

CS 425 / ECE 428
Distributed Systems
Fall 2016

Indranil Gupta (Indy)

Sep 8, 2016

*Lecture 6: Failure Detection and
Membership, Grids*

All slides © IG

A CHALLENGE

- You've been put in charge of a datacenter, and your manager has told you, "Oh no! We don't have any failures in our datacenter!"
- Do you believe him/her?
- What would be your first responsibility?
- Build a failure detector
- What are some things that could go wrong if you didn't do this?

FAILURES ARE THE NORM

... not the exception, in datacenters.

Say, the rate of failure of one machine (OS/disk/motherboard/network, etc.) is once every 10 years (120 months) on average.

When you have 120 servers in the DC, the **mean time to failure (MTTF)** of the next machine is 1 month.

When you have 12,000 servers in the DC, the MTTF is about once every 7.2 hours!

Soft crashes and failures are even more frequent!

TO BUILD A FAILURE DETECTOR

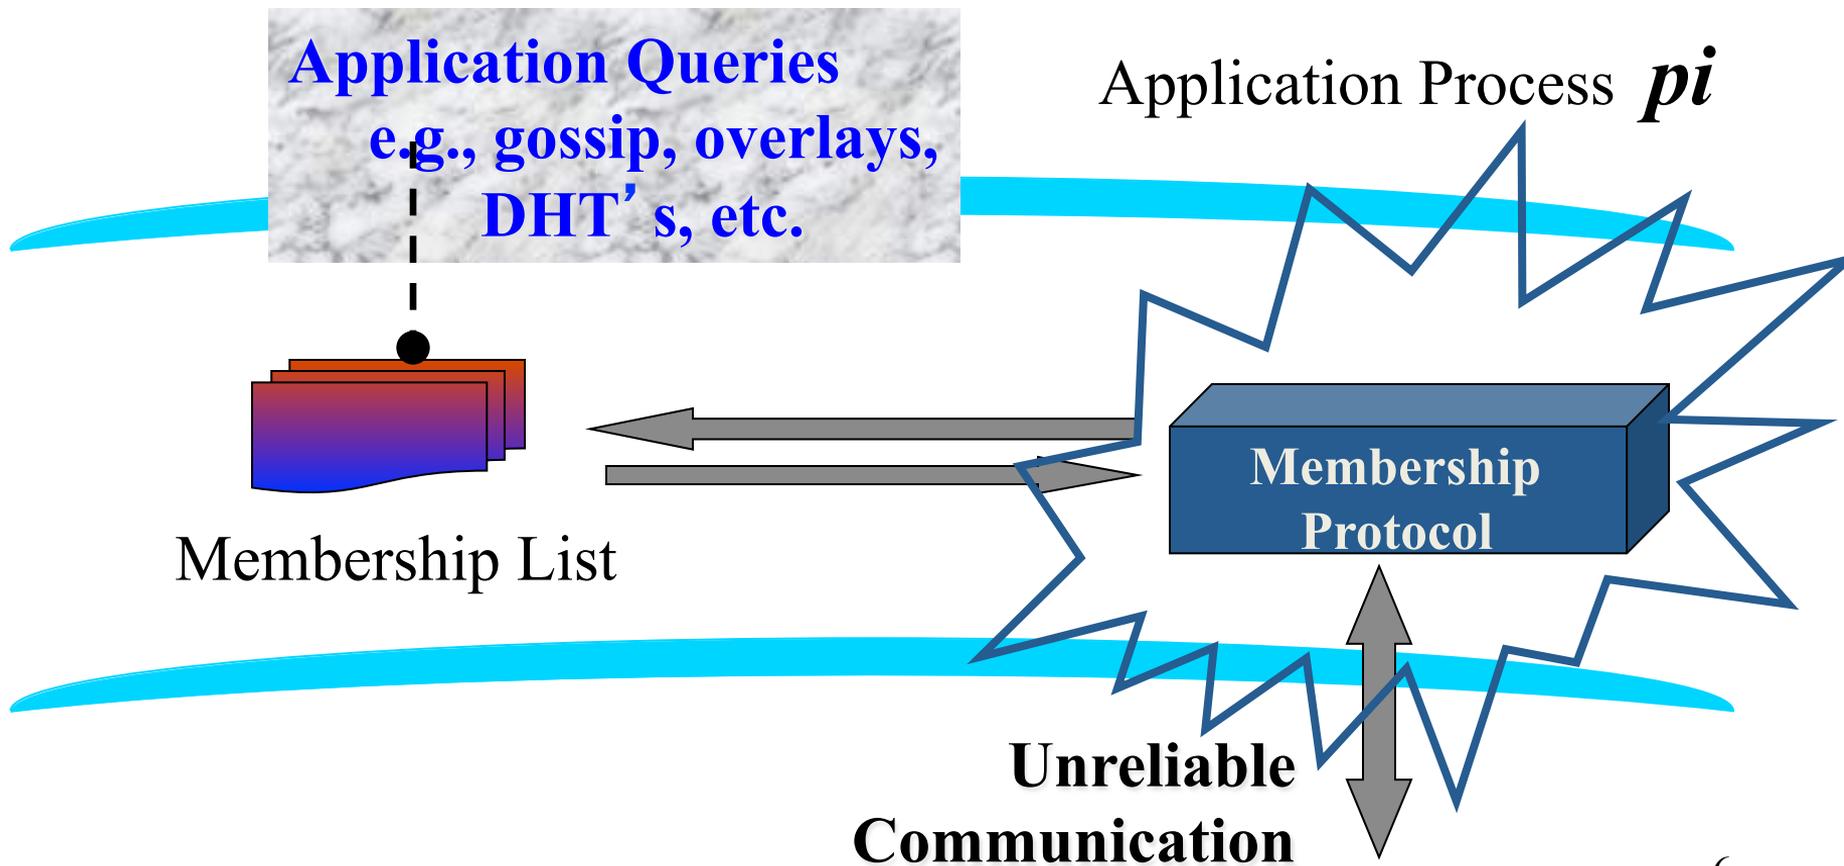
- You have a few options
 1. Hire 1000 people, each to monitor one machine in the datacenter and report to you when it fails.
 2. Write a failure detector program (distributed) that automatically detects failures and reports to your workstation.

TARGET SETTINGS

- Process ‘group’ -based systems
 - Clouds/Datacenters
 - Replicated servers
 - Distributed databases

- Fail-stop (crash) process failures

GROUP MEMBERSHIP SERVICE

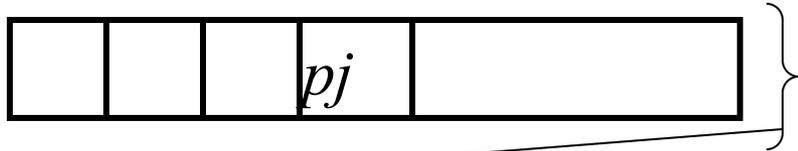


TWO SUB-PROTOCOLS

Application Process pi

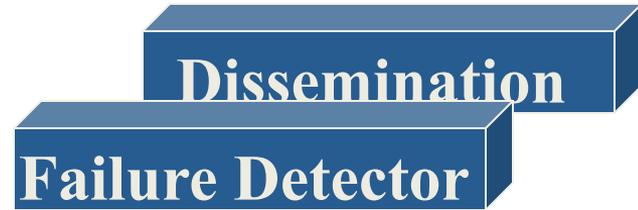
Group

Membership List



- **Complete list all the time (Strongly consistent)**
 - Virtual synchrony
- **Almost-Complete list (Weakly consistent)**
 - Gossip-style, SWIM, ...
- **Or Partial-random list (other systems)**
 - SCAMP, T-MAN, Cyclon, ...

Focus of this series of lecture

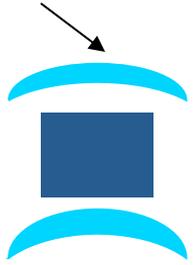


Unreliable
Communication

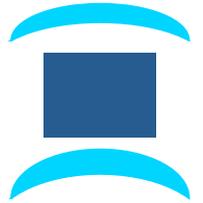
LARGE GROUP: SCALABILITY A

GOAL

this is us (pi)



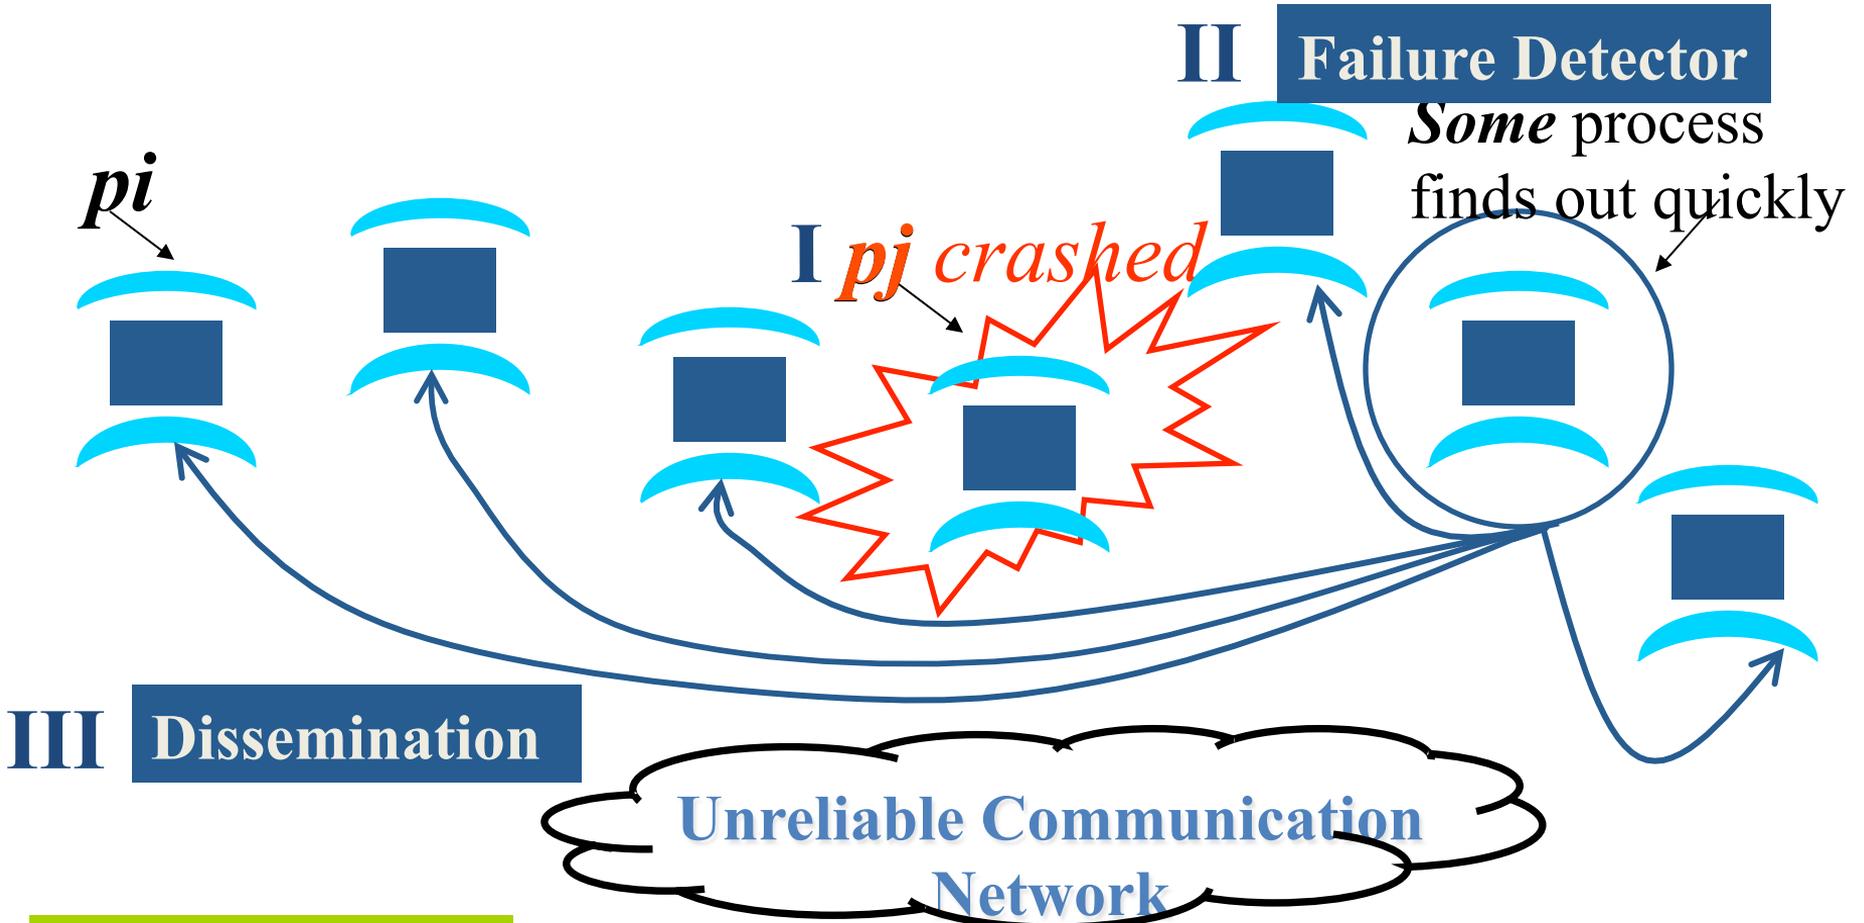
Process Group
"Members"



