Mobility-aware Topology Manager for Publish-Subscribe Networks

Borislava Gajic, Jesús Palenzuela, Janne Riihijärvi and Petri Mähönen
RWTH Aachen University, Institute for Networked Systems
Kackertstrasse 9, D-52072 Aachen, Germany
Email: {gbo, jep, jar, pma}@inets.rwth-aachen.de

***Abstract*—In this work we analyze the effect of mobility on publish-subscribe based topology management and the computation of the data delivery paths. We take as a reference the WiMAX architecture enhanced with publish-subscribe features. In order to optimize the network operation we enrich the topology management functionality of the publish-subscribe network with a set of mobility prediction mechanisms. We elaborate various factors that influence the performance of the prediction function. We indicate the critical aspects of applying the prediction function and discuss the impact of such issues on overall system performance. Moreover, we examine the trade-off between the advantages offered by mobility prediction in the publish-subscribe context, compared to the signaling and complexity overhead. Despite the challenges posed by mobility, our work shows that the users' throughput in publish-subscribe networks can be improved by incorporating intelligent mobility prediction strategies into topology management functionality.**

I. Introduction

The end-to-end principle of the present Internet and its sender-oriented nature have introduced various problems by leaving large space for user misbehavior. The characteristic of the current networks is to deliver the data to the destination address specified by the sender. In such sender-driven way of communication the receiver is left without direct control over the traffic flow. As a consequence, the attacks targeted to the specific end-point in the form of unwanted traffic are rather simple to achieve. Such vulnerability of end users, as well as tight connection between the address and the location, have become one of the main concerns of the present Internet. Additionally, the user requirements towards the network have evolved over time. The users are primarily interested in the content they are willing to retrieve, regardless of the source from which the data originates, leading to the emergence of an information-centric type of data exchange.

Due to its information-centric characteristics, the publish-subscribe [1] pattern appears as a more flexible communication model than the traditional host-centric one. The communication is carried out by means of publications and subscriptions to the content. The publications serve to the network entity as an artifact for announcing the available data, whereas the subscriptions indicate interest in receiving particular content. A dedicated rendezvous system is accountable for observing the publication and subscription events, finding the match between them and commencing the process of its handling. The actual conveyance of data from the publisher to the subscriber is performed upon the appearance of the matching event. This enables complete control of the data reception, diminishing the probability of receiving undesired content. A large number of architectures have already been inspired by the publish-subscribe idea of departing from current end-point oriented inter-networking, giving the major role to the content itself [2]–[7]. The PURSUIT project [8] targets an entire restructuring of the present Internet architecture, grounding the new model of the Internet completely on the publish-subscribe fabric. Our work builds upon the PURSUIT architectural concepts and it is easily applicable to any information-centric architecture. The main building blocks of PURSUIT architecture are rendezvous, topology and forwarding modules. After the user interests have been matched on the rendezvous system, the rendezvous issues the request for the generation of optimal forwarding paths from the publisher to the subscriber. Upon the required paths being obtained from the topology module, i.e., topology manager, the forwarding system performs the final data transmission.

Optimization of data delivery paths in topology manager is a critical operation since it directly influences the efficiency of the entire system. The path from the publisher to the subscriber can be optimized with respect to a wide range of network parameters and thus affects the performance of the system in different ways. Moreover, in mobile environments the network topology changes frequently. The nodes come and leave and the links between them alter correspondingly. The topology manager needs to keep track of all modifications in the network configuration and to update the forwarding paths accordingly. This is not a trivial task, since the network changes can occur frequently and cause significant packet loss unless the topology manager comprises the mechanism to prevent it.

In order to address this issue a variety of protocols have been introduced in the present IP-based Internet [9], [10]. However, the tight connection between addresses and locations imposes the need of rather complex mechanisms in order to achieve satisfying performance of the system. Due to its data-centric characteristic the publish-subscribe communication pattern provides a firm basis for mitigation of mobility issues. Furthermore, the topology management in publish-subscribe architectures represents the key point for handling the mobility-related events.

