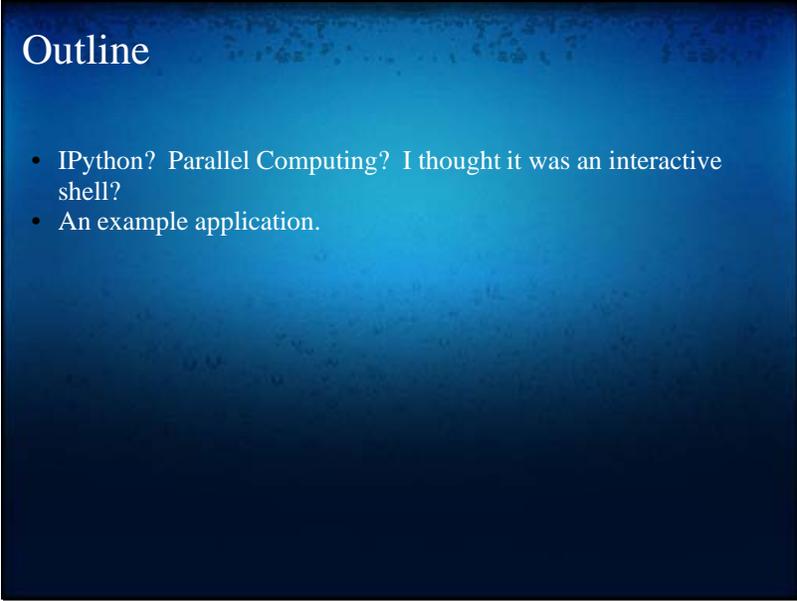


Slide 1

Parallel computing with IPython: an application to air pollution modeling

Josh Hemann, Rogue Wave Software, University of Colorado

Brian Granger, IPython Project, Cal Poly, San Luis Obispo

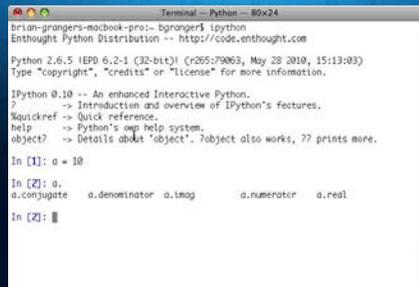


Outline

- IPython? Parallel Computing? I thought it was an interactive shell?
- An example application.

IPython Overview

- Goal: provide an efficient environment for exploratory and interactive scientific computing.
- Cross platform and open source (BSD).
- Two main components:
 - An enhanced interactive Python shell.
 - A framework for interactive parallel computing.



```
Terminal -- Python -- 80x24
brian-grangers-macbook-pro:~ bgranger$ ipython
Enthought Python Distribution -- http://code.enthought.com

Python 2.6.5 |EPD 6.2-1 (32-bit)| (r265:79063, May 28 2010, 15:13:03)
Type "copyright", "credits" or "license()" for more information.

IPython 0.10 -- An enhanced Interactive Python.
?                -> Introduction and overview of IPython's features.
?quickref       -> Quick reference.
help            -> Python's own help system.
object?        -> Details about 'object'. ?object also works, ?? prints more.

In [1]: a = 10

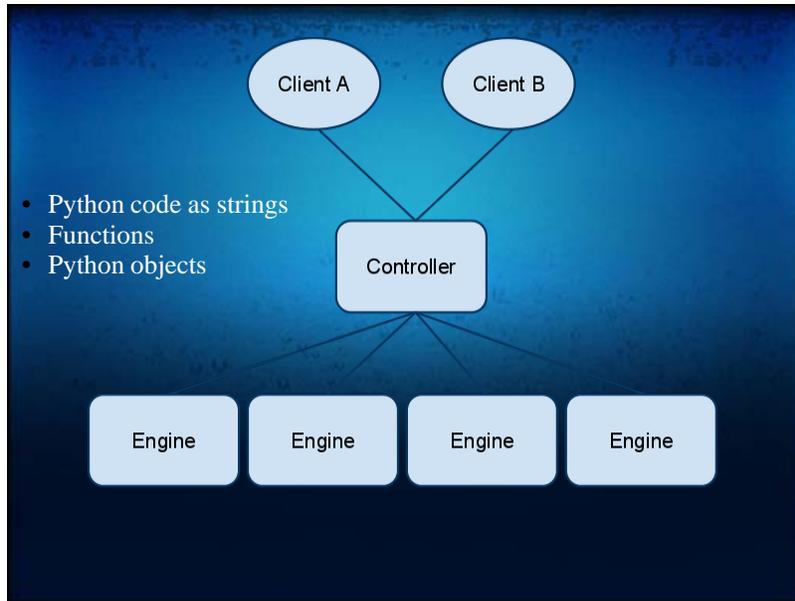
In [2]: a
a.conjugate      a.denominator  a.imag          a.numerator      a.real

In [2]:
```

