

Disruption Tolerant Networking

Dagstuhl Seminar 05142 – Executive Summary

Marcus Brunner¹, Lars Eggert¹, Kevin Fall², Jörg Ott³ and Lars Wolf⁴

¹ NEC Europe Ltd., Network Laboratories
Kurfürstenanlage 36, 69115 Heidelberg, Germany
{brunner|eggert}@netlab.nec.de
<http://www.netlab.nec.de/>

² Intel Research Berkeley
2150 Shattuck Ave #1300, Berkeley, CA 94704, USA
kfall@intel.com
<http://www.cs.berkeley.edu/~kfall/>

³ Helsinki University of Technology, Dept. for Electrical & Communications Engineering
P.O. Box 3000, FIN-02015, Helsinki, Finland
jo@netlab.hut.fi
<http://www.netlab.hut.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

During the past 40+ years, numerous architectures were developed for network communication, including the ISO OSI reference model and its related protocol specifications and – of course – the Internet architecture. These network architectures have been designed with some implicit assumptions about specific target applications and deployment scenarios. Among the most important assumptions are specific characteristics of the underlying network (= link layer) technologies, such as relatively short transmission delays, low error probability and the existence of end-to-end paths.

In certain scenarios, these assumptions no longer hold. Examples of such scenarios include networks with frequent connectivity disruptions, extremely long transmission delays or unstable and variable connectivity. Consequently, the existing network architectures fail to support efficient communication in these scenarios, resulting in either significant inefficiencies or complete loss of connectivity.

Disruption Tolerant Networking (DTN) is a new area of research in the field of networking that deals with extending existing protocols or inventing new ones in a coordinated, architecturally clean fashion, to improve network communication when connectivity is periodic, intermittent, and/or prone to disruptions.

Among the challenges of this field of research are potentially large transmission delays. These may result either from physical link properties or from extended periods of network partitioning. A second challenge is efficient routing in the presence of

frequently disconnected, pre-scheduled, or opportunistic link availability. In some cases, an end-to-end path may not even exist at any single point in time. From a mobility perspective, DTN relaxes the “always on” paradigm, which would be extremely costly or even impossible to realize in challenged environments. A third challenge is that high link-error rates make end-to-end reliability difficult. Finally, heterogeneous underlying network technologies (including non-IP-based internetworks) with very different communication characteristics may need to be embraced.

These challenges can decrease the reliability and performance of communications at essentially all layers of the protocol stack, ranging from packet-based forwarding and routing, to reliability and other features provided at the transport layer, to the application protocols (and applications) themselves. The possibly resulting high transmission delays, errors rates, and the lack of an end-to-end path require different approaches to application interactions, reliability and security mechanisms. In addition, traditional mobility approaches may have to be revisited to accommodate users in networking environments prone to connectivity disruptions.

2 Seminar Topics

Numerous research activities over the past three years have focused on various facets of communications in challenged environments. Architectural concepts have been devised, prototype implementations were developed and research results are available from analysis, simulations and real-world experiments. The *Dagstuhl* seminar brought together researchers working in otherwise at least partly disjoint areas and established an intense dialogue across the variety of application domains.

A key realization of the seminar was that most participants mainly worked in the scope of one of two general areas of disruption tolerant networks. One group of participants is investigating solutions for networks with extremely long communication delays, such as the Delay-Tolerant Networking Architecture investigated within the Delay-Tolerant Research Group (DTNRG) in the Internet Research Task Force (IRTF). These new approaches often build on the paradigm of asynchronous interactions and introduce additional inter-internetworking layers that spawn multiple, specialized, internetworks of different characteristics.

Into this first group of presentations fell Bengt Ahlgren’s talk on applications for asynchronous networking, Ben Hui *et al.*’s talk on “pocket-switched” networks and Srinivasan Keshav’s extensions to the DTNRG architecture.

