DMTCP for Checkpoint-Restart: its Past, Present and Future

Gene Cooperman

College of Computer and Information Sciences
Northeastern University

November 6, 2013

DMTCP: Distributed MultiThreaded Checkpointing

- As easy to use as:

```
dmtcp_launch ./a.out
dmtcp_command --checkpoint
dmtcp_restart ckpt_myapp_*.dmtcp
```

- A Quick Demo!
- The project is now 9 years old.
- (. . . and now also funded by Intel Corporation)

Undergraduates who have contributed to DMTCP

- Jason Ansel (class of 2007
  - DMTCP, multiprocess distributed checkpointing.
  - CRA Award[1] finalist.
  - Currently at MIT, Ph.D. program.
- Alex Brick (class of 2011)
  - Providing MTCP as a checkpoint service for OpenMPI.
  - Currently at Amazon.
- Tyler Denniston (class of 2012)
  - Record-replay and python infrastructure for FReD, the reversible debugger.
  - CRA Award*: Honorable mention.
  - Currently at MIT, Ph.D. program.
- Gregory Kerr (class of 2013)
  - The first version of InfiniBand support for DMTCP.
  - Currently at Apple Corporation.

[1] Computer Research Association: Award for best undergraduate computer science researchers in North America

DMTCP Users

http://sourceforge.net/projects/dmtcp/files/stats/map?dates=20

DMTCP Architecture

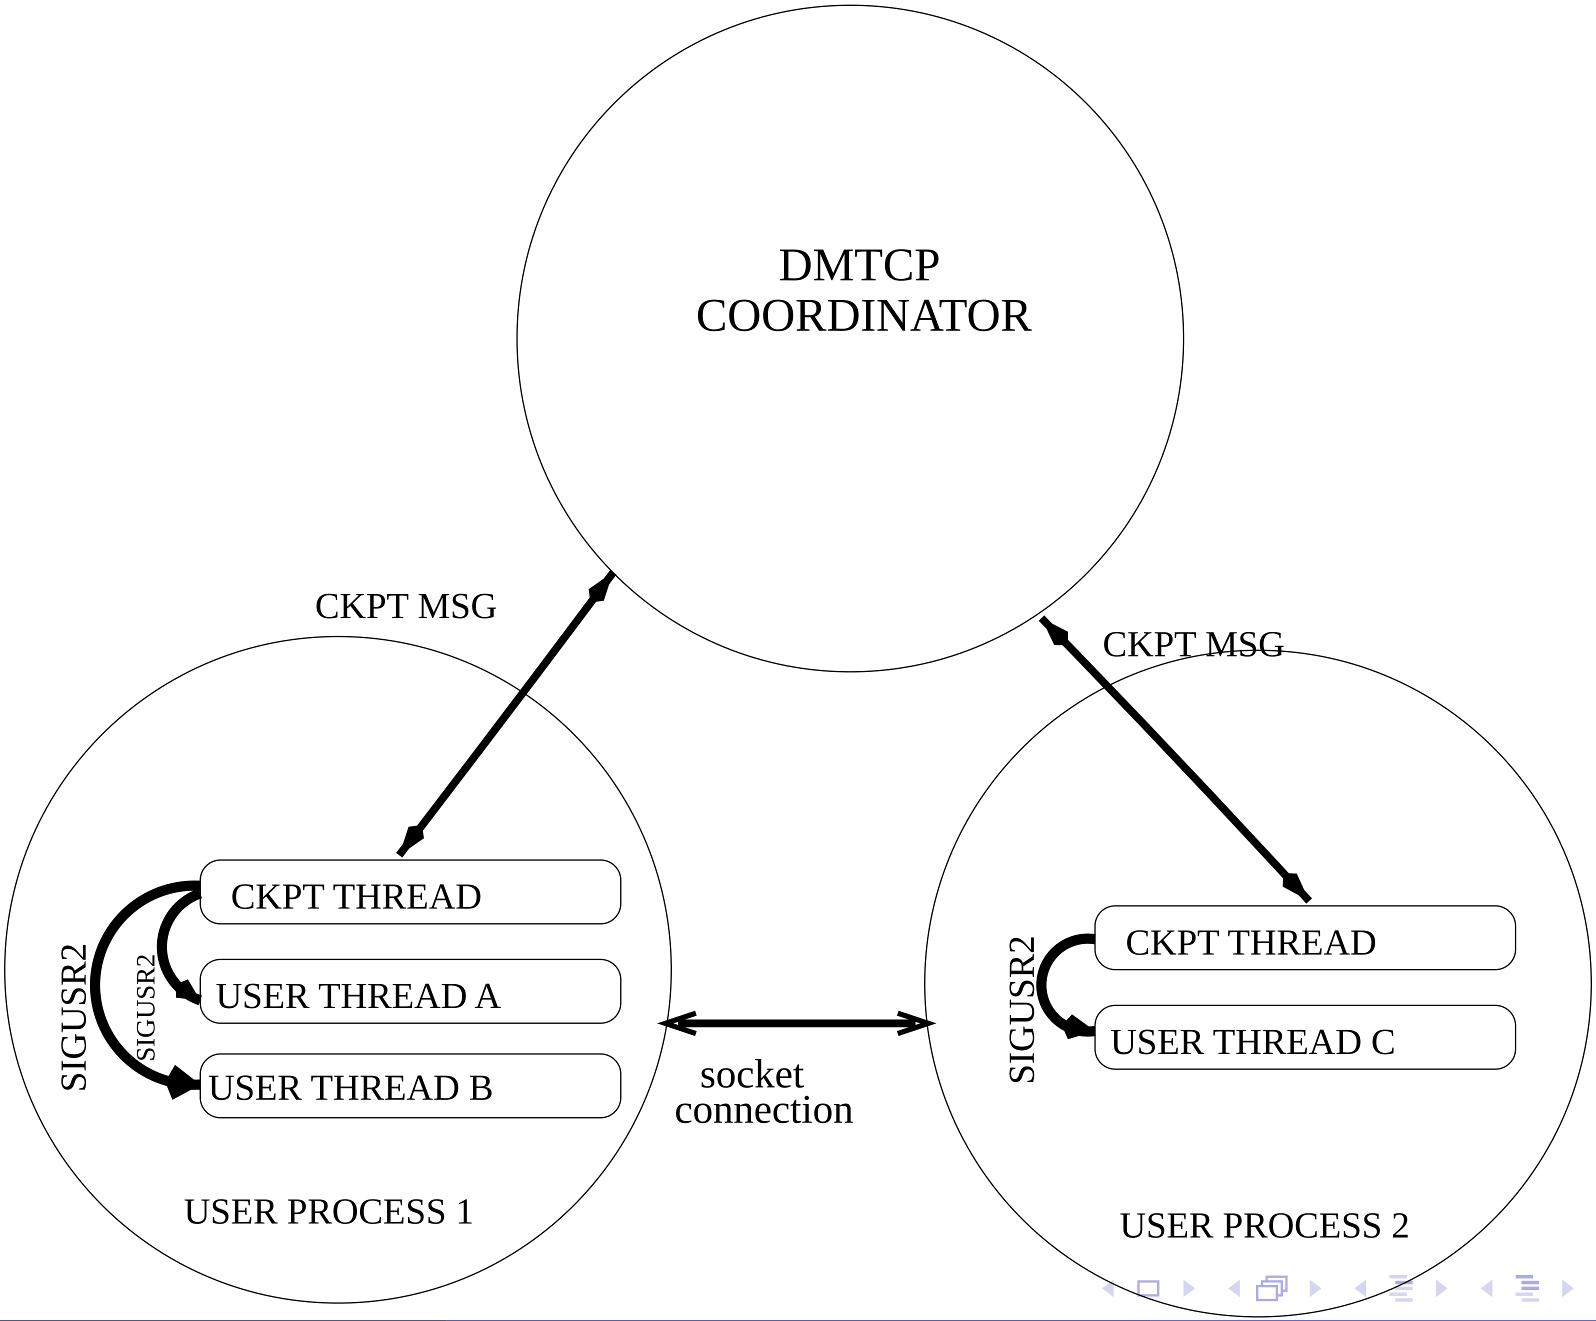

Ultimate Goal

- **Vision:** Why do we need program and data input?
- Why not just save the state of the process and consider them as as an object that can be modified, replicated, migrated, etc?
- For example, bring the process to the data, instead of bringing the data to the process.
- To accomplish this vision, we face a series of technical challenges.
- Here we present them as a series of puzzles with existing and proposed solutions.

Virtualizing the Process Id

- **PRINCIPLE:**
  The user sees only virtual pids; The kernel sees only real pids