IPython's Parallel Framework

- Goal: provide a high level interface for executing Python code in parallel on everything: multicore CPUs, clusters, supercomputers and the cloud.
- Easy things should be easy, difficult things possible.
- Make parallel computing collaborative, interactive.
- A dynamic process model for fault tolerance and load balancing.
- Want to keep the benefit of traditional approaches:
 - Integrate with threads/MPI if desired.
 - Integrate with compiled, parallel C/C++/Fortran codes.
- Support different types of parallelism.
- Based on processes not threads (the GIL).
- Why parallel computing in IPython?
 - R(EEEEE...)PL is the same as REPL if abstracted properly.

Slide 5



Architecture details

- The IPython Engine is a Python interpreter that executes code received over a network.
- The Controller maintains a registry of the engines and a queue for code to be run on each engine. Handles load balancing.
- Dynamic and fault tolerant: Engines can come and go at any time.
- The Client is used in top-level code to submit tasks to the controller/engines.
- Client, Controller and Engines are fully asynchronous.
- Remote exception handling: exceptions on the engines are serialized and returned to the client.
- Everything is interactive, even on a supercomputer or the cloud.

MultiEngineClient and TaskClient

- MultiEngineClient
 - Provides direct, explicit access to each Engine.
 - Each Engine has an id.
 - Full integration with MPI (MPI rank == id).
 - No load balancing.
- TaskClient
 - No information about number of Engines or their identities.
 - Dynamic load balanced queue.
 - No MPI integration.
- Extensible
 - Possible to add new interfaces (Map/Reduce).
 - Not easy, but we hope to fix that.

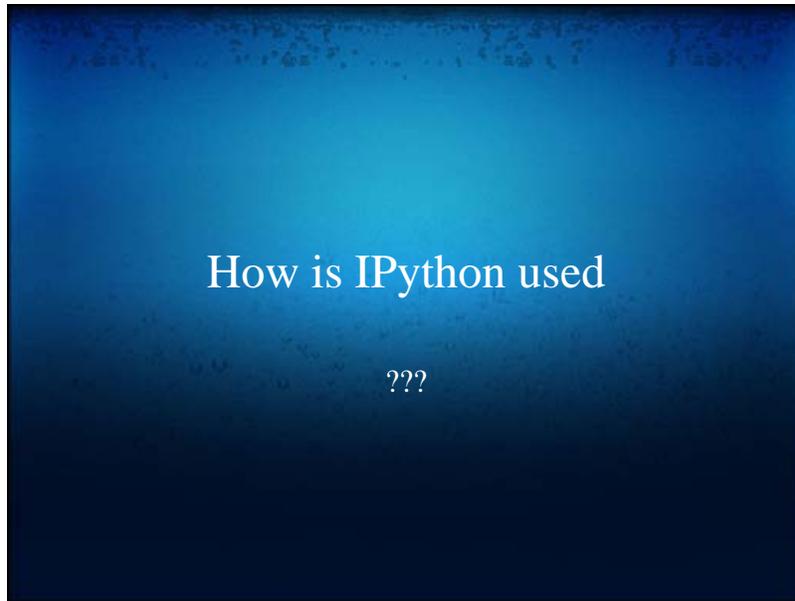
Job Scheduler Support

To perform a parallel computation with IPython, you need to start 1 Controller and N Engines. IPython has an **ipcluster** command that completely automates this process. We have support for the following batch systems.

