Grid Mashups

Gluing grids together with Condor and BOINC

Mark Silberstein, Artyom Sharov, Assaf Schuster, Dan Geiger
Technion – Israel Institute of Technology

Problem ....

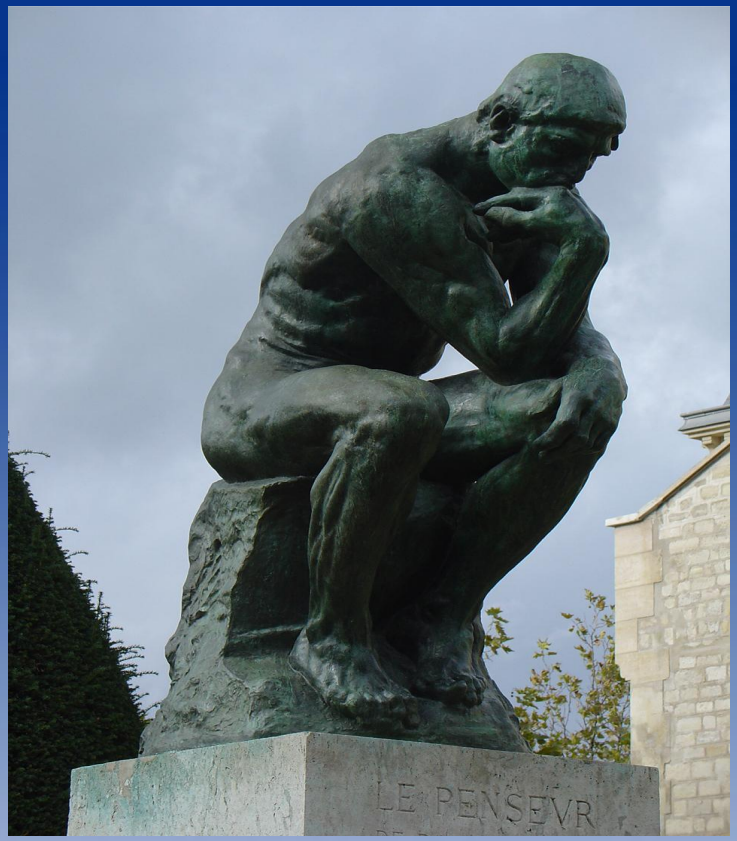

Problem ....

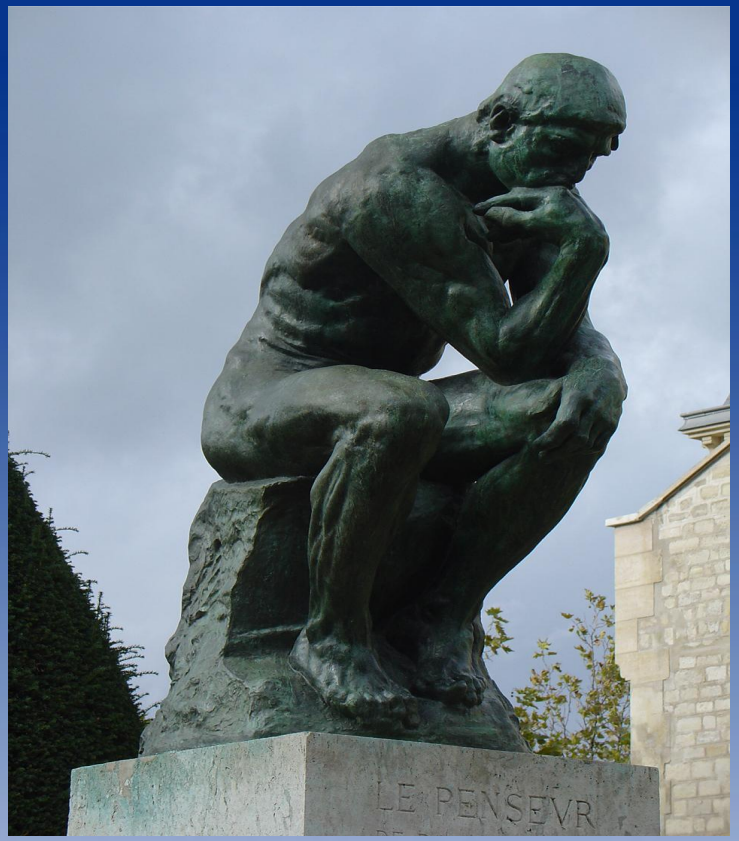

Problem ....

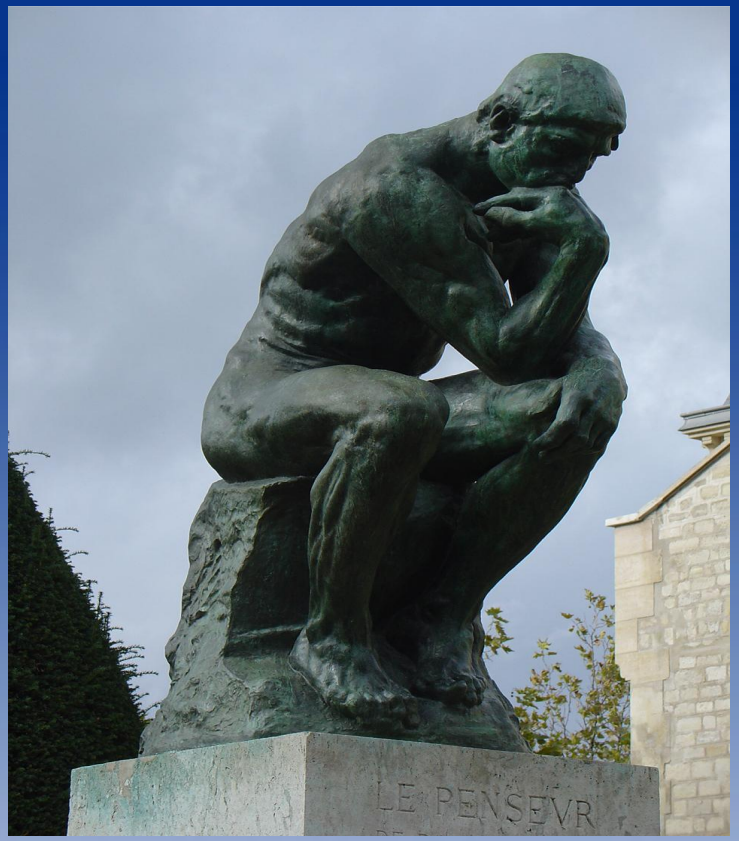

Parallelization

From tens to millions of subtasks

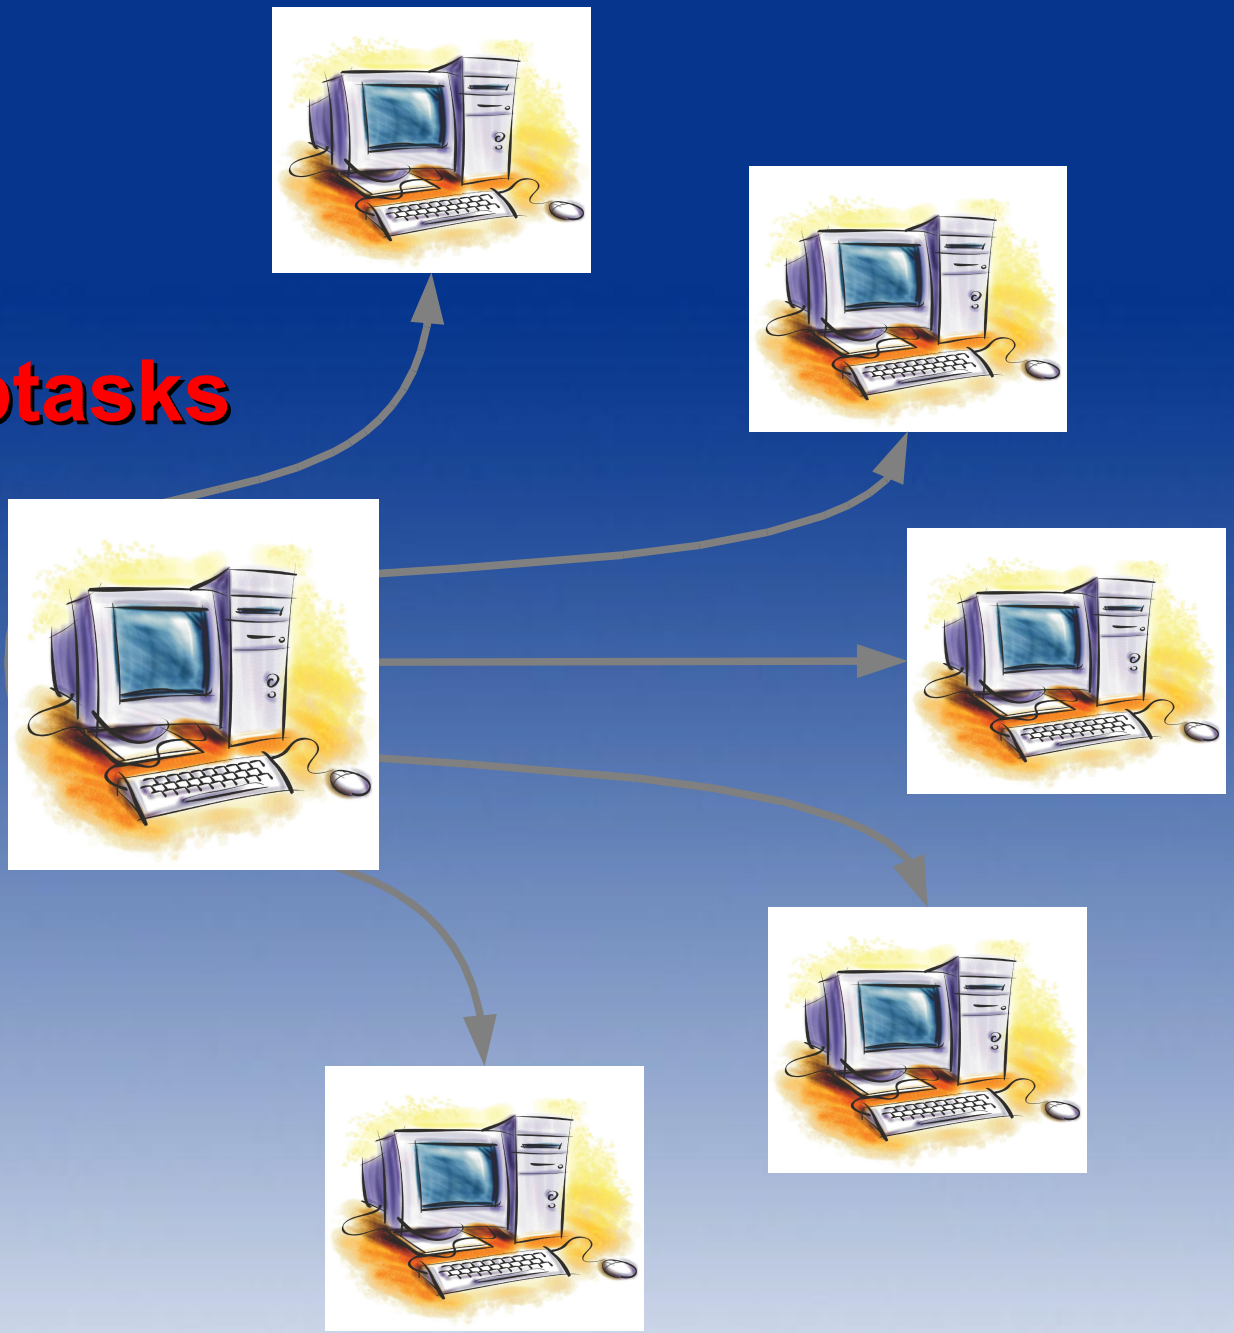

Parallelization

From tens to millions of subtasks

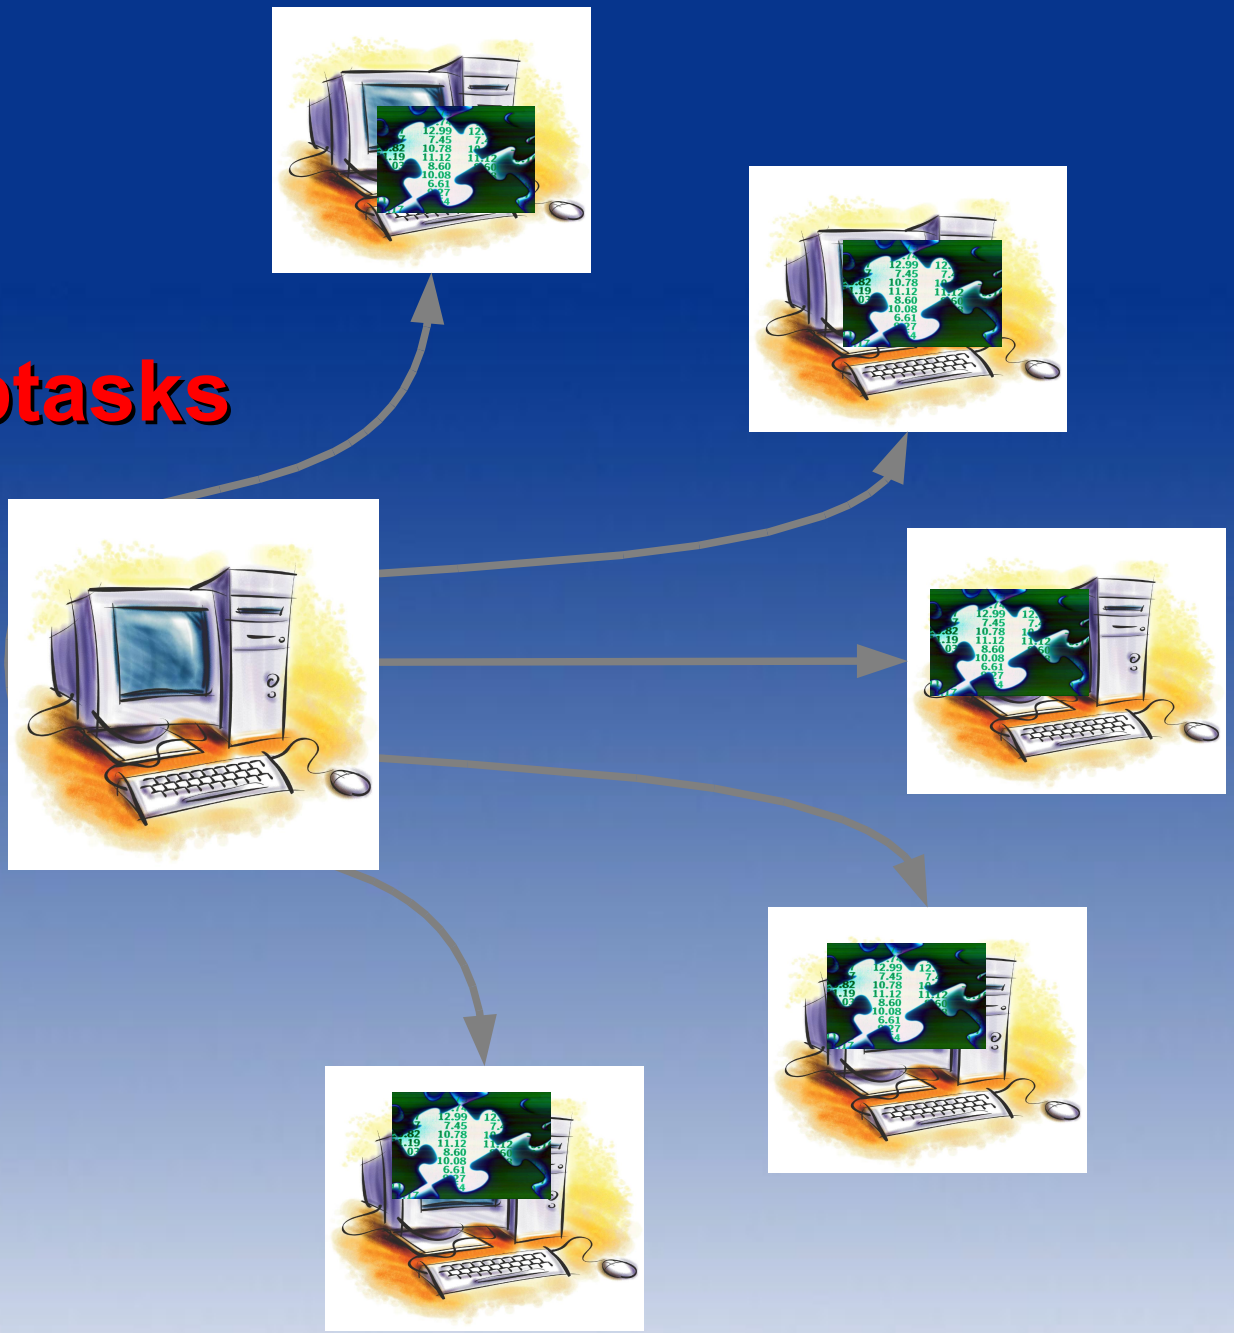

Where to find the computers?????

