Volunteer Computing in the Next Decade

David Anderson

Space Sciences Lab
University of California, Berkeley

4 May 2012

The consumer digital infrastructure (CDI)

Consumer

- 1.5 billion PCs
  - GPU-equipped
  - paid for by owner
- 5 billion mobile devices
- Commodity Internet

Organizational

- ~ 5 million cluster/cloud nodes
- supercomputers
- Research networks

Why not use the CDI for scientific computing?

- How to get access to PCs?
  - Poll: would you pay $5/month to help
    - do cancer research
    - predict climate change
    - find the Higgs boson
  - about 5% of people answer “yes”
    - 5% of 1.5B = 75M

Many applications have two levels of parallelism

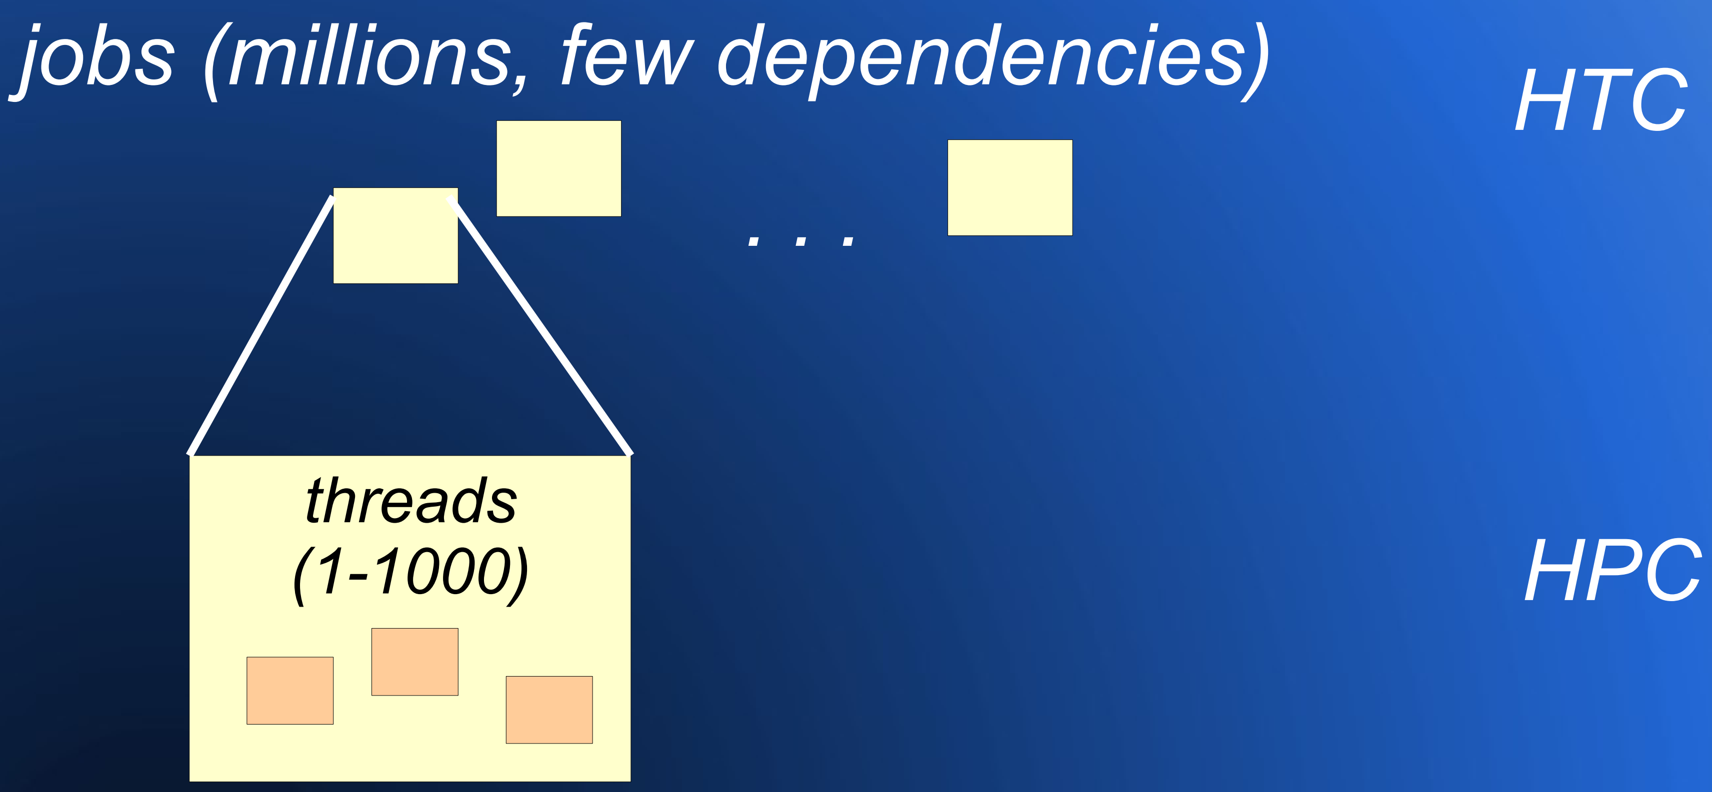

- Don’t need to have all your FLOPS in one cabinet or room

Examples

- Analysis of large data
  - radio/optical/gravitational instruments
  - particle colliders
  - genetic mapping and analysis
- Physical simulations
  - molecular, global, cosmic
- Biology-inspired computing
  - genetic optimization algorithms

Volunteer computing

- PC owners *volunteer* to let science projects use their computing resources
- Your PC can
  - simulate collisions for CERN
  - discover gravitational waves
  - discover binary pulsars
  - research human diseases
  - predict global climate change
  - etc.

BOINC: middleware for volunteer computing

- Supported by NSF since 2002
- Open source (LGPL)
- Based at University of California, Berkeley
- http://boinc.berkeley.edu

Issues handled by BOINC

- Heterogeneity
  - multiple platforms
- Handle untrusted, anonymous resources
  - Result validation
    - replication, adaptive replication
  - Credit: amount of work done by a host or user
- Consumer-friendly client

Supported resource types

- Desktops, laptops, cluster nodes, servers
  - Windows, Mac OS X, Linux, other Unix
  - Can use GPUs as well as CPUs
- Tablets and smart phones
  - Android
- Video game consoles
  - Sony Playstation 3