| 1000' s of processes |

Unreliable Communication
Network

GROUP MEMBERSHIP PROTOCOL



Fail-stop Failures only

NEXT

- How do you design a 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

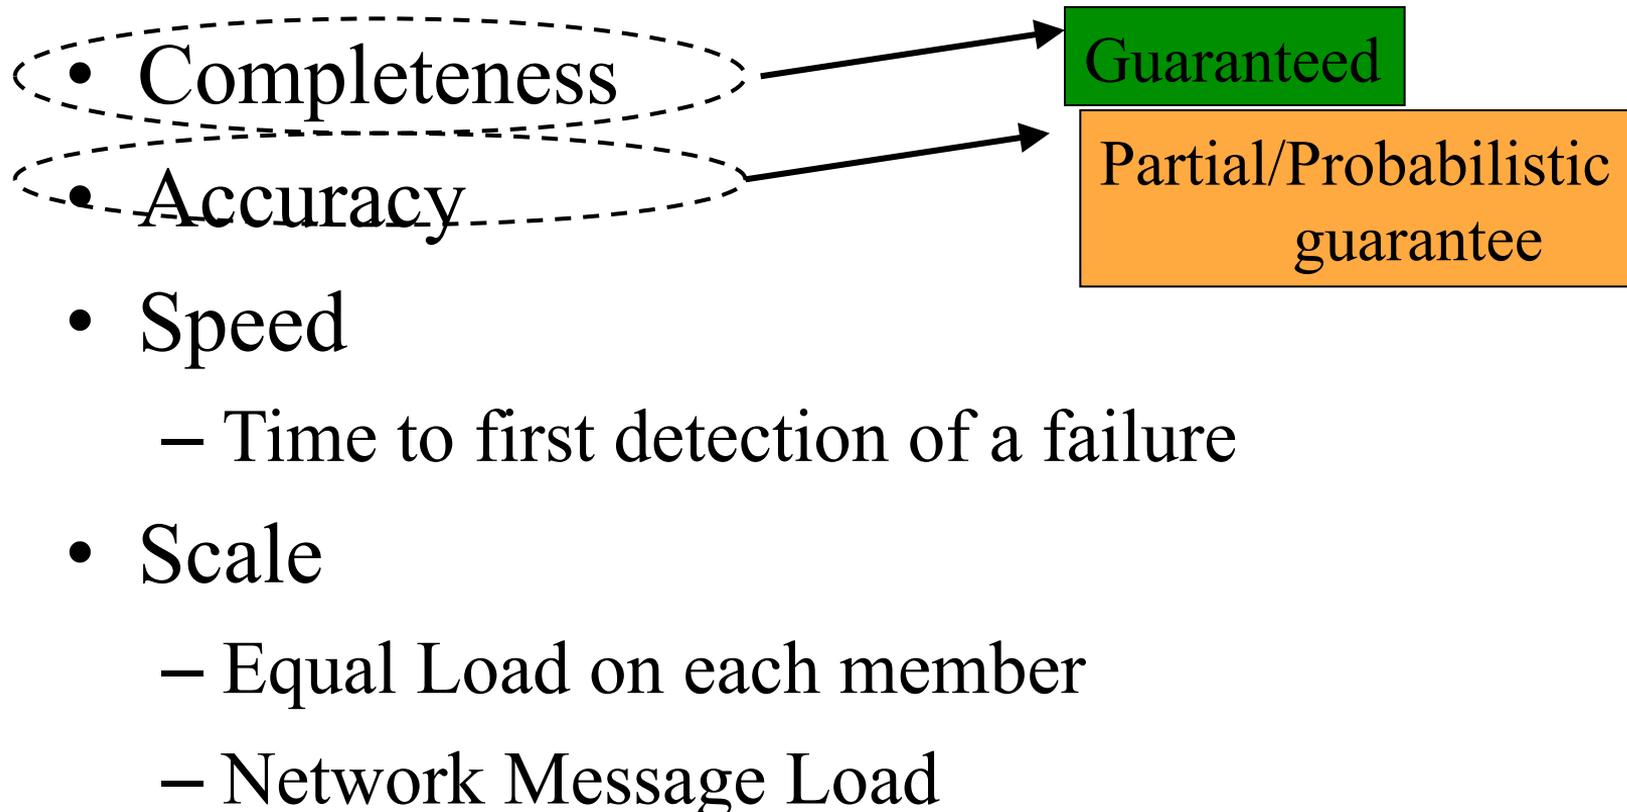
DISTRIBUTED FAILURE DETECTORS: PROPERTIES

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

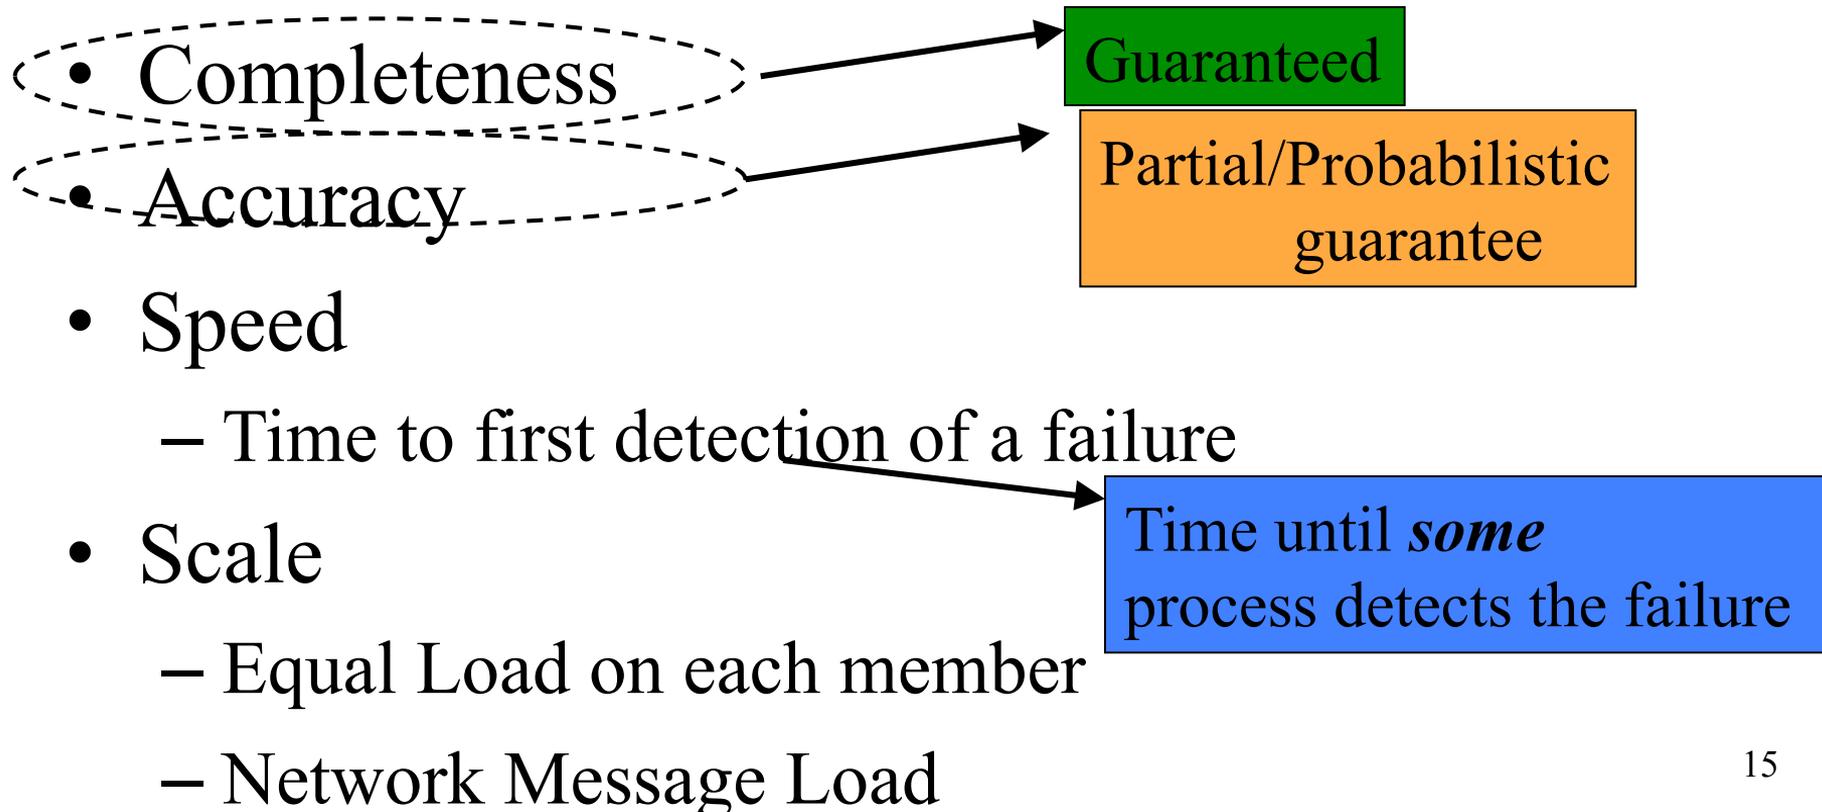
Impossible together in lossy networks [Chandra and Toueg]

If possible, then can solve consensus! (but consensus is known to be unsolvable in asynchronous systems)

