

Bulk Data Transfers through an Airline Delay-Tolerant Network

Christina M. Malliou, Nikolaos Bezirgiannidis, Vassilis Tsaoussidis
Democritus University of Thrace
Xanthi, Greece
e-mail: {cmalliou, nbezirgi, vtsaousi}@ee.duth.gr

Abstract— In the era of big data, the Internet engineering community is searching for solutions to alleviate the issues caused by the constantly increasing data traffic. In this paper, we attempt to revive the sneakernet paradigm as a possible solution for non-real-time bulk data transfers. We propose a sample network architecture that takes advantage of the existing worldwide airline infrastructure, and leverages Delay-Tolerant Networking architecture to transfer data over the air in an automated way. We exploit Contact Graph Routing algorithm, which utilizes flight schedules to route bundles based on delivery delay or cost minimization. We examine the applicability of our proposal in a scenario that includes bulk space-data transfers between ESA data centers. Through simulations, we illustrate that the proposed scheme can deliver data efficiently between connected data centers, while the achieved throughput increases with the amount of data transmitted.

Keywords—*Delay/Disruption Tolerant Networking; Airport network; Sneakernet; Bulk data; Contact Graph Routing.*

I. INTRODUCTION

Internet growth and digital data production are typically considered as two symmetric but also reciprocally-influenced aspects of our digital era: Internet boosts digital data production and digital data growth is served through the growing Internet. However, data privacy, ownership or processing parameters occasionally call for local access of data – typical Internet speeds cannot accommodate timely transfer of huge data to local ends.

Every year a vast amount of digital information is produced, with a rate that grows exponentially. This includes non-real-time data, such as data center backups and scientific data. For example, according to [1] the European Bioinformatics Institute (EBI) in Hinxton, UK, currently stores 20 petabytes of data and back-ups; in 2016 the Large Synoptic Survey Telescope, in Chile, will create 140TBytes of data every five days [2] and the European Space Agency (ESA) expects production of 850Gbit of compressed data per day from a single mission (Euclid) [3] NASA in [4] claimed that planned missions will easily stream more than 24TBytes a day. Also, every year, particle-collision events in CERN's Large Hadron Collider generate around 15 petabytes of data [1]. Frequently, these data have to be transferred around the world, in order to be elaborated locally by various research centers, occasionally with different scope. All such non-real-time transmissions of scientific data, plus other potential information sources of medical, meteorological or social nature, that produce huge data, constitute the major challenge of the proposed approach. Until now, such data was delivered either using costly dedicated networks or via physical delivery of hard copy elements between the source and the end-user.

Our architectural approach exploits two major properties of future Global Internet (see [5]): (i) the ability to accommodate data-in-flight in storage, following the Delay-Tolerant Networking (DTN) paradigm, and (ii) the Contact Graph Routing algorithm (CGR), which optimizes routing when deterministic contacts are scheduled – with or without probabilistic influence.

In particular, we explore the potential to incorporate a Delay-tolerant network within the network of airports around the world, and build a Contact Graph Routing algorithm using the scheduled flight connections, aiming at carrying bulk data between data centers located near airports. We evaluate this potential based on throughput gains and cost expenditures; however, given the increasing amount of big-data applications and the digitalization of everything, we consider that our approach to evaluate impact is rather modest. In order to present realistic results, we use the high-capacity dedicated internal network of ESA [6], as reference for comparison.

Practically, we consider end users unaware of network characteristics. Certainly, users are indeed delay-tolerant since they expect to transfer huge amounts of data towards a far end; therefore, our assumption for an end-to-end delay-tolerant application for delivering huge data across the globe is not unrealistic. Instead of using a transmission link, the network uses the physical airline connection. Thus, the aircraft becomes the transmission link. Clearly, the Delay X Bandwidth product with the typical link versus the airline-as-the-link differs; the real issue is when the data-to-transfer or the storage capacity of the airplane balances the trip delay.

In the context of the physical architecture constraints, we show that a significant improvement in terms of delivery time/throughput could be achieved; the significance grows when data amount grows and the importance of the solution grows when the geographical distribution of the end users expands. Beyond that, one can focus also on the importance of automating the manual data delivery service, alone.

Our work is structured as follows: we discuss the context of our approach within the framework of related work in Section 2 and present the architectural constraints in Section 3. We detail the evaluation methodology in Section 4 and show the results in Section 5. We conclude in Section 6 along with our remarks for complementary and future work.

