ABCD-LMA, August 16, 2018

"Why Do Medical Students Need a 10,000-Core Cluster?"

Questions about HMS Research Computing

Amir Karger
Research Computing Group
Harvard Medical School Information Technology

Whom Does Research Computing Support?

- 700 MD students
  - Few do research, so who needs RC?
- 800 PhD students doing basic research
- 1500 post-docs "on Quad"
- Thousands more:
  - HSDM, HSPH
  - 16 Affiliate hospitals
  - Collaborators around the world

What Does RC Support?

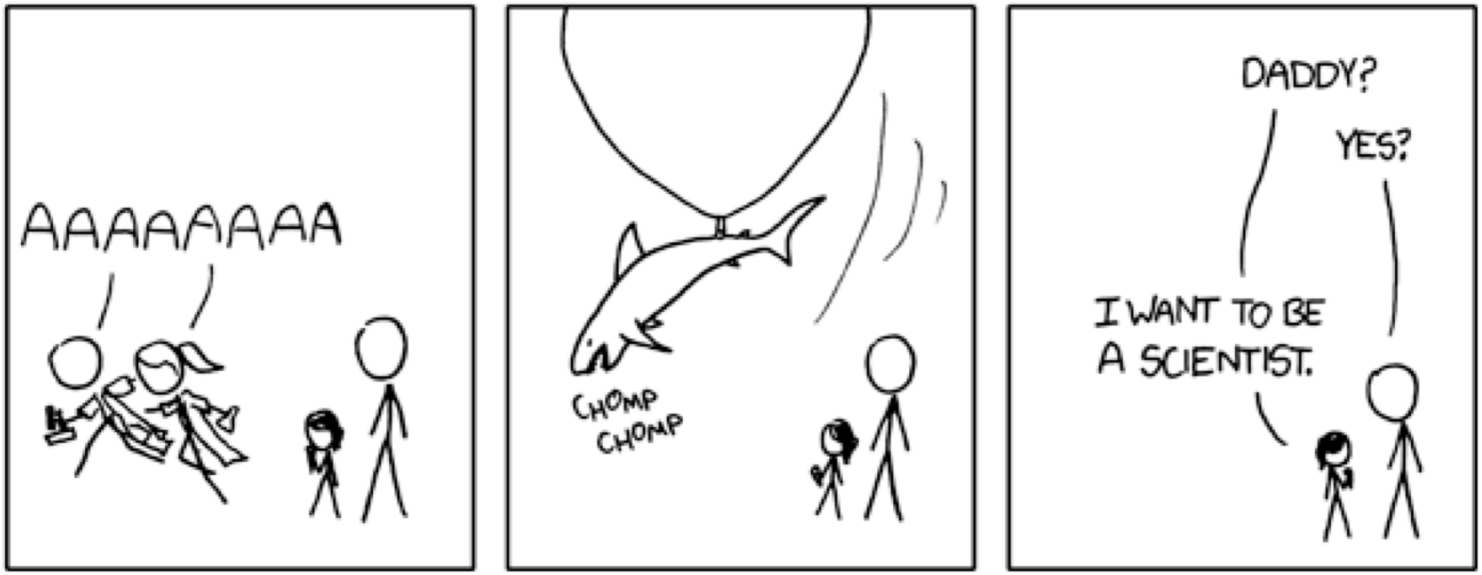

XKCD #585

Science!

What Does RC Support II?

- Basic research: bacteria, worms, mice
- Football Players' Health Study (HIPAA!)
- Tracing drug-resistance in tuberculosis
  - Farhat Lab, GHSM + DBMI
- Opiods: duration more serious than dose
  - Kohane Lab, Biomedical Informatics
- Mapping neurons flies use to smell
  - Lee Lab, Neurobiology
- Pretty pictures at the end

Who Are You?

