ROSIE

progress report and future directions

rosie.rosettacommons.org

Sergey Lyskov GrayLab@JHU

rosie.rosettacommons.org

Welcome | Queue | About | Documentation | Support | Login | Create an account

Recommend 0 Send +1 11

Rosetta Protocols opened for academic users:

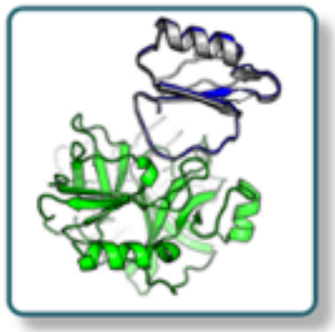

[Docking]

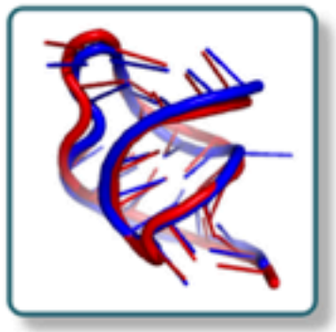

[Rna_denovo]

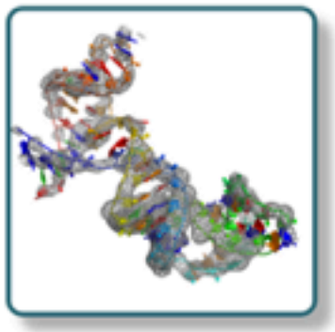

[Erraser]

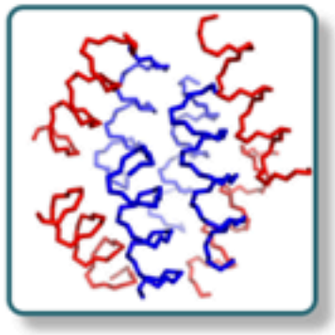

[Beta_peptide_design]

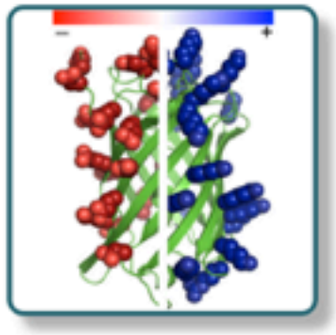

[Supercharge]

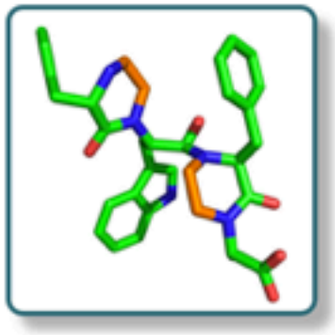

[Ncbb_design]

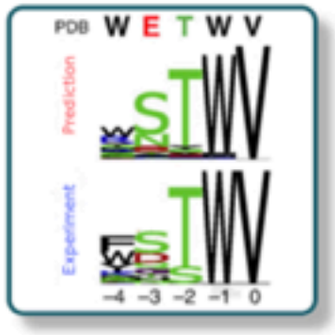

[Sequence_tolerance]

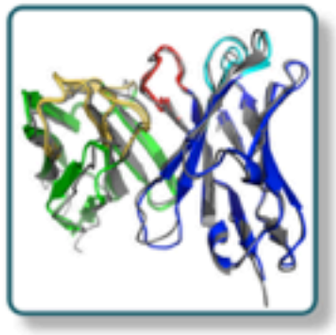

[Antibody]

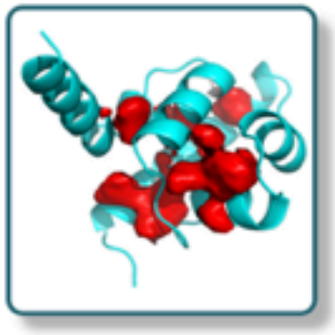

[Vip]

ROSIE stats (24hrs):

Users:	1158	+2
Jobs:	2746	+8
CPU hours:	254,273	+2,794

See more info at our About page.

Get Started with ROSIE

- ROSIE Documentation – Server related documentation and info.
- Rosetta Forums This is a list of forums for Rosetta users to discuss problems with running Rosetta and is monitored by Rosetta developers.

Using ROSIE

- Rosetta Manual Latest Rosetta User Guide.

9 Apps developed by:

(3 more in development)

Gray@JHU

Das@Stanford

Kuhlman@UNC

Bonneau@NYU

Kortemme@UCSF

Havranek@WUSTL

Maintaing so many applications could take significant resources so we were working on how to simplify it:

```
# Something like:

docking2 = App(
    name  = 'docking2', display_name = 'Docking',
    input_= [
        FileInput(...),
        StringInput(...),
        IntInput(...),
        ],

    commands = [
        docking_prepack(...),
        dock(...),
        top_n_models(...),
    ],

    output = [
        JobHeader(...),
        TopModels(...),
        ScorePlot(...),
    ],

    documentation = 'Docs here...',
)
```

Since March 2012 to July 2013: ~1,158 users registered

+ 4 each day!

Users:

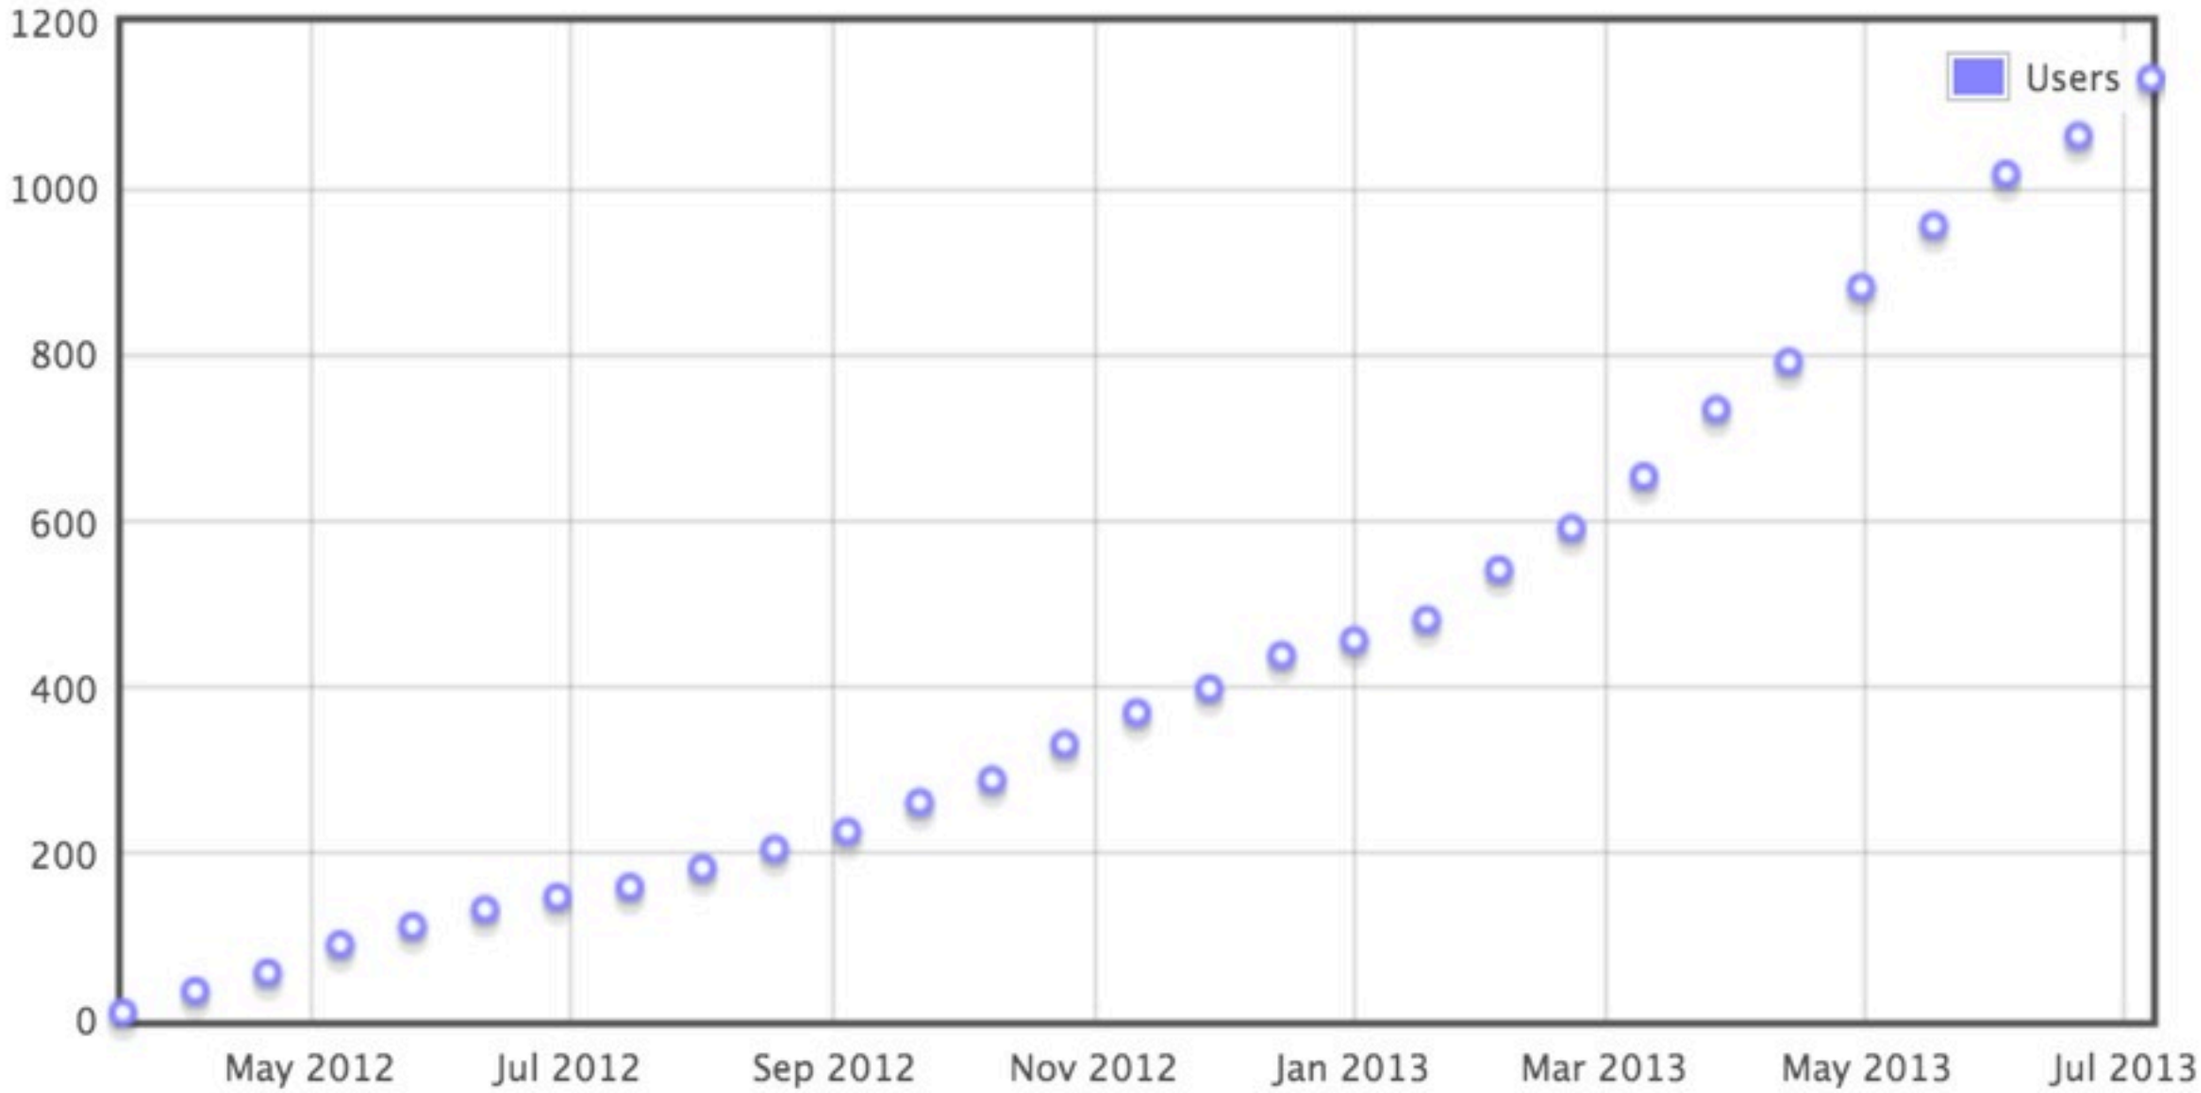

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~2,700 jobs was submitted

+ 10 each day!

