Microsoft Research

Brief Announcement :

Partition Approach to Failure Detectors for *k*-Set Agreement

Wei Chen, Microsoft Research Asia

Jialin Zhang, Tsinghua University

Yu Chen, Microsoft Research Asia

Xuezheng Liu, Microsoft Research Asia

Background

- System model:
  - asynchronous message-passing,
  - processes may fail by crashing
- *k*-Set Agreement:
  - *n* distributed processes, each proposes one value
  - agree on at most *k* different proposed values
  - *k*=1 is the classic consensus problem
  - *k*-set agreement is impossible when *k* processes may fail
- Failure detectors
  - abstract synchrony and failure conditions

Failure Detectors for *k*-Set Agreement

- Several FD variants known, among them:
  - [Mostefaoui, et.al, PODC'06]: Ω_k, the weakest among the known FDs that solves *k*-set agreement (with a majority of correct processes)
  - Weakest FD for *k*-set agreement (for any *k*>1) is an open problem: "… remains one of the greatest research challenges in the fault-tolerant asynchronous computing theory community." [Raynal, Travers '06]
  - $\Omega_k \times \Sigma$: the weakest one known in the message-passing model

Our Contribution

- We introduces the partition approach
  - A general approach to weaken FDs for *k*-set agreement
- We apply the partition approach to $(\Omega_k \times \Sigma)$ in the message-passing model, and define:
  - Π_k: statically partitioned FD, strictly weaker than $\Omega_k \times \Sigma$
  - Π^S_k: splittable partitioned FD, strictly weaker than Π_k
  - Π_k and Π^S_k strong enough to solve *k*-set agreement

Defining Partitioned FDs via the Partition Approach

- For $\prod_k$, Informally,
  - FD decides a static partition of processes $\{P_1, \ldots, P_s\}$ in each run
  - For all P_i, FD on P_i satisfies all safety properties of $\Omega_{k\_i} \times \Sigma$ restricted on P_i, such that $k\_1 + k\_2 + \ldots + k\_s \leq k$.
    - to guarantee at most k decisions for k-set agreement.
  - There exists a P_j (live component), FD on P_j satisfies all liveness properties of $\Omega_{k\_j} \times \Sigma$ restricted on P_j
    - to guarantee that eventually processes make decisions.
- For $\prod^S_k$, informally allow dynamic splitting

