

Tailoring a broadcast algorithm to particular network conditions*

Carl Ellis
Lancaster University
UK
c.ellis1@lancaster.ac.uk

Hugo Miranda
University of Lisbon
Portugal
hmiranda@di.fc.ul.pt

François Taïani
Lancaster University
UK
f.taiani@lancaster.ac.uk

ABSTRACT

The broadcast of a message in Mobile Ad Hoc Networks requires its retransmission by multiple devices, consuming both bandwidth and power. The goal of broadcast algorithms is to select the most adequate nodes to retransmit so that the cost of the operation is reduced. Pampa is a broadcast algorithm that has been shown to perform well when the node distribution is uniform. This position paper identifies and discusses some cases where the performance of Pampa can be improved and presents the research directions that will be pursued to address them.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

Keywords

Mobile Ad Hoc Networks, Broadcast Algorithms

1. INTRODUCTION

Due to the decentralised nature of Mobile Ad Hoc Networks (MANETs), many services, like route discovery [5,10], reputation systems [9] or code propagation for sensors [7], require the delivery of some messages to every node. This operation is commonly referred as a broadcast. In some cases, the underlying infrastructure may provide tools to efficiently broadcast messages. This is the case, for example, of spanning trees provided by multicast routing protocols for ad hoc networks. This paper will focus on networks where these tools are not available.

In these cases, the most common implementation of broadcast is by flooding the network. With flooding, all nodes retransmit the message when they receive it for the first time.

*The work described in this paper was partially supported by the ESF under the MINEMA project and by Fundação para a Ciência e Tecnologia (FCT) and FEDER under project Probabilistically-Structured Overlay Networks (P-SON), PO5C/EIA/60941/2004.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MINEMA '08, April 1, 2008 Glasgow, Scotland
Copyright 2008 ACM 978-1-60558-122-4/08/04 ...\$5.00.

Flooding creates a large number of redundant transmissions. Many nodes receive multiple copies of the message, each transmitted by a different node. Therefore, it wastes a non-negligible amount of bandwidth and power. On the sending side, each retransmission consumes resources independently of its usefulness. On the receiving side, recipients also spend a non-negligible amount of energy receiving redundant messages [1] and deciding if they should be retransmitted.

Although the redundant reception of messages cannot be completely avoided, not all participants should be required to retransmit. The minimal number of nodes required to retransmit a message depends on factors outside the control of any broadcast algorithm, like the transmission range of the devices, the location of the source, the size of the region covered by the nodes or their geographical distribution. The role of broadcast algorithms is to devise a subset of nodes to retransmit that simultaneously: *i*) is minimal and *ii*) provides the largest coverage, measured by the proportion of nodes that receive the message.

The Power-Aware Message Propagation Algorithm (PAMPA) [8] is a broadcast algorithm that has shown interesting performance results in simulations with an uniform distribution of the nodes. Pampa was designed to be used in a large number of applications and therefore presents a minimal set of constraints. For example, the algorithm does not require devices to be aware of their location. This design philosophy causes Pampa to neglect some opportunities to improve its performance, that appear in situations typically associated with application of MANETs. In this paper, we discuss two cases:

- In some of the scenarios envisioned for MANETs a uniform distribution of the nodes should not be expected. This is the case for example of rescue operations where teams are possibly distributed uniformly over the affected region but members of the same team are expected to remain in proximity.
- Pampa ranks nodes according to their distance to the source of a message. This operation uses the Received Signal Strength Indication (RSSI) made available by the device driver. Below, we show that the RSSI is highly inaccurate. The increasing ubiquity of location devices (e.g. GPS) which can be easily connected to mobile computers suggests that in the future, one should expect some MANETs to be exclusively composed by devices which can be aware of their precise location.

In this position paper, we present on-going work to improve the performance of Pampa when each of the above

conditions is found. Each is being addressed in a different project. The goal of D-Pampa is to improve coverage in the presence of odd or non-uniform topologies. The G-Pampa project aims at finding more efficient and reliable ranking algorithms for location-aware MANETs.

The paper is organised as follows. For self-containment, Sec. 2 briefly describes Pampa. The impact of an uneven distribution of nodes in Pampa and the proposed solutions are discussed in Sec. 3. The extension of Pampa to networks composed of location-aware devices are the focus of Sec. 4. Finally, Sec. 5 concludes the paper by highlighting future work directions.

2. PAMPA

Pampa is a fully distributed algorithm that prompts nodes more distant to the source to retransmit first. In an ideal environment, and independently of the node distribution, this tend to favour the retransmissions that provide the highest additional coverage possible [12]. Pampa improves previous work [2,3,12] by removing some of the randomness used by these algorithms to select the nodes that retransmit.