A second group of participants focused on approaches for improving Internet-based communication in scenarios where connectivity disruptions are frequent. Although this case can be generalized to the former – a connectivity disruption can be seen as a long communication delay – the dynamic change from short to long communication delays when a disruption occurs deserves special consideration, especially because communication efficiency should remain close to current Internet levels when connectivity is present and delays are short. This second group of presentations included Marc Bechler’s and Holger Füßler’s different protocol modifications for vehicular *ad hoc* networks, Simon Schütz’s TCP modifications for disrupted access links, Carsten Bormann *et al.*’s talk on “near end” DTN solutions, Aaron Falk’s pres-

entation on military satellite communication, Dirk Kutscher *et al.*'s talk on "Drive-thru" Internet access and Lavy Libman's disruption prediction for public transportation,.

Finally, a third group of presentations was orthogonal to this division. They include Per Gunningberg's experimental testbed for DTN solutions and Hannes Tschofenig's security considerations for DTN.

3 Research Questions

One major open question is whether the two areas of work described above – research related to DTNRG's bundle-based long-delay architecture on one hand and modifications to extend current Internet protocols for disruptive environments on the other hand – are in the end similar enough to be pursued within a combined effort. Although there are significant overlaps in mechanisms and approaches, the base characteristics and supported applications differ. Whereas the DTNRG and its related work focus on scenarios that may never support interactive applications, the main focus of the latter approaches is to improve the operation of existing, interactive Internet applications and protocols in situations where network connectivity is intermittent.

One challenge with improving Internet functionality is that the design options are limited, because of the need to remain compatible with deployed infrastructure and applications. The DTNRG architecture, on the other hand, has no such restrictions, because legacy applications are simply unsupported. Consequently, the designers of DTNRG-related mechanisms are free to evaluate and adopt name/locator, security or signaling mechanisms that cannot be used in approaches that extend traditional Internet protocols due to compatibility problems.

Specific questions arose around the issues surrounding whether modifications of only TCP (called various "TCP hacks") are an adequate solution to a subset of the problems of DTN. Although this may be useful for some situations, extended outages (including those that may span a system reset) will probably not be adequately addressed solely based on such modifications. Nonetheless, such modifications may operate in concert with other techniques in addressing more severe disruptions.

The DTNRG has been working on an architecture designed to accommodate a very wide range of network types, including those with potentially very long delays. The question was raised as to whether it is realistic to believe this architecture will truly be able to span such a large variety of networks. It would seem evident that further experience with the DTNRG architecture may help to answer this question. It may be instructive to recall, as well, that the Internet architecture has been adopted by a very wide range of network types and performance characteristics.

The DTNRG architecture generally provides routing based on names, represented as some form of string. The question was raised as to the difference between names and addresses. When addresses are not derived from a numbering space that is tied to the network topology (*e.g.*, cell phone numbers that can roam), names and addresses can be considered to be effectively equivalent.

Some discussion focused on the issue of how to provide security in networks of this kind. Most systems, including PKI schemes, are made even more difficult to

deploy due to the inability to obtain network credentials on demand. Although there is interest in identity-based cryptosystems, this is nascent and the experience with such systems is fairly limited.

In systems with mixed traffic (*i.e.*, that may include asynchronous traffic along with quasi-synchronous traffic), some facility for indicating the quality-of-service requirements associated with the data may be important. The issue of signaling for this purpose (and possibly others) remains largely unexplored. It was noted that the general problem of signaling has gone in and out of IETF over the years (partly in the form of middle-box communication), but that it still remains out of standardization.

In summary, this *Dagstuhl* seminar has sharpened the understanding of the very different perspectives from which researchers approach the problem space of disruption-tolerant networking, their assumptions and requirements, and the short- and long-term solutions they envision. This has broadened the view on DTN at large and contributes further issues to the present DTN research topics such as naming, security, service differentiation and efficiency. Assuming the traditional well-connected Internet architecture and its (interactive) applications as one extreme and the DTNRG architecture for purely asynchronous communications as another, the middle ground of mobile and partly (dis)connected operation may be approached from either edge. Future research will need to determine how far the DTNRG architecture can and should reach towards traditional Internet applications while maintaining its architectural integrity.