How it works

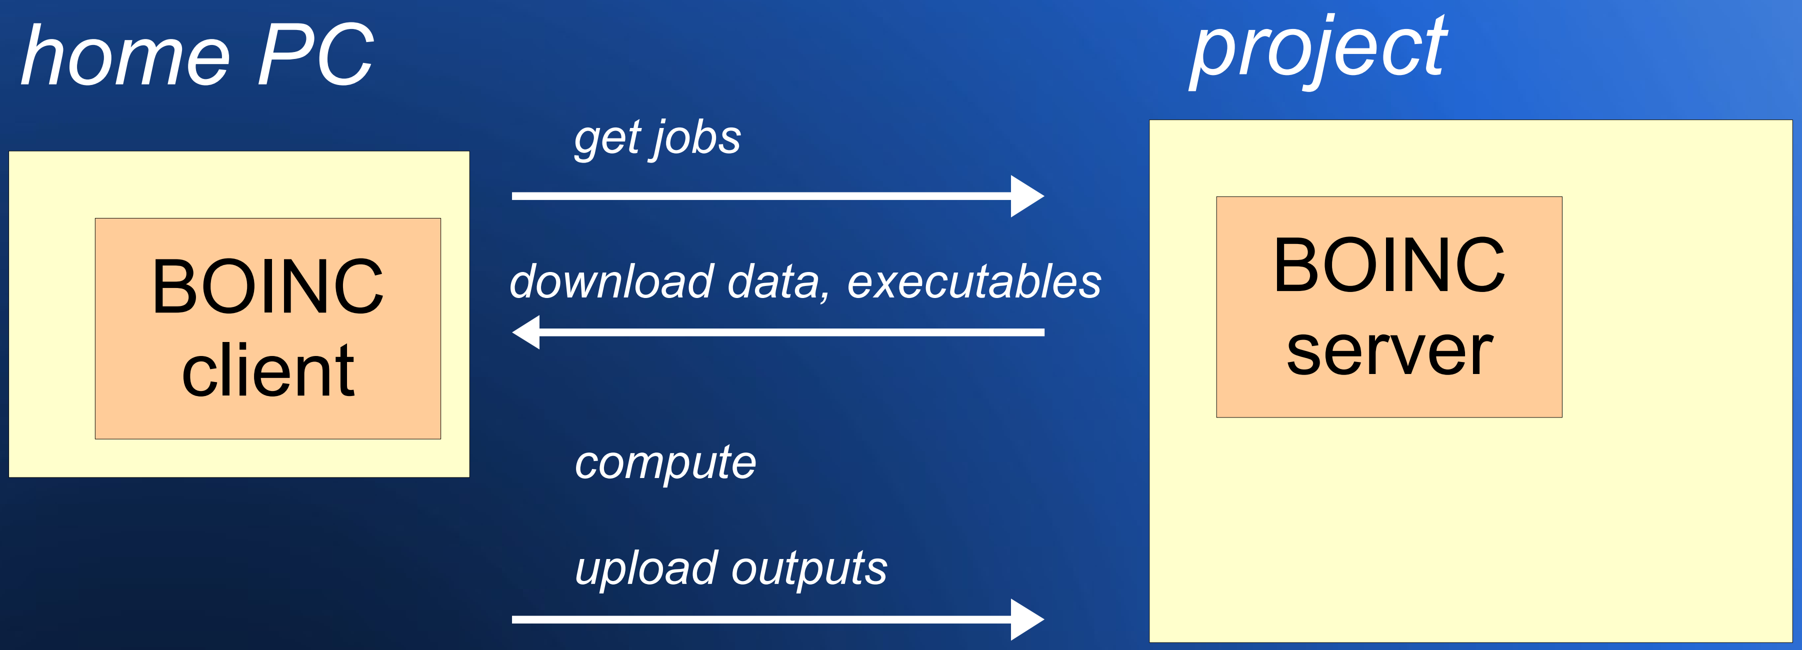

- Can attach to multiple projects
- Works “in the background”; configurable
- Interruptible

Volunteer computing today

- 700,000 active computers
  - 70% availability
  - 40% with usable GPUs
- 12 PetaFLOPS actual throughput
- Projects
  - CAS@home
  - IBM World Community Grid
  - Einstein@home
  - ... ~ 50 others

Cost

The cost of 10 TeraFLOPS for 1 year:

- CPU cluster: $1.5M
- Amazon EC2: $4M
  - 5,000 instances
- Volunteer: ~ $0.1M

The future of the CDI

- Positives
  - Tech R&D is driven by the consumer market
  - Ever-increasing numbers of devices
- Concerns
  - thin clients
  - mobile (battery-powered) devices
  - energy-saving mechanisms
  - proprietary software environments
    - Apple iOS, Windows RT, Xbox, Playstation

The capacity of the CDI in 5 years

- Participation: 50M hosts?
- Processing
  - CPU: 50M * 1 TFLOPS = 50 ExaFLOPS
  - GPU: 50M * 20 TFLOPS = 1000 ExaFLOPS
- Storage:
  - 50M * 10 TB = 500 ExaBytes
- Network bandwidth
  - 50M * 100Mbps = 5 Petabit/sec

Using GPUs

- BOINC detects and schedules GPUs
  - NVIDIA, AMD, soon Intel
  - multiple/mixed GPUs
  - various hardware capabilities
  - various software capabilities (CUDA, OpenCL, CAL)
- Problems
  - non-preemptive GPU scheduling
  - no paging of GPU memory

Using VM technology

- CDI platforms:
  - 85% Windows
  - 7% Linux
  - 7% Mac OS X
- Developing/maintaining versions for different platforms is hard
- Even making a portable Linux executable is hard

BOINC VM support

- Create a VM image for your favorite environment
- Create executables for that environment

