Delay and Disruption-Tolerant Networking (DTN) II

Dagstuhl Seminar 09071 – Executive Summary

Kevin Fall[1], Cecilia Mascolo[2], Jörg Ott[3] and Lars Wolf[4]

[1] Intel Research Berkeley
2150 Shattuck Ave #1300, Berkeley, CA 94704, USA
kfall@intel.com
http://www.cs.berkeley.edu/~kfall/

[2] University of Cambridge, Computer Laboratory
15 JJ Thomson Avenue, Cambridge, CB3 0FD United Kingdom
cecilia.mascolo@cl.cam.ac.uk
http://www.cl.cam.ac.uk/users/cm542

[3] Helsinki University of Technology, Dept. for Communications and Networking
P.O. Box 3000, FIN-02015, Helsinki, Finland
joerg.ott@tkk.fi
http://www.netlab.tkk.fi/~jo/

[4] TU Braunschweig, Institut für Betriebssysteme und Rechnerverbund
Mühlenpfordtstr. 23, 38106 Braunschweig, Germany
wolf@ibr.cs.tu-bs.de
http://www.ibr.cs.tu-bs.de/

1 Introduction

Today's Internet architecture and protocols, while perfectly suitable for well-connected users, may easily experience serious performance degradation and entirely stop working in more challenged networking environments. Such environments are manifold, ranging from mobile users experiencing occasional or frequent disconnections to communication services for remote areas, to vehicular network communication in large areas, sensor networks to habitat or wildlife monitoring, and to space and underwater communications. These scenarios all share two commonalities: that an end-to-end path between two communicating nodes may not exist at any single point in time and that communication delay may be significant. Luckily, in most cases, delay in the delivery of the data can be tolerated. However, with the continued expansion of the Internet into new areas and the increasing penetration of communication technologies into more areas of life and technology, these environments become commonplace and are no longer restricted to exotic sensing applications but are quickly becoming relevant to consumers in everyday life.

Many attempts over recent years of incrementally fixing the Internet protocols in a bottom up fashion have only achieved partial successes: while mobile IP, HIP, transport, session, and cross-layer approaches may support changes of network attachments and short-term disconnections, a more fundamental approach is needed to ad-

Dagstuhl Seminar Proceedings 09071
Delay and Disruption-Tolerant Networking (DTN) II
http://drops.dagstuhl.de/opus/volltexte/2009/2357

dress networking environments in which delays and disconnections may last for significant periods of time, and are the rule rather than the exception.

Delay-tolerant Networking (DTN) has taken a more encompassing approach to dealing with virtually all types of connectivity challenges, from bit rate to errors to delays to disruptions. By providing a novel communication abstraction that relies exclusively on asynchronous hop-by-hop message passing with no need for instant end-to-end connectivity, DTN concepts enable communications even under adverse conditions. This comes, however, at the cost of interactivity of communications, rendering any kind state synchronization or validation more difficult and raising new challenges. These include *routing protocols* – that need to operate under often unknown future conditions, *security mechanisms* – that can no longer carry out instant key derivation or validation even if a security infrastructure was in place, and *application protocols and paradigms* – that can no longer rely on simple lower layer abstractions promising (mostly) instant and reliable interactions.

2 Seminar Topics

For the Dagstuhl seminar DTN II, the organizers identified the aforementioned three aspects as key elements towards successful future deployment of DTN technologies, drawing on the – exploding – research activity and publications in this relatively young field. The Dagstuhl seminar brought together researchers working on applying DTN concepts to different "domains", from space communication to sensor networks to mobile and social networks, joining theoretical fundamental research, experimental analytical studies, and practical application design, implementation, and trials. It is commonly understood that only an encompassing coverage of all aspects from theory to practice, well-founded on understanding the true characteristics of the target deployment environments and accompanied by joint engineering, implementation, and standardization efforts can lead to an impact in the real-world (that other promising technologies such as MANETs failed to achieve). To this end, this Dagstuhl seminar offered a venue for intense exchange of ideas between academia and industry.

The Dagstuhl seminar was organized around short presentations given by the participants summarized in this section. They were complemented by group work addressing issues that were brought up by the talks and discussions of the first day and developing these further towards identifying potential research questions and directions in plenary sessions (next section). The talks covered the guiding themes of the seminar, but their contents is best classified into four topics: 1) social interactions and networks, 2) testbeds, 3) applications, and 4) general protocol design.

Recent research in DTNs has increasingly emphasized communication between humans using their mobile devices. This requires, among many other aspects, understanding behavioral, contact and interaction patterns among humans, so that DTN research has increasingly interacted with (tools from) social sciences. The talks by Eiko Yoneki addressed the issue of analyzing contact information and inferring social structures. Augustin Chaintreau investigated creating "navigable" DTN network graphs from local knowledge about the social environment and Mirco Musolesi's presentation complemented this by their work connectivity models based upon social

interactions for validating DTN protocols. John Solis exploited social interactions for message confidentiality in DTNs.