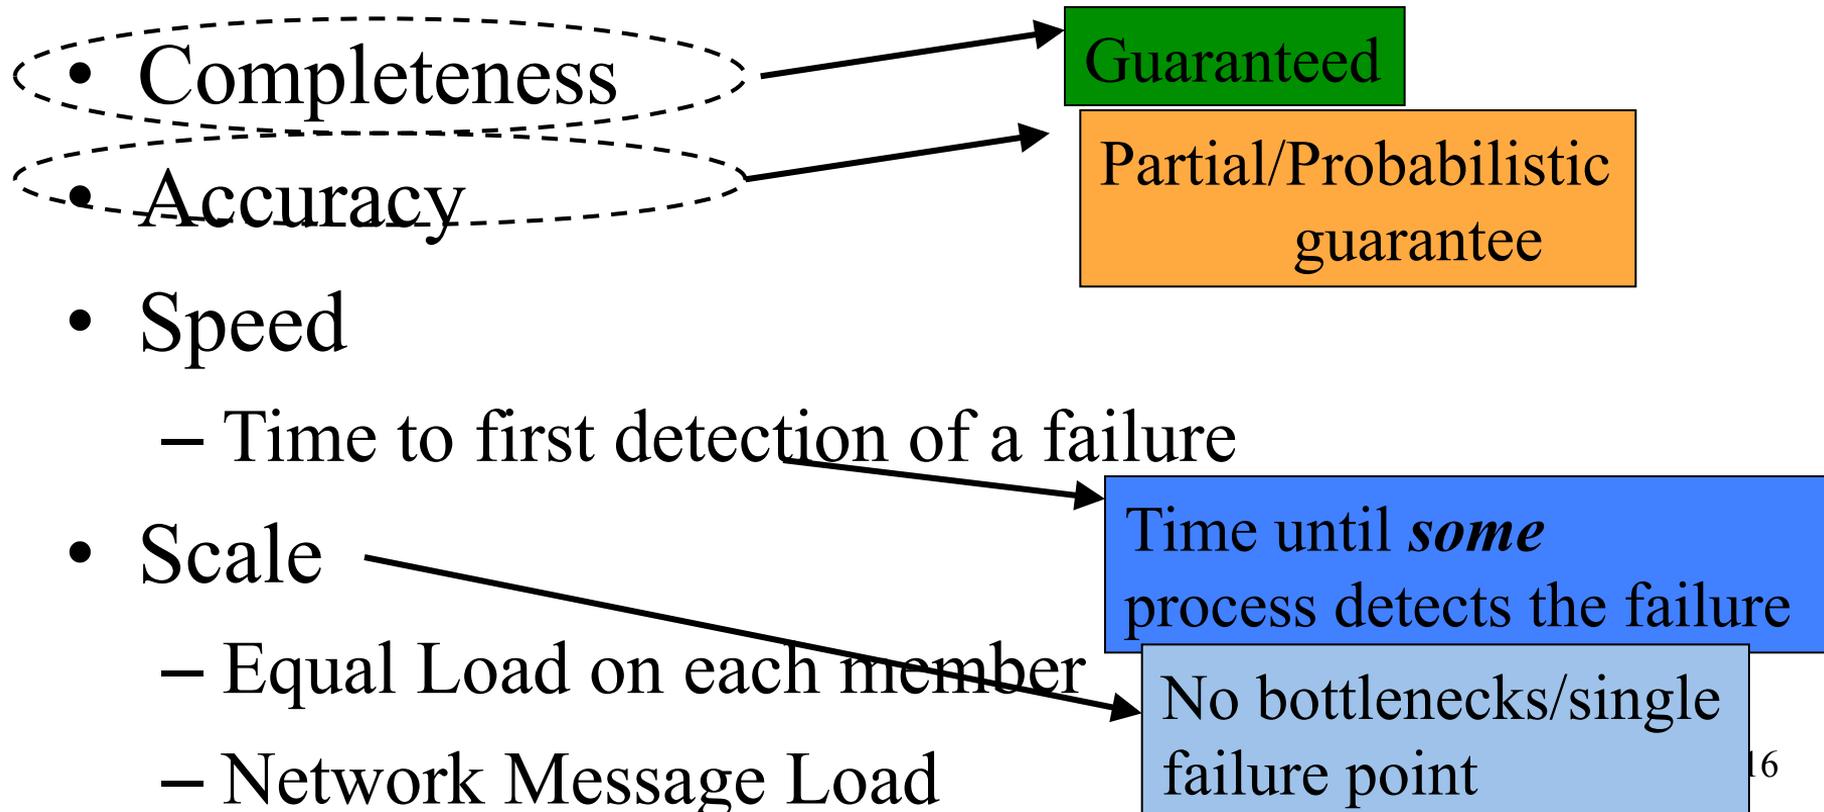
WHAT REAL FAILURE DETECTORS PREFER



WHAT REAL FAILURE DETECTORS PREFER



WHAT REAL FAILURE DETECTORS PREFER

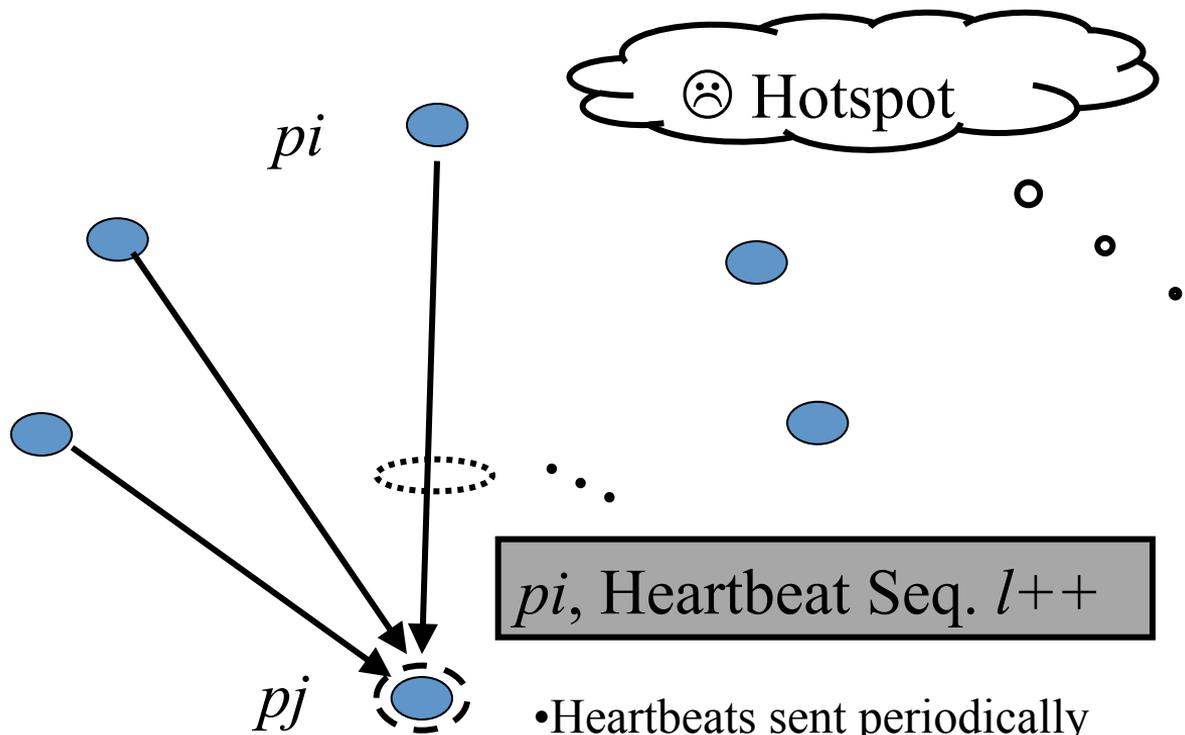


FAILURE DETECTOR PROPERTIES

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

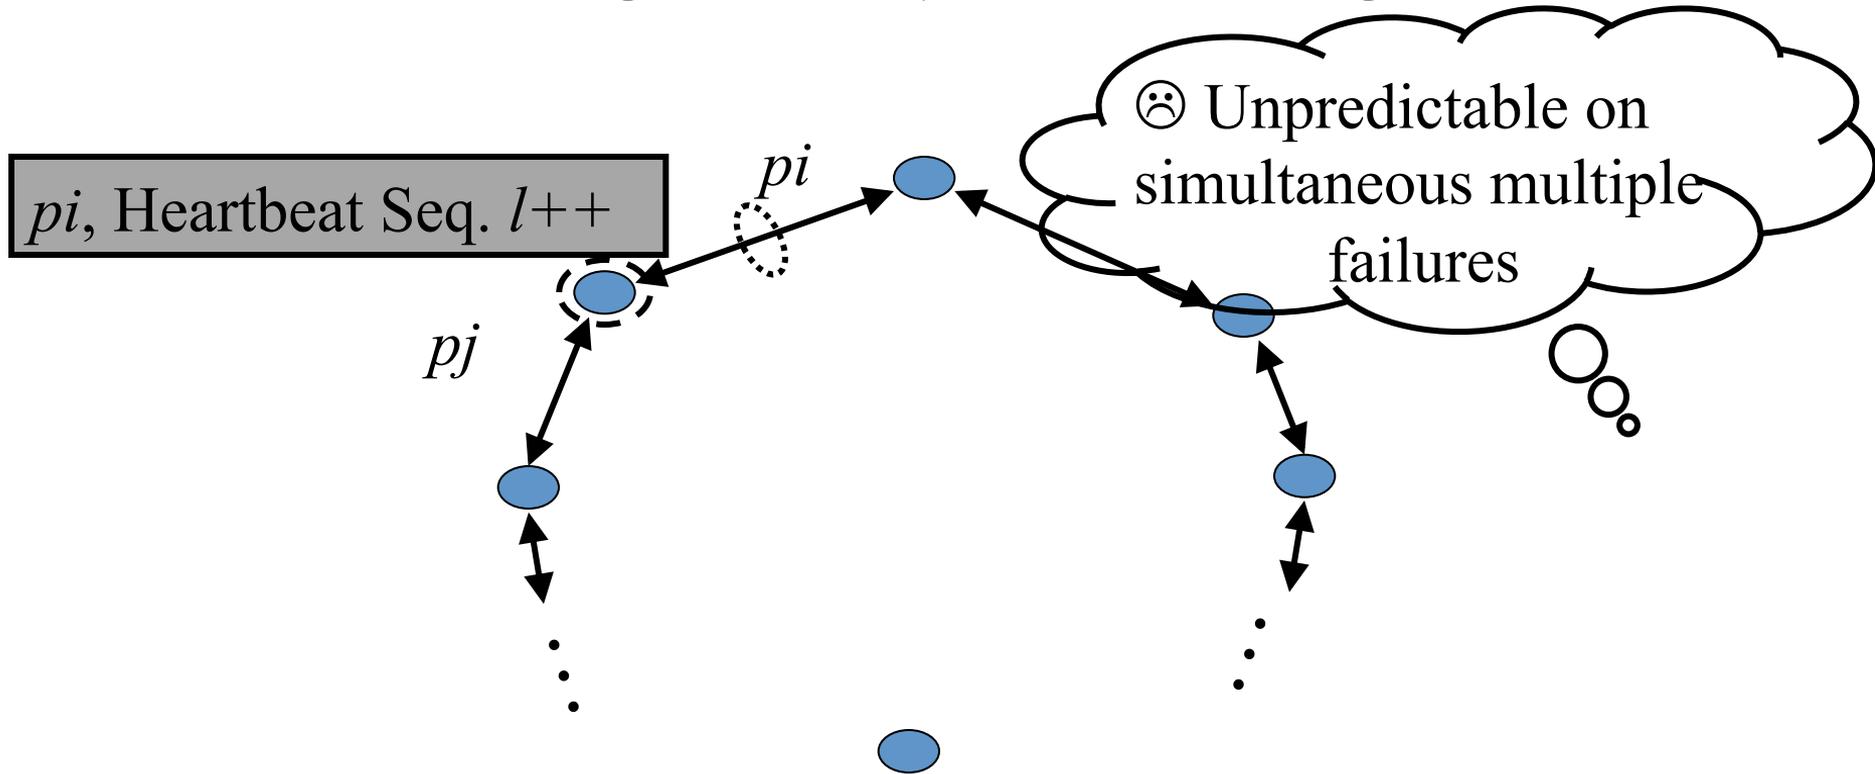
In spite of
arbitrary simultaneous
process failures

CENTRALIZED HEARTBEATING



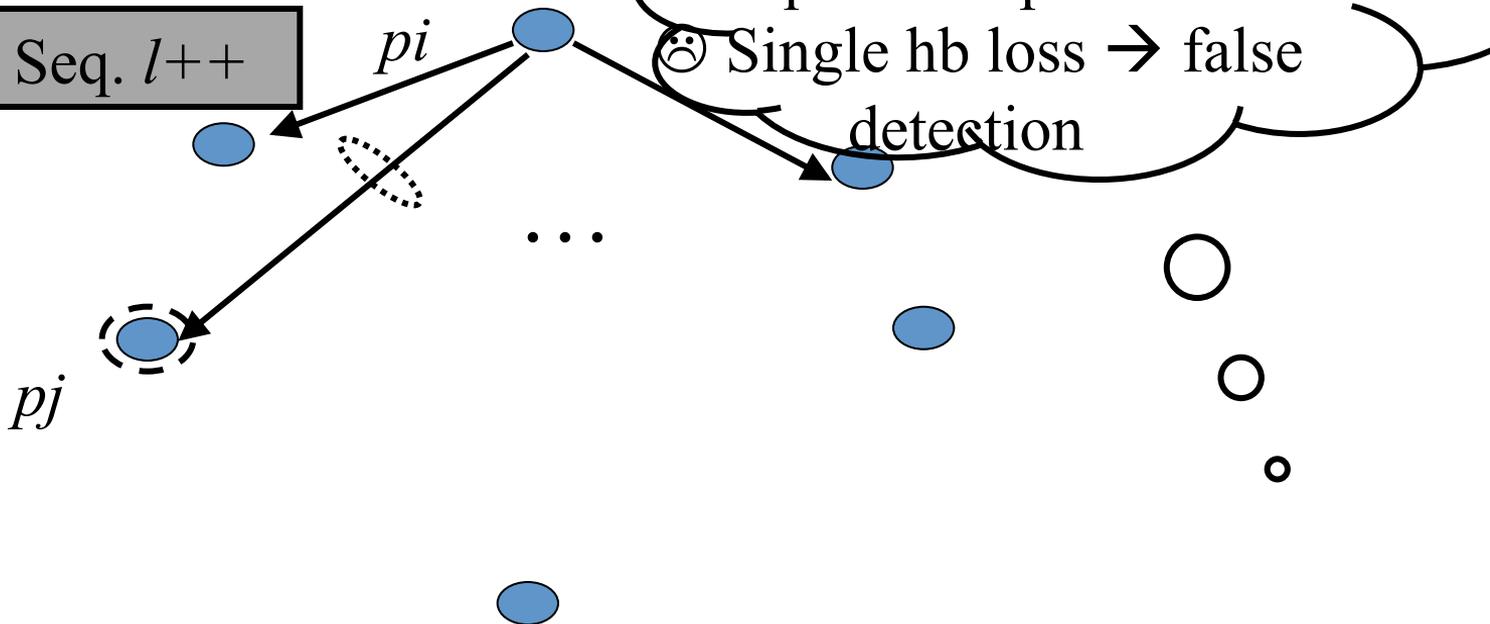
- Heartbeats sent periodically
- If heartbeat not received from p_i within timeout, mark p_i as failed

