

QuePaxa: Escaping the tyranny of timeouts in consensus

Pasindu Tennage
EPFL

Cristina Basescu
EPFL

Lefteris Kokoris-Kogias
ISTA and Mysten Labs

Ewa Syta
Trinity College


Philipp Jovanovic
UCL

Vero Galiñanes
EPFL

Bryan Ford
EPFL

Published in SOSP 2023, Koblenz, Germany


Tyranny of timeout challenges



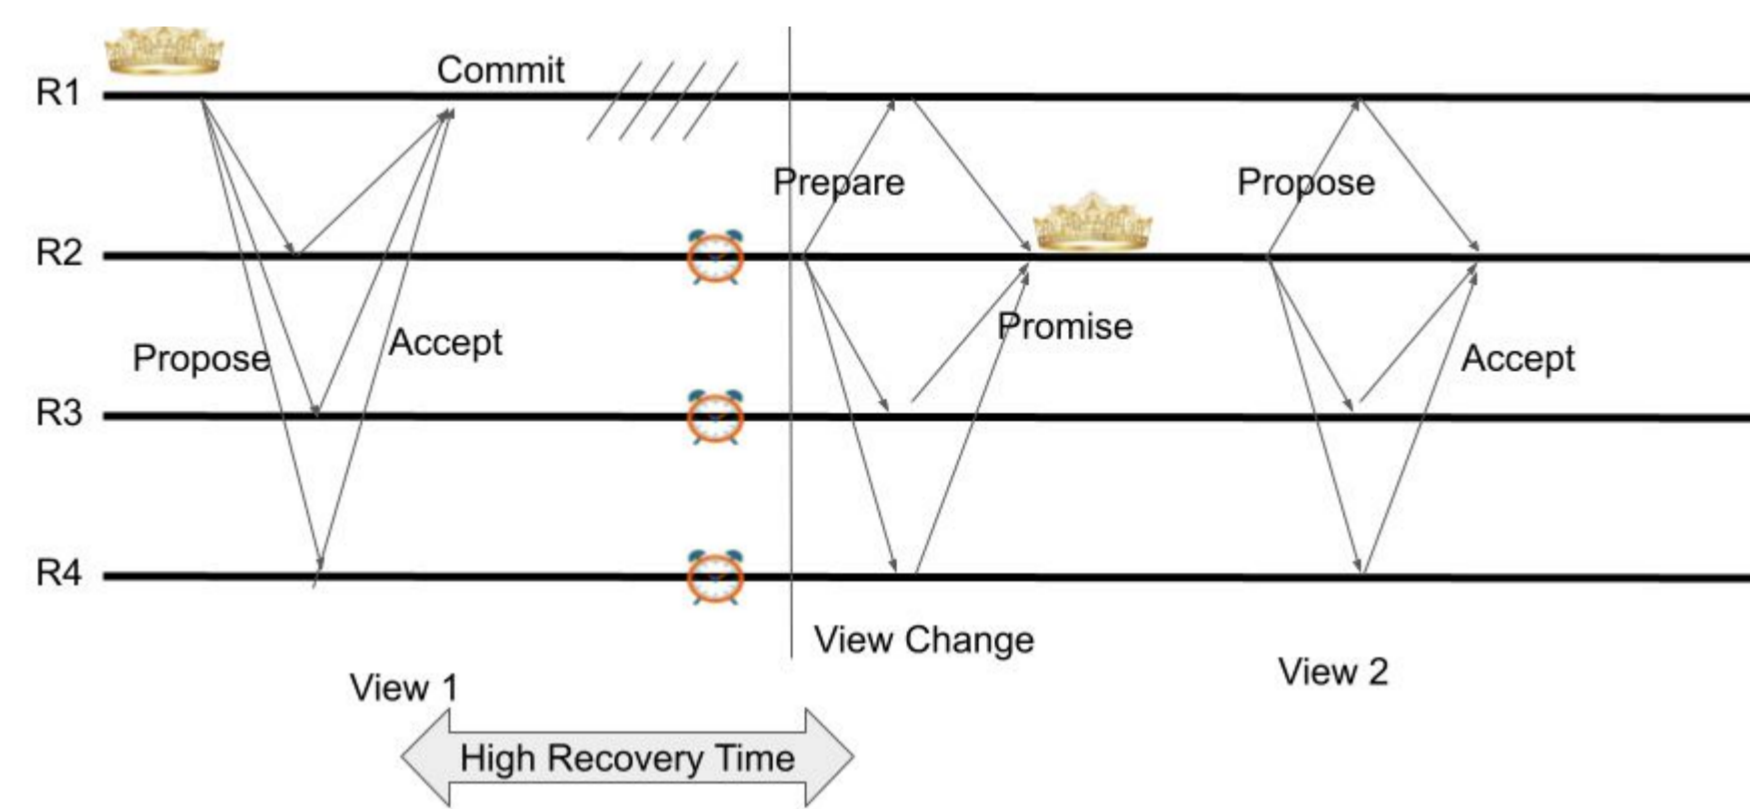
Conservative timeouts



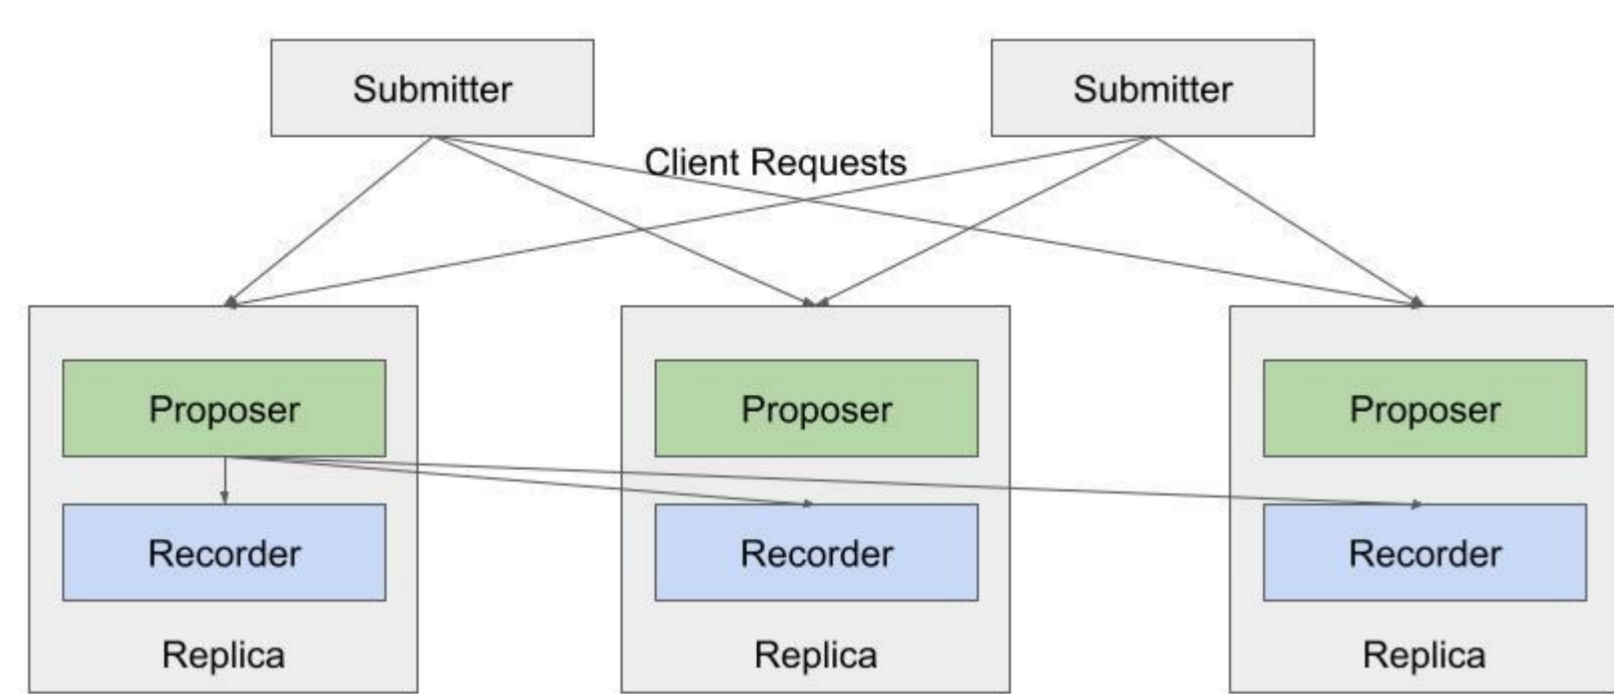
Low timeouts lead to liveness loss



High timeouts increase recovery time



QuePaxa architecture




Abstract QuePaxa algorithm


```

Algorithm 1: Abstract QuePaxa consensus algorithm
Input:  $v \leftarrow$  value preferred by this replica
repeat // iterate through rounds
   $p \leftarrow \langle v, \text{random}() \rangle$  // prioritized proposal
   $(P, \_) \leftarrow \text{tcast}(\{p\})$  // propagate our proposal
   $(E, P') \leftarrow \text{tcast}(P)$  // propagate existent sets
   $(C, U) \leftarrow \text{tcast}(P')$  // propagate common sets
   $v \leftarrow \text{best}(C).\text{value}$  // next candidate value
  if  $\text{best}(E) = \text{best}(U)$  then // detect consensus
    deliver( $v$ ) // deliver decision
    
```