II. RELATED WORK

Researchers have proposed different types of vehicular networks, such as trains, buses, cars and airplanes, for data transportation. In TrainNet [7], storage devices are placed in trains and stations for delivering data from one station to another. As in our approach, this method provides high bandwidth link that could be used to deliver non real-time data. The

authors propose the alternative of fiber optics to connect trains and stations, however they evaluate it using human interference, i.e., the transmission of disk cases between trains and stations. Furthermore, they focus on data queue management policies, and consider only single-hop transmissions, without any routing or automated forwarding policies.

In various works, researchers have proposed the use of buses to extend Internet service to disconnected areas. In [8], buses travel according to a schedule between Internet kiosks and opportunistically exchange data. A similar functionality is presented in [9], where data transfers exploit a connectivity plan routing protocol. In contrast to the network proposed here, the aforementioned architectures strongly rely on the Internet, and couldn't be used for bulk data transmissions.

Also, cars are used in [10] for solving the transmission problem of bulk data. The whole idea exploits the existing worldwide road infrastructure for moving huge amounts of data between geographically distributed locations. In this network, unlike our work, cars do not follow regular prescribed schedules.

The idea of exploiting air flights has been proposed in [11] where the authors suggest a method to send messages between airports, based on the scheduled flight connections. This network is used for delivering small size messages from one airport to another, by using the passengers' mobile devices, where the messages are loaded depending on their destination, while passengers are waiting for their flight. By contrast, our proposed approach targets the transmission of bulk data exploiting an infrastructure installed in airports and airplanes. Furthermore, we propose the use of Contact Graph Routing algorithm, which exploits the scheduled connections, and achieves 100% delivery ratio, assuming adequate storage network capacity.

In [12], the authors suggested the combined use of the Internet together with the postal system to send a part of the data using hard-drives. However, this approach lacks automation, due to the fact that it strongly relies on the human interference.

Here, we attempt to exploit the airline network to create a high-capacity, automated system for bulk data transmissions, which operates in a transparent fashion (i.e., without any need to physically transfer storage media), and routes data according to the predefined air flight schedules.

III. SYSTEM ARCHITECTURE

In this Section, we present the architecture of the proposed system, which can be used to transport massive quantities of non-real-time data between two distant geographical locations. The system combines the existing worldwide airline infrastructure with the concept of delay-tolerant networks, providing a reliable manner to distribute bulk data in an automated way.

In particular, the proposed architecture exploits the DTN paradigm as an alternative or complementary network layer to IP, for the data transmissions. Applications that are used to transmit data, either in a pull or in a push function, function on top of the DTN architecture in an automated and transparent way, similar to the Internet infrastructure.

Unlike the Internet protocol suite, however, DTN does not necessitate end-to-end connectivity; instead, it can deliver data in the presence of communication disruptions, or through intermittent links. Therefore, it constitutes an ideal candidate for transmitting data through a store-and-forward airplane net-

work. Moreover, DTN may function on top of the Internet architecture as well [13], hence providing a hybrid mode of operation, where data can be routed either over Internet links or airplane links, depending on the routing objective.

Our network model comprises a *sending node* (e.g., research center where scientific data are stored), a *destination node*, (e.g., research center where a scientist requests data download), a set of *interconnected airports* with persistent DTN storage capabilities, which are in the vicinity of the research or data centers, and a *fleet of airplanes*, which have persistent storage to carry the bulk sets of data onboard (see Figure 1).

In a data transmission scenario, a bulk data set, originated from the sending research center, is transferred to the nearby airport, where it is stored in the data storage. When the airplane - selected by the data routing function - arrives, the data set is transferred to the airplane storage. For large data sets that do not fit in a single airport, different bundles (i.e., subsets) of data are routed via different airplanes. The airplane(s) travel(s) with data onboard and, upon arrival at destination airport, data are offloaded into the destination airport's persistent storage. From that point, the data bundles may continue with transmission over consecutive flights, until the arrival at the destination airport (i.e., the airport that supports the destination research center), which will forward the data to the destination research center.

In this early work, we focus on a generic architecture scheme and on the evaluation of different routing alternatives. Therefore, we do not study the hardware components and network infrastructure extensively, but, instead, propose a sample infrastructure implementation. In our proposal, the research centers are connected with airports via fiber optics links; the airports have installations of high-capacity network drives where the data are transferred on a parallel mode; airplanes have boxes of storage disks (e.g., Solid-State Drives); and data are transferred between airport and airplane storage drives through high capacity Ethernet (10 /100 Gigabit Ethernet) on parallel mode. Since the storage hardware components are being continuously improved, the infrastructure components may evolve or be updated to higher-capacity, higher-speed, state-of-the-art hardware, and provide further improved data rates and storage capabilities than the ones proposed here.