Let's build a cluster!!!

Let's build a cluster!!!

Let's use several clusters

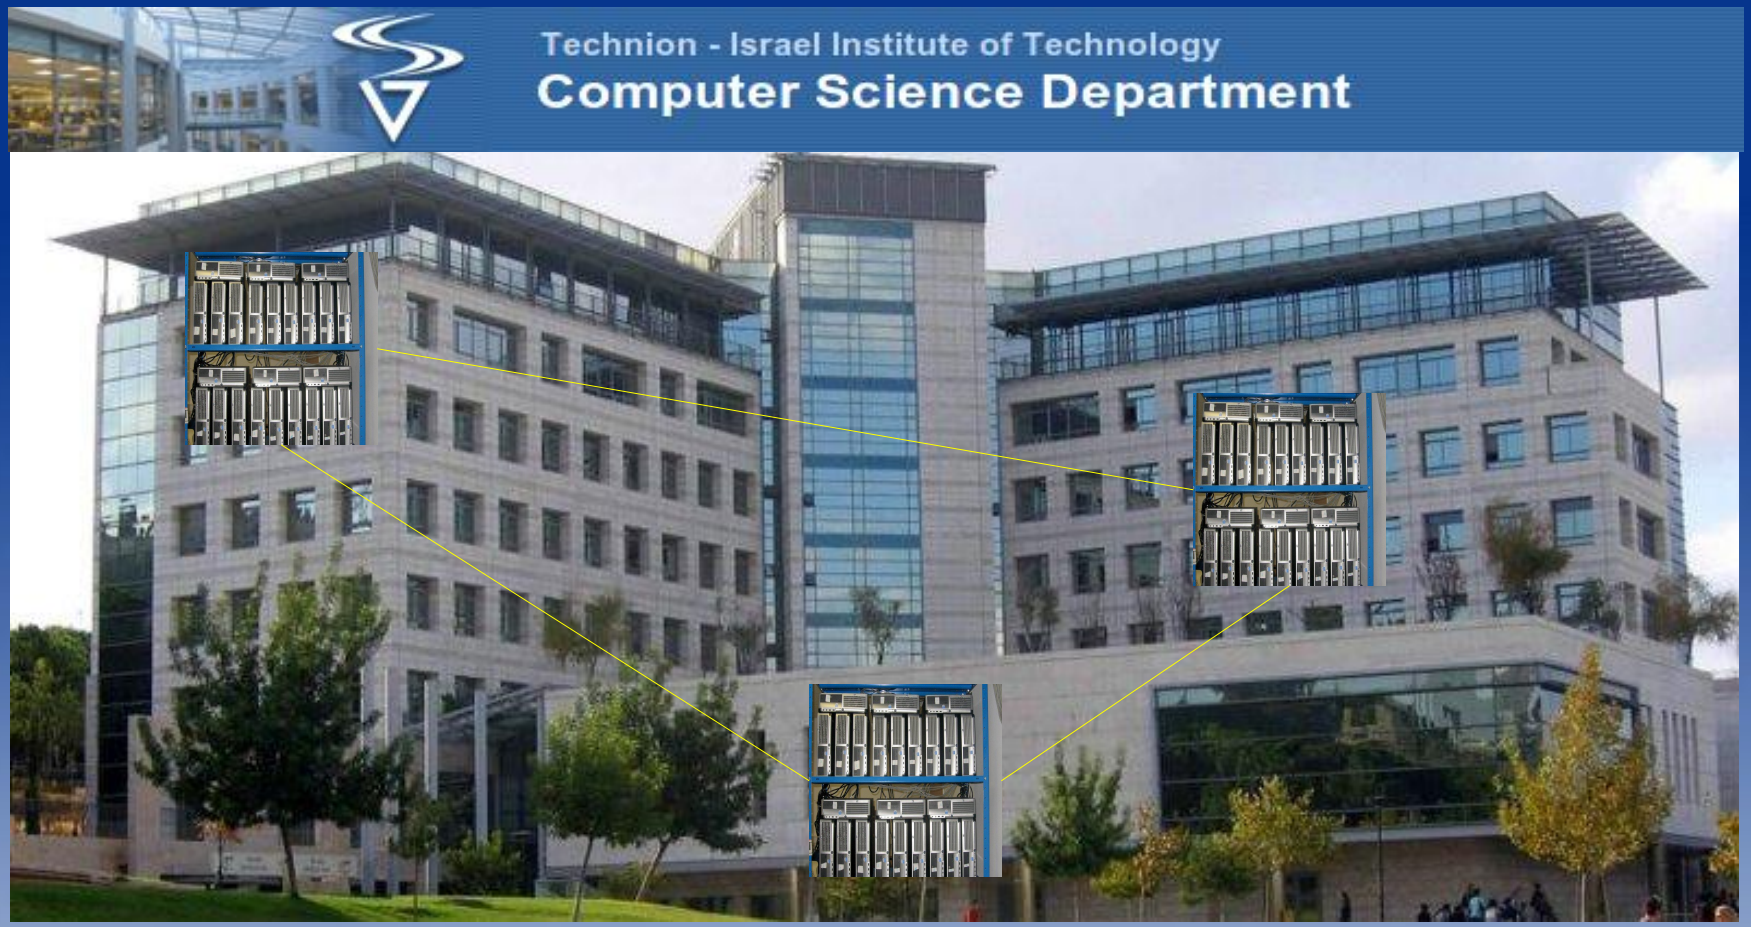

Let's use several clusters

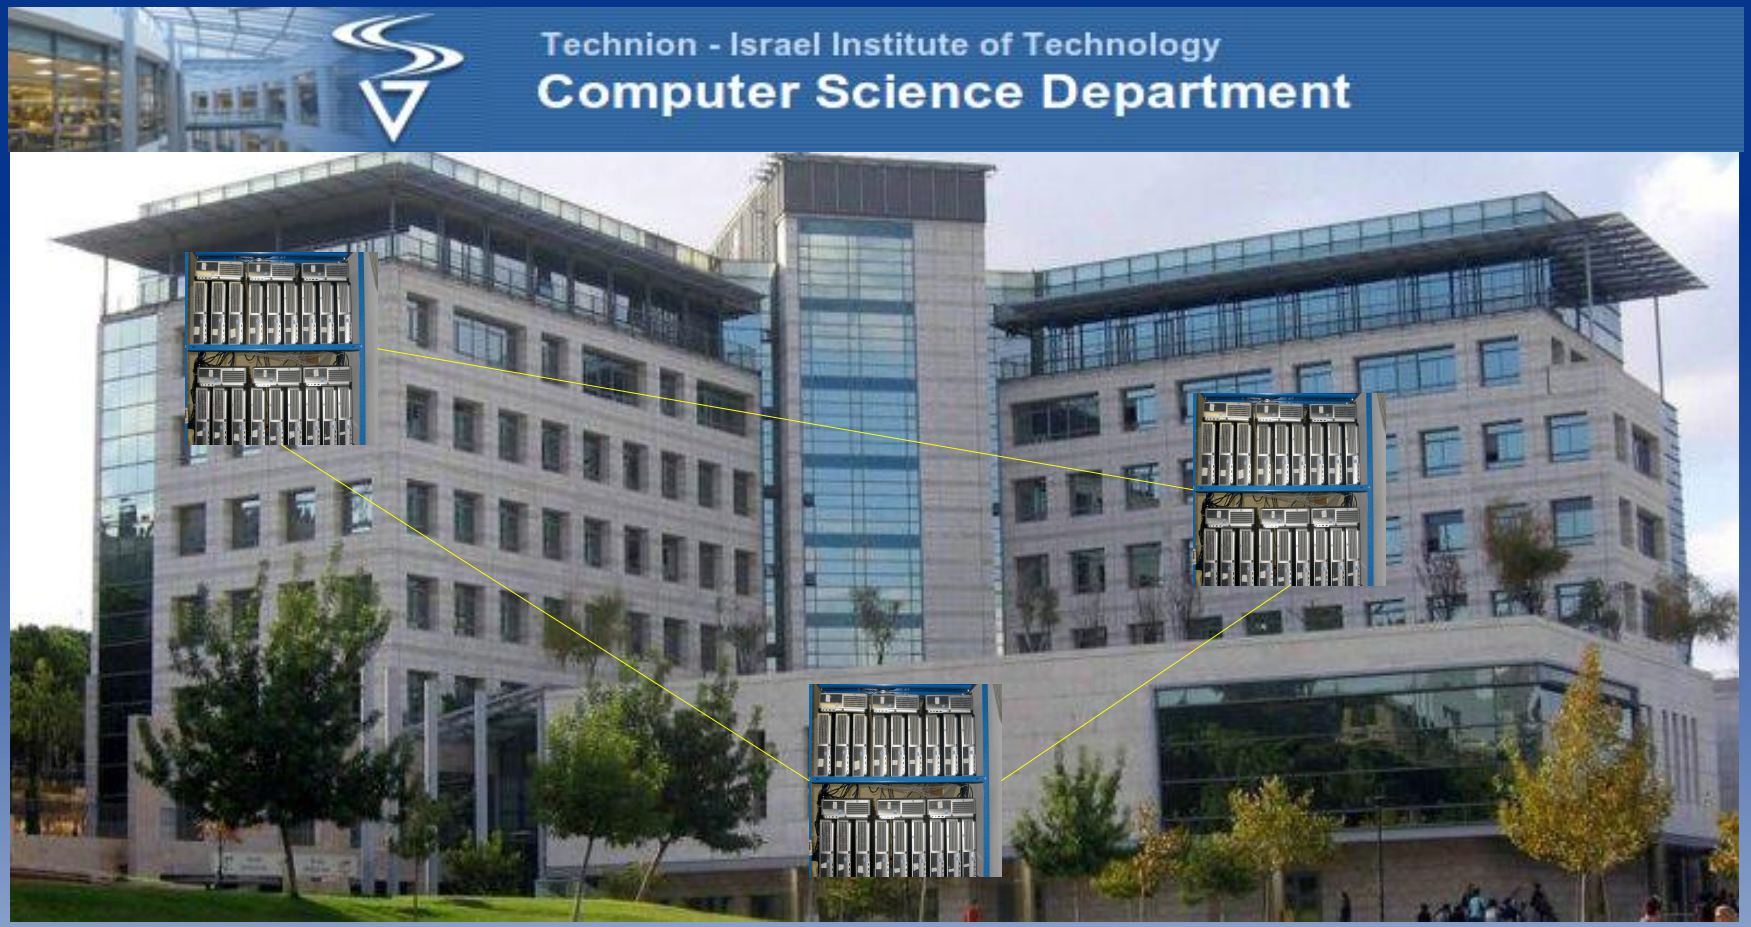

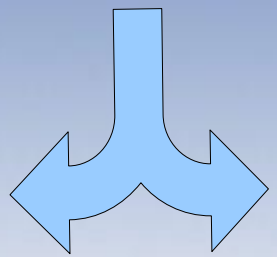

Let's also ask for help from

Let's also ask for help from

@

THE UNIVERSITY of WISCONSIN MADISON

Let's also ask for help from @HOME community!