QuePaxa decision modes



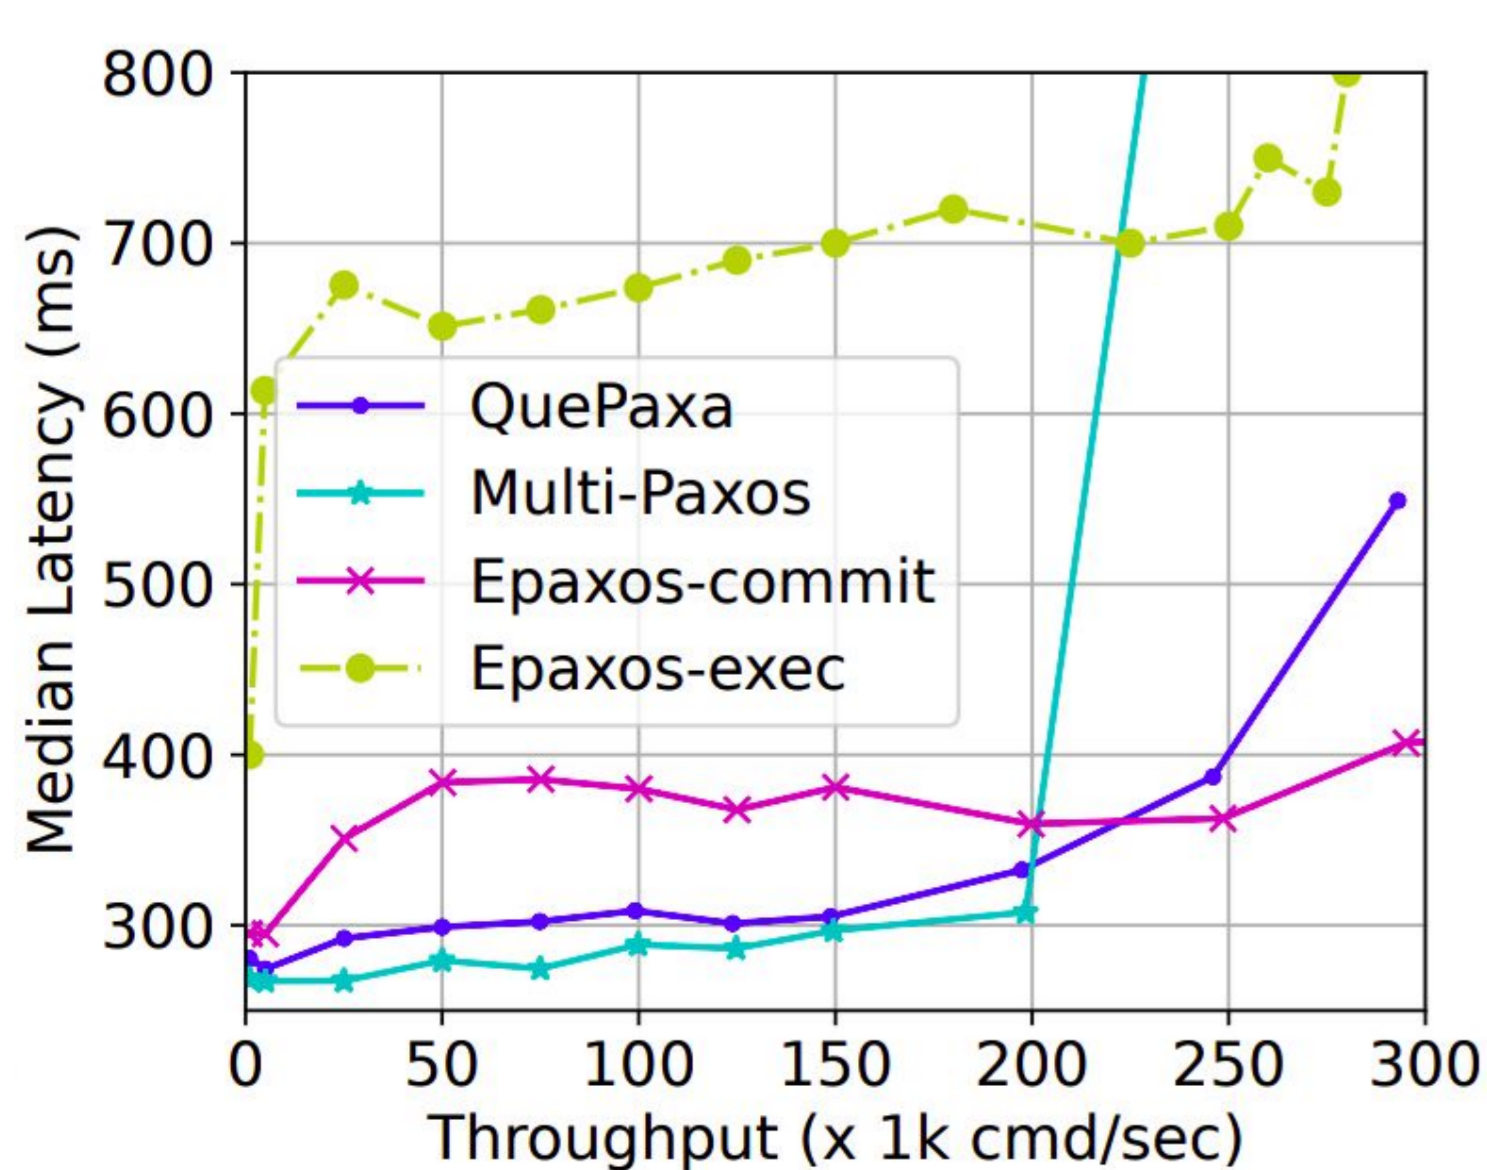
Hedging in QuePaxa



