CLOUD COMPUTING CONCEPTS

with Indranil Gupta (Indy)

MEMBERSHIP

Lecture B

FAILURE DETECTORS

LARGE GROUP: SCALABILITY A GOAL

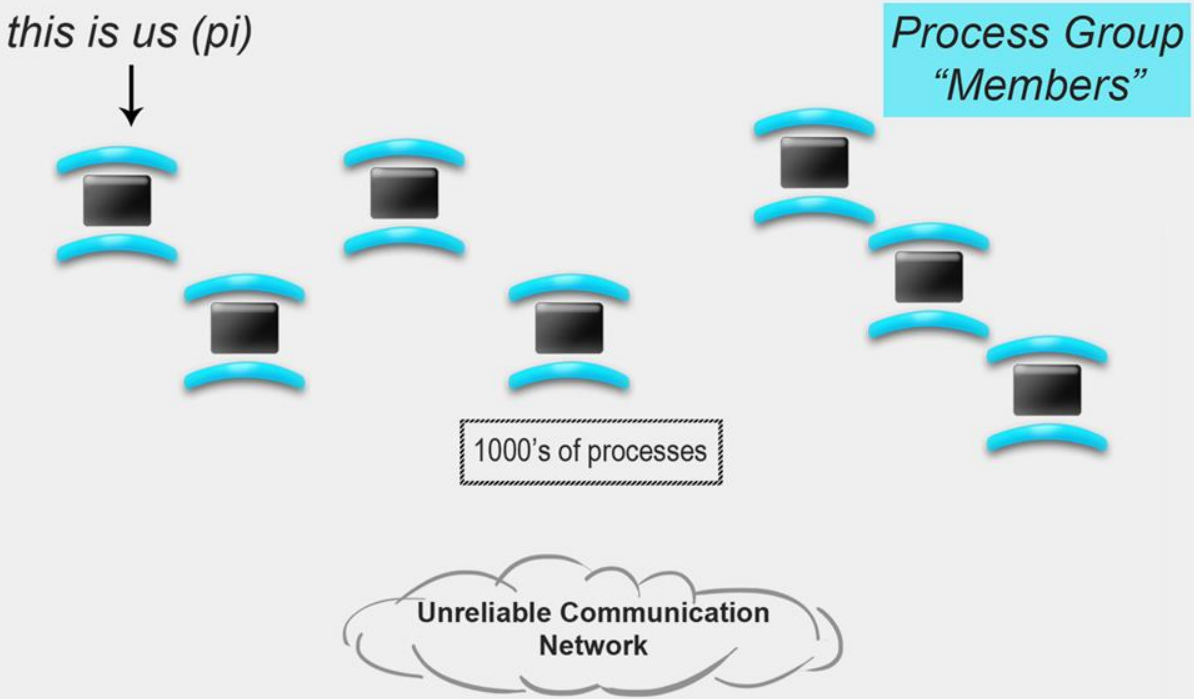

GROUP MEMBERSHIP PROTOCOL

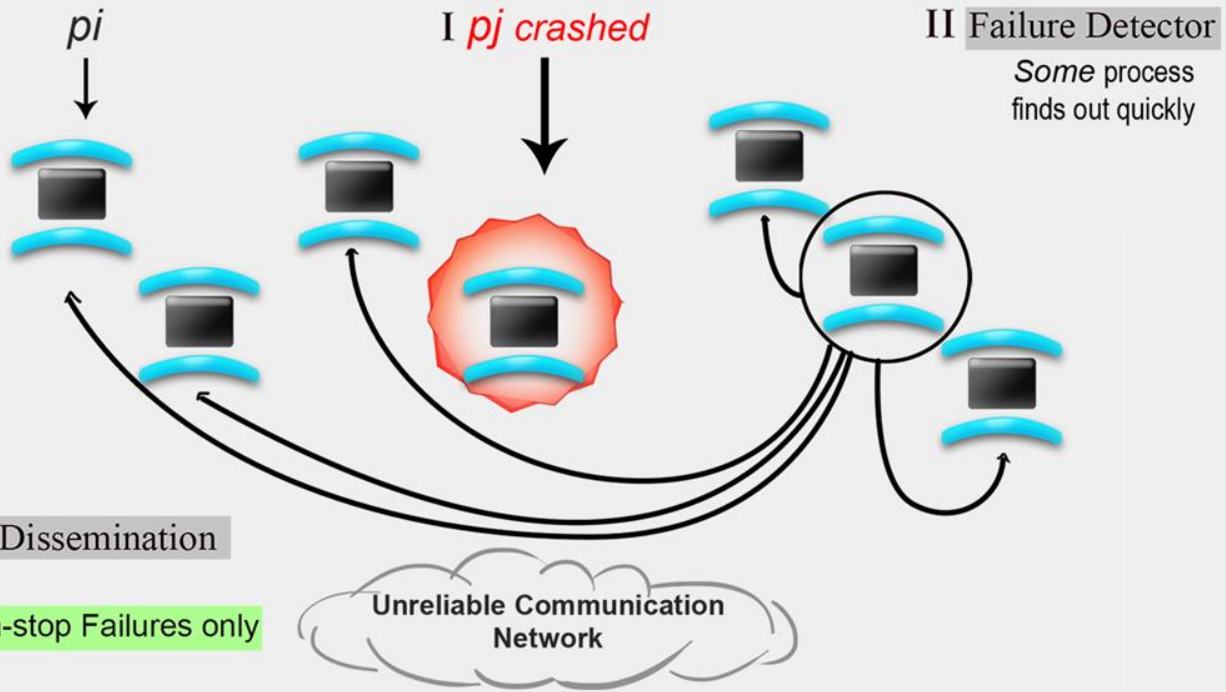

I. *pj* CRASHES

- Nothing we can do about it!
- A frequent occurrence
- Common case rather than exception
- Frequency goes up linearly with size of datacenter

II. DISTRIBUTED FAILURE DETECTORS: DESIRABLE PROPERTIES

- Completeness = each failure is detected
- Accuracy = there is no mistaken detection
- Speed
  - Time to first detection of a failure
- Scale
  - Equal Load on each member
  - Network Message Load

DISTRIBUTED FAILURE DETECTORS: PROPERTIES

- Completeness
- Accuracy

Impossible together in lossy networks [Chandra and Toueg]

If possible, then can solve consensus!

- Speed
  - Time to first detection of failure
- Scale
  - Equal Load on each member
  - Network Message Load

WHAT REAL FAILURE DETECTORS PREFER

- Completeness — Guaranteed
- Accuracy — Partial/Probabilistic guarantee

- Speed
  - Time to first detection of failure
- Scale
  - Equal Load on each member
  - Network Message Load

FAILURE DETECTOR PROPERTIES

- Completeness — Guaranteed
- Accuracy — Partial/Probabilistic guarantee
- Speed — Time until some process detects the failure
  - Time to first detection of failure
- Scale
  - Equal Load on each member
  - Network Message Load

FAILURE DETECTOR PROPERTIES

- Completeness — Guaranteed
- Accuracy — Partial/Probabilistic guarantee
- Speed
  - Time to first detection of failure — Time until some process detects the failure
- Scale
  - Equal Load on each member
  - Network Message Load — No bottlenecks/single failurepoint

FAILURE DETECTOR PROPERTIES

- Completeness
- Accuracy

- Speed
  - Time to first detection of failure
- Scale
  - Equal Load on each member
  - Network Message Load

