Work in Progress: Evaluation of generic Bundle Transmission Scheduling strategies in Vehicular Disruption Tolerant Networks

Michael Doering and Lars Wolf
Institute of Operating Systems and Computer Networks
Technische Universität Braunschweig
Braunschweig, Germany
(mdoering|wolf)@ibr.cs.tu-bs.de

ABSTRACT

Due to the intrinsic properties of vehicular disruption tolerant networks, contacts between nodes exist only for a very limited amount of time. Therefore, for good performance it is important to transmit messages efficiently, and to minimize the waste of capacity by finishing the transmission before the contact ends to lower the amount of data that is discarded or fragmented at the end of a contact. We investigate the performance improvements of various bundle transmission schedulers in simulations with real-world mobility. Moreover, we propose a new generic approach to scheduling. Although our approach has a positive effect, the evaluation also implies that the influence of transmission scheduling in realistic scenarios is lower than generally expected.

1. INTRODUCTION

Limited transmission time during a contact is often a bottleneck in vehicular disruption tolerant networks. Contacts only last for a short time, since both nodes usually are moving, and because the radio range is limited. Therefore, it is vital for a DTNs performance to use transmission time efficiently. A breaking contact causes incomplete bundle transmission. As a consequence, the bundle protocol agent on the receiving node either discards the incomplete bundle or stores a fragment. Unfortunately, reactive fragmentation is problematic with security (e.g. via bundle authentication blocks) and would cause significant overhead in replicating routing algorithms. Therefore, it is preferable to minimize the amount of discarded or fragmented data. A promising approach is to reorder the bundles in the transmission queue. To the best of our knowledge this bundle transmission scheduling has not yet been sufficiently researched. Most routing schemes use a FIFO approach, although there is previous work [5] on sorting the queue using routing specific metrics in order to increase delivery probability. Our work focusses on the evaluation of generic scheduling approaches that are independent of routing specific metrics, and therefore can be applied to various existing routing schemes. For this purpose we have implemented several different scheduling schemes for the The Opportunistic Network Environment [4], which we use along with a real-world vehicular mobility trace from our previous work [1]. In the evaluation the performance of schedulers is compared with the theoretical optimum and complete random scheduling as baseline. Simulations comprise several routing algorithms and a broad range of different configurations and scenarios. Moreover, based on the assumption that contact durations are predictable to a certain extent [2] (e.g. because past average/median values are known), we design a scheduler that makes use of these properties.

The remainder of this paper is structured as follows: First, steps during bundle transmission are discussed and optimization potentials are identified. Additionally, we give an overview of previous work in this area. In Section 2, various existing scheduling strategies are investigated and our own hybrid solution is introduced. Moreover, our approach regarding baseline and optimum schedulers for evaluation purposes is presented. In Section 3 the evaluation methodology and the simulation setup is explained. Furthermore, evaluation results of various existing scheduling and routing approaches are presented and compared to our own solution. In the conclusions we discuss the impact of our findings.

2. GENERIC BUNDLE TRANSMISSION SCHEDULING STRATEGIES

Generic scheduling strategies are independent of metrics that are specific to a certain routing algorithm. A bundle transmission scheduling strategy allocates transmission time to bundles for the duration of a contact. It is assumed that the contact duration is not deterministic and may end before all bundles are transmitted. A scheduling strategy is advantageous if it reduces bundle delivery delay or/and increases bundle delivery rates. For this purpose a maximum channel utilization during bundle transmission is required (i.e. no transmission time is wasted). It can be assumed that this basic requirement is fulfillable by any well implemented scheduling algorithm. Another starting point is to minimize the overhead that is caused by bundles that are not yet completely transmitted when the contact ends. An incomplete bundle is usually discarded, i.e. the transmission time is wasted. Some routing algorithms are able to deal with fragmentation, but this causes at least additional headers for each fragment and additional overhead if fragments are

ExtremeCom2012, March 10-14, 2012, Zürich, Switzerland.

forwarded and replicated by other nodes.

In this paper we focus on scheduling strategies that are generic, i.e. we evaluate strategies that are independent of routing specific parameters. These have two advantages. First, they can easily be integrated to various existing routers; and second, this allows evaluation across a range of different routing algorithms. In the following, we introduce several practical strategies which serve as specific examples. Moreover, two further strategies are presented for benchmark purposes.

