

Hw3 Experiment Write-Up

SJSU Students

Problem: Using the history command (h) conduct experiments to see how quick the routes converge for different nodes in your network of a given total number of nodes. Conduct further experiments to see the effect of shutting down a node on how long it takes for the network to re-stabilize.

Hypothesis:

We hypothesize that as a lined network becomes increasingly populated with routers, the time for those routers to converge to the same routing table will also increase. This is because there are more routers in the graph that need to be accounted for when sending packets and calculating path hops. We also expect to see some variation for multiple trials on the same network since the sender thread sleeps randomly for 3 to 4 seconds.

As for when an end node on the network is shutdown, we hypothesize that the time for the network to converge and restabilize will take longer than that of normal, previous run, since the time to live for each packet is depended on when a node shuts down. Thus, the TTL of each packet will indirectly increase the time for the network to restabilize after the shutdown. We expect to see similar variation as in the previous run.

Procedure:

Used a python function to generate topology files of routers in a line for three different sizes: 5, 10, and 20.

```
def generate_topology_file(base_port, num):  
    num -= 1  
    print(f'{base_port} {base_port+1}-1')  
    for i in range(1, num):  
        print(f'{base_port+i} {base_port+i-1}-1 {base_port+i+1}-1')  
    print(f'{base_port+num} {base_port+num-1}-1')
```

```
generate_topology_file(11370, 5)
```

For the startup experiment we first started a simulation with one of the topology files then would repeatedly use the 'h' command until the tables contained all of the correct information. We would then record the time.

For the shutdown experiment we again started a simulation with one of the topology files and then wait for the routing tables to converge. We then stop a router on one end of the line of routers using the 's' command. Then we wait for the router on the other side of the line to properly update its routing table. We recorded the difference in times between the routing table being correct and the router being stopped.

Tim to converge on startup:

Num Routers	Trial 1 Seconds	Trial 2 Seconds	Trial 3 Seconds
5	7s	10s	7s
10	24s	13s	13s
20	33s	36s	39s

Analysis

Above is data collected from performing 3 trials for each of the three topology files we generated using the program above for a lined network, without stopping a router. The times are recorded when the network has fully converged routes. The main reason why there is variation in each run is due to the random time (between 3 to 4 seconds) that the sender thread sleeps. This might cause the routers to receive packets in a less or more optimal order.

Time to converge after shut down router (we chose to always shutdown a router on the edge):

Num Routers	Trial 1 Seconds	Trial 2 Seconds	Trial 3 Seconds
5	42s	50s	49s
10	50s	52s	49s
20	68s	75s	71s

Analysis

The same method is applied as the first experiment, except this run includes stopping a router on the edge of the lined network, and waiting for the one on the opposite side to converge to the correct routing table without the stopped router. The time to converge after a node shuts down is dependent on the TTL that is set for each packet. In our simulation we have the starting TTL set to 10 seconds. This can be adjusted by changing the START_TTL variable.

Conclusion:

In general, our hypothesis was confirmed in that we observed the time to converge increasing when the network became increasingly populated with routers. For the multiple trials of the different topologies, there was minimal variation in the converge times, and an almost linear progression between the number of routers and times for the normal network. As for the node shutdown, our hypothesis was also confirmed in that the time for the routers to recognize there is a missing node and update their routing tables will take significantly longer than a normal run. We observed that the node shutdown times were noticeably longer than the normal network

times, most likely due to the packets' TTL times when a node shuts down. In conclusion, our hypothesis was mostly supported by the experiments. More routers meant longer time to converge both initially and after a shutdown.