Relationship Lattice

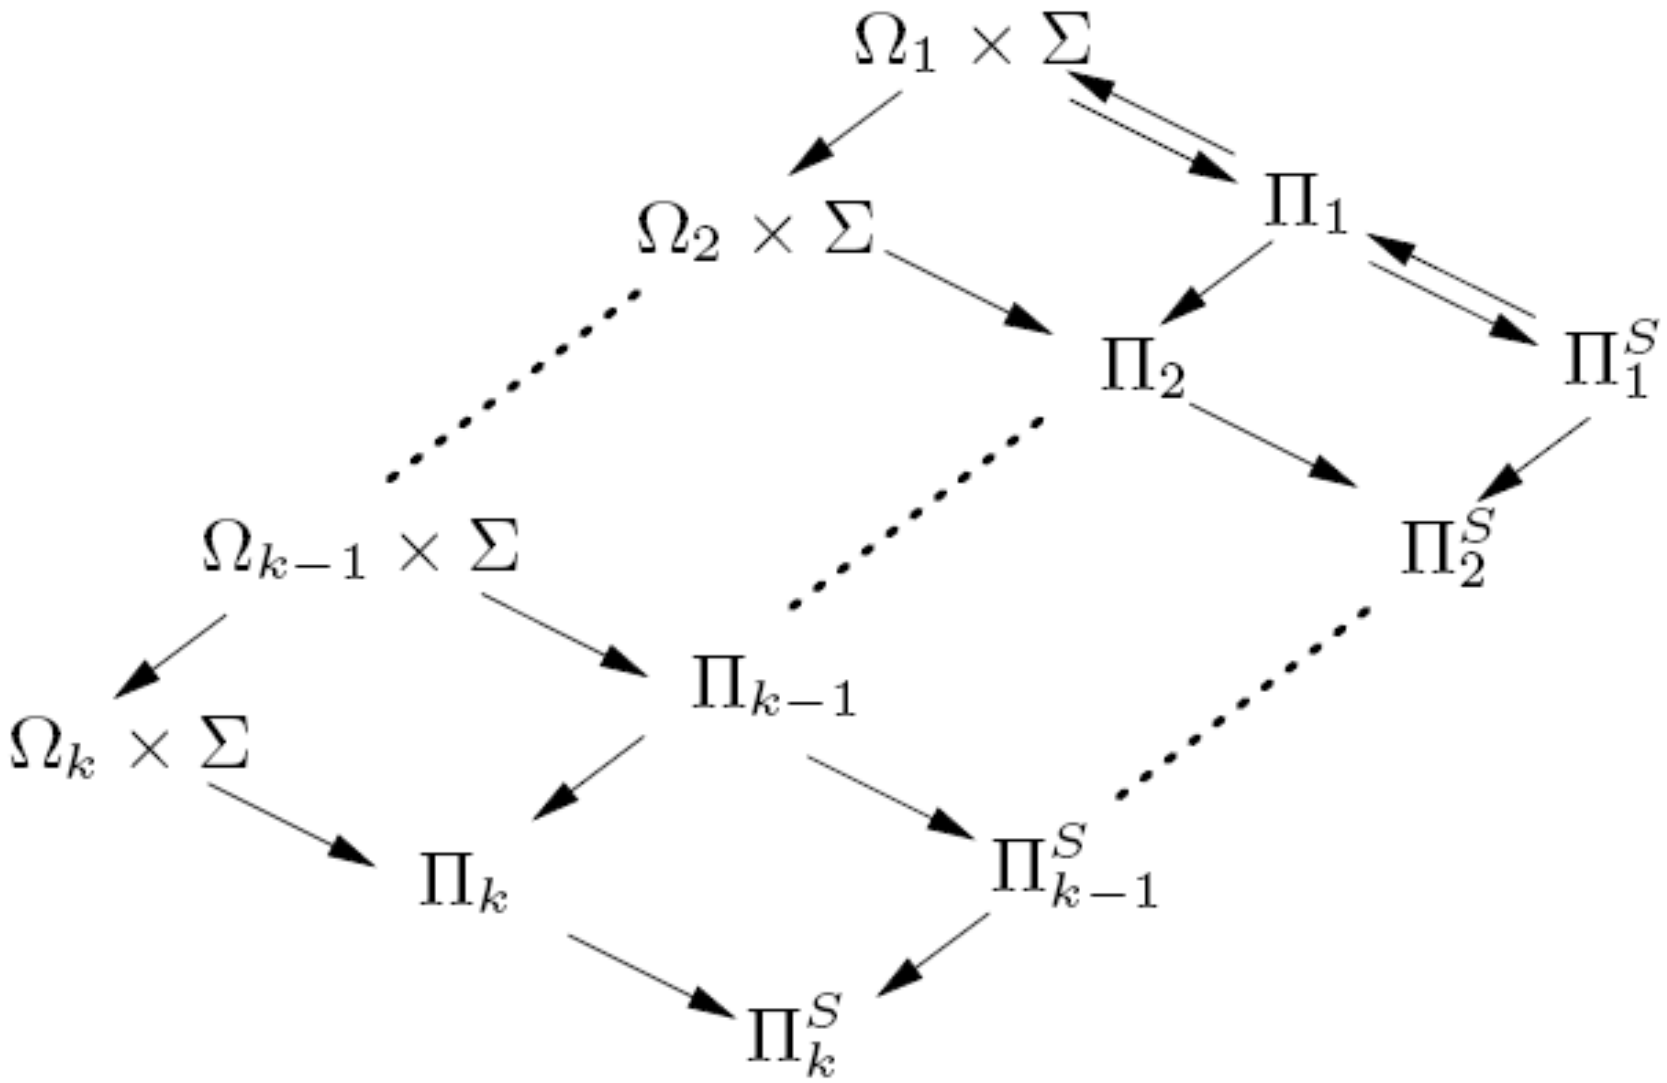

Relationship Lattice

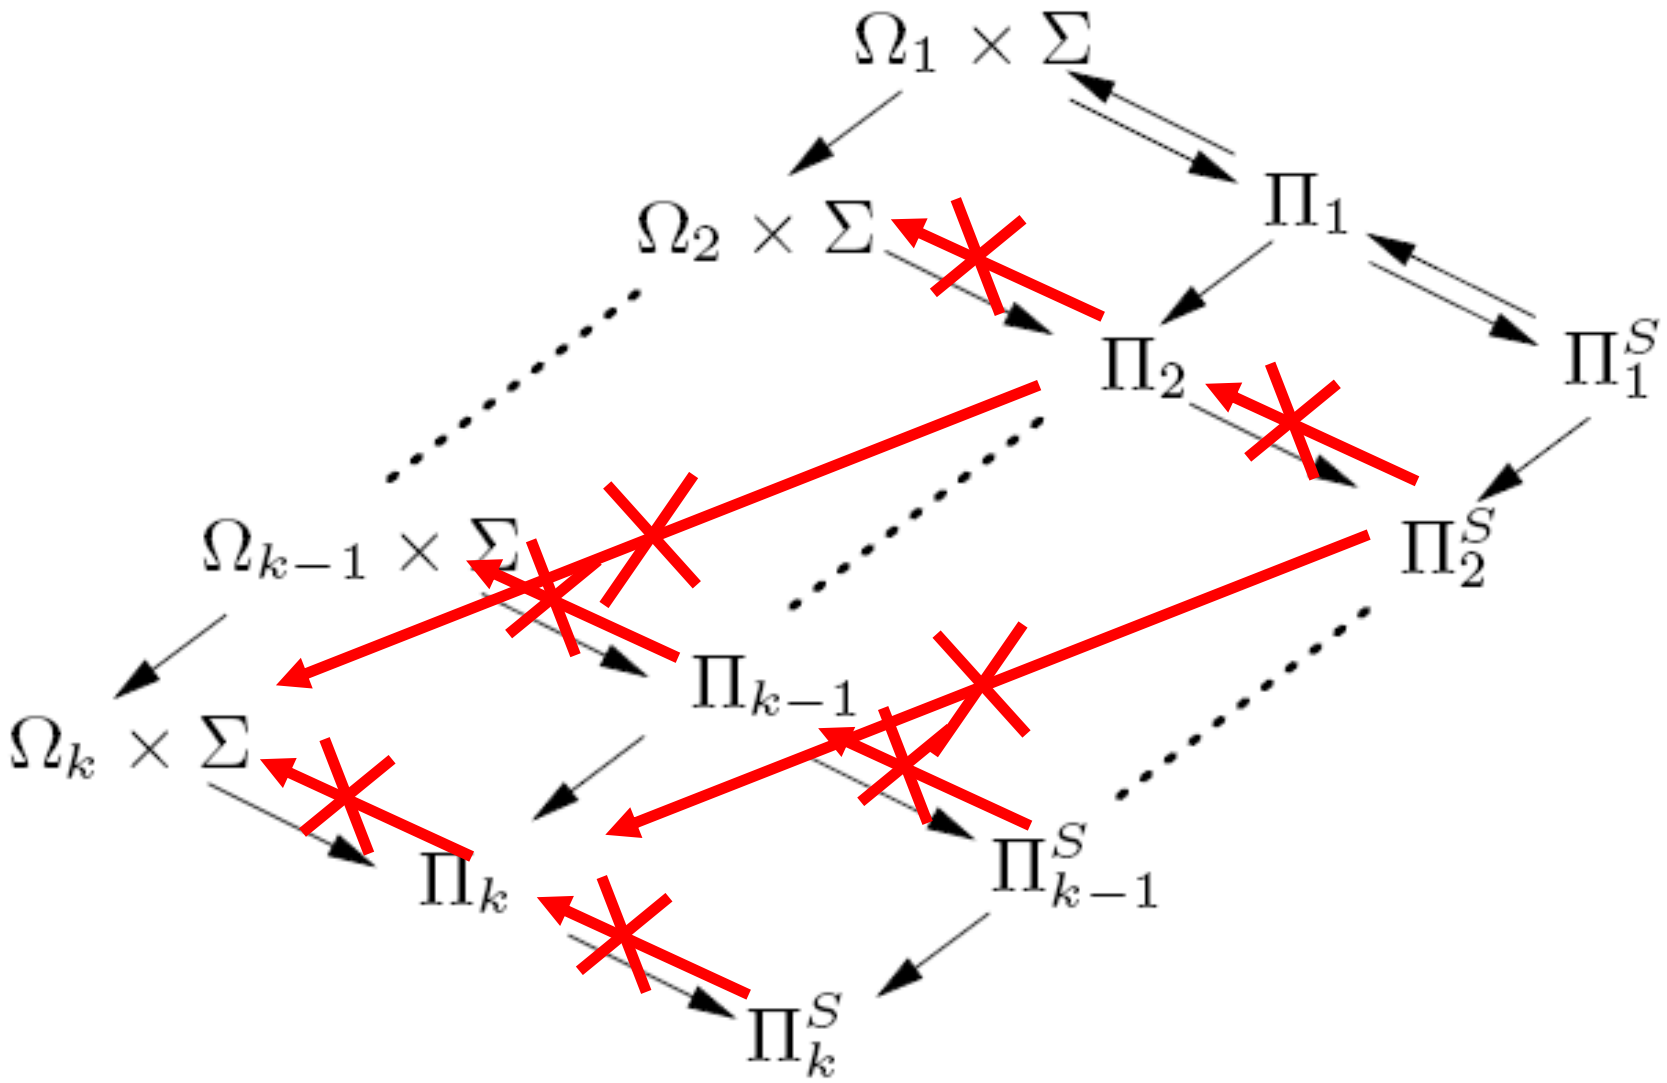

Summary

- Introduce a general partition approach
- Apply the approach to weaken existing FDs in the message-passing model
- We also have results in shared-memory model (DISC'06)
- Open a new dimension in studying weak FDs for *k*-set agreement
- Future work:
  - Formally define partitioned FDs and study the weakest partitioned FDs
  - Implementation of partitioned FDs to match network partitions