DevOps (5 people)

- HPC cluster administration
- Automating sysadmin (“big red lever”)
- Build Level 4 systems
- OMERO
- SQL and web hosting

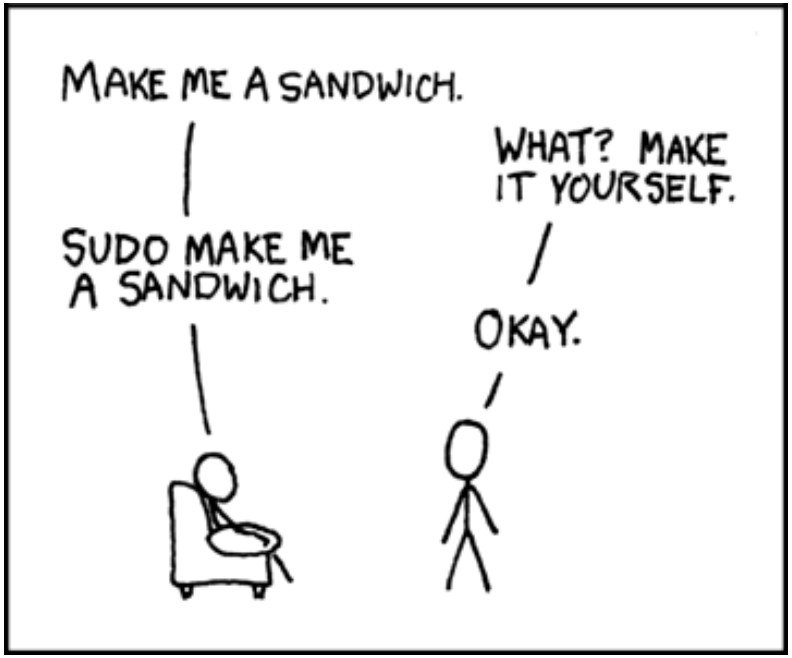

Who Are You II?

RCCs – Research Computing Consultants (6)

- Science degrees
- Programming experience
- Knowledge of HMS computing environment

What Are You Doing?!

- Consulting! – Listen, think, advise
- Science – better, cheaper, faster
  - Use the right resources, programs, statistics
  - Run workflows in the most efficient way
  - Scale up workflows
- Most services are free (!)
  - We (almost) never analyze data for users
  - We show users how to do it
  - Or send to other groups (like HCBC)

How Can We Empower Researchers?

- Desktop/workstation -> O2
- Single gene analysis -> thousands
- Parallelization, GPU
- R, Perl, Python, bash, programming
- SQL Database hosting
- Hosting research websites
- Data ~~manipulation~~ ~~hacking~~ ~~wrangling~~ science
- Do bioinformatics, biostatistics **correctly**

How Can RC Scale?

- Showing is more efficient than doing
- Hands-on classes lasting about 2 hours
  - Intro to Linux & O2
  - Introduction to NGS Analysis
  - Perl, Python
  - R/Biostatistics and Bioconductor
  - Parallel Computing
  - Matlab
  - git and Github
  - Ad-hoc classes, bigger rooms as needed

What is Your Quest?

Orchestra: entered production 2004

- 5000 cores 7+ years old, 12-core nodes, 96 GB RAM
  - (+2000 faster cores, 256 GB RAM)
- Old LSF scheduler
  - $2M upgrade + ongoing support needed
  - No help in Google
  - Lacking features – priority bumps difficult
- 12 years of clutter
  - 250 queues, old software versions, config…

What is Your Quest II?

O2 went into production 9/2017

- 5000 new cores (+ 2000 stolen from Orchestra)
  - Fast-ish processors, 32-core nodes, 256 GB RAM
- Slurm scheduler:
  - Open, popular
  - Free, with cheap support
  - Feature-ful – more knobs and dials
- 12 partitions
- New OS/software, puppet, …

What's the Difference?

- New scheduler commands
  - `bsub -q short -W 1:0 myprog in.txt` ->
    `sbatch -p short -t 1:0:0 myprog.batch in.txt`
  - `bjobs -u ak150` -> `squeue -u ak150`
- New paths for some filesystems
- New module names & versions
- Some new libraries
- Users need to change commands, scripts, …
  - How hard could it be to convince them?

What's the Worst that Could Happen?