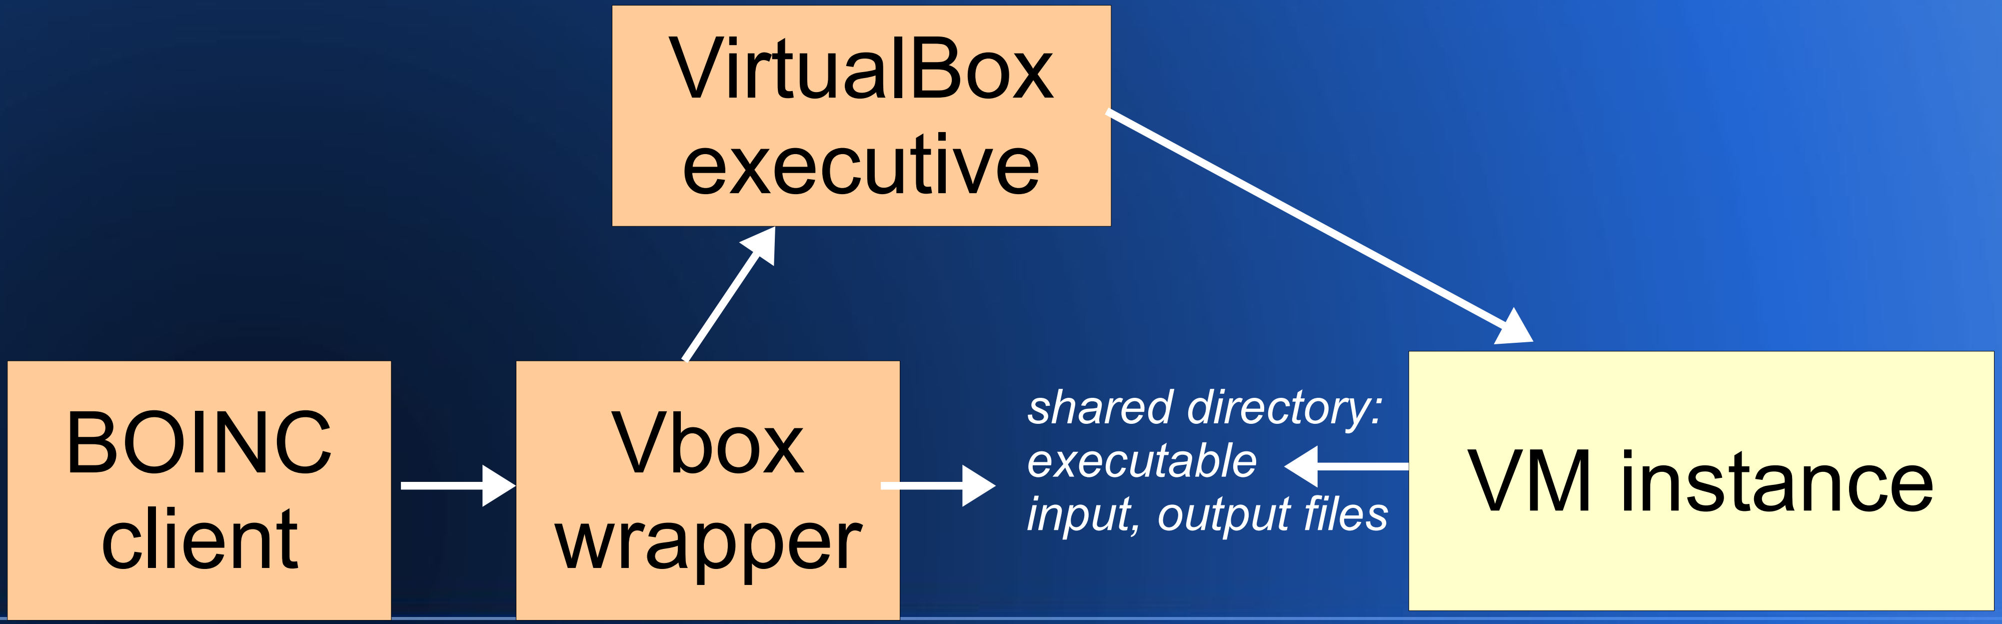

Data-intensive computing

- Examples
  - LHC data analysis
  - Square Kilometer Array
  - Wide-field survey telescopes
  - Genetic mapping and analysis
  - Storage of simulation results
- Performance issues
  - network
  - storage space on clients