- PBS
- ssh
- mpiexe/mpirun
- SGE (coming soon)
- Microsoft HPC Server 2008 (coming soon)

Work in progress

- Much of our current work is being enabled by ØMQ/PyØMQ. See SciPy talk tomorrow and www.zeromq.org
- Massive refactoring of the IPython core to a two process model (frontend+kernel). This will enable the creation of long awaited GUI/Web frontends for IPython.
- Working heavily on performance and scalability of parallel computing framework.
- Simplifying the MultiEngineClient and TaskClient interfaces.



ROGUE WAVE
SOFTWARE

UNIVERSITY OF COLORADO
MECHANICAL ENGINEERING

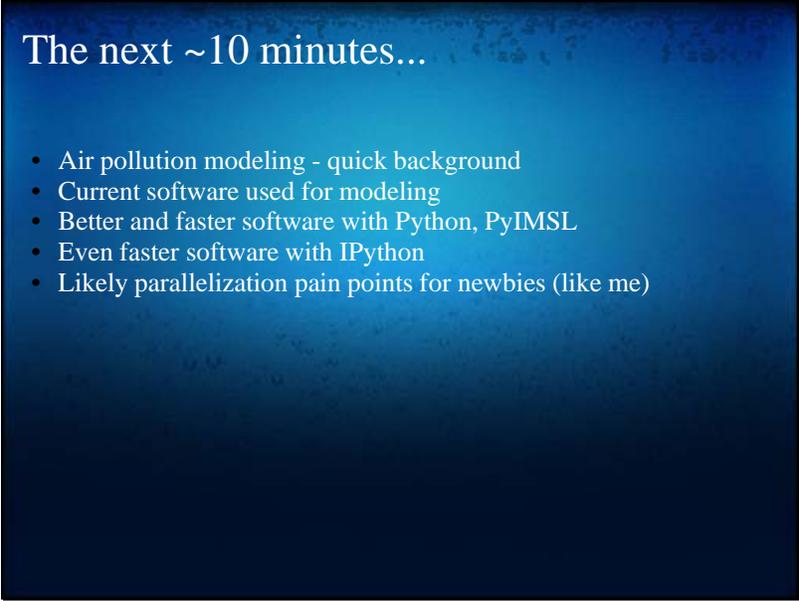
DASH
Sources & Health

Acknowledgements:

- U of Colorado**
Mike Hannigan
Jana Miford
Shelly Miller
Steve Dutton
Josh Hemann
Greg Brinkman
Ricardo Piedrahita
Rebecca Bryant
Brett Casao
Lisa Glenke
Brian Cone
Teresa Coons
Jessica Garcia
Bobby Imiger
Kally Krohn
Jaime Lahner
Stacey Louie
Fatimah Matakah
Toni Newville
Mary Beth Oshnack
Brendan Rudack
James Schroeder
Cathy Vos
Dan Williams
- Colorado State**
Jennifer Peel
- U of Washington**
Sverre Vedal
- U of Wisconsin**
Jamie Schauer
Martin Shafer
Jeff Deminter
- USGS**
Larry Barber
Greg Brown
Paul Schuster
- CDPHE**
Pat McGraw
Bradley Rink
- DPS**
Joni Rix

The views expressed in this poster are those of the authors and do not necessarily reflect the views or policies of the U.S. Environmental Protection Agency.

