Congestion Control in Disruption-Tolerant Networks: A Comparative Study for Interplanetary Networking Applications

Aloizio P. Silva
Inst. Tec. de Aeronáutica
São José dos Campos, Brasil
aloizio@ita.br

Scott Burleigh
Jet Propulsion Laboratory,
California Institute of Technology
Pasadena, California, USA
scott.c.burleigh@jpl.nasa.gov

Celso M. Hirata
Inst. Tec. de Aeronáutica
São José dos Campos, Brasil
hirata@ita.br

Katia Obraczka
UC Santa Cruz
Santa Cruz, California, USA
katia@soe.ucsc.edu

ABSTRACT

Controlling congestion is critical to ensure adequate network operation and performance. That is especially the case in networks operating in "challenged"- or "extreme" environments where episodic connectivity is part of the network's normal operation. Our goal is to study congestion control mechanisms that have been proposed for these so-called disruption tolerant networks, or DTNs. In this paper, we conduct a performance study comparing existing DTN congestion control mechanisms for the specific case of interplanetary networking (IPN) applications. Our results indicate that congestion control helps increase message delivery ratio, even in highly congested network scenarios. Our study also suggests that good design principles for congestion control in IPN scenarios include: using a combination of reactive- and proactive control, using local information instead of relying on global knowledge, and employing mechanisms that are routing protocol independent.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Network Architecture and Design; C.2.1 [**Network communications**]: Store and forward networks—*wireless communication*

General Terms

Performance, Experimentation

Keywords

Delay and Disruption Tolerant Networks; Interplanetary Networks; Congestion Control; Network Performance

1. INTRODUCTION

In the last 10 years, applications such as environmental sensing, habitat monitoring, emergency response, disaster recovery, bridging the digital divide, to name a few have raised great interest in so-called challenged network environments. In such environments, also known as delay and disruption tolerant networks, or DTNs, continuous end-to-end connectivity cannot be guaranteed and the communication channel may be subject to arbitrarily long delays and high error rates. Under these conditions, participating nodes must store data they are transmitting or forwarding in persistent storage until a "contact" opportunity, i.e., the node has a suitable next-hop neighbor that can receive the data, arises.

Congestion control in challenged environments is thus critical in order to ensure nodes are congestion-free and can serve as relays when needed so messages can be delivered end-to-end. Because DTNs violate the fundamental assumptions underlying the TCP/IP protocol architecture, namely the existence of an end-to-end path between nodes and short delays, they cannot employ the principles underlying Internet congestion control mechanisms.

Due to its impact on performance and the challenges it raises, DTN congestion control has motivated a number of research efforts [3] [6] [9]. This work focuses on DTN in deep space communications, also known as Interplanetary Internet (IPN) [1] [2]. We conduct a study comparing the performance of different congestion control schemes in IPN scenarios. Our goal is to understand the performance tradeoff issues raised by state-of-the-art DTN congestion control mechanisms and how these mechanisms behave when different routing protocols are used.

The paper is organized as follows. In Section 2, we provide a brief overview of the congestion control schemes we studied. Section 3 describes our experimental methodology

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHANTS'14, September 7, 2014, Maui, Hawaii, USA.
Copyright 2014 ACM 978-1-4503-3071-8/14/09 ...$15.00.
http://dx.doi.org/10.1145/2645672.2645684.

while Section 4 presents the results of our comparative study. Finally, Section 5 concludes the paper.

2. BACKGROUND: SELECTED DTN CONGESTION CONTROL MECHANISMS

For our comparative study, we picked a subset of DTN congestion control mechanisms from the schemes we surveyed in [3] that we consider representative of the current DTN congestion control state-of-the-art: RRCC [8], AFNER [10], SR [7] and CCC [5]. We consider the performance of each of these mechanisms in the context of three different routing protocols - Epidemic, PRoPHET, and Spray-and-Wait - except that AFNER is only considered in an Epidemic routing environment, as it is specific to that routing protocol.

Retiring replicants Congestion Control (RRCC) [8] looks to understanding the global behavior of congestion in intermittently connected networks and uses this knowledge to control congestion at a local level. RRCC adjusts replication rates at individual nodes to maximize delivery rates.

Average forwarding number based on epidemic routing (AFNER) [10] uses a message's *forwarding number*, defined as the message's number of copies to decide whether to drop a copy of the message. A node randomly drops a message whose forwarding number is larger than the average forwarding number of the network. This average forwarding number is defined as the mean forwarding number of all the messages currently in the network. Although AFNER uses the average forwarding number, this quantity is typically not easily computable in a real DTN. In the simulation point of view, we have computed this parameter consulting each node buffer in the network.

