Title: 14/15 PSI - Mathematica

Date: Sep 05, 2014 10:30 AM

URL: http://pirsa.org/14090040

Abstract:

Multi-Tiered Strategy

- Tier 1: Personal Devices
  - Laptop, desktop
- Tier 2: Scientific Computing Environment (SCE)
  - Shared workstations with remote access
  - Virtual machines
- Tier 3: High Performance Computing (HPC)
  - HPC system “Titan”, access via batch queue
- Tier 4: Support for using outside resources
  - Sharcnet, Scinet, XSEDE, ...

Research Technologies

- Liaison between IT support group and researchers
- Consulting and support e.g. for:
  - Questions on Latex, Mathematica, ...
  - Particular software package
  - Assistance with programming (C, C++, Fortran, Python, Perl, ...)
  - Consulting on numerical algorithms
  - Collaboration on projects/publications

- consult@perimeterinstitute.ca,
  or office 350 (new wing, next to elevator)

Scientific Computing Environment

- 7 “regular” Linux workstations, located in a server room
- Names: compute, compute1, …, compute6
- Running Ubuntu, with many additional packages installed
  - 8 Intel cores with 2.4 GHz
  - 40 GByte RAM
  - Using Perimeter home directory
- Also 1 GPU workstation named “Nvidia”, same software, has Nvidia GPU
- Access via ssh:
  - Login: ssh yourlogin@compute.pi.local
  - Copy files: access via network drive
- Mathematica: can use remote desktop

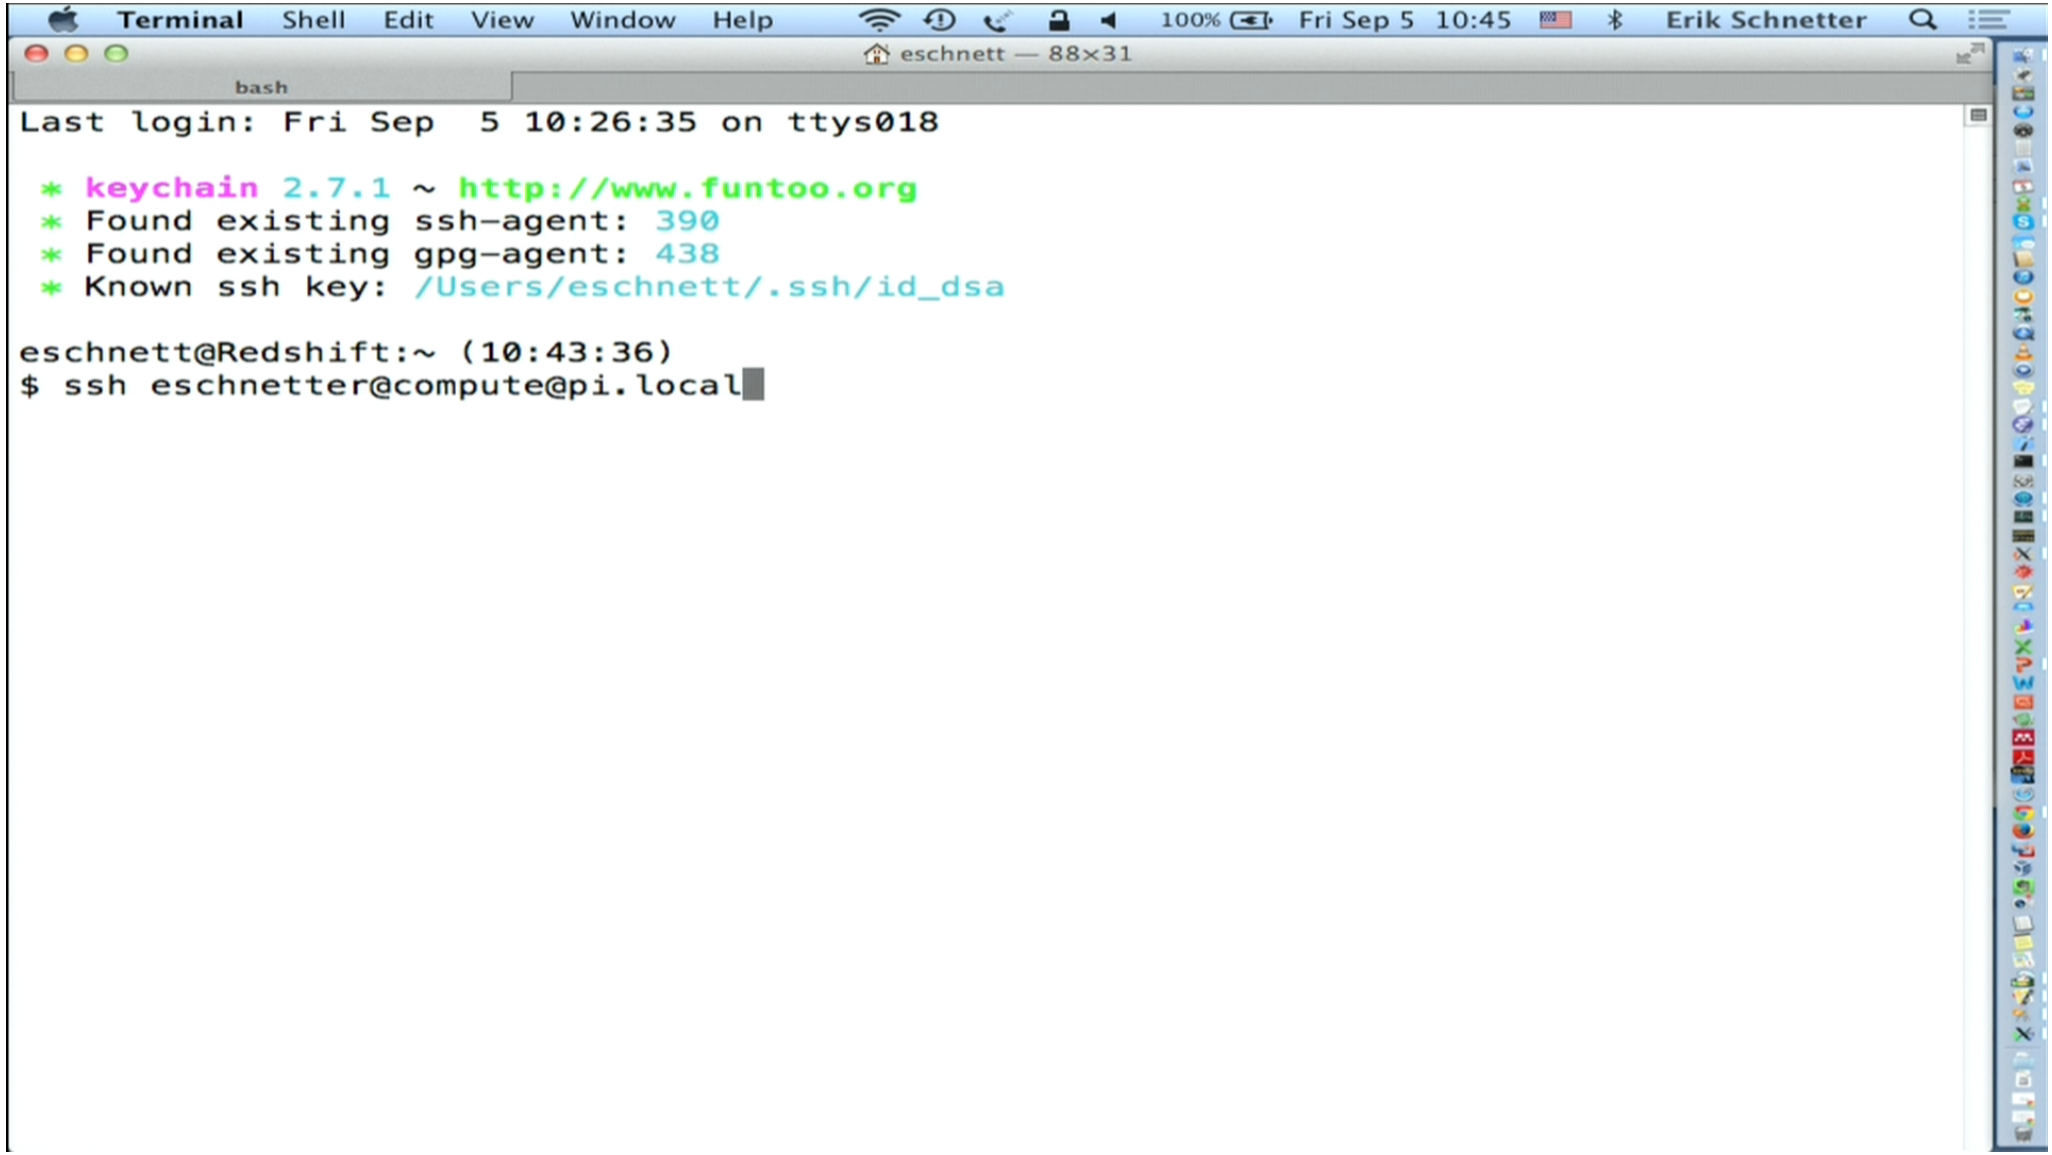
Terminal Shell Edit View Window Help
100% Fri Sep 5 10:45 Erik Schnetter
eschnett — 88×31
bash
Last login: Fri Sep 5 10:26:35 on ttys018
* keychain 2.7.1 ~ http://www.funtoo.org
* Found existing ssh-agent: 390
* Found existing gpg-agent: 438
* Known ssh key: /Users/eschnett/.ssh/id_dsa
eschnett@Redshift:~ (10:43:36)
$ ssh eschnetter@compute@pi.local

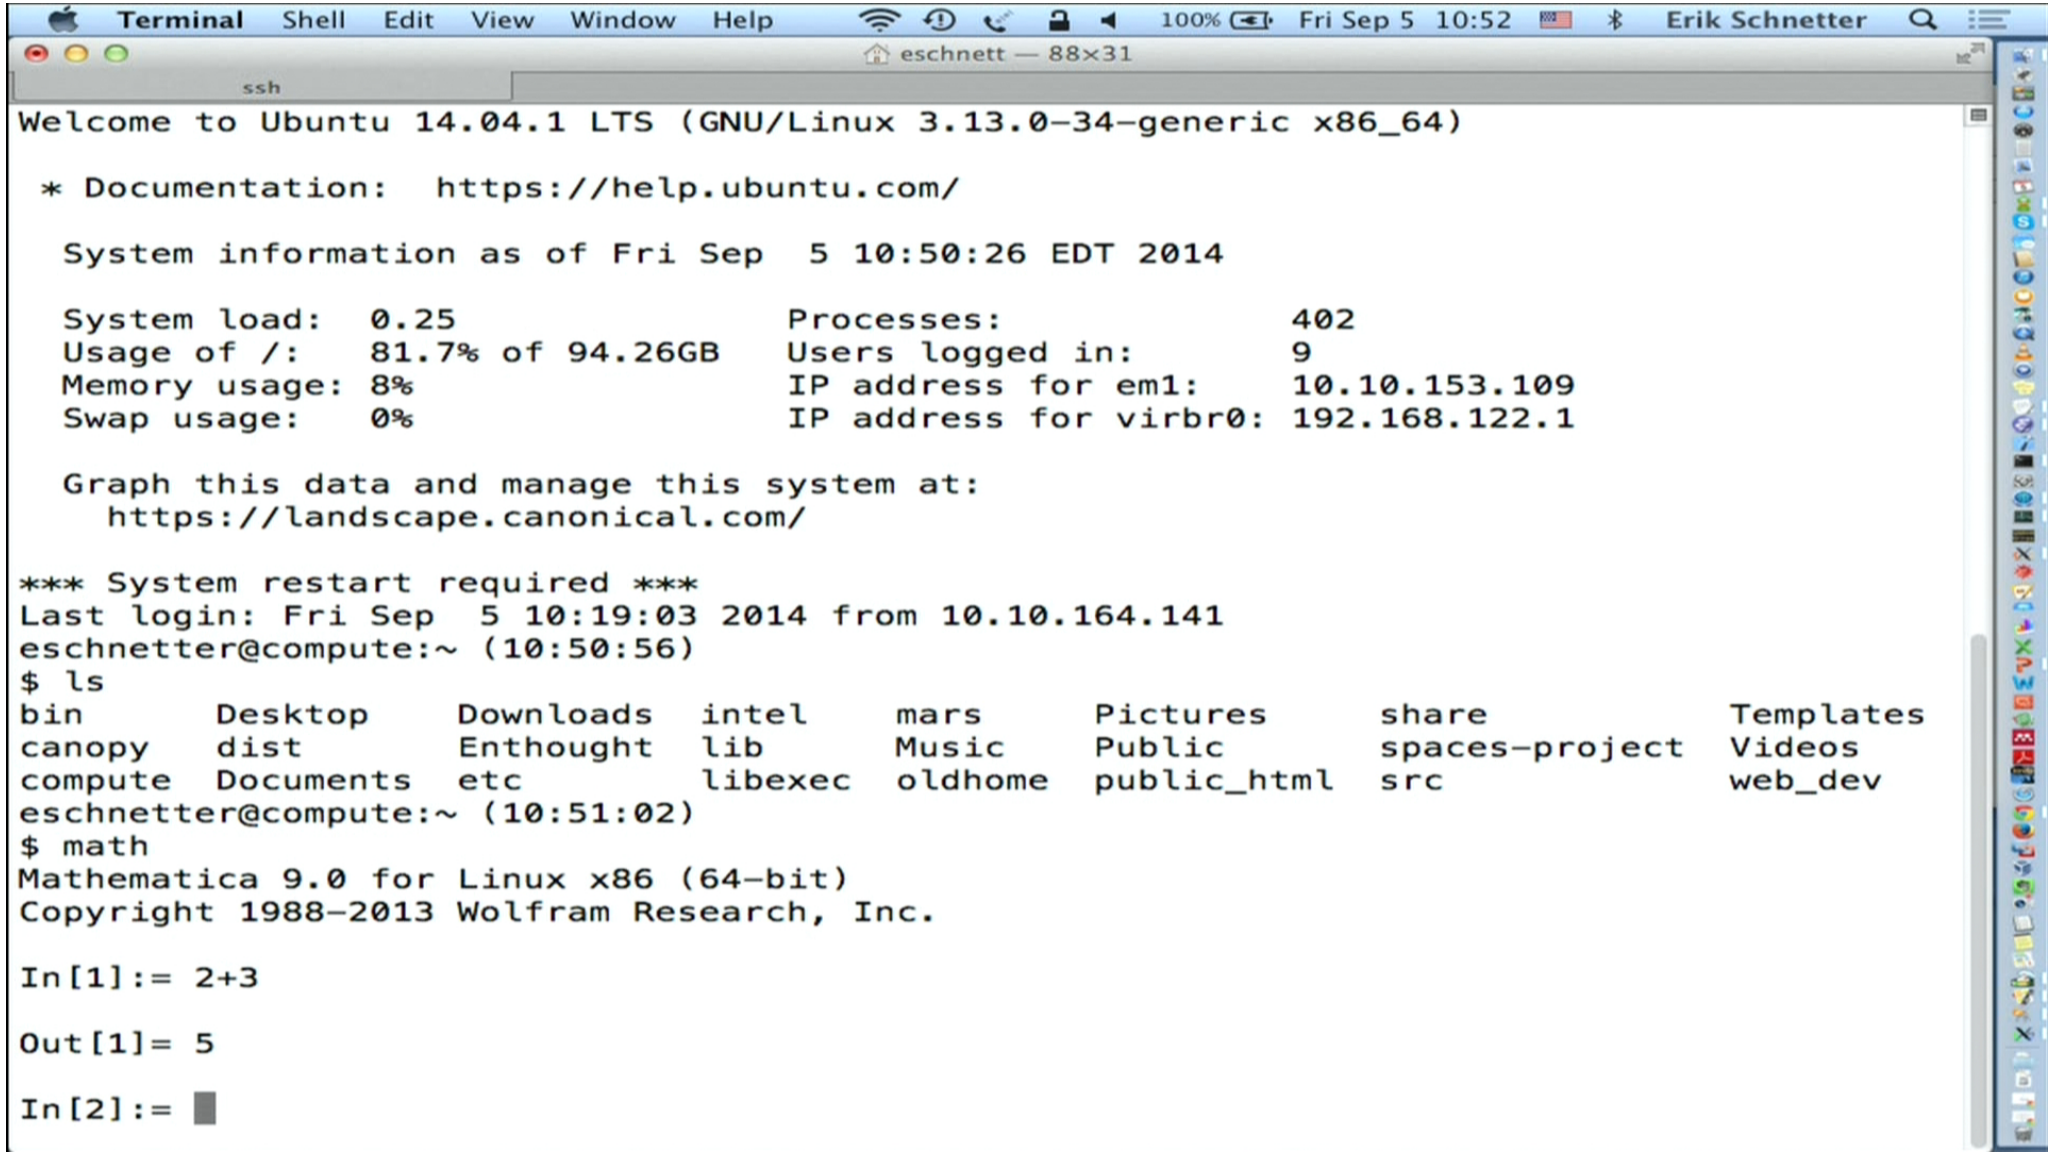
Terminal Shell Edit View Window Help
100% Fri Sep 5 10:52
Erik Schnetter
eschnett — 88×31
ssh
Welcome to Ubuntu 14.04.1 LTS (GNU/Linux 3.13.0-34-generic x86_64)
* Documentation: https://help.ubuntu.com/
System information as of Fri Sep 5 10:50:26 EDT 2014
System load: 0.25
Usage of /: 81.7% of 94.26GB
Memory usage: 8%
Swap usage: 0%
Processes: 402
Users logged in: 9
IP address for em1: 10.10.153.109
IP address for virbr0: 192.168.122.1
Graph this data and manage this system at:
https://landscape.canonical.com/
*** System restart required ***
Last login: Fri Sep 5 10:19:03 2014 from 10.10.164.141
eschnetter@compute:~ (10:50:56)
$ ls
bin Desktop Downloads intel mars Pictures share Templates
canopy dist Enthought lib Music Public spaces-project Videos
compute Documents etc libexec oldhome public_html src web_dev
eschnetter@compute:~ (10:51:02)
$ math
Mathematica 9.0 for Linux x86 (64-bit)
Copyright 1988-2013 Wolfram Research, Inc.
In[1]:= 2+3
Out[1]= 5
In[2]:=