In spite of arbitrary simultaneous process failures

CENTRALIZED HEARTBEATING

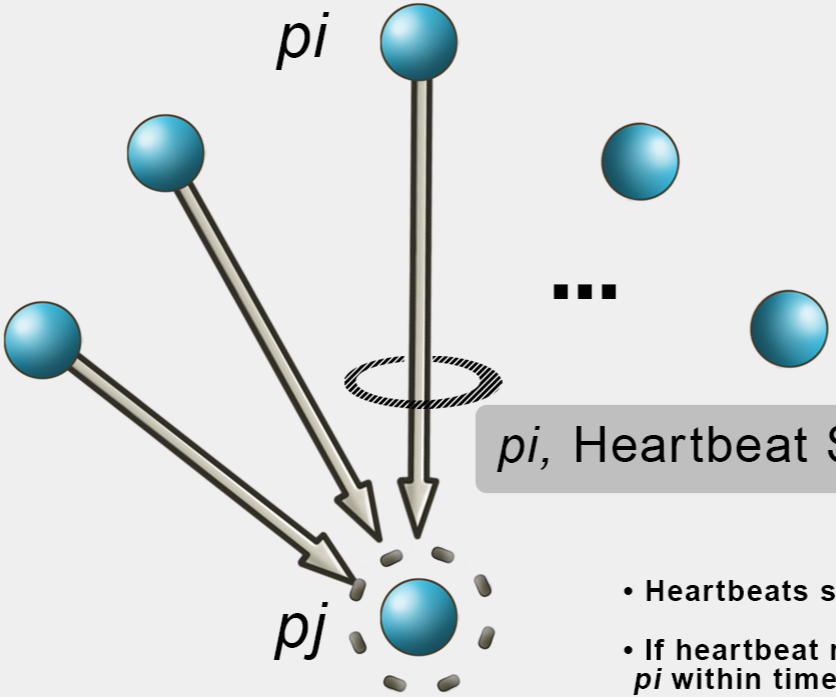

- Heartbeats sent periodically
- If heartbeat not received from *pi* within timeout, mark *pi* as failed

CENTRALIZED HEARTBEATING

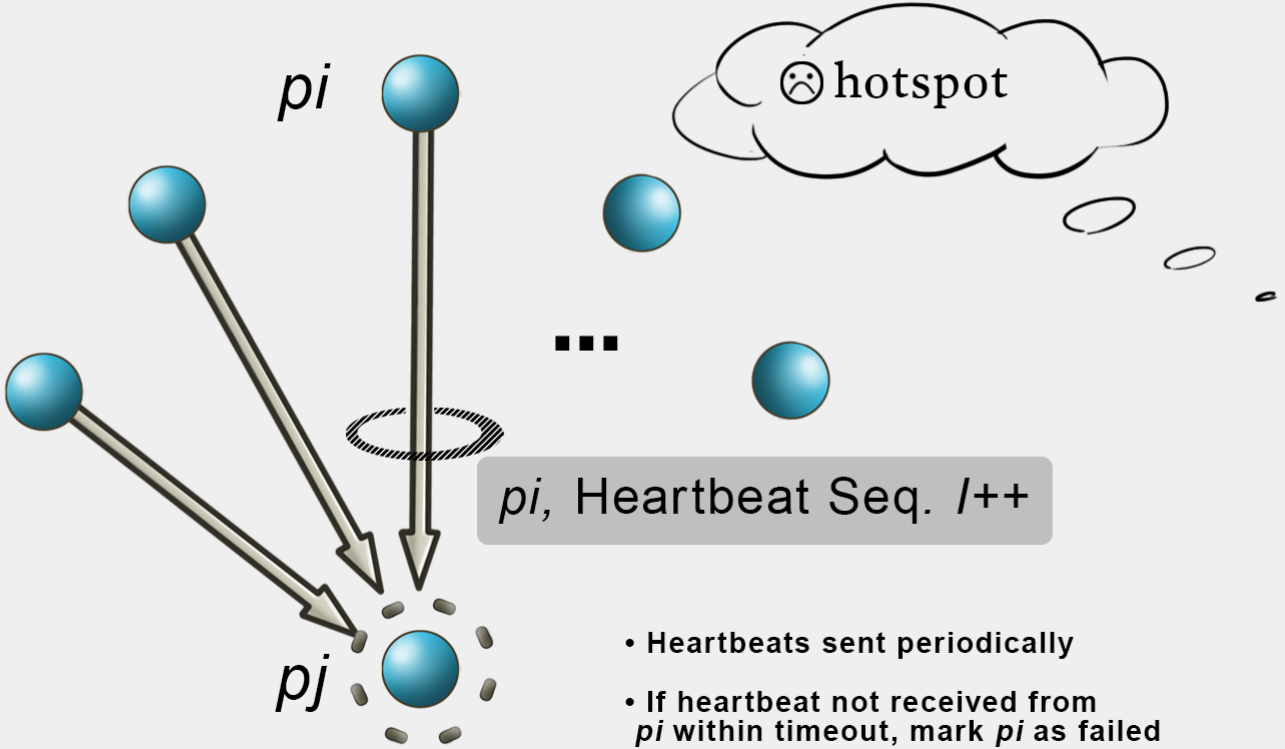

- Heartbeats sent periodically
- If heartbeat not received from *pi* within timeout, mark *pi* as failed