Storage routing (SR) [7] migrates messages to alternate storage locations under congestion. By having congested nodes send a set of messages out to available neighbors, SR operates as a local routing protocol diverting messages from their conventional routing path for later forwarding.

Credit-Based Congestion Control (CCC) [5] employs a reactive, heuristic-based congestion control policy. In order to yield high delivery rate with low number of replicas, CCC tries to delete messages when congestion builds up at a node. Messages are dropped when deemed obsolete by their *time-dependent credit*.

3. EXPERIMENTAL METHODOLOGY

In our study, we use the ONE DTN simulator [4] and the deep space networking environment depicted in Figure1.

3.1 IPN Scenarios

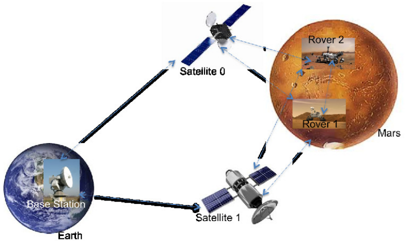

Figure 1: Interplanetary network scenario.

The experimental scenario consists of an IPN which is characterized by high latencies (because of astronomical distances) and *scheduled* contacts, i.e., node encounters that are known a priori. There are five nodes representing a Base Station on the surface of the Earth which sends data to the Rovers in Mars through two satellites located near Mars. Simulation parameters are set to correspond to realistic conditions. As such, links between the satellites and the rovers on the Martian surface are set to 1s propagation delay while the links between the base station on Earth and the satellites at Mars are set to 240s propagation delay.

We use a scheduled contact table (see Table 1) which specifies when a connection between two nodes is established "up") and when the contact between the nodes ceases to exist ("down"). For example, according to the first line of Table 1, at time $2000s$, the connection between *Base Station* and *Satellite 0* is "up" but goes "down" at time $3000s$. Then, the same nodes are again in contact at time $10,000s$ (line 5) for $1,000s$ (line 6).

Time (s)	Identifier	Initial Node	End Node	State
2000	CONN	Base Station	Satellite 0	up
3000	CONN	Base Station	Satellite 0	down
6000	CONN	Base Station	Satellite 1	up
7000	CONN	Base Station	Satellite 1	down
10000	CONN	Base Station	Satellite 0	up
11000	CONN	Base Station	Satellite 0	down

Table 1: Sample of scheduled contact table

Different scheduled contact tables (Contact1, Contact2, Contact3, Contact4, and Contact5) were used to vary the time between contacts.

The parameters of the ONE simulator and their values are listed in Table 2. A scenario without congestion control is considered as baseline for our comparative study. Every pair of nodes that are in communication range of one another according to the scheduled contacts can transfer data between them if they have data to send.

Nodes generate messages according to the $Events1.interval$ parameter (see Table 2). For example, when $Events1.interval$ is between 1 and 5 seconds, one new message is generated every 1 to 5 seconds. We vary the message generation rate according to the values listed in Table 2 to show how this parameter affects the performance of the different congestion control schemes.

4. SIMULATION RESULTS

The main performance metric for our experiments is the *Delivery ratio*, defined as the ratio between the number of received messages at destination nodes to the number of created messages. We use 95% confidence intervals and present as baseline results when no congestion control is employed. Furthermore, each configuration based on Table 2 was run 5 *times* and the statistical values (average, standard deviation) were obtained. Note that some graphics do not show the standard deviation either because the value is zero or approximately zero.

To evaluate the impact of congestion control, we generate enough load to congest the network. Figure 2 shows message delivery ratio as a function of the message generation period. As expected, for longer message generation periods, i.e., lower message generation frequencies, the delivery ratio increases. However, delivery ratios are quite low since newly generated messages cause older messages to be drop-

Table 2: Simulation parameters and their values

Parameters		
Name	Description	Value
Scenario.endTime	simulation time	43200 seconds
btInterface.transmitSpeed	bandwidth	[0.5, 1.0, 1.5, 2.0, 2.5] Mbps
btInterface.transmitRange	transmitting range	150 m
Group.router	routing protocol	[EpidemicRouter, ProphetRouter, SprayAndWaitRouter]
Group.movementModel	mobility model	StationaryMovement
Group.bufferSize	node buffer size	[1000, 2000, 3000, 4000, 5000, 6000, 7000, 8000] KB
Group.msgTTL	message time to live	43200 seconds
Group.nrofHosts	number of nodes in network	5
Movimentmodel.worldSize	area where simulation takes place	6km x 6km
Events1.size	message size	{50, 100} KB
Events1.interval	i.e. one new message every 1 to 5 seconds	[1-5, 1-10, 1-20, 1-30, 1-40, 1-50] seconds

