Brief Announcement: The Fault-Tolerant Cluster-Sending Problem

Jelle Hellings Mohammad Sadoghi

Exploratory Systems Lab,
Department of Computer Science,
University of California, Davis, CA, USA

UC DAVIS
UNIVERSITY OF CALIFORNIA

Vision: resilient high-performance data processing

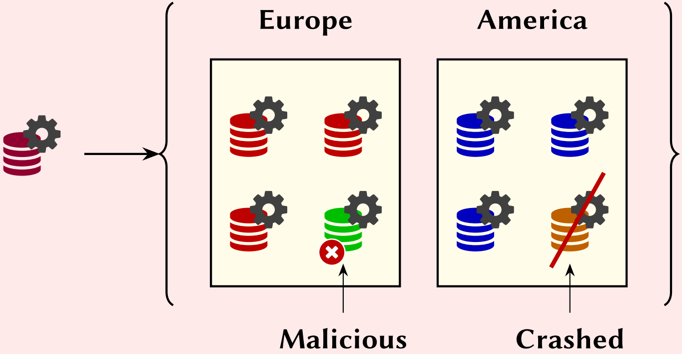

Requirement

Fault-tolerant communication between clusters!

The need for cluster-sending

Definition

The *cluster-sending problem* is the problem of sending a value v from $\mathcal{C}_1$ to $\mathcal{C}_2$ such that:

1. all non-faulty replicas in $\mathcal{C}_2$ *receive* the value v;
2. only if all non-faulty replicas in $\mathcal{C}_1$ *agree* upon sending the value v to $\mathcal{C}_2$ will non-faulty replicas in $\mathcal{C}_2$ receive v; and
3. all non-faulty replicas in $\mathcal{C}_1$ can *confirm* that the value v was received.

Solution (crash failures)

Pair-wise broadcasting with $(\mathbf{f}_1 + 1)(\mathbf{f}_2 + 1)$ messages.

Lower bounds for cluster-sending: Example

$$\mathbf{n}_1 = 15 \qquad \mathbf{f}_1 = 7$$

$$\mathbf{n}_2 = 5 \qquad \mathbf{f}_2 = 2$$

Claim (crash failures)

Any correct protocol needs to send at least 14 messages.

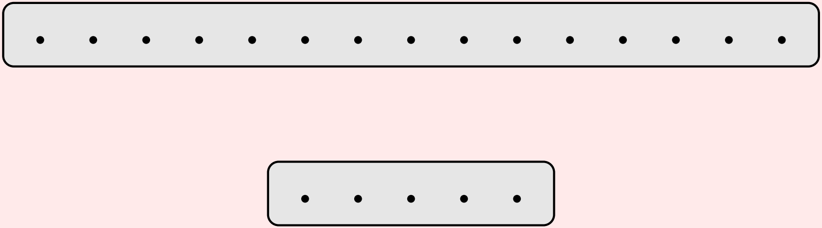

Lower bounds for cluster-sending: Example

$\mathbf{n}_1 = 15$ $\mathbf{f}_1 = 7$

$\mathbf{n}_2 = 5$ $\mathbf{f}_2 = 2$

Claim (crash failures)

Any correct protocol needs to send at least 14 messages.

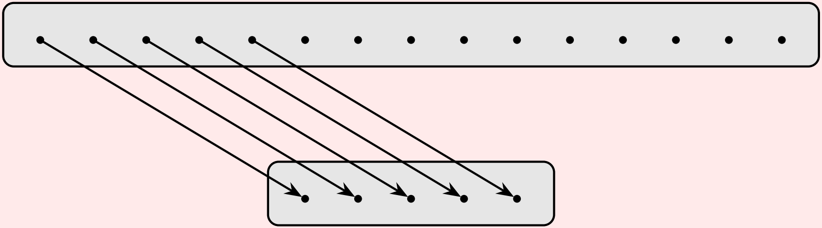

Lower bounds for cluster-sending: Example

$\mathbf{n}_1 = 15$ $\mathbf{f}_1 = 7$

$\mathbf{n}_2 = 5$ $\mathbf{f}_2 = 2$

Claim (crash failures)

Any correct protocol needs to send at least 14 messages.

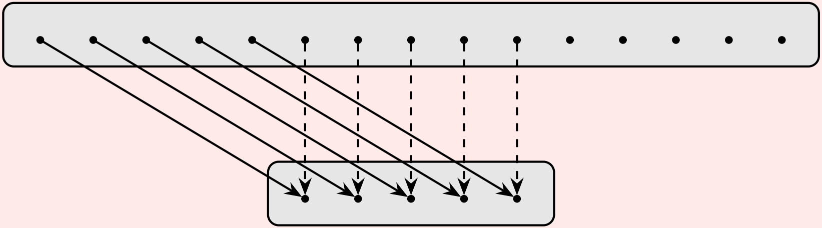

Lower bounds for cluster-sending: Example

$\mathbf{n}_1 = 15$ $\mathbf{f}_1 = 7$

$\mathbf{n}_2 = 5$ $\mathbf{f}_2 = 2$

Claim (crash failures)

Any correct protocol needs to send at least 14 messages.