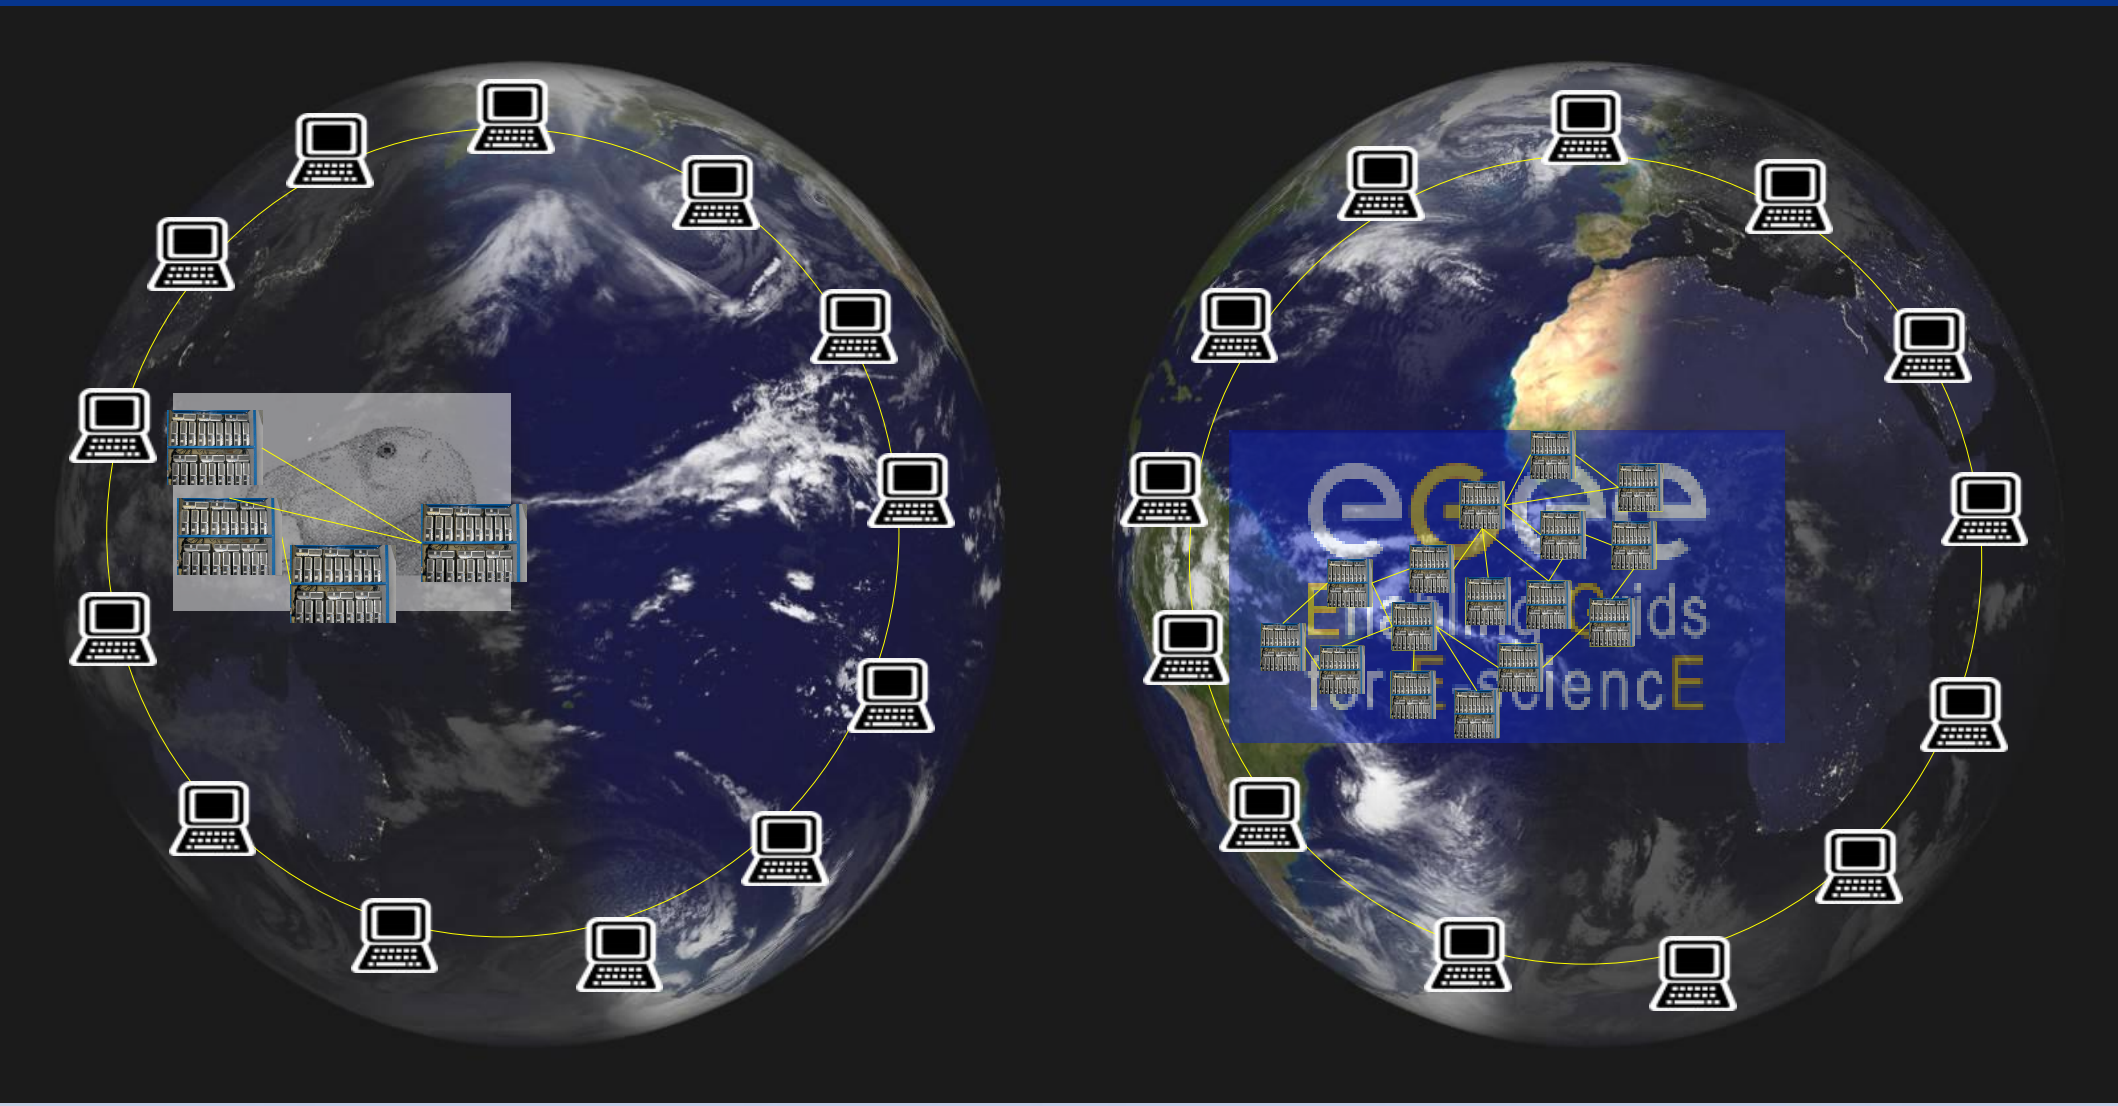

Can we glue these grids together ?

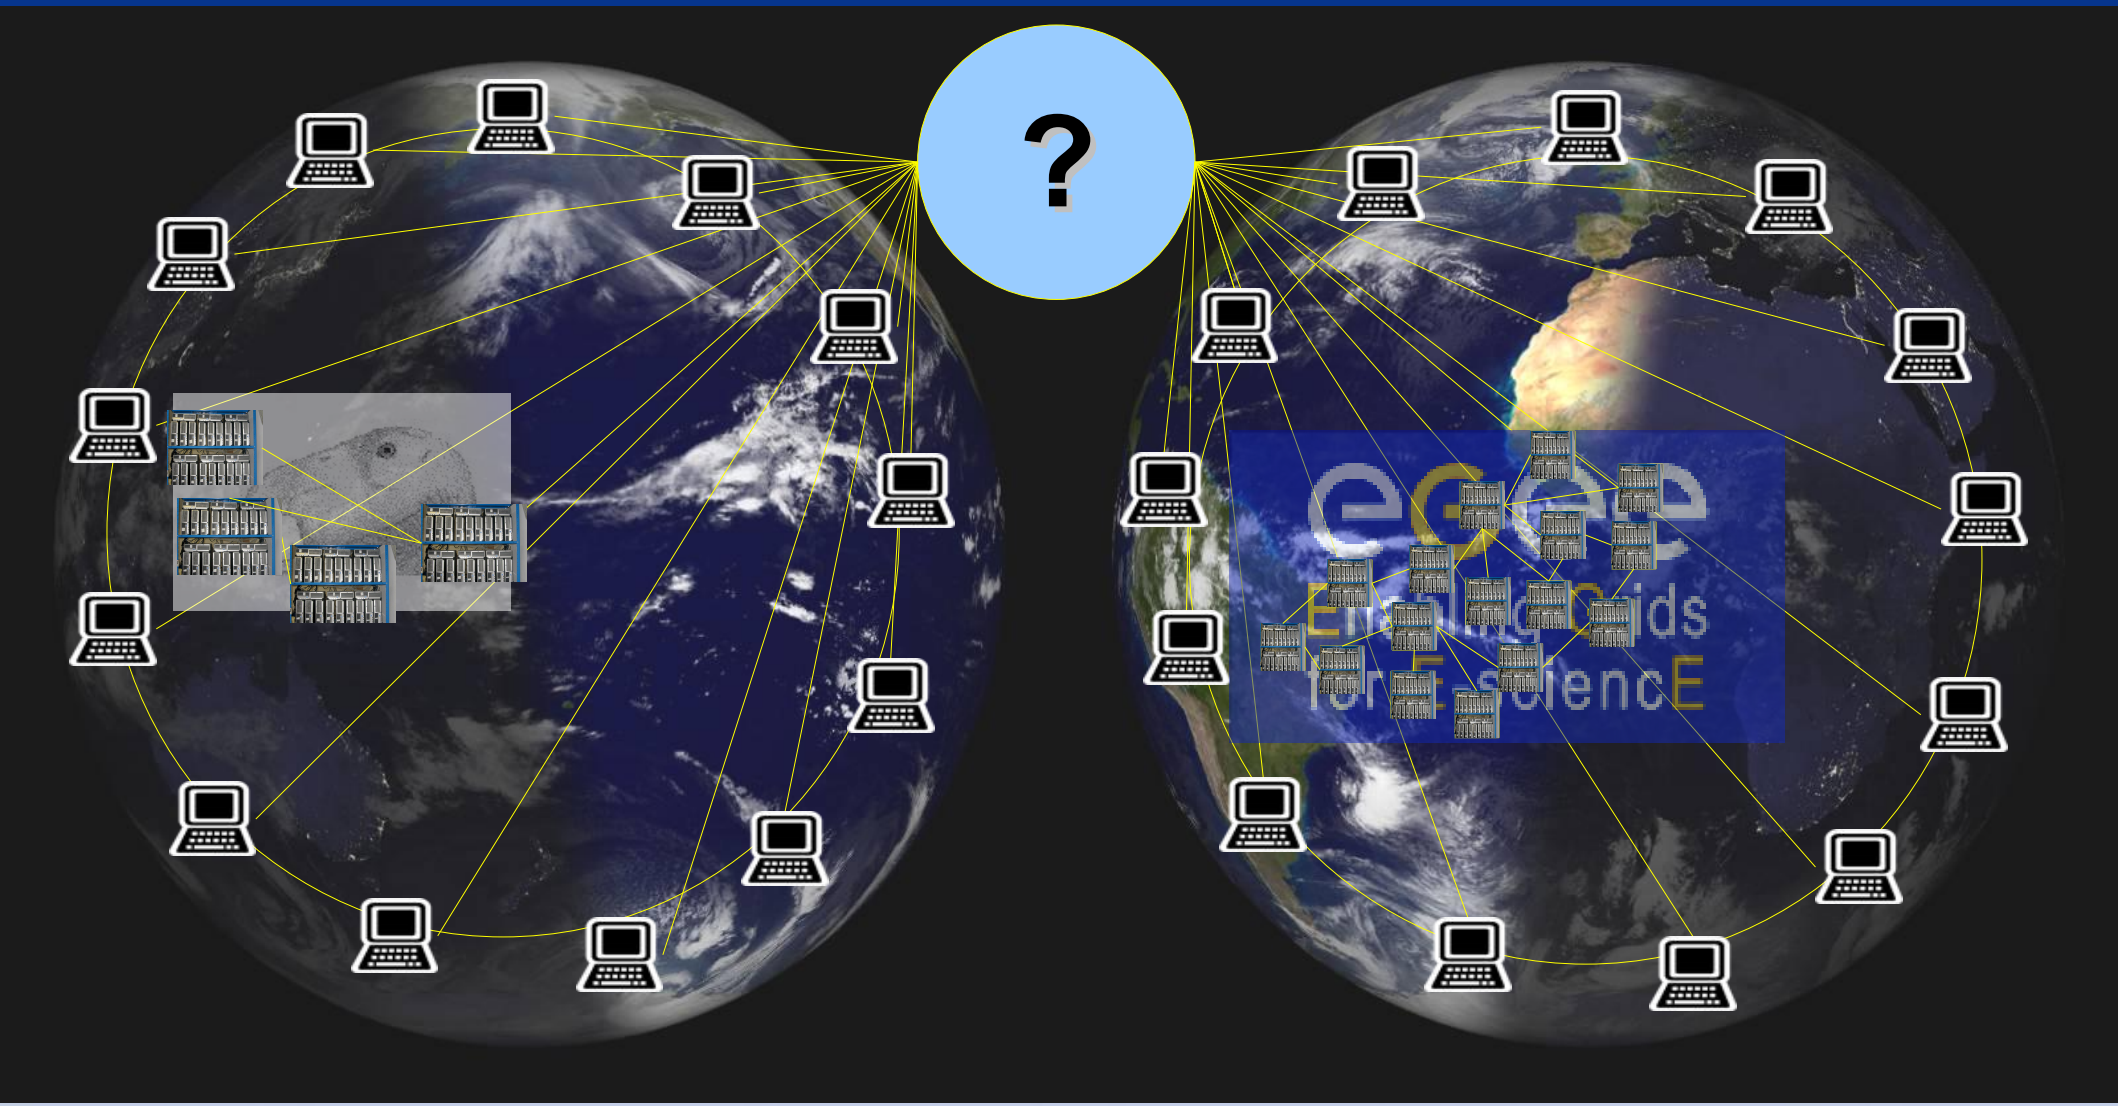

Can we glue these grids together ?