- RC too distracted to build the new system
  - Maintaining 2 systems and migrating users
- Labs refusing to leave Orchestra
  - Can't migrate crucial workflow
- Labs delaying leaving Orchestra
  - "Orchestra works just fine"
- Labs insisting we bring Orchestra back
  - Didn't know about move, didn't test

What Do You Expect?

- Complaining or oblivious users
  - Carrots only encourage 75% of users
  - When do we employ sticks?
- Details missed by busy RC staff
- Timeline:
  - Expectation: 1 year
  - Fear: 18 months? More?
  - Reality: ???

How Do You Get There From Here?

- Minimize differences to lower barriers
  - Same partition names, limits, module "API"
- Notify users early and often
- Document changes fully and concisely
- Run migration workshops
- Hire term employee focusing on migration
- Use every ticket as an opportunity (to pester)
- Move fast nodes to O2

Are We There Yet?

- "All carrots, (almost) no sticks" … worked!
- ~1000 users migrated in **7 months**
  - Not including reeeeeally long beta
- Labs refusing to leave Orchestra: 0
- Orchestra turned off 4/2018 (5000 days)
- Labs asking to bring Orchestra back: 0
- Um, some websites kept running on Orchestra
  - Minimal usage, many nodes turned off

Did You Get Very Far?

- O2 is our flagship service
  - Many types of life science research
  - Batch and interactive work (even GUI!)
- Lots of usage
  - ~7000 faster cores, 32 cores per node, 256 GB
  - GPU & high memory machines
  - 200 users, 70k jobs per day (30% > Orchestra)
  - 3000 cores, more GPU+highmem coming soon
- Efficient usage is still a struggle

What Supports RC?

Storage

- 1 PB Scratch disk on O2
- 1 PB "no_backup" (growing)
- 21/27 PB Tier 1
- 1.5/8 PB Tier 2 (growing a lot)
  - Cloud storage with local "stubs"
- Tier 3 to come
  - Glacier et al.

What Supports RC II?

Network

- 1 Gb for end user devices
  - 10 Gb for some microscopes, etc.
- 4 x 10 Gb to building switch, 2 x 100 to core
- Multiple 100 Gb to/between data centers (Markley, Lowell)
- 2 x 10 Gb (only!) to LMAnet -> Intertubes
- 30 TB/day to/from LMAnet, 100 MDC

Who Supports RC?

- Infrastructure Ops (IOps) – network/storage
- Client Services Group – desktop support
- Chan Bioinformatics Core – analysis, training
- RITS, SBGrid, WARP, etc. – specialized support

- The Dean – "hard money" from IT

- And now … pretty picture time!