Networking landscape

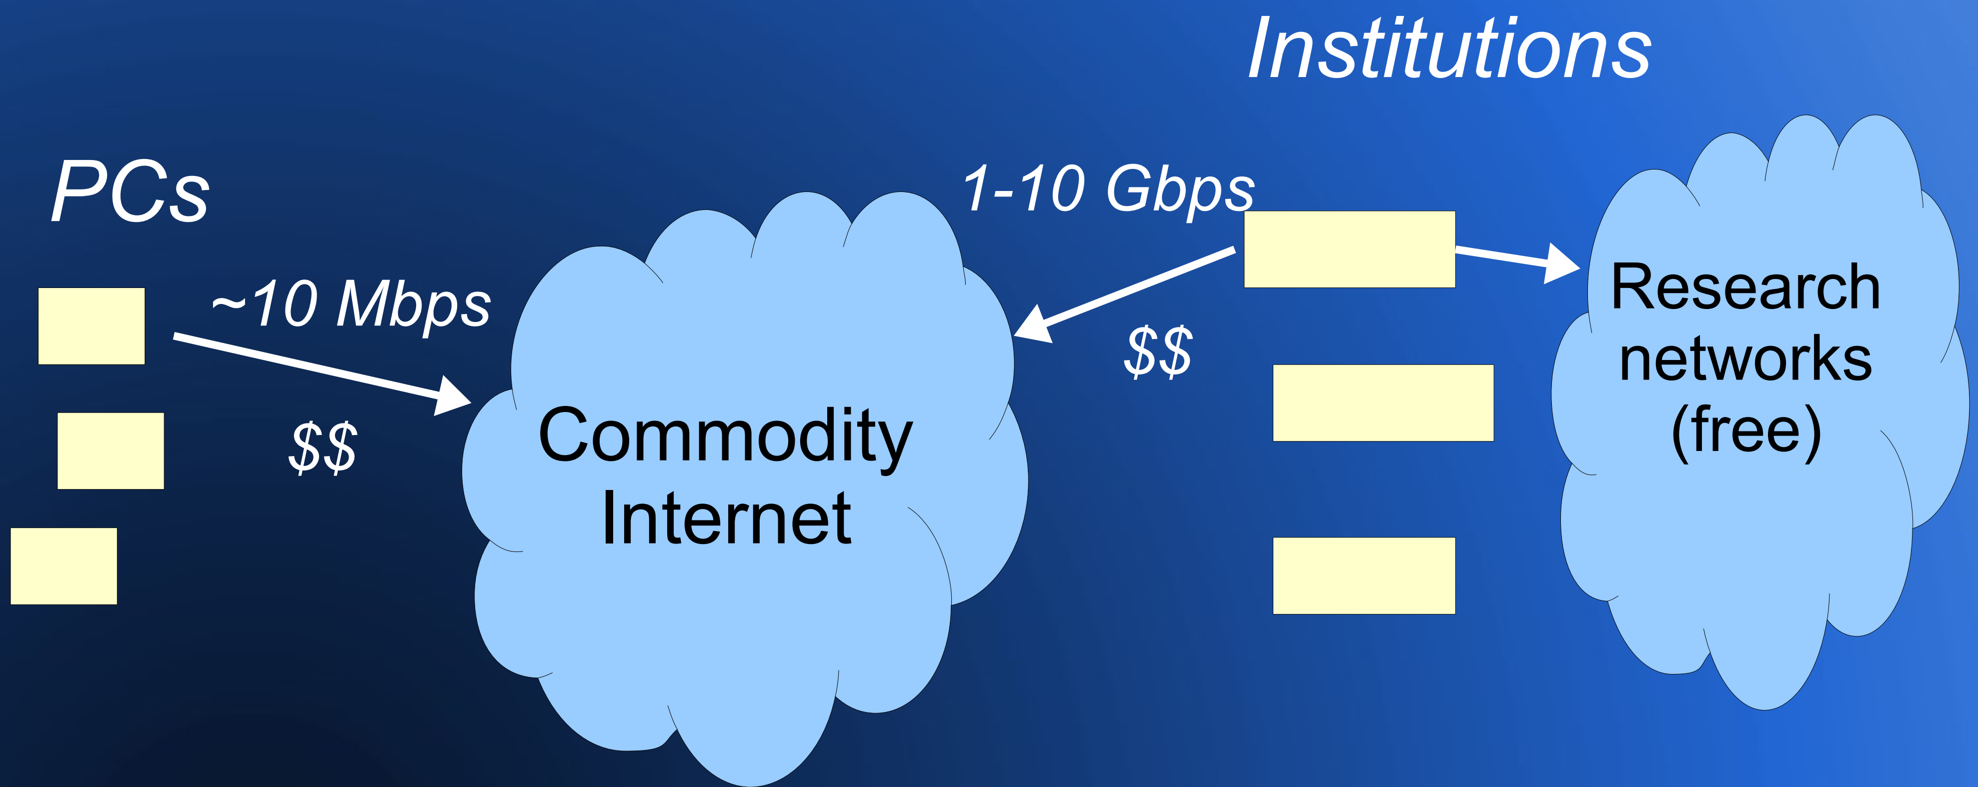

- Most PCs are behind NATs/firewalls
  - only use outgoing HTTP

Properties of clients

- Availability
  - hosts may be turned off
  - hosts may be unavailable by user preference
    - time of day
    - PC is busy or not busy
- Churn
  - The active life of hosts is exponentially distributed with mean ~100 days
- Heterogeneity
  - wide range of hardware/software properties