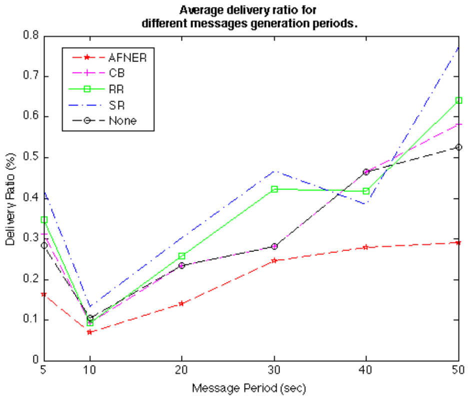

Figure 2: Message delivery ratio as a function of the message generation period. We use Epidemic routing, buffer size of 8000k and transmit speed of 2.5Mbps.

ped before they have time to be delivered. At this level of congestion, even though message delivery ratio is quite low, congestion control plays a role and is able to improve delivery.

4.1 IPN Scenario Without Congestion Control

We report here the results obtained using an IPN scenario that does not include any congestion control mechanism. Tables 1 and 2 list the simulation parameter settings.

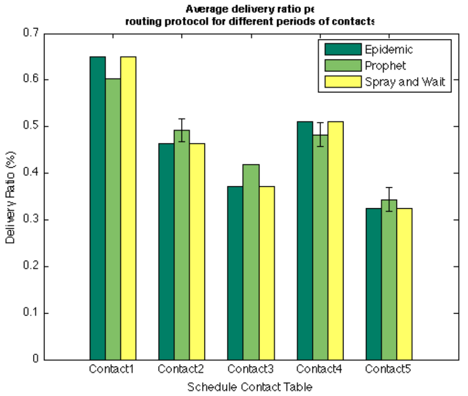

Figure 3: Delivery ratio for IPN scenario without congestion control (different periods between contacts and routing protocols).

Figure 3 confirms that for longer periods of disconnection between contacts, the delivery ratio decreases. This happens because of the fewer transmission opportunity which result in longer data queues and consequently higher probability of data being dropped as node buffers fill up. Note that different routing protocols do not seem to have significant impact on performance.

4.2 IPN Scenario With Congestion Control

4.2.1 Different Contact Periods

Figure 4 shows the average delivery ratio for different contact periods, where the slim bars refer to delivery ratio values when no congestion control is used. We observe that CCC's performance in terms of delivery ratio (Figure 4b) is quite similar to SR's (Figure 4d). Note that RRCC is routing protocol independent since its performance is not affected by using different routing protocols (see Figure 4c). On the other hand, CCC's and SR's delivery ratio see a slight increase when ProPHET is used as the underlying routing protocol. Our hypothesis is that, since ProPHET bases its routing decisions on past contact history, it benefits from scenarios where contacts are scheduled, which is the case of IPN environments. We argue that, due to the high congestion levels, the increase in delivery ratio is not substantial and is not consistent across all contact scenarios. We will run additional experiments to verify these hypothesis.

It is worth pointing out that, except for AFNER, for longer disconnection periods, the congestion control mechanisms yield higher delivery ratios when compared to no congestion control. AFNER uses the network's average forwarding number to mitigate congestion. Thus when the periods between contacts increase, more messages are awaiting to be forwarded. As a result, buffers fill up and AFNER starts to discard messages based on the average forwarding number. This leads to lower delivery ratios.

5. CONCLUSIONS

This paper evaluated the performance of four DTN congestion control mechanisms which represent the state-of-the-art on DTN congestion control based on the survey we present in [3]. Our comparative study focuses specifically on deep space communication scenarios and, to our knowledge, is the first of its kind. We examined two performance metrics, namely: delivery ratio and latency and also evaluated the congestion control schemes in terms of their routing protocol independence.

Our results indicate that congestion control helps increase message delivery ratio, even in highly congested network scenarios. Our study also indicates that good design principles for congestion control in IPN scenarios include: using a com-