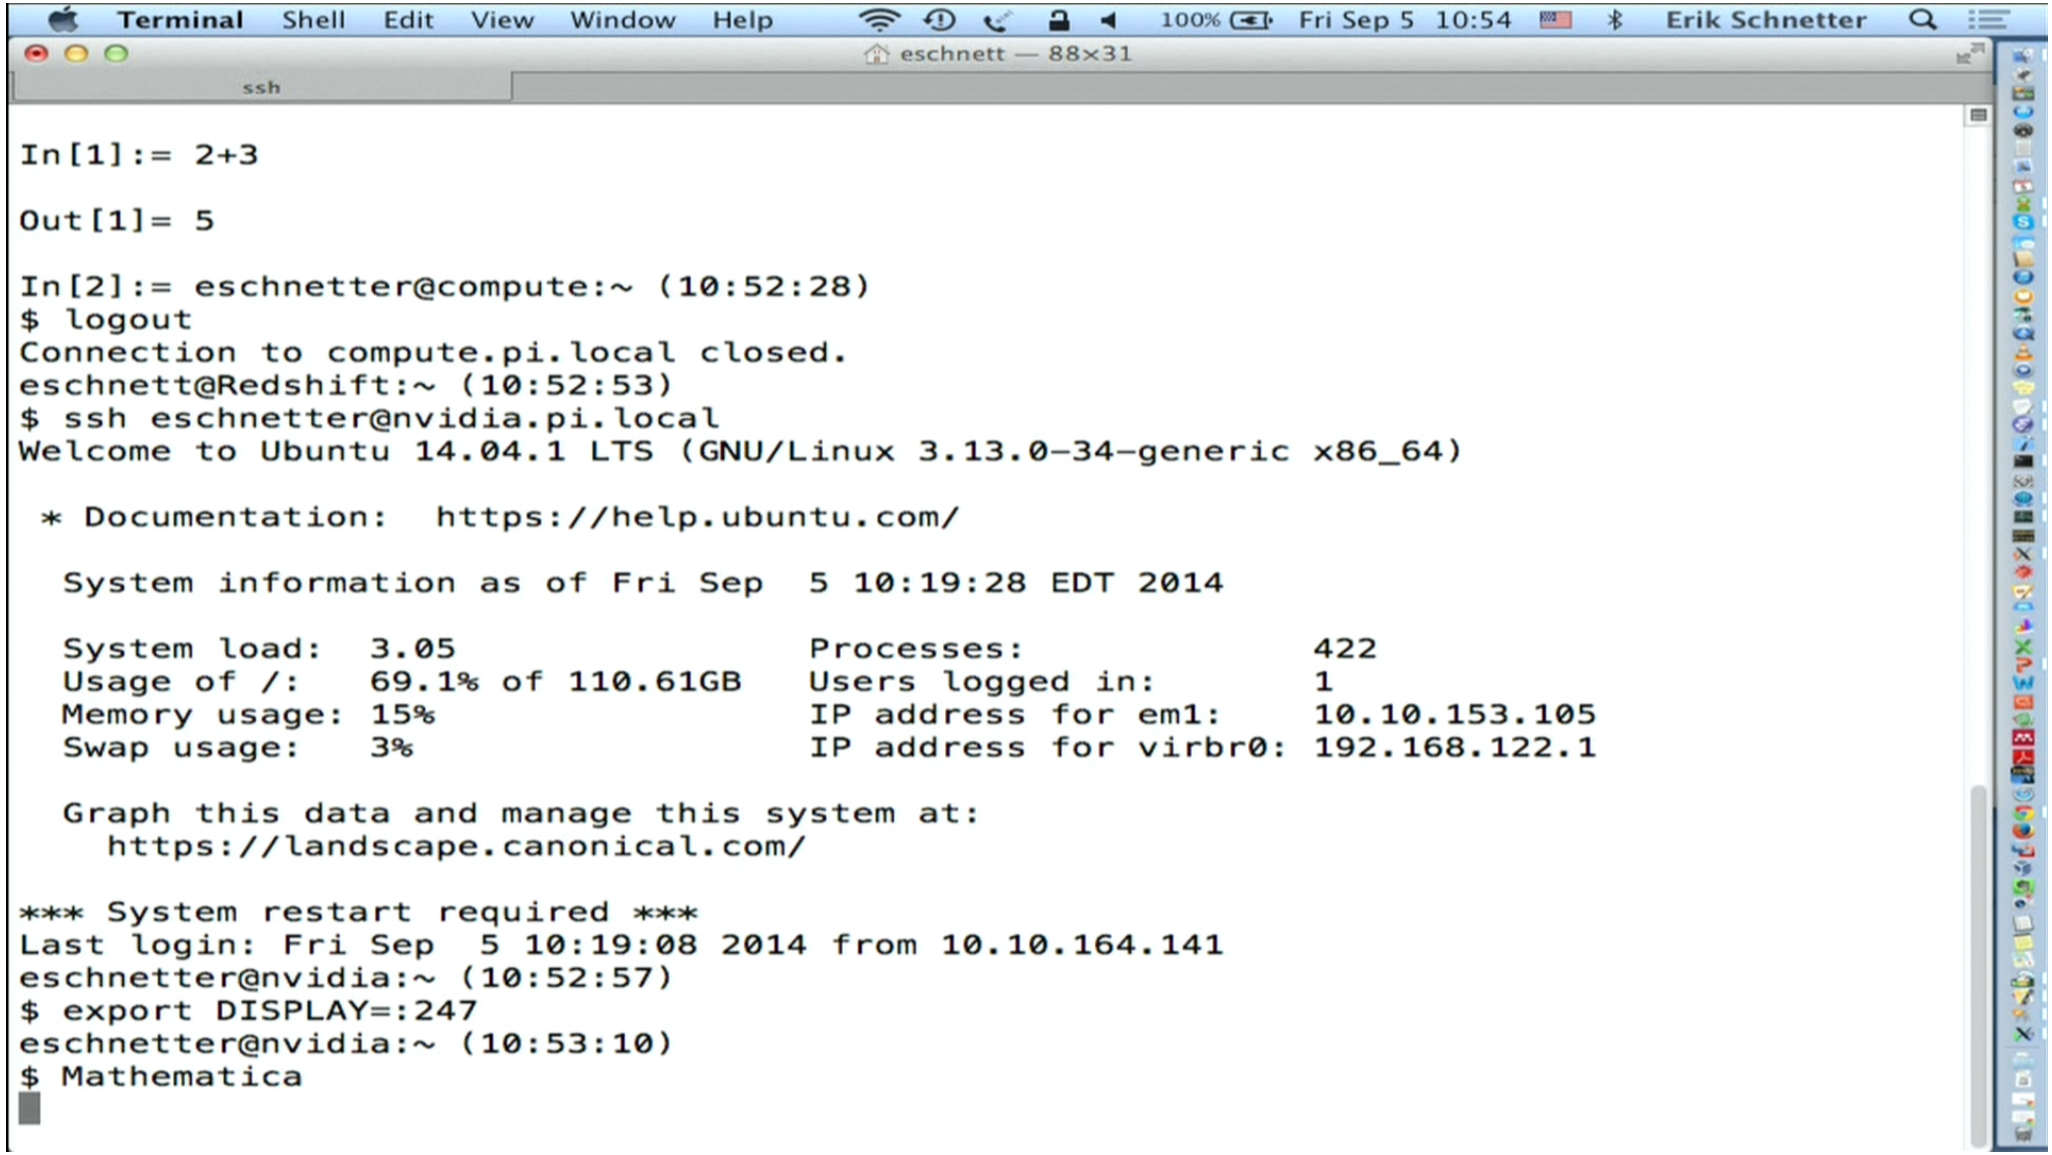
Terminal Shell Edit View Window Help
100% Fri Sep 5 10:54 Erik Schnetter
eschnett — 88×31
ssh
In[1]:= 2+3
Out[1]= 5
In[2]:= eschnetter@compute:~ (10:52:28)
$ logout
Connection to compute.pi.local closed.
eschnett@Redshift:~ (10:52:53)
$ ssh eschnetter@nvidia.pi.local
Welcome to Ubuntu 14.04.1 LTS (GNU/Linux 3.13.0-34-generic x86_64)
* Documentation: https://help.ubuntu.com/
System information as of Fri Sep 5 10:19:28 EDT 2014
System load: 3.05 Processes: 422
Usage of /: 69.1% of 110.61GB Users logged in: 1
Memory usage: 15% IP address for em1: 10.10.153.105
Swap usage: 3% IP address for virbr0: 192.168.122.1
Graph this data and manage this system at:
https://landscape.canonical.com/
*** System restart required ***
Last login: Fri Sep 5 10:19:08 2014 from 10.10.164.141
eschnetter@nvidia:~ (10:52:57)
$ export DISPLAY=:247
eschnetter@nvidia:~ (10:53:10)
$ Mathematica

Scientific Computing Environment

- 7 “regular” Linux workstations, located in a server room
- Names: compute, compute1, …, compute6
- Running Ubuntu, with many additional packages installed
  - 8 Intel cores with 2.4 GHz
  - 40 GByte RAM
  - Using Perimeter home directory
- Also 1 GPU workstation named “Nvidia”, same software, has Nvidia GPU
- Access via ssh:
  - Login: ssh yourlogin@compute.pi.local
  - Copy files: access via network drive
- Mathematica: can use remote desktop

Software Packages on SCE

- Mathematica, Maple, Matlab, Magma, Sage, Ipython, ...

- Python, Perl
- GNU C, C++, Fortran
- CUDA, OpenCL
- Intel compilers (C, C++, Fortran, MKL, debugger/profiler)
- RNPL

- gnuplot
- SuperMongo
- DataVault
- VisIt (visualisation)

- Many more – ask if you need something

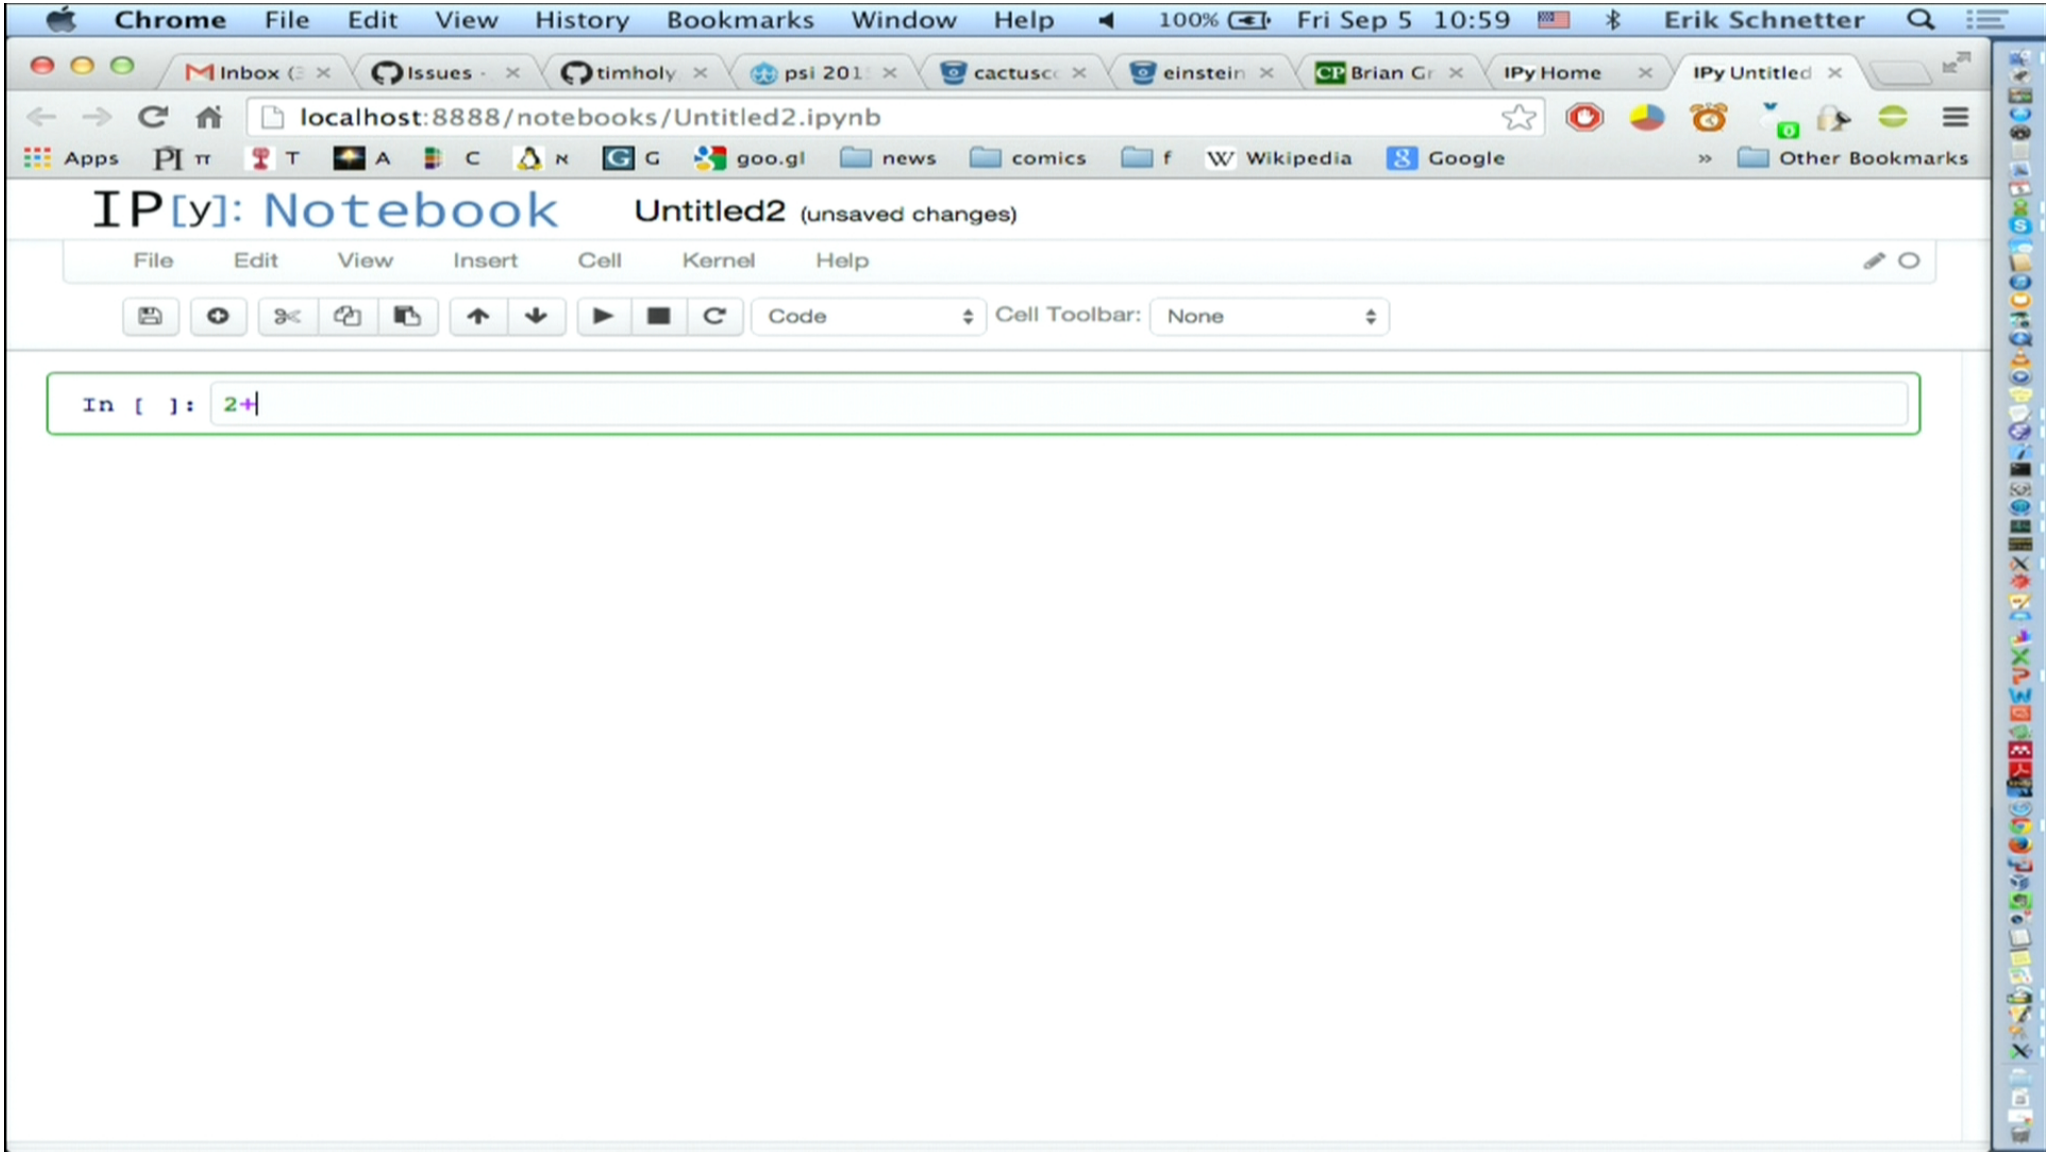
Chrome File Edit View History Bookmarks Window Help 100% Fri Sep 5 10:59 Erik Schnetter
Inbox (3
Issues
timholy
psi 201
cactusc
einstein
Brian Gr
IPy Home
IPy Untitled
localhost:8888/notebooks/Untitled2.ipynb
Apps
goo.gl
news
comics
Wikipedia
Google
Other Bookmarks
IP[y]: Notebook
Untitled2 (unsaved changes)
File Edit View Insert Cell Kernel Help
Code
Cell Toolbar: None
In [ ]: 2+

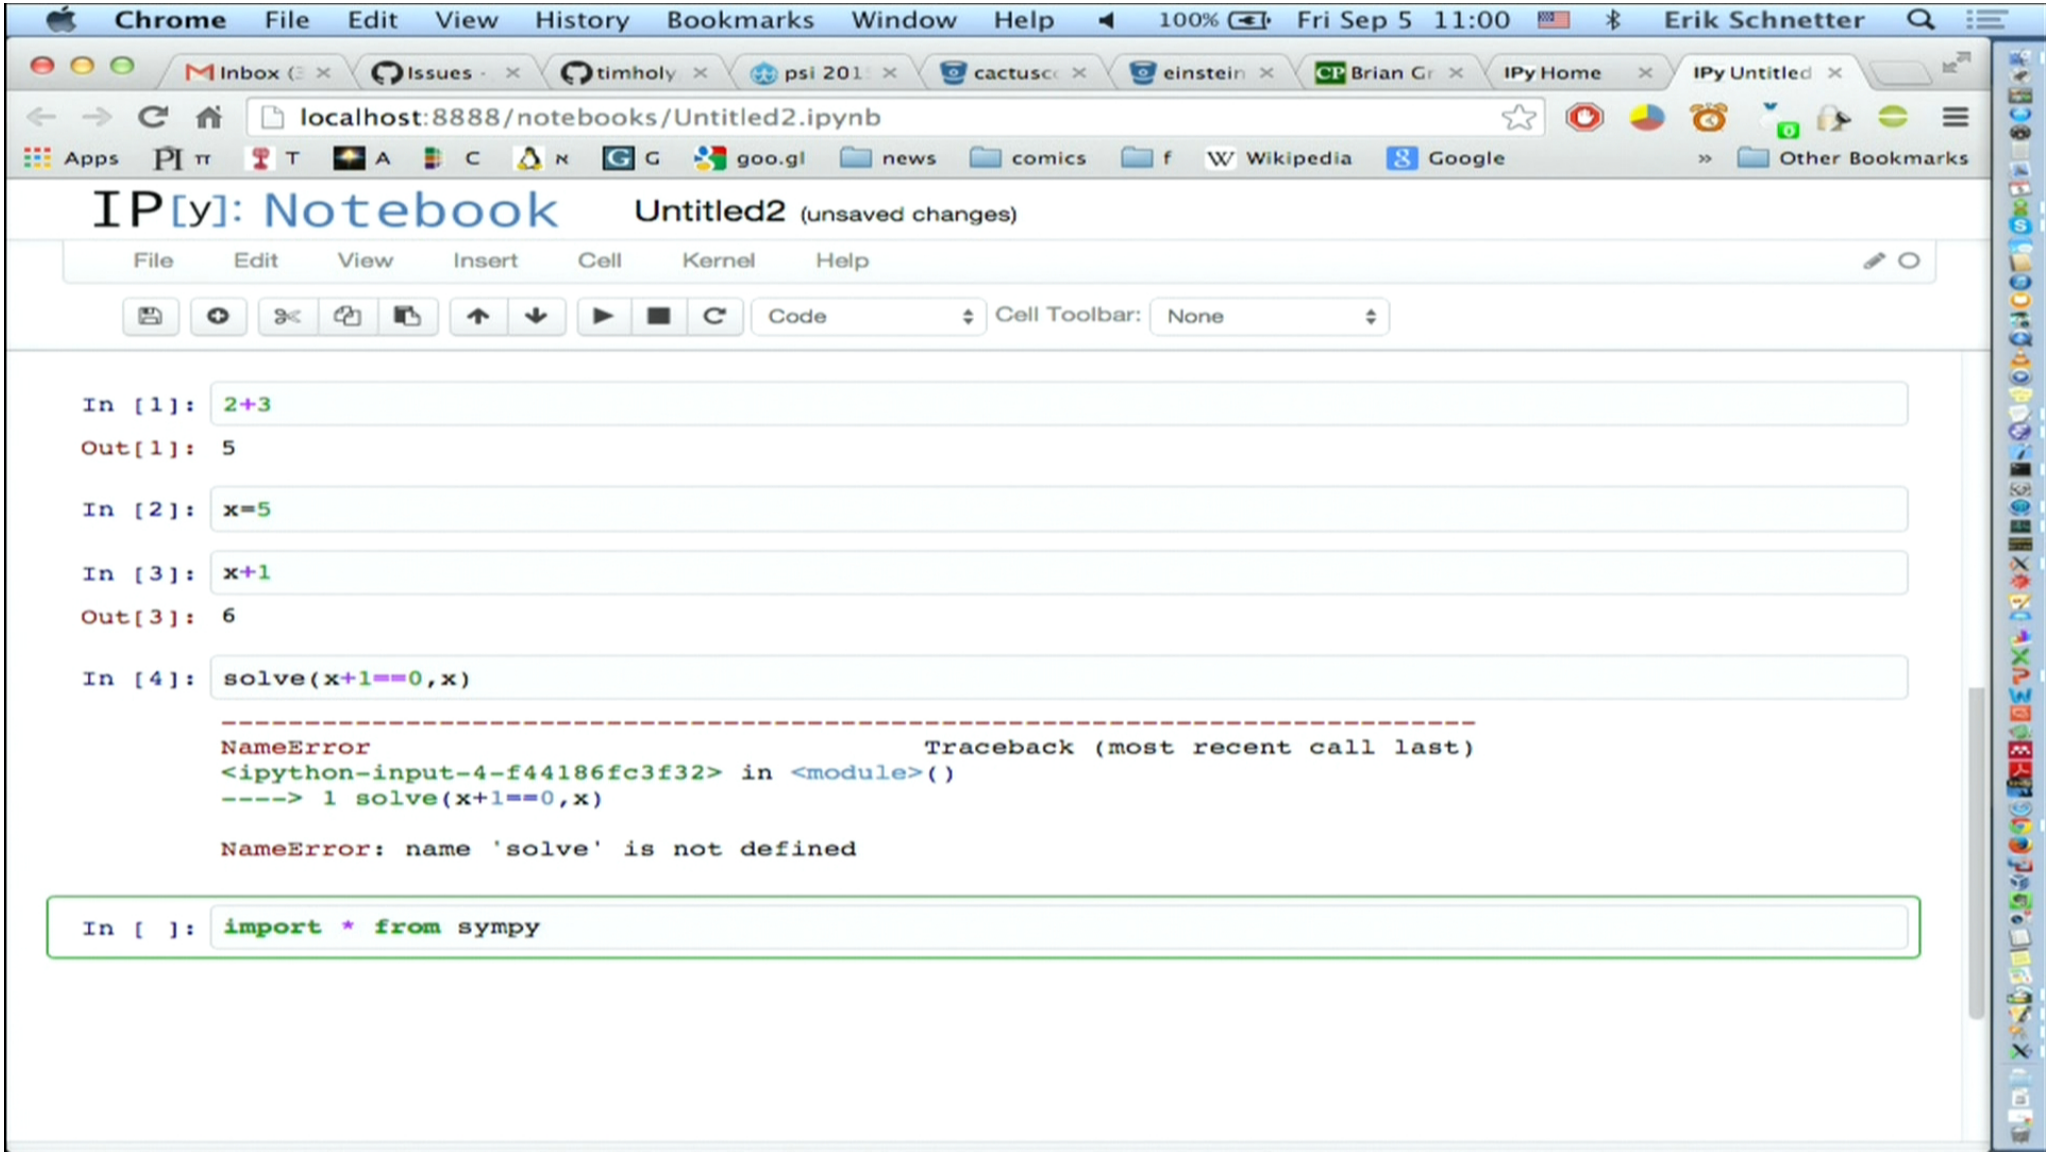
Chrome File Edit View History Bookmarks Window Help 100% Fri Sep 5 11:00 Erik Schnetter
Inbox Issues timholy psi 201 cactusc einstein Brian Gr IPy Home IPy Untitled
localhost:8888/notebooks/Untitled2.ipynb
Apps goo.gl news comics Wikipedia Google Other Bookmarks
IP[y]: Notebook Untitled2 (unsaved changes)
File Edit View Insert Cell Kernel Help
Code Cell Toolbar: None
In [1]: 2+3
Out[1]: 5
In [2]: x=5
In [3]: x+1
Out[3]: 6
In [4]: solve(x+1==0,x)
NameError Traceback (most recent call last)
<ipython-input-4-f44186fc3f32> in <module>()
----> 1 solve(x+1==0,x)
NameError: name 'solve' is not defined
In [ ]: import * from sympy

