

The Future of IPython: Interactive Parallel and Distributed Computing



Brian E. Granger
Tech-X Corporation

CU BOULDER Fernando Perez
University of Colorado

Benjamin Ragan-Kelly (student)
Santa Clara University

SciPy'06
Caltech

IPython: An Enhanced Interactive Python Shell

- IPython is an enhanced interactive Python shell
- It is becoming the de facto shell for scientific computing in Python
- Included with most major Linux distributions
- Capabilities:
 - User extensible syntax <http://ipython.scipy.org>
 - GUI integration (wx, Qt, GTK, etc.)
 - Seamless system shell access
 - Dynamic object/namespace introspection: docstrings, attributes, methods, source code
 - Numbered input/output prompts with command history
 - Session logging and restoring
 - Embeddable

Goals

Create an open source architecture that enables parallel and distributed programs to be developed, monitored, executed and debugged interactively and collaboratively.

- All of IPython's capabilities will be available over the wire.
- Easy things should be easy.
- It should integrate well with existing C/C++/Fortran/MPI code.
- It should support many different styles of parallelism: message passing, task farming, distributed memory.
- Fully interactive work with up to 256 processors (latency < 0.1 sec).
- Wide range of hardware: laptops to NERSC supercomputers

Realities of Distributed Scientific Computing

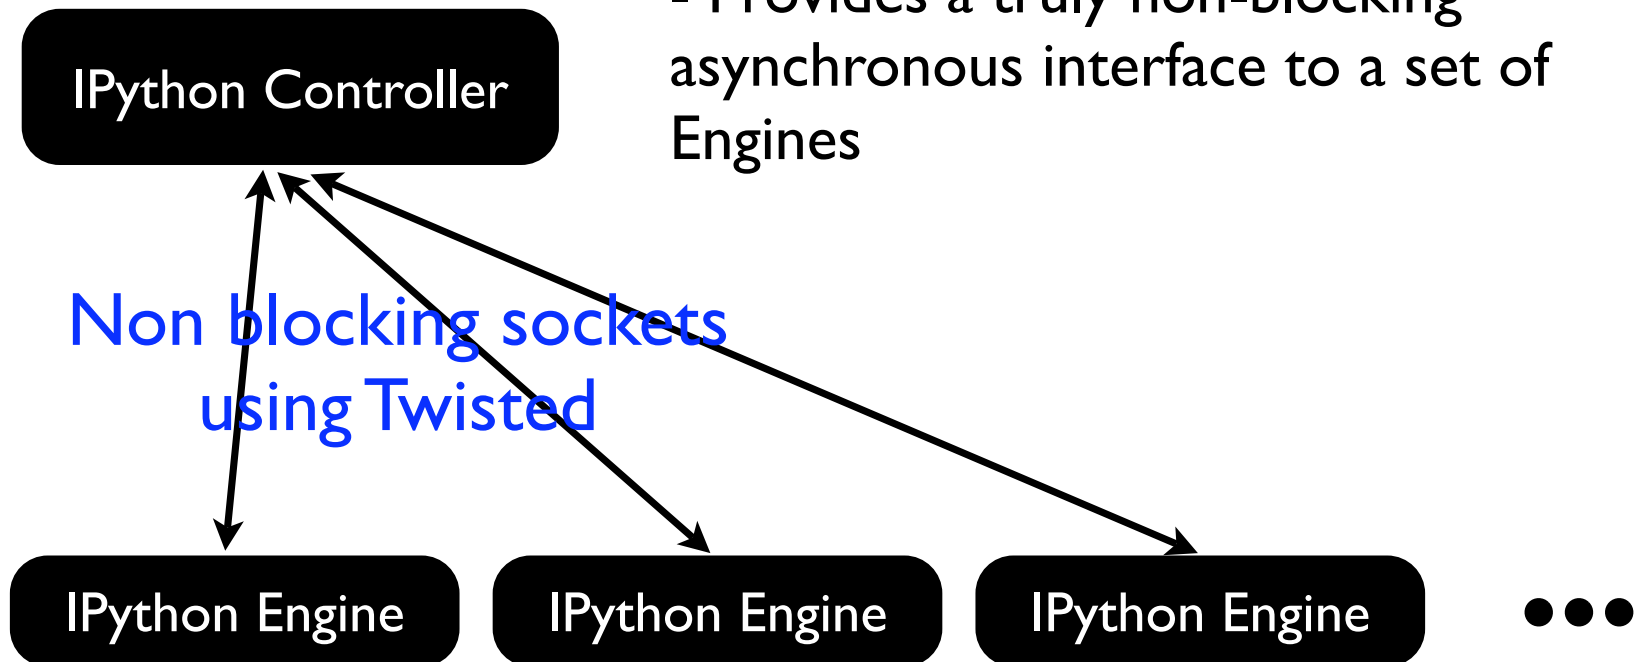
- Scientific code is often written in compiled languages (C/C++/Fortran).
- This code takes a long time to execute.
- While this code is running, nothing else (such as network communications) can happen in process.
- This is completely orthogonal to the asynchronous nature of distributed, network based computing.
- Because of the Global Interpreter Lock (GIL) threads don't provide a general solution to this problem.