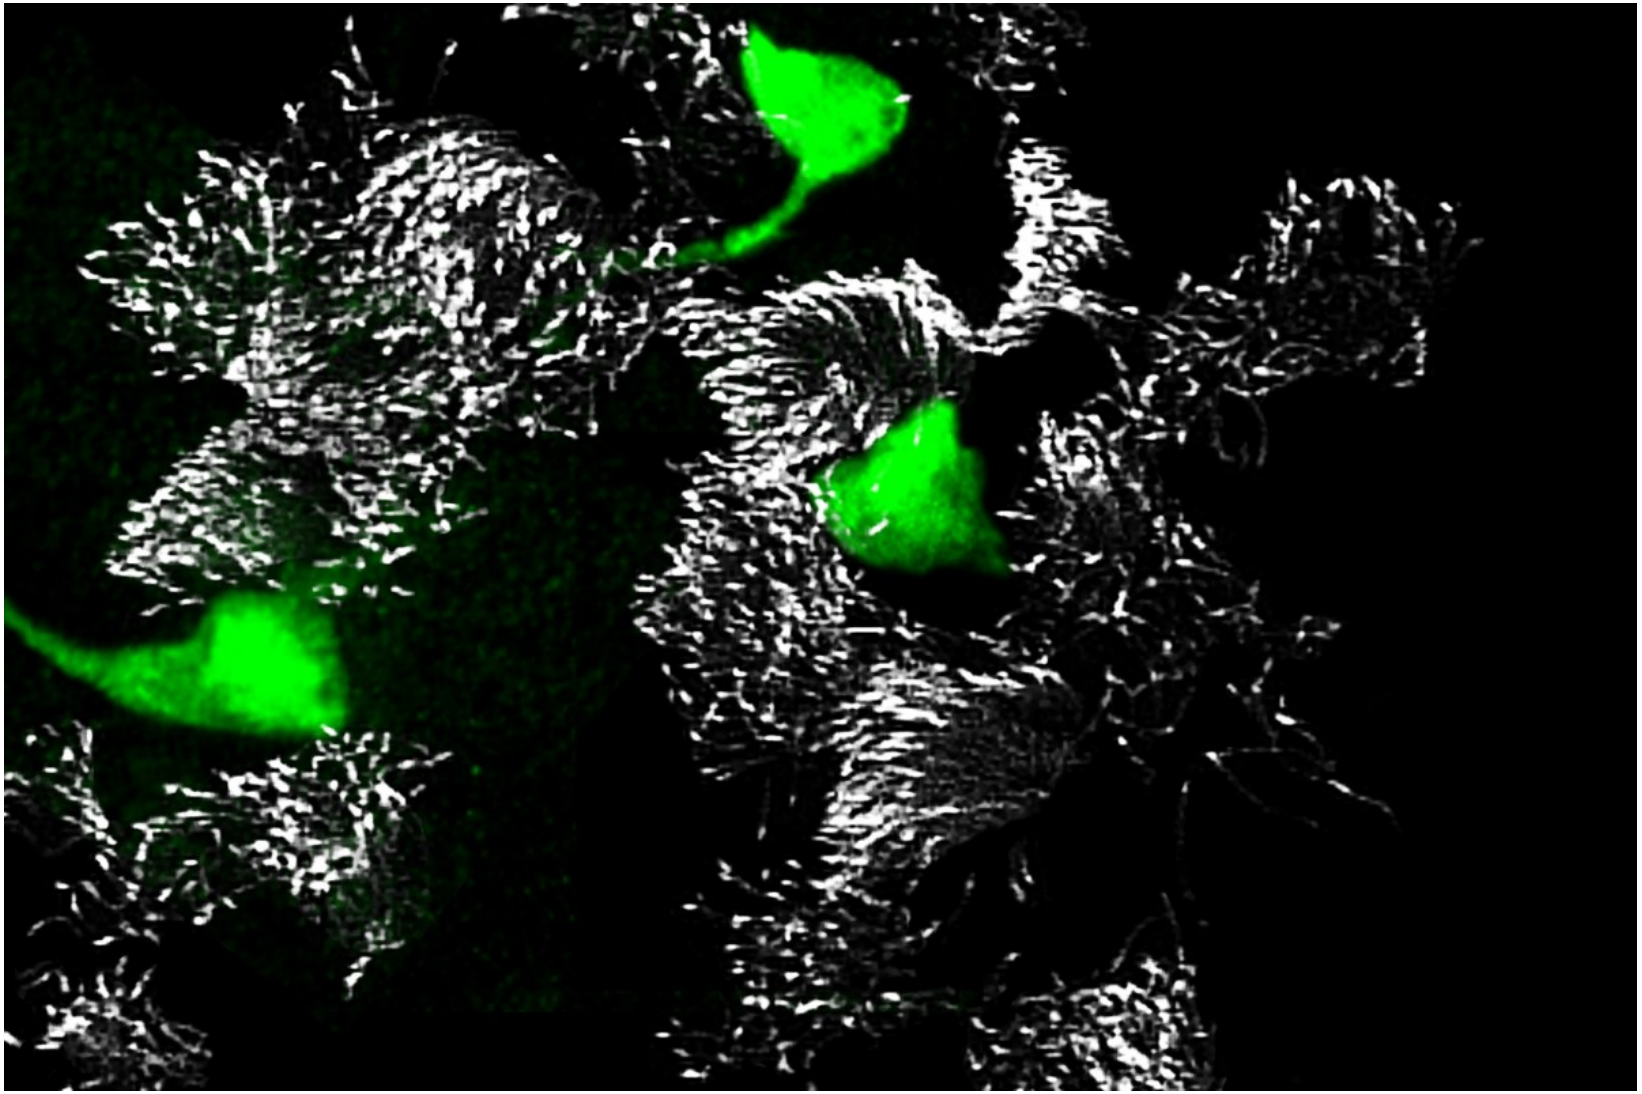

Klein Lab, Systems Biology (Nature, 8/2018)

