Dagstuhl Seminar on Disruption Tolerant Networking*

Marcus Brunner[1], Lars Eggert[1], Kevin Fall[2], Jörg Ott[3], Lars Wolf[4]

[1] NEC Europe Ltd. Network Laboratory, {brunnerleggert}@netlab.nec.com
[2] Intel Research Berkeley, kfall@intel.com
[3] Helsinki University of Technology, Networking Laboratory, jo@netlab.hut.fi
[4] TU Braunschweig, Institut für Betriebssysteme und Rechnerverbund, wolf@ibr.cs.tu-bs.de

ABSTRACT

Disruption Tolerant Networking (DTN) is a new area of research to improve network communication when connectivity is periodic, intermittent, and/or prone to disruptions. A seminar on DTN was held at at Schloß Dagstuhl, Germany, from 3 to 6 April 2005. Researchers from different fields discussed their approaches to dealing with delays, intermittent connectivity, and the potential non-existence of an end-to-end path in a number of different environments. The two major areas identified were: (1) dealing with delay and disruption in the present Internet in the context of wireless, mobile, and nomadic communications, supporting existing applications and (2) addressing new applications with a focus on exploiting discontinuous connectivity and opportunistic contacts for asynchronous communications. This article briefly reviews the seminar presentations and discussions.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Network Communications

General Terms

Performance, Design, Reliability, Security, Standardization

Keywords

Delay-tolerant networking, disconnected operation, mobility, ad-hoc networks, sensor networks, interplanetary Internet

1. INTRODUCTION

Over the past 40+ years, numerous architectures were developed for network communication, including the OSI reference model and protocol specifications following it and – of course – the Internet architecture. These network architectures were 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 and network topologies, such as relatively short transmission delays, low error probability, and the existence of end-to-end paths.

In an increasing number of today's communication scenarios (which become increasingly commonplace) these assumptions no longer hold. Examples of such *challenged networks* include networks with frequent connectivity disruptions, extremely long transmission delays or unstable/variable connectivity – as may be experienced, for example, by mobile users, in remote locations, and in ad-hoc communication environments. The established network architectures often fail to properly support communications in such settings, resulting in either significant inefficiencies or complete loss of service.

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 intermittent or prone to other disruptions.

Among the challenges of this field of research are potentially large transmission delays. These may result either from physical link properties or 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 allows relaxing 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 functions provided at the transport layer, to the application protocols (and applications) themselves. The possibly resulting high transmission delays, error 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.

To discuss these issues and to advance this field of research, a seminar on Disruption-tolerant Networking (DTN) was held at Schloß Dagstuhl, Germany, from 3 to 6 April 2005. Researchers from different fields discussed their respective approaches to dealing with delays, intermittent connectivity, and the potential non-existence of an end-to-end path in a number of different environments. This article briefly reviews the seminar topics, the presentations and discussions as well as the research questions which have been identified and discussed at the seminar. Further information about the seminar is available from [1].

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

* This seminar was supported by the International Conference and Research Center for Computer Science, Schloß Dagstuhl in Wadern, Germany, http://www.dagstuhl.de/.

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 typically build on the paradigm of asynchronous interactions and introduce additional inter-internetworking layers that span multiple specialized internetworks with different characteristics [2].

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 [3].

Two presentations were orthogonal to this division: Per Gunningberg reported on an experimental testbed for Ad-hoc Protocol Evaluation (APE) encouraging experimental validation in addition to simulation approaches. And Hannes Tschofenig discussed security considerations for DTN environments.

2.1 Applications for Asynchronous DTN

Into this first group or presentations, *i.e.*, those dealing with long delays, fell Bengt Ahlgren's talk on applications for asynchronous networking, Ben Hui *et al.*'s talk on "pocket-switched" networks, Per Gunningberg's notes on the Sámi Network Connectivity (SNC) project bringing network access to Lapland, Margaret Martonosi's presentation on the ZebraNet wildlife tracker, Stephen Hailes' talk on context-aware adaptive DTN routing in the RUNES project, Kevin Fall's presentation on applying DTN to oceanographic measurements, and Srinivasan Keshav's extensions to the DTNRG architecture.