Minimal Requirements for Good Parallelism in Python

- Multiple processes.
- Non-blocking sockets.
- Asynchronous error/fault handling.

Architecture

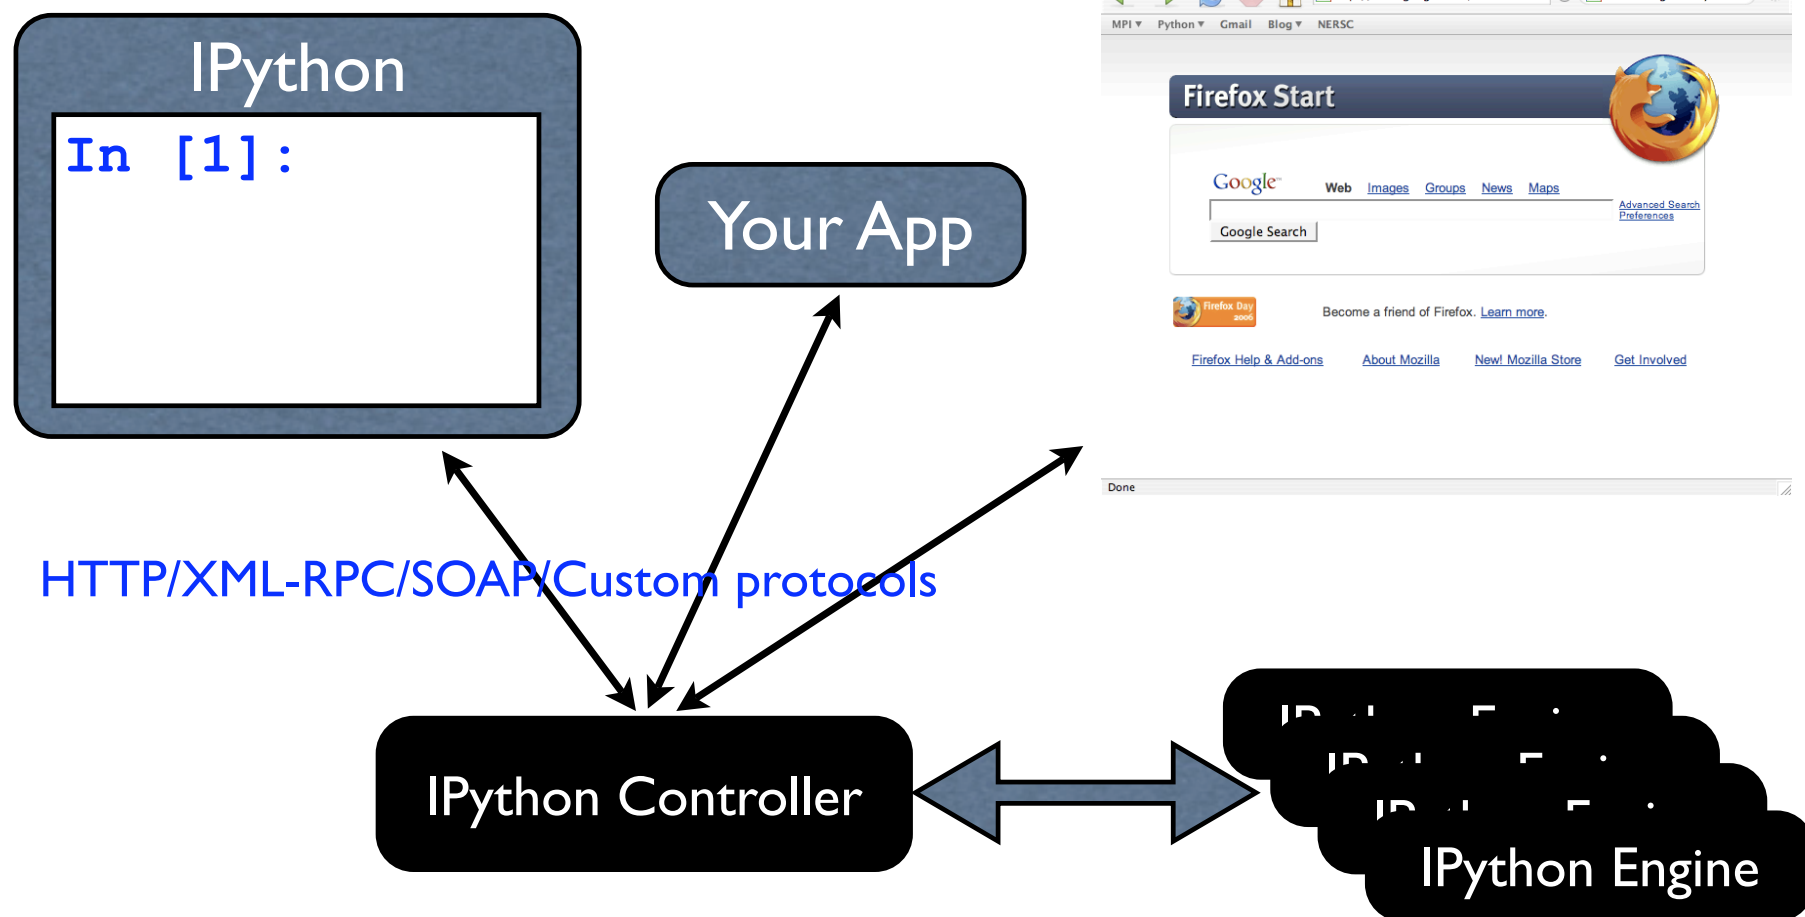
- Manages the Engines
- Queue for each Engine
- Provides a truly non-blocking asynchronous interface to a set of Engines