Due to its highly important role in publish-subscribe architectures we implement the topology management on top of

978-1-4673-0229-6/12/$31.00 ©2012 IEEE

PURSUIT network and investigate the different optimization techniques for improving the topology management functionality. We focus on utilization of publish-subscribe data-centric characteristics, in particular the flexibility of providing the same data from disparate locations. The users are not restricted to data delivery from a single and unchanged network entity throughout all the communication, but the data can be fetched from any point in the network. This gives a certain level of freedom to the system to alter the data suppliers according to the current network state and the position of users. The possibility of such dynamic changes enables the better utilization of network resources and decrease of data latency. Furthermore, we aim at employing prediction techniques of the user movements. By collecting the information related to the users' flow, it is possible to gain knowledge about the usual movement pattern and to use it for forwarding paths creation. Instead of ordinary re-computation of forwarding paths upon the topology changes due to the node's allocation, our objective is to secure the communication in such a dynamic environment by generating backup forwarding paths and bicasting the data based on the prediction of node's next position. Moreover, we optimize the prediction algorithm by dynamically changing the prediction-relevant parameters according to the current network state.

The rest of this paper is organized as follows. In Section II we outline the related work in the domain of publish-subscribe communication with an accent on mobile scenarios and topology creation. Section III explains our main design principles in the implementation of topology management solution using ns-3 network simulator. We proceed by illustrating evaluation results and discussing the performance of the proposed model in Section IV. Finally, in Section V we draw the conclusions.

II. Background and Related Work

The architecture of the present Internet was designed with a strong focus on the sender-oriented way of communication. However, over time end-user requirements changed and the data content has become of central importance. This initiated the work towards content-centric communication [11]. Due to the asynchronous characteristics of content-centric communication, grounding it on the concepts of publish-subscribe model appears as a natural approach. Numerous data-centric notification models have been developed within the publish-subscribe fabric [12], [13]. Furthermore, various research projects are targeting the redesign of the current Internet and building it over the publish-subscribe principle [6]–[8].

Special attention has been devoted to mobility support in publish-subscribe networks. The issue of dynamic networks and frequent change of topology has been addressed from various perspectives. Muthusamy et al. [14] analyze the impact of publisher mobility on system performance. The work is focused on finding the solutions to diminish the delivery latency due to the frequent rebuilding of the multicast tree. The proposed algorithms rely on the use of proxies representing the fixed points in the network responsible for intermediate storage of publications. Due to the tight connection of data delivery with the system of proxies, the authors concentrate on the mobility of the publisher and operations related to handling of such events on proxies. Our mobility model and performance analyses are unrelated to the type of moving entity. Regardless of the case of a moving publisher or subscriber, the subscriber always acquires the information from the closest location.

The work in [15] exploits the advantages offered by caching in order to improve the network performance in mobile conditions. The model is based on a system of brokers accountable for dedicated network regions. The brokers are fixed and play the role of a potential caching point. Having the caching points distributed all over the network as a storage of data replicas eases the data delivery in highly dynamic environments where packet loss occurs frequently. Therefore, in our model we utilize the idea of caching the data over the delivery path in order to more easily retrieve the packets lost during transmission.

The work in [16] focuses on the impact of various mobility parameters on the behavior of a publish-subscribe system in a cellular network. The frequency of connection/disconnection requests, the available bandwidth and the size of the messages are some of the analyzed parameters. Additionally, the authors evaluate the performance of handoff solutions based on the publish-subscribe model. In highly dynamic wireless scenarios the broker based publish-subscribe middleware handoffs show significant performance degradation. Due to the prediction of the next possible location our model shows satisfying results even in frequent handover scenarios.

SocialCast [17] examines the publish-subscribe in a mobile environment from the perspective of social relationship between users. The authors use the idea of sporadically-connected human networks formed due to the existence of social relationships among users. A frequent collocation of socially related users is utilized in forwarding decisions in order to improve the routing fabric. The applied routing algorithm takes into account social associations between users in order to predict the future topology setup and the best information carrier. Our work incorporates prediction aspects as well. Instead of basing our prediction decisions on the social aspects of the networks, we employ the tracking mechanism of the most likely future location of moving nodes. While the tracking-based prediction functions have been applied in the cellular context in research literature before, the part of the novelty of our work is in its application to publish-subscribe information centric networking, i.e., directly into the topology management functionality [18]. We build the publish-subscribe based intelligent topology management module capable of performing a range of optimization operations for overcoming the mobility issues.