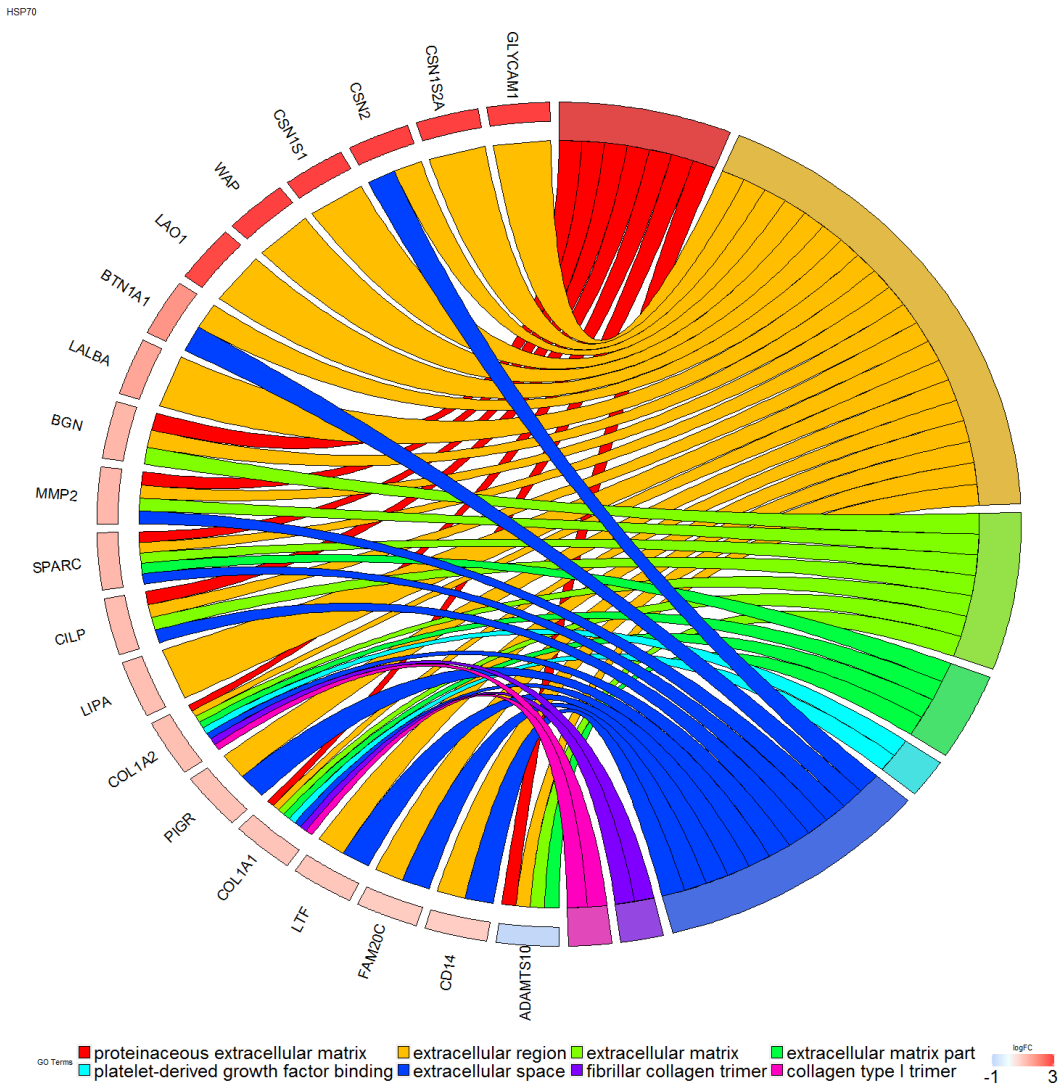

Calderwood Lab, BIDMC (In press)

