Basilic: Resilient Optimal Consensus Protocols With Benign and Deceitful Faults

Alejandro Ranchal-Pedrosa[†,‡] Vincent Gramoli[†,§]

[†]University of Sydney, Sydney, Australia

[‡] Protocol Labs

[§]EPFL, Red Belly Network

Small Council

5 people, 2 Byzantine -> lose throne

Small Council

5 people, 1 Deceitful, 1 non-responsive -> remove deceitful, 4 with 1 non-responsive

Byzantine Generals Problem

Consensus problem:

- Agreement
- Termination
- Validity

Impossibilities [LSP82, DLS88]

- Consensus only possible if t<n/3 (partial synchrony)
- Byzantine faults? meaning?
  - Worst type of fault
  - If non-responsive is worse for protocol -> non-responsive
  - If protocol-specific disagreement attack -> then that
  - Byzantine faults are important, but what if...

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Previous heterogeneous models

- Crash-faults and Byzantines

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Previous heterogeneous models

- Crash-faults and Byzantines
- Byzantine-altruistic-rational Model

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Previous heterogeneous models

- Crash-faults and Byzantines
- Byzantine-altruistic-rational Model
- (k,t)-robust equilibria

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Previous heterogeneous models

- Crash-faults and Byzantines
- Byzantine-altruistic-rational Model
- (k,t)-robust equilibria
- Commission and omission faults

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Previous heterogeneous models

- Crash-faults and Byzantines
- Byzantine-altruistic-rational Model
- (k,t)-robust equilibria
- Commission and omission faults
- Alive-but-corrupt model

Heterogeneous Faults

- What if not all faults in the system are the worst possible fault?

Goal

- Exploit potential heterogeneity of faults for greater tolerance
- Backwards compatibility: $t<n/3$ if only Byzantines must be ensured

Previous heterogeneous models

- Crash-faults and Byzantines
- Byzantine-altruistic-rational Model
- (k,t)-robust equilibria
- Commission and omission faults
- Alive-but-corrupt model
- No previous works make a disjoint distinction between faults that attack agreement and faults that attack termination

Byzantine-deceitful-benign (BDB) model

- Byzantine faults $t \rightarrow$ arbitrary
- Deceitful faults $d \rightarrow$ target agreement
  - Can prevent termination if trying to cause disagreement and failing, but always reply.
- Benign faults $q \rightarrow$ can only prevent termination
  - Crash-faults, invalid messages etc.
- quorum size $h \rightarrow$ greater for agreement, lower for termination

BDB Impossibilities

- Impossible to tolerate t Byzantine, d deceitful and q benign processes if $n \leq 3t + d + 2q$.

BDB Impossibilities

- Impossible to tolerate t Byzantine, d deceitful and q benign processes if $n \leq 3t + d + 2q$.
- At most $d + t < 2h - n$ and $q + t \leq n - h$, with $h \in (n/2, n]$.

BDB Impossibilities

- Impossible to tolerate t Byzantine, d deceitful and q benign processes if $n \leq 3t + d + 2q$.
- At most $d + t < 2h - n$ and $q + t \leq n - h$, with $h \in (n/2, n]$.

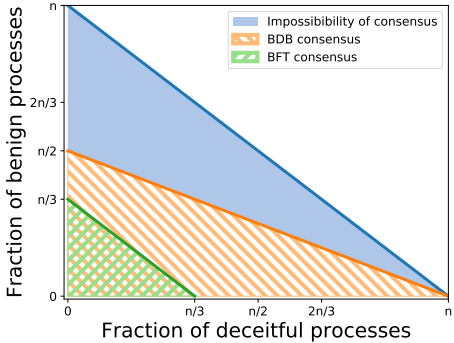

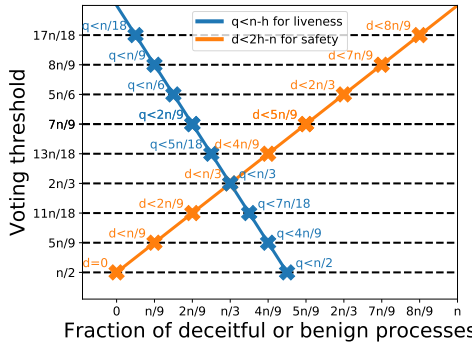

Basilic

Accountability

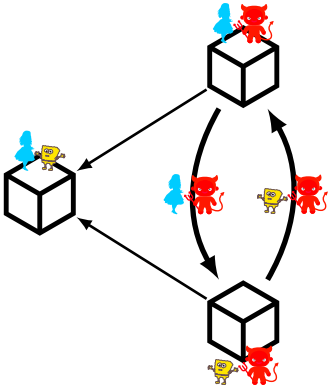

Basilic

Accountability

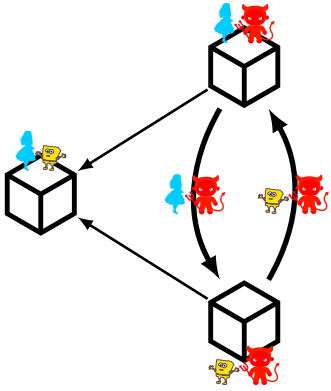

If  attacks agreement property,

then  is caught. But... it could be too late.