BOINC storage architecture

Storage applications

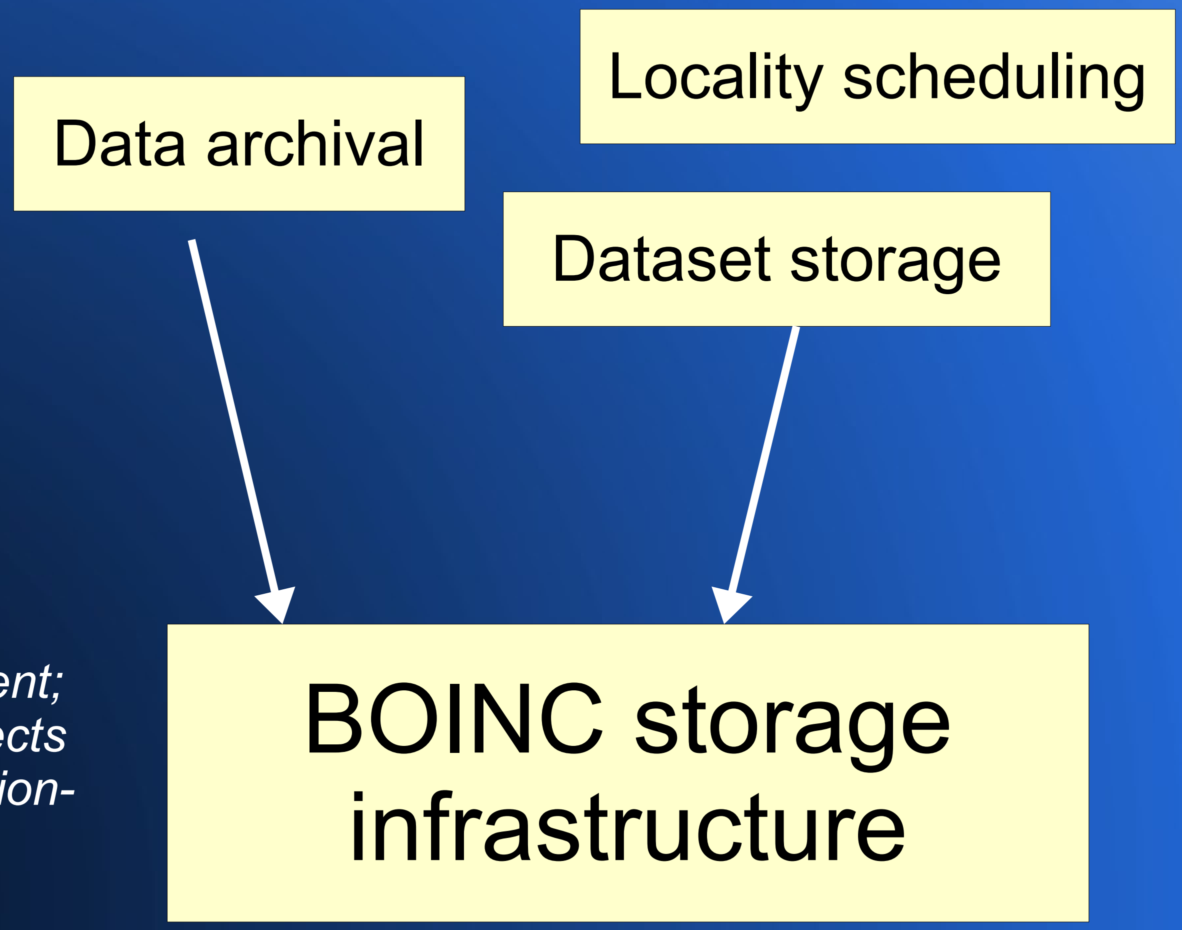

Disk space management;
allocation among projects
framework for application-specific policies
file transfer

Storage applications

- Storage datasets
  - static
  - streaming
- Locality scheduling
- Archival of simulation results
- Data archival

Static dataset storage

- Goal:
  - Minimize turnaround time of queries against a dataset cached on clients
- Scheduling policies
  - whether to use a given host
  - how much data to store on a given host
    - should be proportional to its processing speed
  - update these decisions as hosts come and go

Streaming datasets

- Goal: maximize the window of data availability, subject to a bound on data loss
- Example: Square Kilometer Array (SKA)
  - ~1000 radio telescopes, 2 Gbps each
  - beam-forming via correlation
  - store data on clients in “stripes” across all telescopes
  - with 10M clients, 5 TB each, can store several years of data and perform a variety of analysis functions

Locality scheduling

- Each file in a large dataset is input for a large number of jobs
- Goal: process the dataset using the least network traffic
- Example: Einstein@home analysis of LIGO data