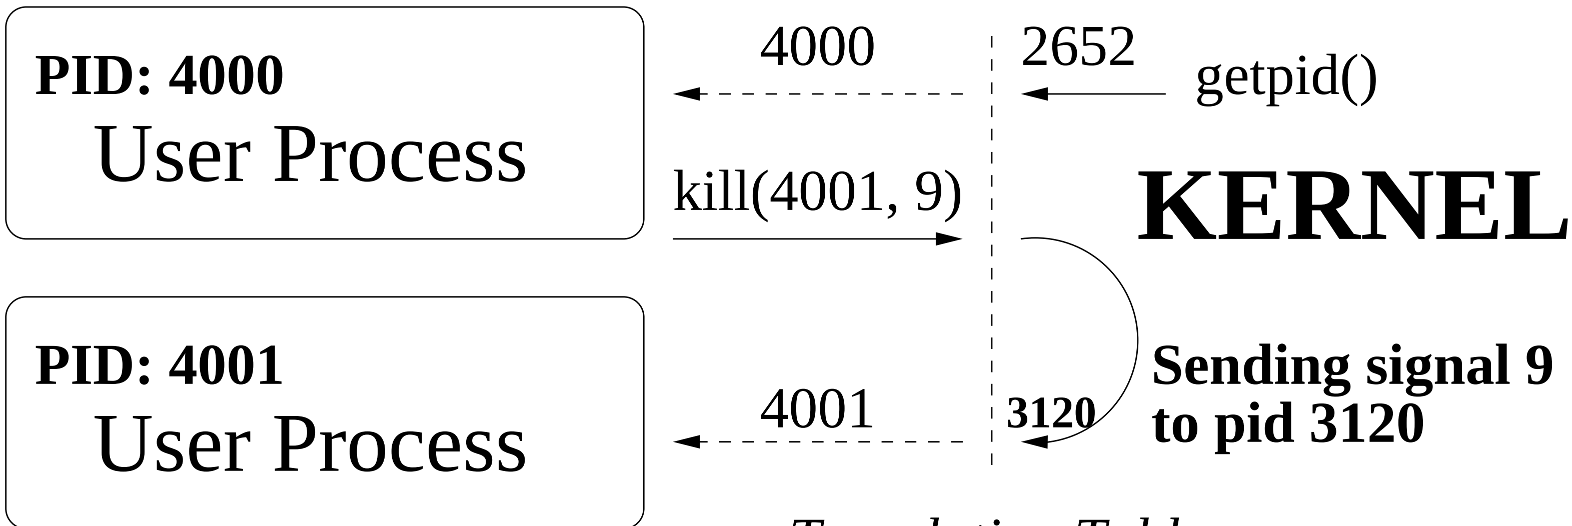

Translation Table

Virt. PID	Real PID
4000	2652
4001	3120

DMTCP and the Cloud

- **The Cloud is the new mainframe:**
  - *on-demand self-service, broad network access, resource pooling, rapid elasticity, measured service*
- But it needs a new scheduler!
- Amazon EC2 and other clouds are inflexible:
  Example: By default, one pays by the hour, even if we don't use it all.
  Why? (... maybe because they don't have a preemptive scheduler)
- **So, let's schedule jobs using DMTCP:**
  - setting priorities
  - suspend jobs
  - migrate jobs
- *Issue:* On restart, the hostname changes, the directory changes, . . .
- *Solution:* Use a plugin to virtualize hostnames and directory names.

DMTCP and the Cloud

- OpenStack

  `http://openstack.org/`
- Massachusetts Open Cloud

  `http://www.massopencloud.org/`

Cloud computing

- Clouds having a dramatic impact:
  - **Consumer**: on-demand access to inexpensive computational capacity, pay for what you use
  - **Producer**: economy of scale, automation
- Like power, most computation will move into **public** clouds.