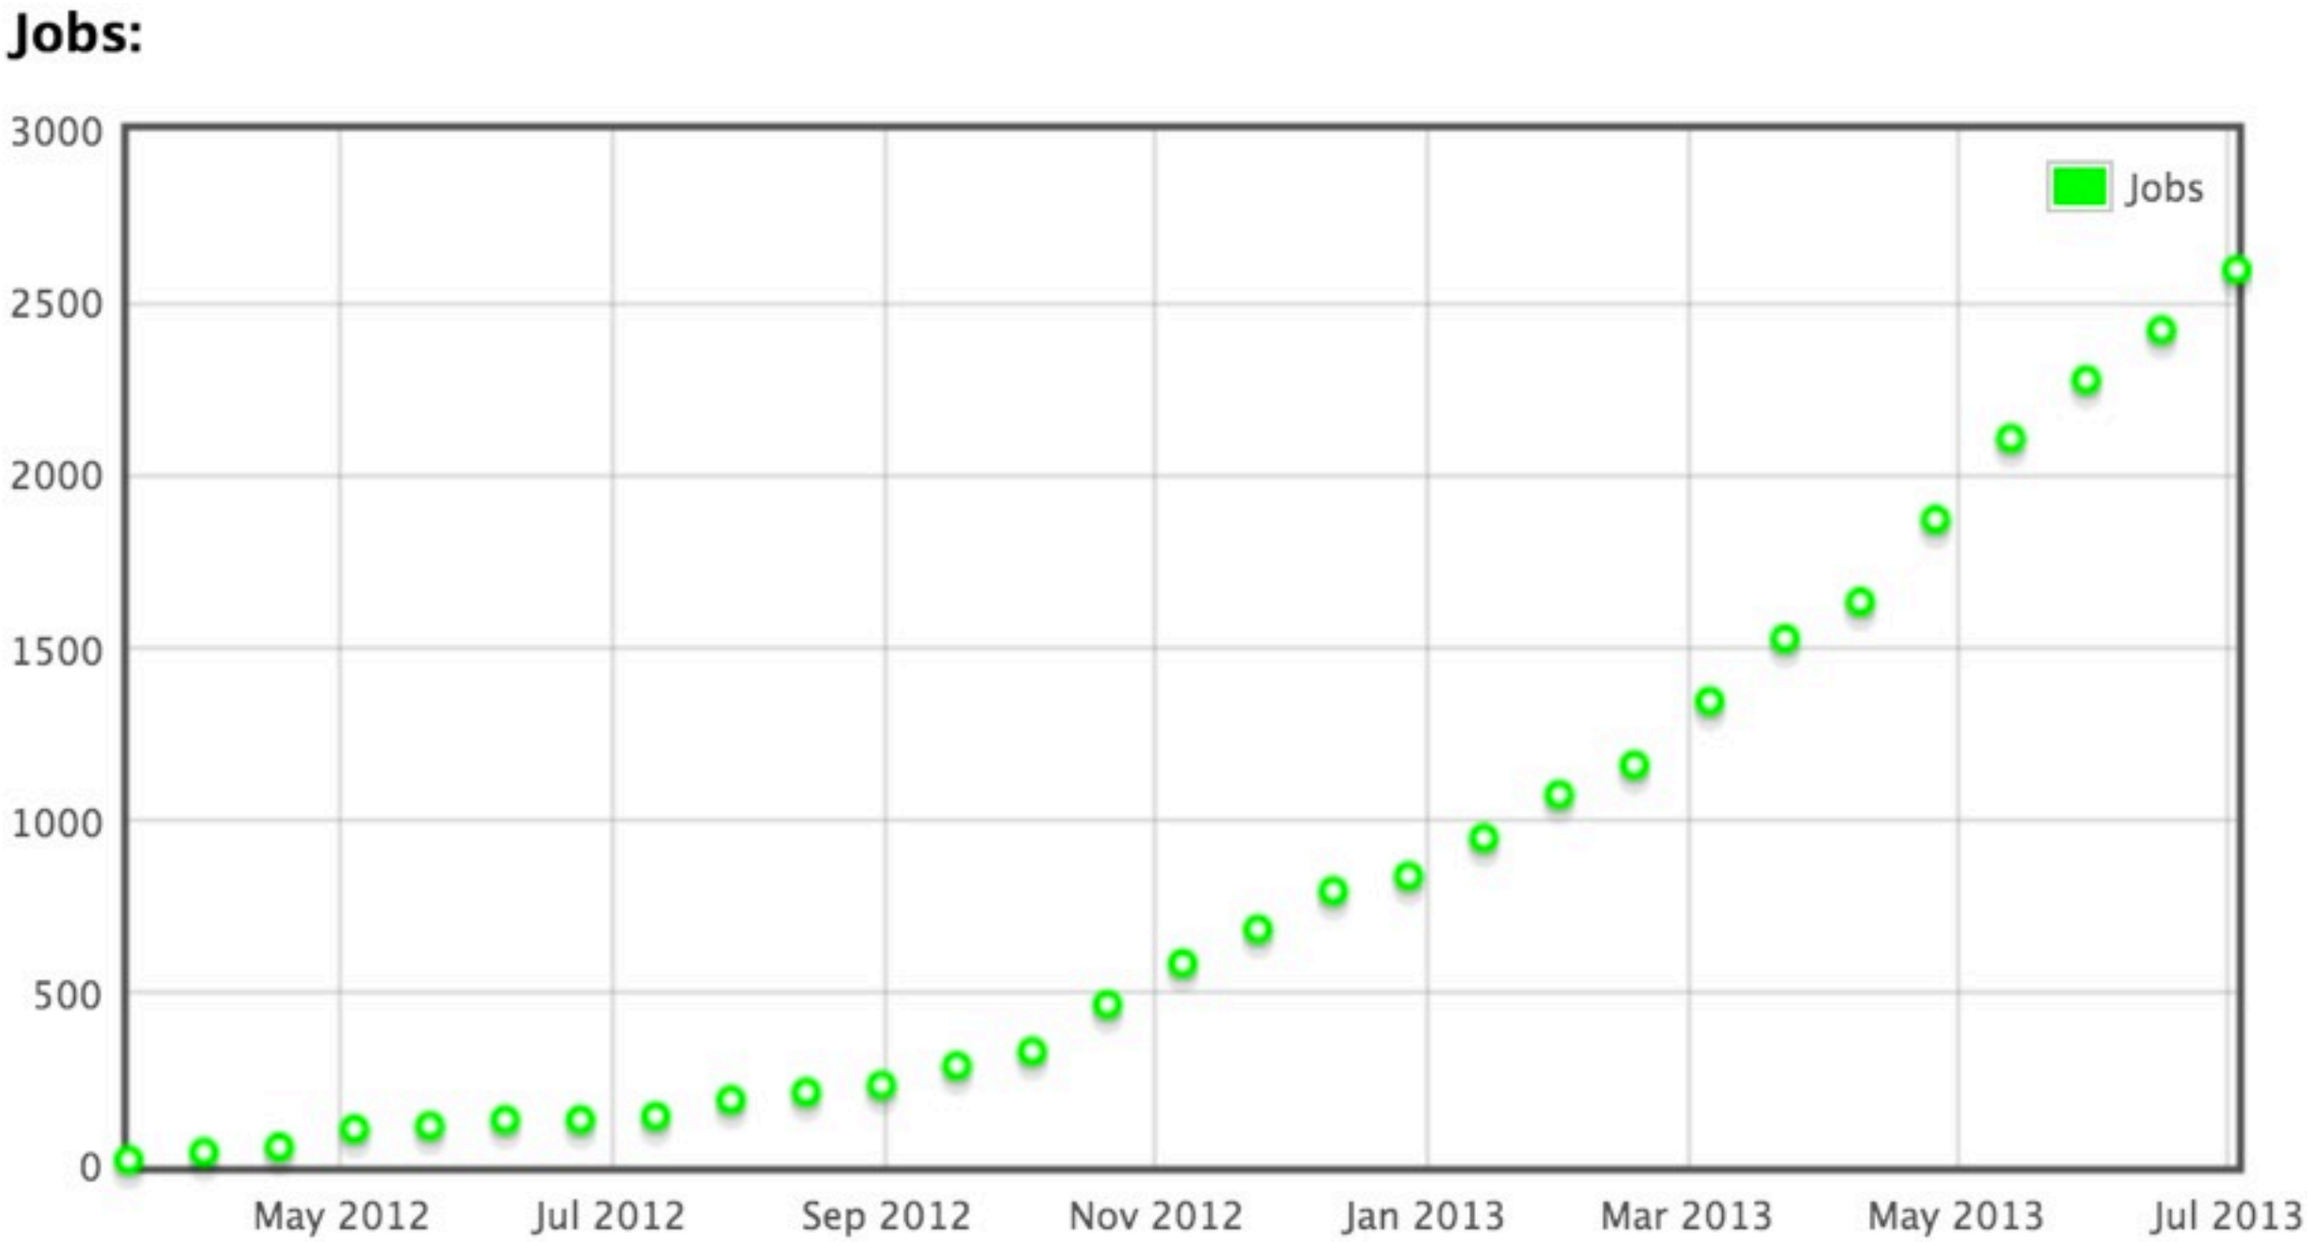

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

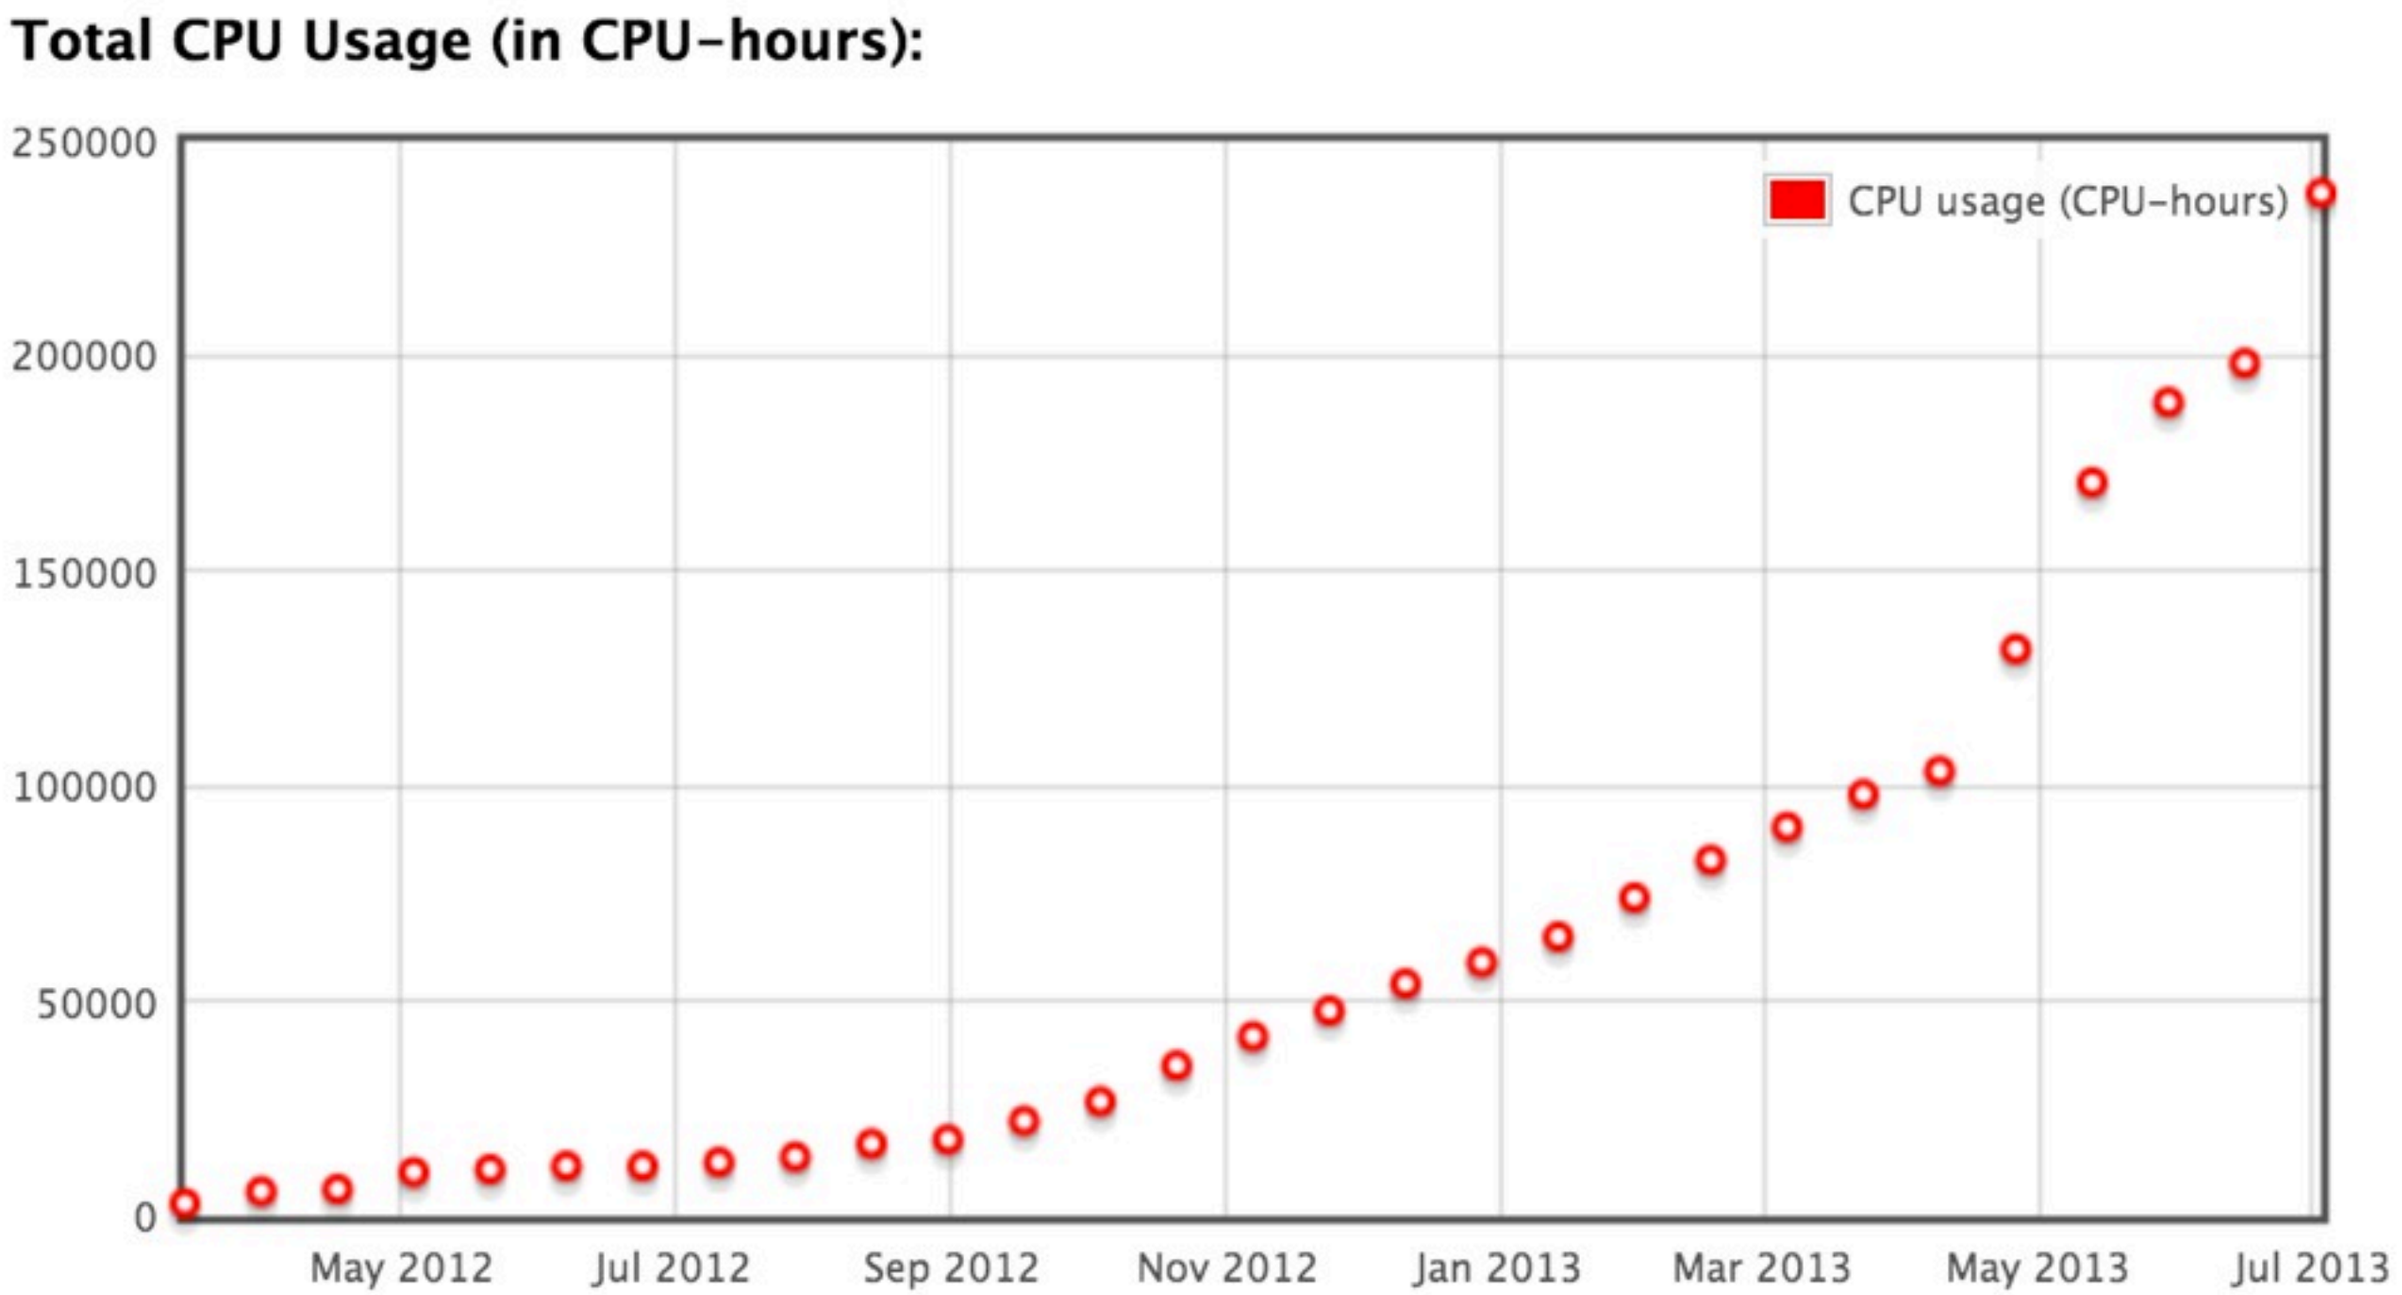

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

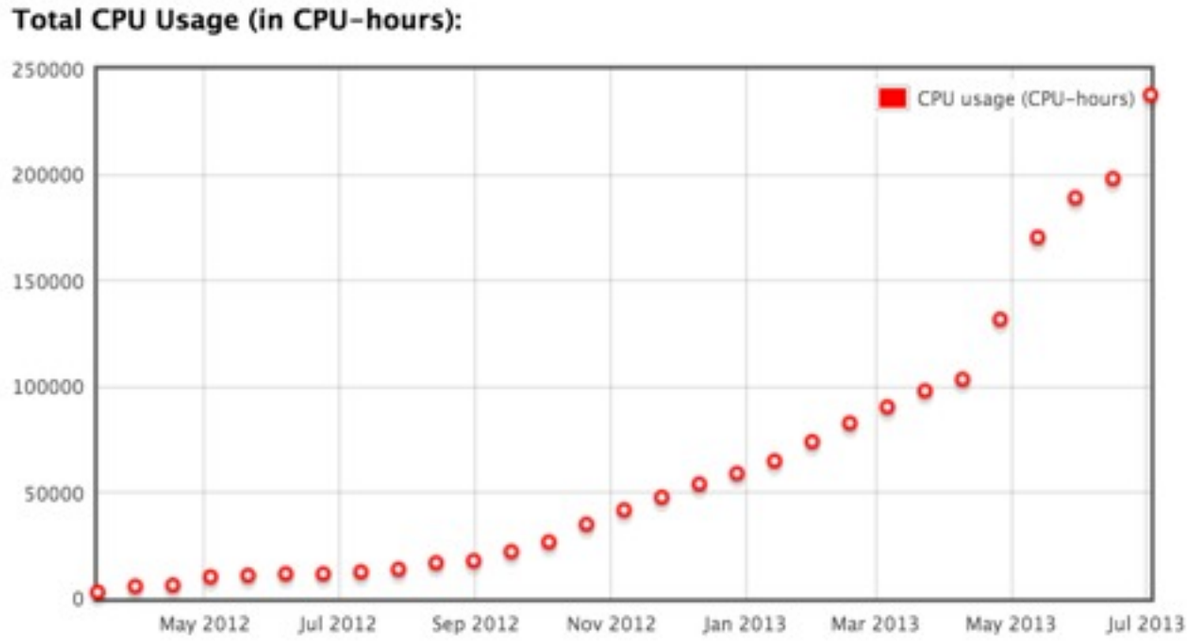

ROSIE Application	Deployment Date	Total Jobs	Total CPU (Jul'13)	Average SU/ Job	Jobs/ Month	SU/Month
Docking	Mar 2012	1,827	77,759	42.5	170	7,225
Antibody	Mar 2013	366	132,447	361.9	44	15,923
RNA Denovo	Mar 2012	217	31,915	147	16	2,352
Erraser	Oct 2012	94	101	1.1	21	23
Supercharge	Nov 2012	91	19	0.2	6	1
VIP	Feb 2013	39	48	1.2	4	5
Sequence Tolerance	Mar 2013	36	2,395	66.5	13	864
NCBB design	Dec 2012	14	3,304	236	1	236
Beta Peptide Design	Nov 2012	6	2	0.3	1	0
Total		2,693	~247,990		276	26,629

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

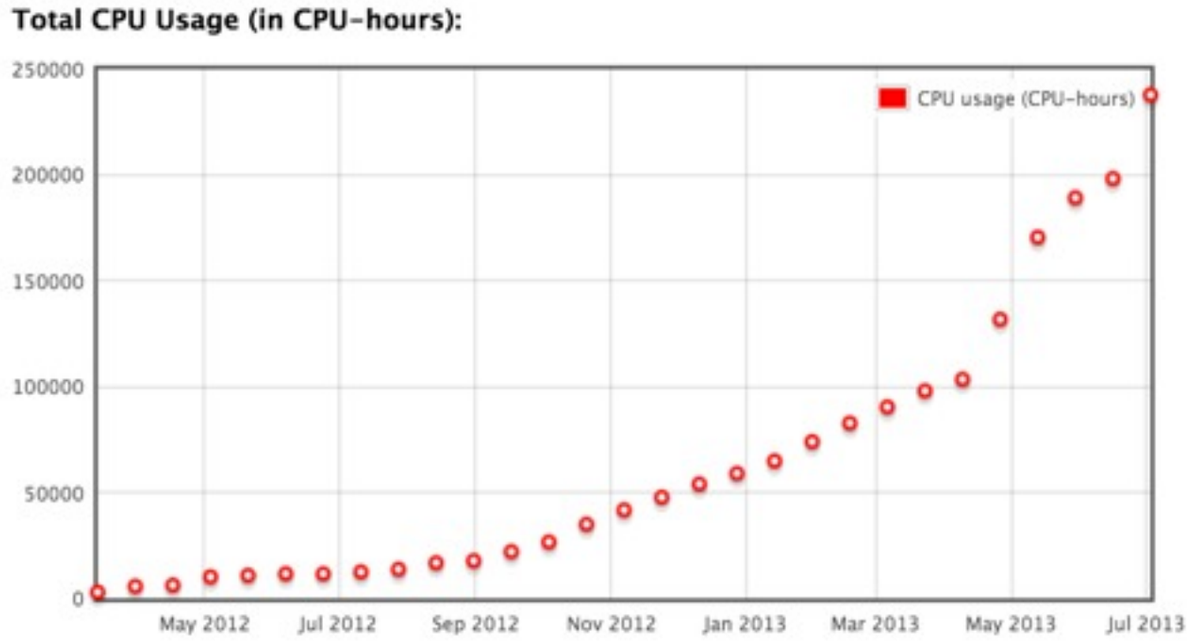

ROSIE Application	Deployment Date	Total Jobs	Total CPU (Jul'13)	Average SU/ Job	Jobs/ Month	SU/Month
Docking	Mar 2012	1,827	77,759	42.5	170	7,225
Antibody	Mar 2013	366	132,447	361.9	44	15,923
RNA Denovo	Mar 2012	217	31,915	147	16	2,352
Erraser	Oct 2012	94	101	1.1	21	23
Supercharge	Nov 2012	91	19	0.2	6	1
VIP	Feb 2013	39	48	1.2	4	5
Sequence Tolerance	Mar 2013	36	2,395	66.5	13	864
NCBB design	Dec 2012	14	3,304	236	1	236