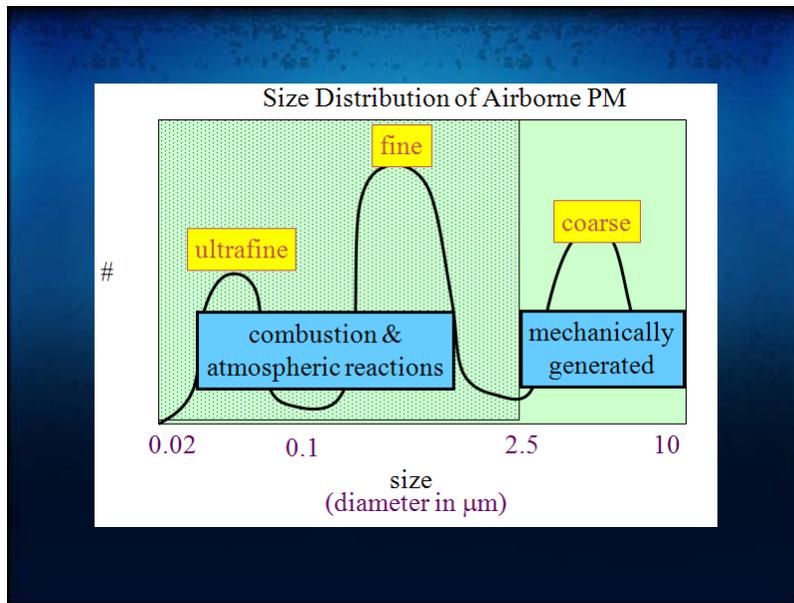
I wear two Python hats. One is for the work I do at Rogue Wave Software on our PyIMSL Studio product. I also work closely with a team of researchers (led by Dr. Mike Hannigan) doing air pollution modeling, and I use Python extensively.



The next ~10 minutes...

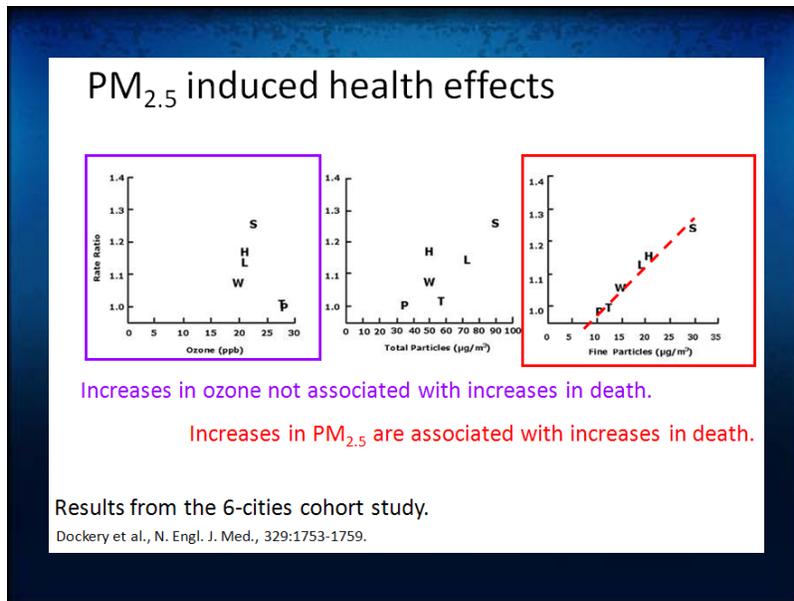
- Air pollution modeling - quick background
- Current software used for modeling
- Better and faster software with Python, PyIMSL
- Even faster software with IPython
- Likely parallelization pain points for newbies (like me)

I was at the SC show (The International Conference for High Performance Computing Networking, Storage, and Analysis) in November of 2009 when I met Wen-ming Ye, who works on Microsoft's High Performance Computing products. He put me in touch with Brian Granger, who helped me get up to speed quickly and parallelize my modeling application. This talk covers some things I have learned along the way.



Since we are responsible for much of the fine and ultrafine particulate matter (through burning stuff like gasoline, coal, wood,...) , we might ask ourselves what are the effects on our health? The study of PM_{2.5} concerns the sources and effects of ultrafine and fine particulate matter.

Slide courtesy of Mike Hannigan



Ozone is associated with health effects, though increased mortality is not one of them. Conversely, fine particulate matter is associated with increased mortality. Sweat the small stuff!

Slide courtesy of Mike Hannigan

What are the sources of the pollution?