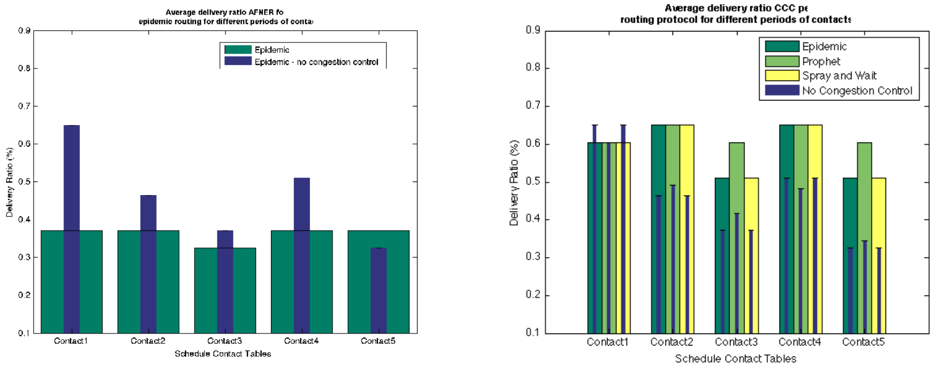

(a) AFNER (b) CCC

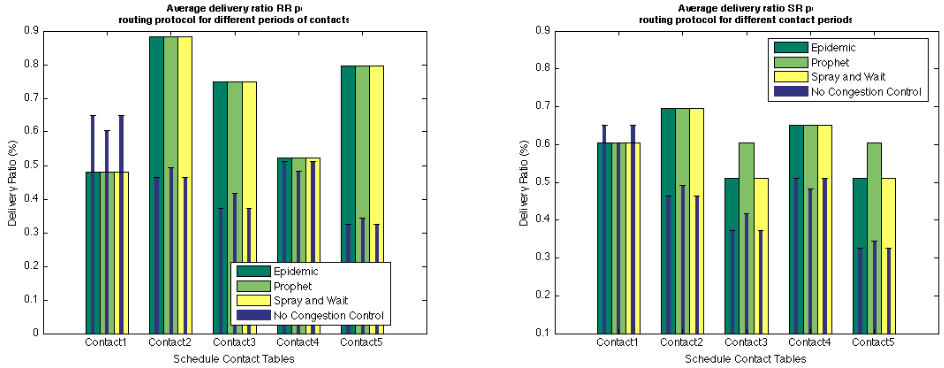

(c) RRCC (d) SR

Figure 4: Average delivery ratio for different periods between contacts.

bination of reactive- and proactive control as well as using local information instead of relying on global knowledge. Additionally, designing congestion control mechanisms that are routing protocol independent helps with interoperability and applicability to a wide variety of DTN scenarios.

6. ACKNOWLEDGMENTS

Part of the research discussed in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. This work was partially funded by NSF under project CNS 1321151.

7. REFERENCES

[1] I. Akyildiz, O. Akan, C. Chen, and W. S. Jian Fang. The state of the art in interplanetary internet. *IEEE Communications Magazine*, 42(7):108–118, July 2004.

[2] G. Araniti, I. Bisio, and M. D. Sanctis. Interplanetary networks: Architectural analysis technical challenges and solutions overview. In *IEEE ICC 2010 Proceedings*, pages 1–5, May 2010.

[3] A. P. da Silva, S. Burleigh, C. M. Hirata, and K. Obraczka. A survey on congestion control for delay and disruption tolerant networks. *Elsevier Ad Hoc Network, Special Issue on New Research Challenges in Mobile, Opportunistic and delay-Tolerant Networks*, 2014. Will appear in Elsevier Ad Hoc Network.

[4] A. Keranen, J. Ott, and T. Karkkainen. The ONE simulator for DTN protocol evaluation. In *The 2nd International Conference on Simulation Tools and Techniques*, 2009.

[5] L. Leela-amornsin and H. Esaki. Heuristic congestion control for massage deletion in delay tolerant network. In *Smart Spaces and Next Generation Wired/Wireless Networking*, August 2010. Third Conference on Smart Spaces and 10th International Conference.

[6] I. Psaras, L. Wood, and R. Tafazolli. Delay-/disruption-tolerant networking: State of the art and future challenges. Technical report, Center for Communication Systems Research, Department of Eletrical Engeneering , University of Surrey, 2009.

[7] M. Seligman, K. Fall, and P. Mundur. Alternative custodians for congestion in delay tolerant networks. In *Proceedings of the 2006 SIGCOMM workshop on Challenged Networks - CHANTS'06*, 2006.

[8] N. Thompson, S. C. Nelson, M. Baknt, T. Abdelzaher, and R. Kravets. Retiring replicants: Congestion control for intermittently - connected networks. In *Proceedings of the IEEE INFOCOM'2010*, pages 1–9, March 2010.

[9] A. G. Voyiatzis. A survey od delay and disruption tolerant networking applications. *Internet Engineering*, 5(1), 2012.

[10] L. Yun, C. Xinjian, L. Qilie, and Y. Xianohu. A novel congestion control strategy in delay tolerant networks. In *Second International Conference on Future Networks*, 2010.