Beta Peptide Design	Nov 2012	6	2	0.3	1	0
Total		2,693	~247,990		276	26,629

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

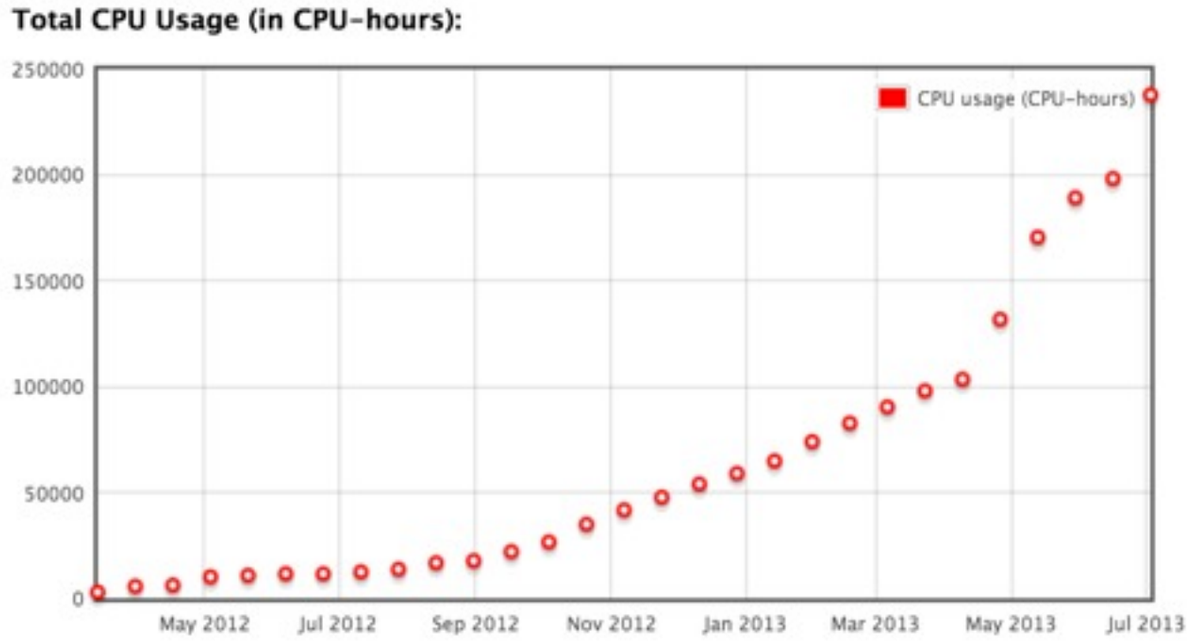

ROSIE Application	Deployment Date	Total Jobs	Total CPU (Jul'13)	Average SU/ Job	Jobs/ Month	SU/Month
Docking	Mar 2012	1,827	77,759	42.5	170	7,225
Antibody	Mar 2013	366	132,447	361.9	44	15,923
RNA Denovo	Mar 2012	217	31,915	147	16	2,352
Erraser	Oct 2012	94	101	1.1	21	23
Supercharge	Nov 2012	91	19	0.2	6	1
VIP	Feb 2013	39	48	1.2	4	5
Sequence Tolerance	Mar 2013	36	2,395	66.5	13	864
NCBB design	Dec 2012	14	3,304	236	1	236
Beta Peptide Design	Nov 2012	6	2	0.3	1	0
Total		2,693	~247,990		276	26,629

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

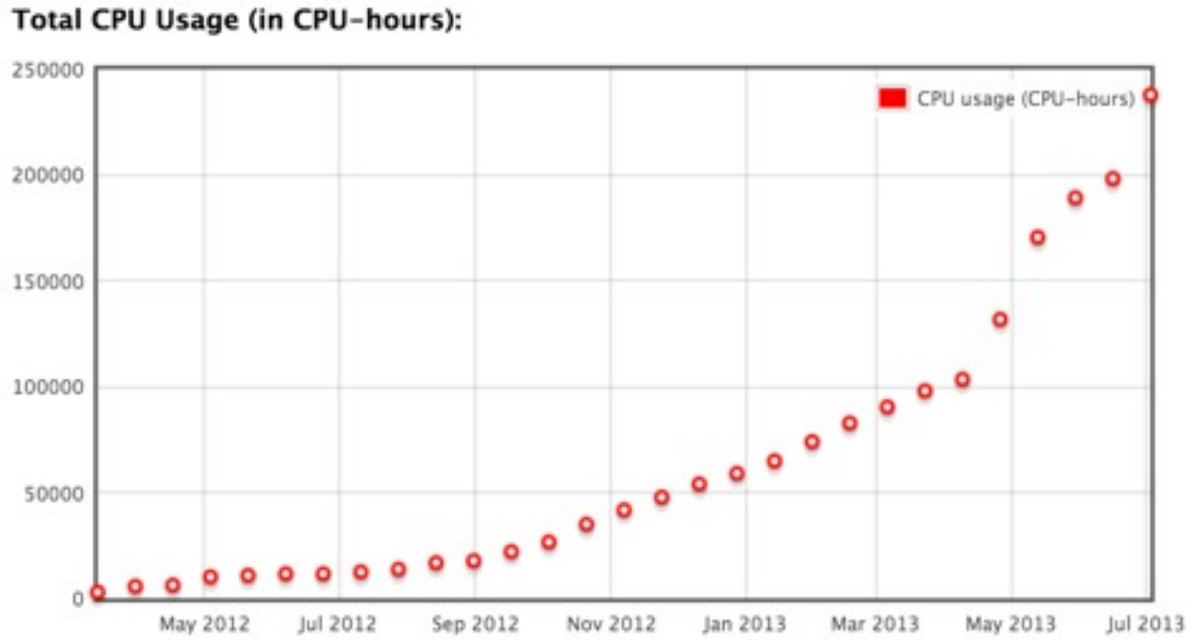

ROSIE Application	Deployment Date	Total Jobs	Total CPU (Jul'13)	Average SU/ Job	Jobs/ Month	SU/Month
Docking	Mar 2012	1,827	77,759	42.5	170	7,225
Antibody	Mar 2013	366	132,447	361.9	44	15,923
RNA Denovo	Mar 2012	217	31,915	147	16	2,352
Erraser	Oct 2012	94	101	1.1	21	23
Supercharge	Nov 2012	91	19	0.2	6	1
VIP	Feb 2013	39	48	1.2	4	5
Sequence Tolerance	Mar 2013	36	2,395	66.5	13	864
NCBB design	Dec 2012	14	3,304	236	1	236
Beta Peptide Design	Nov 2012	6	2	0.3	1	0
Total		2,693	~247,990		276	26,629