III. Model description

Our simulation scenarios are based on the ns-3 simulation tool and the WiMAX standard. We deploy a large number of base stations covering a metropolitan area and having numerous mobile stations in their range. Our simulation setup is built according to the distribution of base stations belonging to large mobile service providers in USA. The base stations are characterized by their positions and transmitting ranges. The mobile nodes are uniformly distributed over the coverage area of the base stations.

Data exchange between mobile stations is conveyed in a publish-subscribe manner. Upon the match between published and subscribed data occurs in the rendezvous, the topology management constructs the routing path from the publisher to the subscriber and signals it to the forwarding modules. The forwarding functions of the system are thereby instructed how to transfer the data. The forwarding modules are capable of caching the content traversing the network; thus the lost data can be easily obtained from the closest cache [19]. We have simulated dynamic scenarios, varying the speed and trajectory of the nodes. The speed of mobile nodes corresponds to a walking speed, speed of a vehicle in an urban environment and high-way speed. The mobility patterns are generated randomly for a given area covered by the set of base stations, in a way that ensures existence of handovers. The topology management stores the information about nodes' movement over time. Any regularity in the movement pattern such as daily repetition of the trajectory for commuting users is used as an input for prediction of the possible future location of the node. Therefore, the topology management is able to construct the backup path which indicates the routing path towards the node on a predicted location.

During the handover phase, while the node is leaving the coverage area of one base station and entering the new coverage area the data is sent to both base stations by bicasting. Thus, the data loss due to handover is minimized. The backup path creation and bicasting is triggered upon the signal power level P_{sig} of the mobile node received by the current base station drops below a certain level, i.e. the prediction triggering threshold. The prediction triggering threshold is slightly higher than the actual handover threshold T_{hand}, thus it serves as an indication of a possible transition into the coverage area of another base station. According to the speed of mobile nodes, the type of the applications, and the importance of the transmitted data, different triggering thresholds could be appropriate. This information about the network is stored in the topology management database. We aim at using this data available to the topology management for setting the most suitable triggering thresholds. Due to its awareness of nodes' movement patterns and optionally the requirements of existing applications with respect to the data transfer, the topology management appears as an appropriate candidate for instructing the system about the prediction triggering levels. Based on the current network state in the simulation scenario the topology manager sets the initial prediction triggering threshold. In order to optimize its value, the topology management calibrates the prediction triggering threshold empirically, by monitoring the system performance. During this process the topology management acquires the knowledge about the most suitable prediction triggering level T_{pred} for the given scenario and stores this information in its database. Algorithm 1 summarizes the applied prediction procedure.

Algorithm 1 Applying the prediction in mobile scenarios

while $T_{hand} < P_{sig}$ **do**
  if $P_{sig} < T_{pred}$ **then**
    Start the prediction process. Calculate the most likely future location of the mobile node and bicast the data.
  else
    Operate without the triggering of prediction process.
  end if
end while
Start the handover process.

Apart from location prediction mechanisms, the topology management utilizes the information-centric nature of the system, thus instructing the forwarding elements to always deliver the data over the shortest path. This implies changing the publisher from which the data is being retrieved, if during its movement a mobile node approaches another publisher of the same content. The topology manager thus changes the traffic flow in order to optimize the data delivery. We investigate the benefits of such content-oriented topology management equipped with prediction mechanisms for nodes' locations.

IV. Evaluation results

In the following, we examine the data-centric and prediction aspects of our model concurrently.

A. The benefit of publish-subscribe data-centric nature

The topology management regenerates the delivery paths as the topology changes, in order to always get the content from the closest location. As long as the number of publishers of the same content remains under a certain level, there is evident benefit in having the same data available at multiple locations. For a large number of publishers the frequent path recalculation and the signaling overhead due to the change of the data origin poses a considerable burden on the system, thus diminishing the benefits of the data-centric communication. Figure 1 illustrates the described impact of the number of publishers on the throughput of the system and the drawback of a large number of publishers.