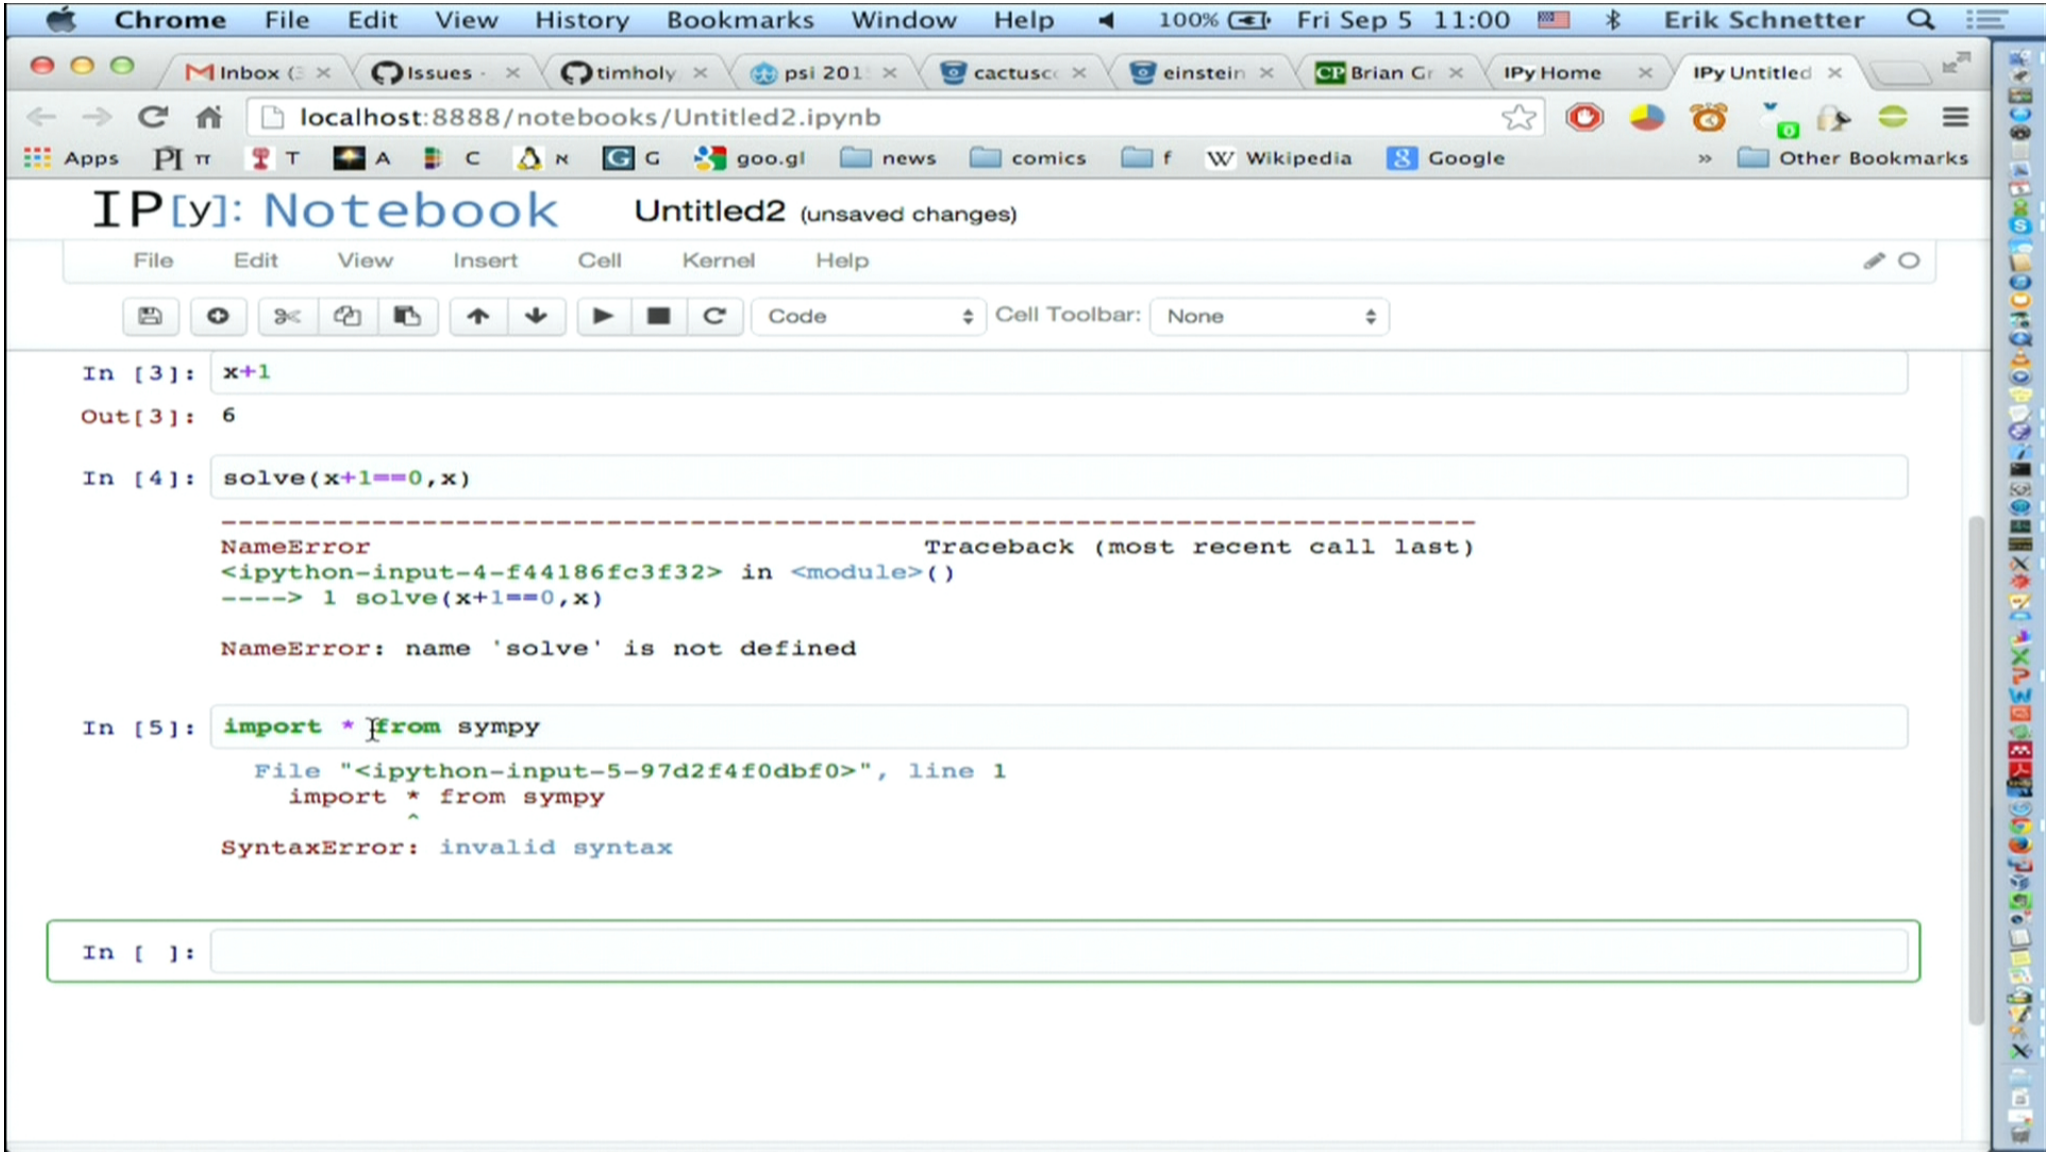
Chrome File Edit View History Bookmarks Window Help 100% Fri Sep 5 11:00 Erik Schnetter
Inbox Issues timholy psi 201 cactusc einstein Brian Gr IPy Home IPy Untitled
localhost:8888/notebooks/Untitled2.ipynb
Apps goo.gl news comics Wikipedia Google Other Bookmarks
IP[y]: Notebook Untitled2 (unsaved changes)
File Edit View Insert Cell Kernel Help
Code Cell Toolbar: None
In [3]: x+1
Out[3]: 6
In [4]: solve(x+1==0,x)
NameError Traceback (most recent call last)
<ipython-input-4-f44186fc3f32> in <module>()
----> 1 solve(x+1==0,x)
NameError: name 'solve' is not defined
In [5]: import * from sympy
File "<ipython-input-5-97d2f4f0dbf0>", line 1
import * from sympy
^
SyntaxError: invalid syntax
In [ ]:

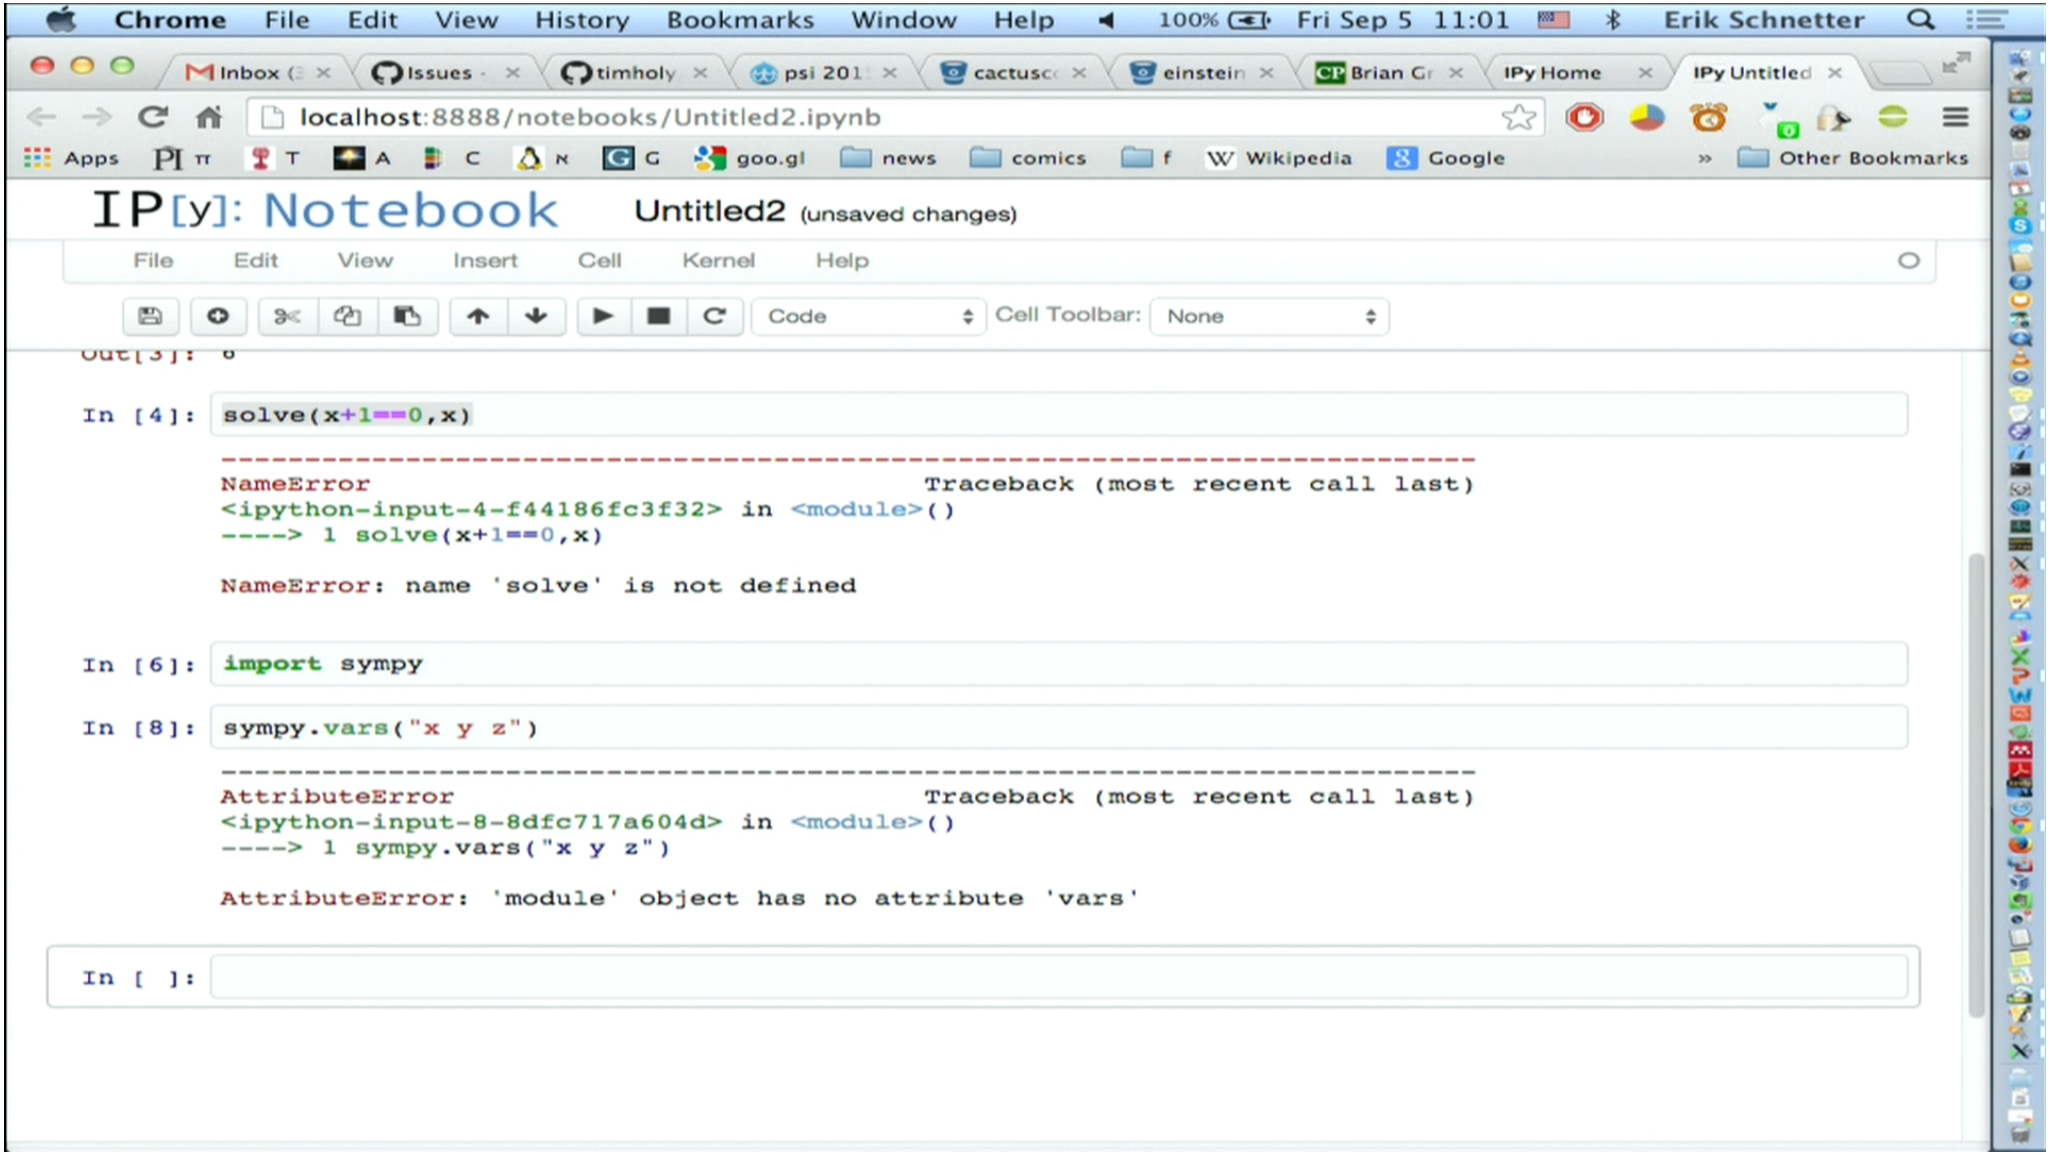
Chrome File Edit View History Bookmarks Window Help 100% Fri Sep 5 11:01 Erik Schnetter
Inbox Issues timholy psi 201 cactusc einstein Brian Gr IPy Home IPy Untitled
localhost:8888/notebooks/Untitled2.ipynb
Apps goo.gl news comics f Wikipedia Google Other Bookmarks
IP[y]: Notebook Untitled2 (unsaved changes)
File Edit View Insert Cell Kernel Help
Code Cell Toolbar: None
Out[3]: 6
In [4]: solve(x+1==0,x)
NameError Traceback (most recent call last)
<ipython-input-4-f44186fc3f32> in <module>()
----> 1 solve(x+1==0,x)
NameError: name 'solve' is not defined
In [6]: import sympy
In [8]: sympy.vars("x y z")
AttributeError Traceback (most recent call last)
<ipython-input-8-8dfc717a604d> in <module>()
----> 1 sympy.vars("x y z")
AttributeError: 'module' object has no attribute 'vars'
In [ ]:

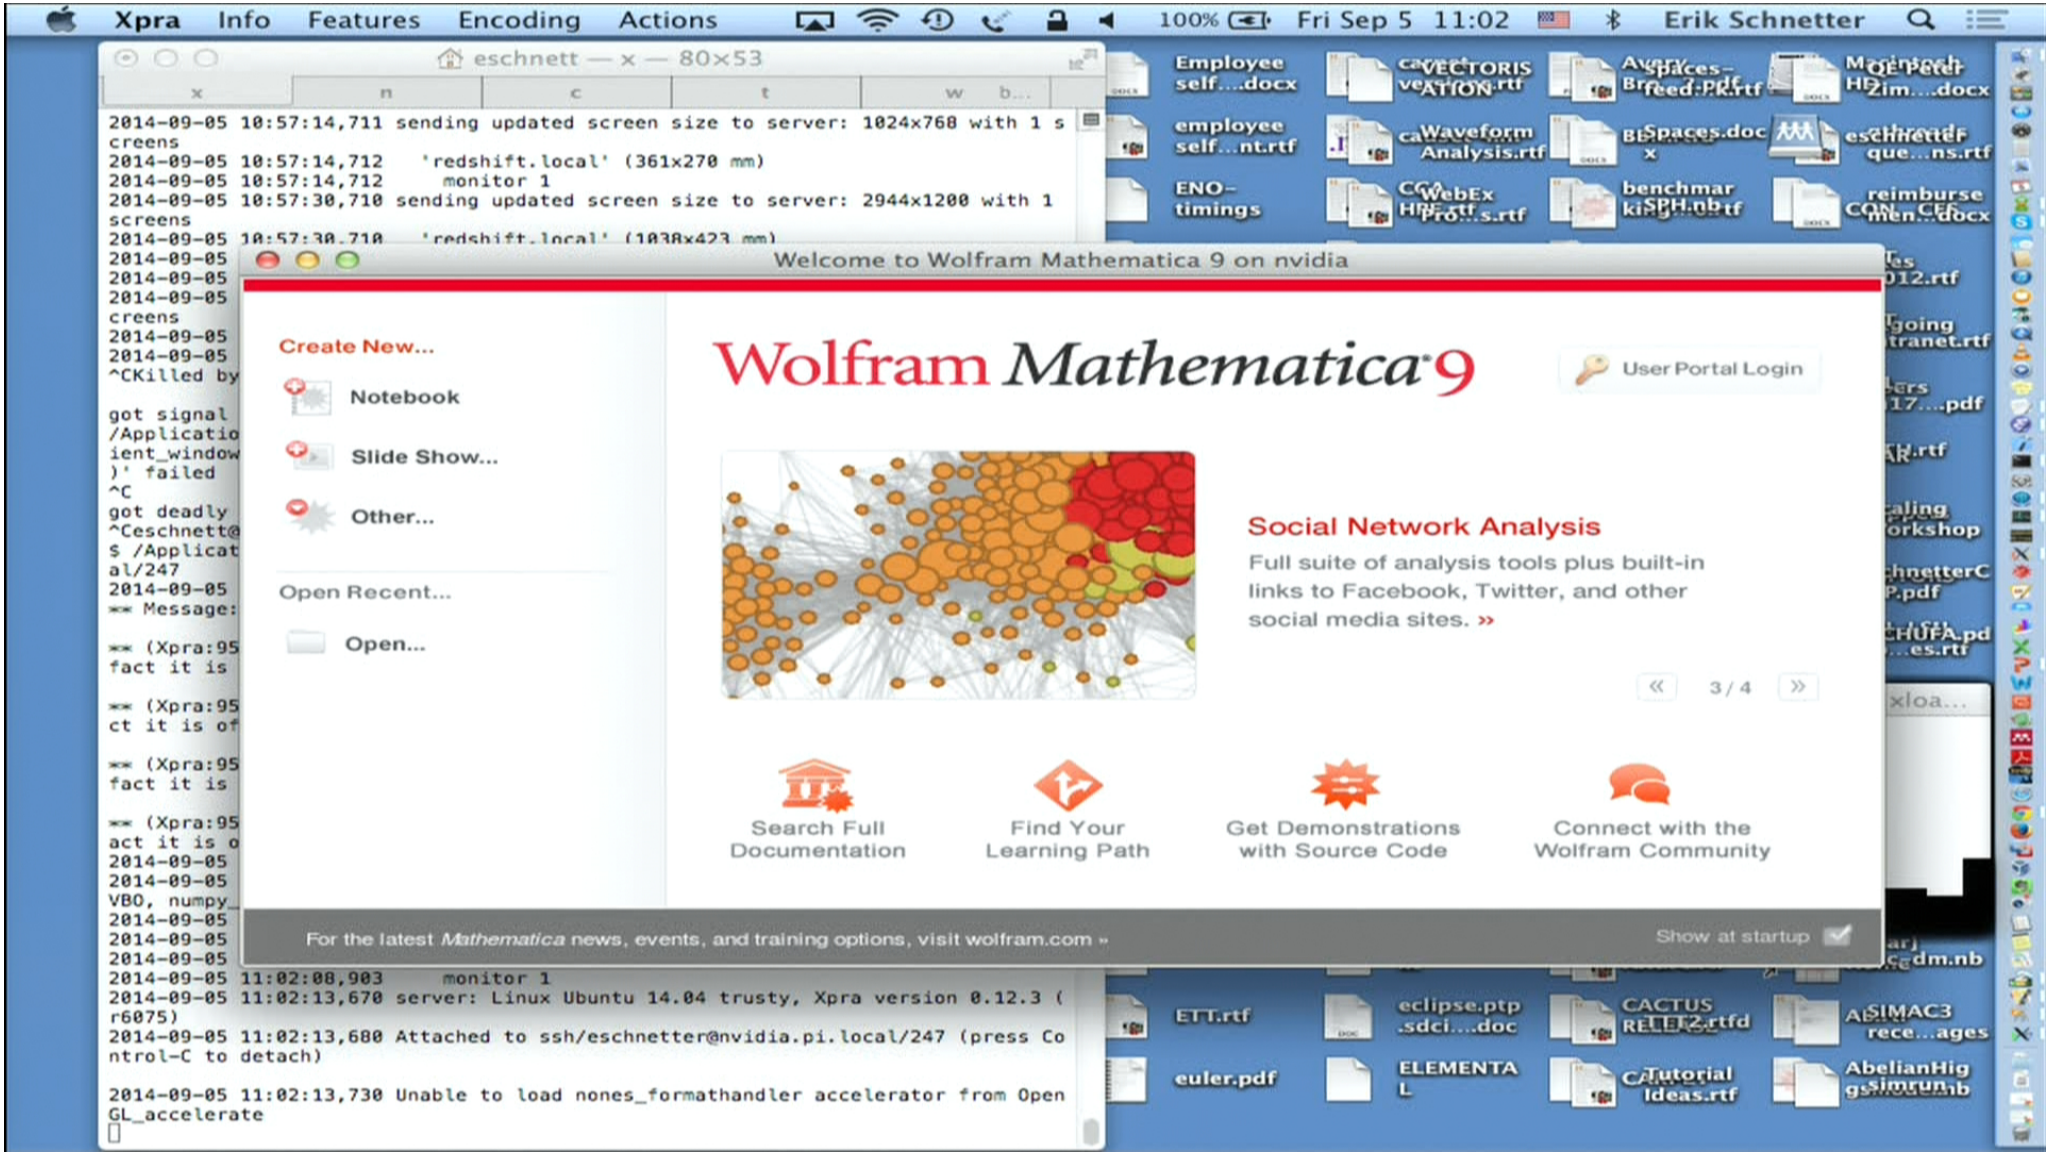
Xpra Info Features Encoding Actions
100% Fri Sep 5 11:02 Erik Schnetter
eschnett — x — 80×53
2014-09-05 10:57:14,711 sending updated screen size to server: 1024x768 with 1 screens
2014-09-05 10:57:14,712 'redshift.local' (361x270 mm)
2014-09-05 10:57:14,712 monitor 1
2014-09-05 10:57:30,710 sending updated screen size to server: 2944x1200 with 1 screens
2014-09-05 11:02:08,903 monitor 1
2014-09-05 11:02:13,670 server: Linux Ubuntu 14.04 trusty, Xpra version 0.12.3 (r6075)
2014-09-05 11:02:13,680 Attached to ssh/eschnetter@nvidia.pi.local/247 (press Control-C to detach)
2014-09-05 11:02:13,730 Unable to load nones_formathandler accelerator from OpenGL_accelerate
Welcome to Wolfram Mathematica 9 on nvidia
Create New...
Notebook
Slide Show...
Other...
Open Recent...
Open...
Wolfram Mathematica 9
User Portal Login
Social Network Analysis
Full suite of analysis tools plus built-in links to Facebook, Twitter, and other social media sites. »
3 / 4
Search Full Documentation
Find Your Learning Path
Get Demonstrations with Source Code
Connect with the Wolfram Community
For the latest Mathematica news, events, and training options, visit wolfram.com »
Show at startup