FIFO/LIFO

These bundle scheduling strategies order the bundles in the transmission buffer by creation time, using the "creation timestamp" field in the bundle header as specified in rfc5050 [6]. FIFO means that the bundles are sorted by age and that the oldest bundle is transmitted first, then the second oldest and so on. This strategy is intuitively "fair" and widely used in DTNs. LIFO is the inverted strategy, it starts by transmitting the youngest bundle first. Therefore, newer bundles outpace older bundles. The expected positive effect is that bundles on shorter paths are delivered with a lower latency, resulting in a lower global buffer space utilization. However, the negative effect is that old bundles are not transmitted at all, if the network is near congestion, because newer bundles are always preferred. On the other hand, this will decrease the overall average latency and increase the overall amount of successfully delivered bundles.

BIG/SMALL

BIG and SMALL order bundles by their size and send large resp. small bundles first. The basic idea behind these strategies is that an ending contact may have different effects depending on the size of the bundle that is in transit when the connection breaks. Because a large bundle requires more transmission time, it is more likely to result in an incomplete transmission. Moreover, this will either result in larger fragments or a larger amount of wasted transmission time because the incomplete bundle is discarded by the receiving node. Based on the fact that the duration of contacts in a vehicular DTN is not evenly distributed (as we have shown in [2]), is can be assumed that reordering the transmission queue by bundle size will have an impact on the overall performance. Again, some bundles may not be transmitted at all, like described above in the LIFO strategy.

SWITCH

SWITCH is a hybrid strategy that combines BIG and SMALL. At the beginning of a contact, this strategy behaves like BIG and starts with the largest bundle. After a certain time (at which the probability that the contact breaks increases disproportionately) the strategy switches to SMALL. Note that this change does not require resorting but just a reversal of the transmission queue. The expected result is a better utilization of transmission time but without discriminating large bundles.

Benchmarks: RANDOM and PERFECT

To understand general limits such as lower and upper bounds, we introduce two further approaches as references. RANDOM is a benchmark scheduler that chaotically reorders the transmission queue based on a uniformly distributed random function. Therefore, it is a baseline scheduler used in the evaluation. Any scheduler which results in a DTN performance that does not exceed this baseline is not worthwhile, because it does not improve anything but still increases computing overhead. PERFECT, on the other hand, is the upper bound of the positive effect that bundle transmission scheduling can have on a DTNs performance. It always makes globally optimal scheduling decisions, based on perfect knowledge (of both contact duration and the future routing path of the bundle).

3. EXPERIMENTAL EVALUATION

3.1 Evaluation Methodology

The purpose of the evaluation is to investigate if (and to which extent) the proposed scheduling strategies have an effect on the performance of exemplary real-world vehicular DTN scenarios. The common metrics are bundle delivery rate and bundle latency. However, for a fair comparision of scheduling strategies not only the amount of delivered bundles must be taken into account, but also the amount of delivered data. Therefore, the sum of payloads of all delivered bundles are chosen as an additional metric for the evaluation. This metric is global, i.e. it is summed up over all delivered bundles during a whole simulation run, and therefore describes the effects on the whole network. Besides a relative comparison of the performance metrics of different scheduling approaches, an absolute value of their quality can be given by using RANDOM (baseline) and PERFECT (optimum) as benchmarks.

Although it would be possible to use common mobility models or synthetic traces, we decided to use simulation parameters that are more realistic. For this purpose, we use the results of our previous work on real world vehicular measurements and mobility traces, and on an in-depht analysis of contact durations in vehicular DTNs. In [1] we reported how a mobility trace of Chicago busses was acquired, and gave an overview of the trace's properties. The analysis in [2] showed that it is possible to clusterize contact durations, and that there are several types of contacts, with different properties resulting from different contact situations. Moreover, it is possible to statistically predict durations if the contact situation is known. We are drawing on these results for the evaluation of SWITCH, in order to investigate if this approach is advantageous in a real-word environment.

3.2 Simulation Setup

For the simulation 'The ONE'[1] [4] is used, along with several included routing protocols that we use for performance comparison. According to our previous measurements [2] we chose 350m radio range and an average channel capacity of 1587.88 Kbytes/s as realistic parameters. The Chicago mobility trace is integrated as an external movement model. Because simulation of some of the routing algorithms included in The ONE with the entire real-world mobility trace would require several weeks of computation time for a single scenario we decided to simulate several subsets. With this approach a large number of smaller scenarios can be simulated, giving more reliable results than just a few big scenarios. Therefore, smaller scenarios are generated with a subset of 45 nodes and a five hour time window. 50 random pairs of source and sink nodes are selected. At the start of the

---

[1] http://www.netlab.tkk.fi/tutkimus/dtn/theone/

simulation bundles are generated at the source, and at the end of the five hour time window The ONE's report files are evaluated. The amount of bundles and their size are varied for different scenarios. Results for one amount/size configuration are given as minimum, maximum, average and median values of 50 simulation runs with different source/sink pairs.