B. The benefit of applied prediction algorithm

With respect to the result shown in Figure 1, our further elaboration of topology management in a dynamic environment is focused on lower density of publishers, i.e. up to 10, which is also the more realistic case in practice. In order to understand the benefits of the applied prediction model we perform a range of tests. The advantage of utilizing the prediction is strongly related to its reliability. According to the collected data about user movements, the prediction function can guess with variable certainty the user's next position. In the case of firmly settled regularities in a user's mobility, e.g., daily commuting between certain places, the prediction function will substantially contribute to performance improvement. On the other hand, if there is an insufficient number of movement samples in the database or if the user does not follow his

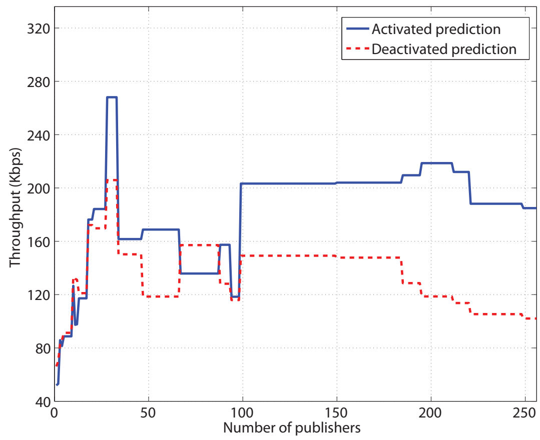

Fig. 1. Impact of the number of publishers on the user's throughput.

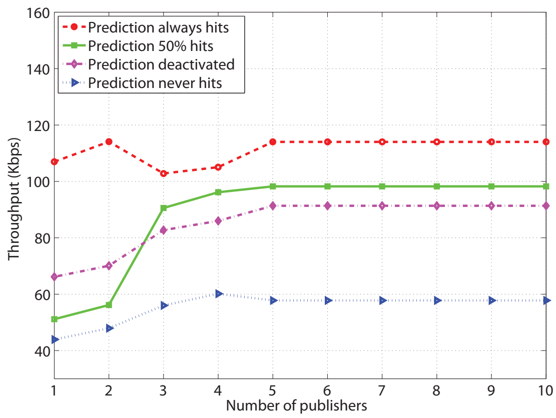

Fig. 2. Increase in throughput due to the prediction. The speed of mobile nodes is $100km/h$, prediction algorithm is activated at $T_{pred} = -89dBm$.

usual pattern, the use of the prediction function can even have negative consequences, placing an additional load on the system by generating unneeded traffic. Figure 2, Figure 3 and Figure 4 illustrate the improvements in system performance with respect to throughput for different accuracy levels of the prediction function and different network setups. It can be observed that the use of prediction, even accounting for 50% of errors in data bicasting, can lead to throughput improvement. However, extremely frequent wrong guesses of most likely future location can diminish the positive effect of prediction. This brings us to the conclusion that employing prediction is appropriate only under certain conditions and only with a guaranteed minimal accuracy.

C. The optimization of the prediction triggering level

Furthermore, we investigate the impact of the prediction triggering level on the performance of the prediction function. Setting the power threshold for activating the prediction is a critical operation. This parameter affects the quantity of

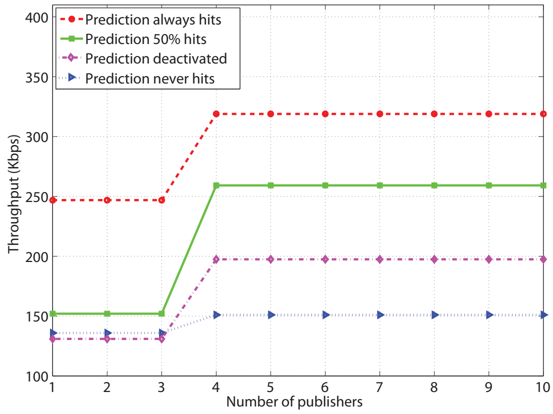

Fig. 3. Increase in throughput due to the prediction. The speed of mobile nodes is $50km/h$, prediction algorithm is activated at $T_{pred} = -87dBm$.