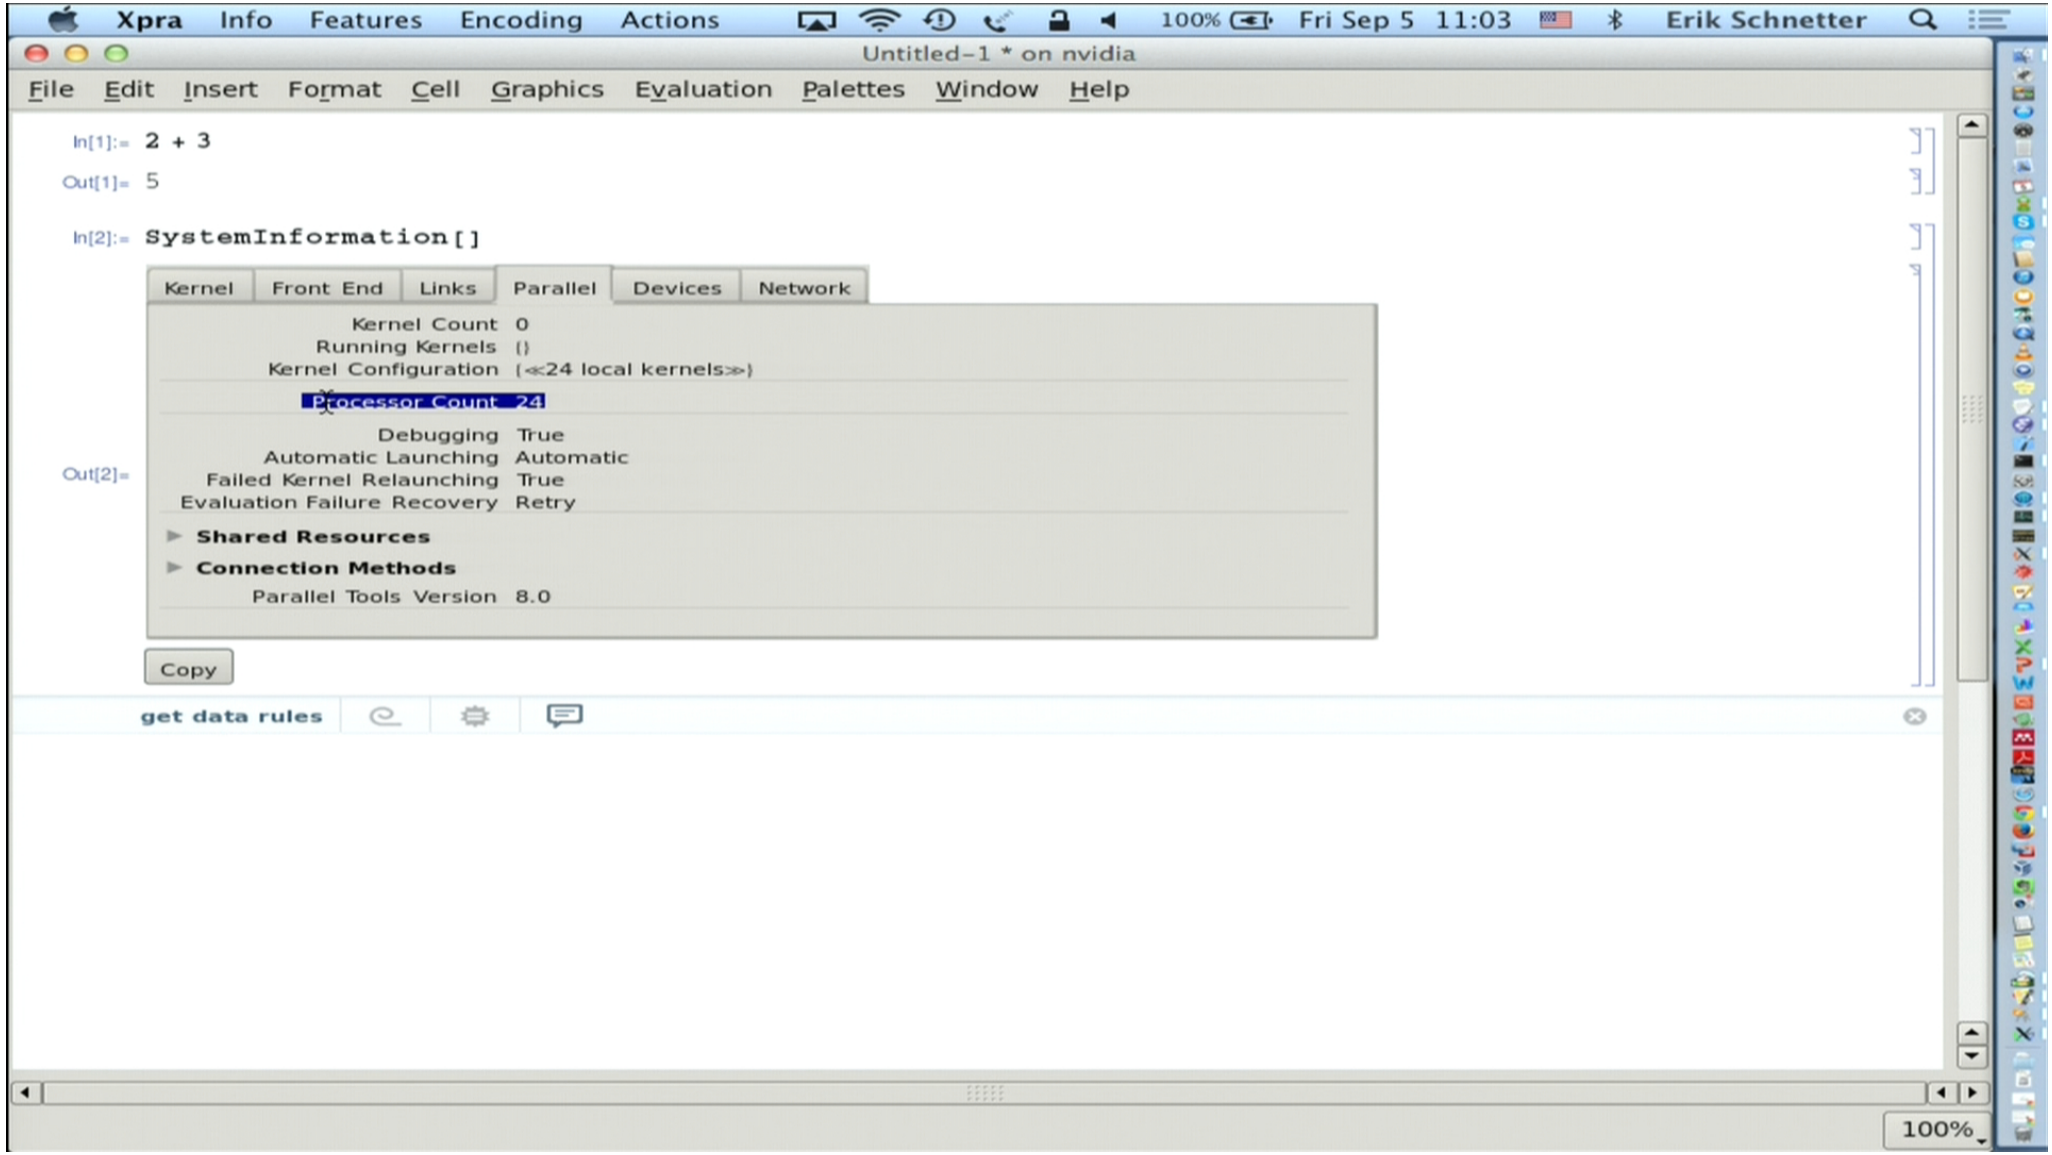
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:03 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
In[1]:= 2 + 3
Out[1]= 5
In[2]:= SystemInformation[]
Kernel Front End Links Parallel Devices Network
Kernel Count 0
Running Kernels {}
Kernel Configuration {≪24 local kernels≫}
Processor Count 24
Debugging True
Automatic Launching Automatic
Failed Kernel Relaunching True
Evaluation Failure Recovery Retry
Shared Resources
Connection Methods
Parallel Tools Version 8.0
Out[2]=
Copy
get data rules
100%

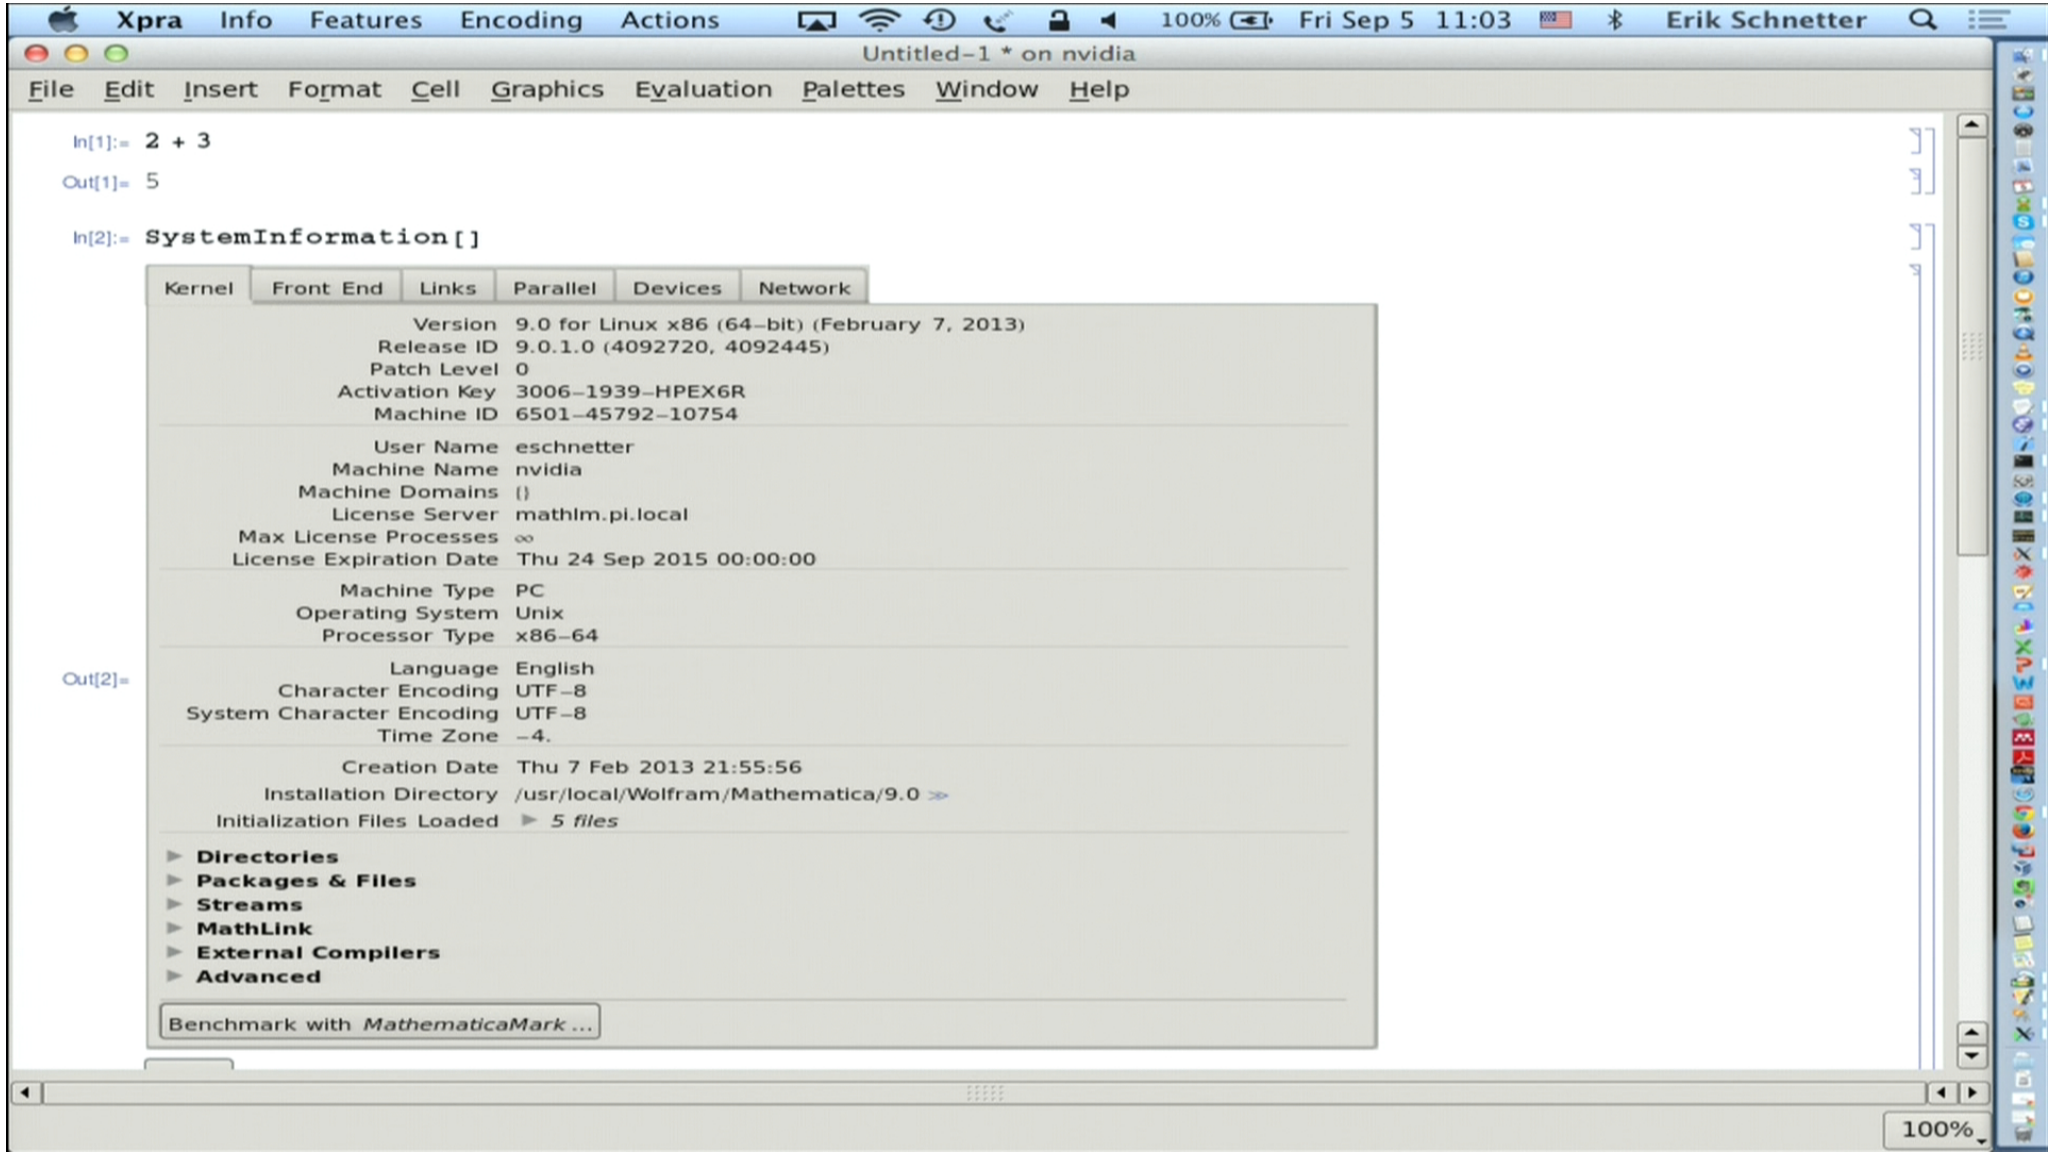
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:03 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
In[1]:= 2 + 3
Out[1]= 5
In[2]:= SystemInformation[]
Kernel Front End Links Parallel Devices Network
Version 9.0 for Linux x86 (64-bit) (February 7, 2013)
Release ID 9.0.1.0 (4092720, 4092445)
Patch Level 0
Activation Key 3006-1939-HPEX6R
Machine ID 6501-45792-10754
User Name eschnetter
Machine Name nvidia
Machine Domains {}
License Server mathlm.pi.local
Max License Processes ∞
License Expiration Date Thu 24 Sep 2015 00:00:00
Machine Type PC
Operating System Unix
Processor Type x86-64
Out[2]=
Language English
Character Encoding UTF-8
System Character Encoding UTF-8
Time Zone -4.
Creation Date Thu 7 Feb 2013 21:55:56
Installation Directory /usr/local/Wolfram/Mathematica/9.0
Initialization Files Loaded 5 files
Directories
Packages & Files
Streams
MathLink
External Compilers
Advanced
Benchmark with MathematicaMark ...
100%