RING HEARTBEATING

pi, Heartbeat Seq. *l++*

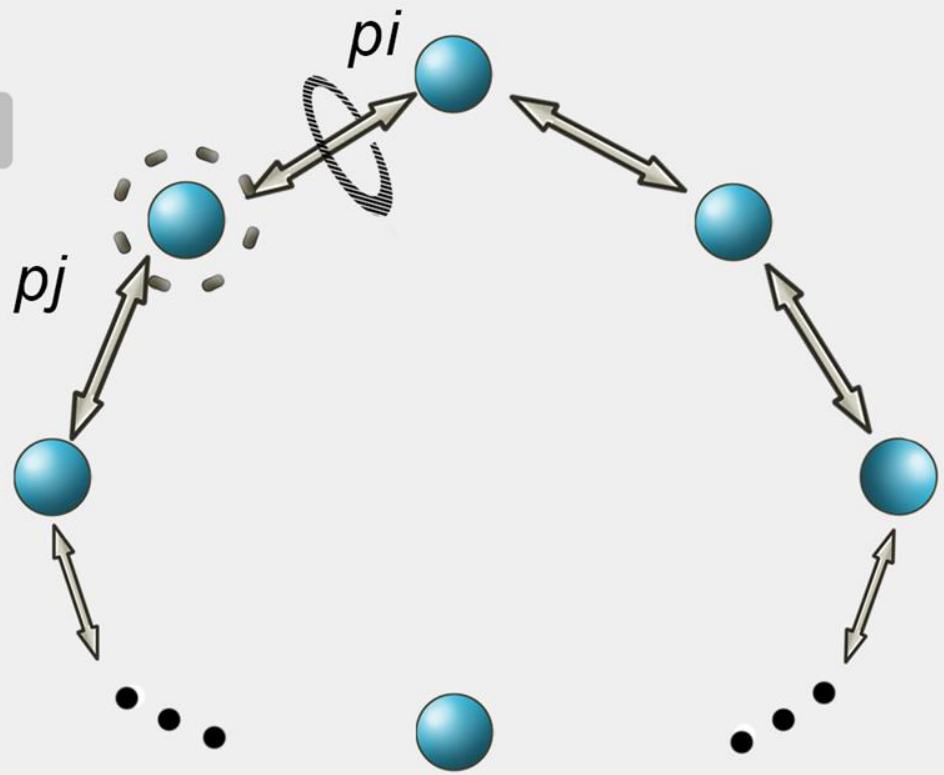

RING HEARTBEATING

pi, Heartbeat Seq. *l*++

