Message efficient Byzantine Reliable Broadcast protocols on known topologies

Tim Anema - *under supervision of* **Jérémie Decouchant**

TU Delft

1 Introduction

Distributed processes need to communicate with each other, and possibly rely on each other to do so. Malicious (Byzantine) processes can hinder this process by dropping or modifying messages, or impersonating other processes.

Protocols to reliably broadcast messages in these harsh environments exist, but use a lot of messages in the process:

1. Dolev [1], which floods messages over disjoint paths in well-connected networks
2. Bracha [2], which only works in fully connected networks but guarantees stronger properties
3. Bracha-Dolev [3], which combines the previous two to get the best of both worlds

2 Research question

Can we reduce the message complexity for BRB protocols such as Dolev, Bracha, and Bracha-Dolev when processes are aware of the network topology?

3 Example

Messages travel over 2F+1 disjoint paths from source to every sink. This example shows the paths from a to f.

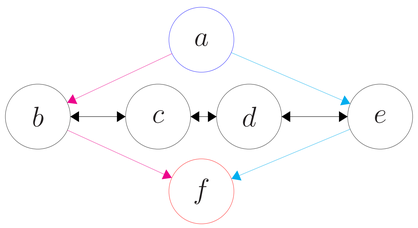

1. Path a-e-f is dropped in favor of a-e-f-b (Contr. 1)
2. Path a-b immediately causes b to deliver (Contr. 2)
3. Paths a-e-d-c travel together to e (Contr. 3)

5 Results

We compared our improved variants to ones using naive routing, on random regular networks, using small (12B) payloads.

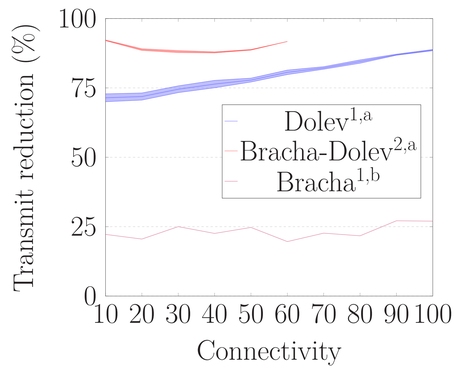

Figure 1: Reduction of message complexity using K-random graphs and fully-connected graphs (Bracha), while varying the connectivity. [1]$N = 150$, [2]$N = 75$, [a]$f = \lfloor \frac{k-1}{2} \rfloor$, [b]$f = \lfloor \frac{k}{4} \rfloor$

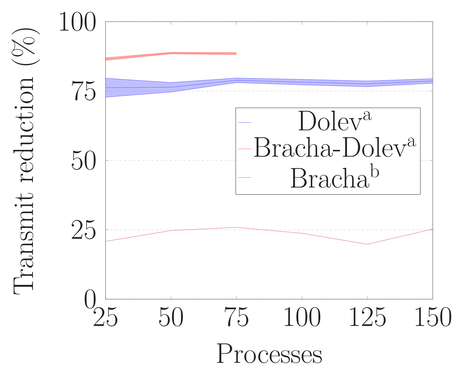

Figure 2: Reduction of message complexity using K-random graphs and fully-connected graphs (Bracha), while varying the number of processes. $k = \lfloor \frac{N}{3} \rfloor$, [a]$f = \lfloor \frac{k-1}{2} \rfloor$, [b]$f = \lfloor \frac{k}{4} \rfloor$

We observed a mean reduction of **79.5%** and **85.86%** for Dolev, **23.3%** for Bracha, and **89.54%** and **92.32%** for Bracha-Dolev in terms of message complexity and bandwidth usage, respectively.

4 Contributions

We introduced the following optimizations:

1. Avoid transmitting subpaths (Dolev)
2. Use a single hop for direct neighbours (Dolev)
3. Merge next hops when possible (Dolev)
4. Reuse paths when possible (Dolev)
5. Merge messages in transit (Dolev)
6. Merge identical payloads (Dolev)
7. Use implict paths (Dolev)
8. Use a subset of processes for agreement (Bracha)
9. Use implicit echo messages (Bracha)
10. Use partial Dolev broadcasts (Bracha-Dolev)
11. Merge related Bracha-Dolev messages (Bracha-Dolev)

6 Conclusion and Future work

We have introduced several optimizations to all three protocols and showed that we can indeed drastically reduce the number of messages transmitted and the network usage when we leverage topology knowledge.

Future iterations of this research might want to focus on:

- Improving evaluation capabilities on real networks
- Optimizing the disjoint path solver
- Combining this research with topology discovery
- Extending contribution 11 by also applying it on the Bracha layer
- Modify our protocol to also work on dynamic networks

Tim Anema
t.p.d.anema@student.tudelft.nl

[1] S. Bonomi, G. Farina, and S. Tixeuil, "Multi-hop byzantine reliable broadcast with honest dealer made practical," 2019.

[2] "Asynchronous byzantine agreement protocols," *Information and Computation*, vol. 75, no. 2, pp. 130–143, 1987.

[3] S. Bonomi, J. Decouchant, G. Farina, V. Rahli, and S. Tixeuil, "Practical byzantine reliable broadcast on partially connected networks," 2021.