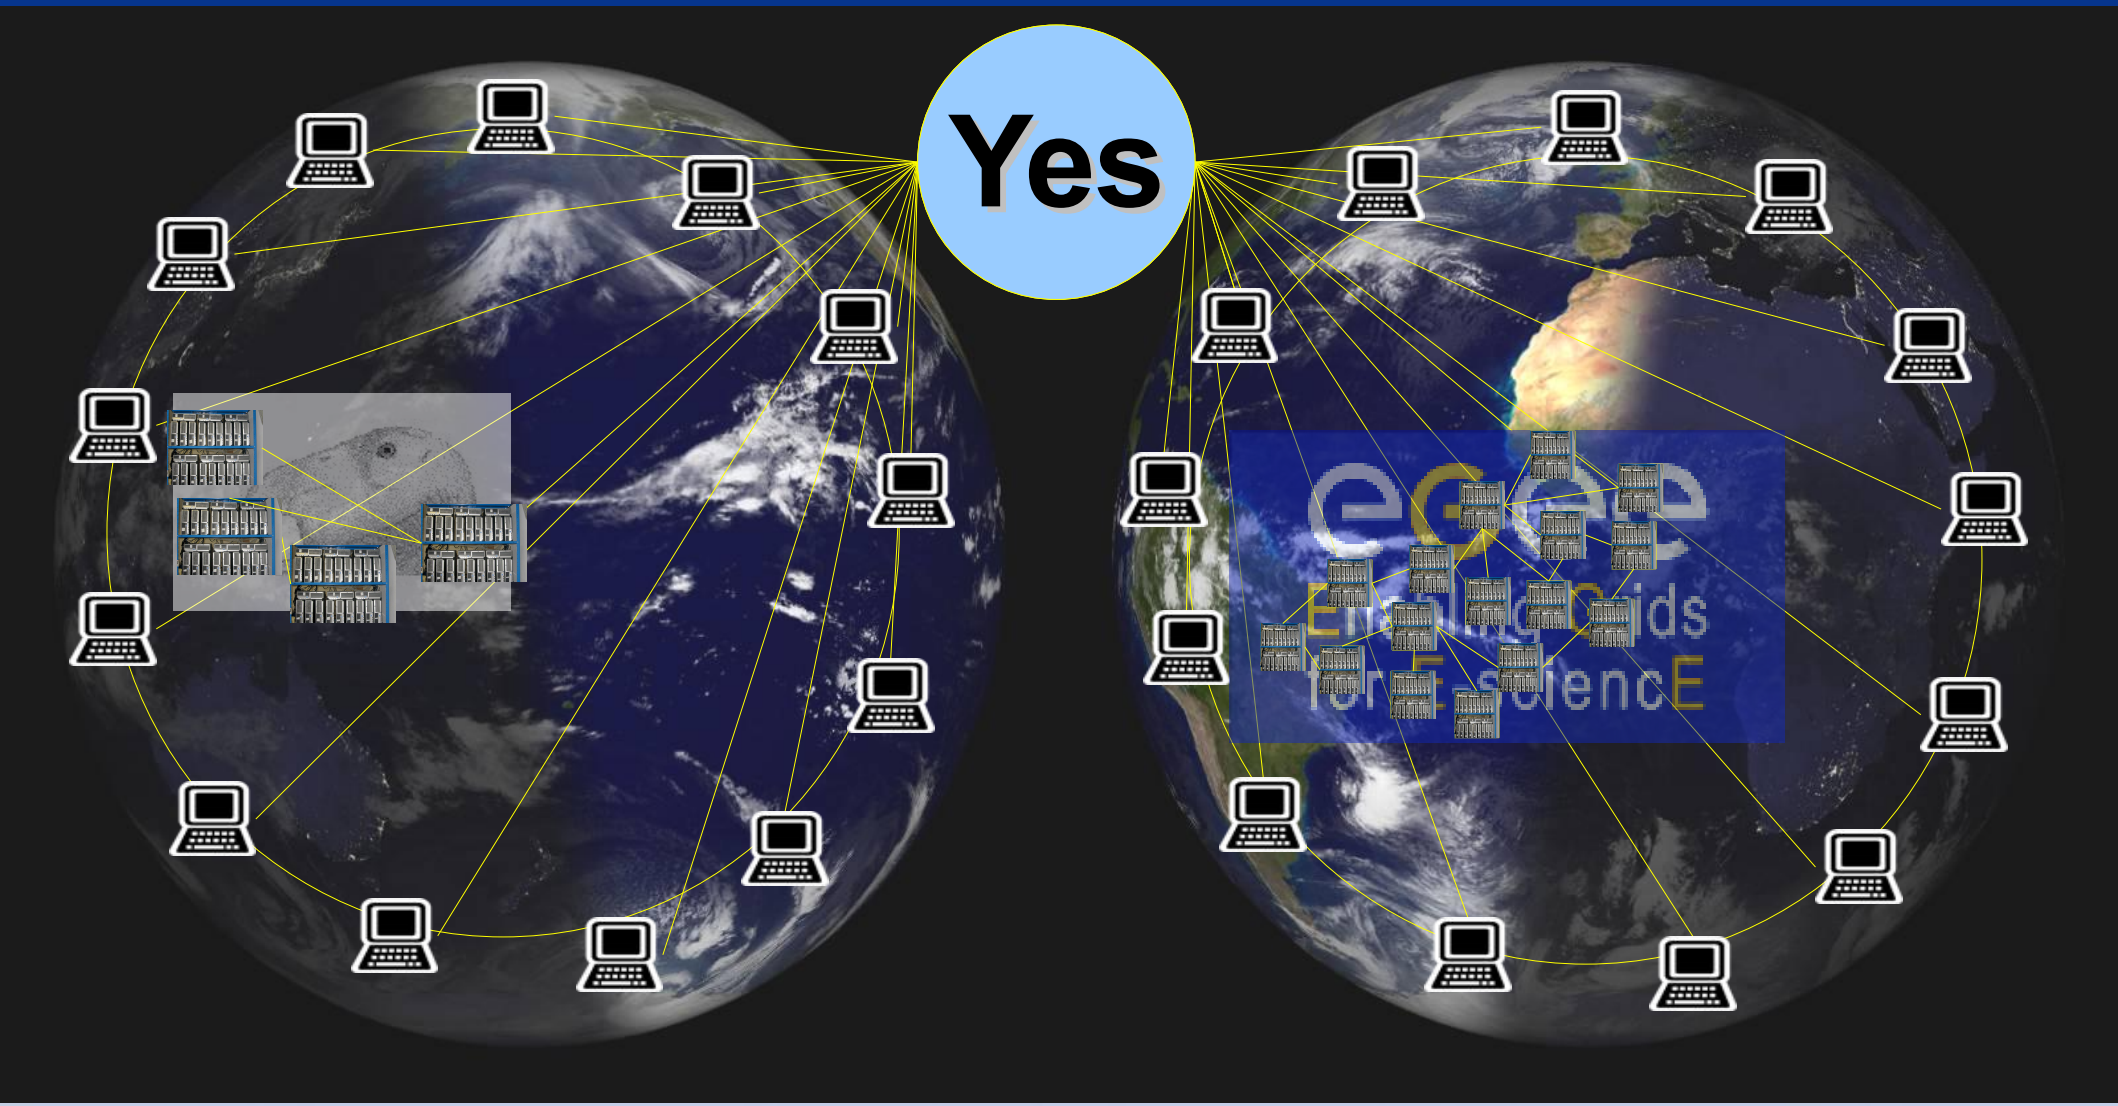

We also want to hide the complexity..

