Considering Capacities in Application-Layer Multicast over Shared Access Networks

Christian Hübsch
Karlsruhe Institute of Technology, Germany
huebsch@kit.edu

I. Introduction

A huge number of application scenarios would greatly benefit from a group communication service in the Internet, but the existing IP Multicast solution still lacks global availability due to several technical and political issues [1]. As a consequence, Application-Layer Multicast (ALM) [2] arose as a promising solution to overcome this missing global multicast deployment. In ALM, forwarding functionality and membership state is exclusively handled on the participating end systems, following the well-known peer-to-peer paradigm. A typical ALM node has to handle higher traffic load than compared with native IP Multicast because packets are duplicated and forwarded on the nodes rather than in infrastructure routers. Therefore, ALM approaches were initially limited to stationary (comparably high-capacity) end systems that could cope with this higher load. However, recent and ongoing developments in the sectors of access technologies (e.g. VDSL, LTE etc.) and end system capabilities lead to scenarios in which employing ALM on mobile systems becomes feasible. Soon, these developments will render two-digit MBit up- and downstream rates possible on every mobile device. This—together with always-on possibilities at an affordable price— will lead to an intense growth in the number of users of such devices. Additionally, modern end system devices are able to communicate via more than one technology, potentially concurrently (compare Figure 1). Examples are cellular networks, WiFi, or even PANs like Bluetooth. We refer to these devices as multi-modal end systems, which we believe to hold great potential for ALM dissemination. At the same time, consumed contents will also grow in order to use the higher technical capabilities of the devices. Although individual capacities grow, some access technologies used are bound to upper capacity limits (e.g. in cellular network cells). Recently in the US, first service providers suffered from service outages triggered by too many iPhone users congesting the shared medium[1]. Although not clear if these congestions occured in the backbone or the access domain, the problem will be amplified in ALM scenarios where data is duplicated in end systems rather than infrastructure routers. Consider a videostreaming scenario, already being one of the most bandwidth-intense applications in the Internet. Since cellular networks will be the technology with the highest availability for mobile users, assume that a number of nodes in these access networks want to receive a single-source videostream (rooted on one of the nodes). With high numbers of participating nodes the involved access domain cells have to face a high load of forwarding the incoming and outgoing video streams. If cell capacity limits are reached, nodes will not be able to participate in the dissemination process any more. In our work we focus on how to preserve such congestions in ALM dissemination scenarios. The solution space is two-fold, such that the problem could be tackled by either consider current capacity usages in the involved shared domains or by balancing data dissemination between orthogonal communication techniques concurrently. In our work we look at a combination of both approaches in order to find a dissemination tree that preserves the access networks from congestions while still offering good service quality with respect to capacity bounds on individual nodes or the overall dissemination delays.

[1] http://www.nytimes.com/2009/09/03/technology/companies/03att.html

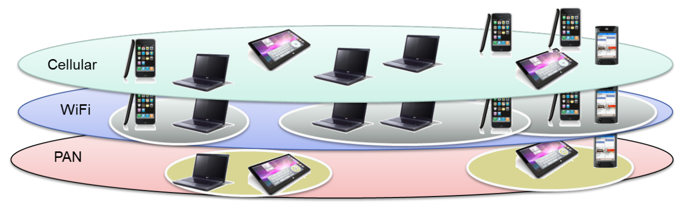

Fig. 1. Example for Access Domains and multiple communication possibilities

II. Approach

Although a multitude of ALM protocols exist that consider network topology in order to reduce or balance traffic, none of them solves the described congestion issues. Existing work mostly focuses on inter-ISP traffic reduction [4] or caching techniques [3] rather than considering shared domain capacities. In our presentation we describe an approach that trys to preserve shared media from congestion in an ALM scenario. Doing so, we also take into account that modern end system like mobile phones are already able to use diverse access technologies, possibly concurrently. This renders even better load-balancing strategies in the considered scenario possible. To explicitly consider capacities in shared domains we introduce the term *Access Domain* to be a subset of all participating nodes that shares the same medium for transmitting data. Examples could be the mentioned cellular network cells, but also more limited areas–like WiFi domains or even PANs (e.g. Bluetooth) (compare Fig.1). We assume that nodes are able to identify the *Access Domains* they reside in (in case of WiFi this may be figured out by hashing the MAC address of the access point and the SSID, for instance).

The scenario can be described as a common graph model which has to be extended by the concept of *Access Domains* and several cost metrics to be considered with respect to ALM data distribution. We focus on the case of single-source videostreaming to a high number of receivers. For sake of simplicity, we initially limit ourselves to a static and linear model without node mobility and with fixed (but diverse) capacities per node and per *Access Domain*. Also, we do not consider cross-traffic issues or interference between different *Access Domains*.