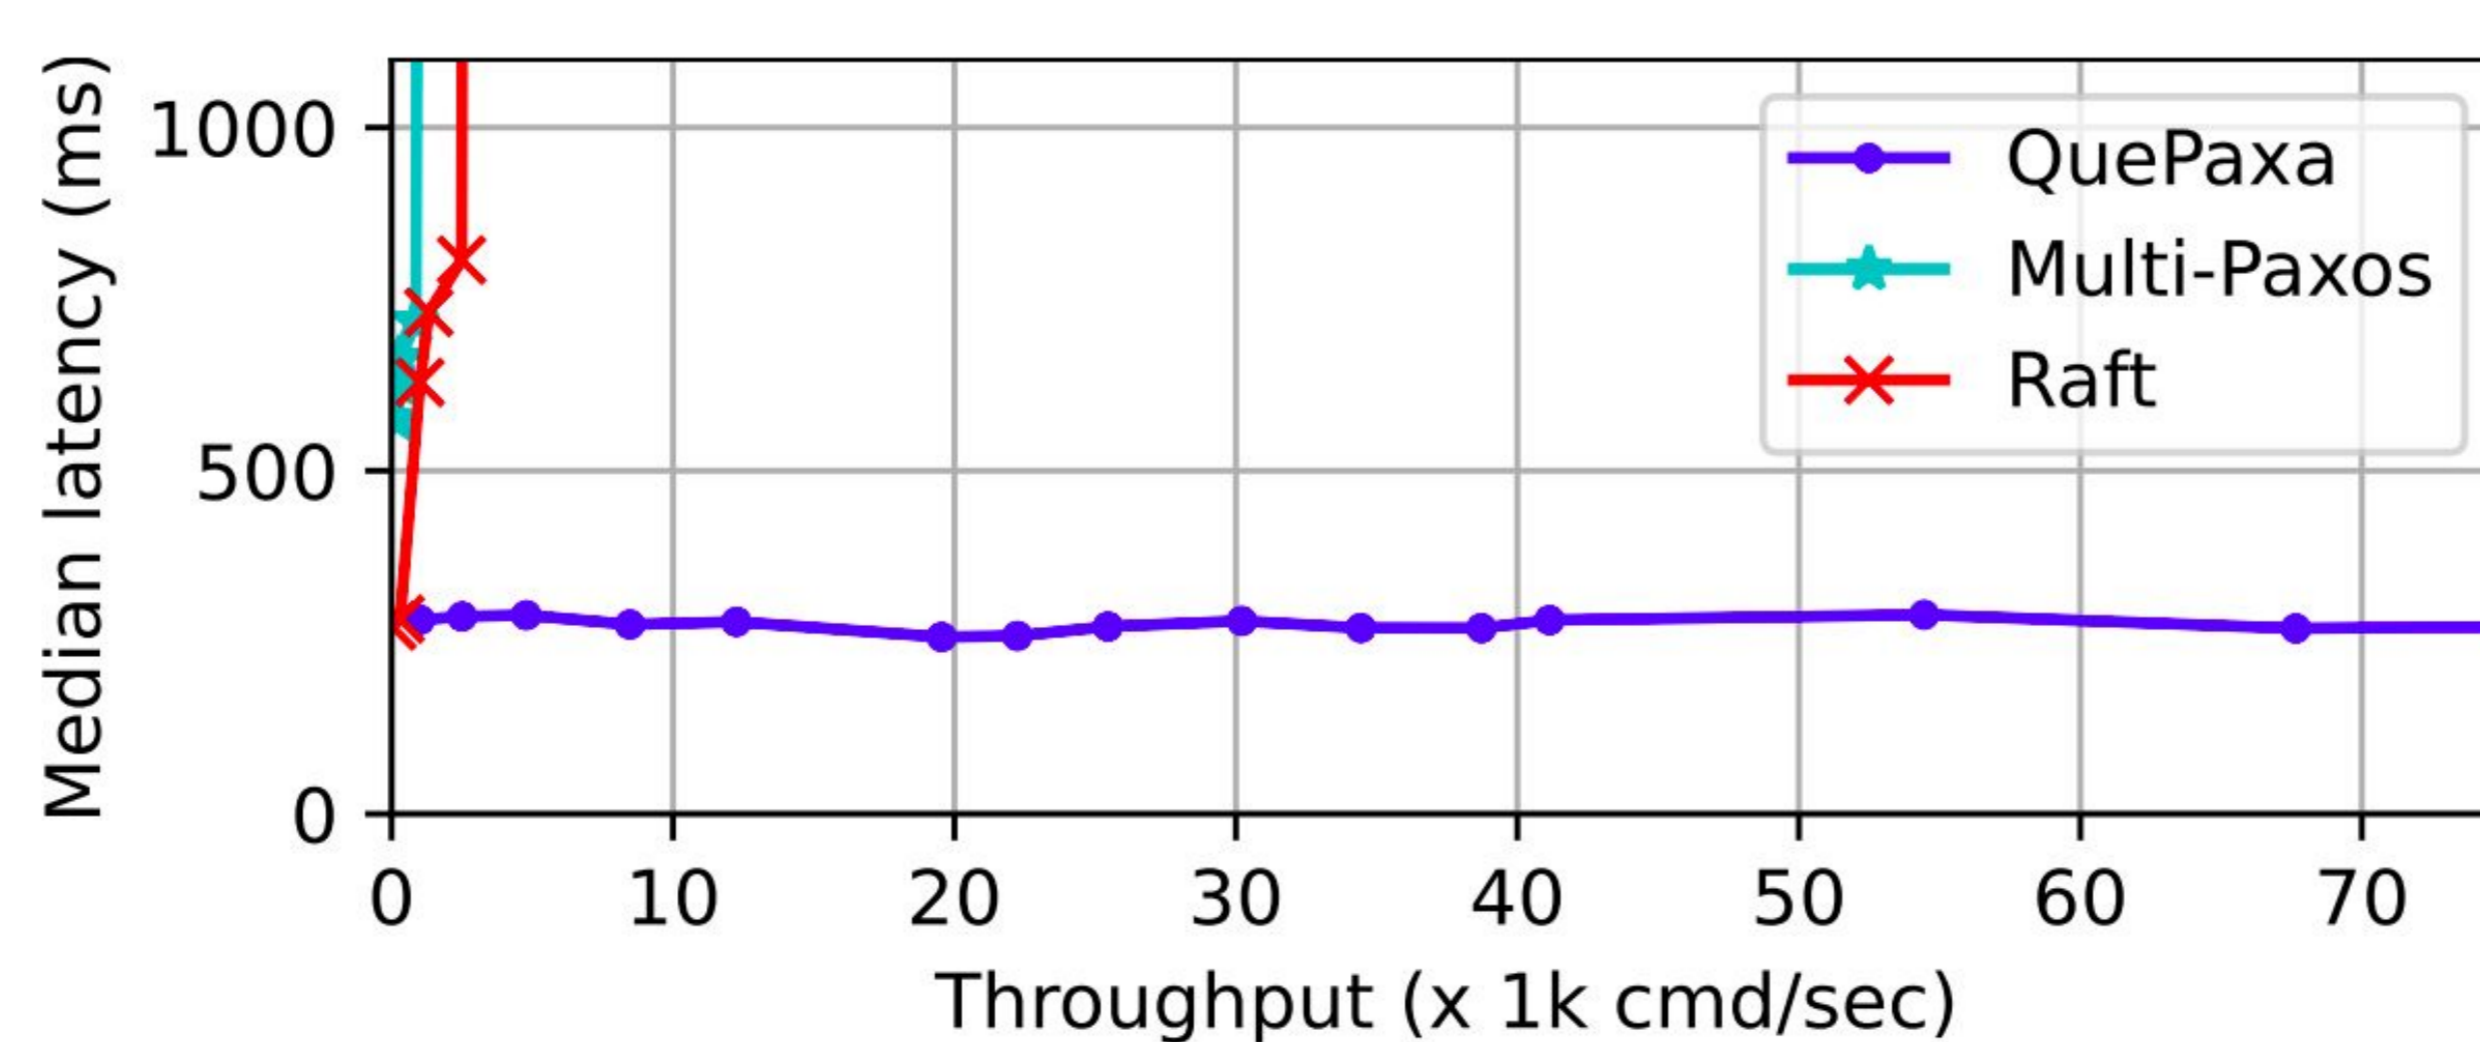
Leader tuning using Multi-armed bandits



Normal case wide-area performance



Performance under active network adversary



Impact of hedging delay for performance and recovery time