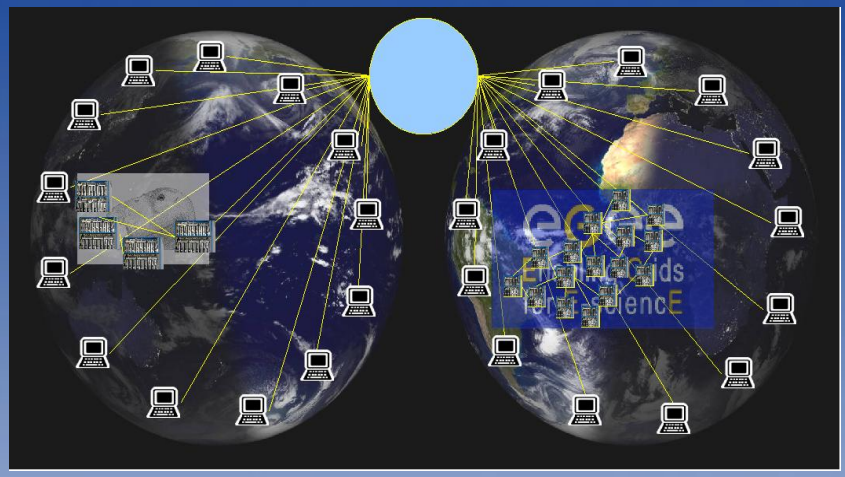

By building virtual supercomputer for domain researchers

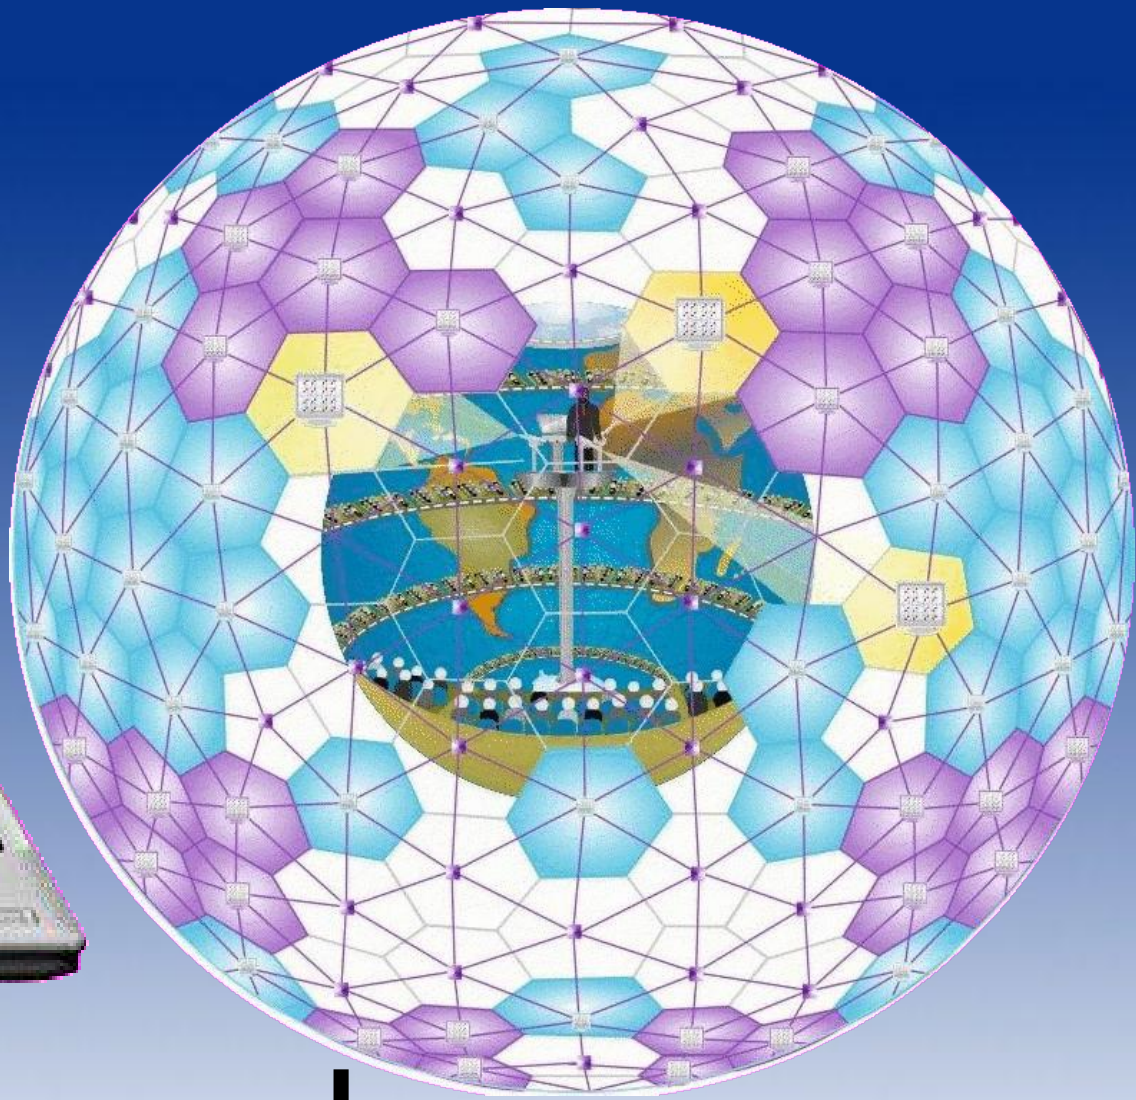

Virtual supercomputer for geneticists

Superlink-online: http://bioinfo.cs.technion.ac.il/superlink-online

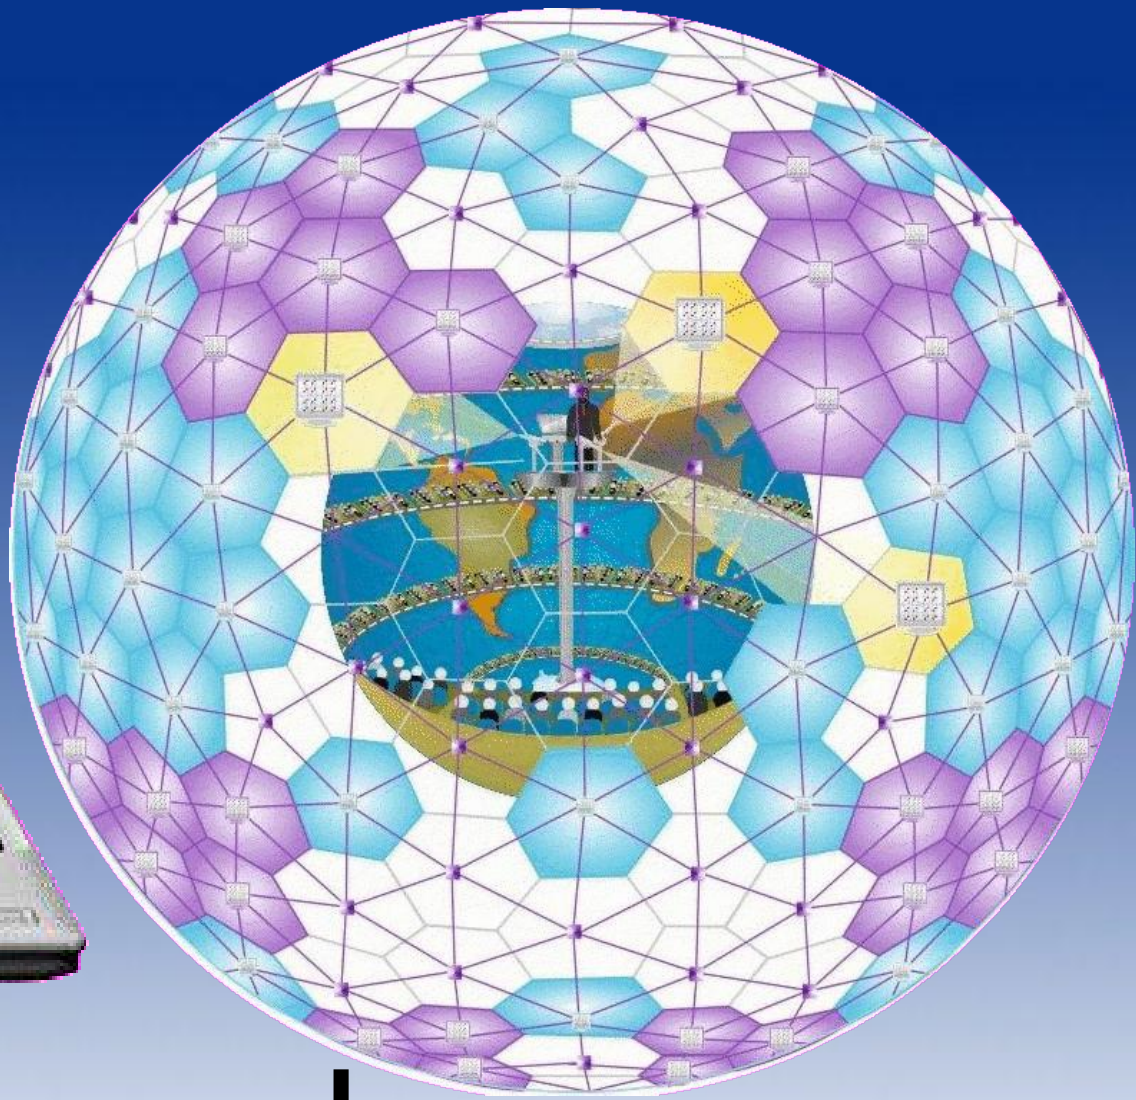

Main components

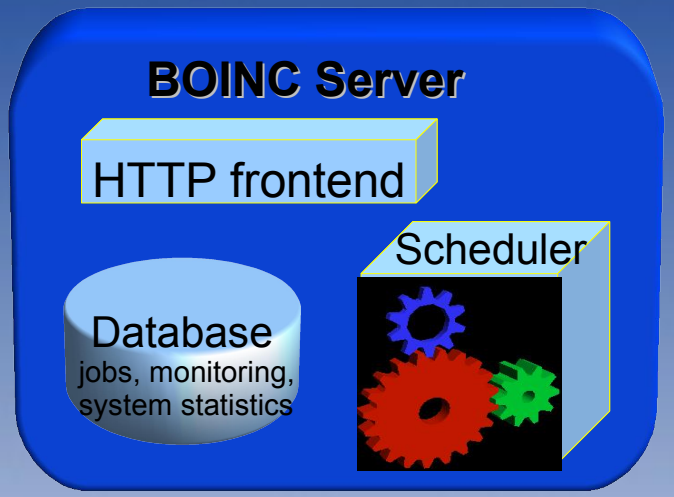

Main components

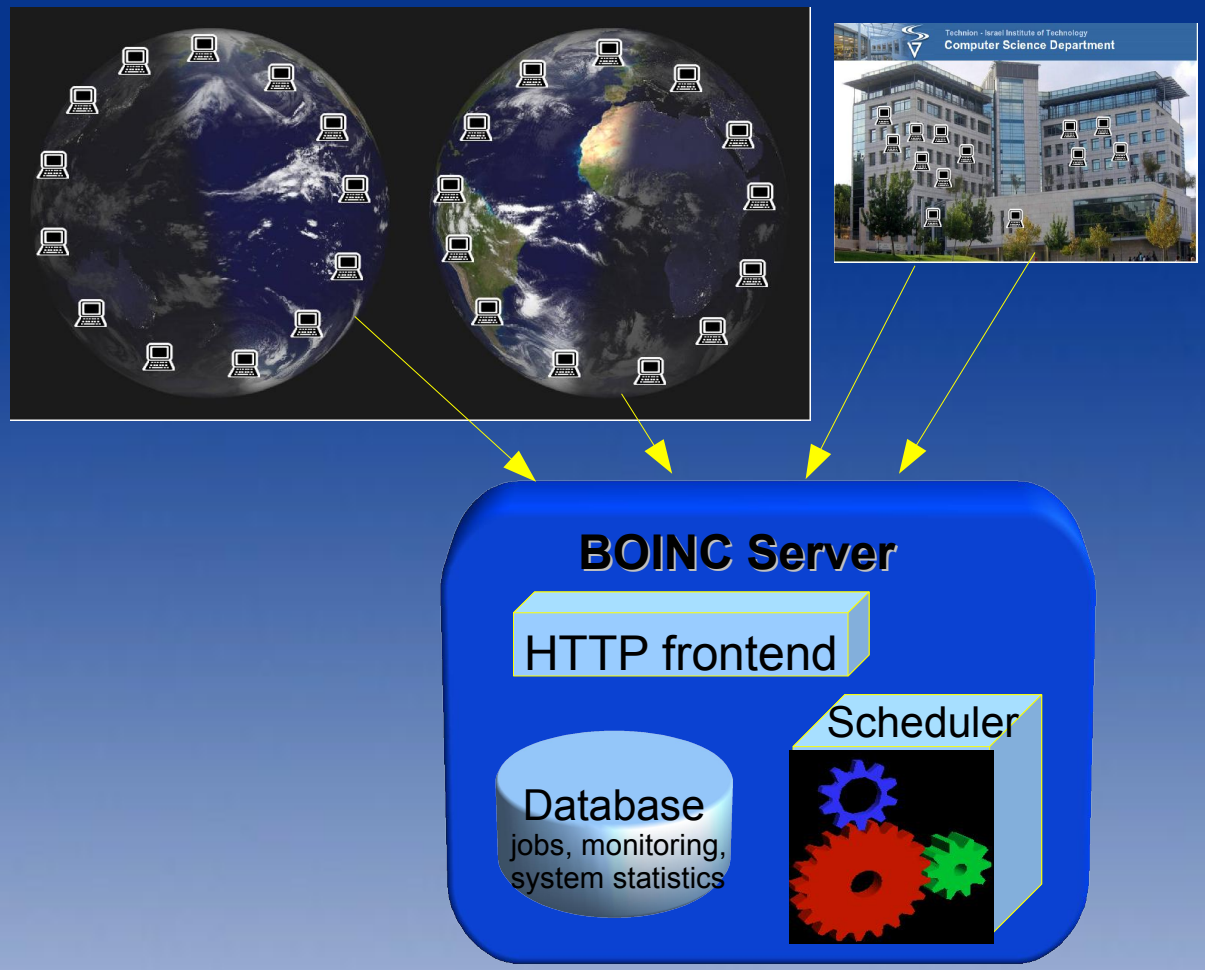

Main components

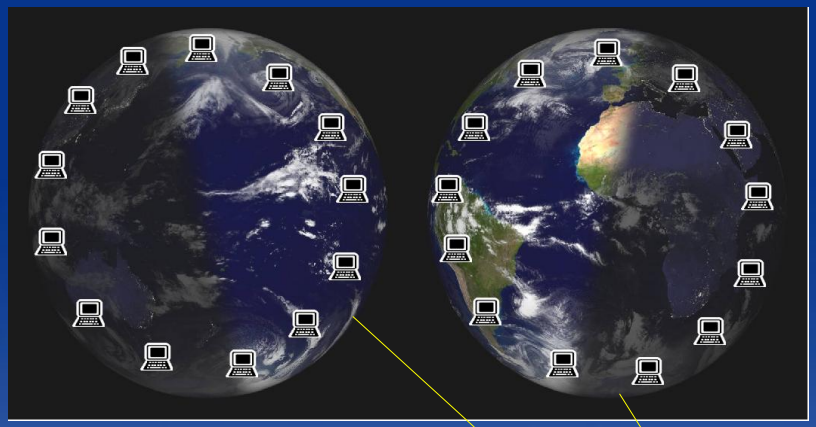

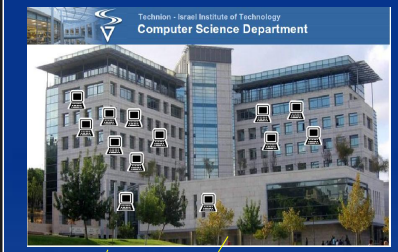

BOINC Server

HTTP frontend

Scheduler

Database
jobs, monitoring, system statistics

BOINC clients submittor for Condor in Madison

BOINC clients submittor for EGEE

Submittor to Technion Condor

Main components

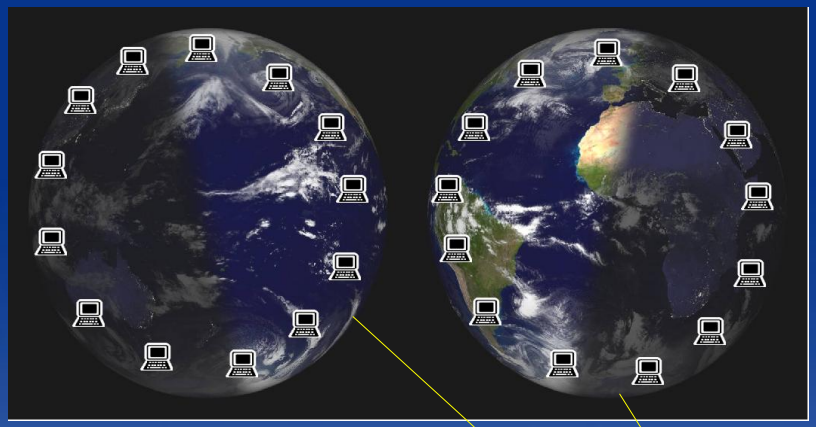

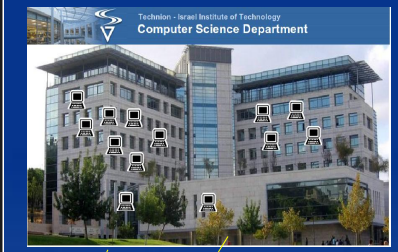

BOINC Server

HTTP frontend

Scheduler

Database
jobs, monitoring, system statistics

BOINC clients submittor for Condor in Madison

BOINC clients submittor for EGEE

Submittor to Technion Condor

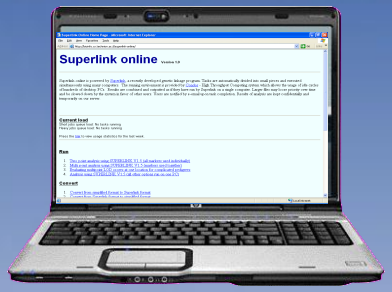

Web Portal

Task state

Main components

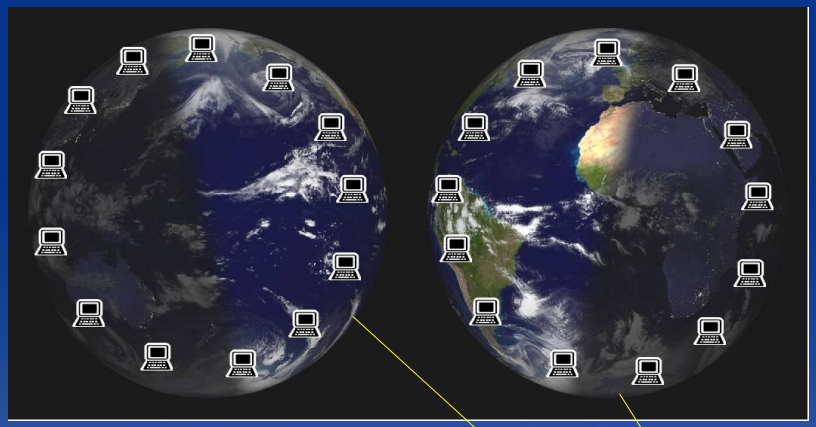

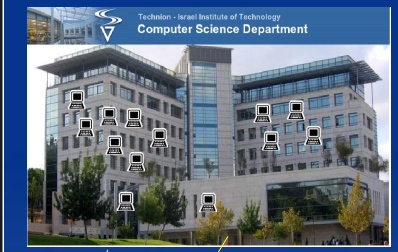

BOINC Server

HTTP frontend

Scheduler

Database
jobs, monitoring, system statistics

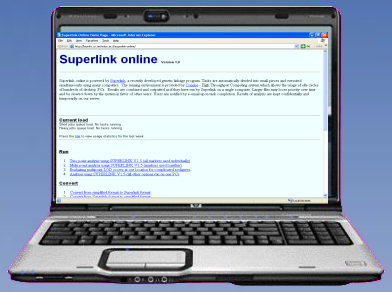

Web Portal

Task state

Task execution and monitoring workflow

BOINC clients submittor for Condor in Madison

BOINC clients submittor for EGEE

Submittor to Technion Condor

Main components

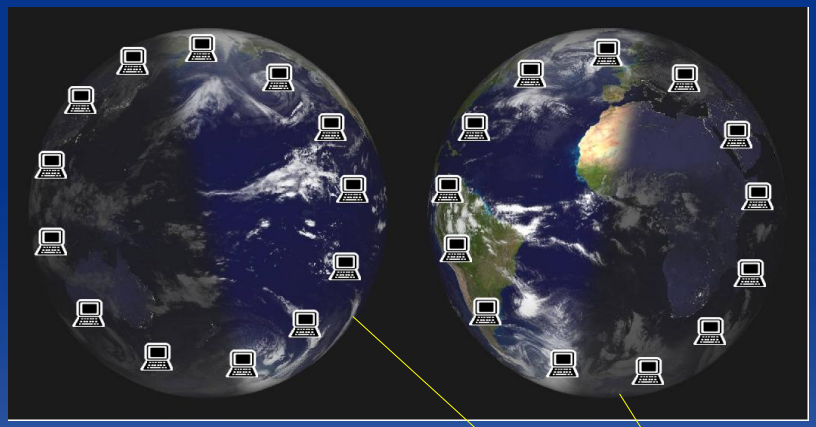

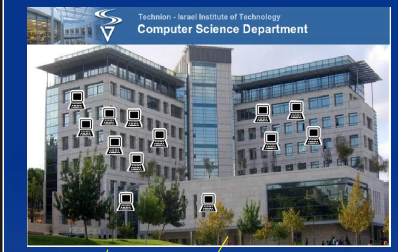

BOINC Server

HTTP frontend

Scheduler

Database
jobs, monitoring, system statistics