RING HEARTBEATING



ALL-TO-ALL HEARTBEATING

p_i , Heartbeat Seq. $l++$

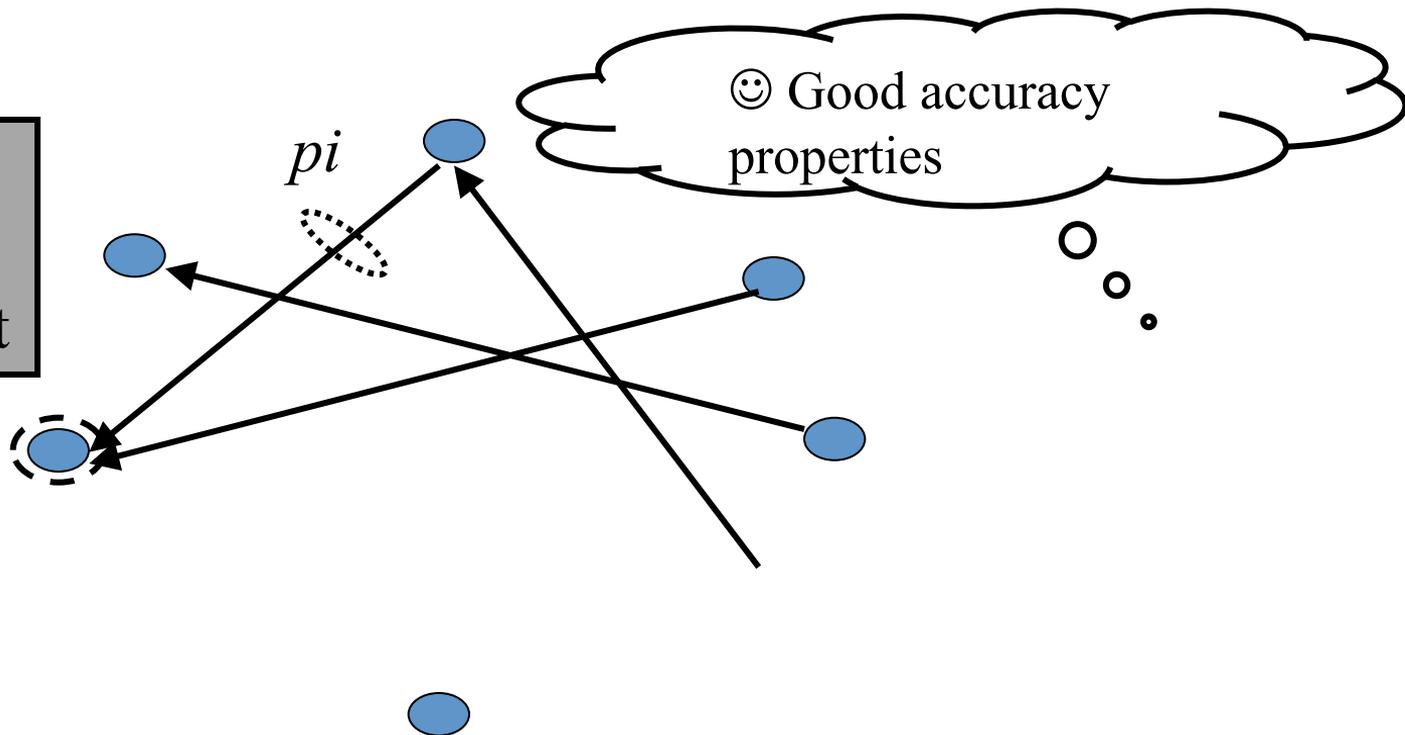


NEXT

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

GOSSIP-STYLE HEARTBEATING

Array of
Heartbeat Seq. l
for member subset

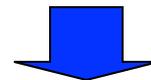


GOSSIP-STYLE FAILURE DETECTION

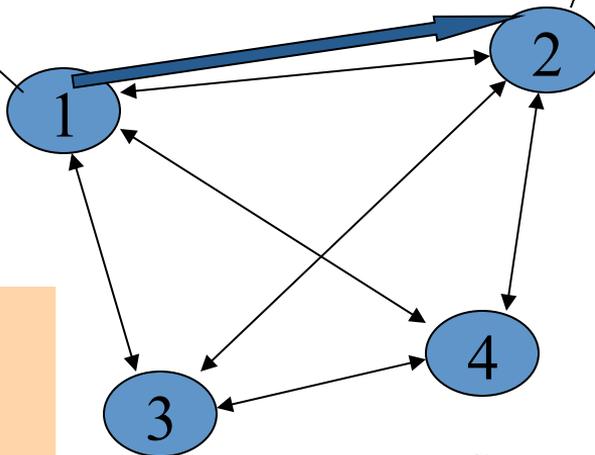
1	10120	66
2	10103	62
3	10098	63
4	10111	65

Address
Heartbeat Counter
Time (local)

1	10118	64
2	10110	64
3	10090	58
4	10111	65



1	10120	70
2	10110	64
3	10098	70
4	10111	65



Current time : 70 at node 2
(asynchronous clocks)

Protocol:

- Nodes periodically gossip their membership list: pick random nodes, send it list
- On receipt, it is *merged* with local membership list
- When an entry times out, member is marked as failed

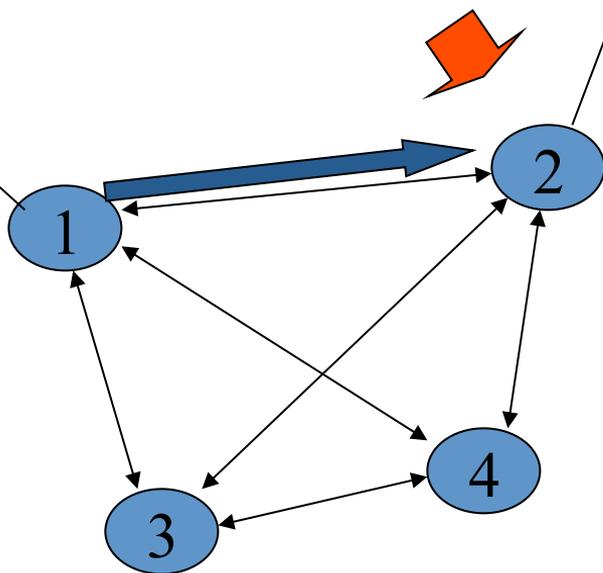
GOSSIP-STYLE FAILURE DETECTION

- If the heartbeat has not increased for more than T_{fail} seconds, the member is considered failed
- And after a further T_{cleanup} seconds, it will delete the member from the list
- Why an additional timeout? Why not delete right away?

GOSSIP-STYLE FAILURE DETECTION

- What if an entry pointing to a failed node is deleted right after T_{fail} ($=24$) seconds?

1	10120	66
2	10103	62
3	10098	55
4	10111	65



1	10120	66
2	10110	64
3	10098	66
4	10111	65

Current time : 75 at node 2

ANALYSIS/DISCUSSION

- Well-known result: a gossip takes $O(\log(N))$ time to propagate.
- So: Given sufficient bandwidth, a single heartbeat takes $O(\log(N))$ time to propagate.
- So: N heartbeats take:
 - $O(\log(N))$ time to propagate, if bandwidth allowed per node is allowed to be $O(N)$
 - $O(N \cdot \log(N))$ time to propagate, if bandwidth allowed per node is only $O(1)$
 - What about $O(k)$ bandwidth?
- What happens if gossip period T_{gossip} is decreased?
- What happens to P_{mistake} (false positive rate) as $T_{\text{fail}}, T_{\text{cleanup}}$ is increased?
- **Tradeoff: False positive rate vs. detection time vs. bandwidth**

NEXT

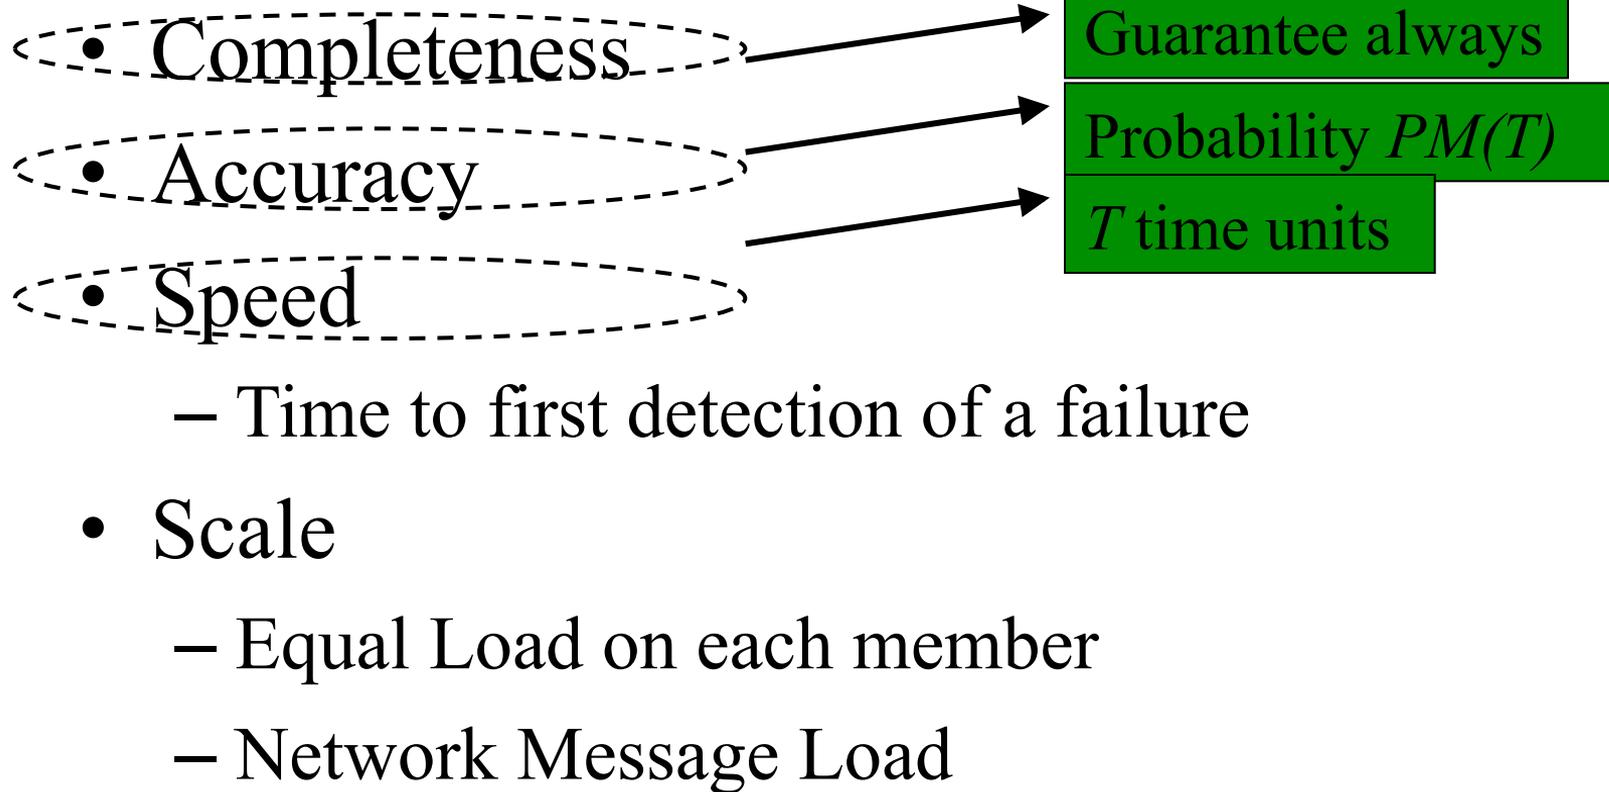
- So, is this the best we can do? What is the best we can do?

FAILURE DETECTOR PROPERTIES ...