Besides real-world analysis and inference, practical deployments and testbeds received significant attention. They are a result of joint international activities (most prominently of the *Delay-tolerant Networking Research Group (DTNRG)* of the *Internet Research Task Force (IRTF)*, of international research projects (such as the EC FP7 N4C project), national activities, e.g., in the US and in Japan, and institutional setups (e.g., from NASA). Especially the talks by Avri Doria and Stephen Farrell focused on various testbeds and the discussion showed strong interest by the community in engaging in such kinds of activities, recognizing testbeds as an important element of practical research. Scott Burleigh presented the JPL DTN testbed for space communication, leveraging the *Deep Impact* experiment. Testbeds rely on interoperable implementations, for which Johannes Morgenroth reported on a RFC5050-compliant implementation of the DTN bundle protocol for embedded systems developed at TU Braunschweig IBR.

Application aspects were covered by a number of presenters who mostly emphasized on content distribution and retrieval within mobile variants of social networks, possibly fed by infrastructure nodes. The talk by Mikko Pitkänen and Gunnar Karlsson addressed rather local content search and content distribution, respectively; related to these thoughts, Michael Doering presented idea for city-wide DTN-based information distribution using various elements of public transportation systems. Janico Greifenberg spoke about wide-area dissemination using satellites and mobile vehicular infrastructure for publish-subscribe style dissemination.

Finally, general protocol design aspects for delay-tolerant networking were addressed: Philip Ginzboorg discussed the questionable(?) value of protocol optimizations, N. Asokan elaborated on resource management in DTNs, and Vania Conan presented considerations on the capacity of DTNs and its increase with growing delay tolerance.

3 Research Directions

The discussion of the aforementioned topics yielded a large set discussion items and numerous associated research questions that were summarized into three group work topics, reflecting different views of the (immediate) future of delay-tolerant networking:

1) The *pessimistic* viewpoint asks if DTN technologies will matter at all and find solid business cases and thus corporate interest – or share the fate of past promising technologies. One important option for a future developed could be niche markets, as presently exists in space and underwater communications as well as in wildlife monitoring and in military communications. Those cases, however, may easily remain domain-specific and lead to a fragmentation of DTN efforts to match the respective specialized needs, without much vision towards a unified communication platform for challenged environments. While the technology would matter, its overall relevance would be quite limited. Obstacles to a broad uptake are seen in the ever-increasing performance of infrastructure-based mobile (cellular) networks and the lack of incen-

tives for operators to support opportunistic communications as an integrated part of mobile communication solutions. Users would not want to put up with segmented solutions and are – to quite some extent – used to and simply expect instant interactions, making the perceived quality of experience of DTN-based applications too poor. On the other hand a discussion about complementing backbone access with opportunistic links to improve bulky content delivery in dense (city) environments, for example, was deemed interesting and work on the integration of heterogeneous networks was solicited.

2) The group taking the *optimistic* viewpoint assumed that DTN communication is feasible for a broader scope, as was suggested by various authors and put forward explicitly by Scott Burleigh in this seminar, echoed by Kevin Fall and Jörg Ott. One important element of the reasoning is that, while messages may be large, delays significant, and connectivity may not exist, the DTN protocols will work perfectly fine in a well-connected environment, too, and could enable efficient synchronous communication for all known Internet applications. DTN concepts add robustness to communication where it was missing before and could thus make important contributions to the design of a *Future Internet*. Following this thinking, the key questions discussed revolved around a future standard implementation and API: will there be *the* reference DTN stack and the *socket*-equivalent API? How would adaptive applications be designed that would help preserving the self-regulating nature that we find in the Internet today (most notably for TCP congestion control)? And how will a high-speed DTN be designed? Further discussion revolved about adapting the concepts and semantics of current Internet applications (that are more real time in nature) to DTNs: e.g., what is the meaning of "presence" in social network applications in DTNs and how do these and further application concepts get extended to embrace delay-tolerance as well.

3) Finally, the *pragmatist* group took up identifying the impending research and engineering problems that need solving, no matter whether DTN becomes a niche or a mainstream market. The three most pressing practical problems identified *naming and addressing* (and the interaction with *routing*), *network management*, and *benchmarking* as the three most pressing issues for practical deployments. Inside the DTN community, various naming and addressing schemes were tried out, from hierarchical naming (that could match DTN topology) to flat name spaces to semantics-/content-based naming. Multiple name spaces may need to co-exist in overlapping networks (and, e.g. using metadata, even inside the same message) and, similarly, support for multiple routing schemes may be required (possibly to be selected by a name). Due to the potentially disconnected nature of the networks and thus the inability to perform instant lookups or translations, it was felt that, generally, more explicit information will need to be embedded inside DTN messages. For network management, common logging and tracing facilities to follow and diagnose network operation are the most imminent needs. Specific challenges arise from the asynchronous operation, leading to potentially delayed information about network status. Configuration and also remote software updates are considered the next two important elements to be addressed, both of them raising significant questions concerning securing the interactions and authorizing operations. Finally, for benchmarking, it was found that a one-size-fits-all solution is unlikely to exist due to the quite diverse requirements: for example, resource cost and reliability demands in space communications differ vastly

from human-to-human interaction using mobile phones. A first step towards addressing benchmarking will be to develop reference scenarios and devise commonly accepted (domain-specific) utility functions for measuring performance.

Overall, the Dagstuhl seminar DTN II has provided the participants with a forum for fruitful discussion of present and future work on emerging networking applications and paradigms. The seminar has contributed to furthering the understanding of the perspectives of future development and real-world deployments of delay-tolerant networking as well as helped identifying issues – as research and engineering directions – to be resolved on this way.