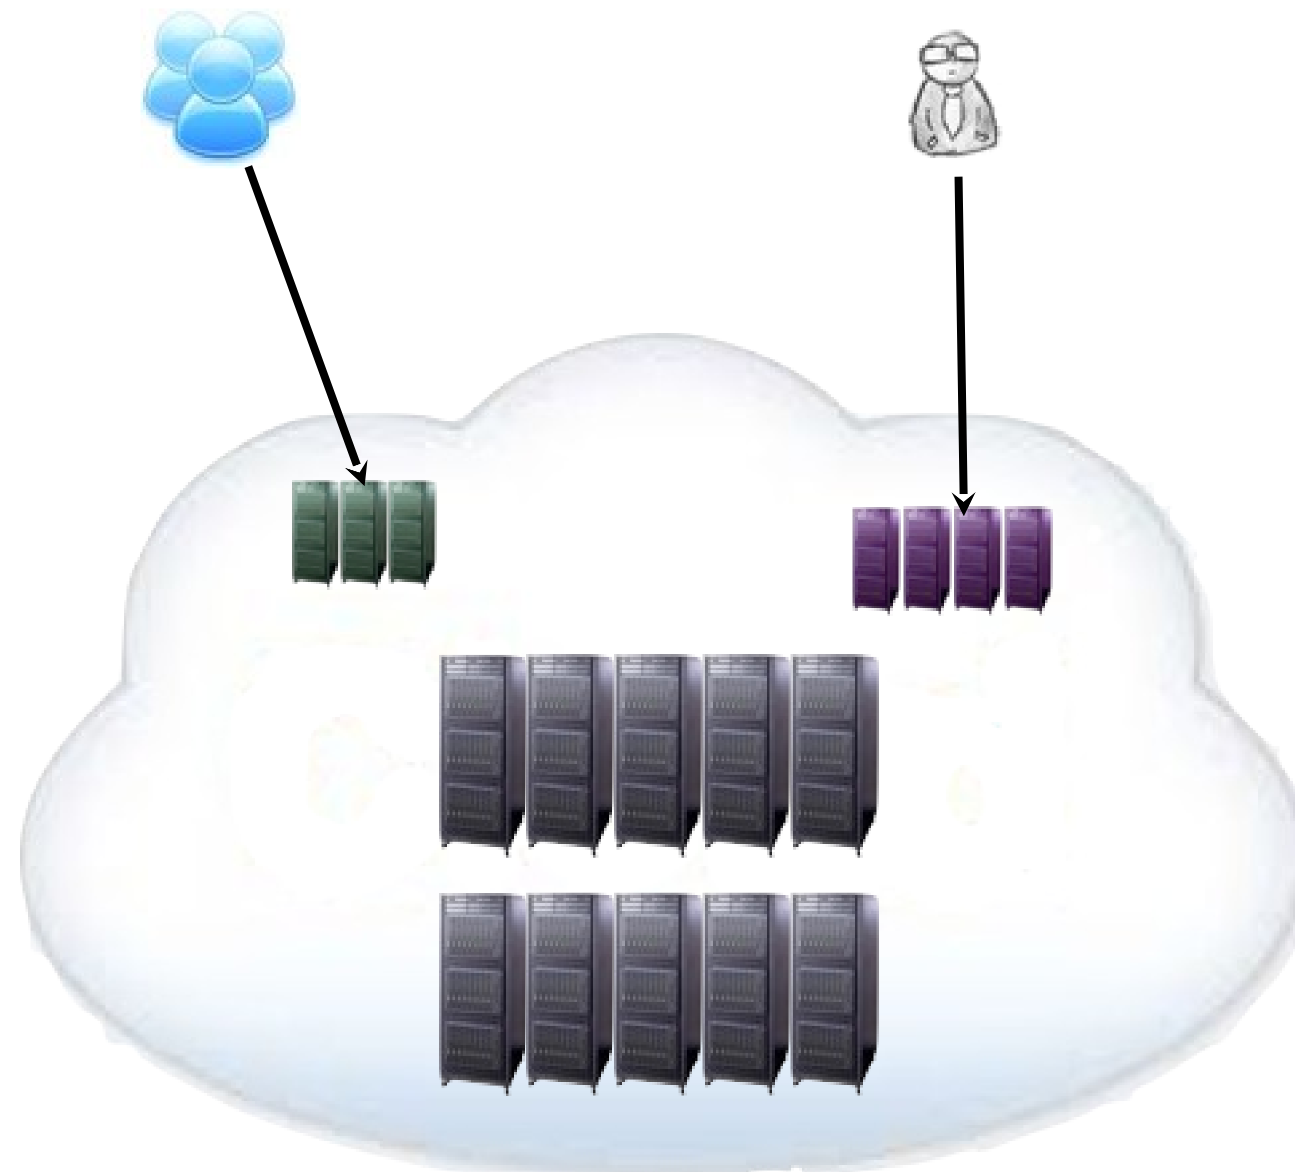

Problems with today's "closed" public clouds

- Highly prescriptive in HW, computational model, economic model; focus on scale-out web applications
- Operational/performance data limited to the single provider
- Limiting research, innovation by third parties
  - ➡ technology companies locked out of public clouds; disconnect with private clouds
  - ➡ difficult for anyone else to efficiently support/innovate Big data platforms,
- No visibility/auditing of internal operations:
  - ➡ Major security challenge for hosting critical datasets
- Accretion of features/services into Provider offering
- Monoculture increasingly dangerous
- Vendor lock in by features, interfaces, and pricing model.

A new model is required: an "open cloud"

- Multiple "partners" participate in implementing and operating cloud
- Each partner determines how to charge for her services
- Operational data visible to stakeholders
- Domain specific "intermediaries":
  - provide customers with simple model
  - enable optimization
- Multi-sided marketplace