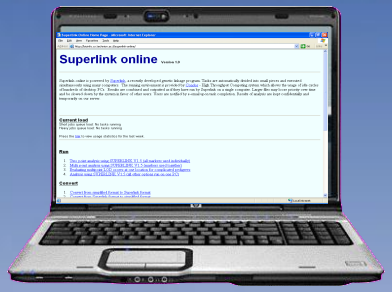

BOINC clients submittor for Condor in Madison

BOINC clients submittor for EGEE

Virtual cluster maintainer

Submittor to Technion Condor

Web Portal

Task state

Task execution and monitoring workflow

Main components

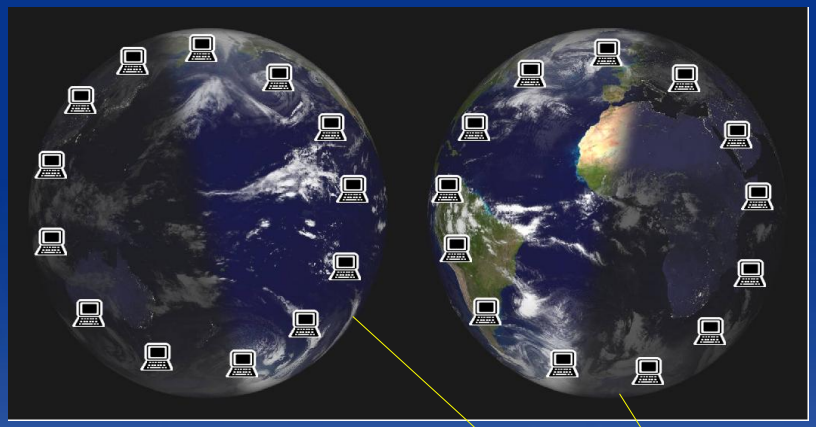

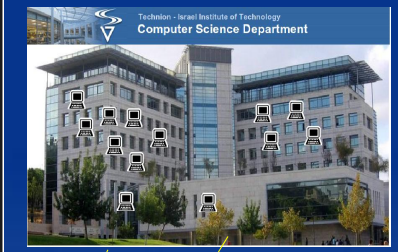

BOINC Server

HTTP frontend

Scheduler

Database
jobs, monitoring, system statistics

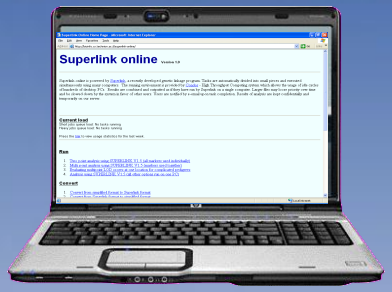

BOINC clients submittor for Condor in Madison

BOINC clients submittor for EGEE

Virtual cluster maintainer

Submittor to Technion Condor

Web Portal

Task state

Task execution and monitoring workflow

Dedicated cluster fallback

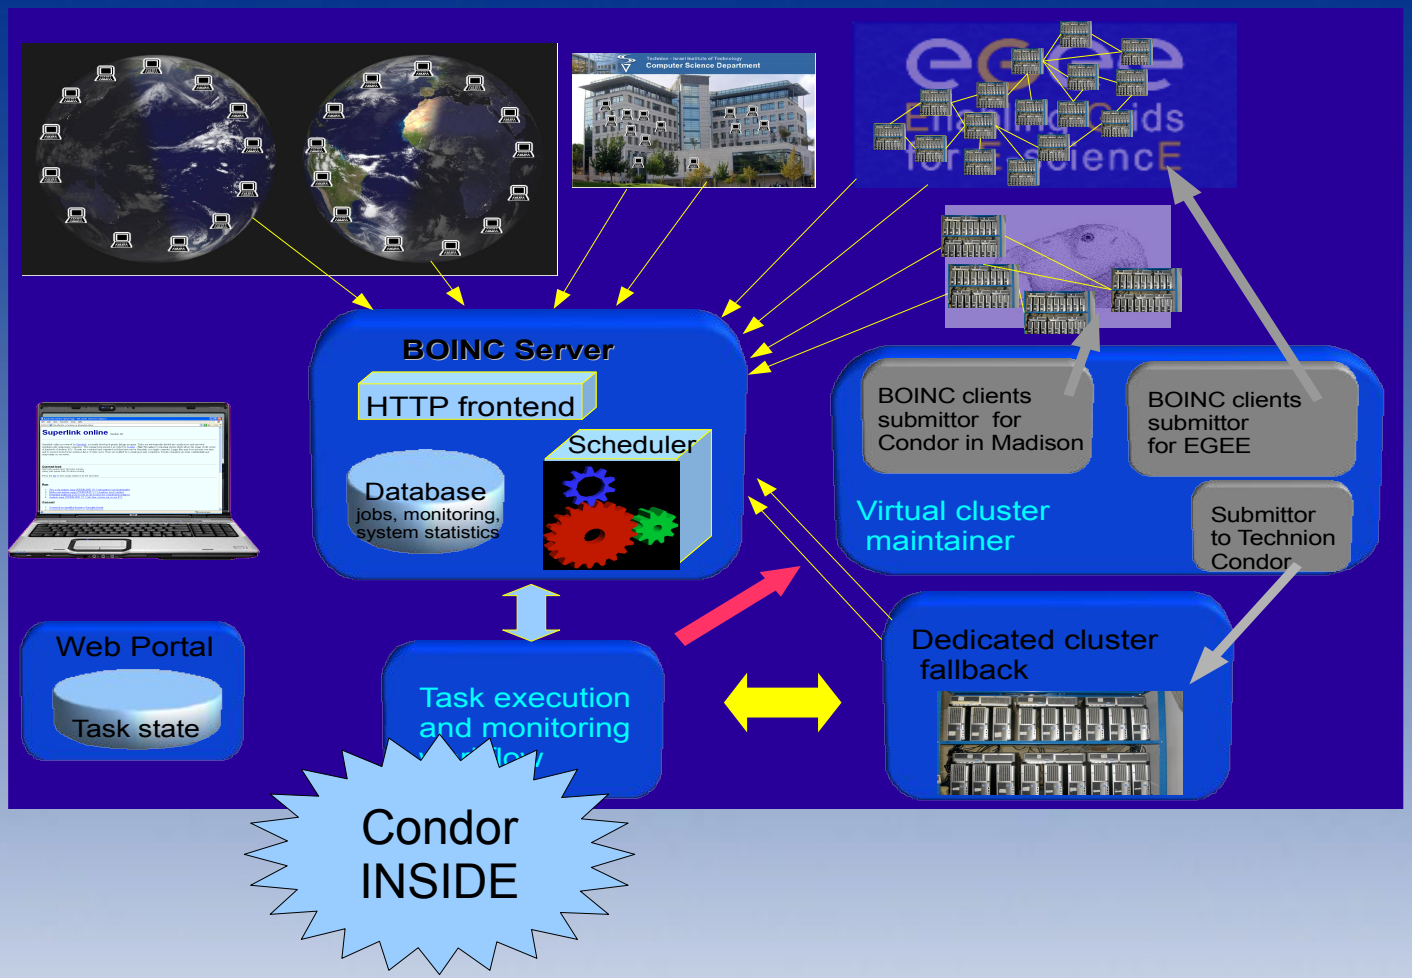

Condor DAGman reliability

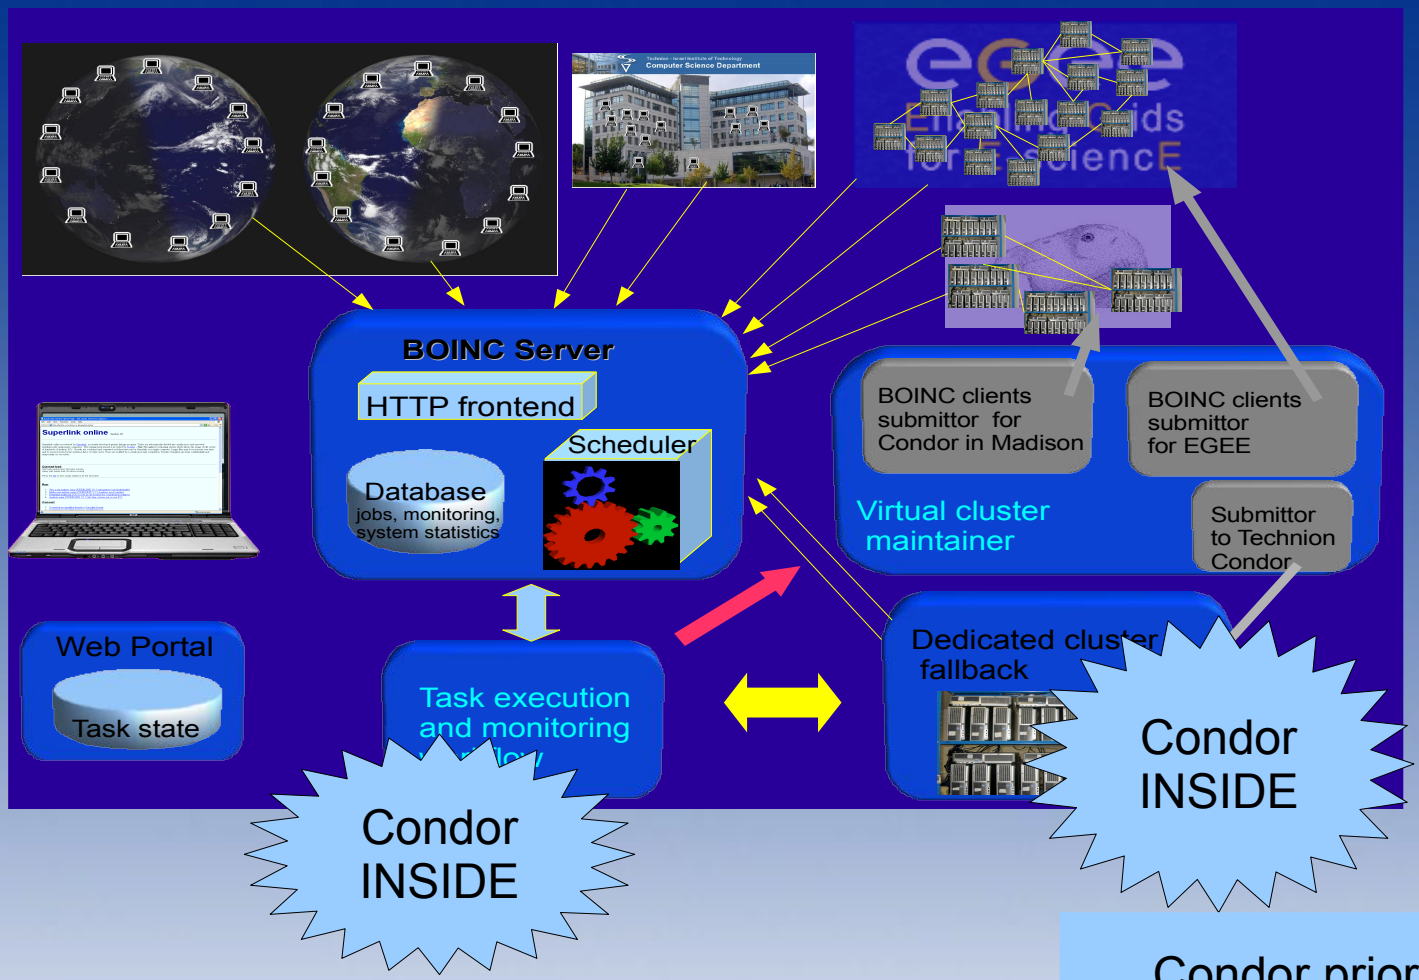

Condor DAGman reliability

Condor priority management

Condor STARTD as a stand-alone resource manager (Condor over BOINC)

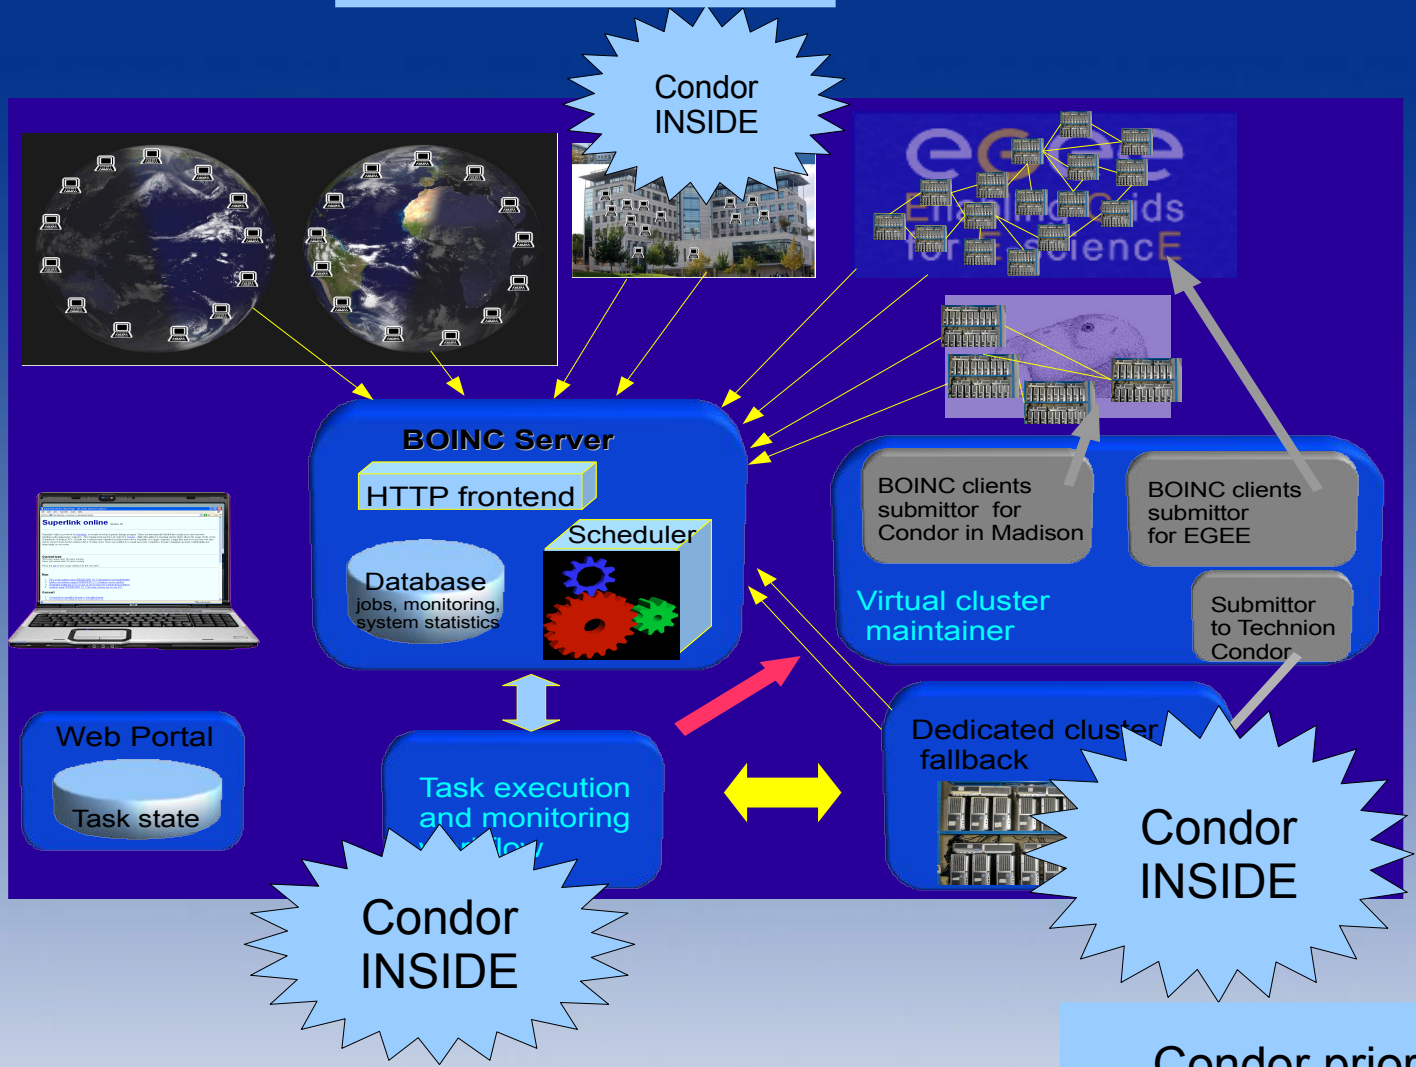

Condor DAGman reliability

Condor priority management

http://cbl-boinc-server2.cs.technion.ac.il/superlinkattechnion

Welcome to **Superlink@Technion!**

Superlink@Technion helps geneticists all over the world find disease-provoking genes causing some types of diabetes, hypertension (high blood pressure), cancer, schizophrenia and many others.
Press here to learn more

To contribute the idle cycles of your computer you need to **download** and install a small client, called BOINC. Please note that this project is still in **beta**, meaning that this site can be shut down for maintenance or work units may fail. We would appreciate your feedback via our forums.

Join now!