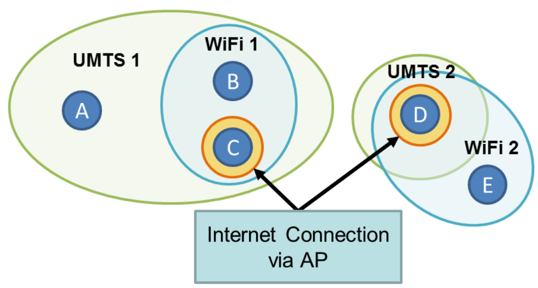

(a) Example Access Domains

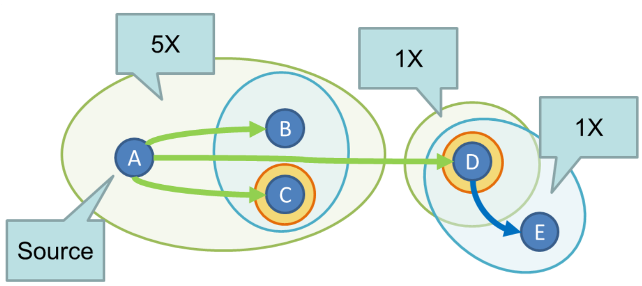

(b) Naive approach via straightforward connectivity

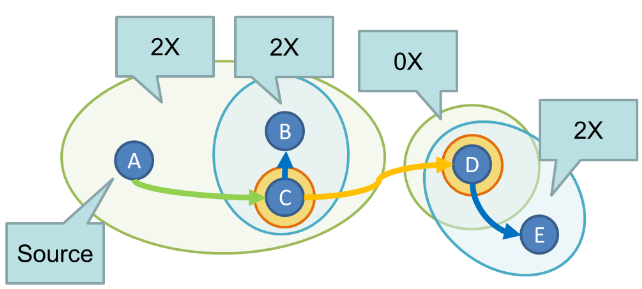

(c) Consideration of load and multi-modality

Fig. 2. Example Access Domain Scenario

Figure 2 shows a simple scenario, consisting of 5 nodes with partially different communication possibilities. Figure 2(a) illustrates the overall setting. Here, node A may communicate via UMTS. Same applies to node B, C and D, which may also use WiFi technology. Node E is limited to WiFi only. An important point is that nodes C and D are also able to reach other nodes through their respective WiFi access point which routes them to the Internet, which is not the case for node B and E, being limited to their local WiFi domain. Given a straightforward *Access Domain*-agnostic approach as shown in Figure 2(b) may lead to source node A sending the videostream to all nodes it can reach via UMTS directly, followed by a forwarding from node D to E. Figure 2(b) sketches the resulting medium occupancy per *Access Domain* in this case. It is clear that the UMTS *Access Domain* in which the source node A resides is subject to a higher traffic load compared to the other *Access Domains*. Figure 2(c), in contrast, illustrates an alternative forwarding strategy that explicitly considers load in the *Access Domain*, trying to find a more balanced solution. Source node A forwards the stream only to node C, which passes the stream to node B via WiFi and to node D via its access point Internet connection. Like before, node D forwards the stream to node E via WiFi (for there is no alternative path here). The occupancies show a much more balanced characteristic, although *Access Domains* are used now that haven't been involved before. Therefore, capacity load peaks may be shifted and evened using alternative communication paths.

In our approach, basic connectivities in the scenario are discovered through the *ariba* underlay substrate that is part of the SpoVNet architecture [5] [6]. We plan to develop a solution as described here in the context of our existing group communication service in SpoVNet. Finding an optimal dissemination tree with several metrics to consider is known to be NP-hard [7], therefore we work out strategies to find a near-to-optimal solution. We present first evaluations of the potential of the approach, our current work in this context and our vision for the future.

Acknowledgments

This work was partially funded as part of the *Spontaneous Virtual Networks* (*SpoVNet*) project by the Landesstiftung Baden-Württemberg within the BW-FIT program.

References

[1] C. Diot and B. N. Levine and B. Lyles and H. Kassem and D. Balensiefen, *Deployment issues for the IP multicast service and architecture*, IEEE Network, Vol. 14, No. 1, Pages 78–88, Feb 2000.

[2] M. Hosseini and D. T. Ahmed and S. Shirmohammadi and N. D. Georganas, *A Survey of Application-Layer Multicast Protocols*, Communications Surveys & Tutorials, IEEE, Vol. 9, Pages 58–74,No. 3, July 2007.

[3] G. Shen and Y. Wang and Y. Xiong and B. Y. Zhao and Z.-L. Zhang, *HPTP: Relieving the tension between ISPs and P2P*, IPTPS, Pages 423–432, Feb. 2007.

[4] F. Picconi and L. Massoulie, *ISP Friend or Foe? Making P2P Live Streaming ISP-Aware*, ICDCS '09: 29th IEEE International Conference onDistributed Computing Systems, Pages 413–422, 2009.

[5] C. Hübsch and C. P. Mayer and S. Mies and R. Bless and O. P. Waldhorst and M. Zitterbart, *Reconnecting the Internet with ariba: Self-Organizing Provisioning of End-to-End Connectivity in Heterogeneous Networks*, ACM SIGCOMM Computer Communication Review, Vol. 40, No. 1, Pages 131–132, Jan 2010.

[6] C. Hübsch and C. Mayer and S. Mies and R. Bless and O. Waldhorst and M. Zitterbart, *Demo Abstract: Using Legacy Applications in Future Heterogeneous Networks with ariba*, Proceedings of IEEE INFOCOM, San Diego, May 2010.

[7] Z. Wang and J. Crowcroft, *Quality-of-Service Routing for Supporting Multimedia Applications*, IEEE Journal of Selected Areas in Communications, Vol. 14, No. 7, Pages 1228–1234, 1996.