We note that, since in this work we focus on transmission of bulk amounts of non-real-time, non-confidential data (e.g., scientific/research data), we do not consider security or confidentiality aspects of the data deliveries.

The aforementioned procedure is depicted in Figure 1.




Figure 1. Sample Network Topology

One of the main features of the proposed bulk data delivery system is that data transmissions are based on the air flight schedules. Hence, we employ CGR [14] [15], an algorithm that bases its routing decisions on some ordered list (named "contact plan") of anticipated connectivity changes, named "contacts".

The contact plan involved in proposed architecture comprises of i) the continuous contacts between research centers and airports; ii) the intermittent contacts between airport storage and airplane storage, with rate equal to the storage writing speed, and contact intervals of one hour prior to the flight time; and iii) the contacts that represent the flights, with propagation delays equal to the flight times. For each of the aforementioned contacts, we have the following structure {FromNode ID, ToNode ID, Contact Start Time, Contact End Time, Transmission rate, Propagation delay}. We assume that each network node has an accurate contact plan knowledge that was obtained using a dissemination protocol, such as the Contact Plan Update Protocol (CPUP) [16], and uses it to build a “routing table” data structure, which is a list of “route lists,” i.e., one route list for every research center in the network.

When a research center initiates the transmission of a data set, it is segregated into multiple bundles, according to the bundle size. For every one of the bundles, the routing algorithm calculates the paths between this research center and destination research center, based on the connectivity plan, and selects the one that achieves the earliest bundle delivery time. The routing procedure is subsequently executed in every network node through the path to destination, where each node recalculates the optimal route towards the data destination, excluding the previously visited node to avoid routing loops.

In this paper, also, we propose a different routing objective, namely cost minimization. Cost represents the transmission value per MB of the data delivered through the entire path from source to destination, including the air flights. If cost for two routes is equal (e.g., two similar flights for two successive days), the routing decision is based on the earliest delivery time. For each of the contacts, we now have the following structure {FromNodeID, ToNode ID, Start Contact Time, End Contact Time, Transmission rate, Propagation delay, Cost}.

IV. EVALUATION METHODOLOGY

In order to evaluate the proposed framework, we consider scenarios with transmissions of bulk space data and compare it with the European Space Agency’s (ESA) internal dedicated network with high-speed connections (1, 2.5, and 10 Gbit/s) [6], which interconnects ESA’s assets inside Europe (shown in Figure 2), with the purpose of transferring a vast amount of space data among them.



Figure 2. ESA’s Data Network, Copernicus [6]

The performance of the proposed architecture is evaluated using the DTN simulator that was originally used in [16]. We assume that the airports in the vicinity of ESA’s research center support the architecture described in Section 3; we also use one extra airport (namely OSL) to support the connection with Svalbard airport. We assemble the contact plan of the formed network, using the FlightStats Web Services API [17], and the set of connected airports, for a certain period of time of 40 days. We assume that the connections between airports and airplanes, as it is mentioned, last one hour prior to the flight time and after landing. Also, we assume that there are no delays in flights, the research centers have continuous connection with the corresponding airport, and the storage capacity in airports is unlimited. The parameters of the proposed architecture are given in Table I.

TABLE I. SIMULATION PARAMETERS

<i>Parameter</i>	<i>Value</i>
Bandwidth of optical fiber	10Gbps
SSD read speed	500Mbps
SSD write speed	377Mbps
Number of ssd	60
Bundle size	10GB

We assess the performance of the proposed architecture in different simulation series, in terms of system throughput achieved, for different data values. At the first set of simulations a single research center in Frankfurt transmits data towards a single receiver in Madrid. The second set includes parallel data transmission from many research centers (specifically 2 and 5) to one. We also evaluate the performance of the proposed system with different routing objectives in mind, alternatively to the earliest delivery delay. Using cost minimization as the routing objective, data routing will not rely on the fastest transmission but on the most economical one. For an initial evaluation, we measure the cost per MB of the flights based on the flight distance; undoubtedly, in a deployed system the cost would depend on more parameters, e.g., the airport infrastructure expenses, the airline policies, the agreement between data providers and airline companies, etc. Finally we compare the data delivery times of the proposed model with ESA’s dedicated network, assuming constant throughput, and with a hybrid model that leverages both internal network and airplane data transmissions. Since the delivery delay (and throughput, respectively) of the airplane network depends mainly on the flight schedules and the transmission initiation time, we used uniformly random transmission start times, at a daily basis, for 30 days, to obtain statistical deviations and associated confidence intervals.