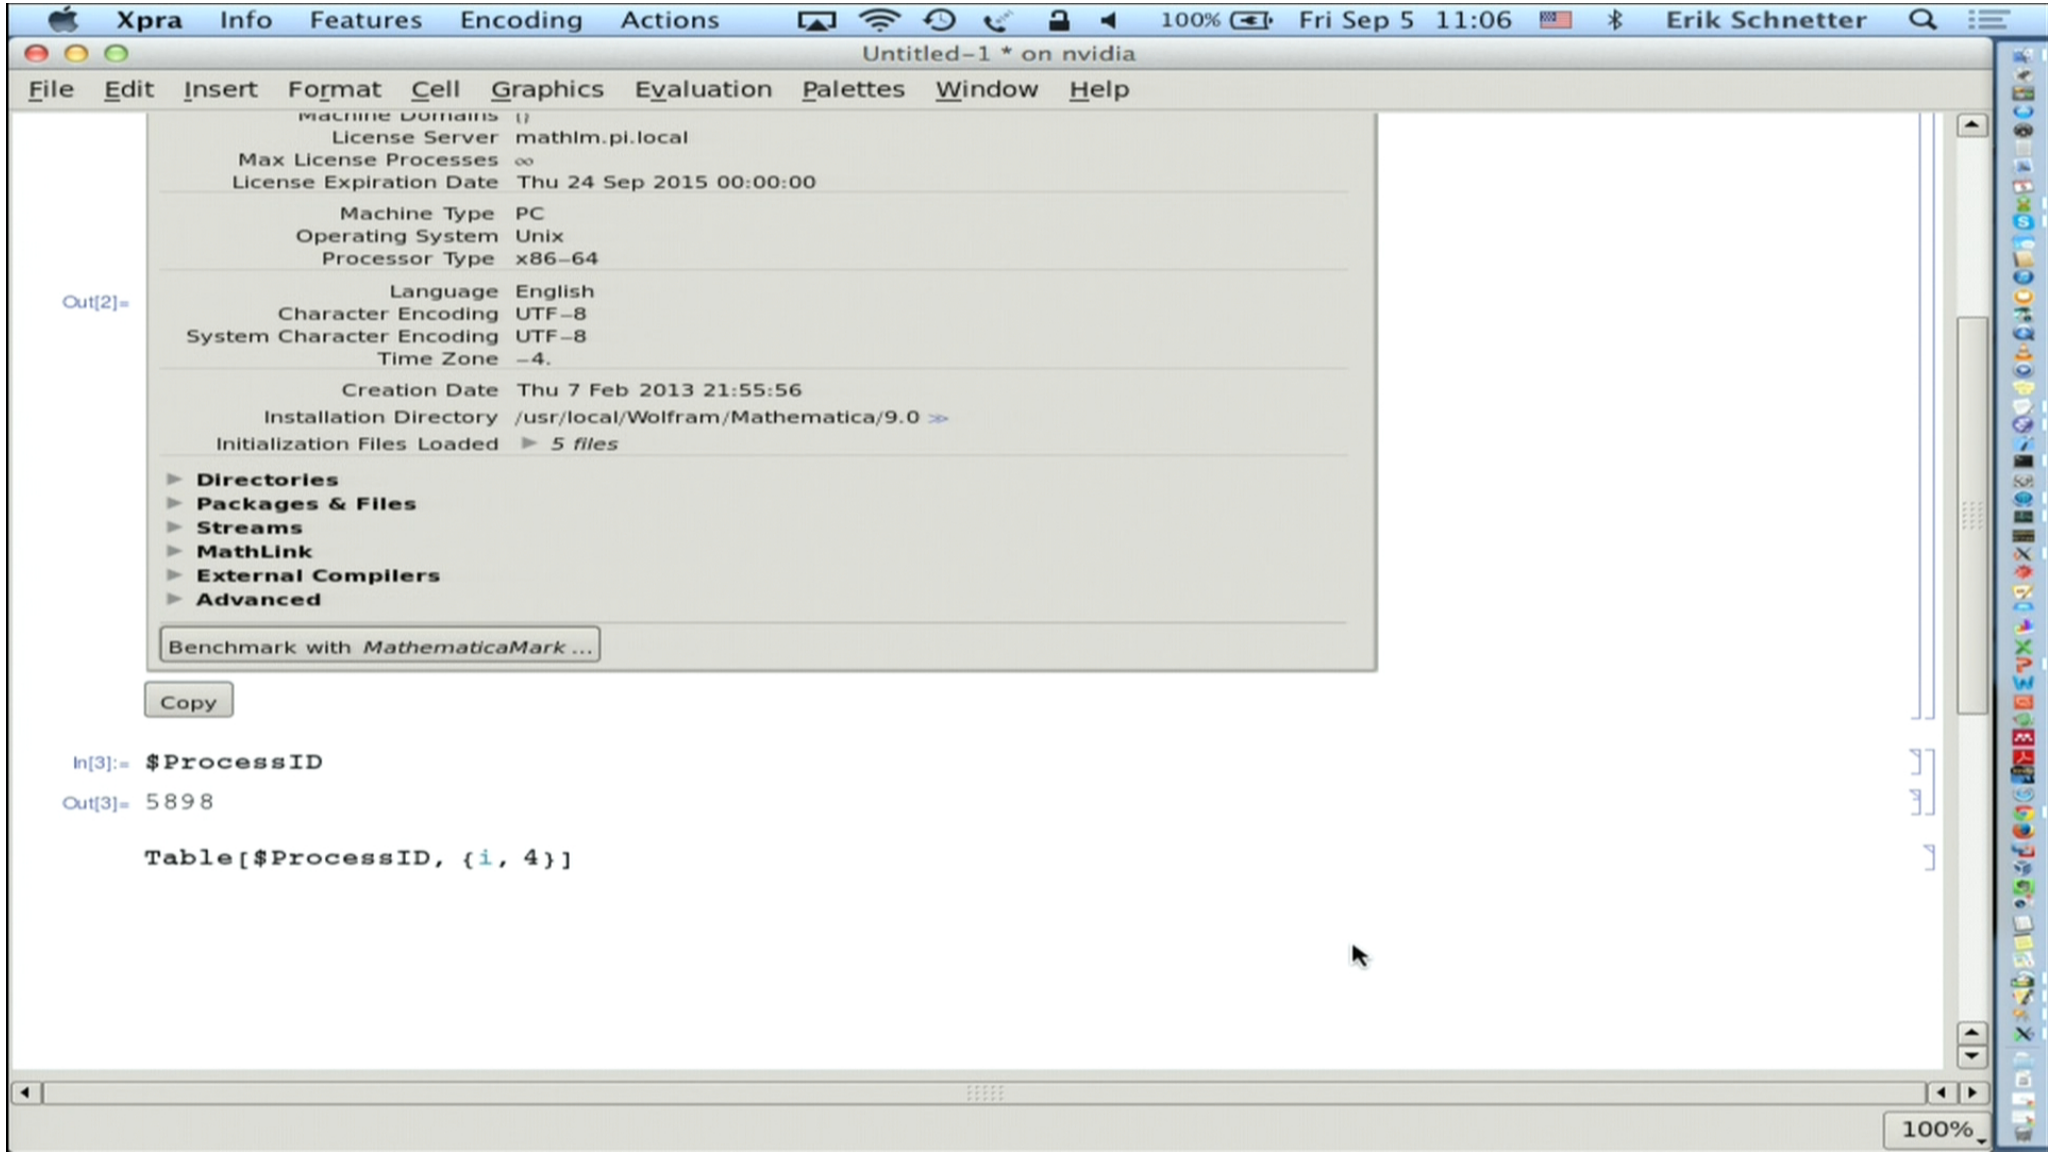
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:06 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
License Server mathlm.pi.local
Max License Processes ∞
License Expiration Date Thu 24 Sep 2015 00:00:00
Machine Type PC
Operating System Unix
Processor Type x86-64
Out[2]=
Language English
Character Encoding UTF-8
System Character Encoding UTF-8
Time Zone -4.
Creation Date Thu 7 Feb 2013 21:55:56
Installation Directory /usr/local/Wolfram/Mathematica/9.0
Initialization Files Loaded 5 files
Directories
Packages & Files
Streams
MathLink
External Compilers
Advanced
Benchmark with MathematicaMark ...
Copy
In[3]:= $ProcessID
Out[3]= 5898
Table[$ProcessID, {i, 4}]
100%

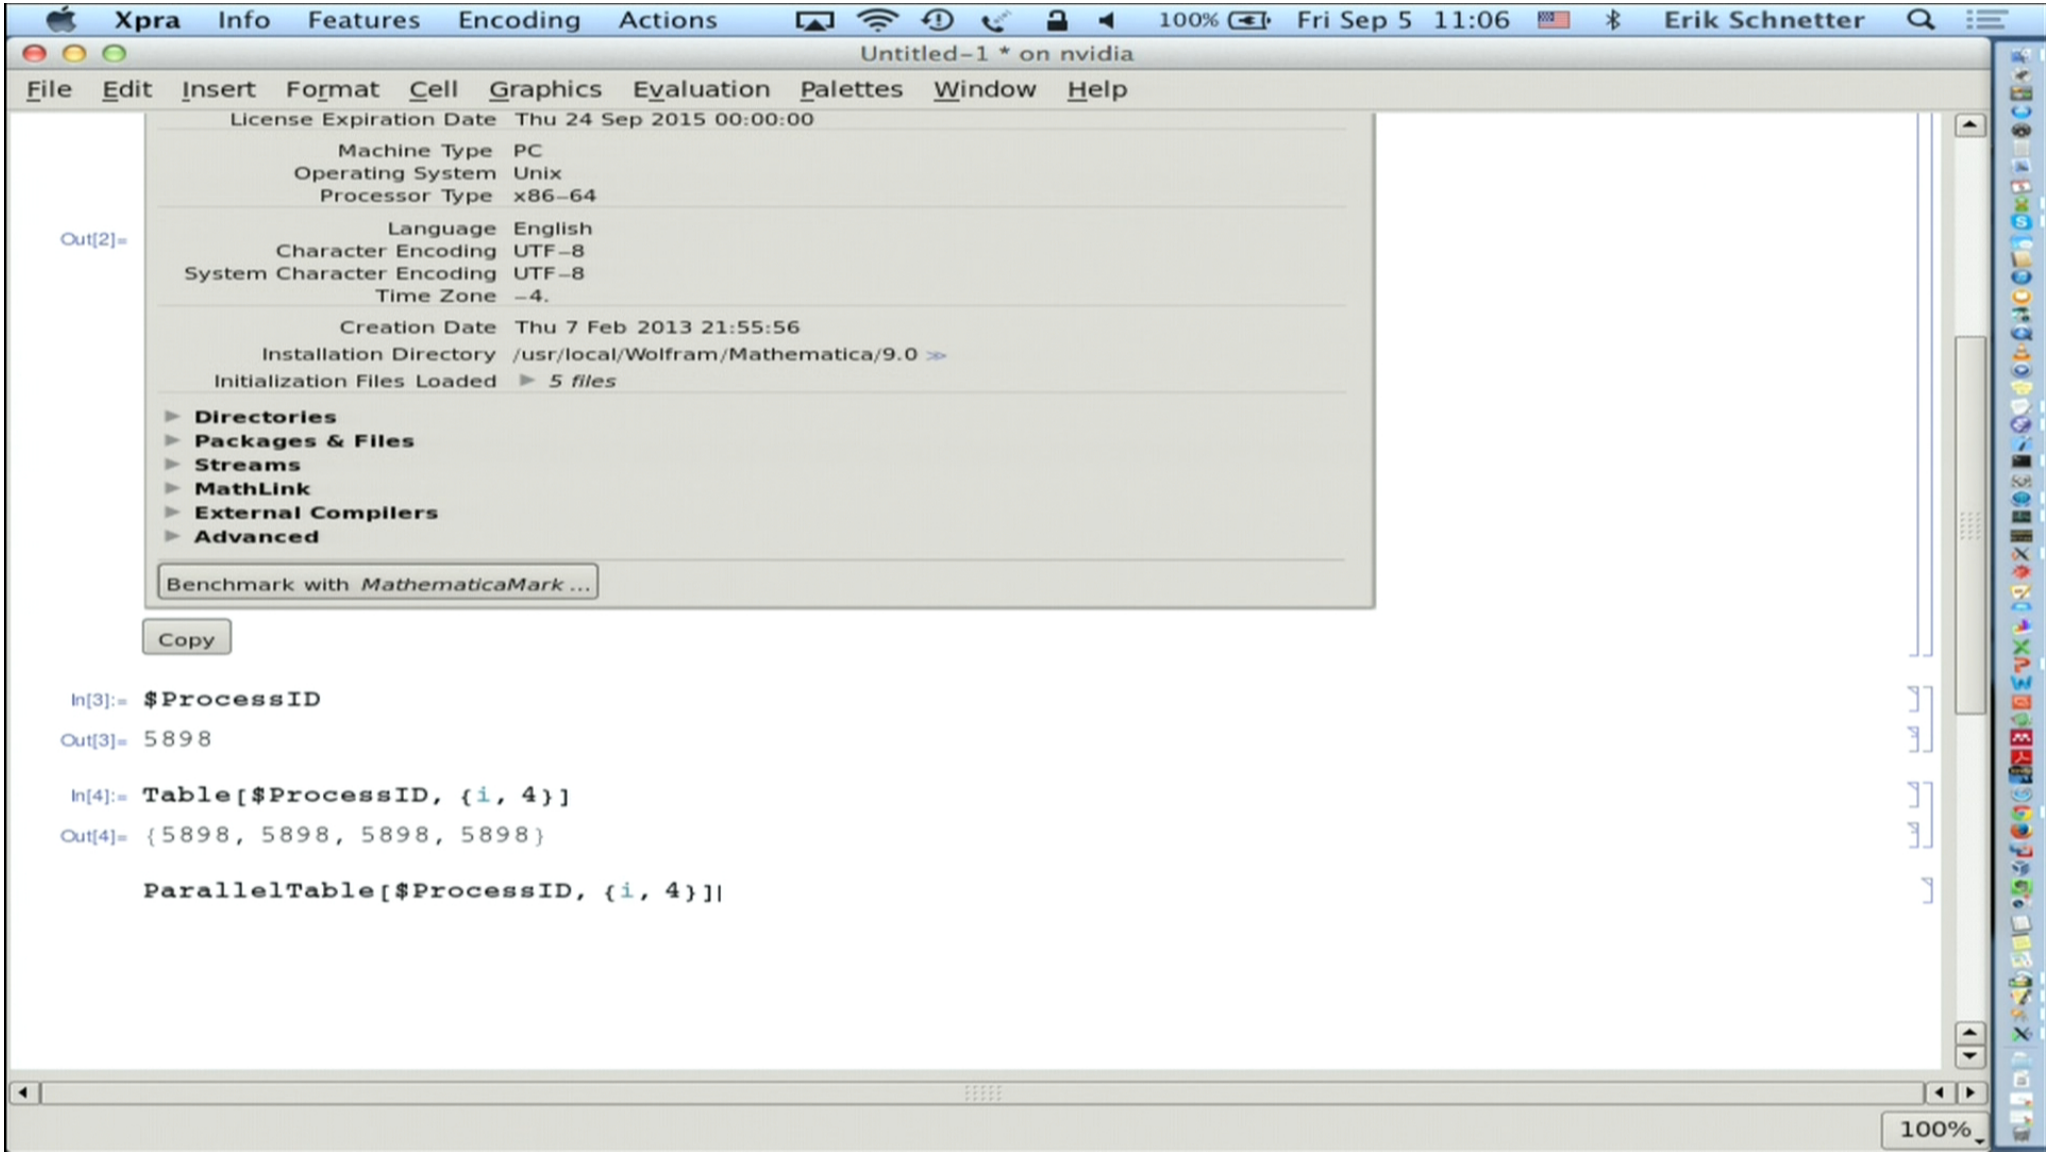
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:06 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
License Expiration Date Thu 24 Sep 2015 00:00:00
Machine Type PC
Operating System Unix
Processor Type x86-64
Out[2]=
Language English
Character Encoding UTF-8
System Character Encoding UTF-8
Time Zone -4.
Creation Date Thu 7 Feb 2013 21:55:56
Installation Directory /usr/local/Wolfram/Mathematica/9.0
Initialization Files Loaded 5 files
Directories
Packages & Files
Streams
MathLink
External Compilers
Advanced
Benchmark with MathematicaMark ...
Copy
In[3]:= $ProcessID
Out[3]= 5898
In[4]:= Table[$ProcessID, {i, 4}]
Out[4]= {5898, 5898, 5898, 5898}
ParallelTable[$ProcessID, {i, 4}]
100%

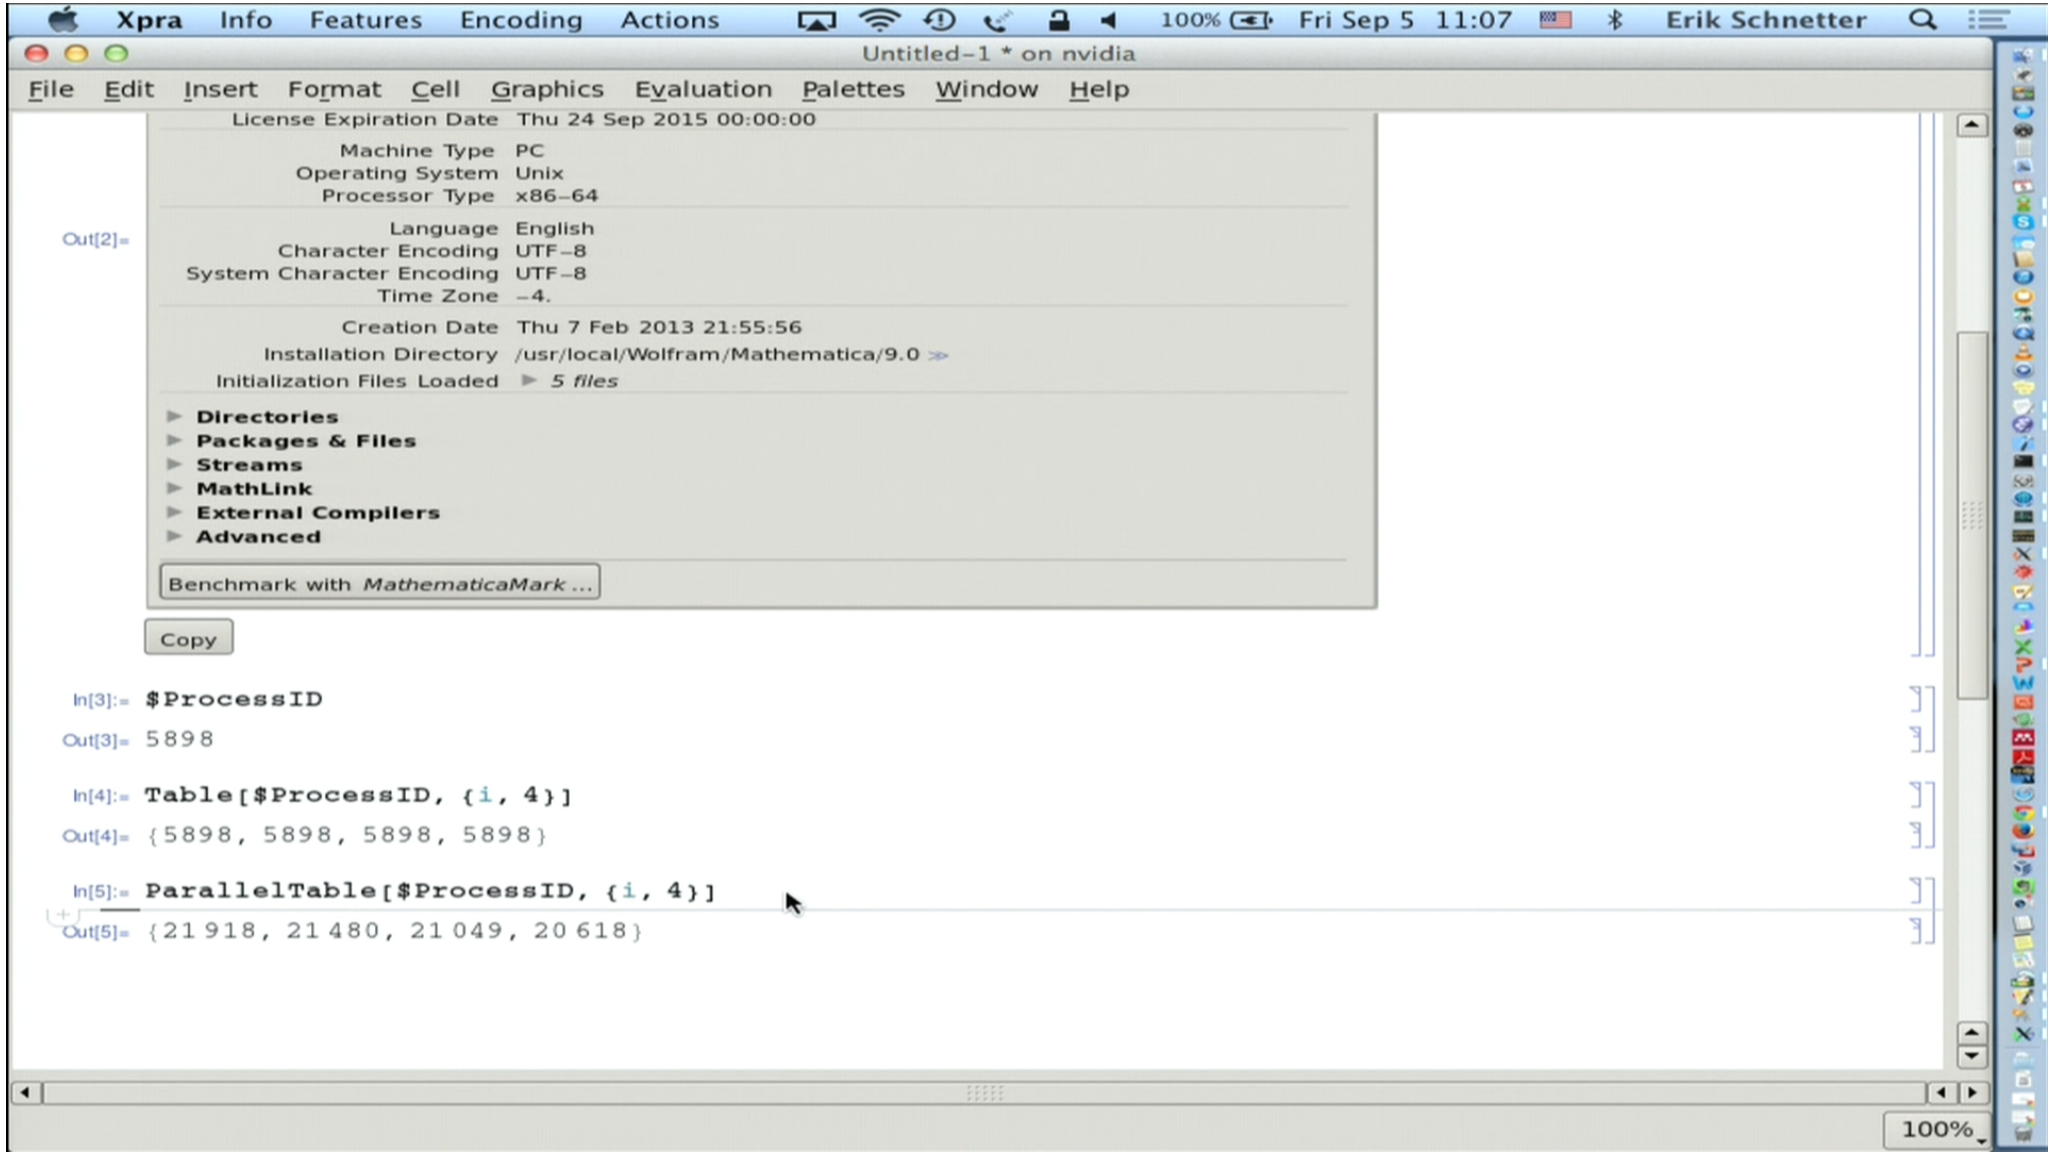
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:07 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
License Expiration Date Thu 24 Sep 2015 00:00:00
Machine Type PC
Operating System Unix
Processor Type x86-64
Out[2]=
Language English
Character Encoding UTF-8
System Character Encoding UTF-8
Time Zone -4.
Creation Date Thu 7 Feb 2013 21:55:56
Installation Directory /usr/local/Wolfram/Mathematica/9.0
Initialization Files Loaded 5 files
Directories
Packages & Files
Streams
MathLink
External Compilers
Advanced
Benchmark with MathematicaMark ...
Copy
In[3]:= $ProcessID
Out[3]= 5898
In[4]:= Table[$ProcessID, {i, 4}]
Out[4]= {5898, 5898, 5898, 5898}
In[5]:= ParallelTable[$ProcessID, {i, 4}]
Out[5]= {21 918, 21 480, 21 049, 20 618}
100%

