

Promoting Tolerance for Delay Tolerant Network Research

Jon Crowcroft, Eiko Yoneki, Pan Hui
University of Cambridge
England

firstname.lastname@cl.cam.ac.uk

Tristan Henderson
University of St Andrews
Scotland

tristan@cs.st-andrews.ac.uk

ABSTRACT

So what is all this DTN research about anyway? Sceptics ask: “Why are there no DTN applications?”, or “Why is DTN performance so miserable?” This article attempts to address some of these complaints. We present suggestions of expectations for applications, and metrics for performance, which suggest a more tolerant view of research in the area.

1. INTRODUCTION

The Internet architecture has enabled new applications, such as e-mail, the World Wide Web and Voice over IP, which have transformed the way in which we access and create information. This architecture, however, makes strong assumptions about connectivity, such as available end-to-end paths and low round-trip times, and high availability to naming, caching and search infrastructures to provide locator-based access such as DNS. This means that it is difficult to use e-mail or VoIP in a so-called *challenged* environment with unstable or high-delay connectivity, such as the dark side of the moon. This has led to the emergence of Delay Tolerant Networks (DTN) [7] as a different communication paradigm; one which is decentralized and distributed over a multitude of devices that are dynamically networked, carried by people, and embedded in everyday-life. For instance, the Huggle Project¹ features a “Pocket-Switched” variant of DTN where people carry devices in their pockets, which communicate directly with other devices within their range or with infrastructure. As people move around, they can exchange messages with nearby devices, carrying a message until it is close to another device. Thus, our DTN-device-carrying astronaut might have their e-mail messages ferried by other DTN users taking their commercial flights to and from the moon.

The last few years have seen an explosion in DTN research, with several papers in the main SIGCOMM conference, and more specific venues such as the CHANTS workshop² which is now in its fourth year. At the same time, however, there remains some scepticism about the need for a new DTN-like architecture. This article is designed to address some of this scepticism. First, we briefly outline a history of DTN research. We then list some of the usual sceptics’ questions. Subsequently, we discuss the future of DTN Architectures and Applications, and what we consider to be the design space for appropriate Metrics for success.

¹<http://www.huggleproject.org/>

²<http://chants.cs.ucsb.edu/>

2. A BRIEF HISTORY OF DTN RESEARCH

DTN research started with Vint Cerf and the Interplanetary Internet initiative [5], which proposed a new architecture that could work over both terrestrial and interplanetary links. This architecture could enable applications such as the remote operation of scientific experiments on other planets, controlled using TCP/IP from Earth.

The Interplanetary Internet ideas were generalised by the IRTF DTN Research Group³, which focusses on any kind of challenged environment where end-to-end connectivity may not always be available, and the DARPA Disruption Tolerant Networking programme⁴, which concentrates on developing protocols for bundling application-layer data units into DTN-layer protocol bundles for transport by DTN nodes.

Opportunistic networks [21] and message ferrying [25] concentrate on mobile ad hoc DTNs as in the previously-discussed Huggle example, where routes are built dynamically between source and destination, and any possible intermediate node can be used *opportunistically* to ferry data as required. At the same time, applied EE/CS researchers, dissatisfied with the inefficacy of MANETs that they were trying to build in the real world, have started to do more and more experimental work, mapping out the world in terms of real radio propagation experiences, and actual mobility traces of genuine users (pedestrian, vehicular, zebras, badgers, albatrosses and so on). An important aspect of DTN route-building is therefore an understanding of mobility patterns, and as such recent DTN research has studied social networks, both in humans [6, 19, 4] and other species [16]. These two strands of research can be seen to have converged in attempts to build realistic models for high-variance density MANETs.