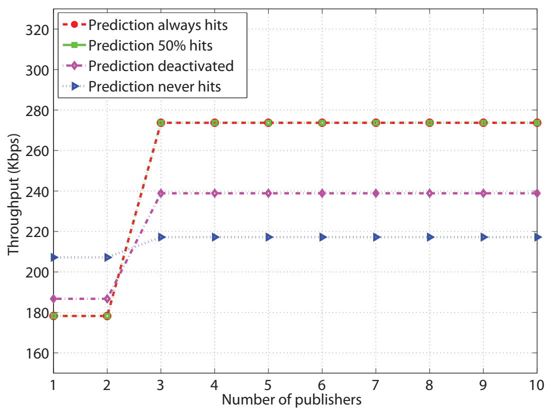

Fig. 4. Increase in throughput due to the prediction. The speed of mobile nodes is $10km/h$, prediction algorithm is activated at $T_{pred} = -85dBm$.

signaling traffic, hence an unsuitable triggering level can drastically reduce the benefits of prediction. The optimal prediction activation level depends on the current topology and a large set of network parameters. In our scenarios the topology management having information about network structure, speed and position of mobile users and constantly monitoring the system performance is used to dynamically optimize the triggering threshold. Figure 5 and Figure 6 illustrate the obtained throughput for different user speeds, for the case of an unoptimized prediction triggering threshold. It is evident that applying the prediction can have an extremely negative impact on the throughput, such that the deactivation of prediction leads to the best system performance. On the other hand, for the same network setting and the speed of mobile nodes the system shows much better performance using the optimized prediction triggering threshold as illustrated in Figure 2 and Figure 4.

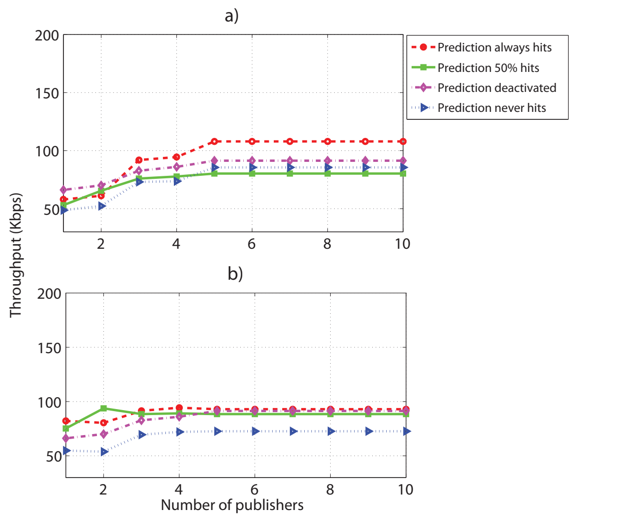

Fig. 5. Impact of unoptimized prediction triggering threshold on the throughput. The speed is $100km/h$. a) The triggering threshold is $-85dBm$ b) The triggering threshold is $-80dBm$

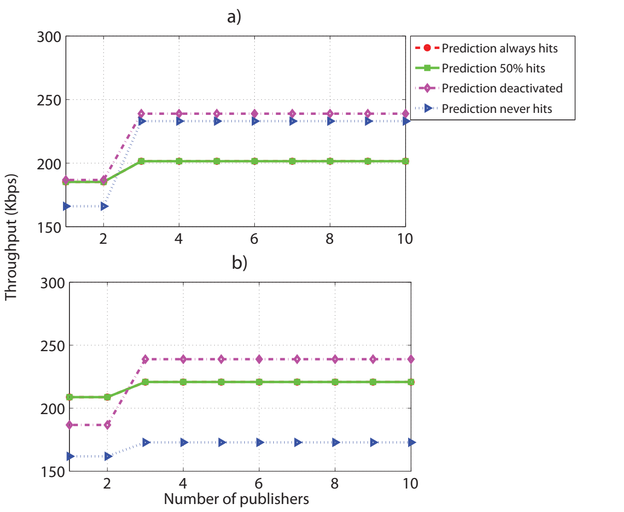

Fig. 6. Impact of unoptimized prediction triggering threshold on the throughput. The speed is $10km/h$. a) The triggering threshold is $-89dBm$ b) The triggering threshold is $-80dBm$

V. Conclusions

We have presented a model of publish-subscribe based topology management enriched by mobility prediction functionality. We have simulated a dynamic WiMAX network in which all nodes communicate using a publish-subscribe model. The movement pattern of nodes is tracked throughout time and used as an input to prediction function. Having the information about most probable future locations of nodes, the topology management can compute the backup forwarding paths for mobile nodes. Therefore, the system becomes more robust and more quickly responsive to sudden change in nodes' positions. Our evaluation results show that the publish-subscribe model equipped with intelligent location prediction can be a winning paradigm in mobile networks. However, for the optimal performance of such a system various parameters need to be adapted. The topology manager having the information about the network structure and nodes' mobility appears as the most natural candidate for performing such adaptation.