(I)Python virtual machines exposed to the network. These can and will block!

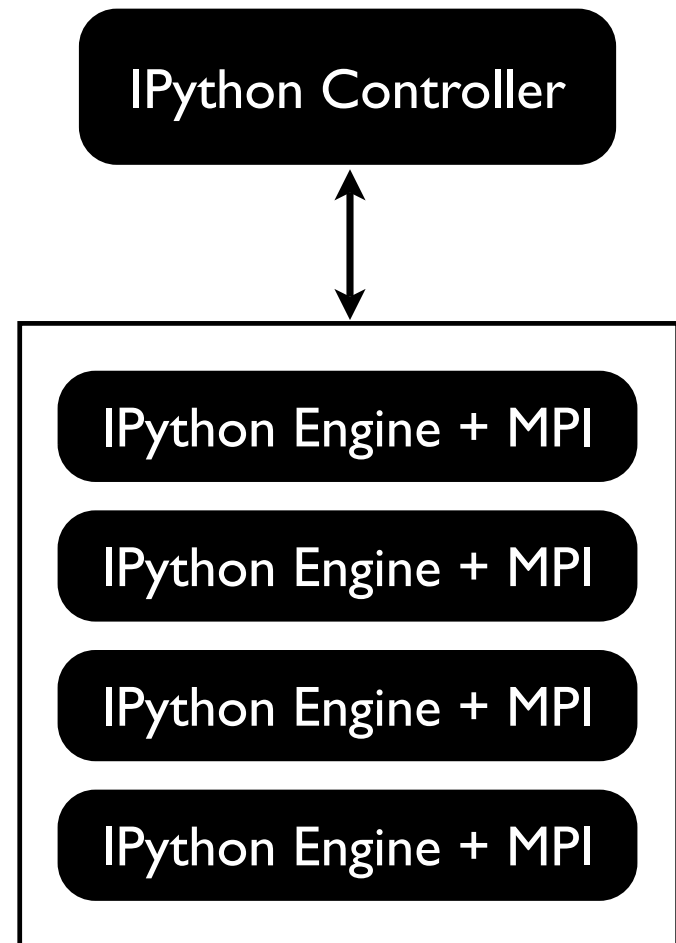
Architecture

Multiple/simultaneous users/apps can connect to the Controller = collaboration

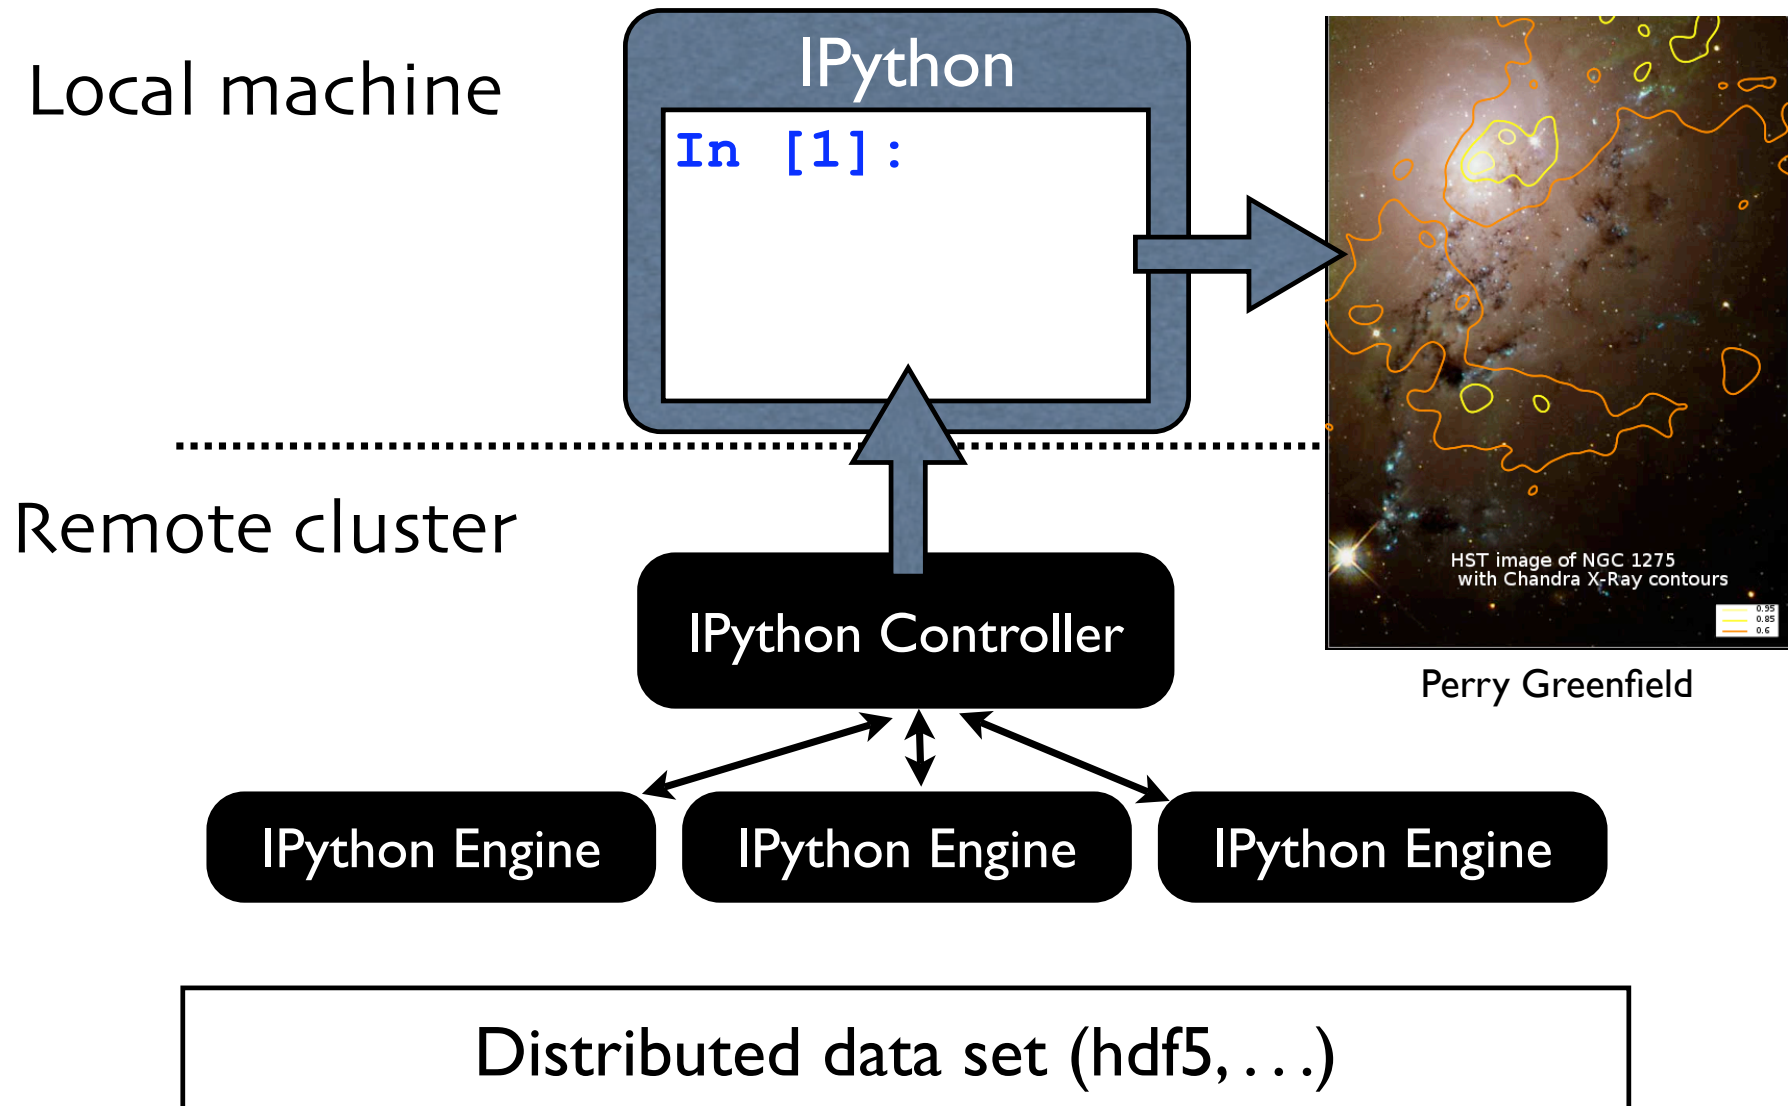


Interactive Steering of Traditional Parallel Codes