Receptor Models

species j measured in sample i
(ambient measurements)

x_{ij}

$$x_{ij} = \sum_{k=1}^p g_{ik} f_{jk} + e_{ij}$$

residual

e_{ij}

mass from source k in sample i
(source contribution)

g_{ik}

mass fraction of species j from source k
(source profile or fingerprint)

f_{jk}

uncertainty of x_{ij}

s_{ij}

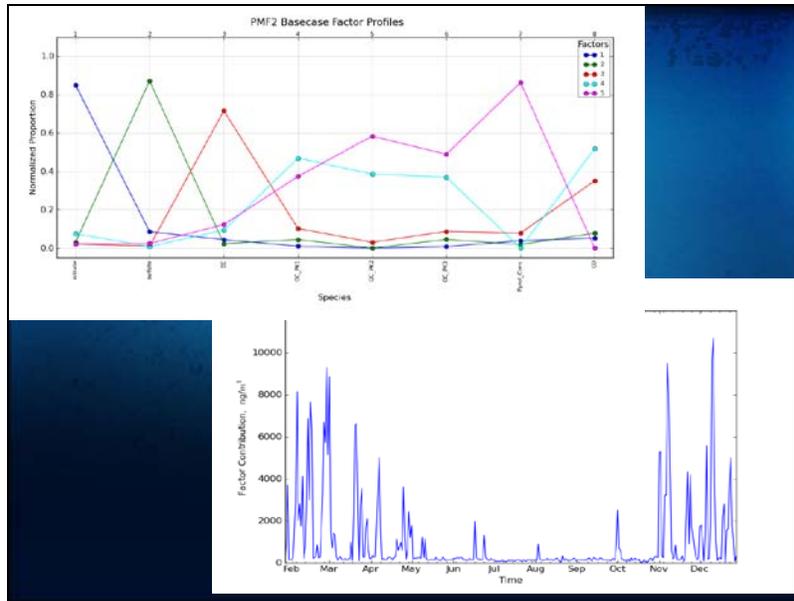
To solve, we find minimum of ...

$$Q = \sum_{i=1}^m \sum_{j=1}^n \left(\frac{e_{ij}}{s_{ij}} \right)^2$$

Receptor models used to apportion measured pollution to environmental sources (“factors” in the model). Elements of F and G are constrained to be non-negative. Non-negative matrix factorization was first developed by Pentti Paatero in the early 90s to handle this issue, later extended and now quite a useful tool in many fields. Model uncertainty: We have to choose p , the number of factors, we think can be extracted from the data. Hard to do a priori => explore different p values.

Note that you can weight specific measurements by specific uncertainties.

Slide 16



Dimensionality reduction.

Top graph: Rows of F, "factor loadings" or "profiles": what chemical species the factor accounts for.

Bottom graph: Columns of G, "factor scores" or "contributions": what the factor is doing in time.

Thus, the measurement matrix X is factorized into constituent matrices characterizing

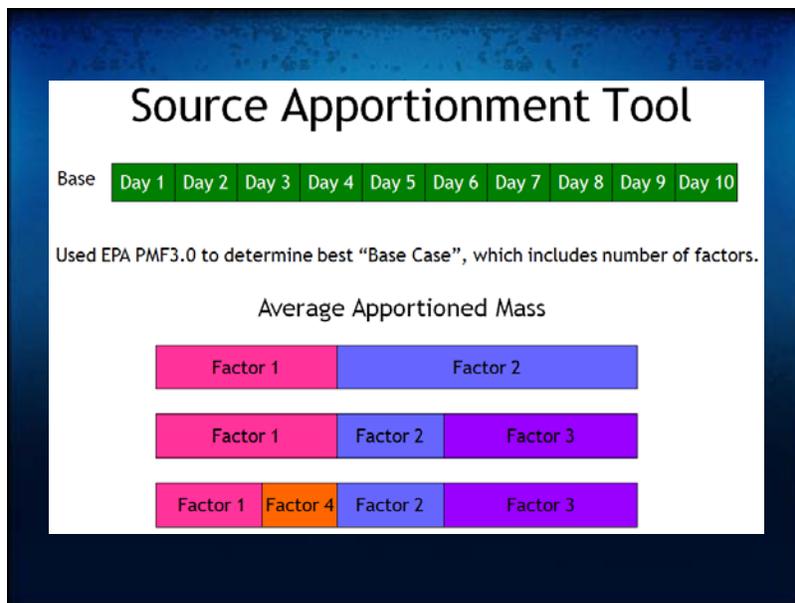
- How much a factor spits into the air
- What is in the stuff it spits out

Little sense of uncertainty can be gleaned from the model results though!