Project status	
Account creation	enabled
Users with credit	1812
Hosts with credit	4130
CPUs with credit	8663
Days of CPU utilized	21827.68
Detailed information	charts
Server state	state

System Requirements

- **RAM:** 1GB of RAM is suggested, as the computation typically requires up to 700MB of memory.
- **Disk space:** insignificant.
- **Supported platforms:** Windows (Vista/XP/2000/NT/Longhorn), Linux (Intel/PowerPC, 32/64 bits) and MacOS (PowerPC and Intel).
- **BOINC client version:** 5.8.x and higher is most preferable

User of the day

Sysadm
a frängischer Jung aus Närnberch

Die Römer wussten bereits Bescheid:...

Superlink@clusters http://cbl-boinc-server1.cs.technion.ac.il/superlinkatclusters

This project is a part of Superlink-online genetic linkage analysis research, conducted jointly by Distributed Systems Lab and Computational Biology Lab. Superlink@clusters serves as a submission platform to clusters at EGEE, UW Madison and Technion, as well as for the use of idle cycles of the computers at the CS department at the Technion.
If you are using one of the BOINC-enabled computers at the CS department and encounter any problem, please contact **Mark Silberstein** via e-mail. Any feedback will be much appreciated.

Project status	
Account creation	disabled
Participants total	7
Computers total	23542
Days of CPU utilized	18087.33
Global charts	charts
EGEE charts	charts
Technion Condor charts	charts
Madison Condor charts	charts
Server state	state

News

Jan 16, 2008 **Merging hosts by name**
Similar hosts are being merged on a daily basis

Oct 22, 2007 **Start of the project**
Superlink@clusters project started

News is available as an RSS feed XML.

Preliminary results

- **110 CPU Years consumed and a few millions jobs completed in 4 months ~ 350 CPUs working for us around the clock**
  - 49 (clusters: EGEE, UW Madison, Technion CS > 20,000 CPUs)
  - 61 (Contributors of Superlink@Technion > 8000 CPUs)

Top participants

Superlink@Technion

Rank	Name	Recent average credit	Total credit	CPU time used (hh:mi:ss)	Country	Participant since
1	zombie67	3,454.76	280,698.87	19079:54:20	United States	25 Jun 2007 4:07:20 UTC
2	The Swordfish	2,781.78	223,758.97	13495:21:38	United States	12 Nov 2007 5:20:42 UTC
3	tng	2,770.58	645,542.36	43501:19:14	United States	11 Jul 2007 12:54:55 UTC
4	Steve of Boinc UK	2,359.91	155,414.76	6239:46:02	United Kingdom	4 Nov 2007 5:16:28 UTC
5	VC/HVH	2,279.92	94,535.72	4621:39:10	Germany	14 Mar 2008 23:37:22 UTC
6	smoked_trout	2,098.63	38,418.90	1583:35:58	United States	17 Apr 2008 16:00:41 UTC
7	thelmores	1,669.28	35,921.30	979:53:26	United States	17 Apr 2008 1:44:15 UTC
8	biancaw	1,559.84	47,191.36	2516:25:01	Germany	31 Mar 2008 19:10:43 UTC
9	Kloedhilf	1,464.10	34,378.61	1443:11:38	Czech Republic	1 Apr 2008 18:14:49 UTC
10	tomba	1,380.41	89,284.65	5567:17:28	United States	13 Jan 2008 17:10:30 UTC
11	UBT-Timby	1,032.69	14,264.51	761:34:54	United Kingdom	30 Jun 2007 19:08:56 UTC
12	ANCHULA-MARK	976.76	22,500.27	2031:35:01	United Kingdom	2 Mar 2008 14:37:28 UTC
13	fellie	861.52	34,276.93	1939:59:33	United Kingdom	15 Mar 2008 22:35:16 UTC
14	BobCat13	810.72	51,659.19	3051:42:48	United States	20 Jun 2007 0:38:20 UTC

Top participants

Superlink@Clusters

Rank	Name	Recent average credit	Total credit
1	egee	19,669.21	2,607,640.88
2	technioncondor	3,580.51	412,529.58
3	farm	2,277.24	272,766.05
4	madisoncondor	2,259.92	274,493.93
5	lccn	805.61	201,605.33
6	dsl	803.94	48,187.52
7	assistant	500.92	51,887.40

Top participants

Superlink@Technion

~CPU time

Rank	Name		Total credit	CPU time used (hh:mi:ss)	Country	Participant since
1	zombie67	3,454.76	280,698.87	19079:54:20	United States	25 Jun 2007 4:07:20 UTC
2	The Swordfish	2,781.78	223,758.97	13495:21:38	United States	12 Nov 2007 5:20:42 UTC
3	tng	2,770.58	645,542.36	43501:19:14	United States	11 Jul 2007 12:54:55 UTC
4	Steve of Boinc UK	2,359.91	155,414.76	6239:46:02	United Kingdom	4 Nov 2007 5:16:28 UTC
5	VC/HVH	2,279.92	94,535.72	4621:39:10	Germany	14 Mar 2008 23:37:22 UTC
6	smoked_trout	2,098.63	38,418.90	1583:35:58	United States	17 Apr 2008 16:00:41 UTC
7	thelmores	1,669.28	35,921.30	979:53:26	United States	17 Apr 2008 1:44:15 UTC
8	biancaw	1,559.84	47,191.36	2516:25:01	Germany	31 Mar 2008 19:10:43 UTC
9	Kloedhilf	1,464.10	34,378.61	1443:11:38	Czech Republic	1 Apr 2008 18:14:49 UTC
10	tomba	1,380.41	89,284.65	5567:17:28	United States	13 Jan 2008 17:10:30 UTC
11	UBT-Timby	1,032.69	14,264.51	761:34:54	United Kingdom	30 Jun 2007 19:08:56 UTC
12	ANCHULA-MARK	976.76	22,500.27	2031:35:01	United Kingdom	2 Mar 2008 14:37:28 UTC
13	fellie	861.52	34,276.93	1939:59:33	United Kingdom	15 Mar 2008 22:35:16 UTC
14	BobCat13	810.72	51,659.19	3051:42:48	United States	20 Jun 2007 0:38:20 UTC

Top participants

Superlink@Clusters

~CPU time

Rank	Name		Total credit
1	egee	19,669.21	2,607,640.88
2	technioncondor	3,580.51	412,529.58
3	farm	2,277.24	272,766.05
4	madisoncondor	2,259.92	274,493.93
5	lccn	805.61	201,605.33
6	dsl	803.94	48,187.52
7	assistant	500.92	51,887.40

Single framework for grid performance analysis

- Superlink@Technion environment is more predictable!

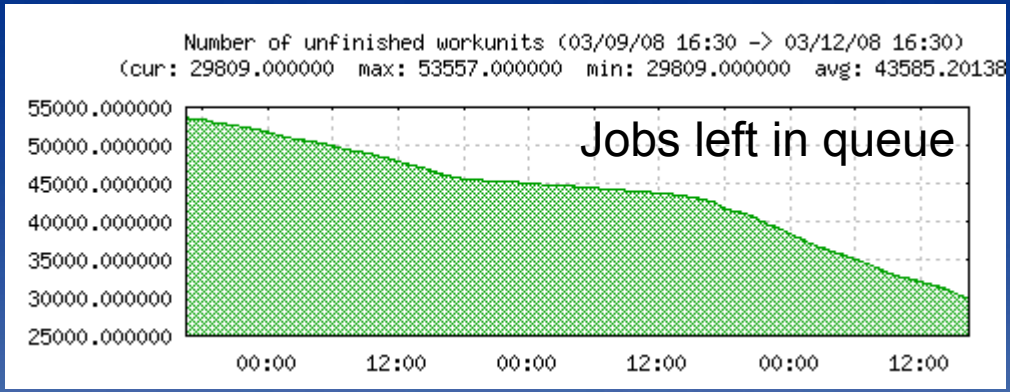

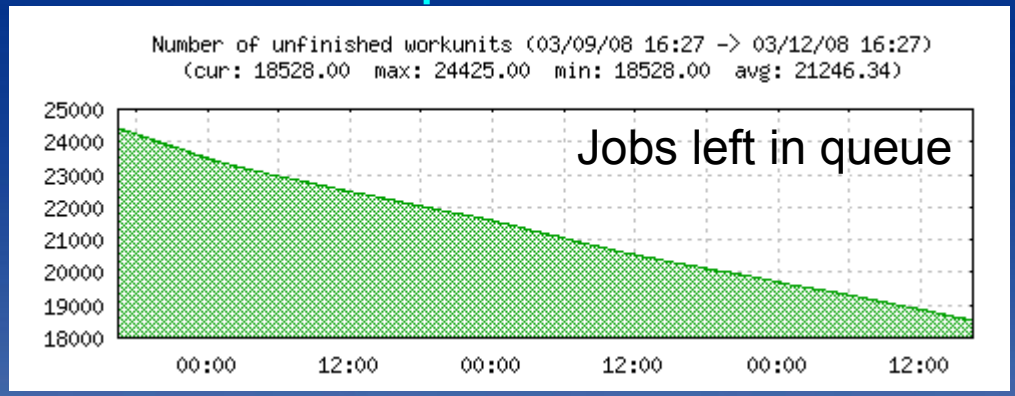

- Since clusters are much more volatile because of policy

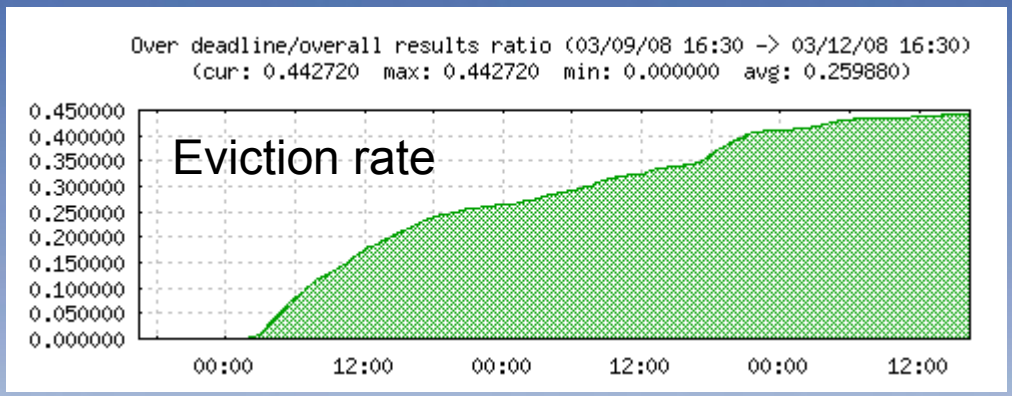

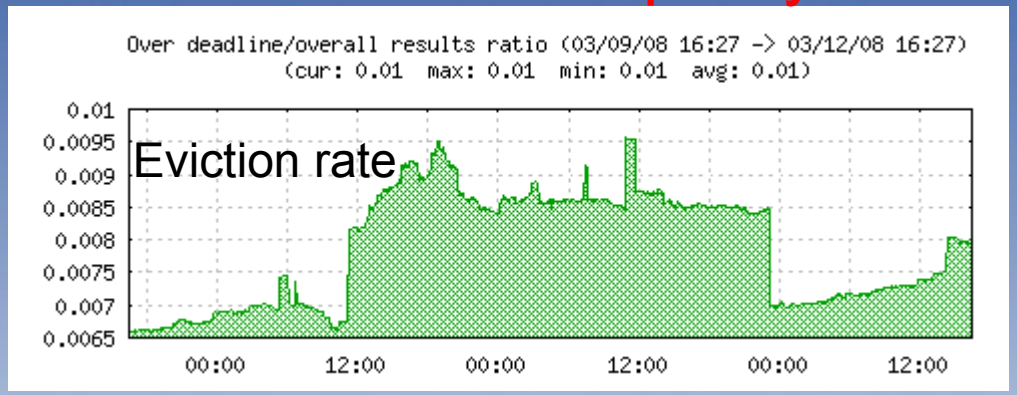

Single framework for grid performance analysis

- EGEE vs. Madison pool
  - EGEE behaves more as @Technion than Madison pool. Why?

High load, complex policies and centralized control result in rapid (statistically unpredictable) changes in resource allocation

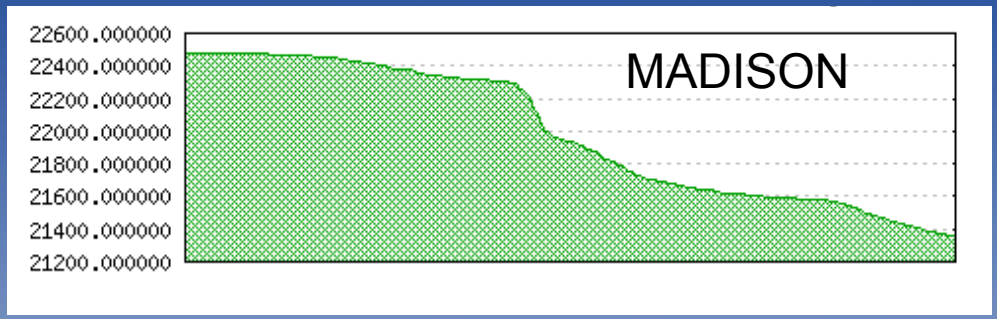

EGEE is much less centralized – evictions are less correlated, hence have steady state

