# Cluster Computing and Raspberry Pis

William Harrell
1/23/2020
LUG @ NC State

# History

- **First “supercomputers” were just regular computers with beefed up components**
  - First use of “supercomputer” was for the CDC 6600, for outperforming most other computers tenfold
- **Cray-2, 1985 was the first parallel supercomputer**
  - 8 liquid cooled CPUs
  - 1.9 gigaflops – about the power of a 2nd gen iPad
  - Used for nuke, sonar, and vehicle crash simulations
  - Had problems reaching peak performance, unlike its single CPU predecessor, but paved the way for parallel computing

# Cray-2

# Supercomputers today

- **Though high performance components are still used, quantity has won out over quality**
- **Currently the most powerful one in the world is Summit at Oak Ridge**
  - 200.795 petaflops – 75 more than 2nd place, about the power of a 3rd gen iPad
  - Studies cosmology, biology, climatology, fusion, AI
  - 3rd most energy efficient in the world
  - 4,608 nodes – 202,752 CPU cores, 27,648 GPUs, 6.451PB RAM
  - Each node uses 2x 22-core Power CPUs and 6x NVIDIA Tesla GPUs, 1.4TB of RAM
  - Built by IBM, runs on RHEL

# Summit Supercomputer

# Distributed Supercomputing

- **Distributing a program across multiple networks/computers**
- **Servers hand tasks to nodes, which report the results back to the servers when the task is complete**
  - Folding@home
  - BOINC

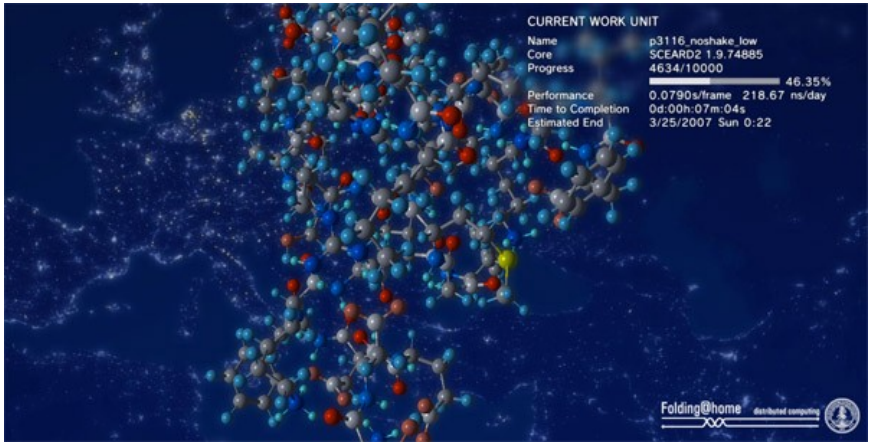


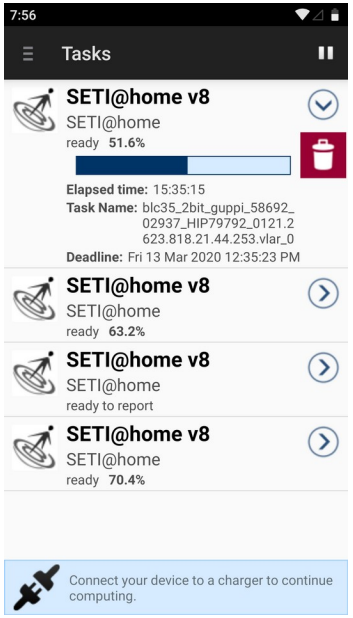

# Building supercomputers out of things you really shouldn't

- **PlayStations**
  - Sony released a "Linux for the PS2" kit
  - A supercomputing organization hooked 60-70 together with "okay" results
  - The PS3 also supported Linux for a time
  - Researchers, including the Air Force, began building PS3 clusters due to their high cost efficiency
  - Condor Cluster (right)
    - 1,700 PS3s
    - Was 35th most powerful supercomputer in world

"Every time you ran this thing, it would cause the kernel on whatever machine you ran it on to kind of go into this weird unstable state and it would have to be rebooted, which was a bummer," Steffen said.