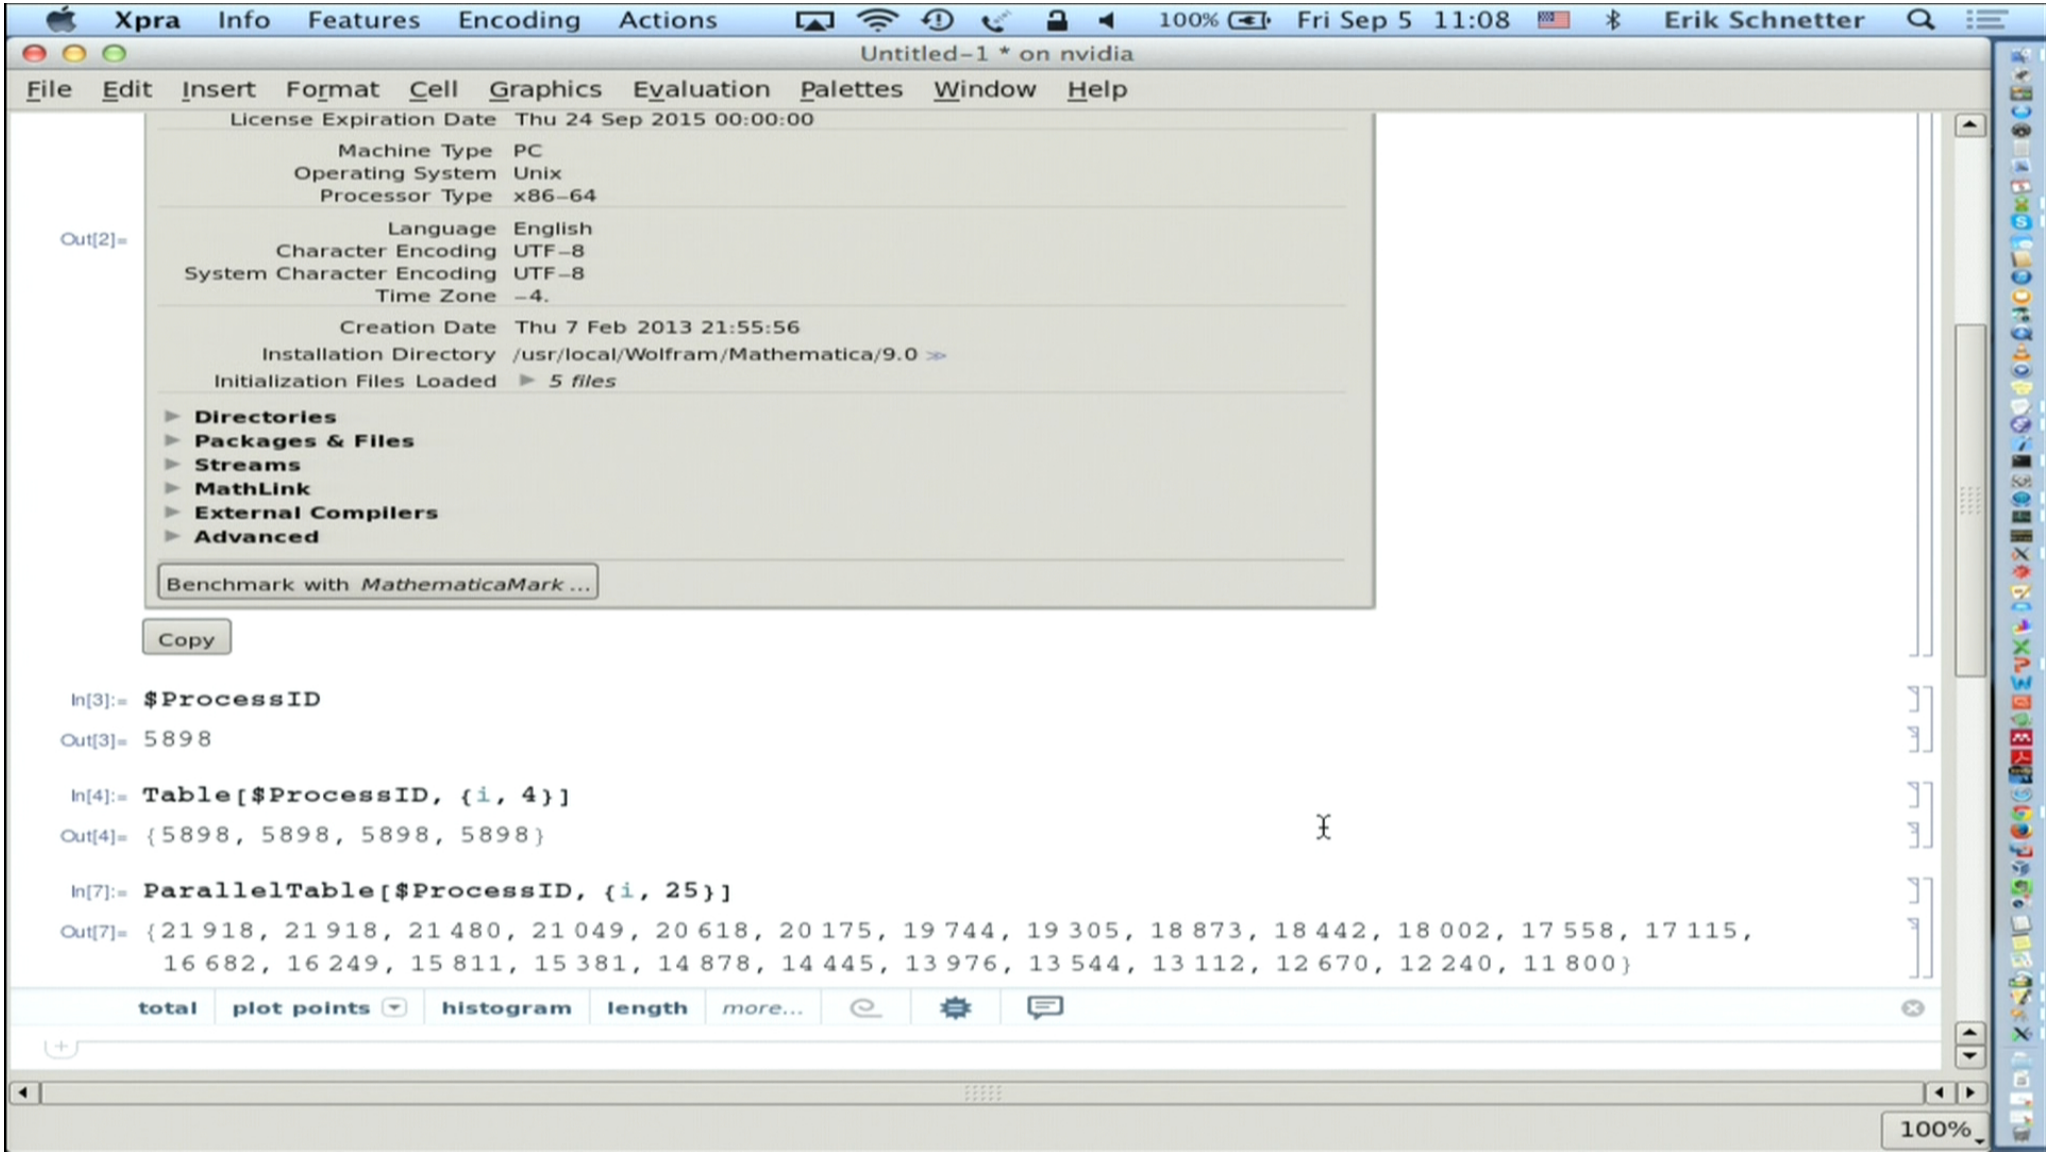
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:08 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
License Expiration Date Thu 24 Sep 2015 00:00:00
Machine Type PC
Operating System Unix
Processor Type x86-64
Out[2]=
Language English
Character Encoding UTF-8
System Character Encoding UTF-8
Time Zone -4.
Creation Date Thu 7 Feb 2013 21:55:56
Installation Directory /usr/local/Wolfram/Mathematica/9.0
Initialization Files Loaded 5 files
Directories
Packages & Files
Streams
MathLink
External Compilers
Advanced
Benchmark with MathematicaMark ...
Copy
In[3]:= $ProcessID
Out[3]= 5898
In[4]:= Table[$ProcessID, {i, 4}]
Out[4]= {5898, 5898, 5898, 5898}
In[7]:= ParallelTable[$ProcessID, {i, 25}]
Out[7]= {21 918, 21 918, 21 480, 21 049, 20 618, 20 175, 19 744, 19 305, 18 873, 18 442, 18 002, 17 558, 17 115,
16 682, 16 249, 15 811, 15 381, 14 878, 14 445, 13 976, 13 544, 13 112, 12 670, 12 240, 11 800}
total plot points histogram length more...
100%

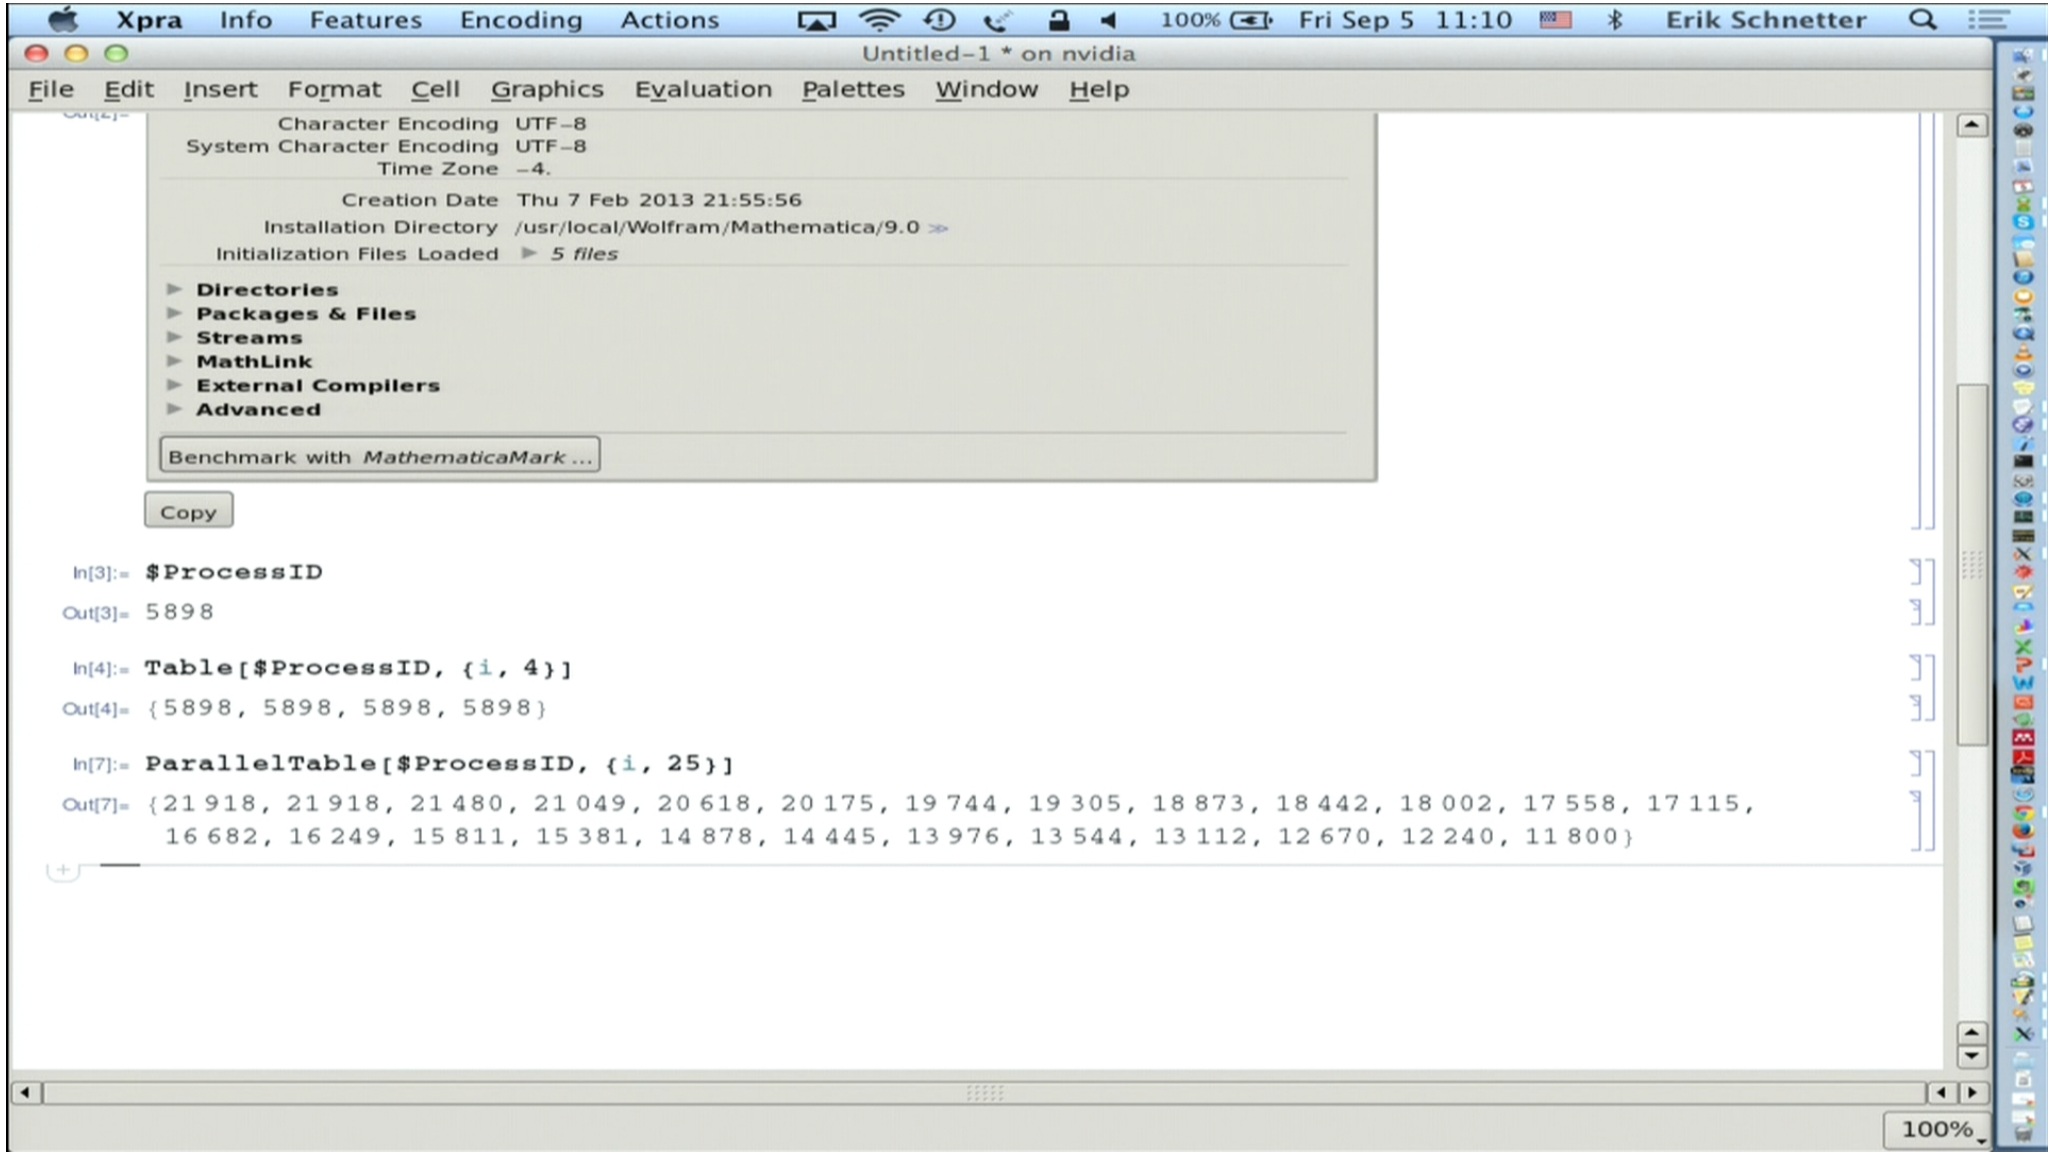
Xpra Info Features Encoding Actions 100% Fri Sep 5 11:10 Erik Schnetter
Untitled-1 * on nvidia
File Edit Insert Format Cell Graphics Evaluation Palettes Window Help
Character Encoding UTF-8
System Character Encoding UTF-8
Time Zone -4.
Creation Date Thu 7 Feb 2013 21:55:56
Installation Directory /usr/local/Wolfram/Mathematica/9.0
Initialization Files Loaded 5 files
Directories
Packages & Files
Streams
MathLink
External Compilers
Advanced
Benchmark with MathematicaMark ...
Copy
In[3]:= $ProcessID
Out[3]= 5898
In[4]:= Table[$ProcessID, {i, 4}]
Out[4]= {5898, 5898, 5898, 5898}
In[7]:= ParallelTable[$ProcessID, {i, 25}]
Out[7]= {21 918, 21 918, 21 480, 21 049, 20 618, 20 175, 19 744, 19 305, 18 873, 18 442, 18 002, 17 558, 17 115,
16 682, 16 249, 15 811, 15 381, 14 878, 14 445, 13 976, 13 544, 13 112, 12 670, 12 240, 11 800}
100%

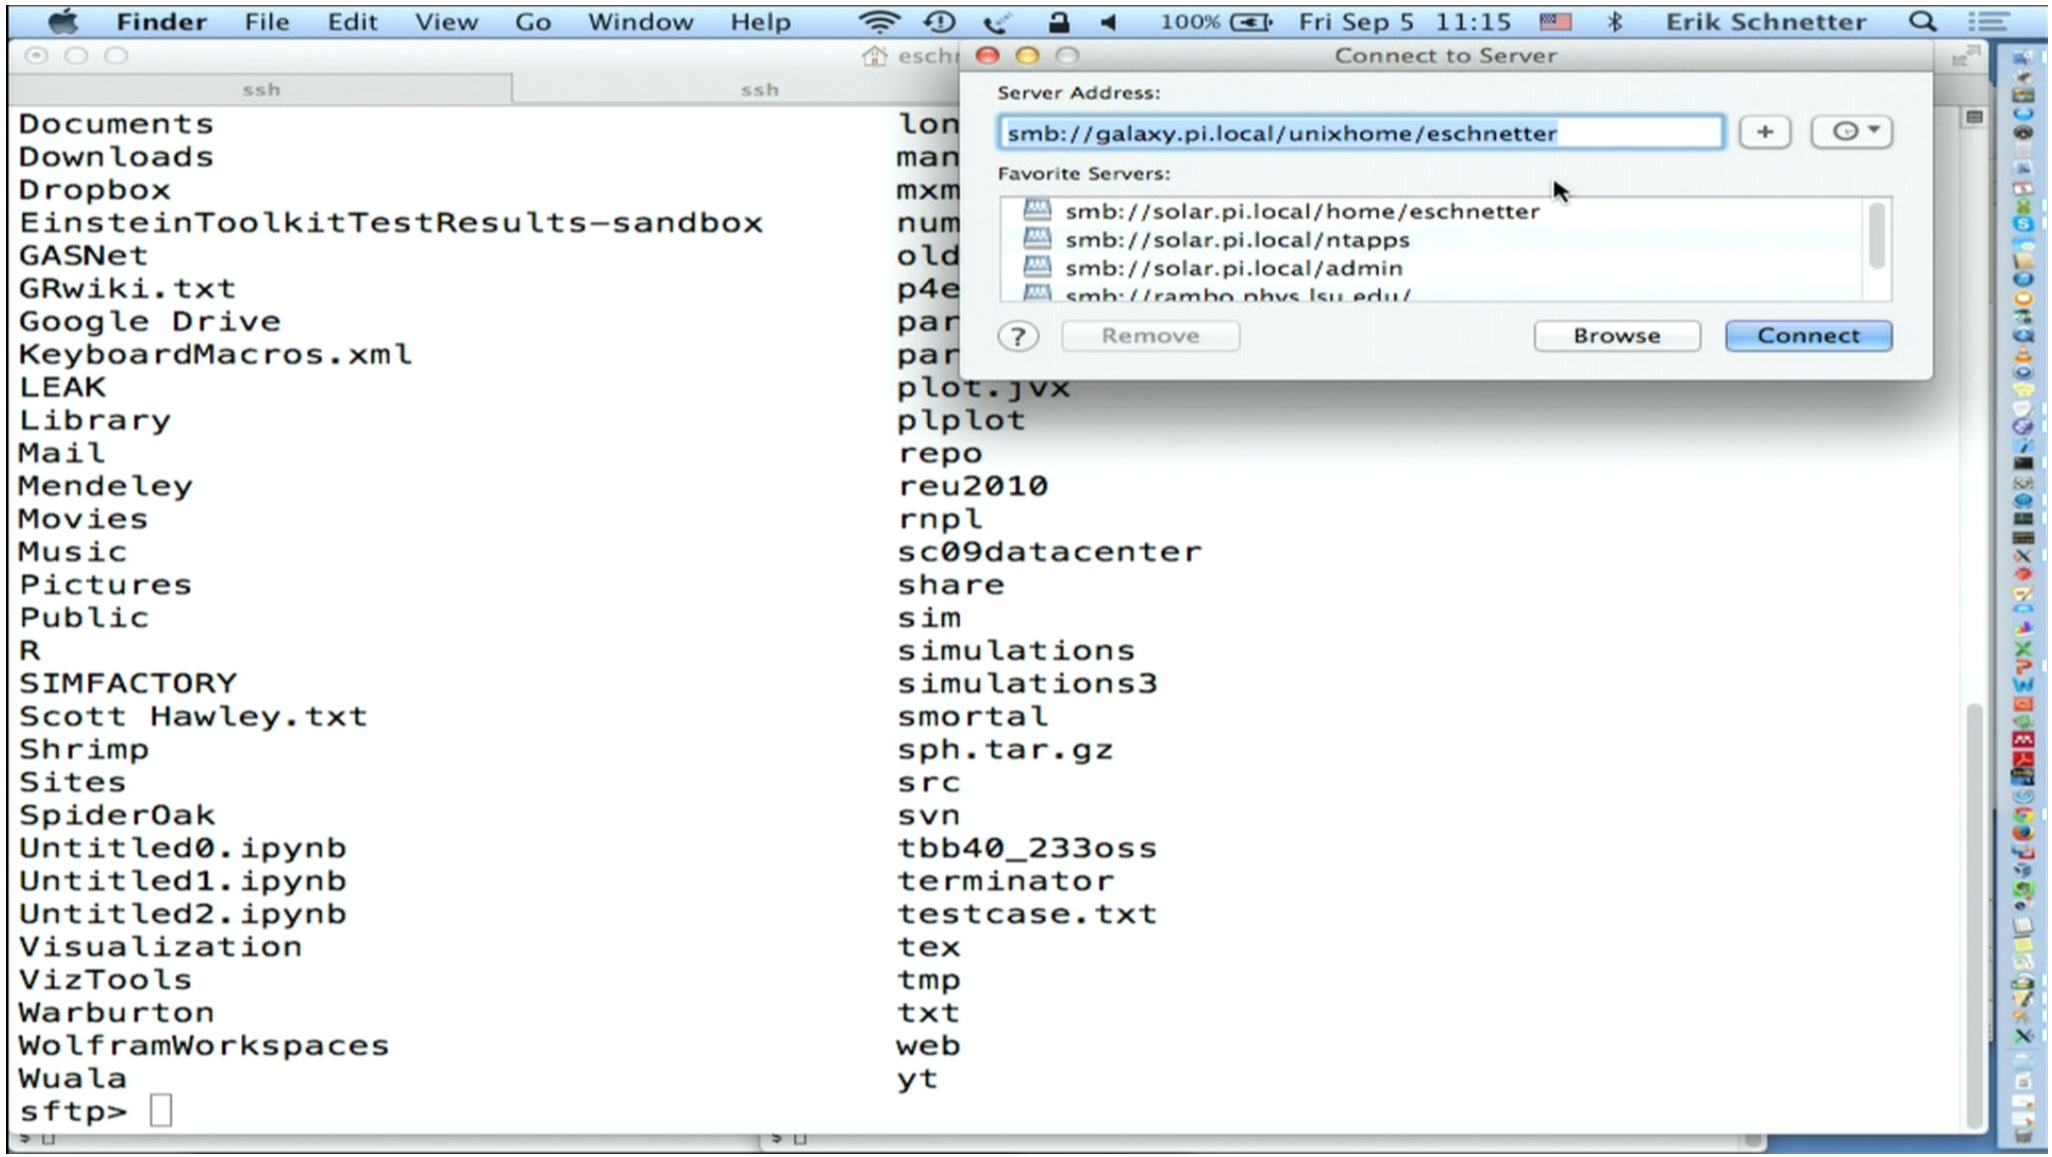
Finder File Edit View Go Window Help
100% Fri Sep 5 11:15 Erik Schnetter
ssh
ssh
Documents
Downloads
Dropbox
EinsteinToolkitTestResults-sandbox
GASNet
GRwiki.txt
Google Drive
KeyboardMacros.xml
LEAK
Library
Mail
Mendeley
Movies
Music
Pictures
Public
R
SIMFACTORY
Scott Hawley.txt
Shrimp
Sites
SpiderOak
Untitled0.ipynb
Untitled1.ipynb
Untitled2.ipynb
Visualization
VizTools
Warburton
WolframWorkspaces
Wuala
sftp>
plot.jvx
plplot
repo
reu2010
rnpl
sc09datacenter
share
sim
simulations
simulations3
smortal
sph.tar.gz
src
svn
tbb40_233oss
terminator
testcase.txt
tex
tmp
txt
web
yt
Connect to Server
Server Address:
smb://galaxy.pi.local/unixhome/eschnetter
Favorite Servers:
smb://solar.pi.local/home/eschnetter
smb://solar.pi.local/ntapps
smb://solar.pi.local/admin
Remove
Browse
Connect