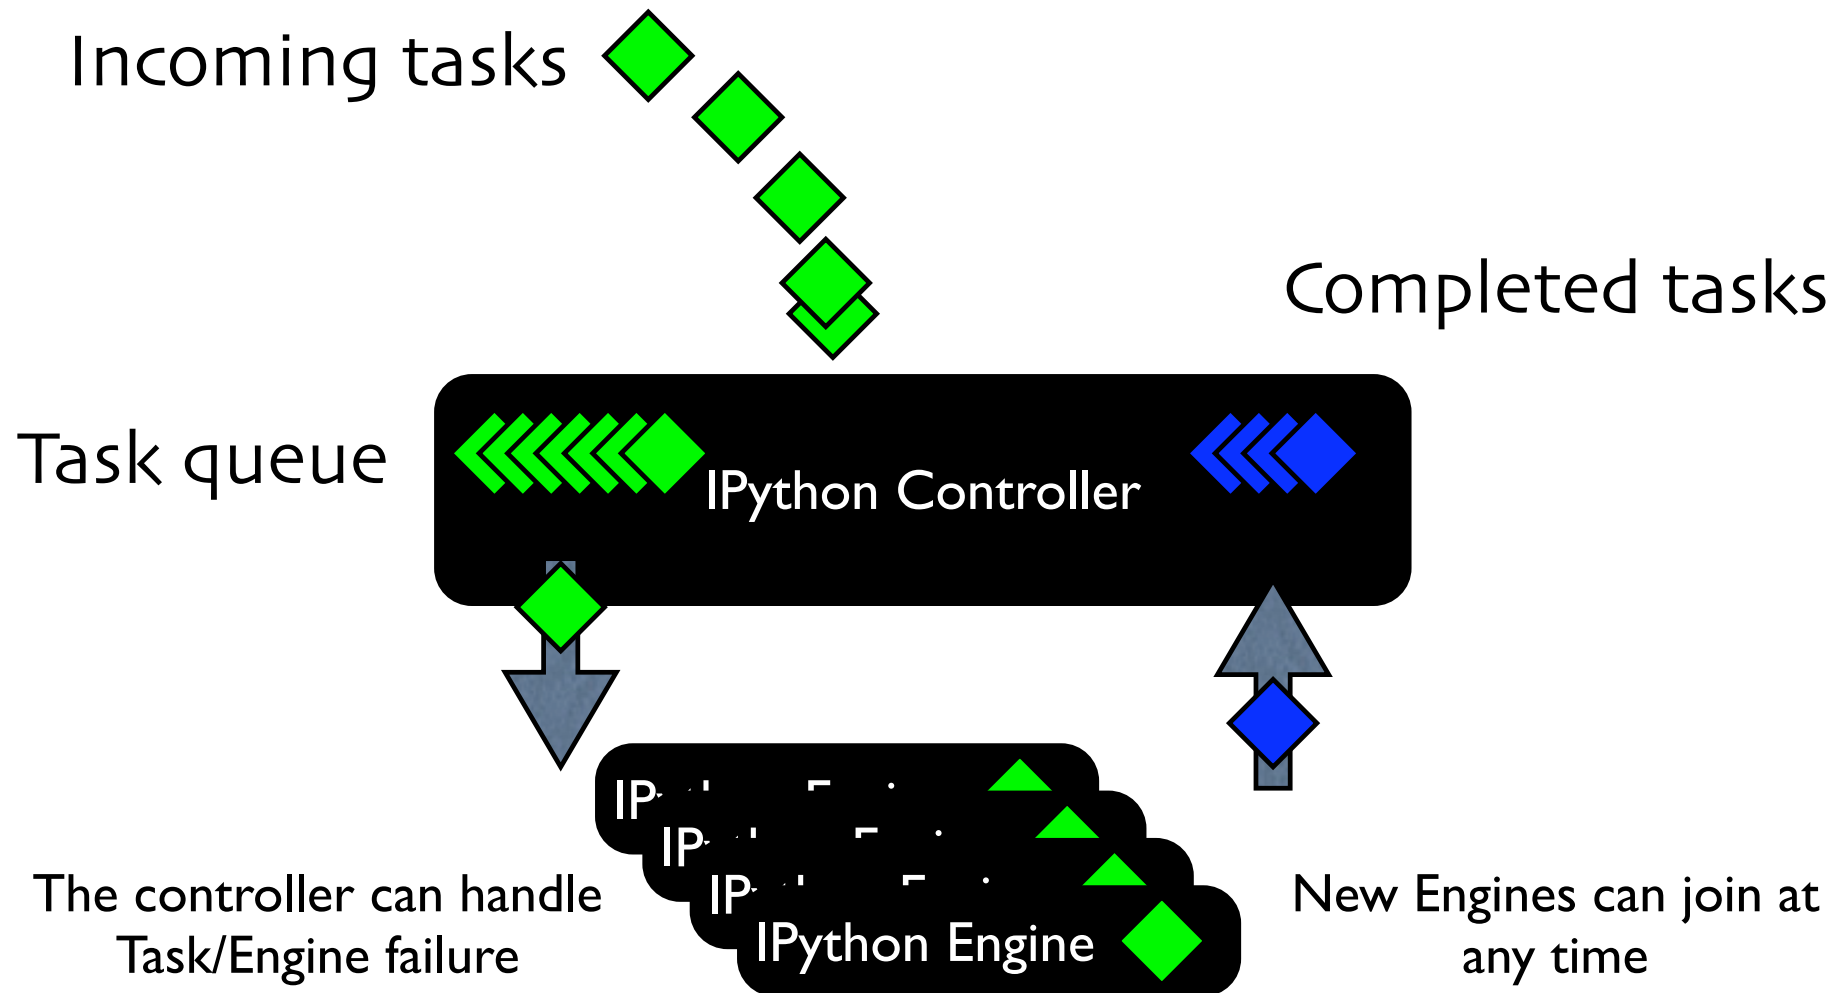
- The IPython Engine can optionally initialize MPI at startup.
- User code can call wrapped C/C++/Fortran code that makes calls to MPI.
- We also support many of the Python/MPI bindings.
- The Controller and Frontend/Clients don't use MPI.
- Full interactive control of a traditional parallel app.



Interactive and Parallel Data Analysis of Massive Distributed Data Sets



Load Balanced, Fault Tolerant Task Farming



Status

- Most of the infrastructure for interactive parallel and distributed is done.
- Working on testing/optimization/documentation.
- Beginning to implement the “IPythonic Features.”
- Watch ipython.scipy.org for updated info.
- Or check it out:

```
svn co http://ipython.scipy.org/svn/ipython/ipython/branches/chainsaw ipython1
```

<http://ipython.scipy.org/moin/Design>