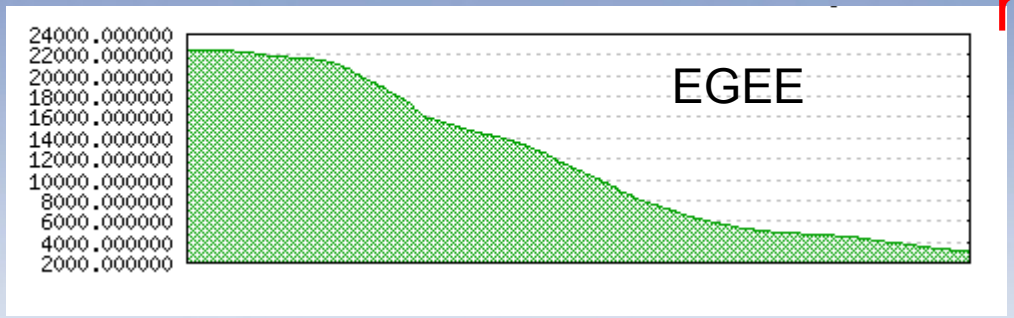

Lesson 1: Track the system performance from the application perspective

- Simple criterion - compare with the time it would take on a dedicated cluster

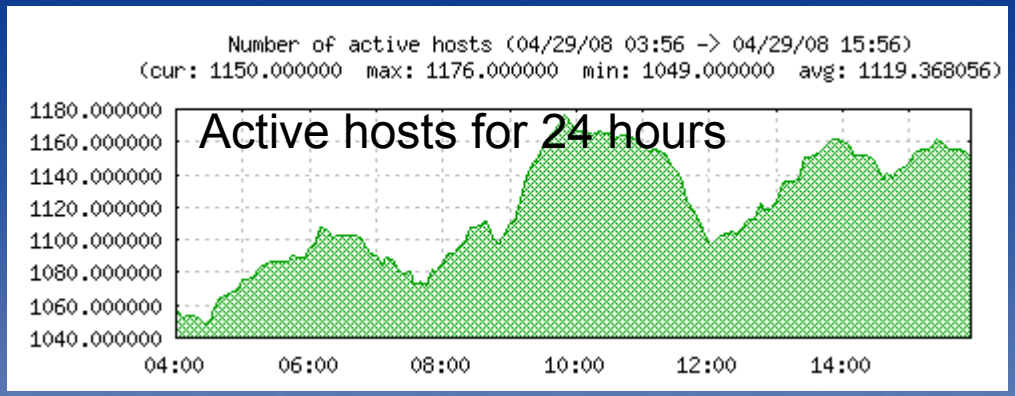

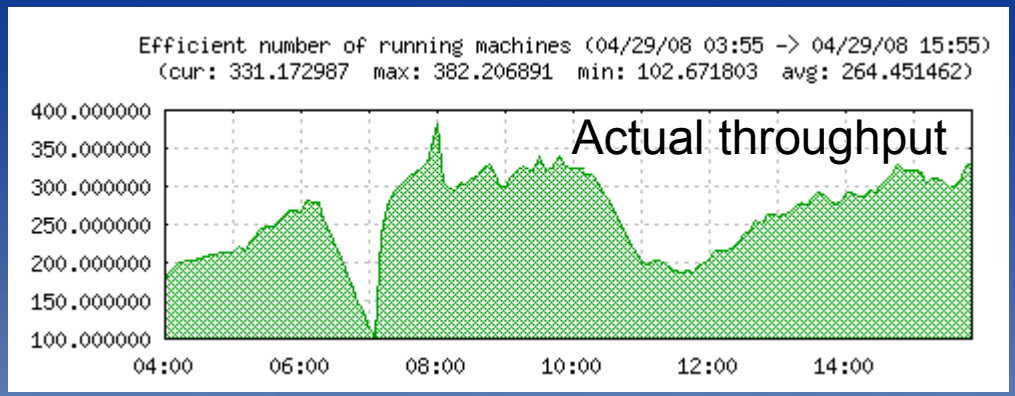

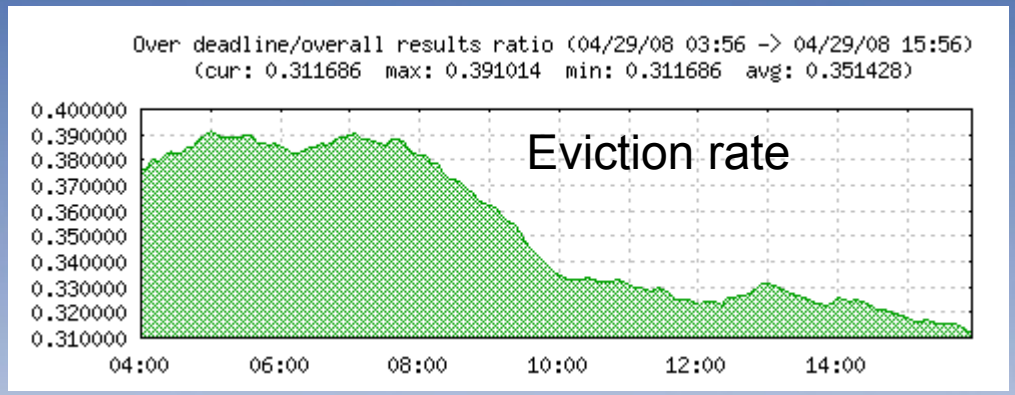

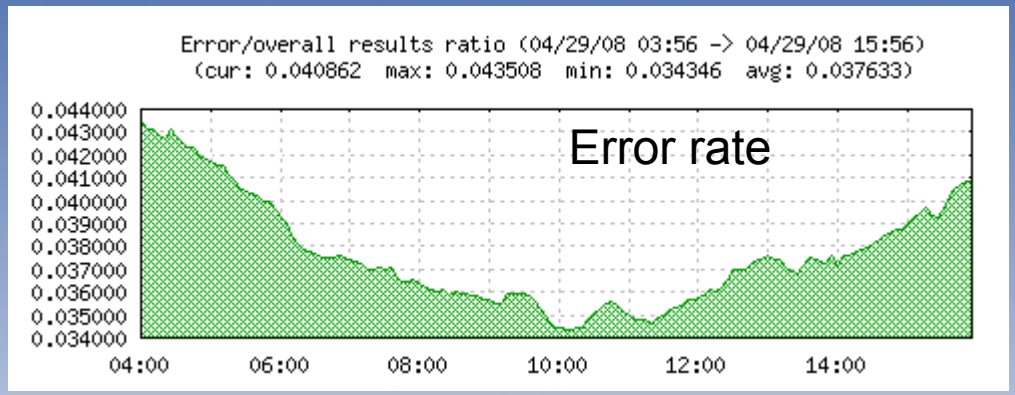

- Using BOINC (glide-ins) helps separate application performance from batch system performance

Lesson 2: Learn the system parameters

- Example: Average running time of a BOINC client until preempted

Wrong job granularity

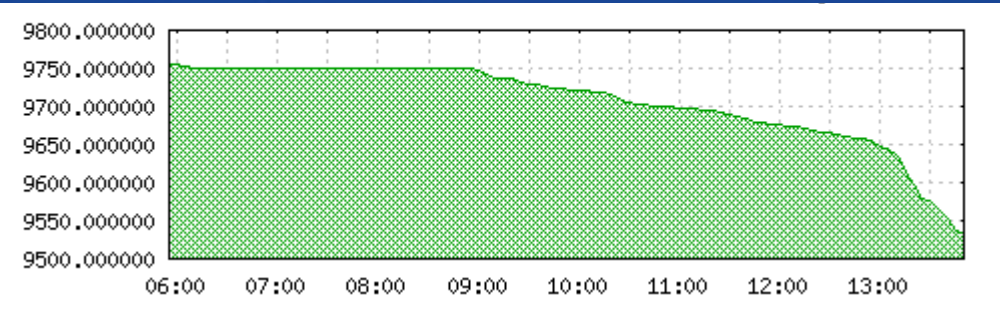

The same run with proper job granularity

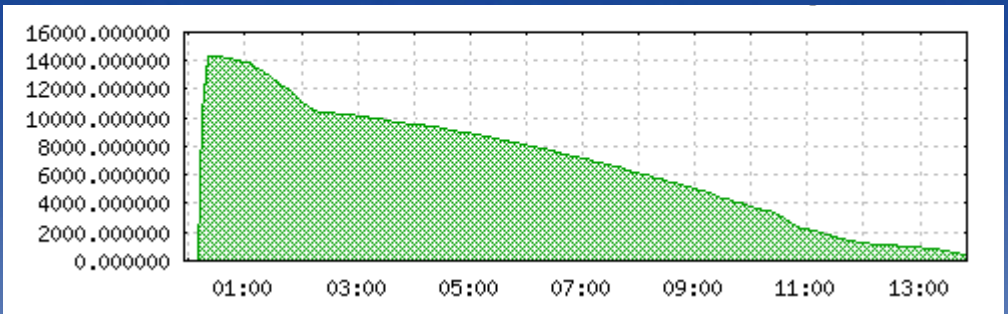

- Using BOINC allows to tailor job size with no additional overhead : few minutes long jobs run as fast as few hours long

Lesson 3: Distinguish between high throughput and high performance runs

- Many-jobs: high throughput runs – number of jobs much more than number of running machines
  - Less sensitive to failures - overlapped with the execution of other jobs
- Few-jobs: high performance runs – number of jobs is about the number of running machines
  - **Performance very sensitive to failures**

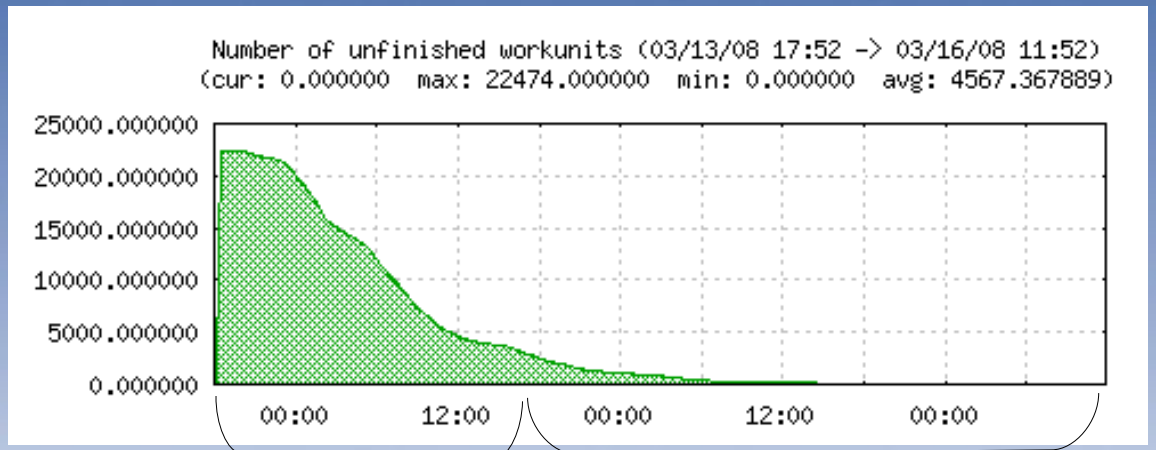

- Dedicated cluster comes handy!!!

Thinking out loud

- Will grids become less useful to opportunistic users when they become easy to use (higher utilized)?
- Should Condor policy be hostile to opportunistic users, or some guarantees are better to be provided?
- And if above is true, and you are going to have a lot of jobs – maybe it's better to buy your own cluster...

Future work.... Superlink@Mars.. Venus.. Moon

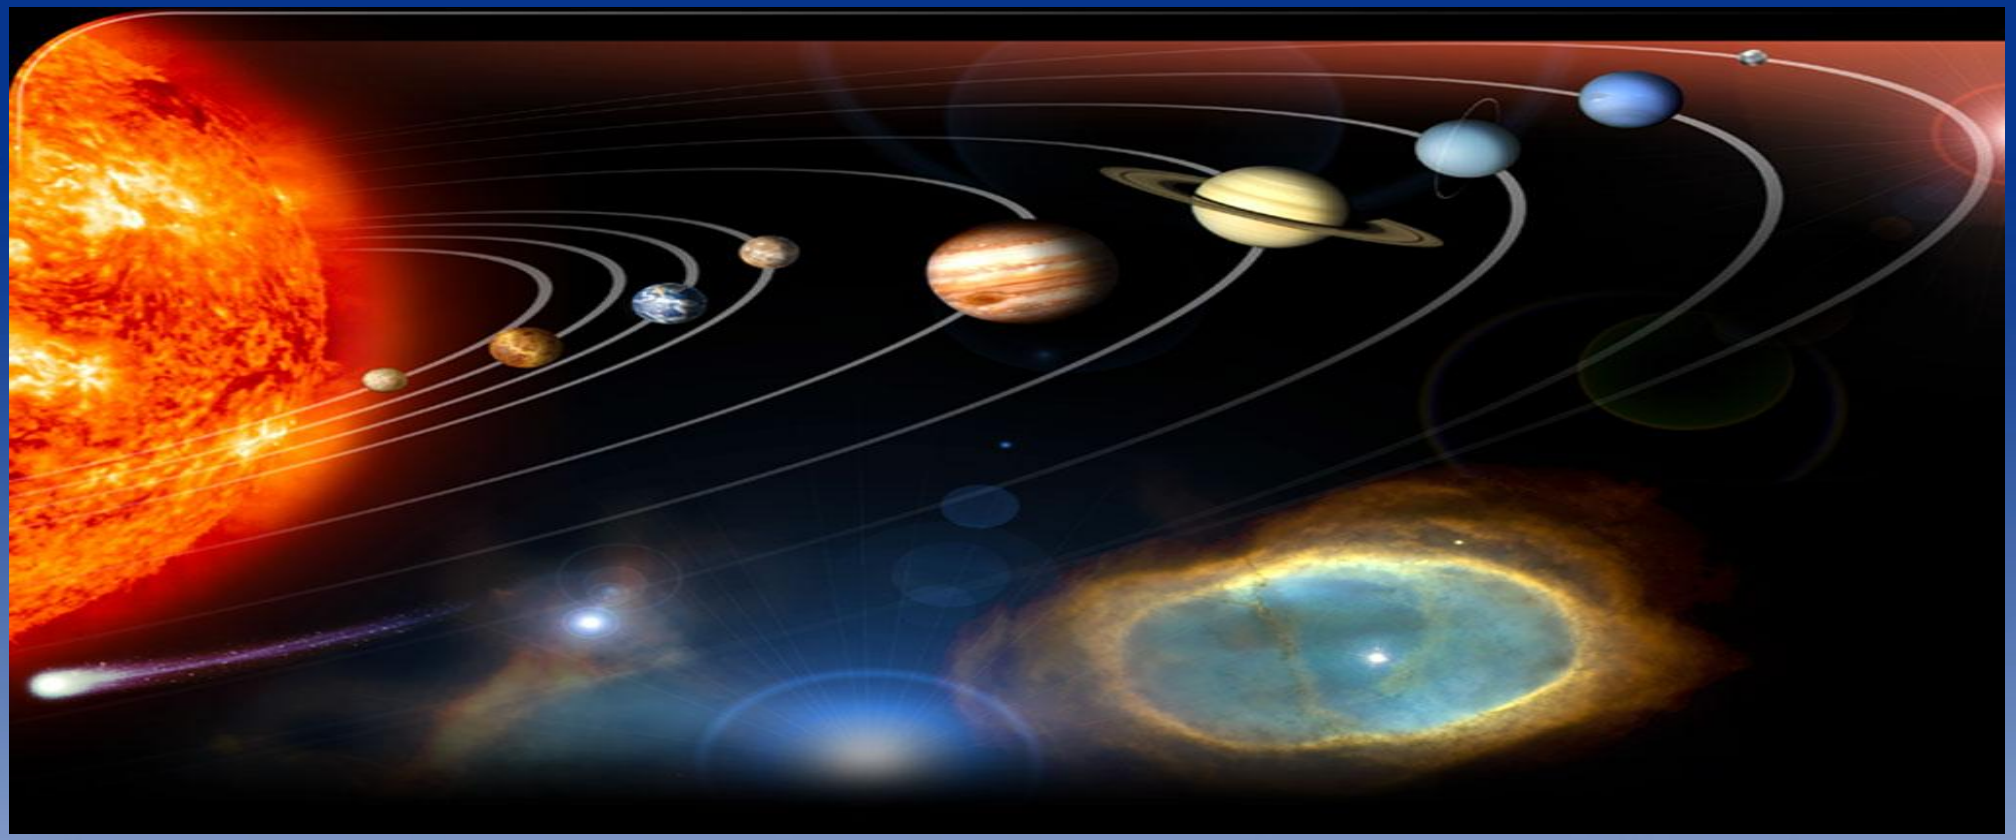

But before that ... If your grid can handle backfill – let us in!

If you want to contribute your PC(s) to the search for disease-provoking genetic mutations –

Join Superlink@Technion!