

CS425, Distributed Systems: Fall 2025

Machine Programming 2 – Distributed Group Membership

Released Date: September 16, 2025

Due Date (Hard Deadline): Sunday September 28, 2025 (Code and report due at 11.59 PM)

Demos on Monday September 29, 2025

The time for this MP is rather short! So please start early! Start now!

You must work in groups of two for this MP. Please stick with the groups you formed for MP1. (See end of document for expectations from group members.)

Your company from MP1 just got acquired by the fictitious FakeCompany Inc. (this fake company's business model is to detect, not generate, fake news). This company liked your previous work in MP1, so they've commissioned you to build a distributed group membership service for them.

You must work in groups of two for this MP.

This service maintains, at each machine in the system (at a daemon process), a list of the group, i.e., other machines that are connected and up. This membership list is a full membership list, and needs to be updated whenever:

1. A machine (or its daemon) joins the group;
2. A machine (or its daemon) voluntarily leaves the group; or
3. A machine (or its daemon) crashes from the group (you may assume that the machine does not recover for a long enough time).

There is only one group at any point of time. Since we're implementing the crash/fail-stop model, when a machine rejoins, it must do so with a node id that includes not only IP and port but also a timestamp or version number that distinguishes successive versions of the same machine (this node id is what is held in the membership lists). This timestamp/version is needed because our model is a fail-stop model and not a fail-recovery model, so subsequent rejoins of the same VM must be distinguishable. (Don't confuse these timestamps/version numbers with suspicion incarnations, which mean something different).

A machine failure must be reflected in at least one membership list within 3 seconds (assuming synchronized clocks), and it must be provided no matter what the network latencies are. A machine failure, join or leave must be reflected within 6 seconds at *all* membership lists, assuming small network latencies – this is called *time-bounded completeness*, and it must be provided no matter what the network latencies are.

You're told that **at most three machines can fail simultaneously**, and after any set of simultaneous (up to 3 machines) have failed, the next set of failure(s) don't happen for at least 15 seconds (i.e., enough time for your system to converge back to a good topology). Your system must ensure completeness for all such failures (up to three simultaneous failures). That is, whenever a process fails, all membership lists that contained the process must be updated within 6 seconds (for 10 VMs).

Your algorithm must be scalable to large numbers of machines. For your experiments however, you may assume that you have $N > 5$ machines in the group at any given time ($N=10$ for the demo). Note that this is not a limit on the set of machines eligible to be group members. Typical runs will involve about 7-10 VMs.

Implement **two variants**:

- I. **Gossip**: Gossip-style heartbeating with a Suspicion mechanism (like SWIM, but without any pinging). Note that the original Gossip does not have Suspicion, so you will have to design it carefully.
- II. **PingAck**: Implement Ping-Ack (as discussed in class in the SWIM protocol) and augment it with Suspicion mechanism. You do not need to implement any other optimizations from SWIM that you do not consider necessary.

In both cases, you will need to implement incarnations – see lecture slides/video (and also page 7 of the SWIM paper from the website). We recommend that you implement the Suspicion mechanism in a modular way so that both Gossip and Ping-Ack can call it. Beware of subtleties in how each of those call Suspicion, and ensure you handle them carefully.

You will also get to compare these variants experimentally (see Report section). Think of the parameter settings you need (heartbeat frequency, frequency of pings, timeouts, etc.) to achieve the time bounds mentioned above.

Note that in gossip-style heartbeating or PingAck with Suspicion, detection time is measured when the entry is finally deleted (after suspicion timeout), not when the failure detection timeout first expires. This is consistent with application (on top of membership) expectations of being told about the failure just once, only after it is confirmed (locally). Once the failure detection layer tells (confirms) an application about a failure, it cannot be rescinded.

Design first, then implement. Keep your design (and implementation) as simple as possible. Use the adage “KISS: Keep It Simple Si...”. Otherwise FakeCompany Inc. may generate fake news about you, give you fake points, and then have KISS (the musical group) sing it in a duet with Taylor Swift (which naturally wouldn't be nice for anyone's sanity in this entire world).

In your report, justify your design choices and why your design meets 6 second completeness. Your algorithm must use a small bandwidth (defined as Bytes per second, and NOT as messages per second) needed to meet the above requirements.

For the failure detection or leaves, you cannot use a leader (archaic: master), since its failure must be detected as well. However, to enable machines to join the group, you can have a fixed contact machine that all potential (joining) members know about, which you already know is called the “introducer”. When the introducer is down, no new members can join the group until the contact has rejoined – but the rest of the group should proceed normally including failures should still being detected, and leaves being allowed.

Pay attention to the format of messages that are sent between machines. Ensure that any platform-dependent fields (e.g., ints) are marshaled (converted) into a platform-independent format. An example is Google’s Protocol Buffers (this is not a requirement, especially since it is not installed on CS VM Cluster). You can invent your own, but do specify it clearly in your report.

Make your implementation bandwidth-efficient. You should use UDP (and not TCP) by default. Make sure you use UDP for the detection (and suspicion) traffic. That is the main requirement. For a few key operations (e.g., introduction) you can use either TCP or UDP.

Create logs at each machine, and use MP1 to debug. These logs are important as we will be asking you to grep them at demo time. You can make your logs as verbose as you want them, but at the least you must log: 1) each time a change is made to the local membership list (join, leave or failure) and 2) each time a failure is detected or communicated from one machine to another. We will request to see the log entries at demo time, via the MP1’s querier. Thus, make sure you integrate MP2 with MP1 to make this possible.