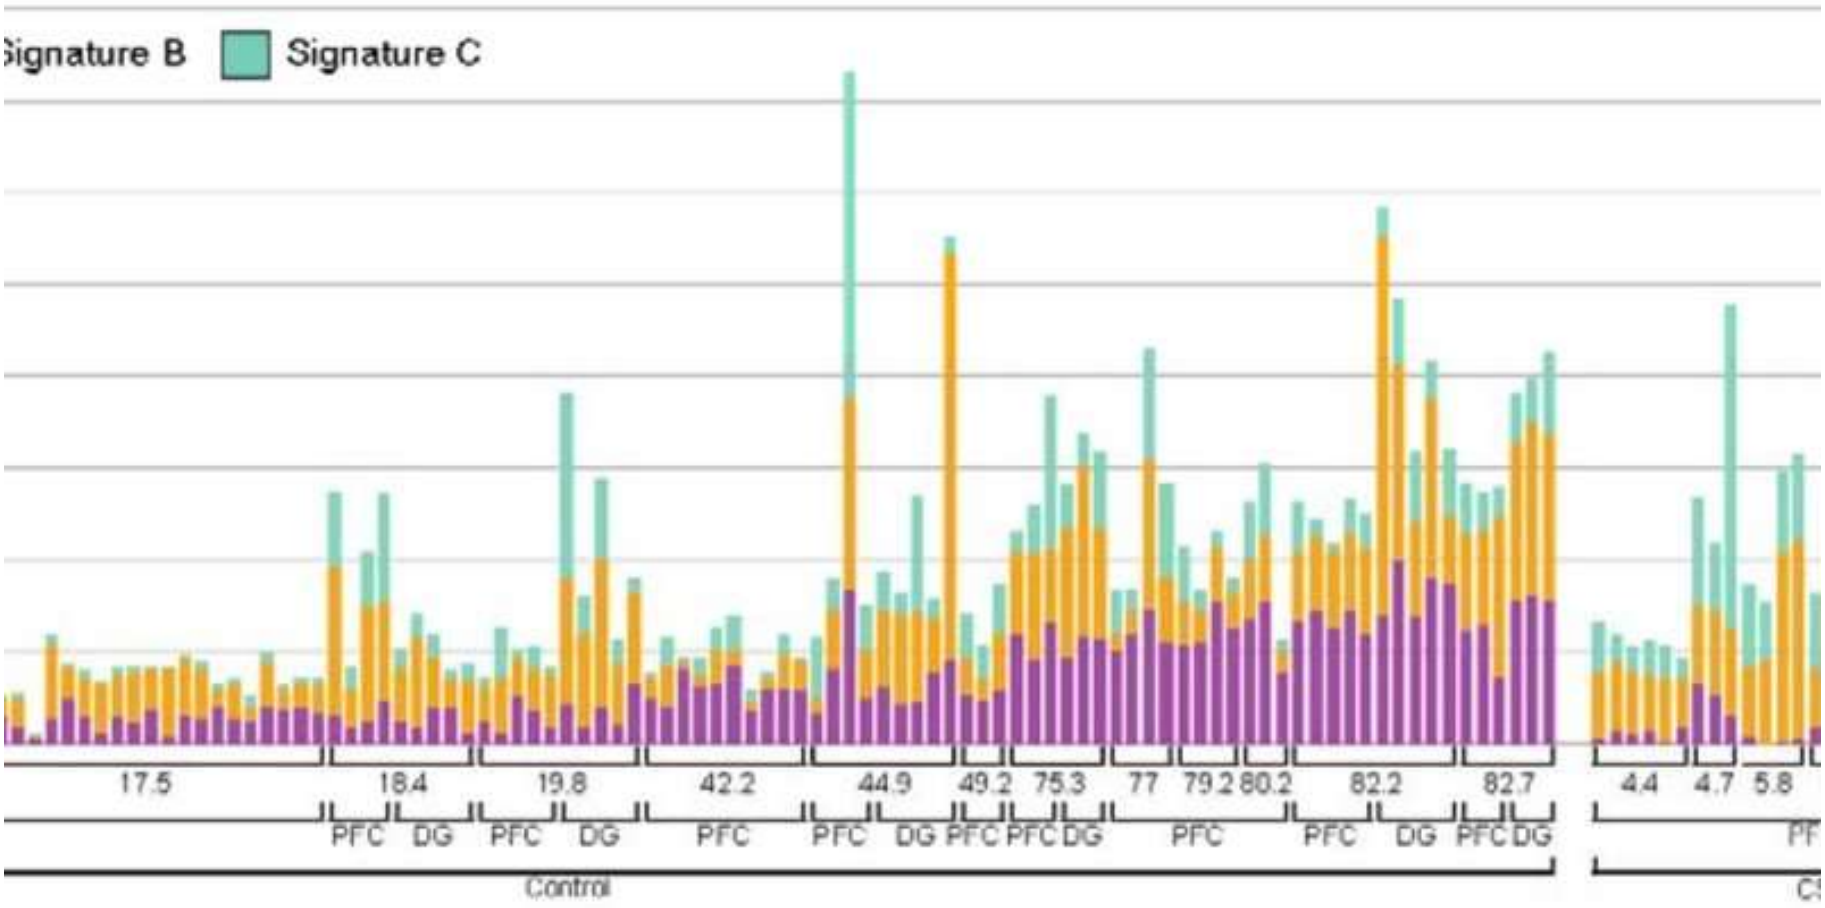

Walsh Lab, Childrens Hospital (Science, 12/2017)