In Pampa, when receiving a message for the first time, a node stores the message and sets a timer for a delay d , given by a function $delay$, to be addressed later. During this period, the node counts the number of retransmissions listened. The message is transmitted if, when the timer expires, the node did not listen to a sufficient number of retransmissions.

The function $delay$ uses the Received Signal Strength Indication (RSSI) of a transmission to compute a delay. This function is expected to map an increasing distance to the source (corresponding to a smaller RSSI) to a smaller delay. Implicitly, the function orders the nodes according to the distance to the source, with nodes more distant to the source expiring their timer first. It should be noted that the function is fully distributed: the algorithm is triggered exclusively by the transmission of the broadcast message and it does not require any coordination between the nodes. Like in the “counter-based scheme” [12], the algorithm prevents excessive redundancy by having nodes to count the number of retransmissions listened. However, Pampa bias the delay such that the nodes refraining from transmitting are usually those that are closer to the source.

The selection of a good $delay$ function is key to the performance of Pampa. We have been using a simple $delay$ function that multiplies the RSSI by a constant k to return the number of seconds that the node should wait before retransmitting. The most adequate value of k is likely to depend of the execution scenario. The problem of finding an adequate constant k mapping the RSSI in an adequate time interval has been discussed elsewhere [8] and is out of the scope of this paper.

3. D-PAMPA

Pampa in its current form is isotropic in that the retransmission rule is only based on signal strength. In particular, it does not take into account the general direction a message is travelling along. In uniformly populated networks, this does not seem to cause any undesirable consequences. We conjecture, however, that in heterogeneously populated networks, Pampa’s lack of ‘directionality’ may prevent the optimal propagation of messages.

Consider for instance the configuration shown on Fig. 1. A broadcast originates from Node A, and is retransmitted by




Figure 1: Non-propagation caused by a heterogeneous distribution




Figure 2: Experimental ns-2 distribution to test D-Pampa

B, which is the node furthest away from A within A’s range. Similarly the message is repeated by C. Both Nodes D and I are within range of the retransmission and receive the broadcast. Being each time closer to the source of the message than the retransmitting node (I is closer to A than B; D is closer to B than C), neither I nor D retransmit any message (in a one-transmission scenario).

D is unfortunately critically located in this example, as it acts as the only available bridge between the right-hand side cluster of nodes {A, B, C, D, I} and the remainder of the network {E, F, G, H}. For D to fulfil its bridging role, Pampa needs an additional mechanism to allow D to (i) detect that it lies at the boundary of a tight-knit cluster, and hence might be the only gateway to some other parts of the network and (ii) to decide to re-transmit the broadcast message in spite of additional retransmission it might have detected.

D-Pampa (for Directional Pampa) is an experimental early work to study the phenomenon described in Fig. 1. We are currently running simulations in ns-2 with heterogeneous node distributions such as the one shown in Fig. 2 to assess whether Pampa’s lack of directionality is indeed hindering propagation. To this aim, we run Pampa in scenarios where nodes happen to be located in ‘bridging’ positions between 2 otherwise disconnected parts of the network (the left and right hand-side ones in Fig. 2).

If we observe that the phenomenon is asserted, and Pampa does not perform as well in such cases, we have sketched an extension to Pampa that keeps track of transmitting hops since the original broadcast. In Fig. 1 for instance, D would record one two-hop retransmission by B, and one three-hop retransmission by C. Our goal is to combine this information with the signal strength to detect ‘boundary’ cases, and prompt an additional retransmitting behaviour in such cases.

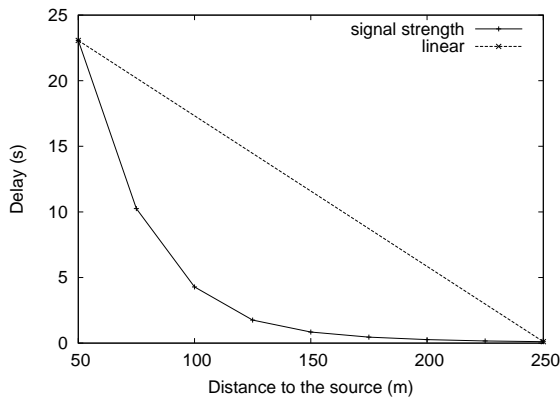


Figure 3: Different *delay* functions

4. G-PAMPA

Pampa relies on the Received Signal Strength Indication (RSSI) to separate nodes according to their distance to the source. However, realistic propagation models show that the signal strength decays rapidly with the distance and is affected for example by obstacles [6, 11]. Figure 3 shows a *delay* function that multiplies by some constant the signal strength. In this example we used the Two Ray Ground propagation model as defined in *ns-2*, v. 2.28. As it can be seen, when the distance to the source is large, there is a minimal difference in the delays to be applied at each node. This may result in concurrent decisions to retransmit by multiple nodes, resulting in a sub-optimal performance of the algorithm. In practice, concurrency can even be increased due to the lack of resolution of the RSSI variable, which according to the IEEE 802.11 [4] standard, may not provide more than 256 values.¹

A clearer separation between the nodes according to their distance would improve the performance of Pampa. An example is also presented in Fig. 3 where the delay decreases linearly with the distance. However, one must also take into account that mobile devices have limited computing power. Therefore, one must refrain from performing complex calculations upon the reception of each broadcast packet.