All the work presented shares the fundamental assumption that "the world is not connected" as noted by Hui *et al.* Instead, typically rare opportunistic (sometimes scheduled or predictable) contacts are available to pass data between nodes which are usually mobile. In addition to traditional communications, node mobility itself is also exploited to move data: conveying a message from a sender to a receiver can be achieved by having intermediate nodes (store and) forward the data, by having the data travel with the sender to the receiver, or by some combination of the above.

This motivates research in routing algorithms for disruption-tolerant ad-hoc networks. Particularly with mobile devices owned and carried by human users, findings show that random walk models are inadequate for describing user movement in a realistic fashion and that stateless approaches to routing may lead to low delivery probability. After all, regardless of whether the mobile devices are carried by humans or by animals, motion usually follows social behavior and thus considering social aspects and communities in mobility models is important.

Different routing algorithms were discussed: SNC uses probabilistic routing that calculates a history of mutual encounters between nodes (decaying over time) as a routing metric; routing in the RUNES project groups nodes into densely populated and internally well-connected *clouds* that use proactive synchronous protocols for intra-cloud and asynchronous mechanisms for inter-cloud routing (with the messages moving physically with the nodes in the latter case). For inter-cloud routing, a routing protocol keeping minimal state and relying on prediction is used: attributes describing a node's context (connectivity, co-location with other nodes, etc.) are maintained, past changes of their values are observed, and future evolution is predicted and resulting (reachability) probabilities are exchanged as part of the routing protocol. ZebraNet has experimented with a combination of a history-based approach to routing and flooding in order to move data between field sensors and a fixed infrastructure node.

The applications discussed in this context are all asynchronous and can be classified into sensor data retrieval, file sharing, general information access, messaging and other inter-personal communication, and tracking of mobile nodes and users. While some applications may benefit from vast amounts of storage available in personal computing devices, sensor network nodes may suffer from memory constraints.

In either case, storage management becomes a crucial issue beyond pure routing decisions that touches upon security: how often to replicate information to ensure ultimate delivery, when to purge data, which nodes to accept as custodian for a message, whom to consider trustworthy, etc. The threshold applied by a user may be highly application-specific (*e.g.*, forwarding the vacation pictures from a digital camera to safe storage vs. sending an email with greetings and the weather of the day).

Infrastructure components (*e.g.*, dedicated nodes) may benefit DTN applications as powerful and well-connected (and possibly even trustworthy) anchors. They may simply serve as persistent stationary storage nodes independent of other users or also act as gateways into a fixed network such as the Internet and thus be able to bridge distances without relying on opportunistic forwarding. In the latter case, if multiple options for mobile network access are available, economic and performance tradeoffs may become relevant as another dimension for routing decisions.

Networking architectures and applications (*e.g.* for information retrieval) may even rely on the Internet as a well-connected backbone (a "special DTN region" as suggested by Keshav) with DTN being primarily considered as a management mechanism for "leaf" networks interconnecting mobile and nomadic users. This allows considering DTN as a means to extend the reach of the Internet for asynchronous applications.

Motivations similar to the latter are inherent to those projects coming "from the Internet side" and enhancing its protocols and system architectures in order to enable Internet access in spite of intermittent connectivity.

2.2 DTN Aspects of Internet Communications

The 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, Dirk Kutscher *et al.*'s talk on "Drive-thru" Internet access, Carsten Bormann *et al.*'s talk on issues to be addressed for "near end" DTN solutions, Aaron Falk's presentation on military satellite communication, Nils Seifert's discussion of a pragmatic short-term approach to DTN and its needs, and Lavy Libman's disruption prediction for public transportation.

