

CSCI-1380: Distributed Computer Systems

Homework #1

Assigned: 02/02/2021

Due: 02/08/2021

1 System Design

In general distributed systems are about understanding and picking appropriate design choices for your scenarios. Ultimately, your knowledge of novel design choices and interesting combination will be based on internal documents or publicly released descriptions of existing techniques.

HW#1 and HW#5 are both related HWs! Both of these assignments are not technical in nature. They are meant to provide you with a mechanism for reading and understanding internal design docs and external system descriptions which you may have to deal with as a new member of a distributed systems team.

As a new member, many terms will be foreign to you. However, fundamentally most distributed system design documents focus on a few properties: performance, availability, correctness, and security. They explore different techniques to try and realize these concepts. As you read the document and description you should be able to start mapping these techniques to the concepts based on the context surrounding the design. For example, if you read section 5.1 from the Cassandra paper you may be unaware of what “consistent hashing” or “rings” are but from the context and descriptions, you can get a high level of their of their purposes.

For both assignments, you will read the original 6-page Cassandra paper from Facebook. Cassandra is a system that identifies and leverages different design choices that we will cover in this class.

1. For HW#1 your goal is to identify concepts or words that you are unfamiliar with – and try to guess what they may mean based on how they are used.
2. For HW#5, you will revisit this paper after we are over 70% through the class and see if there are specific concepts or ideas that are still foreign to you.

What to do: Read sections 1,3,5 from the original Cassandra paper. You can find the paper here: [Cassandra](#).

What to Turn in: A list of terms that you are unfamiliar with.

1. For each term, try to determine what properties the term provides (1) performance, (2) reliability (i.e., failure), (3) correctness, (4) other (or unknown).
2. For each term, try to guess the definition (or purpose) of the term based on the context in which it is defined.

Note: Given the varying levels of expertise, we do not expect everyone to have the same list of terms. Moreover, given some of the private questions at least one or two of you are familiar with material in the first third of the class.