The image shows a screenshot of the EPA website for the Human Exposure and Atmospheric Sciences section, specifically the EPA Positive Matrix Factorization (PMF) 3.0 Model page. The page includes a search bar, navigation links, and a detailed description of the model. A software interface window titled 'EPA PMF v3.0 2.2' is overlaid on the page, showing various input and output fields for running the model. The interface includes sections for Input Files (Concentration Data File, Uncertainty Data File), Missing Value Indicator, Output Files (Output Folder, Output File Type), and Program Configuration (Configuration File). The EPA website text explains that receptor models provide scientific support for current ambient air quality standards and for implementation of these standards by identifying and quantifying contributions for source apportionment. It also notes that the EPA PMF 3.0 model works on Windows XP and Windows Vista, and requires a 2.0 GHz processor and at least 1 GB of memory.

Tools exist to perform this pollution source apportionment. The US EPA provides a MATLAB-based tool called EPA PMF that is used by many practitioners modeling air pollution. Real-world impact: policy decisions are shaped by the results of our tools, but EPA PMF has some room for improvements...

<http://www.epa.gov/heasd/products/pmf/pmf.html>



Experiment with different numbers of factors. How many factors can we reasonably extract from the data? How stable are the results if the data are perturbed?

In the middle row, assume Factor 1 looks like "automobile fuel combustion". If you add another factor to the solution, Factor 1 might split out to gas and diesel combustion. If you add too many factors though the mathematical factors may not have real-world analogues, and the degenerate case is just one factor per chemical species...

Slide courtesy of Mike Hannigan

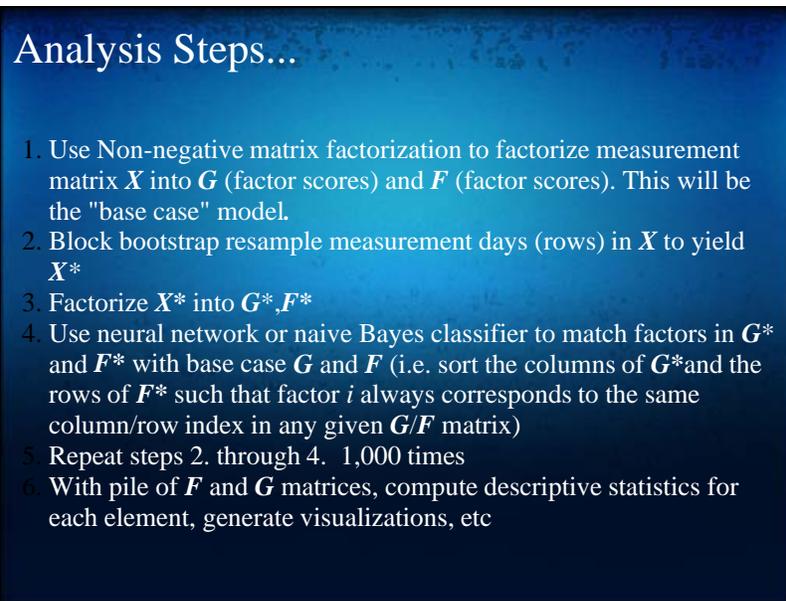
Source Apportionment Tool

Base	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	Day 10
Bootstrap 1	Day 9	Day 10	Day 3	Day 1	Day 4	Day 7	Day 3	Day 6	Day 3	Day 7
	Created our own tool using Python and PyIMSL Studio				⋮					
Bootstrap 99	Day 3	Day 5	Day 3	Day 6	Day 7	Day 1	Day 5	Day 5	Day 1	Day 2

Critical to success of the bootstrap ...
Matching factor 4 from bootstrap 86 to a
factor in the Base Case

Mathematical "factor 1" is row 1 of F and column 1 of G, and say it looks like gasoline combustion. In a bootstrap solution this same mathematical factor might appear but just happen to be in row 2 of F, column 2 of G. We need to sort the rows/columns of the solutions such that "factor 1", i.e., row 1 and column 1 in every F,G matrix always corresponds to gasoline combustion.