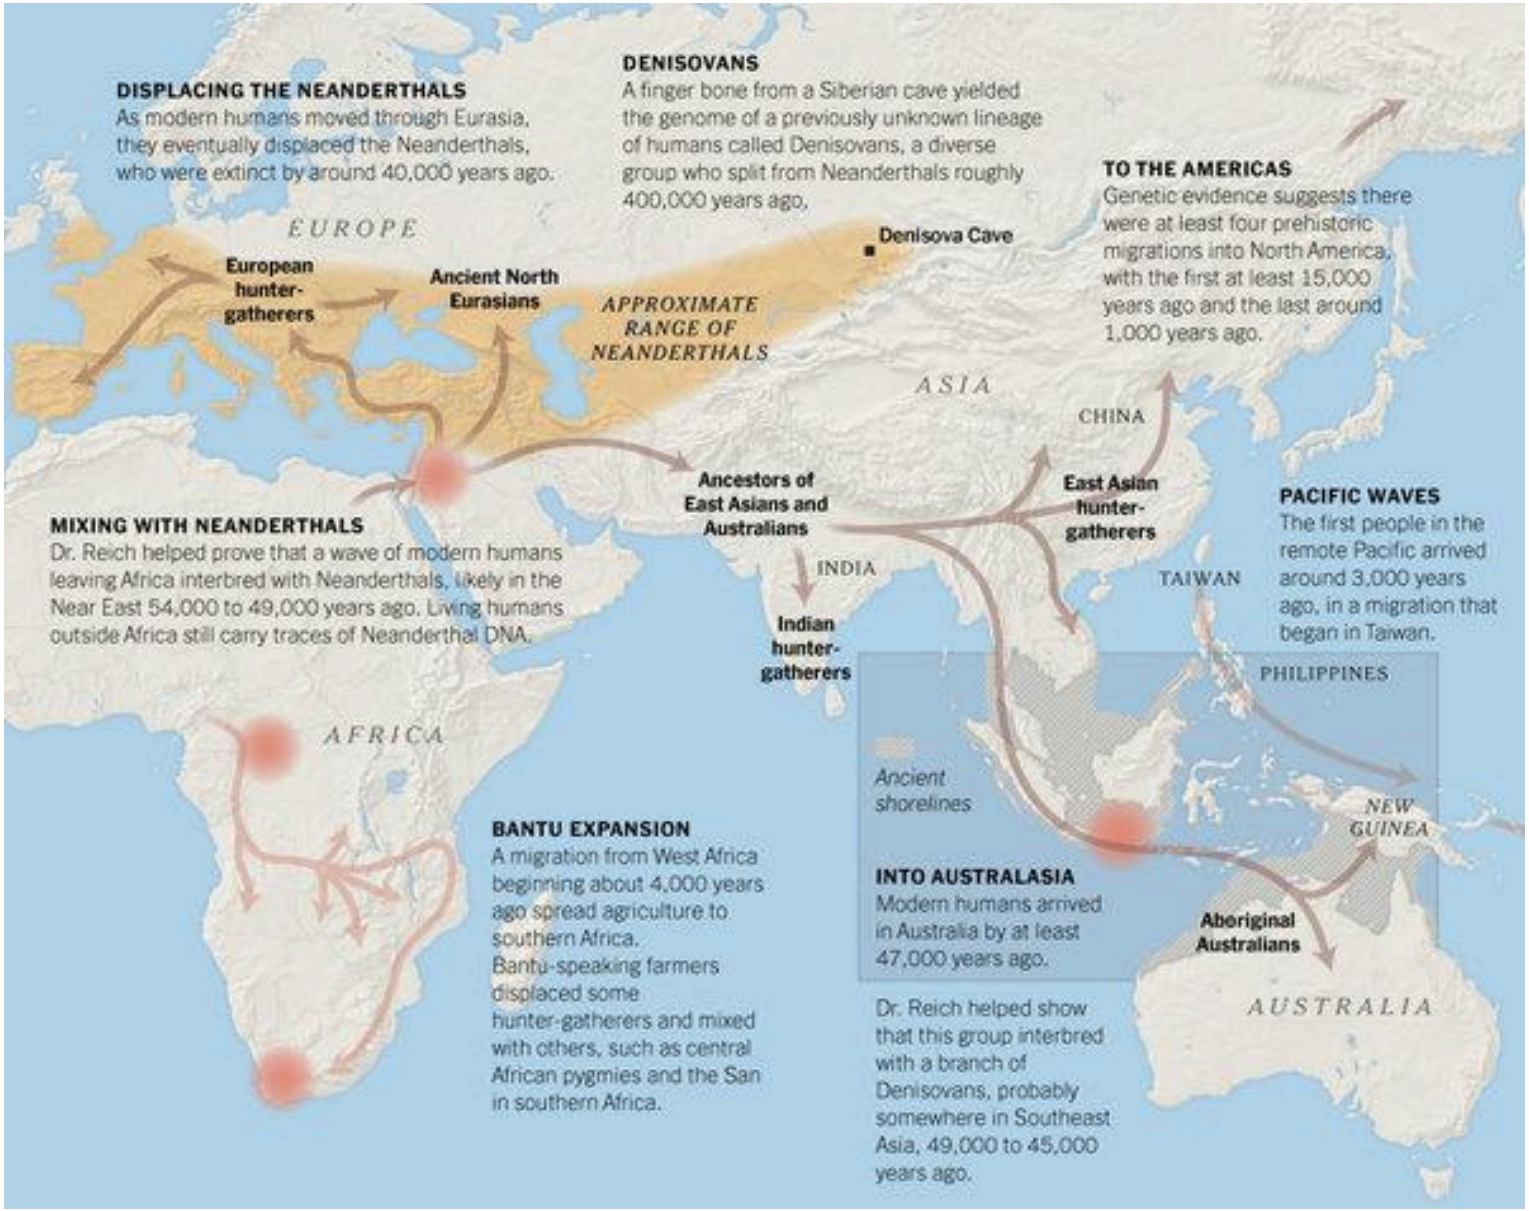

Reich Lab, Genetics (New York Times, 3/2018)

Any (of Your) Questions?

- http://rc.hms.harvard.edu