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

...ARE APPLICATION-DEFINED

REQUIREMENTS



...ARE APPLICATION-DEFINED

REQUIREMENTS

• Completeness

• Accuracy

• Speed

– Time to first detection of a failure

• Scale

– Equal Load on each member

– Network Message Load

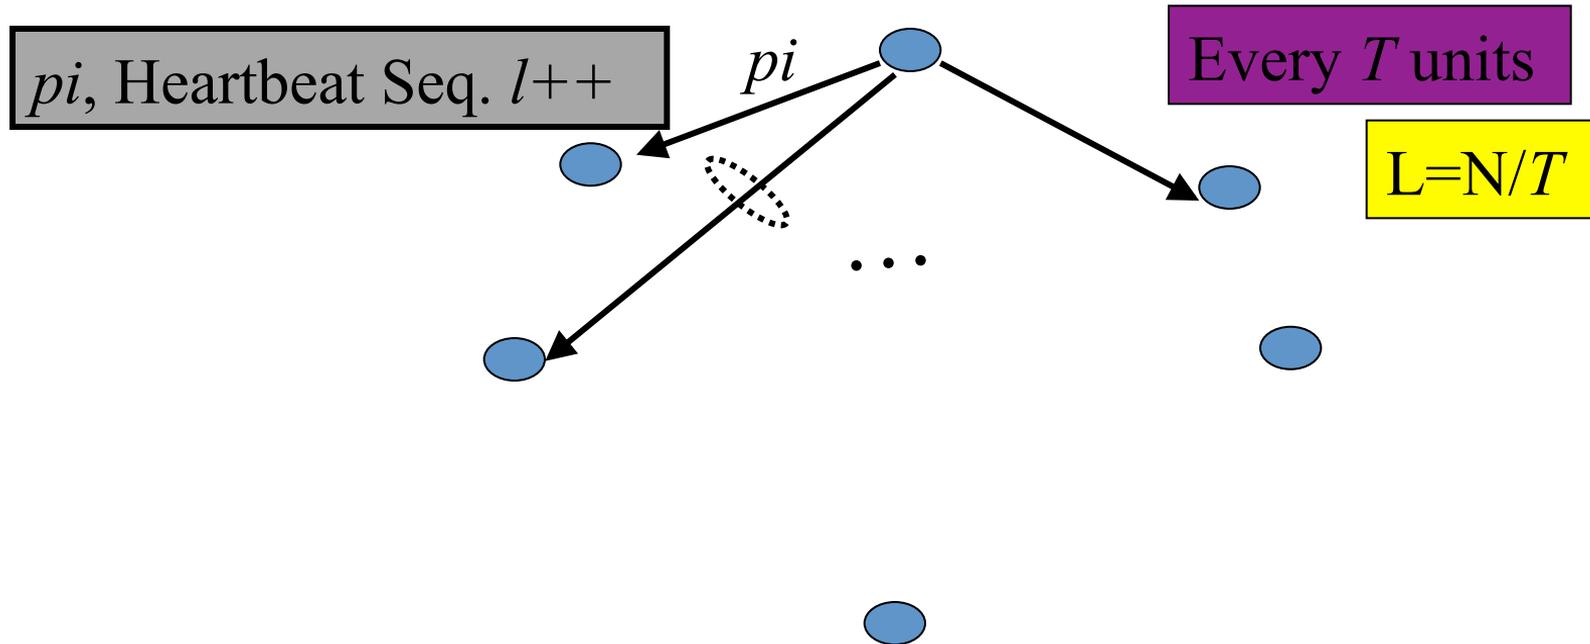
Guarantee always

Probability $PM(T)$

T time units

$N*L$: Compare this across protocols

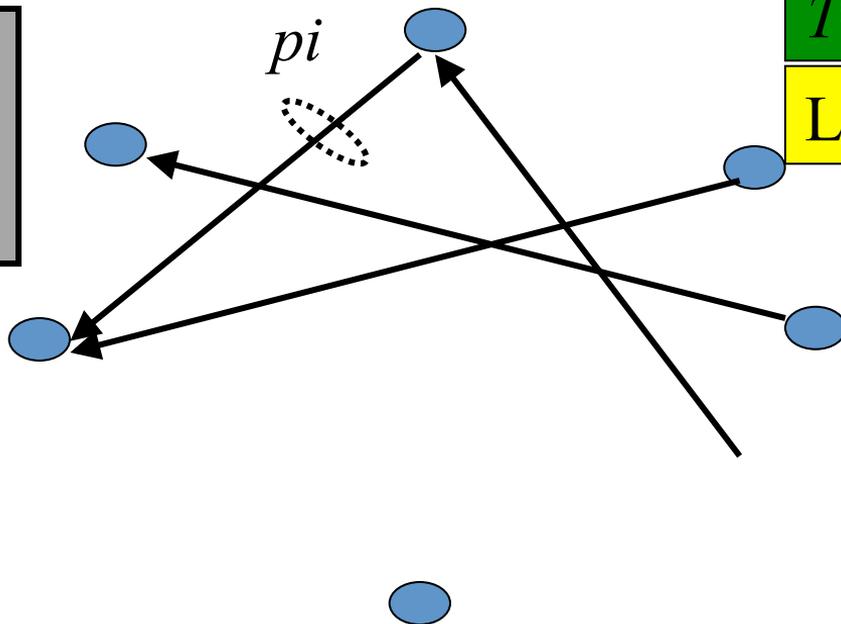
ALL-TO-ALL HEARTBEATING



GOSSIP-STYLE HEARTBEATING

Array of
Heartbeat Seq. l
for member subset

Every tg units
=gossip period,
send $O(N)$ gossip
message



$$T = \log N * tg$$

$$L = N/tg = N * \log N / T$$

WHAT'S THE BEST/OPTIMAL WE CAN DO?

- *Worst case* load L^* **per member** in the group (messages per second)
 - as a function of T , $PM(T)$, N
 - Independent Message Loss probability p_{ml}

- $$L^* = \frac{\log(PM(T))}{\log(p_{ml})} \cdot \frac{1}{T}$$

HEARTBEATING

- Optimal L is independent of N (!)
- All-to-all and gossip-based: sub-optimal
 - $L=O(N/T)$
 - try to achieve simultaneous detection at *all* processes
 - fail to distinguish *Failure Detection* and *Dissemination* components

⇒ Can we reach this bound?

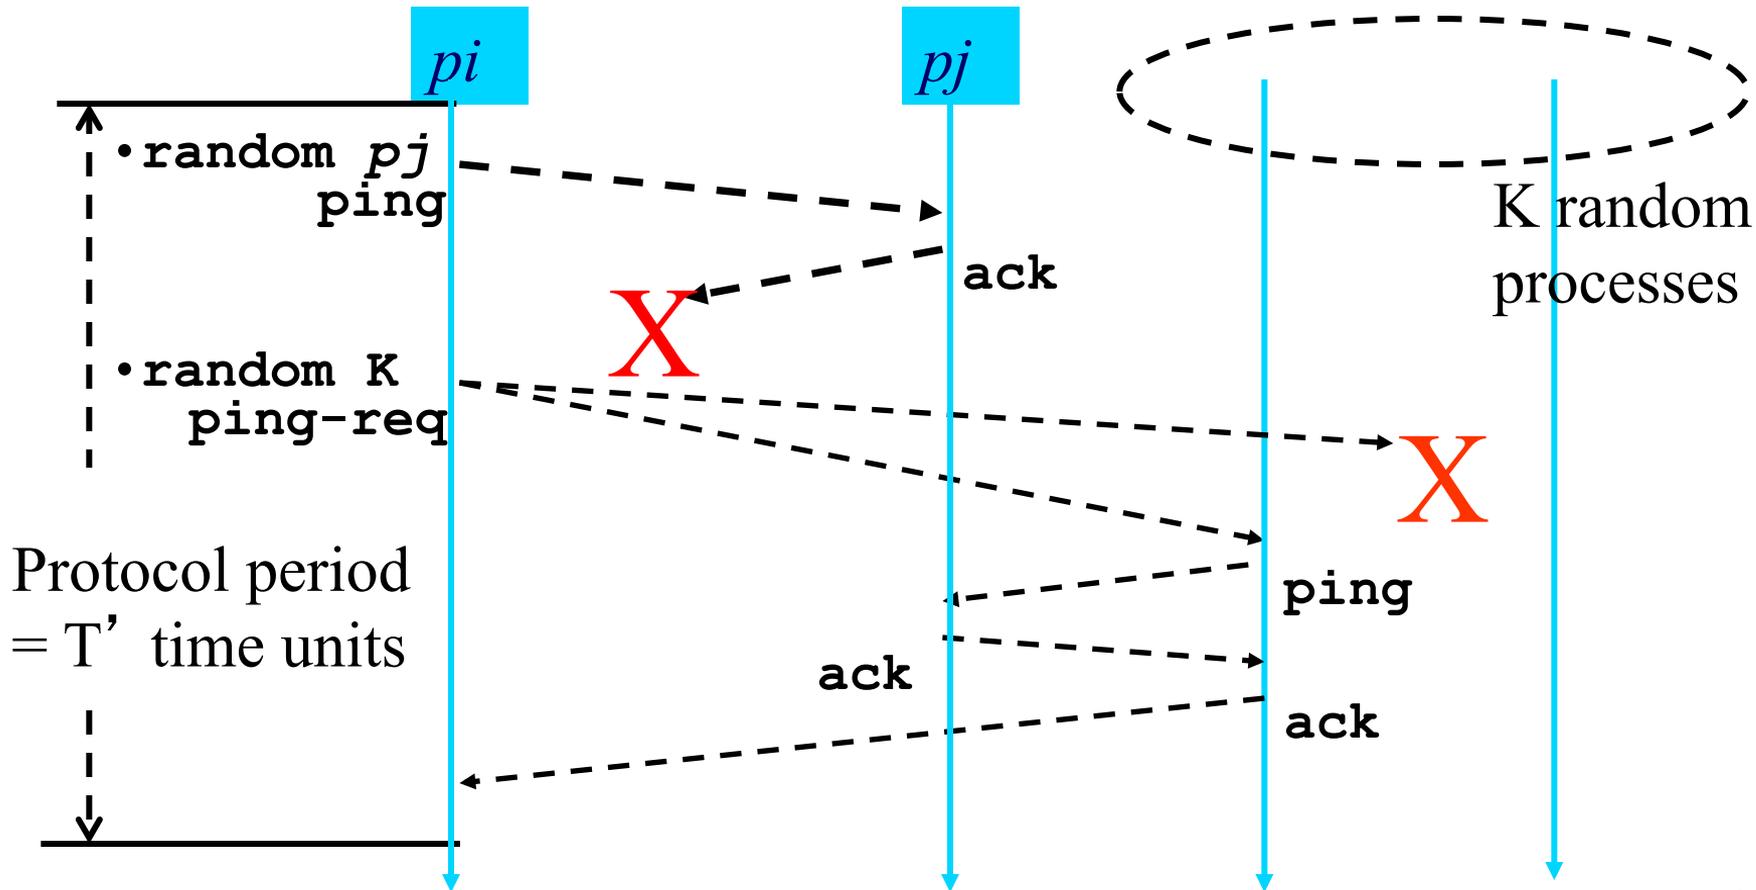
⇒ Key:

- Separate the two components
- Use a non heartbeat-based Failure Detection Component

NEXT

- Is there a better failure detector?

SWIM FAILURE DETECTOR PROTOCOL



DETECTION TIME

- Prob. of being pinged in $T' = 1 - \left(1 - \frac{1}{N}\right)^{N-1} = 1 - e^{-1}$
- $E[T] = T' \cdot \frac{e}{e-1}$
- Completeness: *Any* alive member detects failure
 - Eventually
 - By using a trick: within worst case $O(N)$ protocol periods