V. EVALUATION RESULTS

Based on the described evaluation method we examine, initially, bulk data transmissions from a single sender to a single receiver. In Figure 3, we illustrate the throughput of the proposed system for different amounts of transmitted data; single-hop transmissions represent the case in which the contact plan

contains only the direct flights between the two airports that reside near the sending and receiving research centers, whereas in multi-hop the contact plan contains the flights between more airports in network, exploiting also intermediate flights. In both cases the achieved throughput increases with the amount of data transmitted, and, for large bulk data sizes (i.e., 100TB), the system throughput approaches the optical fiber speed. As expected, further data increase wouldn't improve the system's throughput, due to the fact that it approaches the optical fiber speed, which is the transmission bottleneck and the throughput's upper bound of the proposed system.

Therefore, since the throughput achieved by the flights is not the bottleneck, there is no discrepancy between single-hop and multi-hop data transmissions.




Figure 3. Throughput over different data sizes with 99% confidence interval

We continue the evaluation of the proposed methodology by considering many-to-one bulk data transmissions, in order to study the effect of multiple, parallel flows on the throughput. In Figure 4, we illustrate the achieved throughput regarding the simultaneous data deliveries from multiple research centers to one; specifically, we have used two and five research centers with each one sending 10TB or 100TB of data towards the same receiving research center. As Figure 4 illustrates, the transmission from multiple senders to one, over the proposed system can be proven quite beneficial; for example the system in case of 2 senders manages to transmit, in more or less the same time, 2 times the amount of data (approximately double throughput) that can be transmitted by only one sender. The main reason for the above enhancement is the usage of an aircraft as transmission link, which allow us to transmit simultaneously (using different aircrafts) a vast amount of data from different resources, where the theoretical limit is the airplane's store capabilities and the corresponding optical fiber between the airport and research center at destination.

In Figure 5, we illustrate the delivery delay per bundle, for 10TB data transmissions; we compare the proposed system with a dedicated Internet link with throughput equal to 1Gbps, as well as with the use of a hybrid system, where the two aforementioned approaches are combined. We observe that the first bundles are delivered faster with the dedicated Internet

link, since the transmission starts right away, rather than wait for the first flight. After 5.8TB, however, data are routed through the air flight link. Exploiting this separation, the bundles can be transferred via the faster available mean, reducing both the transmission time and the congestion on the dedicated Internet link.




Figure 4. Throughput over different number of senders for 10TB and 100TB with 99% confidence interval




Figure 5. Destination delivery times per bundle

In Figure 6, we present the transmission cost and throughput for different amounts of transmitted data, applying cost and delivery time as the routing objective (referred to as CGR_COST and CGR, respectively). The throughput, depicted in bars, is approximately the same in both cases, with a minor improvement at the case where routing decisions are made based on delivery time. The corresponding cost is illustrated by the two lines across different data samples. It is worth noticing that, in cases where we use as routing criterion the cost, the method achieves to maintain a constant cost/MB ratio, in contrast to the second case where the routing criterion does not have any relation with the flights' cost, thus, as depicted in

Figure 6, the overall transmission cost may vary, depending on the specific test case's available flights.




Figure 6. Throughput over different data sizes and corresponding cost per MB with 99% confidence interval

However, the cost in second case seems to have only small fluctuations, in comparison to the first one, for mainly two reasons: First, the evaluation is based on the distances of the European airports that form the aforementioned ESA network, which renders costs among them quite similar. Second, the number of scheduled flights is restrictive, due to the usage of only specific airports. In other words, the bundles do not have the opportunity to follow multiple routes in order to decrease the delivery delay at the expense of cost.

VI. CONCLUSION

In this paper, we have described a network architecture for bulk data transfers, by using airline infrastructure. This architecture exploits Delay-Tolerant Networking and Contact Graph Routing algorithm providing an automated way to efficiently transmit large amounts of data. We evaluate the proposed network via a set of simulations and compare the throughput of the proposed architecture with a dedicated network of ESA, in order to show that an acceptable level of service (in terms of throughput) can be provided. Furthermore, a possible combination of the proposed system, along with the existing, internal network infrastructure could further reduce transmission times and congestion on the dedicated links. Finally, we demonstrate the impact of the proposed approach when motivation is the minimization of data transmission cost; we incorporate it as an alternative objective of our routing policy.