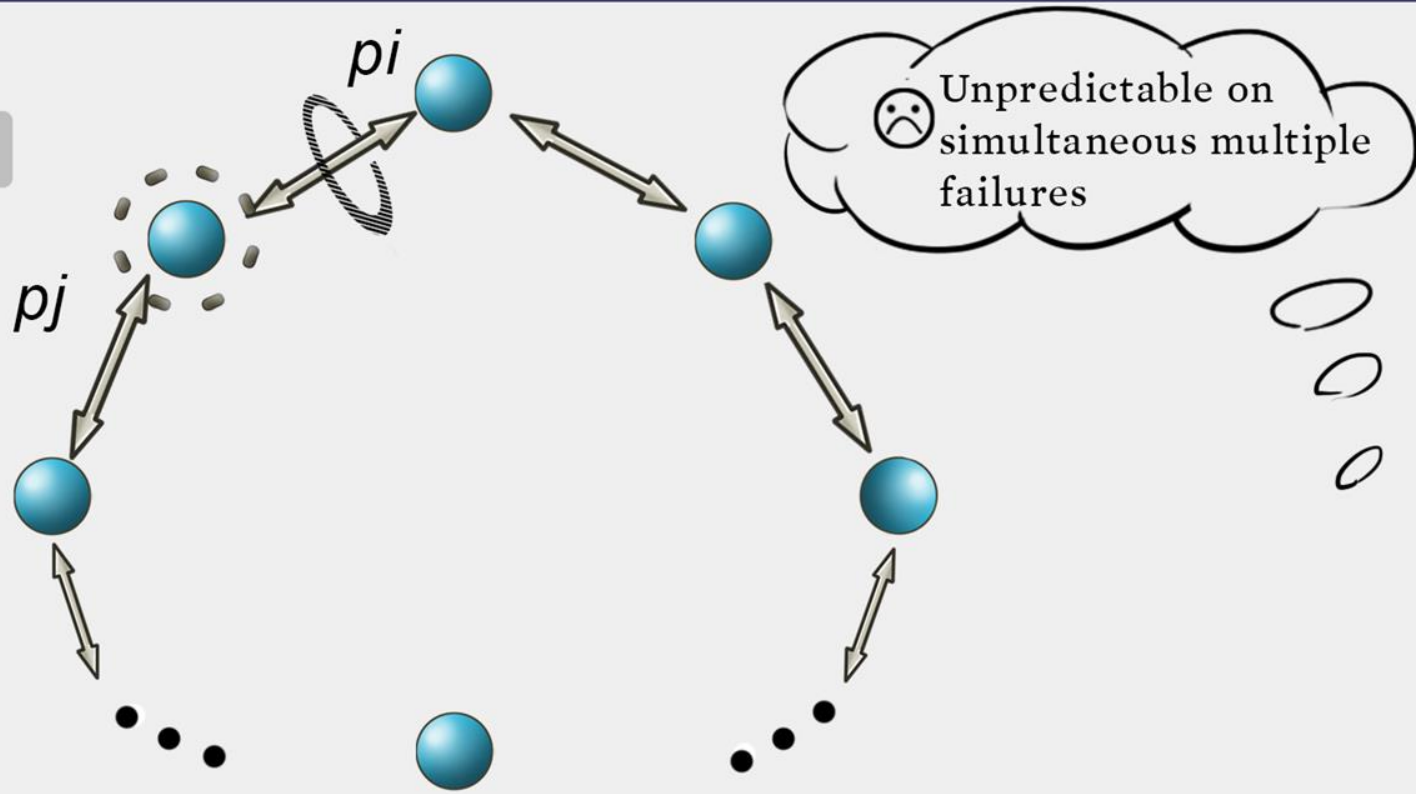

ALL-TO-ALL HEARTBEATING

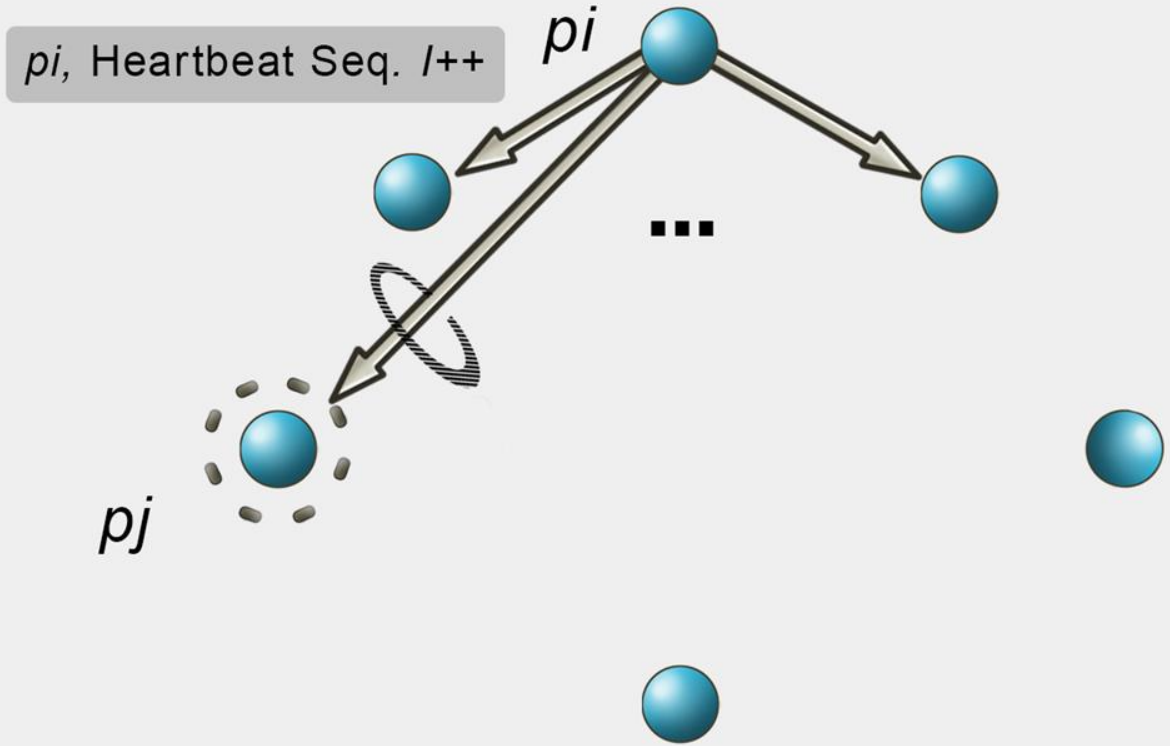