Locality scheduling

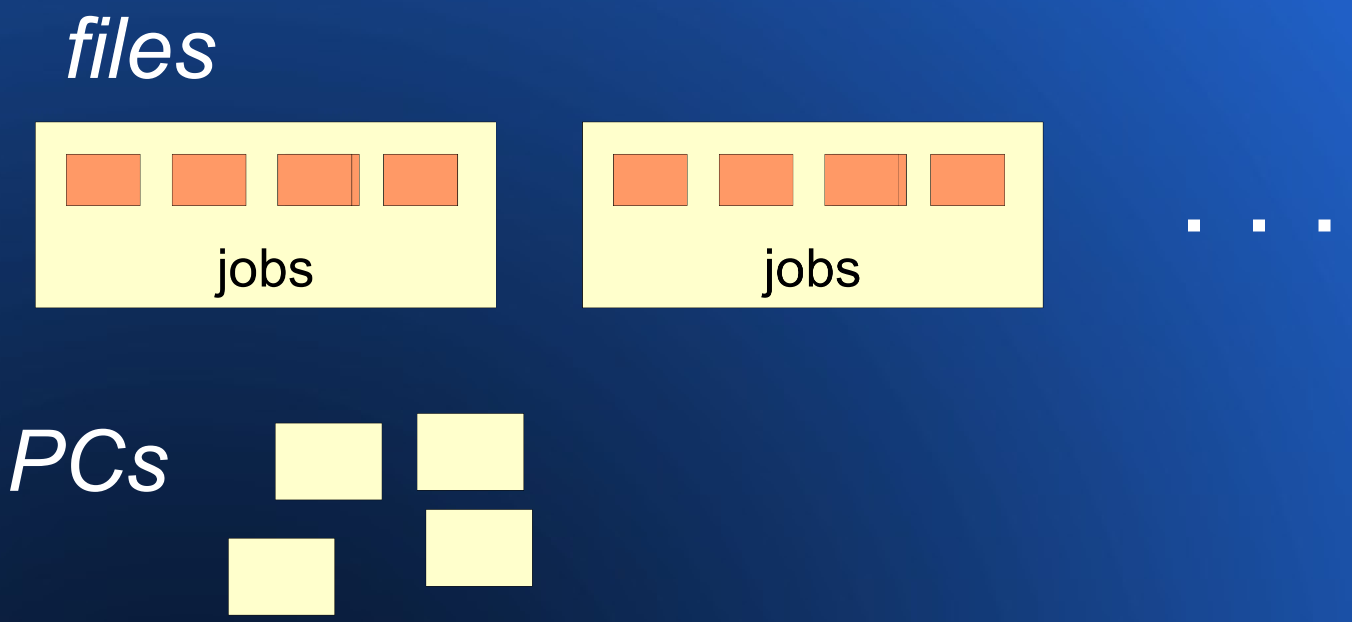

- Processing jobs sequentially is pessimal
  - every file gets sent to every client

Locality scheduling: ideal

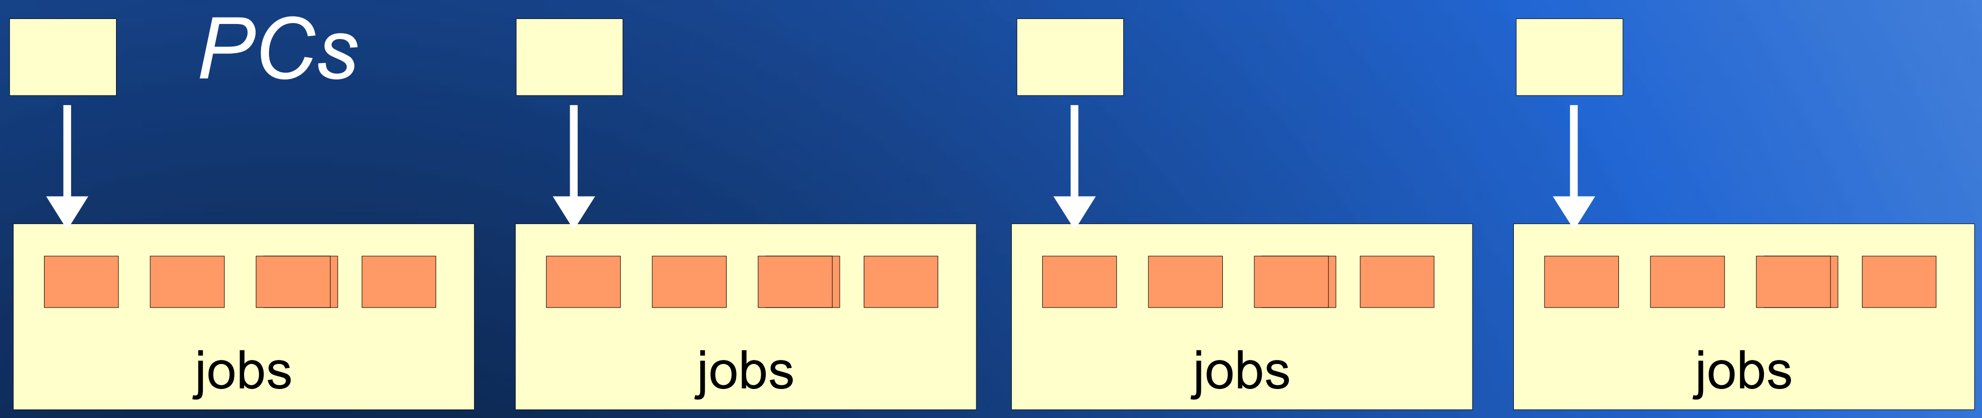

- Each file is downloaded to 1 host
- Problems
  - Typically need job replication
  - Widely variable host throughput