Our future work has several dimensions. One critical dimension is to expand the architecture per se and cancel the restrictive assumption that research centers are located nearby airports – this is doable via a combination of bus or train transportation service. Another recent dimension is derived by the motivations of Future Internet and, in particular, of UMOBILE project [18], where localized access to data necessitates, in some occasions, transmissions of large amounts of data over

different network architectures, complementary or alternative to IP.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Union's (EU) Horizon 2020 research and innovation programme under grant agreement No 645124 (Action full title: Universal, mobile-centric and opportunistic communications architecture, Action Acronym: UMOBILE). This paper reflects only the authors' views and the Community is not liable for any use that may be made of the information contained therein.

REFERENCES

- [1] V. Marx, "Biology: The big challenges of big data." *Nature* 498.7453, pp. 255-260 (2013).
- [2] The Economist, "Data, data everywhere". [Online]. Available from: <http://www.economist.com/node/15557443> 2015.09.25
- [3] European Space Agency (ESA), Euclid mission. [Online]. Available from: <http://sci.esa.int/euclid/46661-mission-operations/> 2015.09.25
- [4] N. Skytland, "Big Data", open.NASA blog. [Online]. Available from: <http://open.nasa.gov/blog/2012/10/04/what-is-nasa-doing-with-big-data-today/> 2015.09.25
- [5] S. Burleigh, V. Cerf, J. Crowcroft, and V. Tsaoussidis, "Space for Internet and Internet for Space", Elsevier Ad Hoc Networks, Special Issue on New Research Challenges in Mobile, Opportunistic & Delay-Tolerant Networks, vol. 23, pp. 80–86, December 2014.
- [6] G. Buscemi, "The new dimension of the Copernicus Data Network", NETSPACE Workshop, 2014.
- [7] A. M. Zarafshan, TrainNet: A Novel Transport Infrastructure for Non Real-Time Data Delivery. University of Wollongong, 2009.
- [8] S. Guo, M.H. Falaki, E.A. Oliver, S. Ur Rahman, A. Seth, M.A. Zaharia, and S. Keshav, "Very low-cost internet access using KioskNet." *ACM SIGCOMM Computer Communication Review* 37.5, pp. 95-100 (2007).
- [9] I. Komnios and V. Tsaoussidis, "CARPOOL: Connectivity Plan Routing Protocol", 12th International Conference on Wired & Wireless Internet Communications (WWIC 2014), Paris, France, May 26-28, 2014.
- [10] R. A. Gorcitz et al., "Vehicular carriers for big data transfers (poster)." *Vehicular Networking Conference (VNC)*, 2012 IEEE. IEEE, 2012.
- [11] A. Keränen and J. Ott, "DTN over aerial carriers." *Proceedings of the 4th ACM workshop on Challenged networks*. ACM, 2009.
- [12] B. Cho and I. Gupta, "Budget-constrained bulk data transfer via internet and shipping networks." *Proceedings of the 8th ACM international conference on Autonomic computing*. ACM, 2011.
- [13] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, Delay-Tolerant Networking Architecture. RFC 4838 (Informational), Apr 2007.
- [14] S. Burleigh, Contact graph routing. Tech. Rep. draft-burleigh-dtnrg-cgr-01, Network Working Group, Jul 2010.
- [15] G. Araniti, N. Bezirgiannidis, E. Birrane, I. Bisio, S. Burleigh, C. Caimi, M. Feldmann, M. Marchese, J. Segui, and K. Suzuki, "Contact graph routing in DTN space networks: overview, enhancements and performance," *Communications Magazine*, IEEE, vol.53, no.3, pp.38,46, March 2015.
- [16] N. Bezirgiannidis, F. Tsapeli, S. Diamantopoulos, and V. Tsaoussidis, "Towards Flexibility and Accuracy in Space DTN Communications", in *8th ACM MobiCom Workshop on Challenged Networks*, (CHANTS'13), Miami, Florida, USA, September 30, 2013.

- [17] Flightstats. [Online]. Available from: https://developer.flightstats.com/products/scheduled_flights 2015.09.25
- [18] Umobile European project. [Online]. Available from: <http://www.umobile-project.eu/index.php/4-the-network-of-excellence-on-internet-science> 2015.09.25