Slide courtesy of Mike Hannigan



Analysis Steps...

1. Use Non-negative matrix factorization to factorize measurement matrix X into G (factor scores) and F (factor scores). This will be the "base case" model.
2. Block bootstrap resample measurement days (rows) in X to yield X^*
3. Factorize X^* into G^*, F^*
4. Use neural network or naive Bayes classifier to match factors in G^* and F^* with base case G and F (i.e. sort the columns of G^* and the rows of F^* such that factor i always corresponds to the same column/row index in any given G/F matrix)
5. Repeat steps 2. through 4. 1,000 times
6. With pile of F and G matrices, compute descriptive statistics for each element, generate visualizations, etc

Step 1 is the computationally time consuming part.

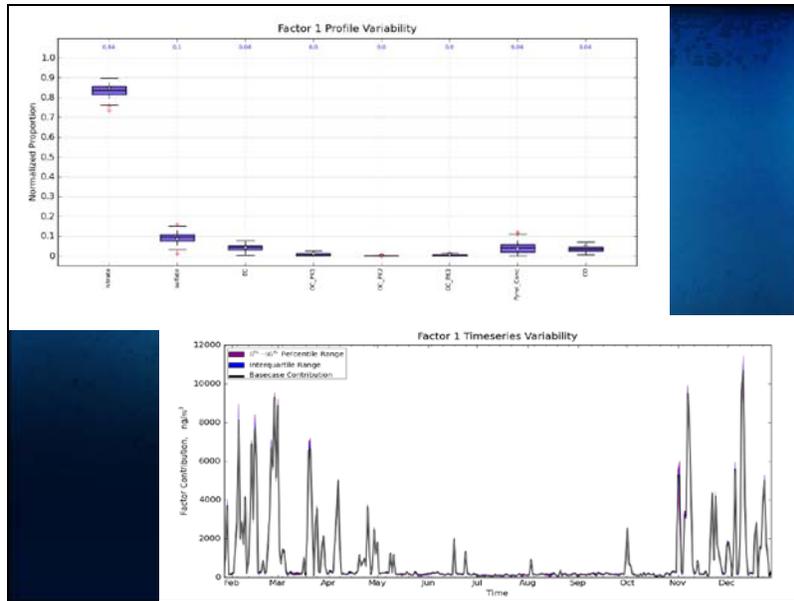
Step 2: Stationary Block Bootstrap. Mean found by looking at median p for AR(p) models fit to each pollutant species' time series.

Step 2: keep track of which days were resampled and where they showed up in the synthetic X !

Step 4 uses pattern classification algorithms in PyIMSL Studio

Embarrassingly parallel analysis

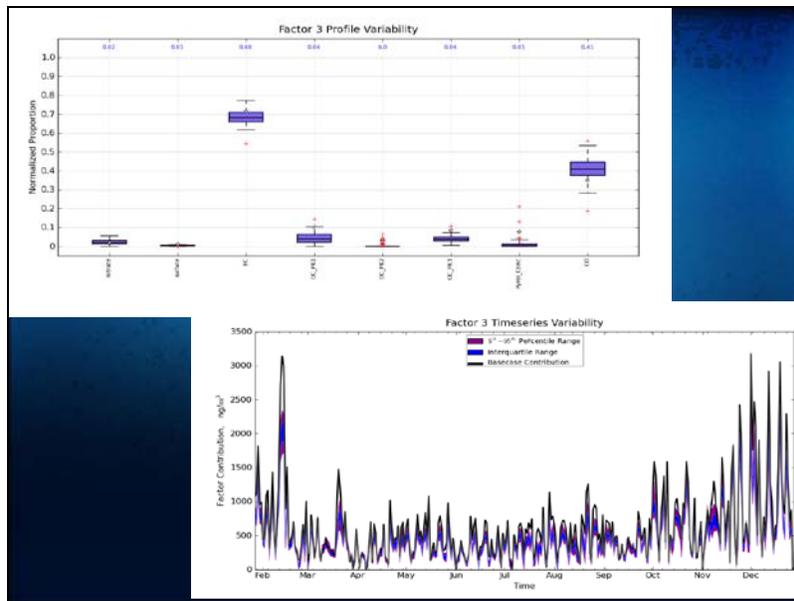
Slide 21



These results show bias and variability in the model results due to sampling error. This is what the huge simulation gets us!