# Raspberry Pis

- **Not very powerful, but very cheap**
  - Quad-core 1.5GHz 64-bit ARM CPU
  - 1GB, 2GB, or 4GB RAM
  - Gigabit Ethernet
  - $35, $45, $55
- **Los Alamos uses them for designing and testing supercomputer architectures**

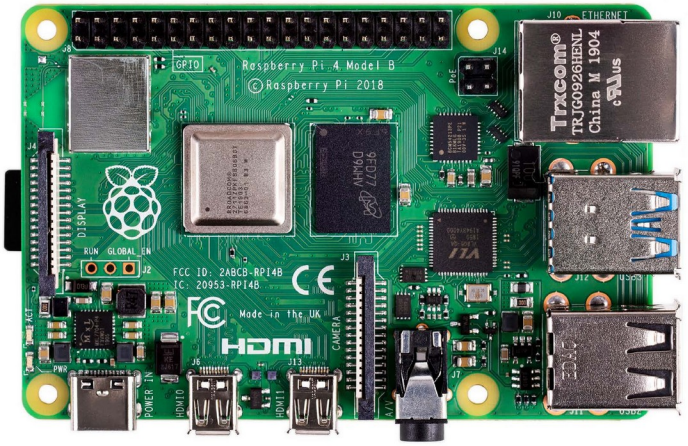


# Act 1: Cracking Really Weak RSA Keys

- **Using a cringeworthy RSA implementation:**
  - Generate a keypair from two primes, p and q
  - To crack a message, take the encrypted message and a known substring
  - Try p and q values until we run out or a pair decrypts the message
  - To distribute, split the list of possible p values into separate lists, and give one to each node

# Message Passing Interface (MPI)

- **Networking standard intended for distributed programming**
- **Program starts on one computer (primary), which calls upon other computers in cluster (secondaries) to help out**
- **Data can be sent one-on-one, broadcast, synchronously, asynchronously**
- **Usually used in C, but we're using mpi4py (Python)**

# Diagram

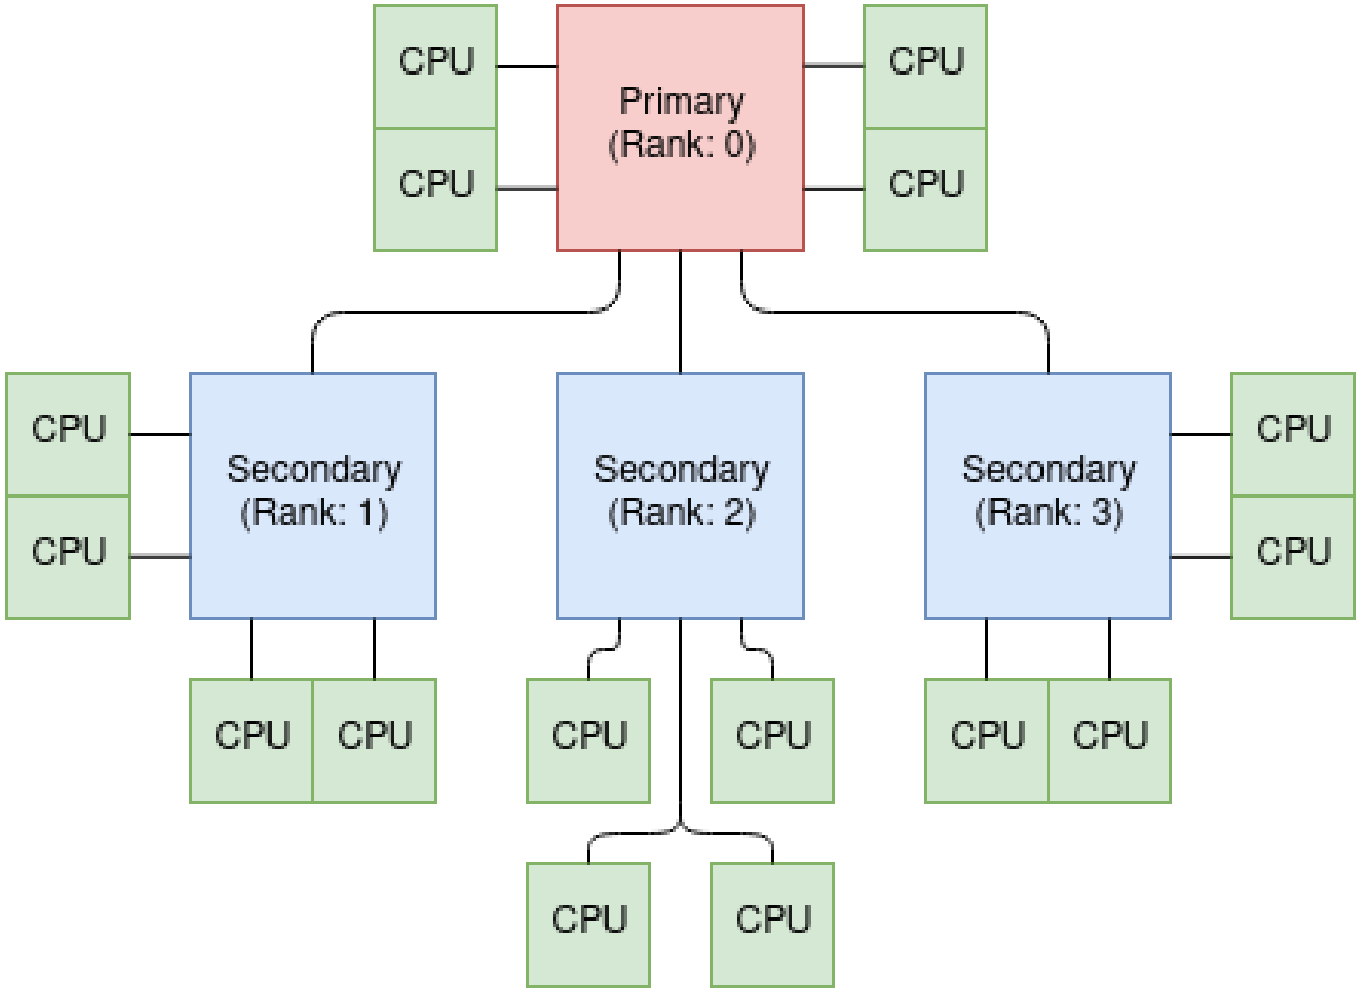
CPU
CPU
Primary
(Rank: 0)
CPU
CPU
CPU
CPU
Secondary
(Rank: 1)
CPU
CPU
CPU
CPU
Secondary
(Rank: 2)
CPU
CPU
CPU
CPU
Secondary
(Rank: 3)
CPU
CPU
CPU
CPU