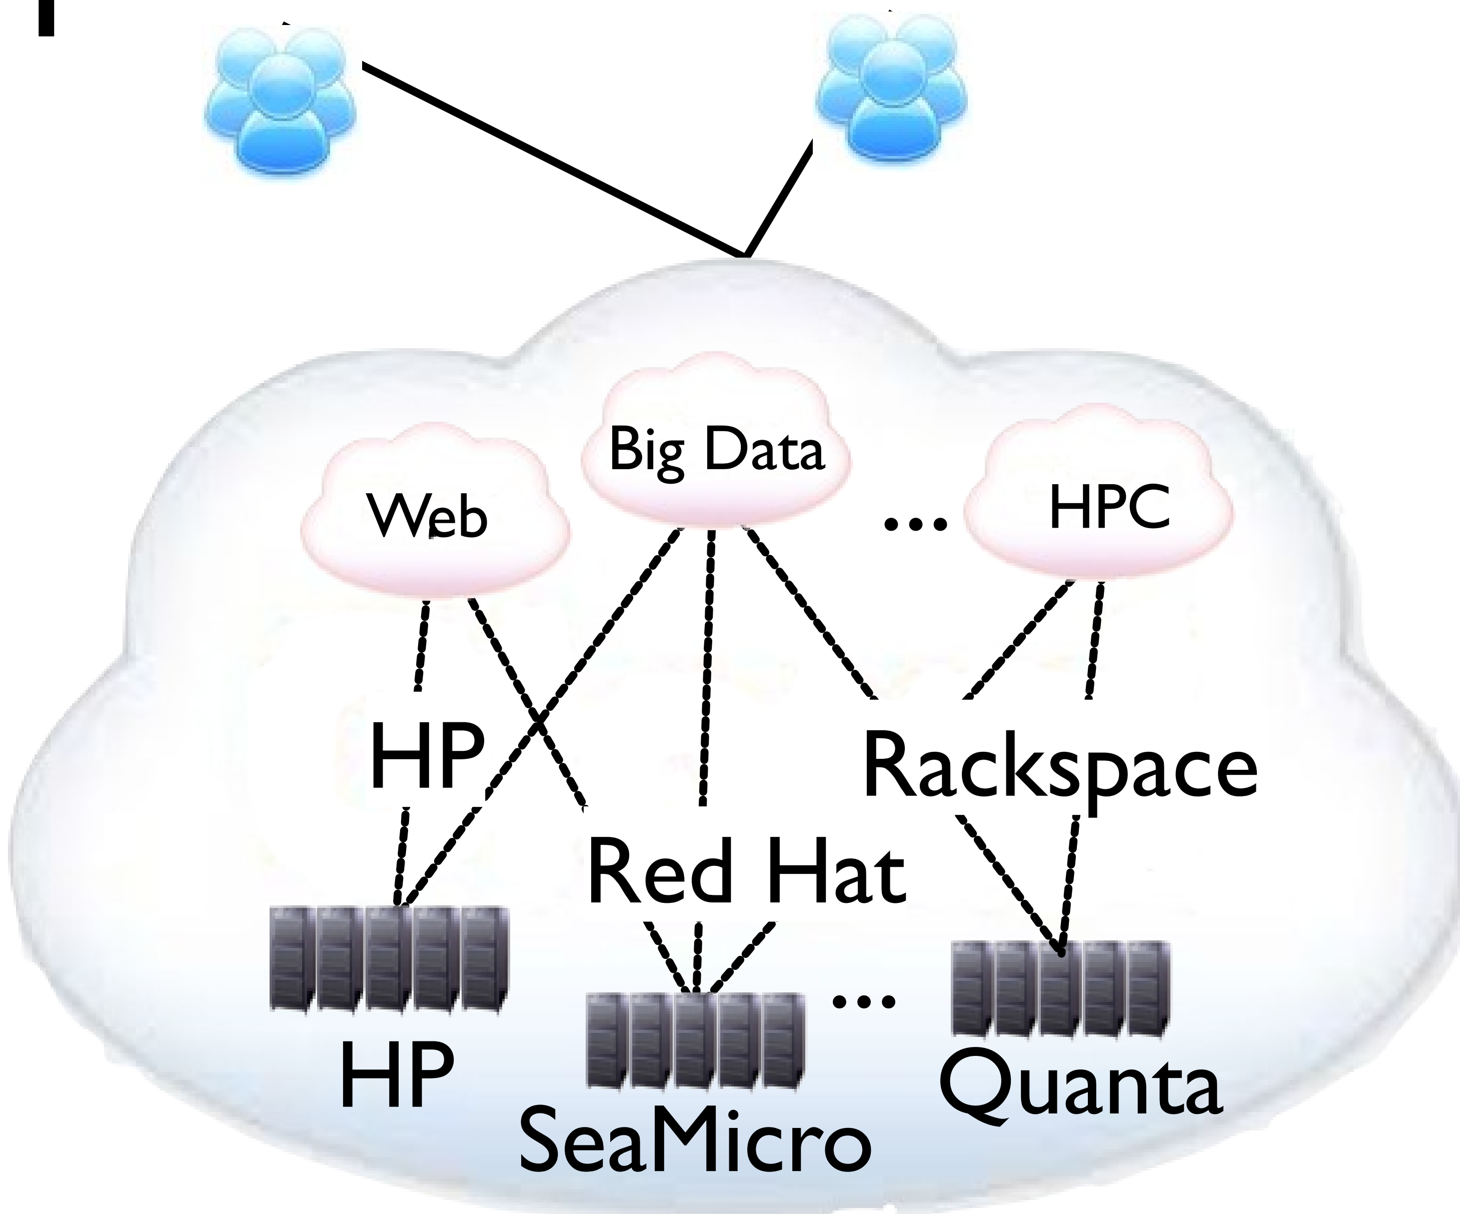

The Opportunity

- OpenStack provides most of what we need:
  - modular structure with multiple independent services and support for plugins
- 15 MW MGHPCC data center, low power cost, excellent network connectivity...
- MGHPCC consortium: BU, MIT, NEU, UMass, Harvard.
  - Operate production cloud capacity for research computation & enable Big Data & HPC users
  - Enable research in Big Data, Cloud Computing
- Incredible regional cluster of technology companies and innovative users of technology
- Commonwealth Big Data Initiative
- Launched attempt to create "Massachusetts Open Cloud (MOC)" as a partnership: State, MGHPCC, Industry

DMTCP and the Cloud

- **The Cloud is the new mainframe:**
  - *on-demand self-service, broad network access, resource pooling, rapid elasticity, measured service*
- But it needs a new scheduler!
- Amazon EC2 and other clouds are inflexible:
  Example: By default, one pays by the hour, even if we don't use it all.
  Why? (... maybe because they don't have a preemptive scheduler)
- **So, let's schedule jobs using DMTCP:**
  - setting priorities
  - suspend jobs
  - migrate jobs
- *Issue:* On restart, the hostname changes, the directory changes, . . .
- *Solution:* Use a plugin to virtualize hostnames and directory names.

DMTCP Plugins: Adapting to the Cloud

- Examples of adapting to the Cloud on restart
  - Virtualize the hostname, network address, and current directory, since the Cloud will assign new values for all of the above.
  - Other changes: directly notify the application by changing its environment variables (e.g., `NUM_RESTARTS`, `CURR_PRIORITY`, etc.)
- **DEMO**:
  `cd test/plugin; make check sleep1; make check sleep2`
- *What else is there:*
  `ls contrib`
- **Changed Environment Variable**: `contrib/modify-environ`

Virtual Machines and the Cloud

- IaaS (Infrastructure as a Service): Translation into English: Put your software in a virtual machine, and the Cloud will run your virtual machine for you.
- *Issue:* The Cloud needs to migrate virtual machines to handle priorities, suspending low-priority jobs, load-balancing of resources, etc.
- *Solution:* virtualize the KVM API Write a DMTCP plugin to checkpoint the KVM virt. machine. *Total lines of code (kvm plugin):* 300 lines