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

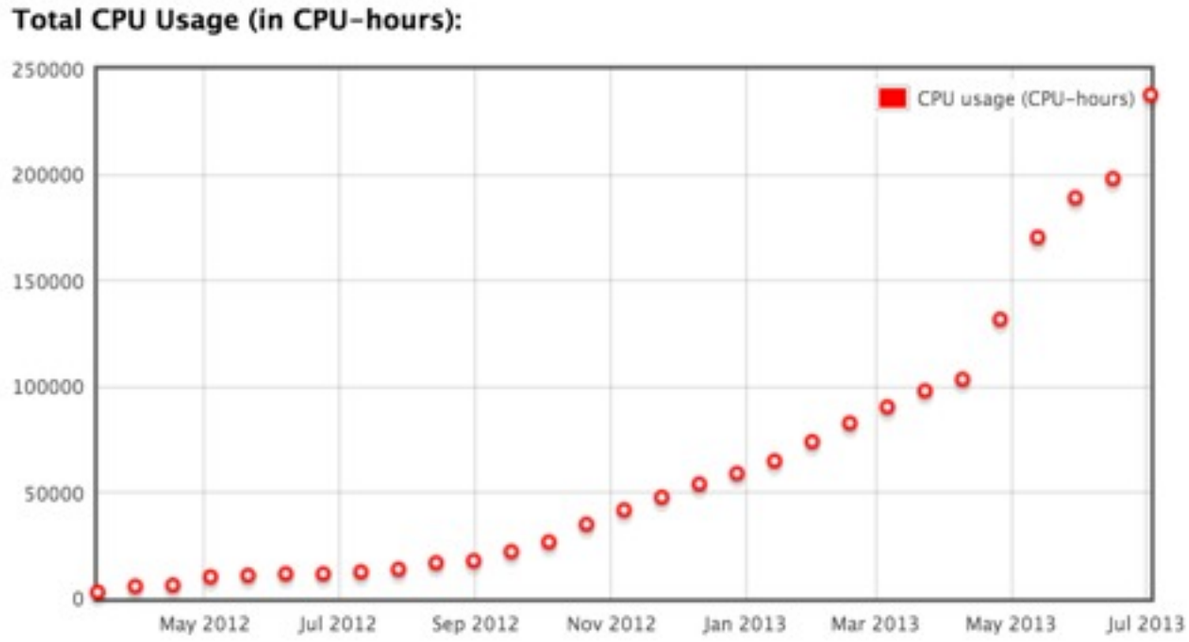

ROSIE Application	Deployment Date	Total Jobs	Total CPU (Jul'13)	Average SU/ Job	Jobs/ Month	SU/Month
Docking	Mar 2012	1,827	77,759	42.5	170	7,225
Antibody	Mar 2013	366	132,447	361.9	44	15,923
RNA Denovo	Mar 2012	217	31,915	147	16	2,352
Erraser	Oct 2012	94	101	1.1	21	23
Supercharge	Nov 2012	91	19	0.2	6	1
VIP	Feb 2013	39	48	1.2	4	5
Sequence Tolerance	Mar 2013	36	2,395	66.5	13	864
NCBB design	Dec 2012	14	3,304	236	1	236
Beta Peptide Design	Nov 2012	6	2	0.3	1	0
Total		2,693	~247,990		276	26,629

rosie.rosettacommons.org/about

Since March 2012 to July 2013: ~254,000 CPU hours was used

+ 900 each day!

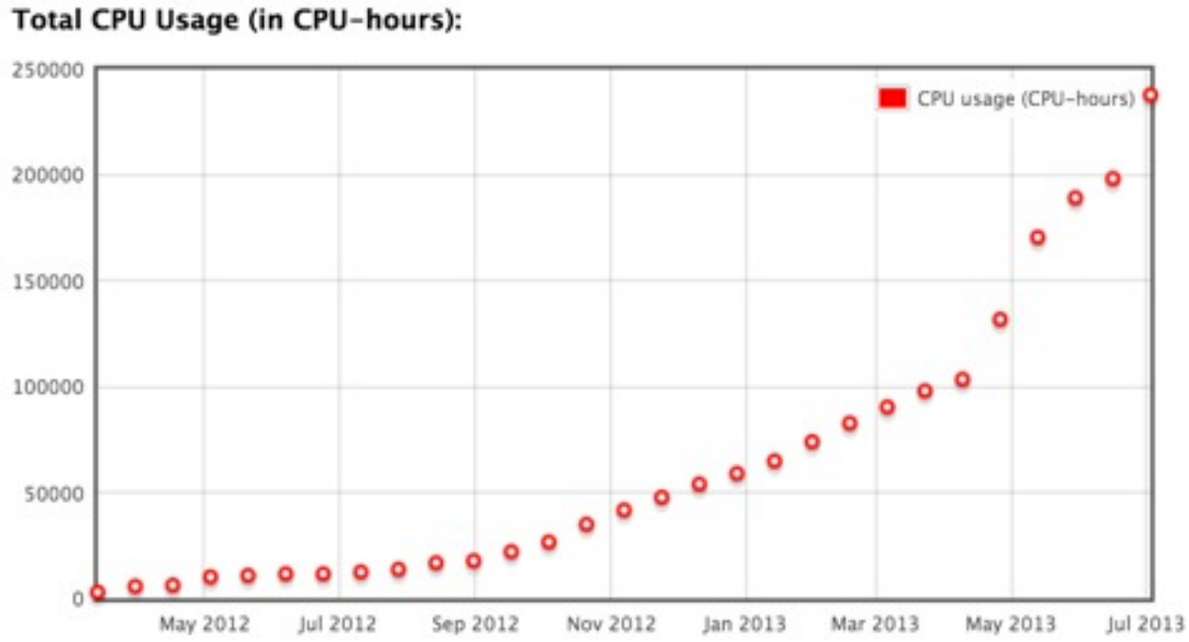

Two HPC backends:

GrayLab cluster, 328 CPU's

Rosetta clusters, 96 CPU's

ROSIE Application	Deployment Date	Total Jobs	Total CPU (Jul'13)	Average SU/ Job	Jobs/ Month	SU/Month
Docking	Mar 2012	1,827	77,759	42.5	170	7,225
Antibody	Mar 2013	366	132,447	361.9	44	15,923
RNA Denovo	Mar 2012	217	31,915	147	16	2,352
Erraser	Oct 2012	94	101	1.1	21	23
Supercharge	Nov 2012	91	19	0.2	6	1
VIP	Feb 2013	39	48	1.2	4	5
Sequence Tolerance	Mar 2013	36	2,395	66.5	13	864
NCBB design	Dec 2012	14	3,304	236	1	236
Beta Peptide Design	Nov 2012	6	2	0.3	1	0
Total		2,693	~247,990		276	26,629

rosie.rosettacommons.org/about

Since introduction of antibody App our monthly CPU usage almost doubled, and our queue is already ~10 days long, so we thought it would be a good idea to use:

Since introduction of antibody App our monthly CPU usage almost doubled, and our queue is already ~10 days long, so we thought it would be a good idea to use:

TACC **TEXAS ADVANCED COMPUTING CENTER**

Powering Discoveries That Change The World

THE UNIVERSITY OF TEXAS AT AUSTIN

Home | About | Resources | User Services | Research & Development | Partnerships | Education & Outreach | News

STAMPEDE

Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors

Resources

HPC Systems
Visualization Resources
Data Storage
Networking
Software

HPC Systems

Stampede
Lonestar
Ranger - *Decommissioned*

Quick Links

About TACC
TACC Software

Contact

Chris Hempel
Associate Director, Resources & Services
hempel@tacc.utexas.edu

Since introduction of antibody App our monthly CPU usage almost doubled, and our queue is already ~10 days long, so we thought it would be a good idea to use:

TACC TEXAS ADVANCED COMPUTING CENTER

Powering Discoveries That Change The World

THE UNIVERSITY OF TEXAS AT AUSTIN

Home About Resources User Services Research & Development Partnerships Education & Outreach News

STAMPEDE

Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors

Resources
HPC Systems
Visualization Resources
Data Storage
Networking
Software

HPC Systems
Stampede
Lonestar
Ranger - *Decommissioned*

Quick Links
About TACC
TACC Software

Contact
Chris Hempel
Associate Director, Resources & Services
hempel@tacc.utexas.edu

- #6 in Top-500

Since introduction of antibody App our monthly CPU usage almost doubled, and our queue is already ~10 days long, so we thought it would be a good idea to use:

TACC TEXAS ADVANCED COMPUTING CENTER

Powering Discoveries That Change The World

THE UNIVERSITY OF TEXAS AT AUSTIN

Home | About | Resources | User Services | Research & Development | Partnerships | Education & Outreach | News

STAMPEDE

Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors

Resources
HPC Systems
Visualization Resources
Data Storage
Networking
Software

HPC Systems
Stampede
Lonestar
Ranger - *Decommissioned*

Quick Links
About TACC
TACC Software

Contact
Chris Hempel
Associate Director, Resources & Services
hempel@tacc.utexas.edu

- #6 in Top-500
- 102,000 cores

Since introduction of antibody App our monthly CPU usage almost doubled, and our queue is already ~10 days long, so we thought it would be a good idea to use:

TACC TEXAS ADVANCED COMPUTING CENTER
Powering Discoveries That Change The World
THE UNIVERSITY OF TEXAS AT AUSTIN

Home | About | Resources | User Services | Research & Development | Partnerships | Education & Outreach | News

STAMPEDE

Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors

Resources
HPC Systems
Visualization Resources
Data Storage
Networking
Software

HPC Systems
Stampede
Lonestar
Ranger - *Decommissioned*

Quick Links
About TACC
TACC Software

Contact
Chris Hempel
Associate Director, Resources & Services
hempel@tacc.utexas.edu

- #6 in Top-500
- 102,000 cores
- +Intel Xeon PHI Co-processors

TACC STAMPEDE

TACC TEXAS ADVANCED COMPUTING CENTER
Powering Discoveries That Change The World
THE UNIVERSITY OF TEXAS AT AUSTIN