As well as being interesting from a systems and experimental point of view, research problems in DTNs, rather like MANETs before them, are also academically fascinating. A large class of models from the world of *network science* and *complex systems* can be extended covering the graph theoretic properties, and using physical analogs for the propagation of information through their dynamic topologies (percolation, diffusion, statistical thermodynamic/entropic models, and many more). In some senses, there is a body of work which stands on its own. In other senses, the reality gap between these theoretical models and the real results can be seen in the differences between predicted and measured performance of test applications. We are not trying to write a “defence of the dark arts of DTNs” syllabus here. Instead,

³<http://www.dtnrg.org/>

⁴<http://www.darpa.mil/sto/strategic/dtn.html>

we concentrate on questioning the realism of DTN research.

3. SCEPTICS' FAQs

Here are some of the accusations that we have recently heard levelled against DTN research. Please feel free to add your own attack (or defence) on CCR Online.

1. *Sceptic:* The world just keeps getting more connected. Yes, we may have frequent dropped phone calls over 2G/3G networks, and 3G and WiFi coverage are not completely ubiquitous yet, but the same can be said of wireline Internet access in the first ten years.

Short answer: However well-connected the world gets, there will still be outages. Moreover, there will still be variation in capacity - it is unrealistic to assume universally consistent coverage. An architecture that seamlessly deals with disconnections, disruption and delay will simply work better than one that doesn't. It may have more capacity (at the expense of larger expected mean and variance of delay), and this may be tolerable or even transparent to some applications.

2. *Sceptic:* You don't have any compelling applications.

Short answer: There was an 11 year gap between the deployment of TCP/IP in 1981 and the very first WWW browser/server in 1992, which we now know to be the Internet's "killer app". Expecting the widespread emergence of new applications during the initial stages of a network research programme seems a little unfair. In any case, there may well be several compelling applications, which we outline below.

3. *Sceptic:* Do people living without electricity really need cell phones or \$100 laptops? Aren't there more important things on which to spend time, money and effort?

Short answer: Evidently! Indeed growth in cell phone ownership, despite poor coverage, is faster in these areas than elsewhere now [22]. DTN systems such as Kiosknet [10] and similar systems such as the OLPC data carousels may prove compelling applications for these areas.

4. *Sceptic:* OK so sensor networks might usefully be architected as DTNs, but aren't sensors a special case?

Short answer: Sensors, and other send-and-forget applications are already being integrated on larger devices, such as the Nokia sports mobile phones. Why build two separate network architectures? Why not just have one architecture that copes in all ranges of the operating spectrum/regimes.

5. *Sceptic:* Provisioning for DTNs makes the APIs harder!

Short answer: Designing DTN applications motivates us to solve the problem of applications being written to survive frequent disconnects. This is already occurring in many web service applications — for instance web-mail providers such as Google incorporate application-layer mechanisms to recover from dropped links. If we can build an architecture that enables such applications to work in DTN environments, then won't they work even better in the "unchallenged" infrastructure environments?

6. *Sceptic:* Many challenged environments involve disaster recovery. Shouldn't we pre-plan for disasters anyhow?

Short answer: The reality is that disasters create unpredictable failures (see New Orleans for instance), so you need a network that can cope in unforeseen circumstances, and moreover allows people to adapt to these circumstances as well.

7. *Sceptic:* Why do you believe that your small-scale experiments have anything to do with what you will get in future large-scale deployments?

Short answer: We believe that there is no reason to say that the scaling we see in the Huggle experiments and related measurements with such low delivery success, low rate and long delay, will not scale up in the multi-hop wifi and denser populations — indeed, measurements of cell phone networks [8] and transportation usage [17] appear to confirm this.

8. *Sceptic:* Why don't DTN theory and practice link up better?

Short answer: You might ask the same of the wireline Internet. "Joining up the dots" between experiments and theory takes a long long time, as evinced by the telephone networks.

These are some short and somewhat glib responses. No doubt all we have done is encourage the sceptics to demand further details. We provide these by classifying the FAQs into three broad categories: architectures, applications, and performance.

4. ARCHITECTURES

So the question here is: do we really need a brand-new architecture, and if so, how many new architectures do we need? We describe the generic DTN architecture and a specific instantiation of this with which we are familiar: Huggle.

