

Homework #1

Due: 6 Feb 2009, 23:59 PST

Homework must be submitted electronically to <csci555@usc.edu>. It should have a subject of "Homework 1." You must submit ASCII text without embedded formatting commands or markup. That means, among other things, no postscript, no Microsoft Word, no PDF, no FrameMaker, no TeX, no groff, no DocBook, no HTML, no XML, and no JavaDoc. Do not submit your homework as an attachment to your e-mail. Do not base 64 encode it. Do not Rot13 encode it. Plain ASCII text. You may PGP sign it, but are not required to do so. Do not PGP encrypt it. If you submit something that is not unmarked-up ASCII, it is functionally the same as turning nothing in.

Homework turned in on the due date is not penalized, one day late is 25% off (that is the grade will be multiplied by 0.75), 2 days late is 50% off, and 3 days late is 75% off. No work will be accepted more than 3 days late. I will generally use the Date: line of the mail, but should the situation merit it, I am not above looking through mail system logs to confirm the submission time. I should not have to mention it, but forging a Date: line to avoid a late deduction is grounds for an F.

Do the work yourself. Computer science is a collaborative science, and I encourage you to talk over the ideas in the homework with other students. *However*, the final submission, that is, *the text of the homework*, must be composed individually by each student. If you hand in homework that is identical to another student, you risk failing the class. (In fact the only way that you would not fail the class in such a circumstance would be if one student had copied another student without the knowledge of the copied student; the copied student would not be penalized.) That is an awfully large risk for 10% of your total grade. Do the work yourself.

As with all work for csci555, this work is subject to the USC code of Student Conduct.¹ Read it, learn it, live it. Should you have any questions on how to apply the code, do not hesitate to contact me or the Office of Student Judicial Affairs and Community Standards.² Should it prove possible, do not plagiarize work from sources outside the class. Plagiarizing homework is grounds for *failing the class*. It is perfectly all right to properly cite external sources, should you find some that are useful.

Answers will not be graded on their beauty of expression. Answers will be graded on whether they show a logical approach and sensible explanation. Short, simple sentences are fine. What is important is that your ideas are clear to the reader, and that they answer the question. Of course, no answer will be penalized because it is beautifully expressed, either.

Each question has equal weight.

Homework

1. In the Mesa paper[Lampson80], we pointed out a case where using Broadcast gives correct results and Notify gives incorrect ones. It is easy to extrapolate that Broadcast is always the best primitive to use, but that is not the case. Construct an example where Notify is both correct and more efficient than Broadcast and explain why.
2. In Section 3.5 of Birrell and Nelson's RPC system[Birrell84] they describe a system for dispatching RPC packets to server processes that includes process identifiers provided by the participants. What is the purpose of these identifiers (in your own words)? How important is it that they are completely accurate, and why? In what way are they like Mesa Notify commands?

¹ <http://web-app.usc.edu/scampus/university-student-conduct-code/>

² <http://www.usc.edu/student-affairs/SJACS/>

3. If one is going to implement a lock in Linda[Carriero85], one must do it using the `in` command. In CSP[Hoare78] one has the choice of using either the send message or receive message primitive. Explain why.

Here are the outlines of two implementations of a lock using CSP constructs. In both cases each process that may use the lock is identified by a positive integer. The process representing the lock is called the lock manager.

In the first implementation, the lock manager is implemented very much along the lines of the semaphore discussed in class, except that the `val` variable is only allowed to be zero or one. This is accomplished by adding a term to the guard that accepts `V()` messages that only accepts them when `val` is 0. Processes acquire the lock by sending a `P()` message to the manager. When that message has been received, the process holds the lock. It returns the lock by sending a `V()` message to the manager.

The second implementation also represents the lock as a process. The lock manager uses a two-phase protocol to assign the lock. When a process requests a lock it sends a `request_lock()` message to the manager. The requester then waits for a `lock_granted()` message from the manager. The requester holds the lock when it receives that message. To release the lock the process sends a `release_lock()` message to the manager, and the lock is released when that message is delivered. The manager keeps a variable with the identifier of the process holding the lock (or -1) in it. It loops always accepting `request_lock()` or `release_lock()` messages, and enqueueing the identifiers making the requests. When the lock is available (i.e., is -1) the manager sends a `lock_granted()` message to a process with a pending request. When the lock becomes free (on receipt of a `release_lock()` message) the manager assigns the lock to another process, if a request is waiting.

In both cases I am omitting details of how the manager enforces the requirement that only the lock holder can release the lock, but that is straightforward. You may want to think it through, however.

Explain an advantage of using the local queue in the second implementation. Describe the costs of the second implementation.

4. JINI[Waldo99] adopts a decentralized layout for both their service location networks and for their authentication (each component decides which keys it trusts to sign code). Athena[Champine90] adopts a centralized model for these choices (Hesiod locates services from a single configuration database, and Kerberos provides central authentication).

Explain the problems that each system was trying to solve that led to that part of the design, and how their design addresses those problems.

References

- [Lampson80] B.W. Lampson and D.D. Ridell, "Experiences with Processors and Monitors in Mesa," *Communications of the ACM*, vol. 23, no. 2, February 1980, 105-117,
- [Birrell84] A. Birrell and B. Nelson, "Implementing Remote Procedure Calls," *ACM Transactions on Computer Systems*, ACM, vol. 2, no. 1, February 1984, 39-59,
- [Carriero85] Nicholas Carriero and David Gelernter, "The S/Net's Linda Kernel," *ACM Transactions on Computer Systems*, ACM Press, vol. 4, no. 2, 1986, 110-129,
- [Hoare78] C. A. Hoare, "Communicating Sequential Processes," *Communications of the ACM*, vol. 21, no. 8, August 1978, 666-677,

[Waldo99] Jim Waldo, "The Jini Architecture for Network-centric Computing," *Communications of the ACM*, ACM, vol. 42, no. 7, July 1999, 76-82,

[Champine90] George A. Champine, Daniel E. Geer, Jr., and William N. Ruh, "Project Athena as a Distributed Computer System," *IEEE Computer*, IEEE, vol. 23, September 1990, 40-50,