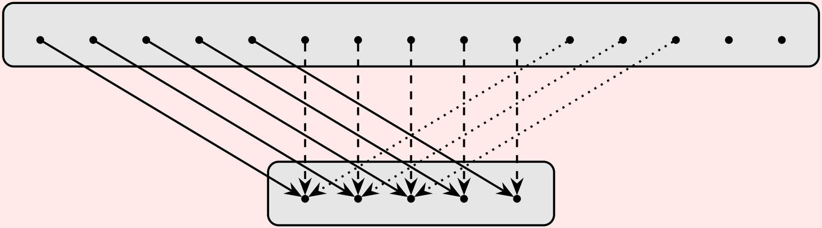

Lower bounds for cluster-sending: Example

$\mathbf{n}_1 = 15$ $\mathbf{f}_1 = 7$

$\mathbf{n}_2 = 5$ $\mathbf{f}_2 = 2$

Claim (crash failures)

Any correct protocol needs to send at least 14 messages.

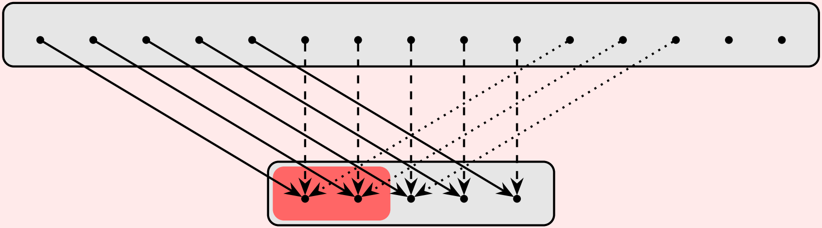

Lower bounds for cluster-sending: Example

$\mathbf{n}_1 = 15$ $\mathbf{f}_1 = 7$

$\mathbf{n}_2 = 5$ $\mathbf{f}_2 = 2$

Claim (crash failures)

Any correct protocol needs to send at least 14 messages.

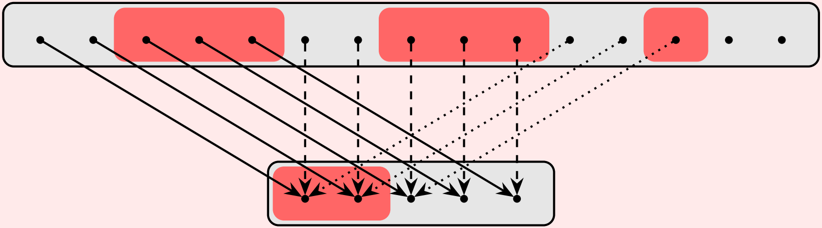

Lower bounds for cluster-sending: Results

Theorem

Assume $\mathbf{n}_1 \geq \mathbf{n}_2$ *and let*

$$\begin{aligned} q &= (\mathbf{f}_1 + 1) \operatorname{div} \mathbf{nf}_2; \\ r &= (\mathbf{f}_1 + 1) \bmod \mathbf{nf}_2; \\ \sigma &= q\mathbf{n}_2 + r + \mathbf{f}_2 \operatorname{sgn} r. \end{aligned}$$

We need to exchange at least σ *messages to do cluster-sending.*

- Similar results for $\mathbf{n}_1 \leq \mathbf{n}_2$.
- Byzantine failures: similar lower bounds on signatures.

An optimal algorithm

Protocol for the sending cluster $\mathcal{C}_1$, $\mathbf{n}_1 \geq \mathbf{n}_2$, $\mathbf{n}_1 \geq \sigma$:

```
Choose replicas 𝒫 ⊆ 𝒞₁ with n_𝒫 = σ.
Choose a n₂-partition partition(𝒫) of 𝒫.
for P ∈ partition(𝒫) do
    Choose replicas Q ⊆ 𝒞₂ with n_Q = n_P.
    Choose a bijection b : P → Q.
    for R₁ ∈ P do
        Send v from R₁ to b(R₁).
```

- Crash failures: $\mathbf{n}_1 > 3\mathbf{f}_1$ and $\mathbf{n}_2 > 3\mathbf{f}_2$.
- Byzantine failures:
  - using signatures: $\mathbf{n}_1 > 4\mathbf{f}_1$ and $\mathbf{n}_2 > 4\mathbf{f}_2$;
  - using threshold signatures: $\mathbf{n}_1 > 3\mathbf{f}_1$ and $\mathbf{n}_2 > 3\mathbf{f}_2$.

Conclusion

More information

`https://jhellings.nl`

Paper: https://doi.org/10.4230/LIPIcs.DISC.2019.45

Technical Report: https://arxiv.org/abs/1908.01455