Locality scheduling: compromise

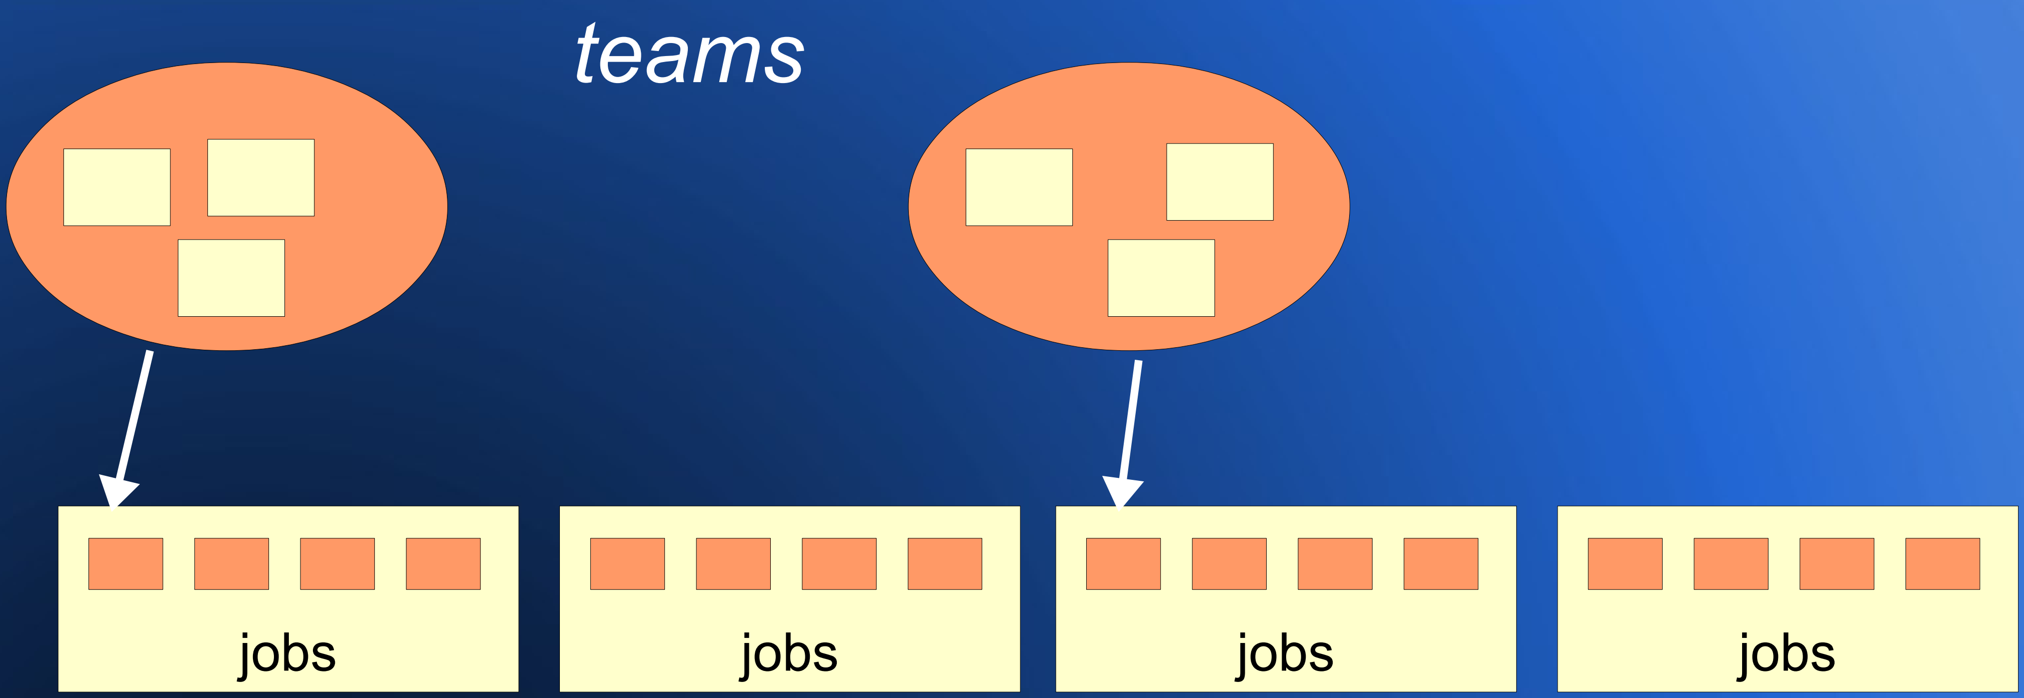

- New hosts are assigned to slowest team
- Teams are merged when they collide
- Each file is downloaded to ~20 hosts

Organizational issues

- What's needed to operate an effective volunteer computing project?
  - learn about BOINC
  - sysadmin, DB admin, application porting
  - publicity/marketing: web, media
  - provide a continuous supply of jobs
- Most research projects can't do all of these

Umbrella projects

- One project serves many scientists
- Examples
  - CAS@home (Chinese Academy of Scientists)
  - World Community Grid (IBM)
  - U. of Westminster (desktop grid)
  - Ibercivis (Spanish consortium)

How to use volunteer computing

- Think about future computing needs
  - what can be mapped to volunteer computing?
  - how could you use unlimited processing power?
- Think about your organizational structure
  - mobilize at the appropriate level
- Plan and gather resources
  - include in your next grant proposal

Conclusion

- Volunteer computing is the shortest path to ExaFLOPS computing
- It’s usable for many applications
- It enables breakthrough research
- It requires a leap of faith