

CMPSC 441
Distributed Systems
Spring 2016

Laboratory Assignment Six: Exploring the Trade-offs of Multicast Communication

Introduction

There are certain situations in which it may not be best to create a client-server-based distributed system that uses unicasting for the exchange of messages. For instance, the need for an alternative to unicasting can arise when a single node in the distributed system must quickly communicate a message to many nodes. Unicasting may also be an unsuitable solution when a client is attempting to locate a server that migrates around a local-area network. This laboratory assignment will explore multicasting as an alternative to unicasting, giving students the opportunity to investigate the trade-offs of this alternative to communicating in a “point-to-point” fashion. In addition to running a multicast system supporting communication between several nodes, you will write a report, using either Markdown or L^AT_EX, examining the trade-offs and applications of multicasting.

Review Your Textbook

Before starting this laboratory assignment, you should review the content in Chapters 1 through 3 of your textbook. If you have a question about this content, then please resolve it before starting this laboratory assignment. Next, you should examine the content in Section 4.1 that discusses the protocols commonly used in distributed systems, paying particularly close attention to the material about connection-oriented and connectionless protocols. It is also important for you to study the content in Section 4.5 that talks about the various approaches to multicast communication. Students who are interested in studying a more advanced treatment of ways to achieve reliable multicasting are encouraged to read Section 8.4. Please see the instructor if you have any questions.

Exploring Multicast Communication

Since you configured your connection to the “share” repository for this course in a previous laboratory assignment, you should be able to change into the `cs441S2016-share/` directory and type the “`git pull`” command to gain access to the two Java classes needed to complete this assignment.

Now, please find the `labs/lab6/` directory in our course’s repository so that you can study and understand the source code for the files called `MulticastSocketReceiver.java` and `MulticastSocketSender.java`. What are the key aspects of how these two Java classes support multicast communication? For instance, you will notice that these Java classes specify a “port” — and, yet, do not give a “host” as you saw in the previous assignments. To learn more about why this is the case and to better understand the trade-offs associated with this approach, you should read the Java documentation for the `java.net.DatagramPacket`, `java.net.InetAddress`, and `java.net.MulticastSocket` classes. In particular, you should investigate what it means for the `DatagramPacket` class to, according to its documentation, “implement a connectionless packet delivery service”. Finally, can you draw a technical diagram that illustrates the general interaction between a client and a server, customized for these two classes that perform multicast communication with sockets? Please see the course instructor if you have questions about these issues.

Once you have carefully studied the two provided Java classes, you should attempt to compile and run them in your terminal window. Please note that this system will only work correctly if you first run the `MulticastSocketReceiver` and you then execute the `MulticastSocketSender`. After running these two programs, you should record the output that they produce. Next, please try to run one, two, three, four, and five separate versions of the `MulticastSocketReceiver` on different computers in the Alden Hall network. You can accomplish this task by using `ssh` to connect to a remote computer in a separate terminal window and then running the `MulticastSocketReceiver` from this location. Next, please find the location in the source code that declares the `String INET_ADDR` variable. What is the purpose of this variable? Is it possible for this program to work if you change this address to a different value? What are some values that seem to work correctly?

There are many real-world applications of multicasting in a distributed system. For instance, the transmission of stock quotes to many listening computers is a natural application for multicasting. Using the ACM Digital Library, please locate and download the article “An Internet multicast system for the stock market” by N. F. Maxemchuk and D. H. Shur. How is the approach to multicasting that this paper presents different than the one that you have used in this laboratory assignment? Would it be a good idea to use the code from this laboratory assignment in a multicasting system for the stock market? Why or why not? In addition to reading and reporting on this paper, you should also download and read one more paper from the ACM Digital Library. Students should focus on finding an additional article that presents either a real-world development and application of multicasting or an empirical characterization of how multicasting is used in a real-world setting. Please see the instructor if you have questions about this part of the assignment.

Summary of the Required Deliverables

This assignment invites you to submit printed and signed versions of the following deliverables. Additionally, all of these deliverables must be in a `cs441S2016-<your user name>` repository that you created for this class; please make sure that you share this repository with the instructor.

1. The well-commented source code of the Java class called `MulticastSocketSender`.
2. The well-commented source code of the Java class called `MulticastSocketReceiver`.
3. Using both text and diagrams, a description of client-server communication with multicasting.
4. A report that explicitly compares and contrasts multicast and unicast remote communication.
5. A document that summarizes the different types of multicasting systems used in production.
6. A paper that responds to the other questions that this assignment poses about multicasting.
7. A reflection on the challenges that you encountered when completing this assignment.

In adherence to the Honor Code, students should complete this assignment on an individual basis. While it is appropriate for students in this class to have high-level conversations about the assignment, it is necessary to distinguish carefully between the student who discusses the principles underlying a problem with others and the student who produces assignments that are identical to, or merely variations on, someone else’s work. With the exception of the provided source code, deliverables that are otherwise nearly identical to the work of others will be taken as evidence of violating the Honor Code. This means that, for instance, all of the other comments, source code, data, and written reports should be the original work of the student completing this assignment.