ACCURACY, LOAD

- $PM(T)$ is exponential in $-K$. Also depends on pml (and pf)
 - See paper

- $\frac{L}{L^*} < 28$ $\frac{E[L]}{L^*} < 8$ for up to 15 % loss rates

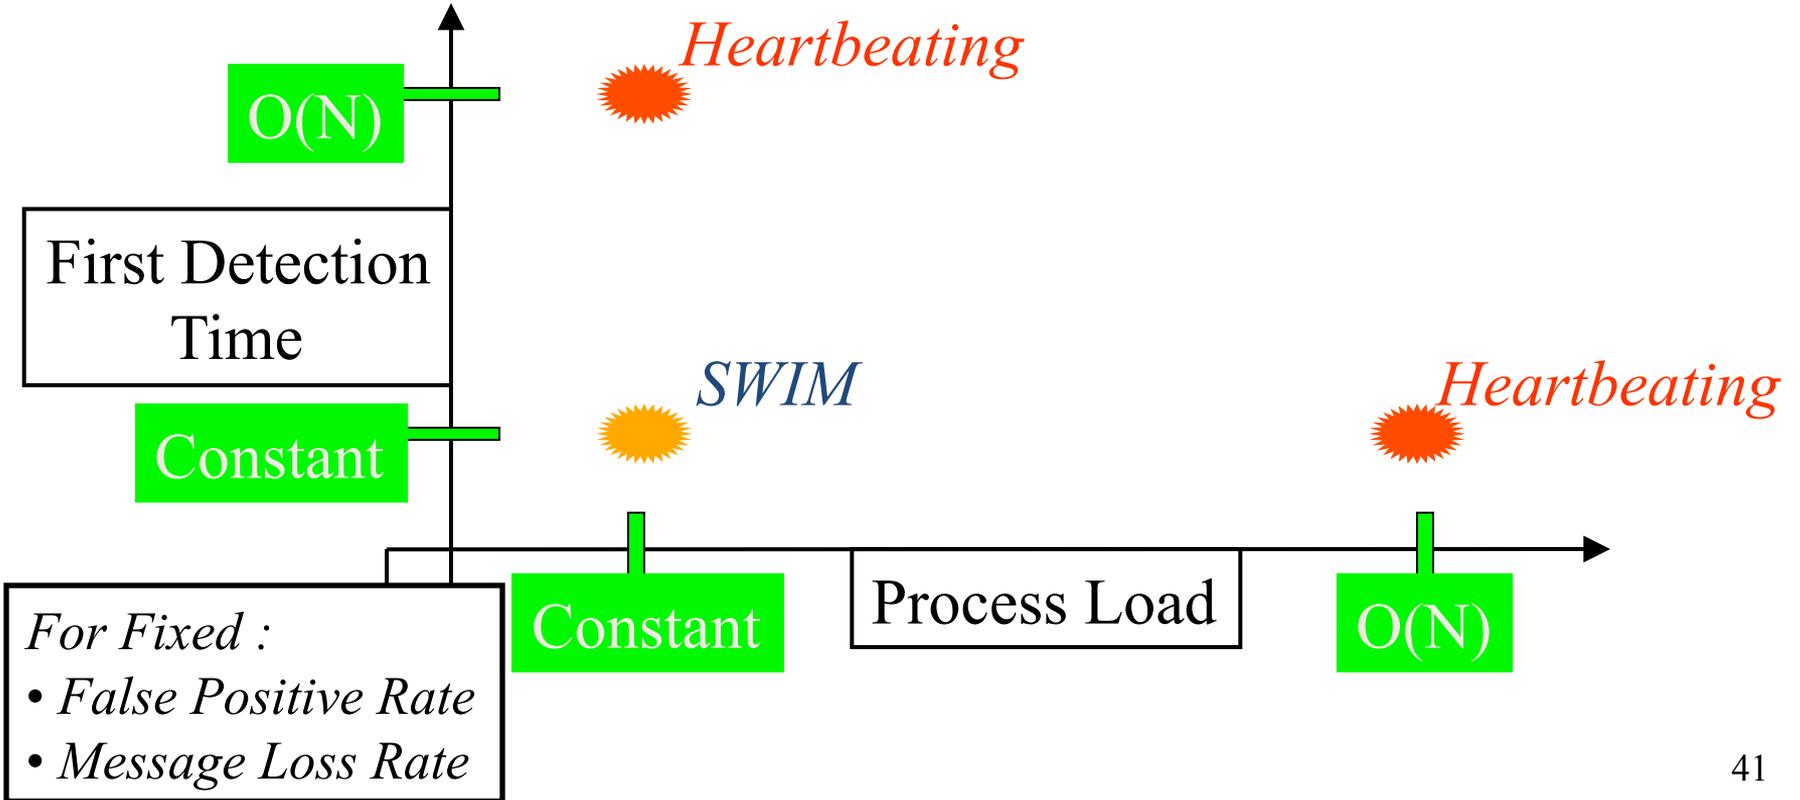
SWIM FAILURE DETECTOR

Parameter	SWIM
First Detection Time	<ul style="list-style-type: none">• Expected $\left[\frac{e}{e-1} \right]$ periods• Constant (independent of group size)
Process Load	<ul style="list-style-type: none">• Constant per period• $< 8 L^*$ for 15% loss
False Positive Rate	<ul style="list-style-type: none">• Tunable (via K)• Falls exponentially as load is scaled
Completeness	<ul style="list-style-type: none">• Deterministic time-bounded• Within $O(\log(N))$ periods w.h.p.

TIME-BOUNDED COMPLETENESS

- Key: select each membership element once as a ping target in a traversal
 - Round-robin pinging
 - Random permutation of list after each traversal
- Each failure is detected in worst case $2N-1$ (local) protocol periods
- Preserves FD properties

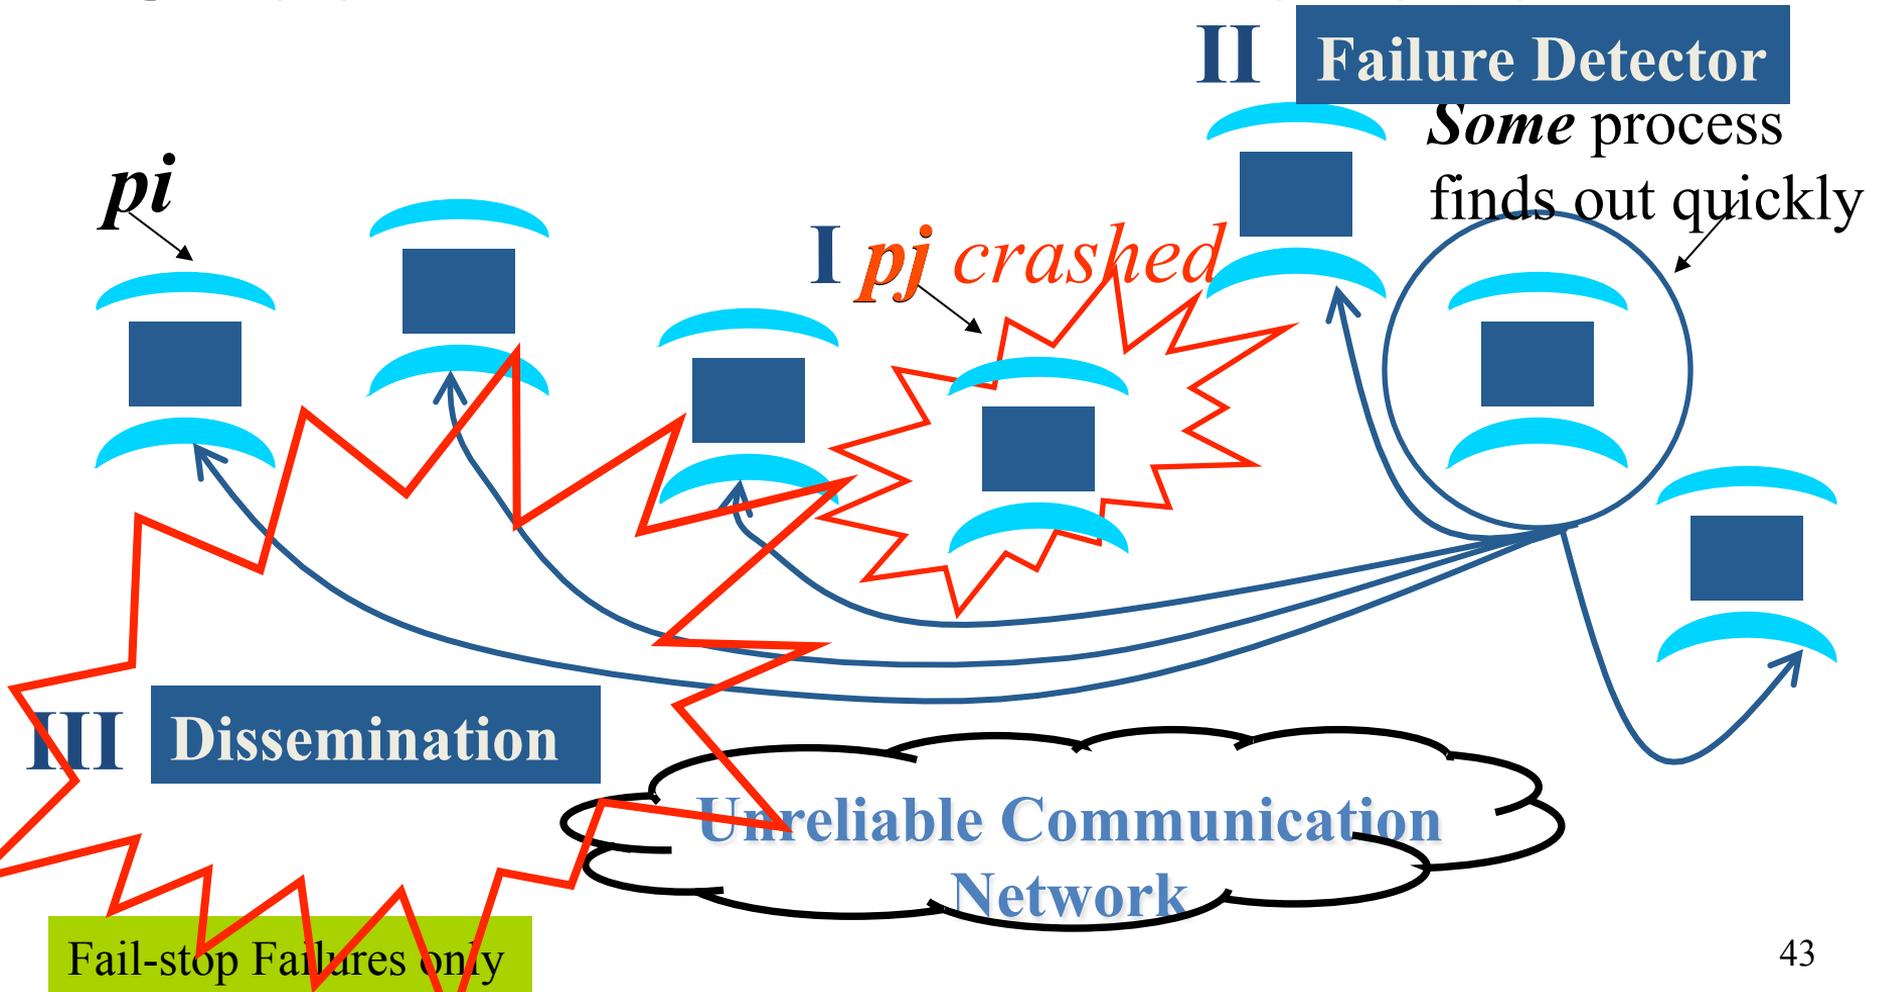
SWIM VERSUS HEARTBEATING



NEXT

- How do failure detectors fit into the big picture of a group membership protocol?
- What are the missing blocks?