Acknowledgements

This work was in part financially supported by RWTH Aachen University through UMIC research centre and European Union (PURSUIT project).

References

[1] R. G. P.Th. Eugster, P. A. Felber and A.-M. Kermarrec, "The many faces of publish/subscribe," *ACM Computing Surveys*, vol. 35, pp. 114–131, June 2003.

[2] T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. Kim, S. Shenker, and I. Stoica, "A Data Oriented (and Beyond) Network Architecture," *In ACM SIGCOMM, Computer Communication Review*, vol. 37, no. 4, pp. 181–192, October 2007.

[3] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking Named Content," *Proceedings of the 5th International conference on emerging networking experiments and technologies ACM CoNEXT*, pp. 1–12, December 2009.

[4] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, and S.Shenker, "ROFL: Routing on Flat Labels," *In Proceeding of the ACM SIGCOMM workshop on future directions in network architecture*, vol. 36, no. 4, September 2006.

[5] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet indirection infrastructure," *Proceedings of the ACM SIGCOMM 2002. conference*, vol. 32, no. 4, pp. 73–86, October 2002.

[6] "PSIRP-Publish Subscribe Internet Routing Paradigm," http://www.psirp.org/, [visited on 02.12.2011.].

[7] "Named Data Networking (NDN) Project," http://www.named-data.net/, [visited on 02.12.2011.].

[8] "EU FP7 PURSUIT project," http://www.fp7-pursuit.eu/, [visited on 02.12.2011.].

[9] C. Perkins, " IP Mobility Support for IPv4," RFC 3344, August 2002.

[10] C. P. D. Johnson and J. Arkko, " Mobility Support in IPv6," RFC 3775, 2004.

[11] V. Jacobson, M. Mosko, D. Smetters, and J. J.Garcia-Luna-Aceves, "Content-centric networking," *Whitepaper describing future assurable global networks. Response to DARPA RFI SN07-12*, 2007.

[12] A. Carzaniga, D. Rosenblum, and A. Wolf, "Design and evaluation of awide-area event notification service," *ACM Transactions On Computer Systems*, vol. 19, pp. 332–383, 2001.

[13] G. Cugola and G. Picco, "REDS, A Reconfigurable Dispatching System," *In Proc. of 6th Inter. workshop on Software Engineering and Middleware*, pp. 9–16, 2006.

[14] V. Muthusamy, M. Petrovic, D. Gao, and H. Jacobsen, "Publisher mobility in distributed publish/subscribe systems," *In 25th IEEE International Conference on Distributed Computing Systems Workshops*, pp. 421–427, 2005.

[15] V. Sourlas, G. S. Paschos, P. Flegkas, and L. Tassiulas, "Mobility support through caching in content-based publish/subscribe networks," *In 5th International Workshop on Content Delivery Networks (CDN 2010) in conjuction with 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2010)*, May 2010.

[16] U. Farooq, E. W. Parsons, and S. Majumdar, "Performance of Publish/Subscribe Middleware in Mobile Wireless Networks," *ACM SIGSOFT Software Engineering Notes*, vol. 29, no. 1, January 2004.

[17] P. Costa, C. Mascolo, M. Musolesi, and G.P. Picco, "Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 5, pp. 748–760, June 2008.

[18] B. Gajic, J. Riihijärvi, and P. Mähönen, "Intra-Domain Topology Manager for Publish-Subscribe Networks ," *18th International Conference on Telecommunications ICT 2011*, May 2011.

[19] S. Arianfar, P. Nikander, and J. Ott, "Packet-level caching for information-centric networking," *Finnish ICT-SHOK Future Internet Project, Technical Report*, 2010.