All of these projects focus on some variant of vehicular (cars, buses, trains, etc.) communications including pure vehicle-to-vehicle communications, direct vehicle-to-fixed network

communications, and multi-hop forwarding via vehicular ad-hoc networks. Users may move as individuals or be part of a mobile network and they largely make use of existing applications accessing today's Internet. It is commonly acknowledged that connectivity disruption is the rule rather than the exception. Disruptions include losing and gaining connectivity for arbitrary durations as well as changes in service quality (from a thick to a thin pipe and vice versa). Mechanisms need to be applied at different layers to deal with intermittent connectivity effectively:

At the link layer, various degrees of diversity can be exploited to improve connectivity in the first place: different access technologies, network operators, and frequency bands may be used to prolong connectivity periods, reduce the error rate, and improve the achievable data rate. Mobile nodes need effective mechanisms for detecting and monitoring link layer connectivity and its quality (without putting undue additional burden on the network itself).

At the IP layer, mobile devices moving between different access networks and service providers require persistent endpoint identifiers. Besides using Mobile IP, Host Identifiers as defined in the Host Identity Protocol (HIP) are one possible solution, (URI-style) identifiers at the application layer are another.

Transport protocols need to freeze their activities during disconnection periods to avoid timeouts and resume quickly when connectivity is regained. Prediction schemes may be applied (*e.g.*, based upon user location, signal strength, and past experience) and, particularly for resumption, link layer indications may provide useful triggers. Instead of suspending and resuming transport connections, they may also be torn down and re-established by some session layer mechanism on top.

Furthermore, application protocols need to be considered and their semantics preserved in spite of connectivity disruptions. As longer disruptions may lead to application layer timeouts, application layer gateways (ALGs) need to be deployed to shield disconnections from the application peers – unless the application protocols themselves can be adapted to become more flexible. To make effective use of connectivity periods and also to decouple the mobile nodes from their (fixed) peer's performance, further application-specific enhancements (such as HTTP prefetching or pro-active caching) are useful options for optimizations, at the cost of generality though. General guidelines for (future) application protocol design in intermittently connected scenarios include: decoupling the application layer protocol functions from the (existence of) underlying transport, *e.g.* for end-to-end security and reliability (as is, *e.g.*, done in Session Initiation Protocol, SIP), allowing for "multithreaded" operation (*i.e.*, avoiding blocking or lock-step operation), and supporting fine-grained suspension and resumption of operations as well as data aggregation at the application layer.

The aforementioned protocol support functions may be performed largely end-to-end (with some assistance only by layer 3 elements) or by using intermediaries (or: proxies) acting at the transport layer and above. In some cases, such transport or application layer gateways may resemble DTN routers in that they accept data from one node destined for another in a store and forward fashion (*e.g.*, email relays or some variant of web caches). The major difference is that such intermediaries are usually optimized for synchronous communications and handle asynchronous operation as a special case while DTN routers support only asynchronous communications.

Finally, application (G)UIs – and ultimately users – need to be able to deal with challenged environments: with disruptive communications, abstracting from the underlying network characteristics becomes infeasible. Instead, the user needs to be informed about progress and temporary suspension of operations, should be able to invoke and queue sequences of operations (of which "tabbed browsing" is a very simple case), and should obtain asynchronous notifications about their completion.

As important as the focus on the existing Internet as infrastructure and its present applications is the capability of incremental deployment without requiring broad roll-out of additional infrastructure (on the fixed and mobile end systems as well as "in the network") – which may suggest end-to-end approaches or overlays. The incremental deployment also demands scalable solutions that are workable with very few nodes in the beginning (already providing benefit) but also when a high penetration is achieved; this is particularly relevant for wireless ad-hoc routing and access links.

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 there are no legacy applications to support. Therefore, 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. Similarly, in contrast to Internet extensions, newly developed DTNRG applications need not be concerned with legacy user experience and expectations.