3.3 Results

Figure 1 gives a first overview of the simulation results. Various combinations of routing algorithms and schedulers were simulated, each in 50 simulation runs with different source/sink pairs. The bars show the minimum, average and maximum delivered data in the whole network. At the beginning of each simulation run, 500 bundles with a random (uniform distribution) size between 1 and 10MB are generated at the source. The simulation ends after the five hour time window of the mobility trace described above. Note that the minimum amount of delivered data is zero for several combinations. This is because there are source/sink-pairs with a very bad DTN connectivity, and due to the large number of simulation runs and the random source/sink selection. Therefore, the set of simulations always contains 'unlucky' source/sink combinations, that inevitably have a very bad performance. The simulation scenario was deliberately chosen to cause congestion (the bottlenecks are the short contact times and a 2GB bundle buffer on each node). This means that even the optimum solution is not able to deliver all bundles because of the overload configuration.

As can be seen in the figure, there are significant differences in the performance of the routing algorithms. From left to right these are Spray-and-Wait [7], First Contact [3], Direct Delivery and Perfect. Perfect is the purely theoretical, omniscient routing algorithm described above and serves as reference for the best possible performance. It is clearly visible and not surprising that it achieves this best performance. It also has remarkable minimum performance values, because its omniscience prevents congestion in some situations. However, comparing the different schedulers for the same routing does not show very distinctive differences. The performance of SWITCH is lower than expected, because the selection of an optimal switch time is still work in progress.

4. CONCLUSIONS AND FUTURE WORK

In this work in progress paper we investigated starting points for the optimization of the bundle transfer phase in vehicular DTNs, and identified transmission scheduling as a promising field. Several scheduling approaches were implemented and simulated with various routing algorithms. We designed SWITCH, a hybrid scheduler, and evaluated it against several other approaches. Baseline and optimum solutions were used for comparison. The results show that the impact of scheduling is only marginal, and that the improvement by SWITCH is also not very significant, although a positive effect is observable. Currently, we are evaluating different switch times and algorithms for the selection of optimal switch times.

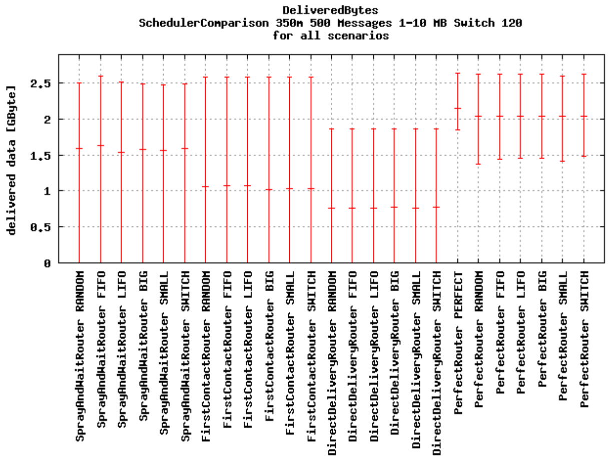

Figure 1: Overview of various routing and scheduling combinations with minimum, average and maximum amount of delivered data, each with 50 different source/sink-pairs

5. REFERENCES

[1] M. Doering, T. Pögel, W.-B. Pöttner, and L. C. Wolf. A new mobility trace for realistic large-scale simulation of bus-based DTNs. In *ACM MobiCom 2010 Workshop on Challenged Networks (CHANTS 2010)*, Chicago, USA, 9 2010.

[2] M. Doering, W.-B. Pöttner, T. Pögel, and L. Wolf. Impact of radio range on contact characteristics in bus-based delay tolerant networks. In *Eighth International Conference on Wireless On-Demand Network Systems and Services (WONS 2011)*, pages 195–202, Bardonecchia, Italy, 1 2011.

[3] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 145–158, New York, NY, USA, 2004. ACM.

[4] A. Keränen, J. Ott, and T. Kärkkäinen. The ONE Simulator for DTN Protocol Evaluation. In *Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, SIMUTools '09, pages 55:1–55:10. ICST, 2009.

[5] A. Lindgren and K. Phanse. Evaluation of queueing policies and forwarding strategies for routing in intermittently connected networks. In *Communication System Software and Middleware, 2006. Comsware 2006. First International Conference on*, pages 1 –10, 0-0 2006.

[6] K. Scott and S. Burleigh. Bundle Protocol Specification. RFC 5050 (Experimental), Nov. 2007.

[7] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In *WDTN '05: Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, pages 252–259, New York, NY, USA, 2005. ACM.