ALL-TO-ALL HEARTBEATING

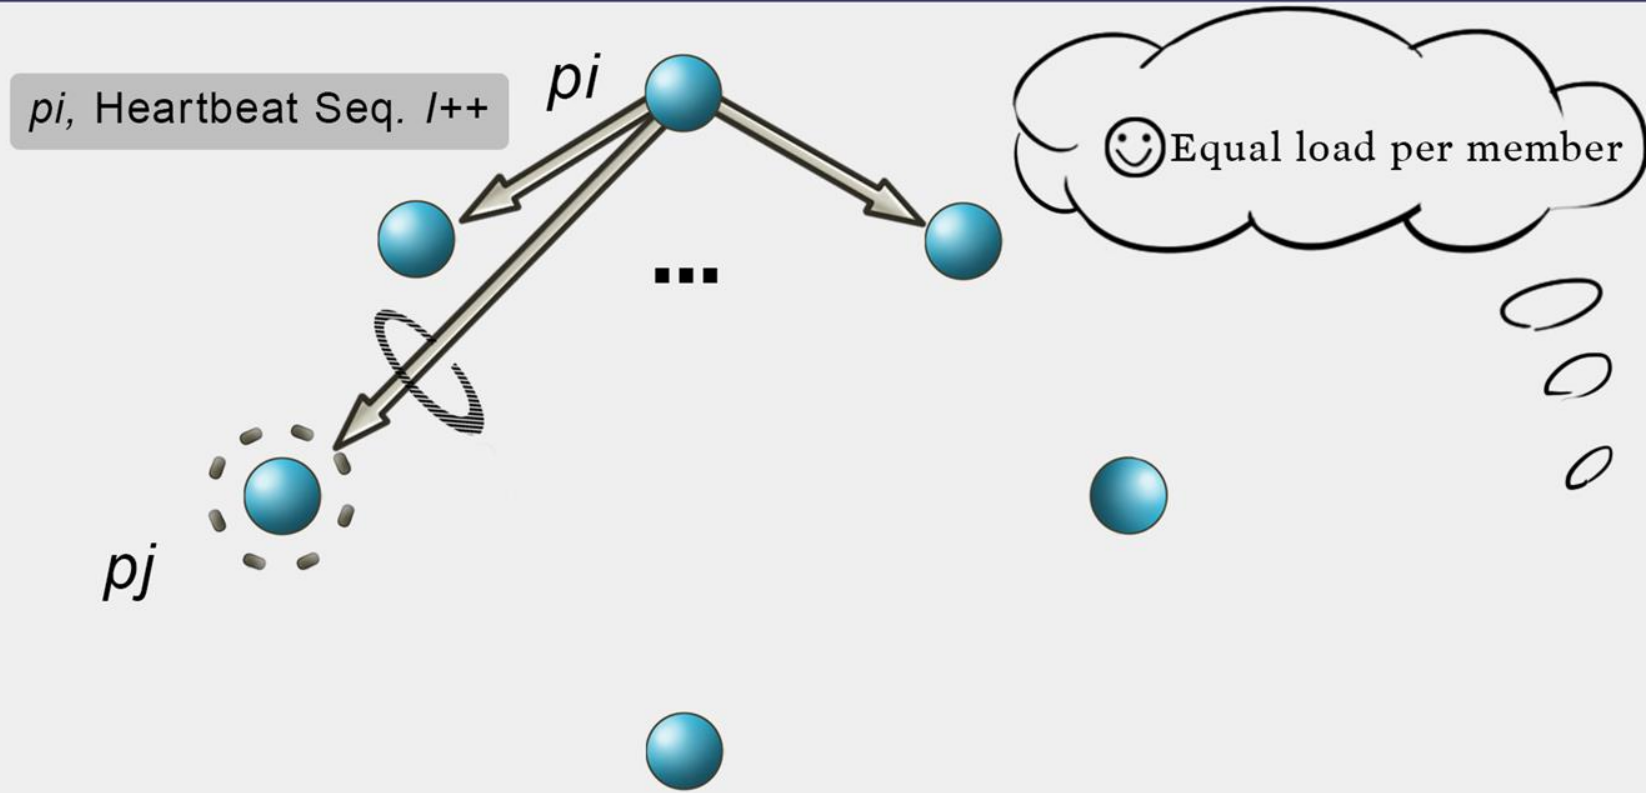

NEXT

- How do you increase the robustness of all-to-all heartbeating?