

Towards an Adaptive Publish/Subscribe Approach Supporting Transitions

Björn Richerzhagen* and Ralf Steinmetz

Multimedia Communications Lab, TU Darmstadt, Germany
{richerzhagen,steinmetz}@kom.tu-darmstadt.de
<http://www.kom.tu-darmstadt.de>

Communication in large-scale distributed systems is mostly based on the publish/subscribe (pub/sub) paradigm, where content creators and interested recipients are decoupled in space, time, and synchronization. With the rise of social applications and the wide distribution of mobile devices highly dynamic scenarios for pub/sub systems become increasingly interesting, requiring the system to be adaptive in order to react to environmental conditions. This paper identifies promising adaptations in pub/sub systems and proposes the novel concept of transitions between multiple pub/sub mechanisms. Transitions enable a pub/sub system to react to a wide range of conditions by switching to the mechanism that fits the current environmental conditions best. This concept leads to a range of interesting research questions, which are presented and addressed in this paper.

1 Introduction

The publish/subscribe (pub/sub) paradigm is the foundation for many applications deployed in the Internet, ranging from full-fledged commercial middleware to conceptual systems presented in the academic world [1]. Prominent applications such as Facebook, Twitter, and Google+ distribute millions of – mostly user-created – content items to interested subscribers, making use of the pub/sub paradigm. At the same time, those platforms provide their services through dedicated applications to a wide range of end user devices, ranging from smartphones over tablets to desktop-PCs. Each of those devices has its own restrictions regarding computational power, energy supply, available network access technologies, and mobility. A pub/sub system needs to take those restrictions into account when forwarding content items and maintaining subscriptions.

Considering, e.g., mobility there exist a range of pub/sub systems that allow publishers and subscribers to be logically and physically mobile when connecting to a fixed broker network [6]. Those solutions, however, do not support mobile brokers, which are essential in the case of pure ad hoc connectivity between devices. Peer-to-peer-based solutions have been proposed for fully decentralized applications on a global scale [2], [3]. Those systems have shown not to be able to deal with higher peer mobility, as they rely on structured overlay topologies.

* This work has been funded by the DFG as part of the CRC 1053 MAKI.

Systems that are specifically designed to support mobile brokers are mostly limited to a local scale and operate as an isolated application [4], [5].

Each of the aforementioned approaches provides a viable solution to a limited set of environmental conditions, for example regarding peer mobility or scale of the overlay network. However, in a real-world application scenario, these requirements can vary significantly, rendering isolated solutions ineffective.

This paper proposes a novel approach towards a highly adaptive pub/sub system that supports heterogeneous devices as well as locality and mobility of publishers, subscribers, and brokers by allowing transitions between different pub/sub mechanisms. The goal is not to reinvent pub/sub, but to provide a methodology that allows existing mechanisms to be adapted and exchanged depending on the current – possibly local – situation, and to understand the impact of different such adaptations on the system. The novel approach of seamless transitions between pub/sub mechanisms rises a number of research questions:

- Between which pub/sub mechanisms can and should transitions occur?
- What are the costs and benefits of these transitions?
- How, when, and by whom are transitions triggered?
- How is a stable system behavior ensured during and after a transition?

In order to tackle these questions, key transitions in a pub/sub system are identified and thoughts on their realization are given in Section 2. A conclusion and first results are presented in Section 3.

2 Transitions in Publish/Subscribe

As stated in the research questions, key transitions are to be identified and studied regarding their cost and benefits. Promising candidates are presented in this section, together with initial thoughts on their realization.

Decentralized vs. Centralized Broker Networks. Current pub/sub broker networks are either decentralized as a peer-to-peer network or to some extent centralized, for example only on trusted machines in a cloud. For most applications, the load on the system can change significantly, depending on the time of day or on events in the real world. Therefore, adaptation possibilities between centralized broker networks and peer-to-peer approaches are to be investigated. A system might provide its service in the centralized manner as long as the workload remains below a certain threshold. As soon as this threshold is exceeded, the system adds additional peers to the broker network in a peer-to-peer manner. Finally, the whole system could adapt into a peer-to-peer overlay, thereby lowering the load on the service provider's infrastructure. A locally bounded topology transition could be of great interest in the dense presence of mobile devices. Such devices could then form an ad hoc overlay that deals especially well with mobile peers to relay messages [7]. One promising approach regarding the construction of overlay networks with arbitrary neighborhood properties is T-Man, proposed by Jelasity et al. [8]. The author's work could serve as a starting point for an abstraction of overlay topology transitions.

Channel-based, Topic-based, or Content-based. The power of the pub/sub paradigm comes from its flexibility regarding subscription and publication semantics. There is a trade-off between filtering overhead for the broker network and expressiveness of the messages, which manifests in three categories of systems: channel-based, topic-based, and content-based. Currently, application designers have to decide which category fits their scenario best. The system itself, however, should provide the abstraction and allow for a flexible filter language to support a wide range of applications. Internally, the system has to find a state that provides the best trade-off between message complexity and computational requirements. This leads to transitions between different kinds of content models with the goal of overall complexity reduction, depending on the current workload and the nature of subscriptions.