Running Persistent Jobs on SCE

- Persistent job:
  - Add “nohup” when starting a job, e.g. “nohup ./mystuff &”
- Persistent terminal:
  - Use ssh to log in, then use “screen” to start a permanent session:
    - ssh yourlogin@compute.pi.local
    - screen [may need to wait a minute]
    - ./mystuff
  - Will continue to run if connection is lost. To re-connect, log in again, then
    - screen –r
- Mathematica:
  - Write .m script, run it via “math”, then see above
  - Use e.g. Xpra for a remote desktop

Introduction to Simulation Science

Erik Schnetter, Perimeter Institute
2014-09-05

Simulations

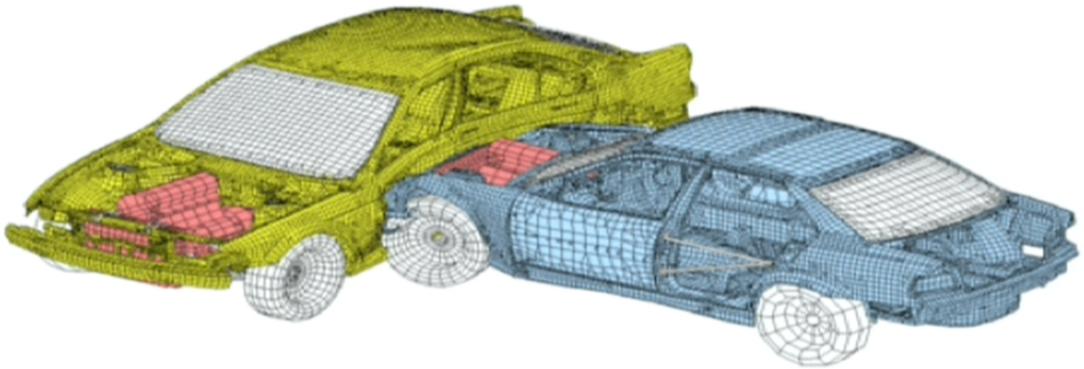

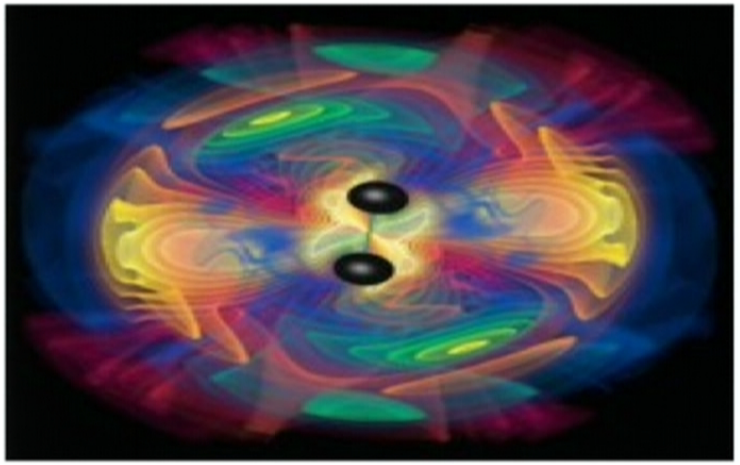

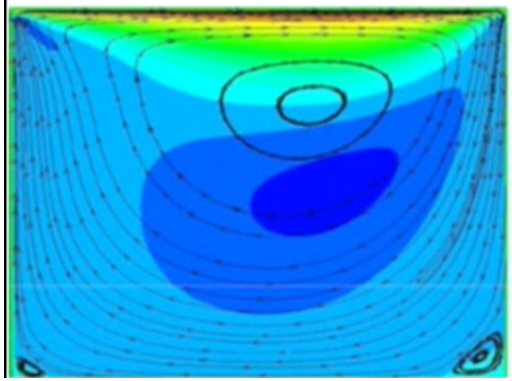

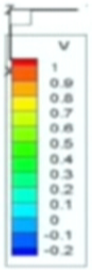

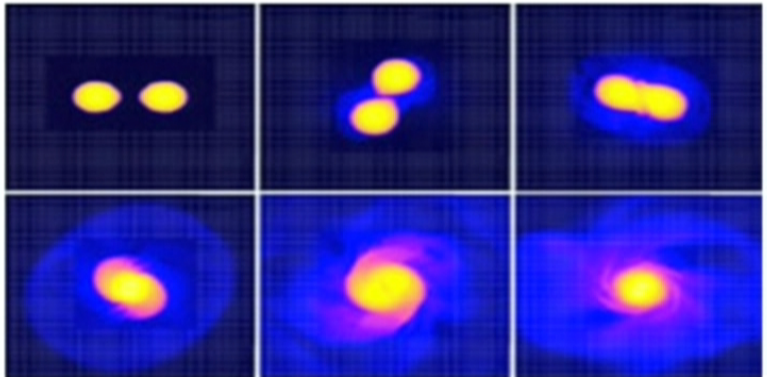

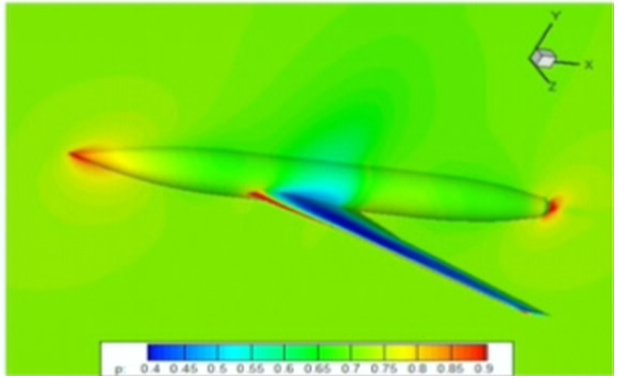

Why Use Simulations?

- Flame propagation in combustion engine: *understand* behaviour that is too fast or too small
- Hurricane modelling: *predict* behaviour
- Car crash testing: *engineer* better devices
- Video games: *create* a fantasy world similar to the real one

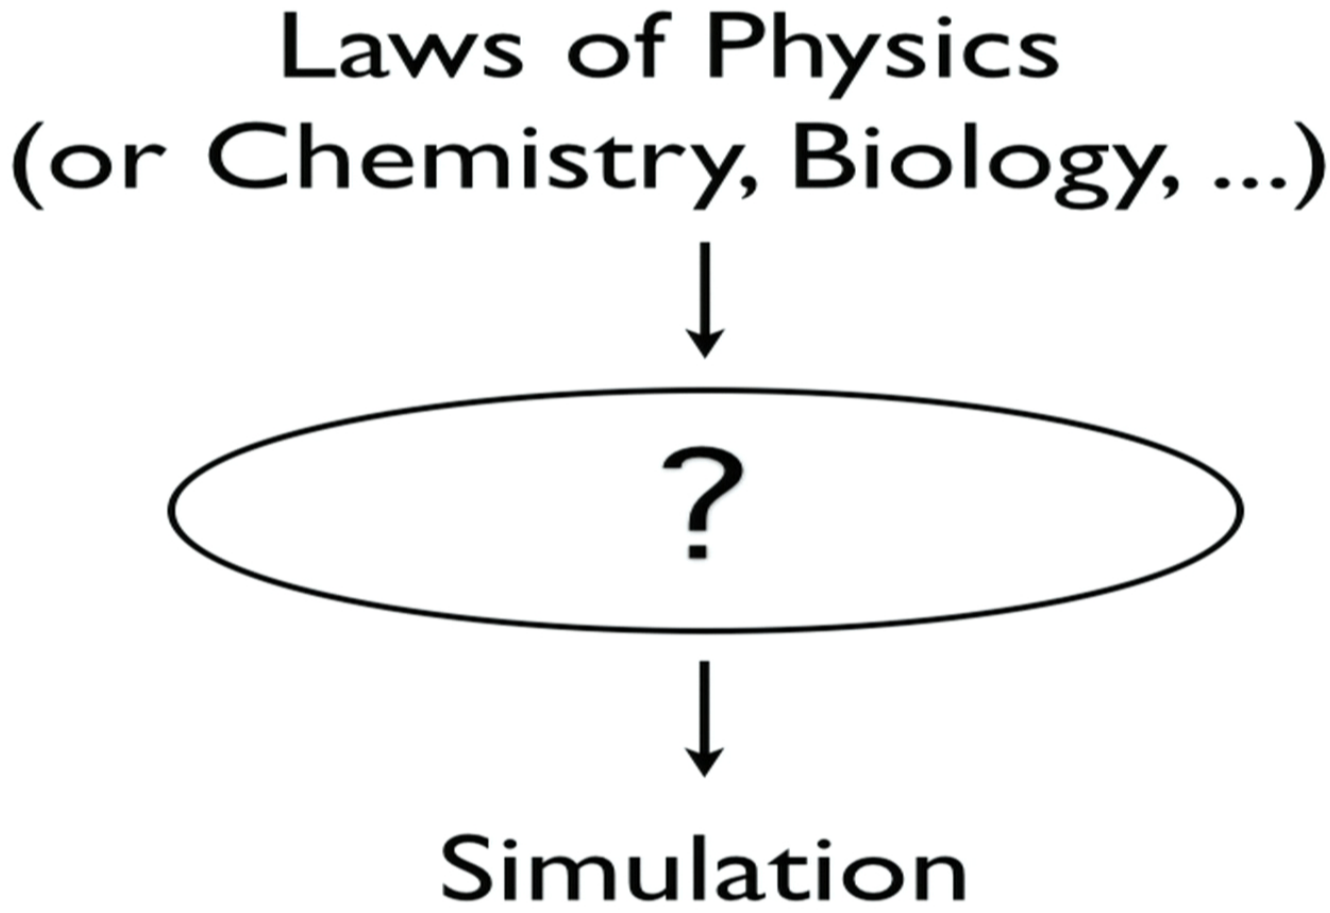
Laws of Physics
(or Chemistry, Biology, ...)
?
Simulation

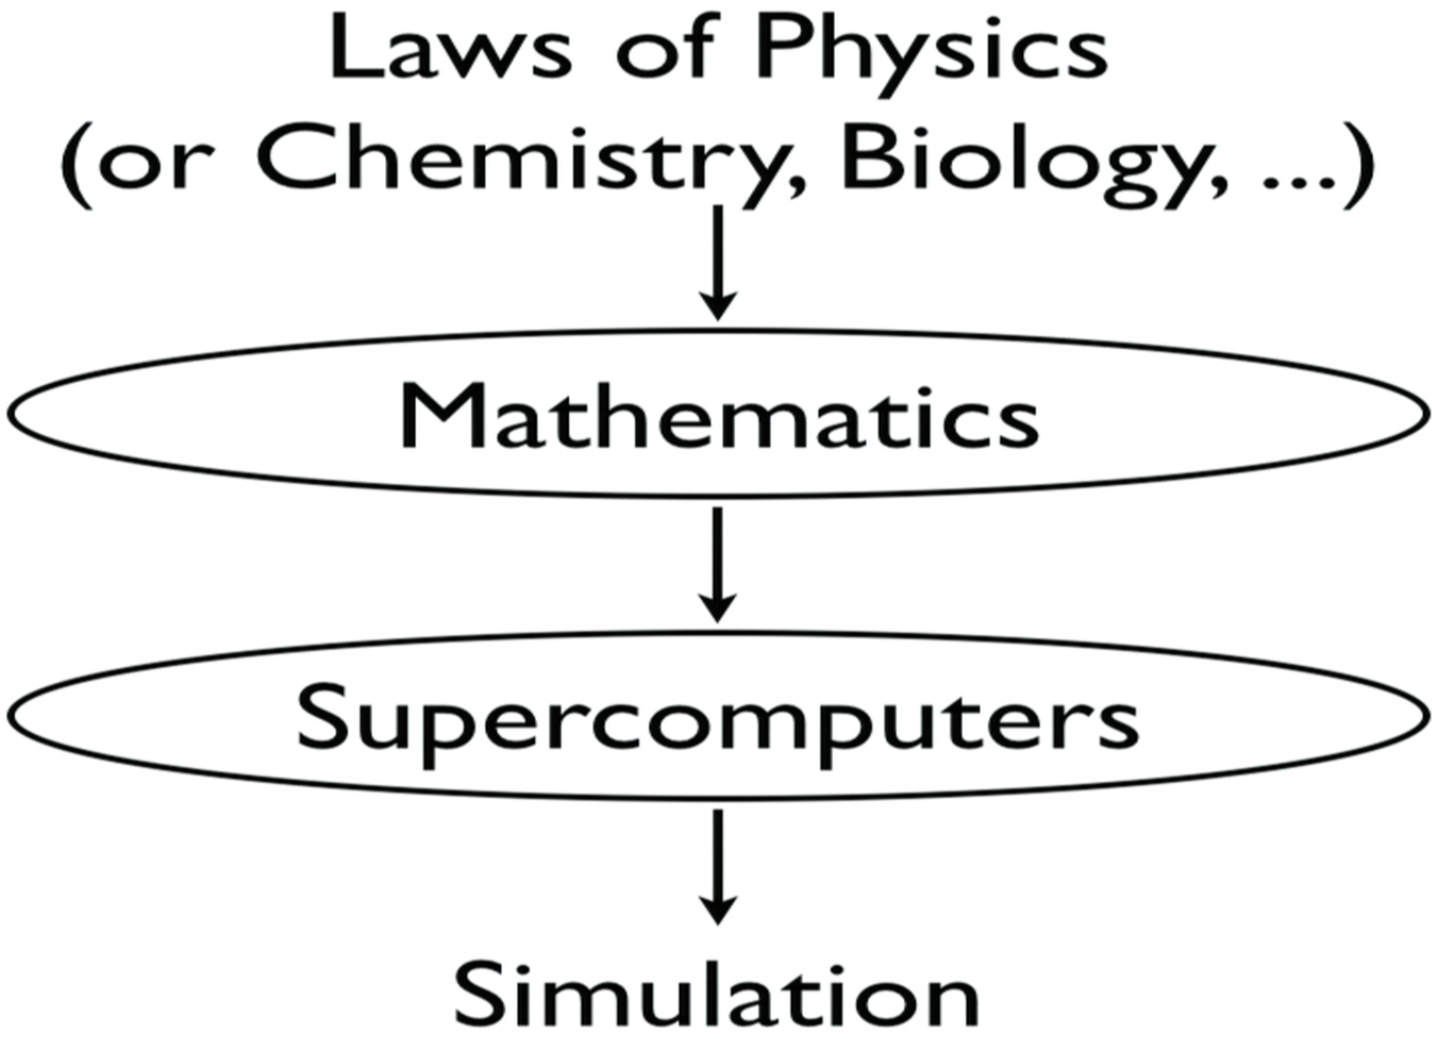
Laws of Physics
(or Chemistry, Biology, ...)
Mathematics
Supercomputers
Simulation

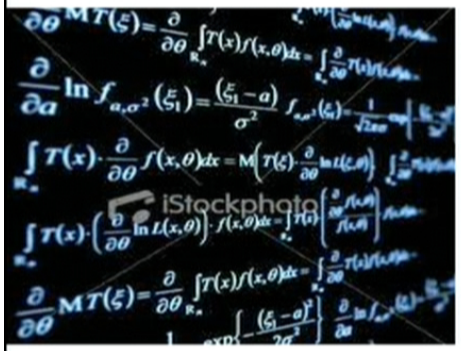

- The physics that is to be simulated is expressed in "the language of Mathematics"
  - Called *Scientific Computing* or *Numerical Analysis*

- The resulting systems of equations are solved on large computers
  - Called *Supercomputers* because they are as large and awkward as a supertanker

Systems and Equations

- The state of a system is described via variables (density, velocity, pressure, etc.)
- Laws of Physics can then often be described via *PDEs* (Partial Differential Equations)
- A PDE describes how a system is *changing* depending on its current *state*

Discretisation

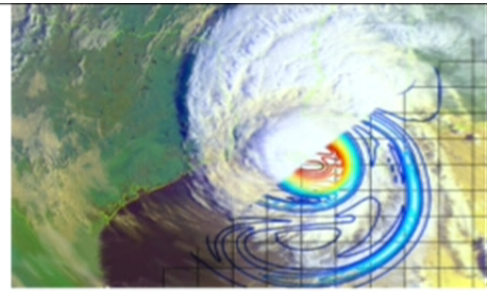

- PDEs describe continuum systems (car body, water, air); these have infinitely many degrees of freedom
- Reduce complexity by approximation via a *discrete system* instead
- Compare e.g. pixels on a TV screen, surface triangulation for visualisation

- Many possibilities:
- finite elements (e.g. small rigid triangles)
- finite volumes (e.g. small cubes)
- finite differences (sample solution on regular grid)
- particles (small chunks of matter)
- many more...