4.1 DTN Architecture

Challenged networks arise primarily as a result of various forms of host and router mobility, but may also come into being as a result of disconnection due to power management or interference. Examples of such networks include Terrestrial Mobile Networks, Exotic Media Networks, Military Ad-Hoc Networks, and Sensor and Sensor/Actuator Networks. The DTN architecture [7] seeks to address the communication needs of these challenged environments through a message-based store-and-forward overlay network that leverages a set of convergence layers to adapt to a wide variety of underlying transports. In addition, the model also espouses novel approaches to application structuring and programming interface, fragmentation, reliability, and persistent state management.

The Huggle architecture builds on the original DTN architecture by providing a data-centric architecture [3] where applications do not have to concern themselves with the mechanisms of transporting data to the right place, since that makes them infrastructure-dependent. By delegating to Huggle the task of propagating data, applications can automatically take advantage of any connection opportunities

that arise, both local neighborhoods opportunities and connectivity with servers on the Internet when available. This is clearly infeasible in the existing TCP/IP architecture.

The Hagggle architecture is designed based on the following four principles.

1. **Data Persists inside Hagggle:** The data on each node in Hagggle must be visible to and searchable by other nodes. In practice, this means that Hagggle must manage persistent data storage for applications, instead of applications storing data in a separate file system.
2. **Networking Protocols inside Hagggle:** Any application-layer networking protocol includes implied assumptions about the type of network available. For example, client-server protocols such as SMTP, POP and HTTP assume that Internet-based servers are contactable.
3. **Name Graphs supporting Late Binding:** Hagggle maintains its own naming repository, with mappings from user-level names to protocol-specific names, which specify the various ways to get to the user-level name. Furthermore, the whole set of mappings (the “name graph”) is transmitted along with the data, allowing even intermediate relay nodes to bind to protocol-specific names as late as possible [2].
4. **Centralised Resource Management:** One role of the networking architecture on every device is to decide what to do with each of its network interfaces *now*. Hagggle contains a centralised resource management component, which decides on a cost/benefit comparison basis what tasks it chooses to perform on each network interface at a given moment.

Hagggle is a clean-slate design tailor-made for mobile devices based on the characteristics of mobile networks (i.e., lack of end-to-end connectivity, availability of local peer-to-peer), and some expected applications as we outline below. At the same time, Hagggle incorporates popular Internet application protocols such as SMTP, POP, and HTTP into its Protocol Manager. Thus, we can see that it is possible for DTN and TCP/IP to be complementary architectures, allowing DTNs to extend to those places where the current Internet cannot.

5. APPLICATIONS

As we have already pointed out, there was a large gap between the development of the Internet architecture and the first compelling applications. This does not mean, however, that we should, or even can, ignore the applications that might benefit from a DTN architecture. We do not claim to be able to predict the killer applications for DTNs; instead, we describe some of the aspects of DTNs that can enable novel applications. Whether these will be a success is up to the market.

- Network Topology

Understanding network topology in a DTN, or in particular a pocket-switched network (PSN), means understanding the human connectivity, social network topology, and also the urban topology. Understanding the physical characteristic about the human networks,

for example the degree distributions, the community structures, social distance, and centrality distributions, can enable new suitable inter-personal communication applications. Knowing the urban topology and human flows can tell us where to install infrastructure to complement the opportunistic communications and also useful for locating the places to distribute contents.⁵

- Social applications

PSNs can enable more social-based applications. Unlike the wired Internet, PSNs not only comprise cables and routers, but mobile humans themselves. By leveraging human connectivity and *social topologies* rather than logical or physical topologies, new mobile social networking applications can be created. At the same time, social and contact information derived from the network itself can be used to improve the performance of both the network (e.g., forwarding), and applications themselves (e.g., more effective ways of adding to a user’s social network).

