

Lab #5: Erlang

Overview

The purpose of this lab is to introduce process spawning and message passing in Erlang.

Step 1

Create an Erlang program that creates a ring of processes and sends messages through this ring. This program should take three arguments: the number of processes in the ring, the number of times the message should be sent, and the message to send. Note that this “program” will consist of a function that takes these initial parameters and any number of helper functions.

Step 2

Instrument your program with code to measure and report the amount of time it takes to send all of the messages through the ring. Do not include the time that it takes to spawn or kill the processes (though you may report these separately if you’d like).

To measure time, you can use the Built-in Function `erlang:statistics/1`. The atoms `runtime` and `wall_clock` are arguments of possible interest.

Step 3

Experiment with the numbers of processes and messages (e.g., spawn 100 processes and send 10 messages; then increment the number of processes and messages independently). How many processes can you reasonably spawn? How does the time grow with the number of processes/messages?