Simulation Procedure

- Choose PDE that describes system well
- Discretise PDE
- Set up initial condition

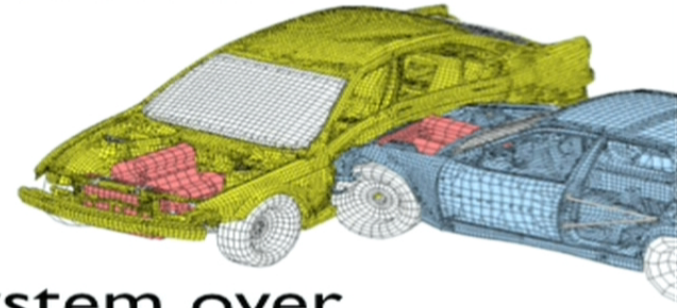

- Follow each element of the system over many many tiny steps
- A simulation can have billions of elements with millions of steps, taking weeks of computing time

Caveat

- Some systems are described not by PDEs but otherwise (e.g. coupled ODEs, discrete transitions)
- Sometimes not time evolution is interesting, but e.g. equilibrium configuration
- Usually (in real life), PDEs and initial conditions are only *approximations or guesses*, and simulation results *may not be reliable*

Supercomputer Architectures

Remote Access

- Supercomputers are located in far away places, need to use ssh/gsissh to access
- Log in is to *front end (head node)* only, usually a large workstation
- Cannot (or should not) use front end to run simulations

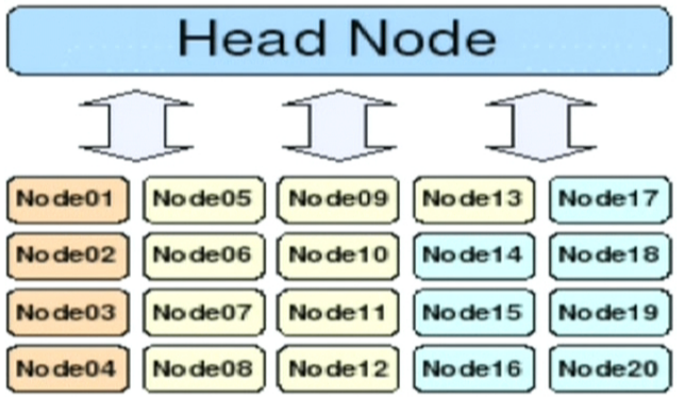

Compute Nodes, Interconnect

- Most supercomputers have a *cluster* architecture with many compute nodes
- Each node has (4 to 32?) cores, similar to a large workstation
- Nodes are connected via a low-latency *communication network* (e.g. Infiniband)
- Overall system has (128 to >8,000?) nodes, or up to 100k cores
- My personal scale:
  <1k cores: small,
  <10k cores: medium
  >10k cores: large

Batch System

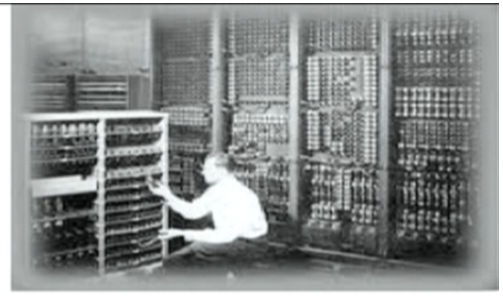

- Cannot (or should not) use compute nodes directly
- Need to submit *job* to *batch system*, requesting *N* nodes...
- ... wait (a few days?) ...
- ... then the job runs
- (... and then one discovers one's errors)
- There is a run time limit, often 24h or 48h
- ... which is inconvenient if one needs to run for 2 weeks: checkpoint/ restart
- Batch systems ensure that a supercomputer is not idle; there are always jobs waiting to be executed

Allocations

- Need to ensure fair use of supercomputer, prevent individual users from monopolising it
- Typically, an *allocation process* decides who can use how much of a supercomputer's time during a year (similar to writing a grant proposal)
- 1 CPU hour costs about 5 cents
- With this metric, Titan (small HPC system at Perimeter) produces about $150k worth of CPU time per year

Software

- Installed/available software is system dependent, not just standard Unix systems
- Therefore cannot just install binaries, need to build software manually (or ask administrators to do that)
- HPC developers often prefer command line tools, don't use GUIs (which may not be available)

Software

- Installed/available software is system dependent, not just standard Unix systems
- Therefore cannot just install binaries, need to build software manually (or ask administrators to do that)
- HPC developers often prefer command line tools, don't use GUIs (which may not be available)

HPC History

- Before MPI: *Vector architectures*, e.g. Cray Y-MP (until ~1992)
- Each instruction acts on 100s or 1000s of data elements "simultaneously" (SIMD)

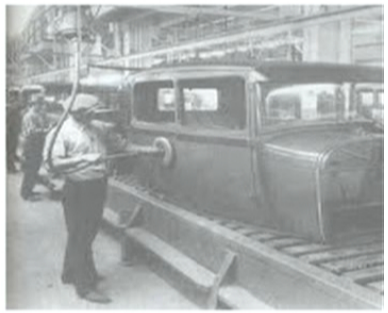

- Much more efficient than scalar processor (compare conveyor belt vs. hand assembly)
- Disadvantages: too inflexible for dynamic data structures, too expensive due to custom-designed (low-volume) hardware

HPC History, Cont'd

- After vector machines, *cluster architecture* became prevalent (Cray T3D, 1993)
- Basic idea: have many simple nodes, connected by high-speed network
- Nodes need to communicate (*exchange messages*) during computation
- Also called *Beowulf architecture*, esp. if only a cheap network is used

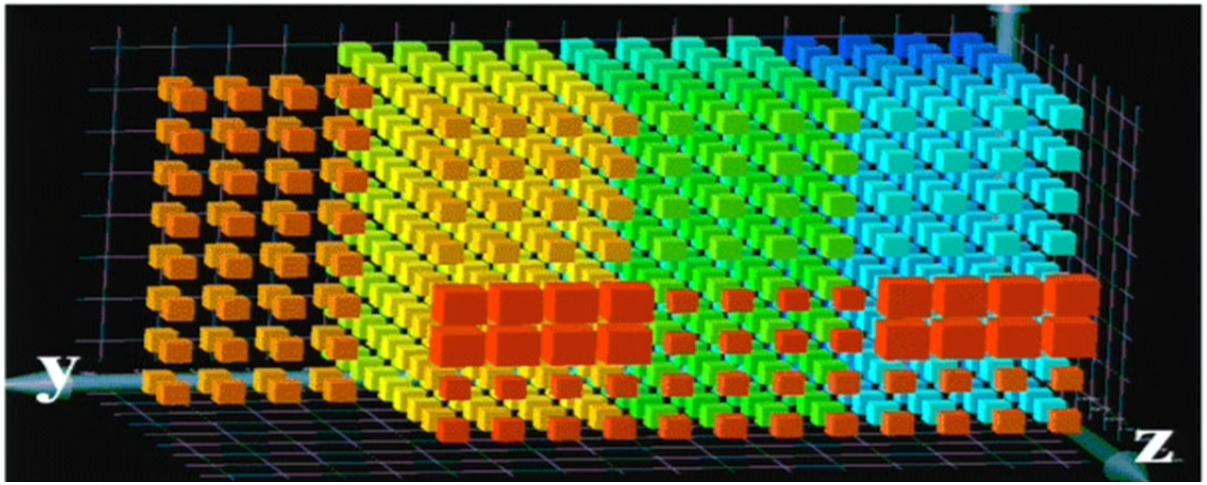

Node connectivity in a Cray T3E

MPI Programming

- Each process runs an independent copy of the program
- Each copy has a unique number assigned to it(0...*N*-1)
- The program needs to divide the total workload into N pieces, and assign one to each process

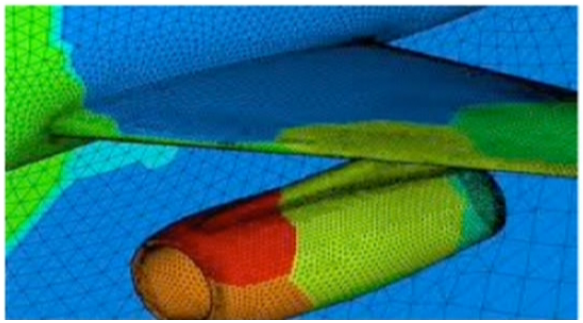

- The processes can talk to each other *only* by exchanging messages
- MPI hides low-level, system-dependent communication details from the programmer
- MPI messages are (by default) *ordered* and *reliable*

Examples of Real Life Message Passing

- Message Passing is very common, even in the real world; for example:
  - Letters via the post office
  - Email
  - Phone text messages
  - Newspapers
  - Chinese whispers game
  - Monopoly
- However, these are NOT message passing – they are *streams* or *interactive* instead:
  - Phone conversation
  - Watching TV
  - Google Docs
  - Charades game
  - WoW

Distributed Array Implementation

- set-element(n,x):

1. determine which process owns element n
2. send message to that process containing index n and value x
3. check whether message has been received
4. if so, set element n to x

- x = get-element(n):

1. determine which process owns element n
2. send message to that process with index n
3. check whether message has been received
4. if so, get element n ...
5. ... and send message back with value x
6. wait for result

In reality, would not send individual elements, but would send *batches* of elements to reduce latency

Amdahl's Law

- When running on *N* processes, not necessarily *N* times as fast – overhead
- Overhead determines maximum possible parallel speedup

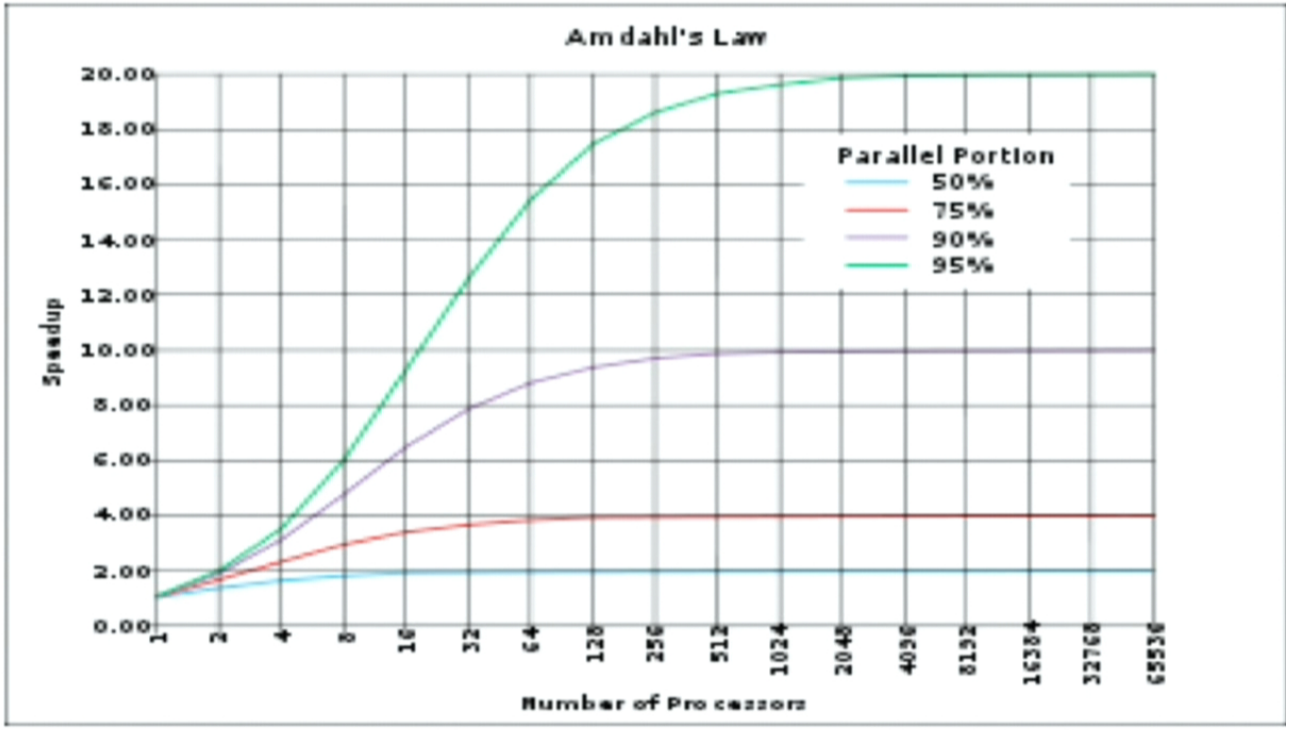

100,000-fold speedup requires >99.999% parallelisation

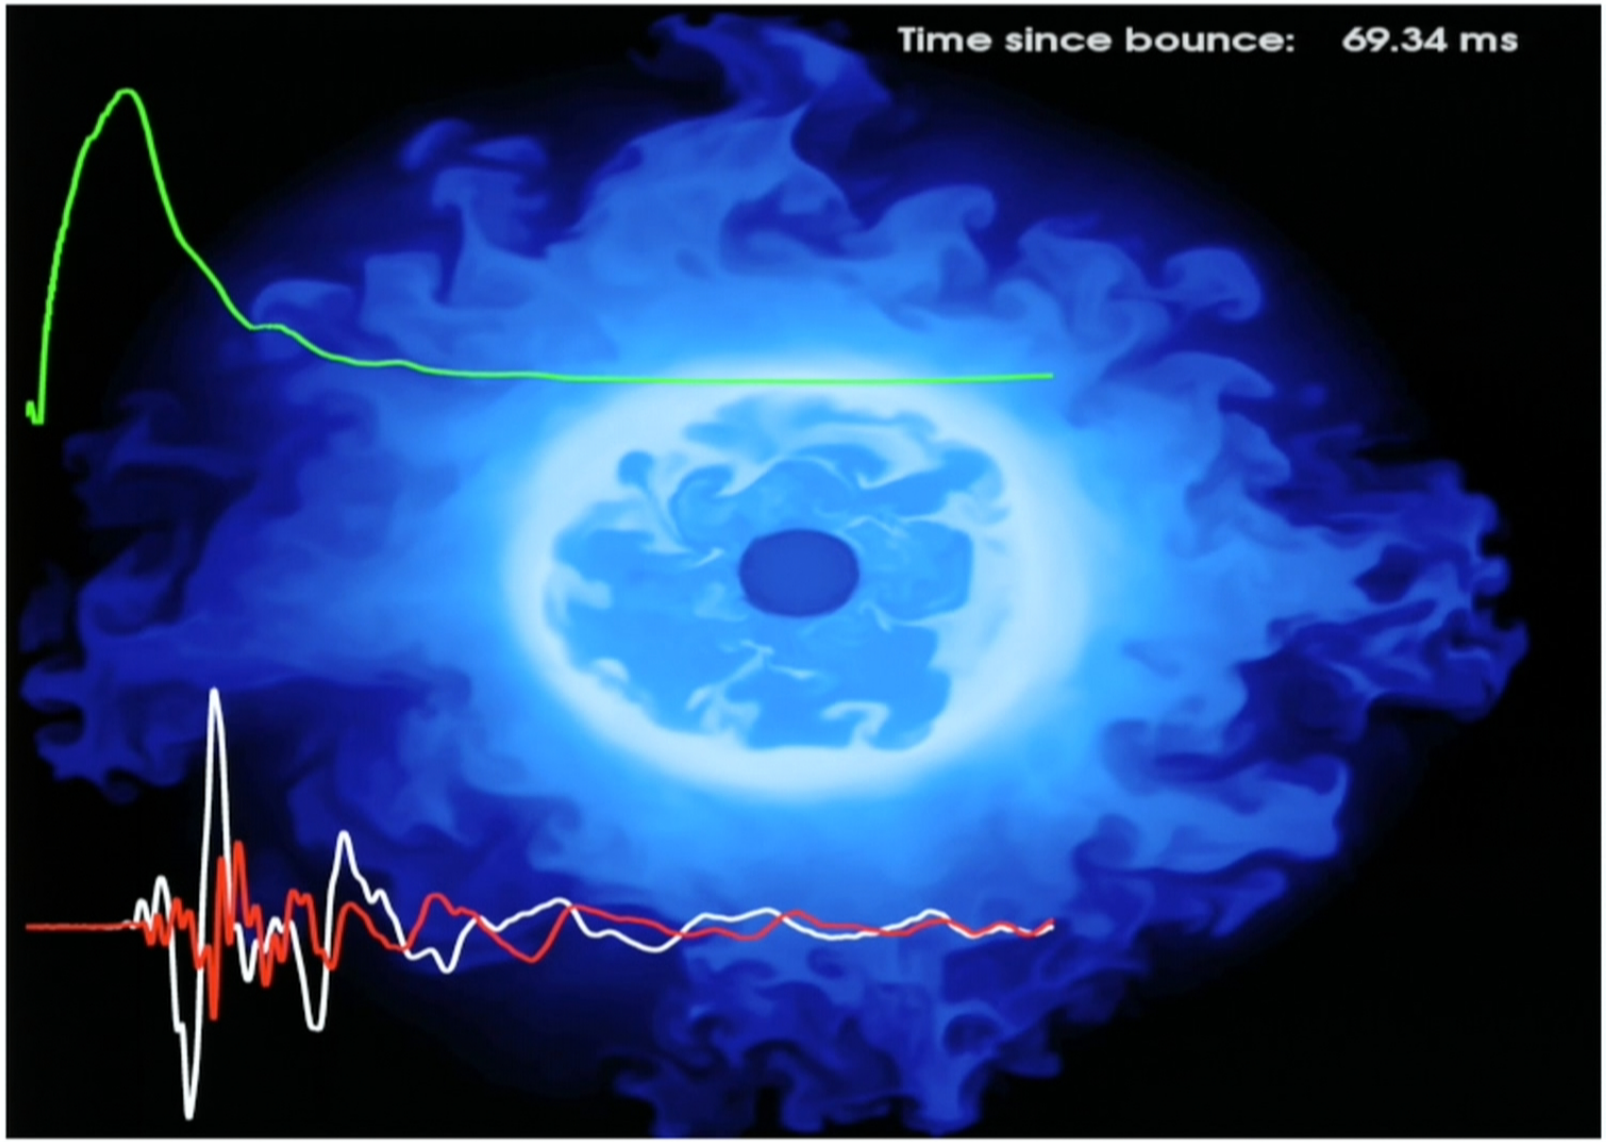
Time since bounce: 69.34 ms

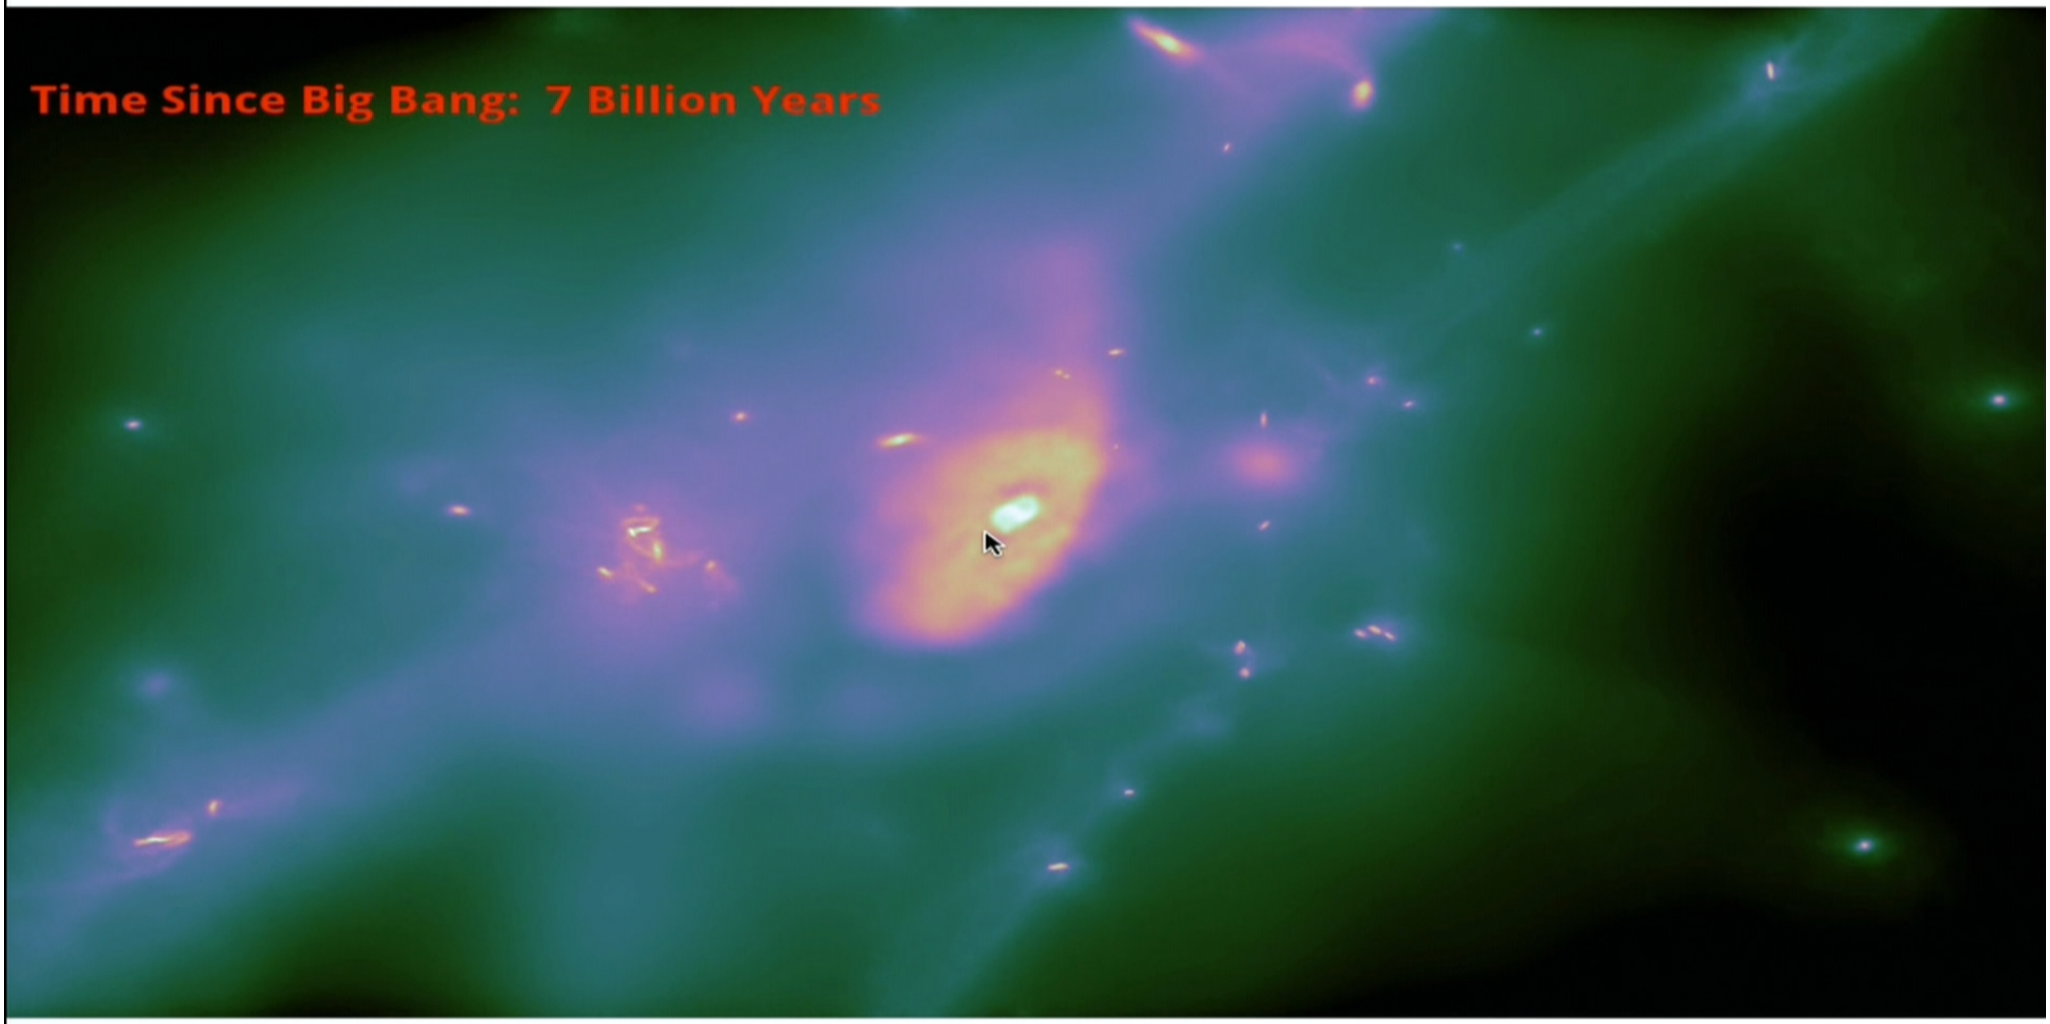
Time Since Big Bang: 7 Billion Years