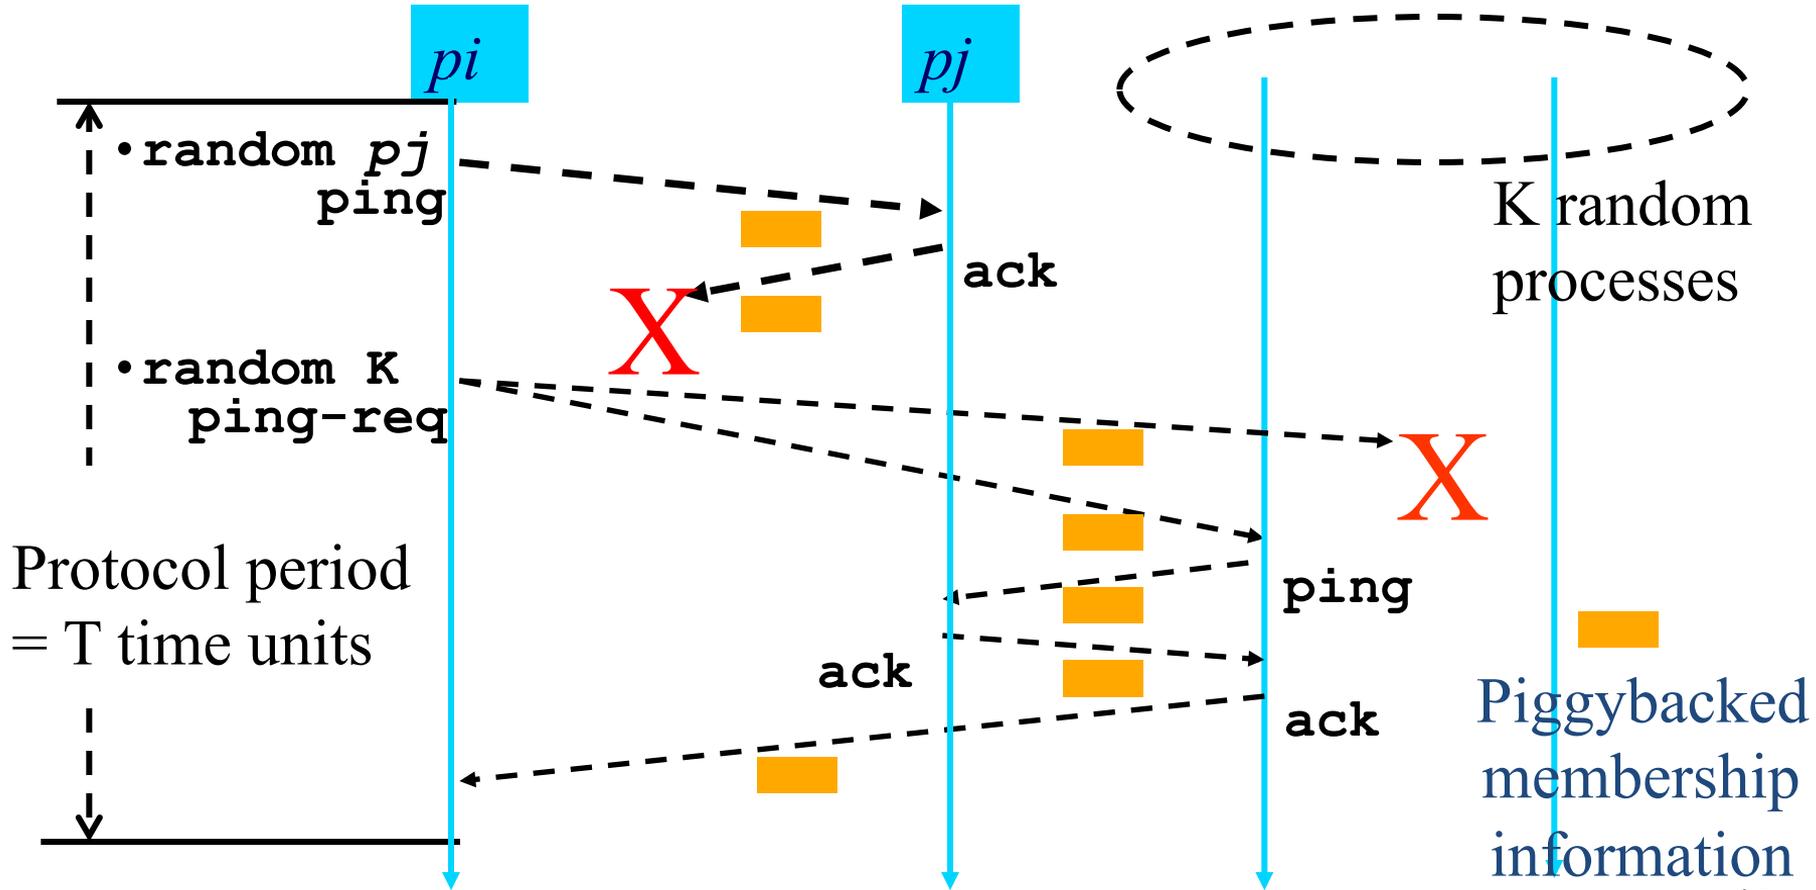
GROUP MEMBERSHIP PROTOCOL



DISSEMINATION OPTIONS

- Multicast (Hardware / IP)
 - unreliable
 - multiple simultaneous multicasts
- Point-to-point (TCP / UDP)
 - expensive
- Zero extra messages: Piggyback on Failure Detector messages
 - Infection-style Dissemination

INFECTION-STYLE DISSEMINATION



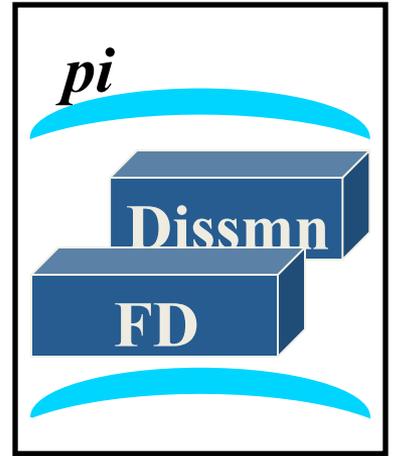
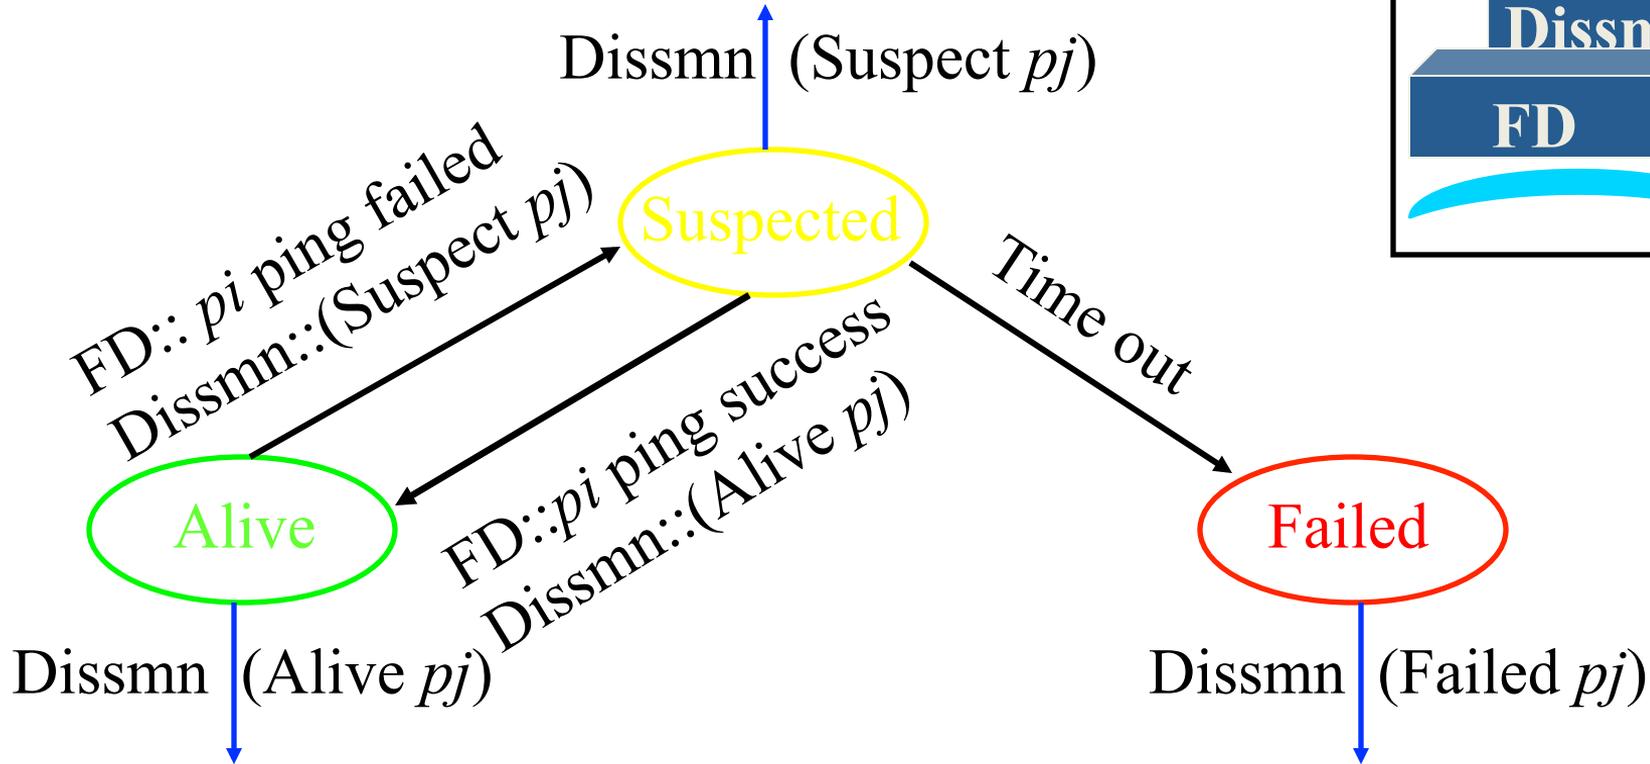
INFECTION-STYLE DISSEMINATION

- Epidemic/Gossip style dissemination
 - After $\lambda \cdot \log(N)$ protocol periods, $N^{-(2\lambda-2)}$ processes would not have heard about an update
- Maintain a buffer of recently joined/evicted processes
 - Piggyback from this buffer
 - Prefer recent updates
- Buffer elements are garbage collected after a while
 - After $\lambda \cdot \log(N)$ protocol periods, i.e., once they've propagated through the system; this defines weak consistency

SUSPICION MECHANISM

- False detections, due to
 - Perturbed processes
 - Packet losses, e.g., from congestion
- Indirect pinging may not solve the problem
- Key: *suspect* a process before *declaring* it as failed in the group

SUSPICION MECHANISM



SUSPICION MECHANISM

- Distinguish multiple suspicions of a process
 - Per-process *incarnation number*
 - *Inc #* for pi can be incremented only by pi
 - e.g., when it receives a (Suspect, pi) message
 - Somewhat similar to DSDV (routing protocol in ad-hoc nets)
- Higher *inc#* notifications over-ride lower *inc#*'s
- Within an *inc#*: (Suspect *inc #*) > (Alive, *inc #*)
- (Failed, *inc #*) overrides everything else

SWIM IN INDUSTRY

- First used in Oasis/CoralCDN
- Implemented open-source by Hashicorp Inc.
 - Called “Serf”
- Today: Uber implemented it, uses it for failure detection in their infrastructure

WRAP UP

- Failures the norm, not the exception in datacenters
- Every distributed system uses a failure detector
- Many distributed systems use a membership service

- Ring failure detection underlies
 - IBM SP2 and many other similar clusters/machines

- Gossip-style failure detection underlies
 - Amazon EC2/S3 (rumored!)

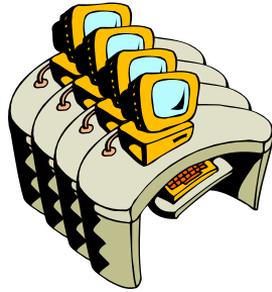
GRID COMPUTING

EXAMPLE: RAPID ATMOSPHERIC MODELING SYSTEM, COLOSTATE U

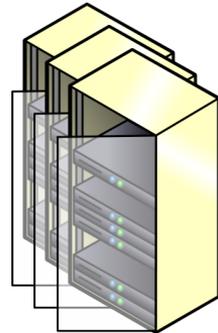
- Hurricane Georges, 17 days in Sept 1998
 - “RAMS modeled the mesoscale convective complex that dropped so much rain, in good agreement with recorded data”
 - Used 5 km spacing instead of the usual 10 km
 - Ran on 256+ processors
- Computation-intensive computing (or HPC = high performance computing)
- *Can one run such a program without access to a supercomputer?*