- Opportunistic Delivery

Several efficient forwarding algorithms for DTNs have been proposed. A majority of the algorithms are based on epidemic routing protocols [24], where messages are simply flooded when a node encounters another node. The optimisation of epidemic routing by reducing the number of copies of the message has been explored. For example, in [23], spray and wait routing assigns a limited number of copies. Many approaches calculate the probability of delivery to the destination node, where the metrics are derived from the history of node contacts, spatial information and so forth. The pattern-based Mobyspace Routing by Leguay *et al.* [14], location-based routing by Lebrun *et al.* [13] and PROPHET Routing [15] fall into this category. PROPHET uses past encounters to predict the probability of future encounters. The transitive nature of encounters is exploited, where indirectly encountering the destination node is evaluated. Message Ferry by Zhao *et al.* [25] takes a different approach by controlling the movement of each node.

Recent research has attempted to uncover hidden stable network structure in social-network-based DTNs. For example, SimBet Routing [6] uses ego-centric centrality and its social similarity. Messages are forwarded towards the node with higher centrality to increase the possibility of finding the potential carrier to the final destination. In [11], BUBBLE exploits further community structures and combines it with a network property of the global betweenness for further improvement of forwarding algorithms. Adapting social context results in a significant reduction in routing overheads. The mobility-assisted Island Hopping forwarding [20] uses network partitions that arise due to the distribution of nodes in space. Their clustering approach is based on the significant locations for the nodes and not for clustering nodes themselves. Clustering nodes is a complex task to understand the network structure for aid of forwarding.

⁵<http://www.amillionpeople.net/>

If, however, we can design effective and efficient mechanisms for opportunistic delivery of DTN application data, this should enable new mobile applications.

5.1 Application Examples

So what kind of new exciting mobile applications could these enablers create?

- Bulk data distribution in urban areas (e.g., free customised newspapers/adverts to cell phones on the metro or tube). This is ideal for a DTN since capacity in a DTN will **always** be much higher than in an infrastructure network [9]. At the same time you will still require some infrastructure, due to the phase transition [12] seen for reachability, and also for the injection of original data, and perhaps for identity and payment mechanisms.
- Sharing of individual contents (rather than bulk data) in urban areas. For example city commuters may stay together in the same bus or train for a good amount of time, using their short-range radio equipped mobile devices, individuals can share content with fellow travelers. Historical colocation and social information can be used to determine the best content sources. [17]
- Disconnected kiosks in rural areas. Rural Internet kiosks in developing regions can provide a variety of useful services such as birth, marriage, and death certificates, land records, and medical and agricultural consulting to the poorest sections of society. KioskNet uses low-cost computing devices as kiosks and use vehicle as mechanical backhaul to move data to Internet gateways [10]. They have run pilot deployment in Anandapuram, a village in South India, as early as in year 2006.
- Mobile location-aware sensing applications. B-MAD is a Bluetooth and WAP push based location-aware mobile advertising system [1]. Bluetooth sensors deployed in the environment discovers the nearby cell mobiles and send the Bluetooth addresses over a WAP connection to the Ad Server. Ad Server will push related advertisements to the end users. This is not really a DTN application, but similar concept can be used in DTN to push bulk data to end users.
- Social mobile applications. Hagggle project has an social mobile application called MobiClique, which uses social network information from Facebook⁶ to bootstrap the data forwarding structure of the DTN, and at the same time social networks are also created from the daily contact patterns. Combining with location services and the assistance of limited amount of infrastructures, DTN can provide a big variety of city-wide social services.

6. PERFORMANCE

Closely related to the issue of potential DTN applications is the issue of metrics. Experimental results may be considered “bad”, but are we measuring “goodness” correctly? Many DTN research projects aim for high delivery ratios, low delivery costs and low delivery delay.

⁶www.facebook.com

Delivery Ratio: The proportion of messages that have been delivered out of the total unique messages created. Usually, we want to achieve higher delivery ratio unless the delivered messages containing malicious materials.

Delivery Delay: The delay experienced by the fastest delivery of the multiple copies of the same message from a source to a destination. DTN is delay tolerant, but most of the time we still want to minimize delays.

Delivery Cost: The total number of messages (include duplicates) transmitted across the air. Normally, we want to reduce the delivery cost to reduce battery consumption and also the contentions of the wireless channel.