# My Epic Homelab

# Part 2: distcc cluster

- Distcc is a distributed C/C++ compiler
- Client PC sends source and library files to hosts to help with compilation
- Typically, this results in much faster compilation
- Although it is not uncommon to use distcc with Raspberry Pis, they are usually not the ones on the host side (though I am not the first to try)

# Preparation

- **Hosts and clients need same gcc and binutils versions, if you actually wanted to use this**
  - This can be done with a cross-toolkit generator, crosstool-ng
- **We don’t need to worry about this since we’re just benchmarking**
- **For a test job, we’ll use coreutils since it still takes awhile to compile but give similar results whether we do compile for AMD64 or ARM**
- **Can 16 ARM cores beat 4 AMD64 cores? (Intel i5-4200U, 2.60GHz) 8 AMD64 cores? (AMD FX-8350)**

# Nope!

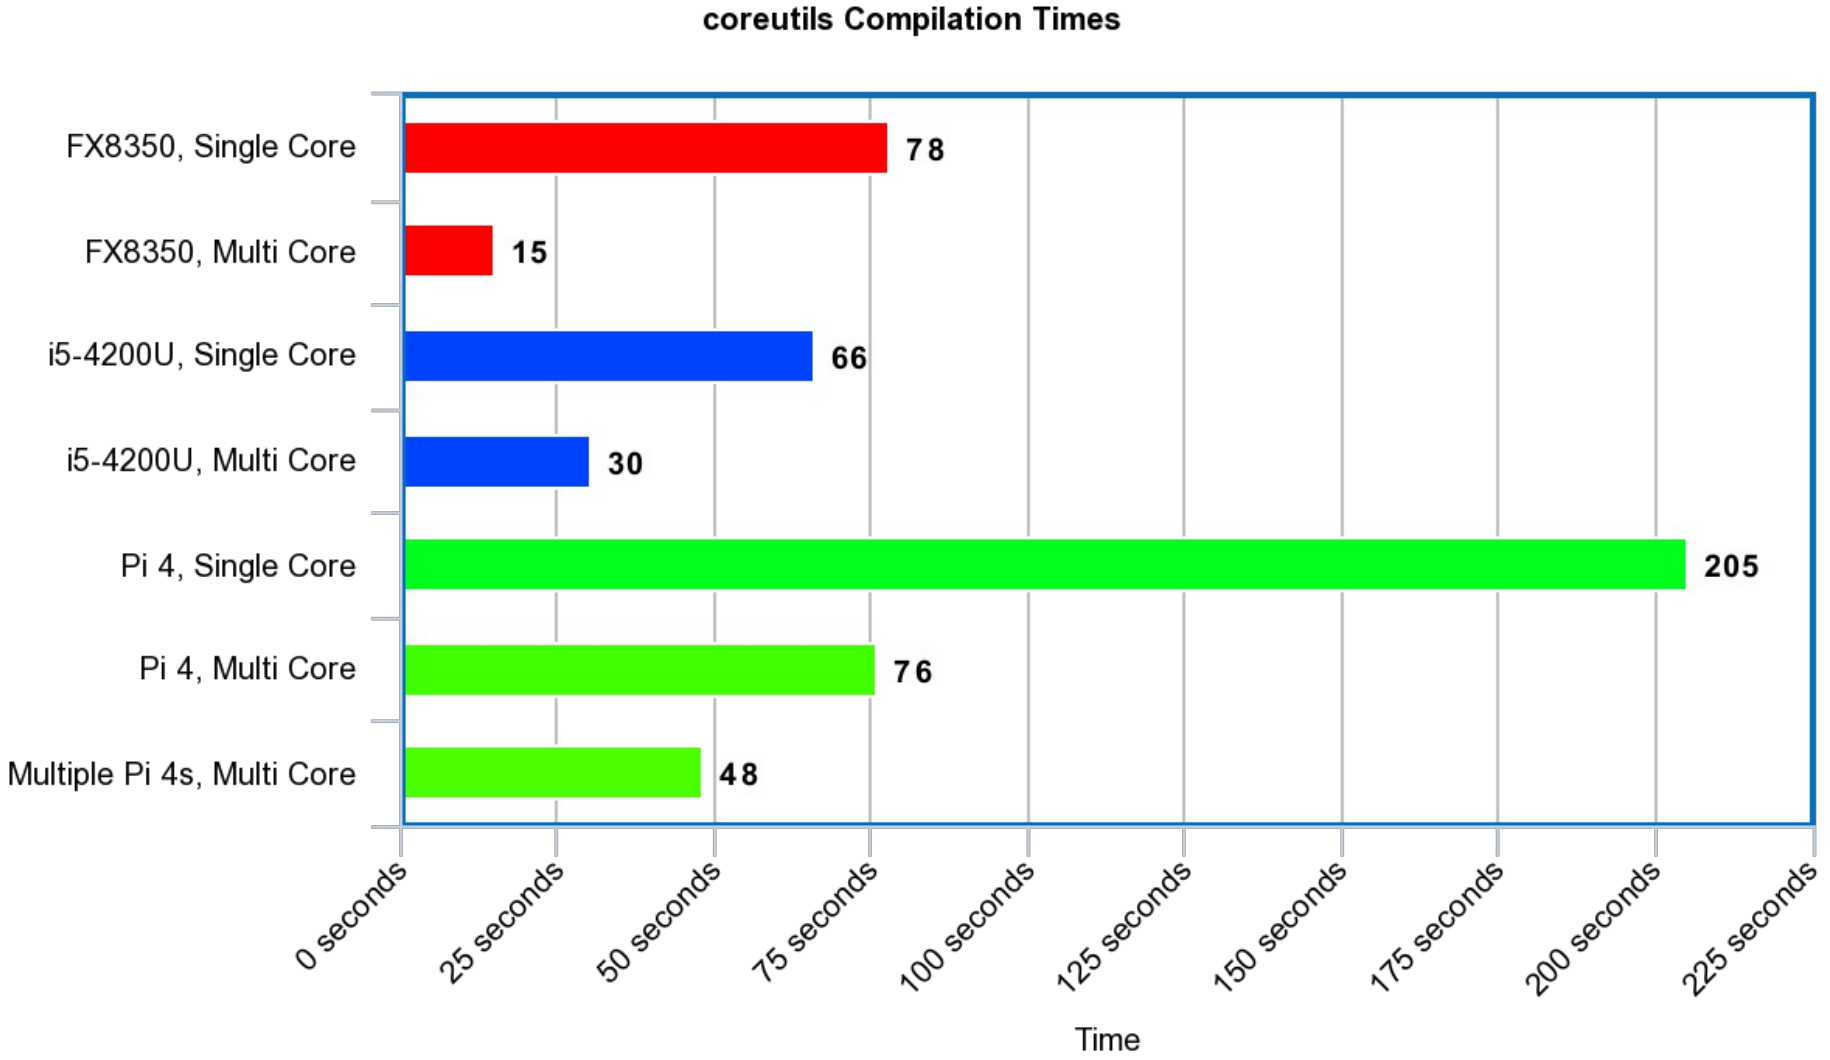
coreutils Compilation Times
FX8350, Single Core
78
FX8350, Multi Core
15
i5-4200U, Single Core
66
i5-4200U, Multi Core
30
Pi 4, Single Core
205
Pi 4, Multi Core
76
Multiple Pi 4s, Multi Core
48
0 seconds
25 seconds
50 seconds
75 seconds
100 seconds
125 seconds
150 seconds
175 seconds
200 seconds
225 seconds
Time