In contrast with location-unaware devices, determining the distance to the source of a message should be trivial if the sender includes its exact location in the packet and the receiver is capable of retrieving its own location, for example from a GPS. As future work, we plan to investigate the benefits of location-awareness by devising a novel *delay* function.

In complement with the work described in Sec. 3, location-awareness can also be used to improve the decision of each node to retransmit. As each node listens to the multiple retransmissions of the same message, he can learn the regions that have been covered by other nodes. The retransmission would then depend of the perspective of the node concerning the coverage of its own region.

5. SUMMARY

Many protocols developed for MANETs require messages to be broadcast. Therefore, the development of resource efficient broadcast algorithms is a fundamental contribution to

¹The range could be even smaller, given that it is only required that RSSI varies between 0 and a vendor dependant parameter `RSSI_Max`.

the successful deployment of MANETs. Pampa is a generic broadcast algorithm in the sense that it requires a minimal set of capabilities from the devices. In this paper, we showed that Pampa trades off some opportunities to improve its performance by its broad application domain and we highlighted some on-going research efforts to improve Pampa's performance when assuming atypical use-case scenarios.

6. REFERENCES

- [1] L. M. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *Procs. of the 20th An'l Joint Conf. of the IEEE Computer and Comm. Societies (INFOCOM 2001)*, pages 1548–1557, 2001.
- [2] Z. J. Haas, J. Y. Halpern, and L. Li. Gossip-based ad hoc routing. In *Proc. of the 21st An'l Joint Conf. of the IEEE Computer and Communications Societies (INFOCOM 2002)*, volume 3, pages 1707–1716. IEEE, 2002.
- [3] Q. Huang, Y. Bai, and L. Chen. Efficient lightweight broadcasting protocols for multi-hop ad hoc networks. In *Proc. of the 17th An'l IEEE Int'l Symp. on Personal, Indoor and Mobile Radio Communications (PIMRC'06)*. IEEE, 2006.
- [4] IEEE Computer Society LAN MAN Standards Committee. Wireless LAN medium access control (mac) and physical layer (phy) specifications. ANSI/IEEE Std 802.11, IEEE, 1999.
- [5] D. B. Johnson, D. A. Maltz, and J. Broch. *Ad Hoc Networking*, chapter DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks, pages 139–172. Addison-Wesley, 2001.
- [6] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proc. of the 7th ACM Int'l Symp. on Modeling, analysis and simulation of wireless and mobile systems (MSWiM'04)*, pages 78–82. ACM, 2004.
- [7] P. Levis, N. Patel, D. Culler, and S. Shenker. Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks. In *Procs. of the 1st USENIX/ACM Symp. on Networked Systems Design and Implementation (NSDI 2004)*, 2004.
- [8] H. Miranda, S. Leggio, L. Rodrigues, and K. Raatikainen. A power-aware broadcasting algorithm. In *Proc. of The 17th An'l IEEE Int'l Symp. on Personal, Indoor and Mobile Radio Communications (PIMRC'06)*, 2006.
- [9] F. Perich, J. Undercoffer, L. Kagal, A. Joshi, T. Finin, and Y. Yesha. In reputation we believe: query processing in mobile ad-hoc networks. In *Procs. of the 1st Annual Int'l Conf. on Mobile and Ubiquitous Systems: Networking and Services (MOBIQUITOUS 2004)*, pages 326–334, 2004.
- [10] C. E. Perkins and E. M. Royer. Ad-hoc on-demand distance vector routing. In *Procs. of the 2nd IEEE Work. on Mobile Computing Systems and Applications*, 1999.
- [11] T. Scott, K. Wu, and D. Hoffman. Radio propagation patterns in wireless sensor networks: new experimental results. In *Proc. of the 2006 Int'l Conf. on Wireless communications and mobile computing (IWCMC '06)*, pages 857–862. ACM, 2006.
- [12] Y.-C. Tseng, S.-Y. Ni, Y.-S. Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. *Wireless Networks*, 8(2/3):153–167, 2002.