In general, DTN forwarding algorithm designers aim to increase delivery ratios, while at the same time reducing delivery delay and cost. But these metrics can prove confusing and even misleading. For instance, some applications may prefer low delivery delay at the expense of a high delivery ratio. Or an application may appear to have a low delivery cost, but only because many messages are dropped. Perhaps we need new metrics that can better demonstrate the “goodput” of a DTN application before we can properly evaluate the results from existing DTN experiments.

Even discarding the issue of metrics, perhaps the sceptics are being unfair about the performance of these experiments. There are many practical difficulties in the way of conducting large scale DTN trials. Hence many experiments are, to date, relatively modest in terms of the size and fraction of a community used, and the duration of the experiment. This also tends, naturally, to limit the kinds of applications considered (or imagined). But it may also have an impact on the perceived performance. For instance, many existing experiments use Bluetooth for measurement, rather than 802.11. Using existing 802.11 in a DTN is difficult because of battery life, but this may change in the future given the forthcoming prevalence of 802.11 in handsets and also IEEE work on low power 802.11. The use of a different link-layer technology

It may also be unfair to evaluate future DTNs based on current experiments because we still do not clearly understand communication patterns. For example many systems assume that every node in the system communicate with each other uniformly. But is this a fair assumption? Assuming uniform traffic and connectivity across the whole population may lead to low delivery ratios, because delivery may rely on a few bridge nodes to carry data between two communities. If, however, it turns out that actual mobility is more community-based, we may observe higher intra-community delivery ratios, with humans communicating more with their friends or “familiar strangers” [18]. We have observed these bridge nodes in small experiments [4] but we do not know yet whether these results generalise.

7. CONCLUSIONS

We have observed some scepticism in our research community about the rationale for delay tolerant networks. This article has attempted to demonstrate that DTNs have the potential to enable exciting new applications that cannot be served by the current Internet. To find out whether this will be the case, we hope that the DTN architecture will be afforded the same time and tolerance as given to TCP/IP.

8. ACKNOWLEDGEMENTS

Thanks to Mark Handley and Costin Raiciu of UCL for

being so skeptical as to prompt us to write this.