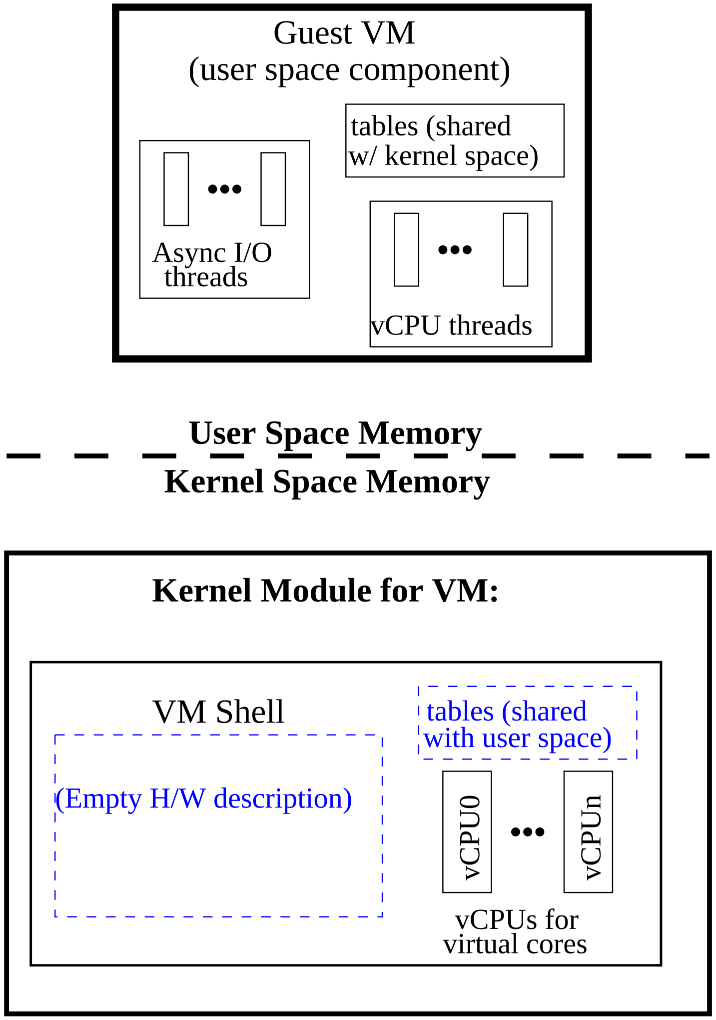

Network of Virtual Machines and the Cloud

- The Cloud needs to also handle parallel or distributed computations. (sometimes called "parallel instances")
- *Issue:* Current virtual machine snapshots cannot also save the state of the network. (Networking virtual machines requires the Linux Tun/Tap kernel module.)
- *Solution:* virtualize the Tun network.
  Write a DMTCP plugin to save the state of the "Tun" network between virtual machines on different physical nodes.
  Total lines of code (tun plugin): 100 lines

Hadoop and Big Data

- Hadoop is already fault tolerant.
- *Problem:* How to preempt a low-priority Hadoop job when a high-priority job is submitted?
- *Solution:*
  - After a "reduce" operation, ask Hadoop to save the back-end data somewhere.
  - Checkpoint the front-end Hadoop application.
  - Use a DMTCP plugin to re-connect the front-end and back-end on restart.

DMTCP and High Performance Computing (HPC)

- HPC workloads require submitting jobs to a batch queues.
- *Problem:* On restart, we may be on a different set of nodes.
  - Directory names, hostnames, files, environment variables, etc. may all change.
- *Solution:* Virtualize directory names, hostnames, etc.

InfiniBand

- Today in HPC, they don't use Ethernet. *(It's too slow!)*
- InfiniBand achieves both high bandwidth and low latency.
- InfiniBand uses RDMA (Remote Direct Memory Access).
- RDMA uses *send queue*, *receive queue*, and *completion queue*