Exploiting Locality of Content and Interest. An emerging class of applications centers around location-based services. In such applications, interest in content is based on the physical position of a user. Exploiting this property has proven beneficial for the performance of distributed applications, as they are able to move content closer towards the user. There exist a range of spatial pub/sub solutions that address content based on its (logical) position [9]. These systems, however, are not intertwined with a global pub/sub solution and operate in a rather confined environment. Therefore, in our scenario, the system should *locally* adapt towards content-based pub/sub if, e.g., ad hoc communication can be utilized. In such a system, nearby users can benefit from each other if they have related interest. They can, for example, reduce their usage of the costly mobile network infrastructure by exchanging information locally in a peer-to-peer fashion. As a possible side effect, communication between nearby users can occur faster than via a centralized broker network, albeit the communication link might be less reliable. The resulting pub/sub system would, thus, be able to meet near real-time constraints on a local scale, while still being able to distribute messages to interested subscribers all over the world.

3 Preliminary Conclusions and Future Work

Current pub/sub systems are isolated applications specifically designed for a set of environmental conditions, based on a rather confined application scenario. In this work, we motivate the concept of transitions between different, well-known mechanisms in order to obtain a highly adaptive pub/sub system. Examples of such mechanisms and the resulting transitions are discussed, as well as potential starting points regarding their realization.

As presented, there exists a range of interesting research questions as to how pub/sub systems can support adaptivity through transitions. Starting with the mechanisms identified above, the costs and resulting benefits of transitions are to be studied in well-defined scenarios. To study the impact of locality and the decentralization of brokers as described beforehand, the scenario of a social

application in a mobile ad hoc network is chosen. Additional to ad hoc communication, devices can use the cellular infrastructure, allowing for a detailed study of the resulting trade-offs. Furthermore, transitions between centralized brokers, i.e., brokers reachable via the cellular link, and brokers distributed throughout the ad hoc network can be studied. One important research question in this context is to how, when, and by whom such transitions should be triggered. Answering this question requires a deep understanding as to how transitions affect the system and how stability of the overall system before, during, and after a transition can be guaranteed. In an earlier work the authors studied the impact of transitions between scheduling mechanisms and overlay topologies in the context of P2P video streaming systems [10]. The evaluation shows that transitions between mechanisms enable the system to work under a wider range of environmental conditions and retain desirable properties even in extreme scenarios.

The scenario of a mobile social application as described above needs to be modeled, and the impact of individual mechanisms and the transitions inbetween is to be evaluated. This should lead to a better understanding on when, how, and by whom transitions should be executed to maintain a high service quality. While initial results are obtained by means of simulation, we later on plan to evaluate the system prototypically.

References

1. Eugster, P.Th., Felber, P.A., Guerraoui, R., Kermarrec, A.: The many Faces of Publish/Subscribe. *ACM Comput. Surv.* 35, 114–131 (2003)
2. Baldoni, R., Marchetti, C., Virgillito, A., Vitenberg, R.: Content-Based Publish-Subscribe over Structured Overlay Networks. In: *IEEE ICDCS* (2005)
3. Bender, M., Michel, S., Parkitny, S., Weikum, G.: A Comparative Study of Pub/Sub Methods in Structured P2P Networks. *LNCS* 4125, 385–396 (2007)
4. Meier, R., Cahill, V.: On Event-Based Middleware for Location-Aware Mobile Applications. *IEEE Trans. on Softw. Eng.* 36, 409–430 (2010)
5. Pongthawornkamol, T., Nahrstedt, K., Guijun W.: The Analysis of Pub/Subs Systems over Mobile Wireless Ad Hoc Networks. In: *IEEE MobiQuitous* (2007)
6. Sourlas, V., Paschos, G.S., Flegkas, P., Tassiulas, L.: Mobility Support Through Caching in Content-Based Pub/Sub Networks. In: *IEEE/ACM CCGrid* (2010)
7. Friedman, R., Kaplun Shulman, A.: A Density-driven Publish Subscribe Service for Mobile Ad Hoc Networks. *Ad Hoc Networks* 11, 522–540 (2012)
8. Jelasity, M., Montresor, A., Babaoglu, O.: T-Man: Gossip-based fast Overlay Topology Construction. *Computer Networks* 53, 2321–2339 (2009)
9. Chen, X., Chen, Y., Rao, F.: An Efficient Spatial Publish/Subscribe System for Intelligent Location-based Services. In: *ACM DEBS* (2003)
10. Richerzhagen, B.: Supporting Transitions in P2P Video Streaming. Master Thesis, www.ps.tu-darmstadt.de/fileadmin/publications/Ric12.pdf (2012)