DISTRIBUTED COMPUTING RESOURCES

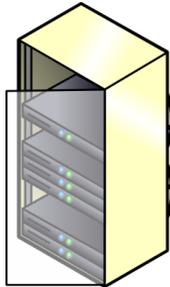
Wisconsin



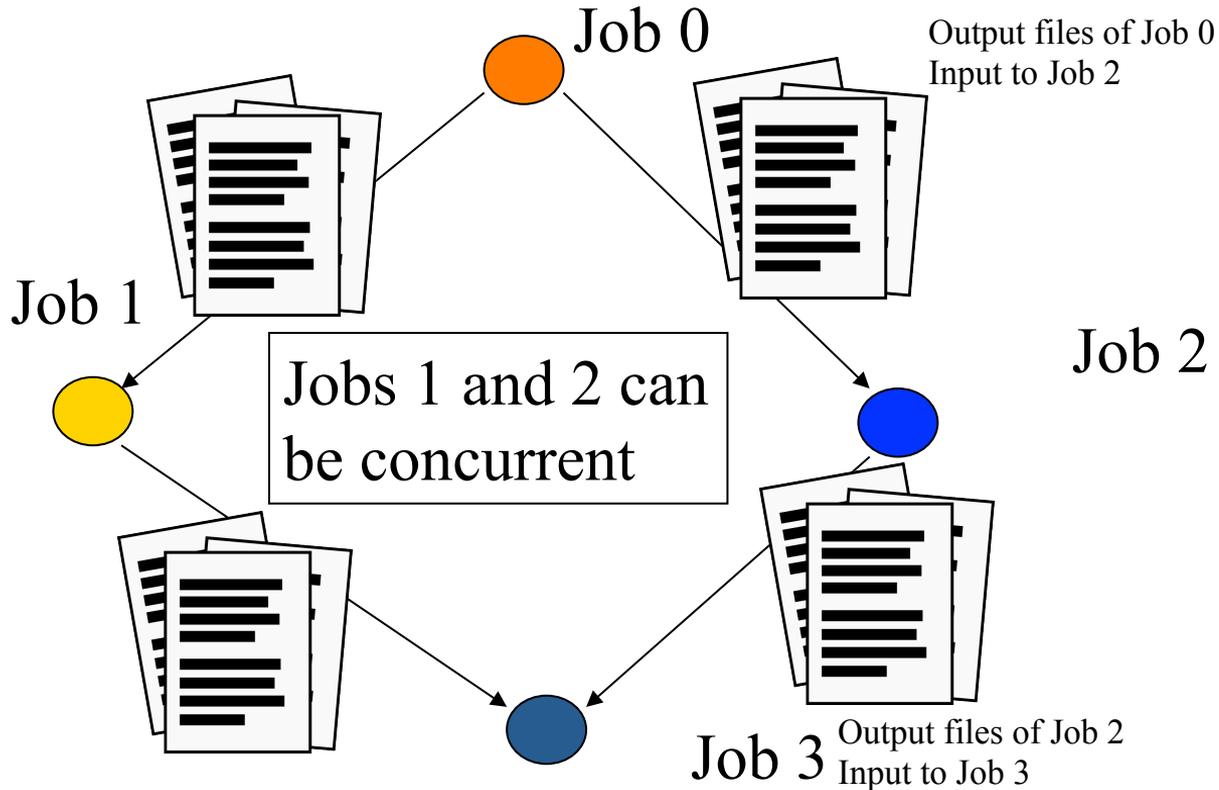
NCSA



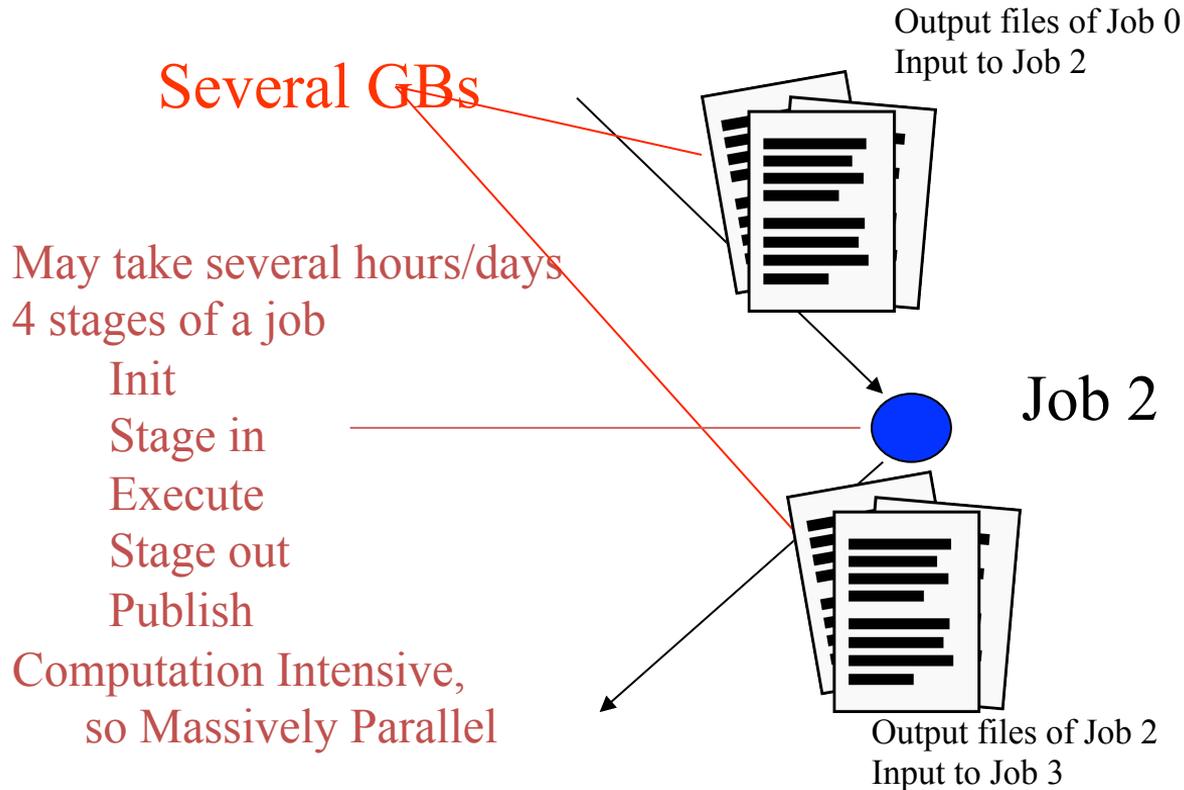
MIT



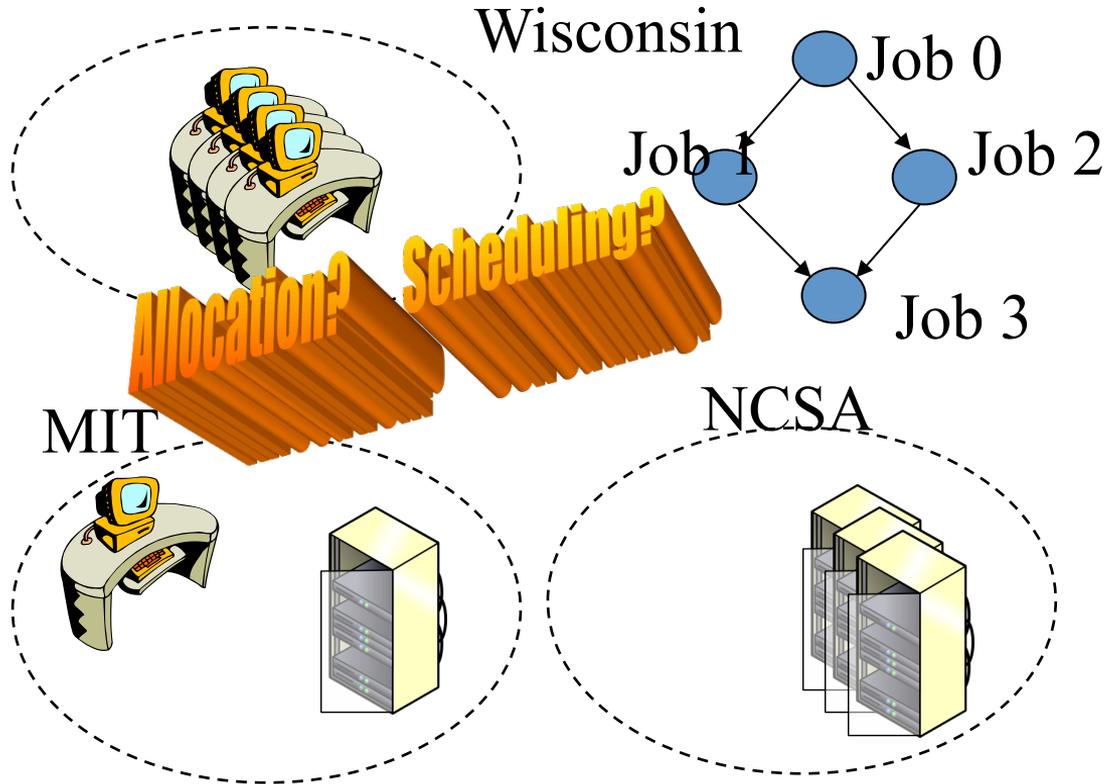
AN APPLICATION CODED BY A PHYSICIST



AN APPLICATION CODED BY A PHYSICIST

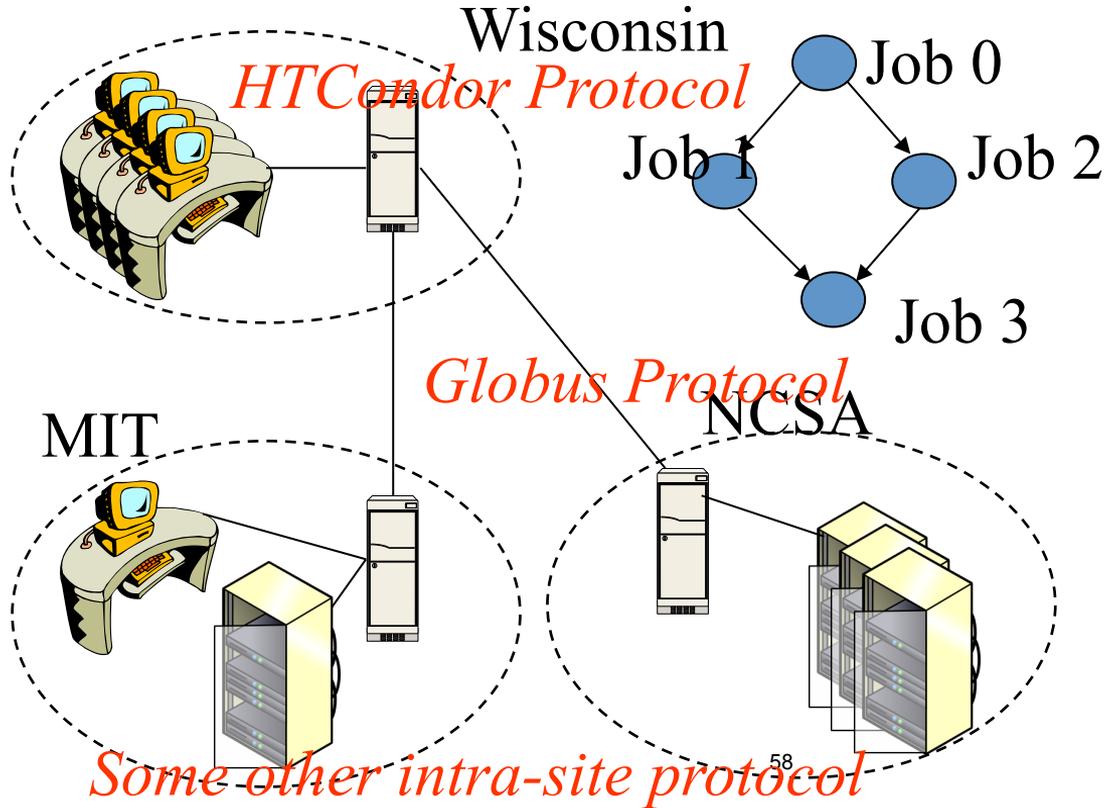


SCHEDULING PROBLEM



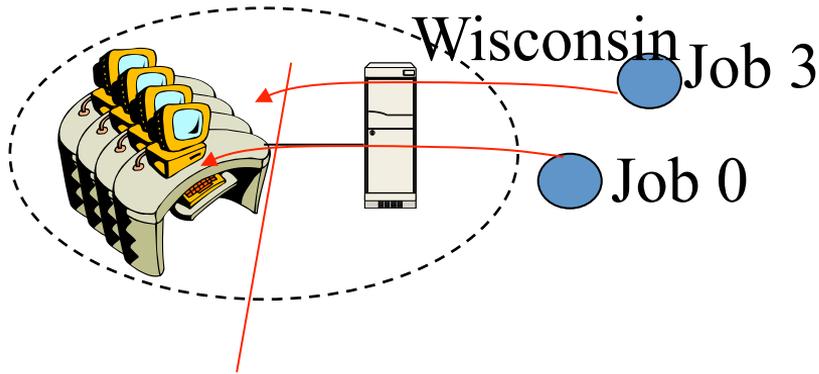
2-LEVEL SCHEDULING

INFRASTRUCTURE



INTRA-SITE PROTOCOL

HTCondor Protocol



Internal Allocation & Scheduling
Monitoring
Distribution and Publishing of Files

CONDOR (NOW HTCONDOR)

- High-throughput computing system from U. Wisconsin Madison
- Belongs to a class of “Cycle-scavenging” systems
 - SETI@Home and Folding@Home are other systems in this category

Such systems

- Run on a lot of workstations
- When workstation is free, ask site’s central server (or Globus) for tasks
- If user hits a keystroke or mouse click, stop task
 - Either kill task or ask server to reschedule task
- Can also run on dedicated machines

GLOBUS

- Globus Alliance involves universities, national US research labs, and some companies
- Standardized several things, especially software tools
- Separately, but related: Open Grid Forum
- Globus Alliance has developed the Globus Toolkit

<http://toolkit.globus.org/toolkit/>

GLOBUS TOOLKIT

- Open-source
- Consists of several components
 - [GridFTP](#): Wide-area transfer of bulk data
 - [GRAM5](#) (Grid Resource Allocation Manager): submit, locate, cancel, and manage jobs
 - Not a scheduler
 - Globus communicates with the schedulers in intra-site protocols like HTCondor or Portable Batch System (PBS)
 - [RLS](#) (Replica Location Service): Naming service that translates from a file/dir name to a target location (or another file/dir name)
 - Libraries like [XIO](#) to provide a standard API for all Grid IO functionalities
 - Grid Security Infrastructure ([GSI](#))

SECURITY ISSUES

- Important in Grids because they are *federated*, i.e., no single entity controls the entire infrastructure
- **Single sign-on**: collective job set should require once-only user authentication
- **Mapping to local security mechanisms**: some sites use Kerberos, others using Unix
- **Delegation**: credentials to access resources inherited by subcomputations, e.g., job 0 to job 1
- **Community authorization**: e.g., third-party authentication
- These are also important in clouds, but less so because clouds are typically run under a central control
- In clouds the focus is on failures, scale, on-demand nature

SUMMARY

- Grid computing focuses on computation-intensive computing (HPC)
- Though often federated, architecture and key concepts have a lot in common with that of clouds
- Are Grids/HPC converging towards clouds?
 - E.g., Compare OpenStack and Globus

ANNOUNCEMENTS

- MP1 –
 - Demo signup sheet available on Piazza
 - Demo details available on Piazza
 - Make sure you print individual and total linecounts
- Check Piazza often! It's where all the announcements are at!