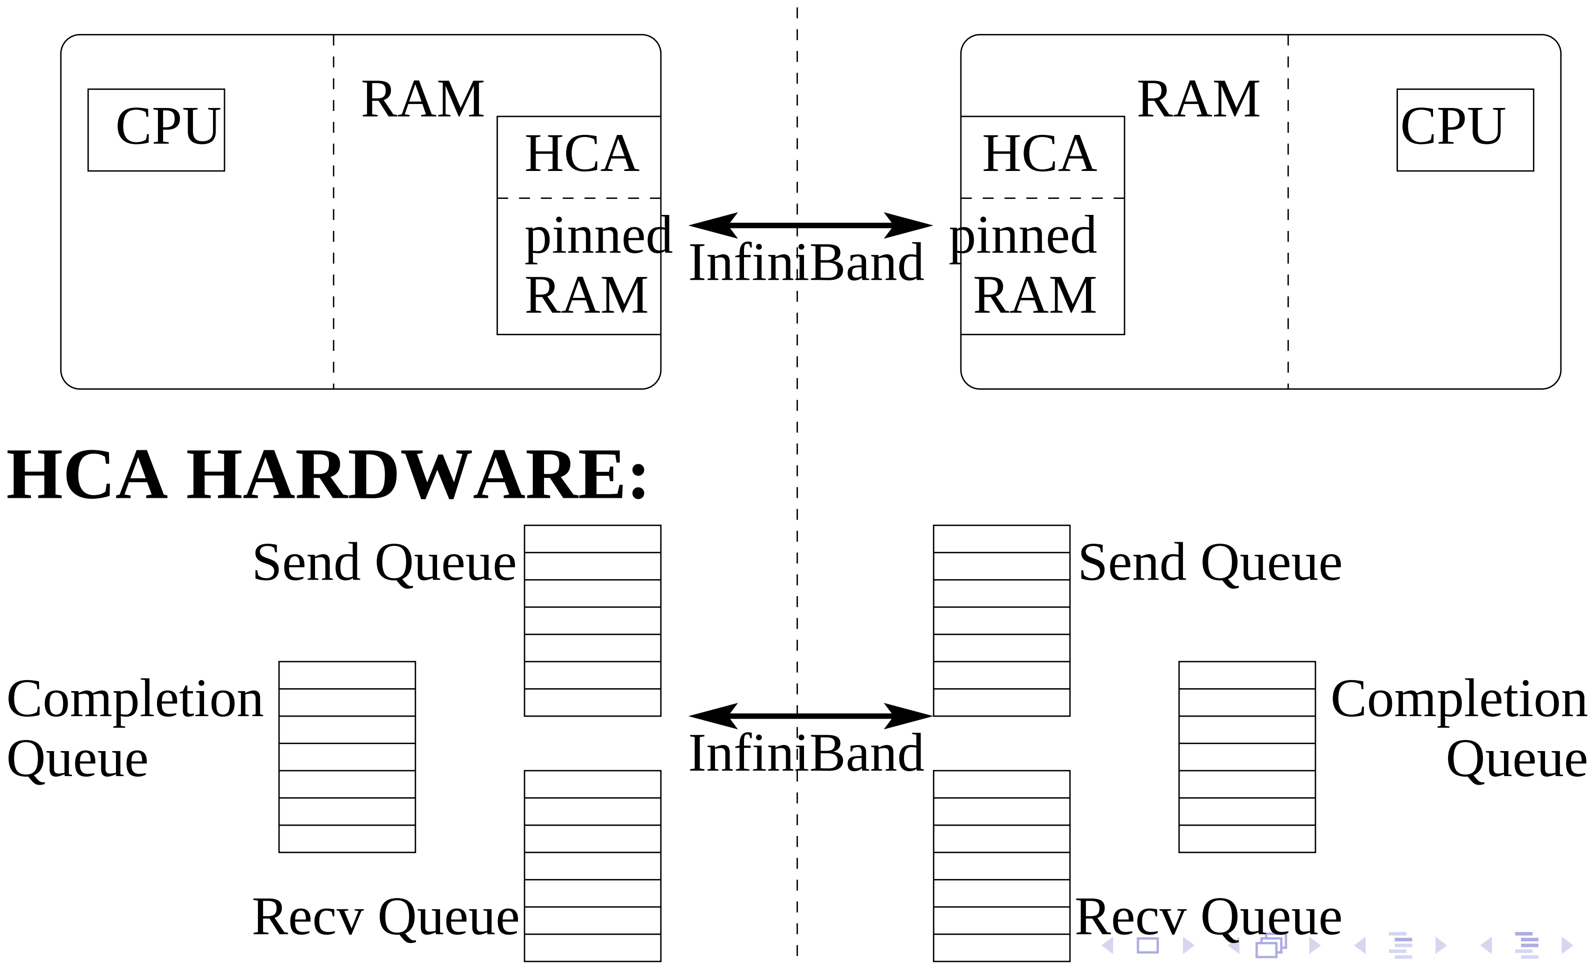

DMTCP and InfiniBand

- *ISSUES:* At restart time, totally different ids and queue pair ids.
- *Solution:* Drain the completion queue and save in memory.
  On restart, virtualize the completion queue:
  - Virtualized queue returns drained completions before returning completions from the hardware.
  - Virtual queue pair ids (id pointing to a hardware).

Supercomputing: the Intel Xeon Phi

- Xeon Phi: 61 Intel 64-bit cores on a single chip
- Stampede (TACC: Texas Advanced Computer Center, 6,400 nodes)
  #6 on list of largest supercomputers
  http://www.tacc.utexas.edu/stampede/
- Tianhe-2 (16,000 computer nodes: 48,000 Xeon Phis)
  #1 on list of largest supercomputers
  http://en.wikipedia.org/wiki/Tianhe-2
- Future supercomputer: Coral (by 2017?, Xeon Phi??)
  Coral: Collaboration, Oak Ridge, Argonne, Livermore
  http://www.hpcwire.com/hpcwire/2013-01-16/doe_to_field_pre

Challenges for DMTCP and Supercomputing

- **Scalability Testing:**
  Initial scalability testing at CERN (recently in the news for discovery of Higgs-Boson particle)
  - Test on thousands of computer nodes at CERN, while the LHC collider is down during the planned upgrades to its performance.
  - *Issue:* Requires DMTCP sub-coordinators (or tree of coordinators)
  - *Solution:* DMTCP plugin
- **Virtualizing the New HPC Networks:**
  InfiniBand must be replaced by customized RDMA network.
  - Tianhe-2 (TH Express-2 interconnect network)
  - Cray XK7 supercomputer (Gemini network)
  - *Issue:* Virtualizing these HPC networks
  - *Solution:* Modify the DMTCP InfiniBand plugin

A New Challenge: DMTCP and OpenGL (3-D graphics)

- Usually a virtual machine cannot take a snapshot of 3-D graphics (cannot snapshot OpenGL applications). This is because the 3-D graphics object are saved in the graphics hardware.
- *Issue:* Same problem as we saw with InfiniBand hardware. What is the solution this time?
- *Solution:* Record, compress, and replay the commands. Virtualize the graphics objects in the graphics hardware accelerater.

Grand Challenge for DMTCP

Gather a user-community to contribute libraries of plugins for adapting software to new environments.

- *After all, we can't write all the plugins.* :-)

Thanks to the many students who have contributed to DMTCP over the last nine years. Among the larger contributors are:

> *Jason Ansel, Kapil Arya, Alex Brick, Jiajun Cao, Tyler Denniston, Xin Dong, Rohan Garg, Samaneh Kazemi, Gregory Kerr, Artem Y. Polyakov, Michael Rieker, Praveen S. Solanki, Ana-Maria Visan*

QUESTIONS?