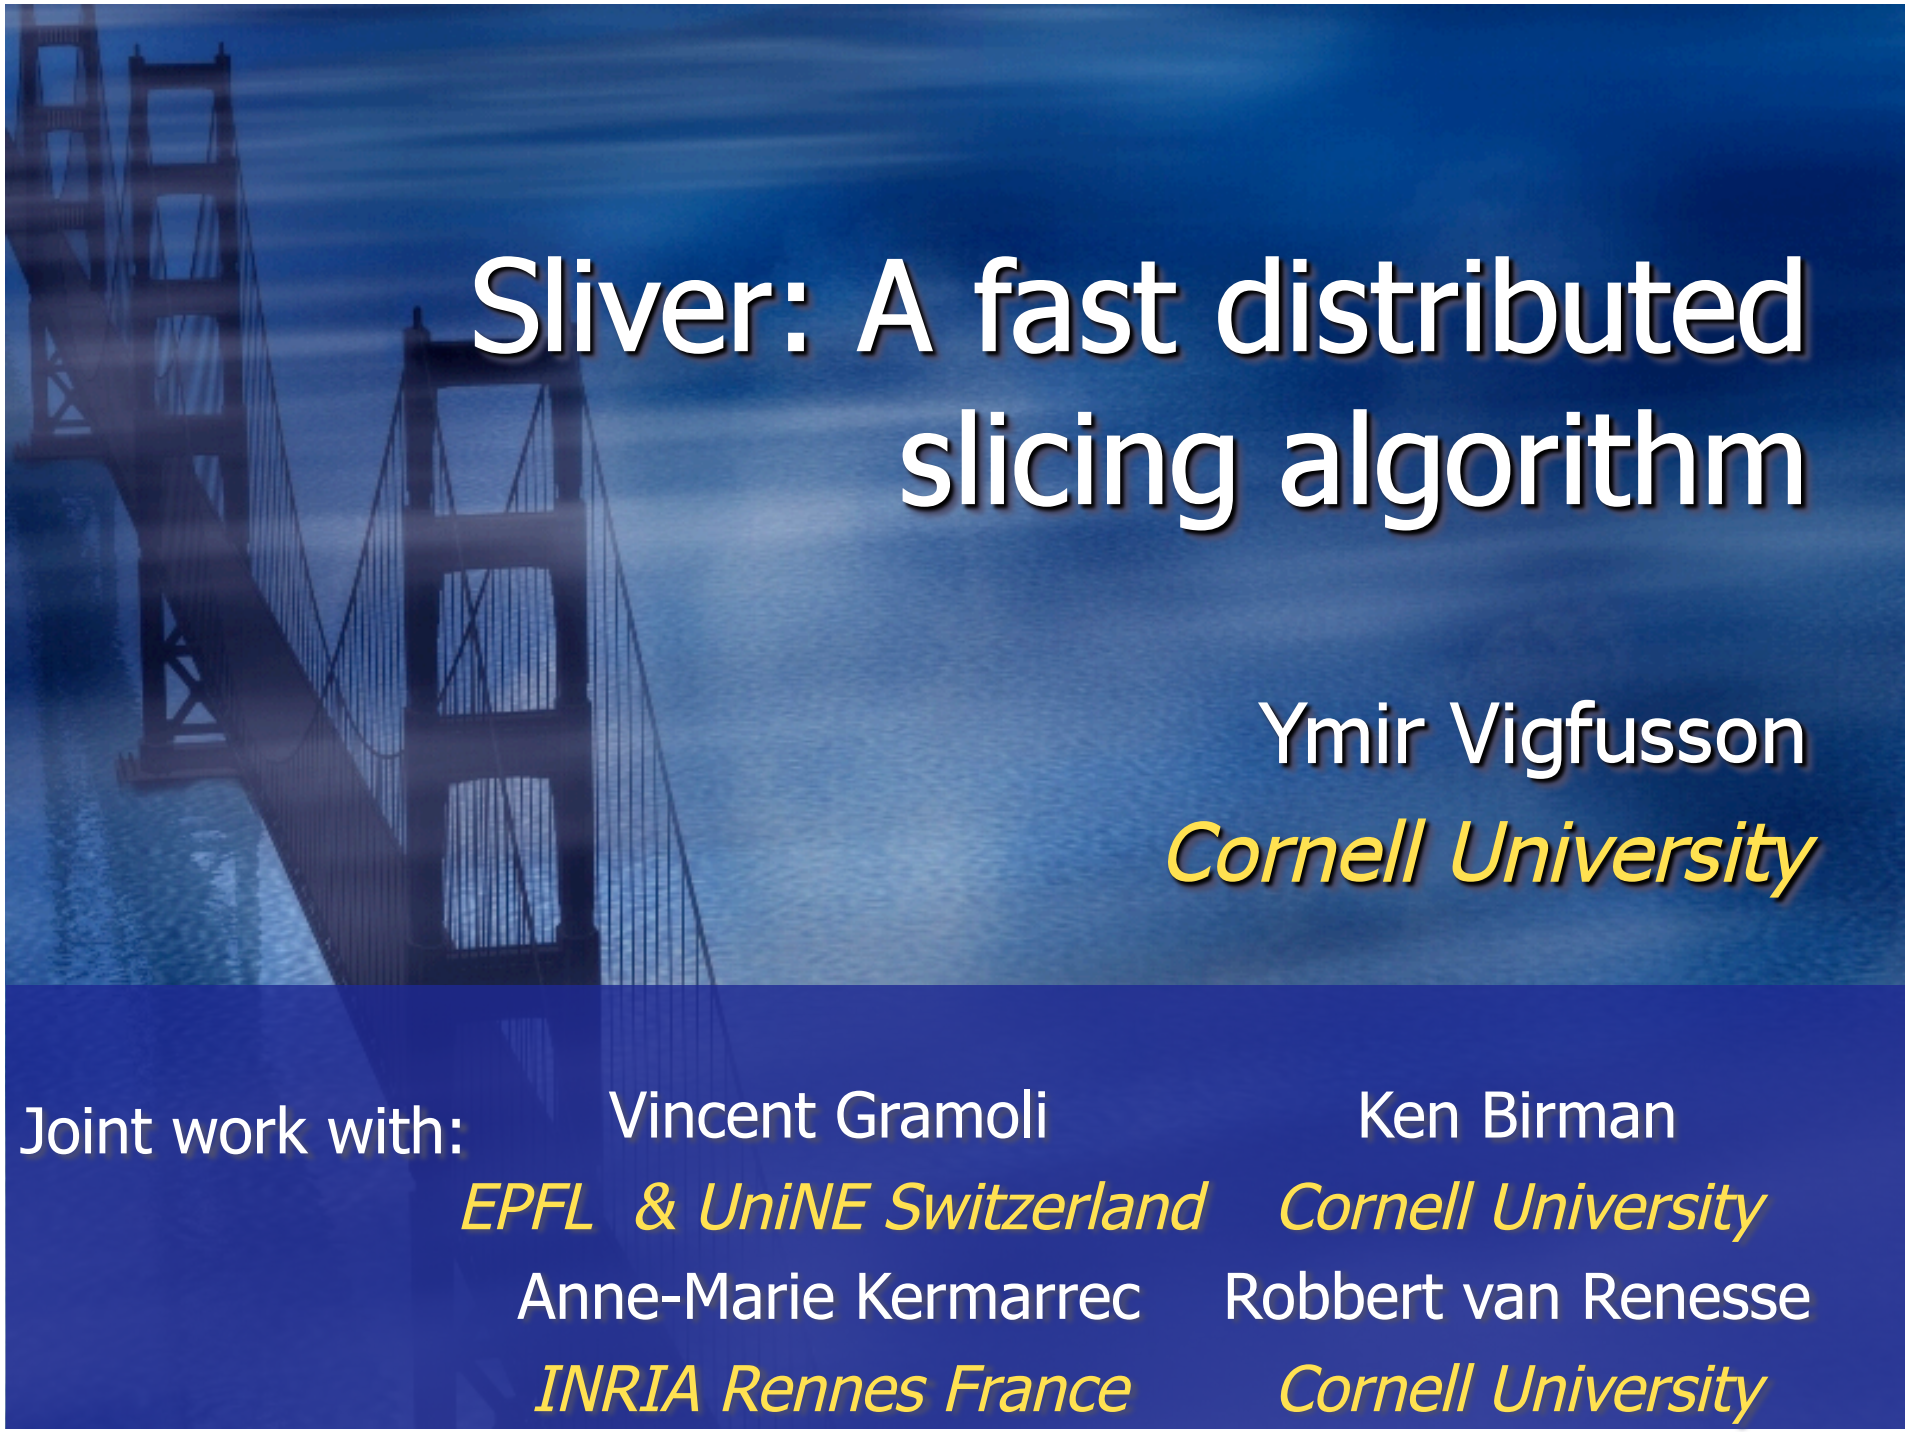

Sliver: A fast distributed slicing algorithm

Ymir Vigfusson
Cornell University

Joint work with:

Vincent Gramoli
EPFL & UniNE Switzerland

Ken Birman
Cornell University

Anne-Marie Kermarrec
INRIA Rennes France

Robbert van Renesse
Cornell University

The Distributed Slicing Problem

- n nodes, each has an attribute value x_i

The Distributed Slicing Problem

- n nodes, each has an attribute value x_i
- Divide the sorted list of x_i's into k slices

The Distributed Slicing Problem

- n nodes, each has an attribute value x_i
- Divide the sorted list of x_i's into k slices
- Each node i wants to independently know to which of the k slices x_i belongs

The Distributed Slicing Problem

- n nodes, each has an attribute value x_i
- Divide the sorted list of x_i's into k slices
- Each node i wants to independently know to which of the k slices x_i belongs

The Distributed Slicing Problem

- Large network with high rate of churn

The Distributed Slicing Problem

- Large network with high rate of churn
- *Example:* Choosing super-peers in the  network

Sliver: Distributed slicing algorithm

- Each node i gossips x_i (and other known values) to c random nodes

Sliver: Distributed slicing algorithm

- Each node i gossips x_i (and other known values) to c random nodes
- Node j keeps track of data it receives
  - Value, sender, expiration time
  - Suppose $\mathcal{B}_j$ values out of m are below x_j

Sliver: Distributed slicing algorithm

- Each node i gossips x_i (and other known values) to c random nodes
- Node j keeps track of data it receives
  - Value, sender, expiration time
  - Suppose B_j values out of m are below x_j
- Node j estimates its slice as $\left\lceil \frac{kB_j}{m} \right\rceil$

Results

- All nodes know their slice within 1 w.h.p. after $O\left(\sqrt{\max\{k, \log n\} \log n}\right)$ rounds in expectation

Results

- All nodes know their slice within 1 w.h.p. after $O\left(\sqrt{\max\{k, \log n\} \log n}\right)$ rounds in expectation
- Experiments on Emulab and simulations on Skype traces indicate rapid convergence and churn-tolerance
- *Conclusion:* Sliver is simple and robust with fast convergence properties.