Common ground is found at the link and the application layers while the mechanisms in-between differ: As both may need to rely on opportunistic communications, efficiently determining availability and loss of link layer connectivity is important, cross-layer integration to make this information available up to the applications, to routing, or to mobility management is considered highly desirable. Application protocols most suitable for intermittent communications bear similarities between both groups, and proxy approaches on one hand and DTN routers on the other share commonalities as well. Finally, end-to-end security mechanisms may need to follow similar approaches as they are not able to rely on underlying transport mechanisms (such as TLS) and both approaches may need to place trust in intermediaries thus requiring strong mutual authentication.

3.1 Extensions to Internet Protocols

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, TCP modifications may operate in concert with other techniques that address more severe disruptions. With increasing disruption duration and frequency, higher layer protocols need to be involved in handling intermittent connectivity and the potential non-availability of an end-to-end path requires introducing intermediaries. Ultimately, the user experience may change noticeably, essentially approaching the DTNRG style of operation.

Solutions developed around TCP hacks may also benefit basic operation using the DTNRG architecture. As noted above, the requirements on lower layer connectivity detection are quite similar, and also DTNRG protocols usually aim at maximizing information exchange at each opportunity. DTNRG defines convergence layers for mapping the abstract protocol operations onto (inter)network-specific services and, among others, one for communication via TCP/IP exists. The design of such convergence layers may draw upon the experience gained from the various TCP hacks to maintain efficient communications in spite of very short outages while dealing with more serious disruptions is left to the DTNRG bundle protocol itself.

3.2 DTNRG Architecture

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. One research question is 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.

In systems with mixed traffic (*i.e.*, including asynchronous traffic along with quasi-synchronous traffic), some facility for an application to indicate its *intentions* associated with the data may be important. This is different from traditional QoS in that it is less about absolute performance requirements but more of a hint to the network and intermediates. The issue of meaningful signaling for this purpose (and possibly others) remains largely unexplored. Similarly, as noted above, applications are expected to benefit from knowing when they are connected. As the disconnected link(s) may be anywhere along the path between two peers, signaling support is also required in the opposite direction.

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. The former notion of *regions* is changing as, unlike the network parts of IP addresses, they do not have topological significance. Instead, they should just be considered *address types* (such as “dns://” or “e164:”). It is up to the DTN routing to determine a forwarding path to a destination identified by such an address URI. This may be done, *e.g.*, by explicit routing table entries or by using source routing information (*hints*) that point to entities that can resolve the address. Hints could be supplied by the originator (if known) or filled in by entities on the path; alternative hints may be provided as backups. The precise shape of hints and the associated mechanisms is subject to discussion.

Some discussion focused on the issue of how to provide security in networks of this kind. The current DTN approach is on hop-by-hop security to protect the infrastructure and prevent theft of service. Each DTN router knows its peer routers and is responsible for its previous link when accepting and forwarding bundles, should validate their identity, and re-sign them so that the next hop can do so, too. This implies that the sender authenticates to the first hop DTN router upon submitting a bundle (which may have some privacy implications). One open issue is telling legitimate duplicates (because of replication or retransmission) from replayed ones (due to DoS attacks). Another one is how to prevent black-holing of bundles, *i.e.*, a DTN router taking custody of a bundle and never releasing or forwarding it. Further points of discussion include the potential penalty associated with bundle authentication and the general idea of policy enforcement at every DTN router. While this kind of access control is supported by the DTN architecture, its actual use in practice depends on the specific network scenario.

Relying on hop-by-hop authentication and thus creating a chain of trust between peer nodes is important for DTNs because most security systems, including PKI, are even more difficult to deploy due to the inability to obtain network credentials on demand (which requires keying material to be provisioned). Although there is interest in identity-based cryptosystems, this is nascent and the experience with such systems is fairly limited. There may also be link layer access control (at least for the really costly links) and application layer security in place from which keying material and trust may be leveraged.

4. CONCLUSION

This Dagstuhl seminar has helped understanding 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.

5. REFERENCES

[1] Website of the Dagstuhl seminar 05142: http://www.dagstuhl.de/05142/

[2] Website of the DTN Research Group: http://www.dtnrg.org/

[3] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby. Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations. RFC 3135. June 2001.