Basilic

Accountability

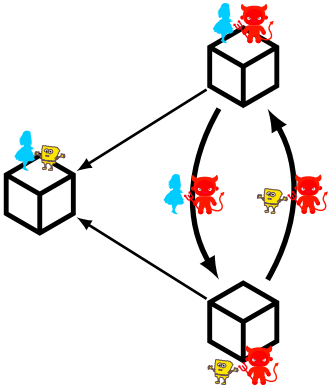

If  attacks agreement property,

then  is caught. But... it could be too late.

Active accountability

- Deceitful faults do not prevent termination

Basilic class

- Basilic: class of consensus protocols
  - Satisfy active accountability:
    - Periodically exchange messages after δ in order to dynamically remove deceitful faults, reducing quorum size accordingly to terminate

Basilic class

- Basilic: class of consensus protocols
  - Satisfy active accountability:
    - Periodically exchange messages after δ in order to dynamically remove deceitful faults, reducing quorum size accordingly to terminate
  - Same code, but protocols of the class change by the initial threshold h_0 given as parameter

Basilic class

- Basilic: class of consensus protocols
  - Satisfy active accountability:
    - Periodically exchange messages after δ in order to dynamically remove deceitful faults, reducing quorum size accordingly to terminate
  - Same code, but protocols of the class change by the initial threshold h_0 given as parameter
  - At any given time, Basilic(h_0) has a dynamic quorum size $h(d_r)=h_0-d_r$

Basilic class

- Basilic: class of consensus protocols
  - Satisfy active accountability:
    - Periodically exchange messages after δ in order to dynamically remove deceitful faults, reducing quorum size accordingly to terminate
  - Same code, but protocols of the class change by the initial threshold h_0 given as parameter
  - At any given time, Basilic(h_0) has a dynamic quorum size $h(d_r)=h_0-d_r$

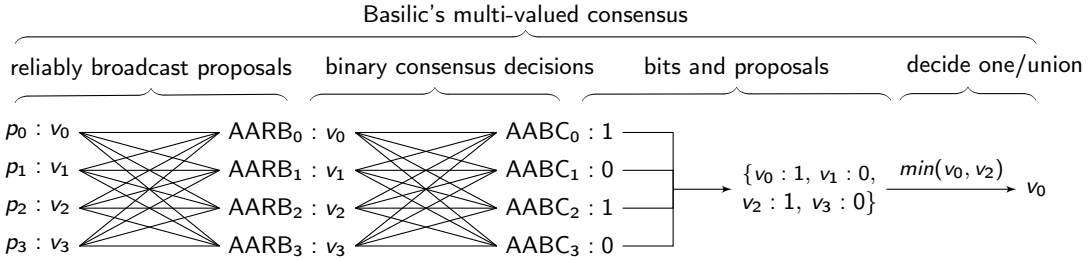

Basilic class' BDB tolerance

Theorem

The Basilic protocol with initial threshold h_0 solves consensus for $d + t < 2h_0 - n$ and $q + t \leq n - h_0$.

Basilic class' BDB tolerance

Theorem

The Basilic protocol with initial threshold h_0 solves consensus for $d + t < 2h_0 - n$ and $q + t \leq n - h_0$.

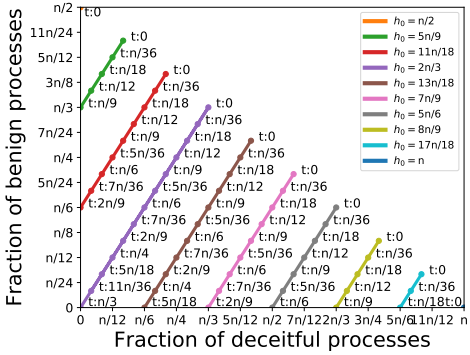

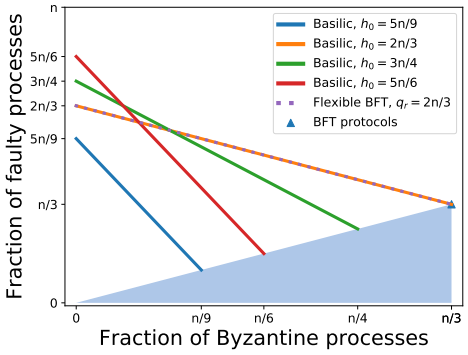

Eventual consensus ($\diamond$-consensus)

Temporary disagreement, but eventual agreement.

Theorem

The $\diamond$-Basilic protocol with initial threshold h_0 solves the $\diamond$-consensus problem if $d + t < h_0$ and $q + t < n - h_0$.

Complexities

- Active accountability has no increase on communication complexity compared to accountability.
- Accountability requires $\mathcal{O}(n^3)$ if deceitful behavior causes disagreement and $\mathcal{O}(n^2)$ otherwise (optimal for consensus).
- Same for active accountability: $\mathcal{O}(n^3)$ if deceitful behavior causes disagreement OR prevents liveness, and $\mathcal{O}(n^2)$ otherwise (optimal for consensus).

Conclusion

- BDB model exploits for heterogeneity of faults, without any real losses in classical BFT model (same complexities, same tolerances, no changes to protocol almost really).
- Basilic class is resilient optimal in both BDB and BFT fault models
- By dynamically removing deceitful faults $\rightarrow$ active accountability
- Customizable depending on quorum size h_0
  - open systems (e.g. Blockchains) $\rightarrow$ greater threshold
  - closed systems (e.g. distributed database) $\rightarrow$ lower threshold

Q/A

alejandro.ranchalpedrosa@sydney.edu.au