Home | About | Resources | User Services | Research & Development | Partnerships | Education & Outreach | News

STAMPEDE

Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors

Resources
HPC Systems
Visualization Resources
Data Storage
Networking
Software

HPC Systems
Stampede
Lonestar
Ranger - *Decommissioned*

Quick Links
About TACC
TACC Software

Contact
Chris Hempel
Associate Director, Resources & Services
hempel@tacc.utexas.edu

TACC STAMPEDE

- We submitted proposal this July, asking for 1,000,000 CPU/hours

TACC TEXAS ADVANCED COMPUTING CENTER
Powering Discoveries That Change The World
THE UNIVERSITY OF TEXAS AT AUSTIN

Home About Resources User Services Research & Development Partnerships Education & Outreach News

STAMPEDE

Dell PowerEdge C8220 Cluster with Intel Xeon Phi coprocessors

Resources
HPC Systems
Visualization Resources
Data Storage
Networking
Software

HPC Systems
Stampede
Lonestar
Ranger - *Decommissioned*

Quick Links
About TACC
TACC Software

Contact
Chris Hempel
Associate Director, Resources & Services
hempel@tacc.utexas.edu

TACC STAMPEDE

- We submitted proposal this July, asking for 1,000,000 CPU/hours
- Dr. Frank Willmore (TACC) working with Andrew to parallelize code for PHI co-processors

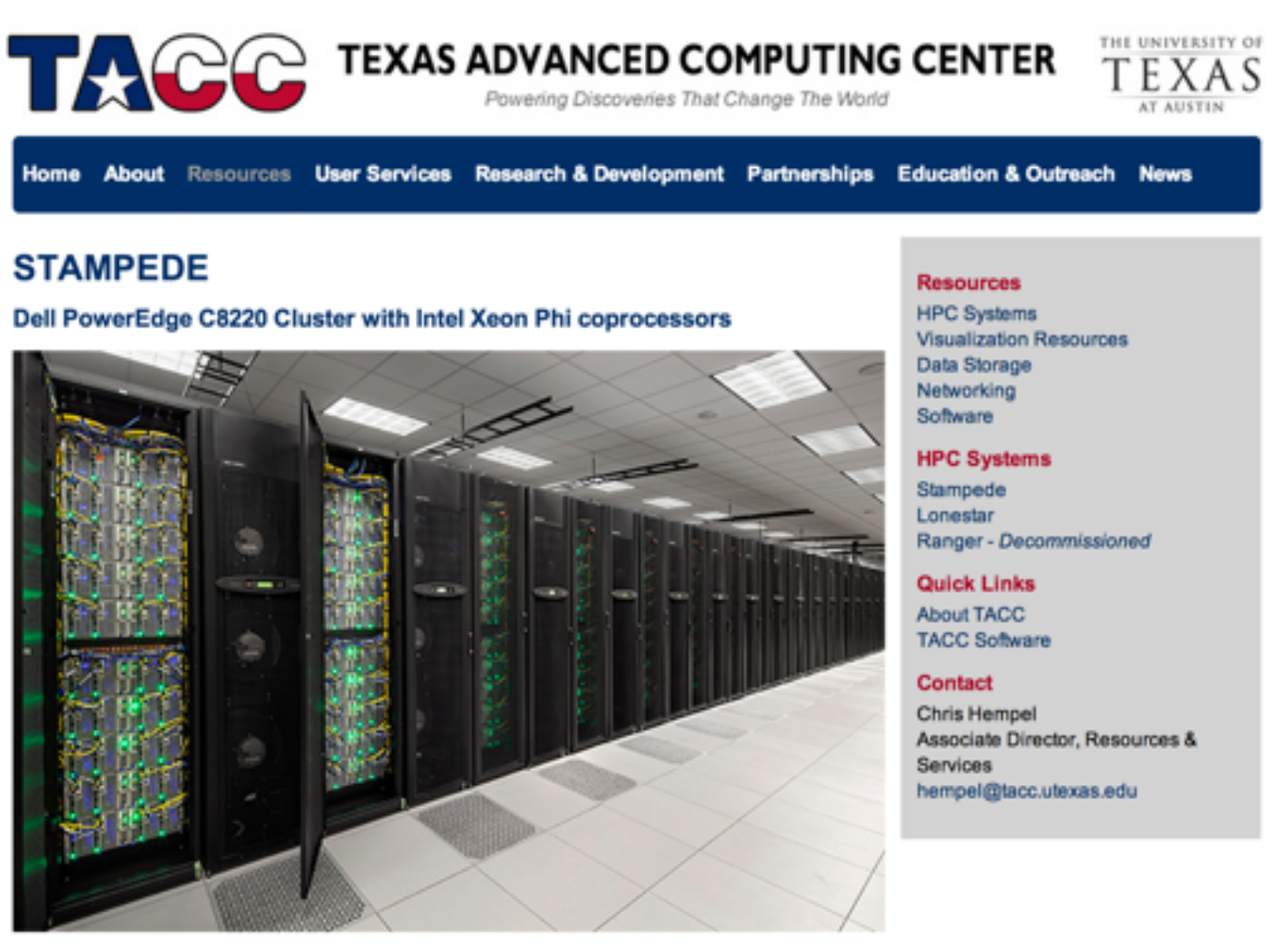

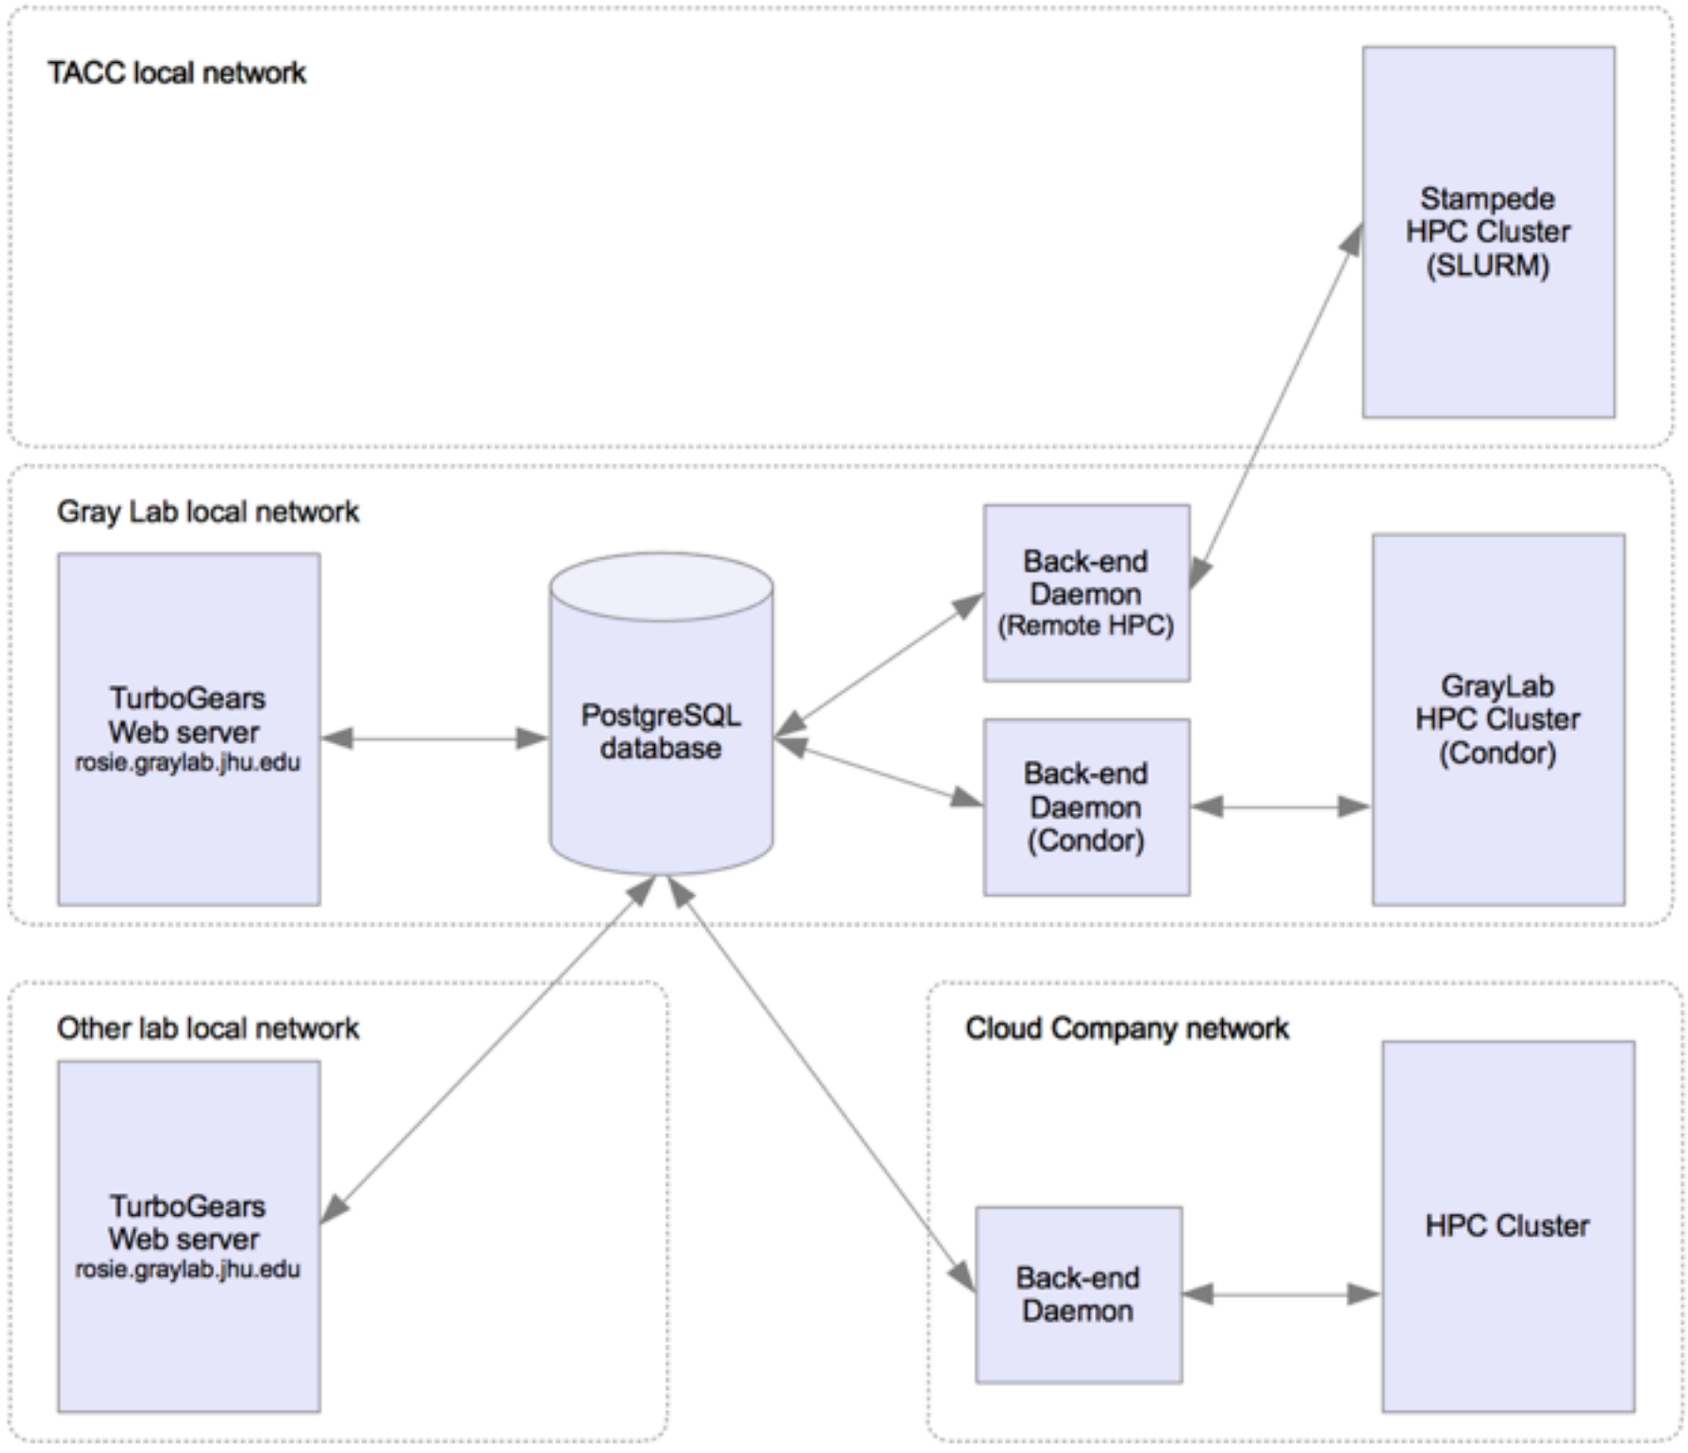
TACC local network
Stampede
HPC Cluster
(SLURM)
Gray Lab local network
TurboGears
Web server
rosie.graylab.jhu.edu
PostgreSQL
database
Back-end
Daemon
(Remote HPC)
Back-end
Daemon
(Condor)
GrayLab
HPC Cluster
(Condor)
Other lab local network
TurboGears
Web server
rosie.graylab.jhu.edu
Cloud Company network
Back-end
Daemon
HPC Cluster