9. REFERENCES

- [1] AALTO, L., GÖTHLIN, N., KORHONEN, J., AND OJALA, T. Bluetooth and wap push based location-aware mobile advertising system. In *MobiSys '04: Proceedings of the 2nd international conference on Mobile systems, applications, and services* (New York, NY, USA, 2004), ACM, pp. 49–58.
- [2] ADJIE-WINOTO, W., SCHWARTZ, E., BALAKRISHNAN, H., AND LILLEY, J. The design and implementation of an intentional naming system. In *Proceedings of SOSP* (Dec. 1999), pp. 186–201.
- [3] AGUILERA, M. K., STROM, R. E., STURMAN, D. C., ASTLEY, M., AND CHANDRA, T. D. Matching events in a content-based subscription system. In *Proceedings of PODC '99* (1999), pp. 53–61.
- [4] BIGWOOD, G., REHUNATHAN, D., BATEMAN, M., HENDERSON, T., AND BHATTI, S. Exploiting self-reported social networks for routing in ubiquitous computing environments. In *Proceedings of the 1st International Workshop on Social Aspects of Ubiquitous Computing Environments (SAUCE)* (Avignon, France, Oct. 2008).
- [5] BURLEIGH, S., HOOKE, A., TORGERSON, L., FALL, K., CERF, V., DURST, B., SCOTT, K., AND WEISS, H. Delay-tolerant networking: an approach to interplanetary internet. *IEEE Communications Magazine* 41, 6 (June 2003), 128–136.
- [6] DALY, E. M., AND HAAHR, M. Social network analysis for routing in disconnected delay-tolerant manets. In *Proceedings of ACM MobiHoc* (Montreal, Quebec, Canada, 2007), pp. 32–40.
- [7] FALL, K. A delay-tolerant network architecture for challenged internets. In *Proceedings of ACM SIGCOMM* (2003), pp. 27–34.
- [8] GONZALEZ, M. C., HIDALGO, C. A., AND BARABASI, A. L. Understanding individual human mobility patterns. *Nature* 453, 7196 (June 2008), 779–782.
- [9] GROSSGLAUSER, M., AND TSE, D. N. C. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking* 10, 4 (Aug. 2002), 477–486.
- [10] GUO, S., FALAKI, M. H., OLIVER, E. A., RAHMAN, U. S., SETH, A., ZAHARIA, M. A., AND KESHAV, S. Very low-cost internet access using kiosknet. *SIGCOMM Computer Communication Review* 37, 5 (Oct. 2007), 95–100.
- [11] HUI, P., CROWCROFT, J., AND YONEKI, E. Bubble rap: Social-based forwarding in delay tolerant networks. In *Proceedings of ACM MobiHoc* (May 2008).
- [12] HUI, P., AND LINDGREN, A. Phase transition of opportunistic communications. In *Proceedings of the ACM Workshop on Challenged networks (CHANTS)* (Sept. 2008).
- [13] LEBRUN, J., CHUAH, C.-N., GHOSAL, D., AND ZHANG, M. Knowledge-based opportunistic forwarding in vehicular wireless ad hoc networks. In *Proceedings of the 61st IEEE Vehicular Technology Conference* (2005), pp. 2289–2293.
- [14] LEGUAY, J., FRIEDMAN, T., AND CONAN, V. Evaluating mobility pattern space routing for DTNs. In *Proceedings of INFOCOM* (Apr. 2006), pp. 1–10.
- [15] LINDGREN, A., DORIA, A., AND SCHELÉN, O. Probabilistic routing in intermittently connected networks. In *Proceedings of Service Assurance with Partial and Intermittent Resources SAPIR* (2004), pp. 239–254.
- [16] LINDGREN, A., MASCOLO, C., LONEGAN, M., AND MCCONNELL, B. Seal2seal: A delay-tolerant protocol for contact logging in wildlife monitoring sensor networks. In *Proceedings of IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS08)* (Oct. 2008).
- [17] MCNAMARA, L., MASCOLO, C., AND CAPRA, L. Media sharing based on colocation prediction in urban transport. In *Proceedings of ACM MobiCom 2008* (Sept. 2008).
- [18] MILGRAM, S. The familiar stranger: An aspect of urban anonymity. In *The Individual in a Social World: Essays and Experiments* (Reading, MA, USA, 1977), S. Milgram, J. Sabini, and M. Silver, Eds., Addison-Wesley.
- [19] MTIBAA, A., CHAINTREAU, A., LEBRUN, J., OLIVER, E., PIETILAINEN, A. K., AND DIOT, C. Are you moved by your social network application? In *Proceedings of the First ACM SIGCOMM Workshop on Online Social Networks (WOSN)* (Seattle, WA, USA, Aug. 2008).
- [20] N. SARAFIJANOVIC-DJUKIC, M. P., AND GROSSGLAUSER, M. Island hopping: Efficient mobility-assisted forwarding in partitioned networks. In *Proceedings of IEEE SECON* (2006), pp. 226–235.
- [21] PELUSI, L., PASSARELLA, A., AND CONTI, M. Opportunistic networking: data forwarding in disconnected mobile ad hoc networks. *IEEE Communications Magazine* 44, 11 (2006), 134–141.
- [22] RAI, S. India leads world in cellphone expansion. *International Herald Tribune* (Sept. 15 2006). <http://www.iht.com/articles/2006/09/15/business/cell.php>.
- [23] SPYROPOULOS, T., PSOUNIS, K., AND RAGHAVENDRA, C. S. Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In *WDTN '05: Proceeding of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking* (2005), ACM Press, pp. 252–259.
- [24] VAHDAT, A., AND BECKER, D. Epidemic routing for partially connected ad hoc networks. Tech. Rep. CS-200006, Duke University, Apr. 2000.
- [25] ZHAO, W., AMMAR, M., AND ZEGURA, E. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of ACM MobiHoc* (New York, NY, USA, May 2004), pp. 187–198.