracking in GPS-less Environments*, Sep. 2008, pp. 1–6.
- [13] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," *Mobile Computing Communications Review*, vol. 9, no. 2, pp. 40–55, Apr. 2005. [Online]. Available: <http://sysnet.ucsd.edu/wtd/>
- [14] D. J. Aldous and W. S. Kendall, "Short-length routes in low-cost networks via Poisson line patterns," *Advances in Applied Probability*, vol. 40, no. 1, pp. 1–21, Mar. 2008.

⁷See Das et al. (Section IV) [7] for the derivation. Our s is their $\frac{1}{3}2^{h-1}\sqrt{2}$.