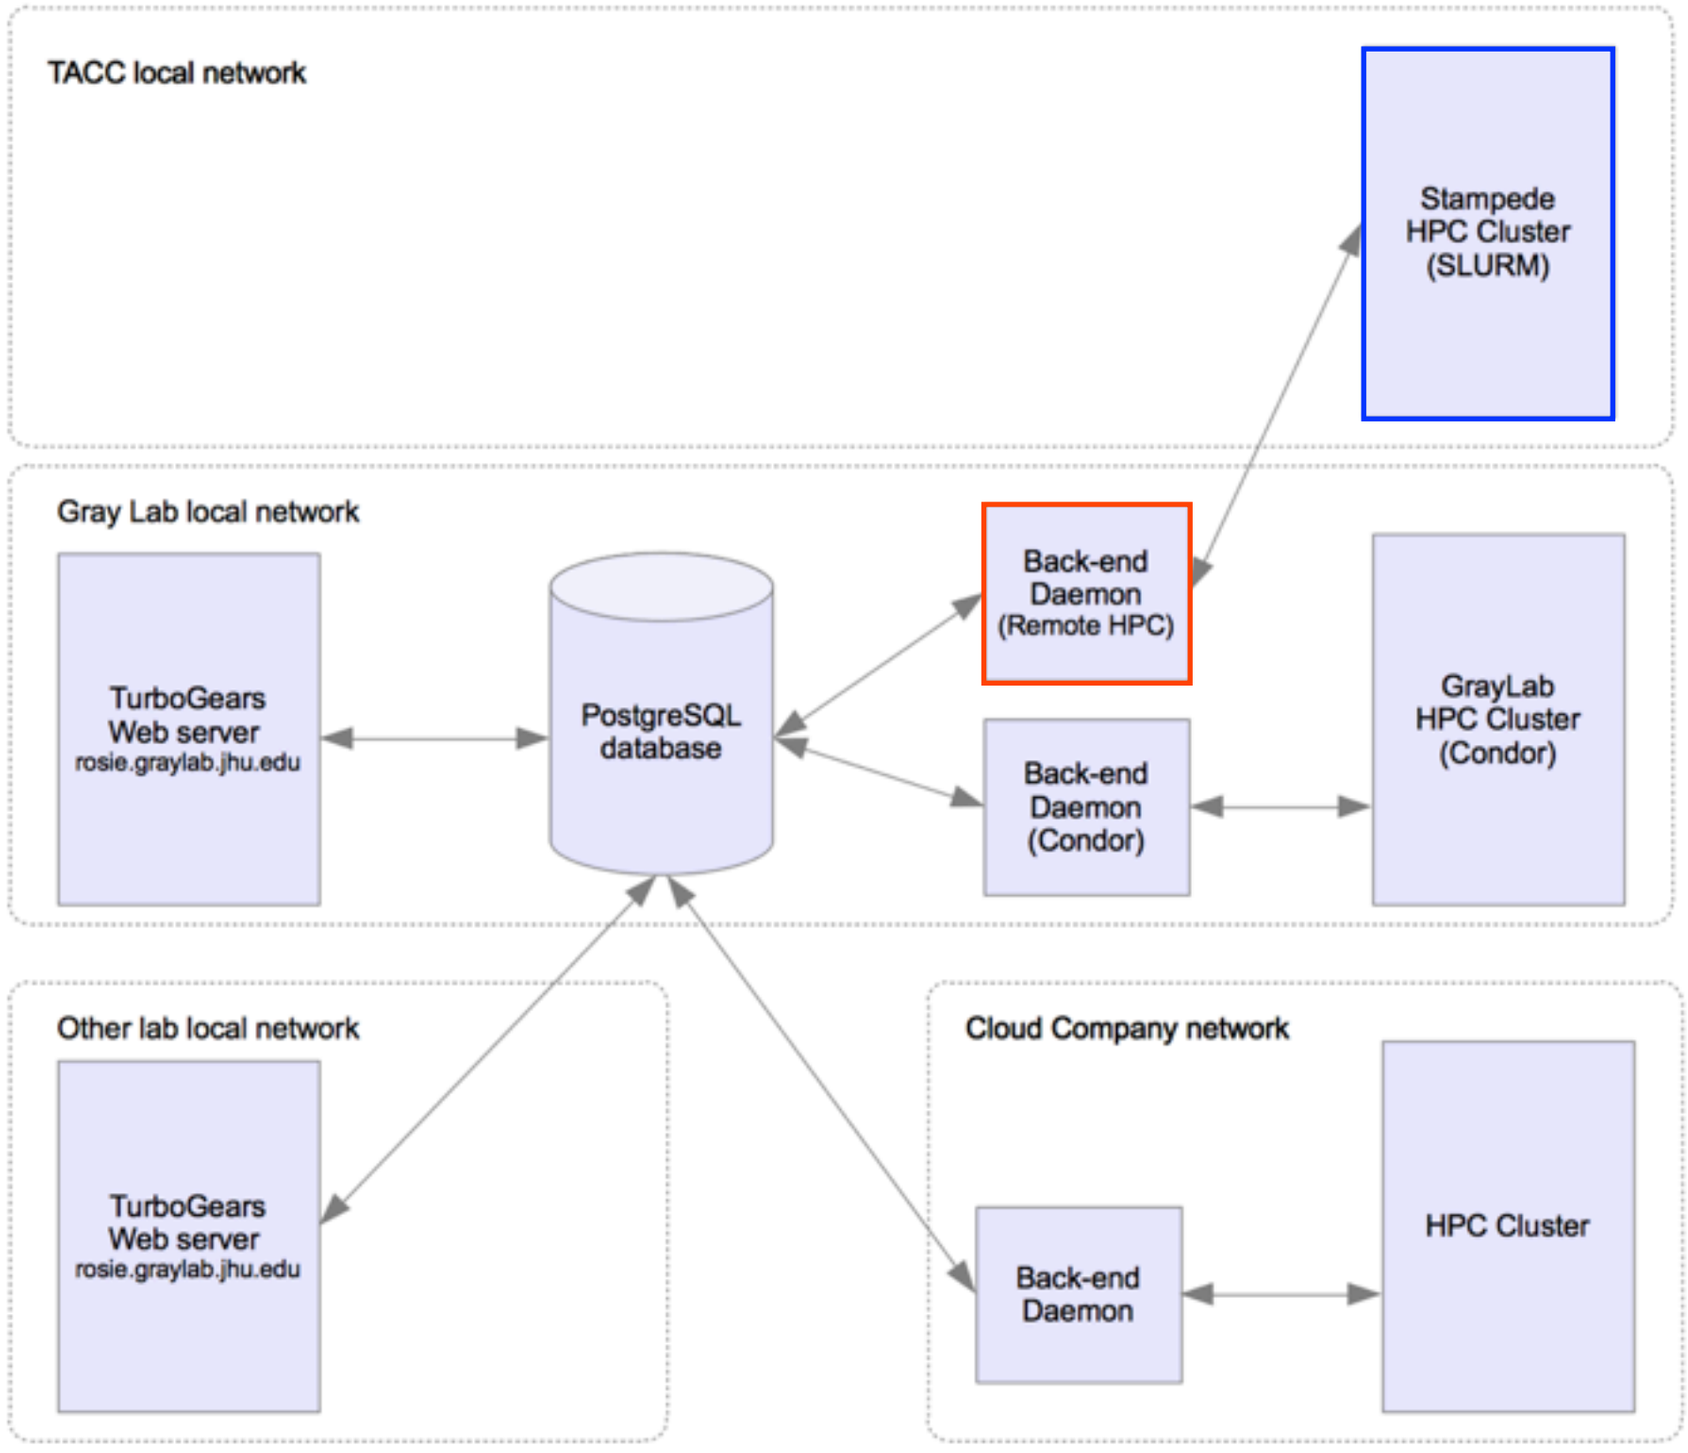
TACC local network
Stampede
HPC Cluster
(SLURM)
Gray Lab local network
TurboGears
Web server
rosie.graylab.jhu.edu
PostgreSQL
database
Back-end
Daemon
(Remote HPC)
Back-end
Daemon
(Condor)
GrayLab
HPC Cluster
(Condor)
Other lab local network
TurboGears
Web server
rosie.graylab.jhu.edu
Cloud Company network
Back-end
Daemon
HPC Cluster

TACC require us to evaluate application scaleability of our code for MPI builds:

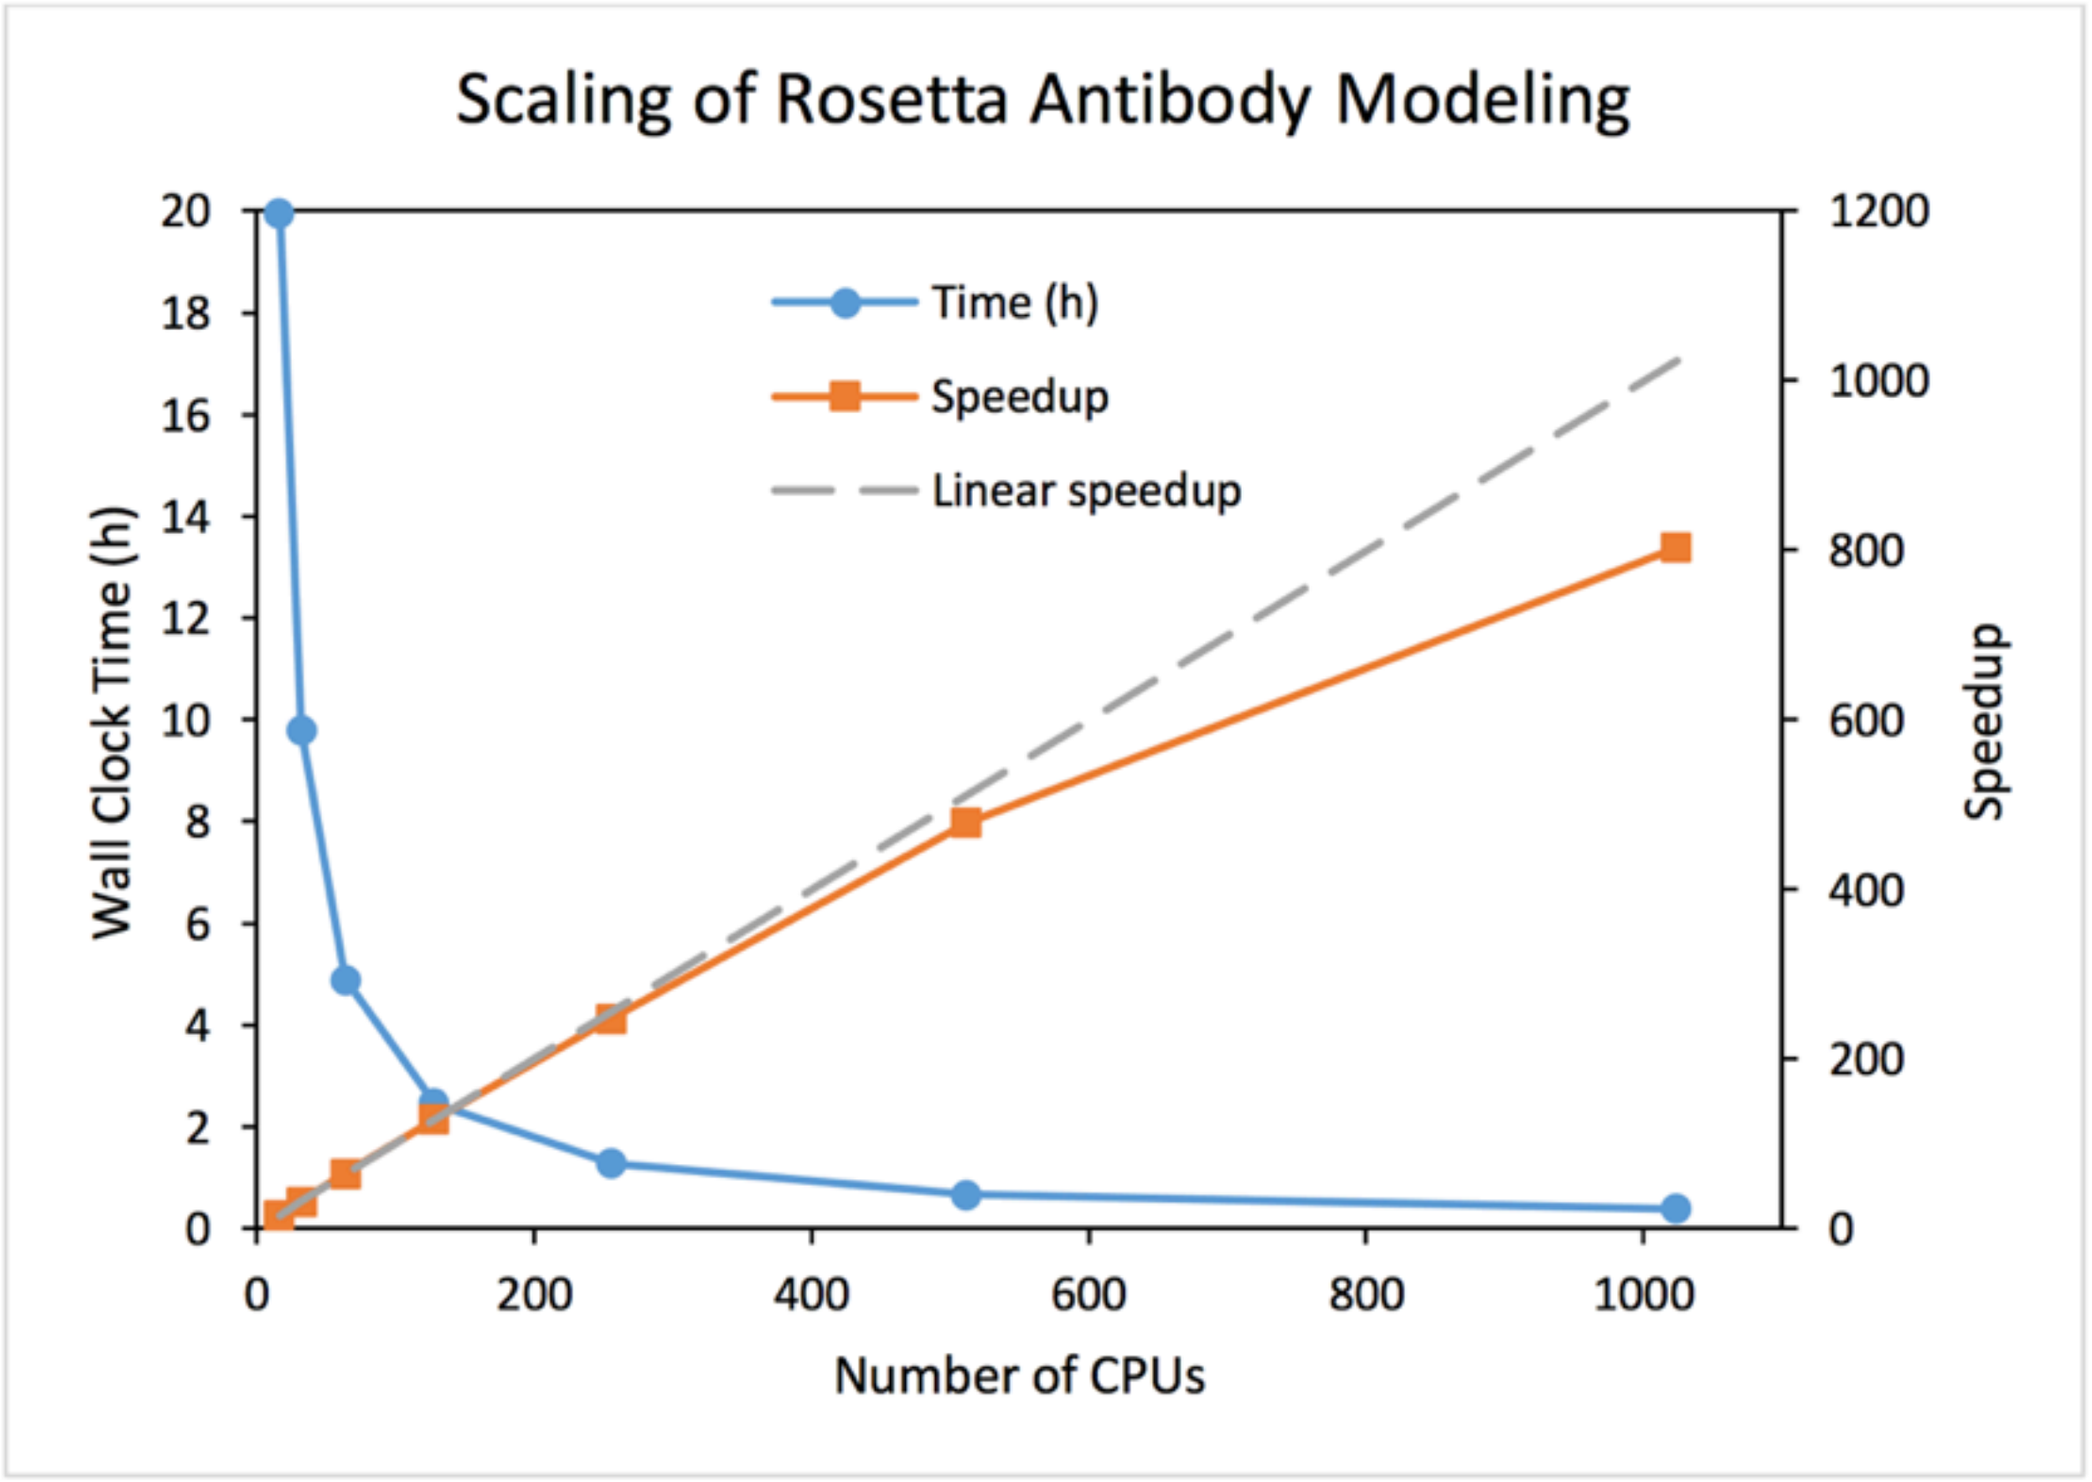

Plans for next year:

Plans for next year:

- Finalize and further simplify App creation

Plans for next year:

- Finalize and further simplify App creation
- 13 more apps from 10 labs:

pKa	Gray@JHU
SymDock	Andre@LU
Capper	Havranek@WUStL
RNA design	Das@Stanford
RNA homology	Das@Stanford
Small loop enumeration	Das@Stanford
FlexPepDock	Furman@HebrewU
PeptiDerive	Furman@HebrewU
pocket_optimize	Karanicolas@KU
ligand docking	Meiler@Vanderbilt
C2H2 ZF prediction	Brandley@FHCRC
BAM	Dunbrack@FCCC
Sequence Tolerance 2.0	Kortemme@UCSF

Plans for next year:

- Finalize and further simplify App creation
- 13 more apps from 10 labs:

Plans for next year:

Plans for next year:

- Finalize and further simplify App creation
- 13 more apps from 10 labs
- Remote HPC driver, so we can use any HPC clusters

Plans for next year:

- Finalize and further simplify App creation
- 13 more apps from 10 labs
- Remote HPC driver, so we can use any HPC clusters
- TACC back-end (if our TACC proposal get approval)

ROSIE and Rosetta benchmark infrastructure is maintained with help of our sys-admin: Matt Mulqueen

ROSIE and Rosetta benchmark infrastructure is maintained with help of our sys-admin: Matt Mulqueen

Thank you!