We also recommend (but don’t require) writing unit tests for each of the join, leave, and failure functionalities. At the least, ensure that these actually work for a long series of join/leave/fail events.

For the demo, please create commands so that: (1) you can seamlessly switch your system between PingAck and Gossip, with the same membership list (i.e., without needing to stop and restart the system, or re-add nodes, or reboot, etc.), and (2) induce a specified message drop rate (default = 0%) – note that drop rates must be introduced at the receiver process (of the message) and not at the sender.

Machines: We will be using the CS VM Cluster machines. You will start by using all 10 VMs for the demo. The VMs do not have persistent storage, so you are required to use git to manage your code. To access git from the VMs, use the same instructions as MP1.

Demo: Demos are usually scheduled on the Monday right after the MP is due. The demos will be on the CS VM Cluster machines. You must use all 10 VMs for your demo (details will be posted on Piazza closer to the demo date). Please make sure your code runs on the CS VM Cluster machines, especially if you've used your own machines/laptops to do most of your coding. Please make sure that any third party code you use is installable on CS VM Cluster. Further demo details and a sign-up sheet will be made available closer to the date. We expect both partners to contribute equivalent amounts of effort during the entire MP execution (not just in the demo).

Language: Choose your favorite language! We recommend C/C++/Java/Go/Rust. We will release "Best MPs" from the class in these languages only (so you can use them in subsequent MPs).

Report: Write a report of less than 2 pages (12 pt font, typed only - no handwritten reports please!). Write your design (1 page max), especially how you handled the integration of Gossip and suspicion, and PingAck and suspicion. Note that bandwidth is always measured in bytes per second, not messages per second. For each of the following questions measure and draw a plot that contains two lines (or two bar graphs), with standard deviation bars: one line/bar for PingAck and one line/bar for Gossip:

1. Given the worst-case detection time bound (as specified earlier in the document) that applies to both PingAck and Gossip (average detection times within 5% of each other is ok) for 10 VMs. Make sure each line in each graph (i.e., each algorithm in each graph) has at least 5 points per line. Draw plots to compare:
 - a. the background bandwidth between PingAck vs. Gossip (i.e., in the no-failure scenario) as a function of number of VMs in the group;
 - b. false positive rate when there are no failures (you may have to introduce artificial message drops to induce false positives - always introduce artificial message drops on the receiving end not on the sending end), as a function of message drop rate. Define false positive rate clearly (there are various definitions); common definitions are false positives per second, or false positives per every correct failure detected, etc. - no matter what you choose, define it clearly. Make sure you run the experiment long enough (e.g., not just until the first false positive, which can give you misleading numbers).

- c. their detection times as a function of number of simultaneous failures.
2. Disable Suspicion in both Gossip and PingAck variants. Now compare all 4 variants to each other (Gossip, Gossip minus Suspicion, PingAck, PingAck minus Suspicion). Draw the same plots as in #1 above. Make sure you run all 4 variants under the same experimental settings as each other – apples to apples, oranges to oranges comparison (“all else being equal, algorithm X performs better than algorithm Y”). Beware of caching effects, and effects of previous experiment trials that pollute subsequent experiment trials. After the plots 2a, 2b, and 2c, comment on two issues:
 - D. which of the two variants (PingAck, Gossip) benefits the most from Suspicion mechanism, and why?
 - E. how do the 2 “minus Suspicion” variants perform compared to each other, and why?

For each plot, choose at least 5 values on x axis (i.e., at least 5 trials per experiment). Plot averages **and** standard deviations (and, if you can, confidence intervals). It is not typical to include each individual trial, but rather average and std dev across trials (for each setting). Discuss your plots, don’t just put them on paper, i.e., discuss trends briefly, and whether they are what you expect or not (why or why not). (Measurement numbers don’t lie, but we need to make sense of them!) You don’t lose points for unexpected results or surprising data; but make sure you explain intuitively *why* you are seeing algorithms behave a particular way. Stay within page limit – for every line over the page limit you will lose 1 point!

Submission: Similar to past MPs. More instructions on Piazza, but the checklist/workflow is the same. Do not miss the deadlines for signing up for the demos (usually the Friday before submission Sunday), or to submit the form and the Gradescope report (both due on Sunday).

When should I start? Start Early. Start NOW. You already know all the necessary class material to do this MP. Each MP involves a significant amount of planning, design, and implementation/debugging/experimentation work. **Do not** leave all the work for the days before the deadline – **there will be no extensions.**

Evaluation Break-up: Demo [40%], Report (including design and plots) [40%], Code readability and comments [20%].

Academic Integrity: You cannot look at others’ solutions, whether from this year or past years. We will run Moss to check for copying within and outside this class – first offense results in a zero grade on the MP, and second offense results

in an F in the course. There are past examples of students penalized in both those ways, so just don't cheat. You can only discuss the MP spec and lecture concepts with the class students and forum, but not solutions, ideas, or code (if we see you posting code on the forum, that's a zero on the MP).

We recommend you stick with the same group from one MP to the next (this helps keep the VM mapping same on EngrIT's end), except for exceptional circumstances. We expect all group members to contribute about equivalently to the overall effort. If you believe your group members are not, please have "the talk" with them first, give them a second chance. If that doesn't work either, please approach Indy.

**Happy Membership (from us and the fictitious
FakeCompany Inc.)!**