This factor is what we might call the "nitrate factor"

Slide 22



More variability and bias seen in Factor 3 than in Factor 1.

This factor is what we might call "diesel combustion", and we should question the base case model results for this factor due to the variability and bias we see under the bootstrap simulation.

How long does this modeling take to run?

1,000 bootstrap replications on my dual-core laptop...

- EPA PMF 3.0
 - ~ 1 hour and 45 minutes
 - Black box, only single core/processor actually used
- Python and PyIMSL Studio
 - ~ 30 minutes
 - MKL and OpenMP-enabled analytics means I don't have to do anything to use both of my cores, It Just Works

Can we make this faster?

Even though key algorithms might be already be multi-threaded, the entire analysis itself can be parallelized too.

Code can actually run slower using MKL and OpenMP due to thread overhead (keywords: OpenMP pragmas and nesting). Can't just blindly turn them on if you are concerned about speed, need to understand nature and scale of your problem. Automatic tuning of libraries on specific hardware, for specific problem types (e.g. ATLAS), is the wave of the future.

```
from IPython.kernel import client

#Set up each Python session on the clients...
mec = client.MultiEngineClient(profile='DASH')
mec.execute('import os')
mec.execute('import shutil')
mec.execute('import socket')
mec.execute('import parallelBlock')
mec.execute('reload(parallelBlock)')
mec.execute('from parallelBlock import parallelBlock')

#Task farm-out the 6 analysis steps...
tc = client.TaskClient(profile='DASH')
numReps = 1000
taskIDs = [ ]
for rep in xrange(1,numReps):
    t = client.MapTask(parallelBlock, args=[rep])
    taskIDs.append(tc.run(t))
tc.barrier(taskIDs)
results_list = [tc.get_task_result(tid) for tid in taskIDs]

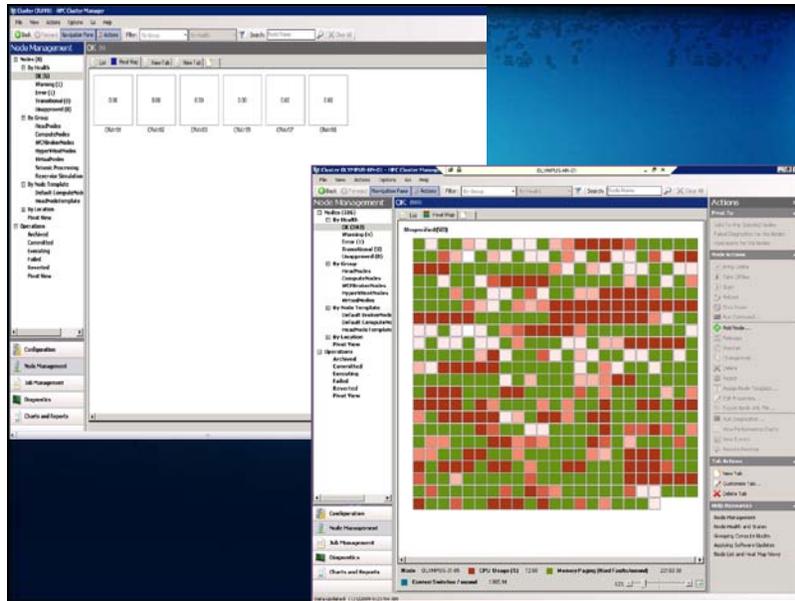
for task, result in enumerate(results_list):
    #Unpack results from each iteration and do analysis/visualization
```

IPython to the rescue...

parallelBlock is just a Python function that performs the six analysis steps outlined a few slides back. Very little boilerplate needed to parallelize the simulation. Basically, we set up the environment on each engine in the cluster, map the six-steps analysis out to the engines, wait until the engines have performed the work, aggregate the results from each engine, and then process the results...

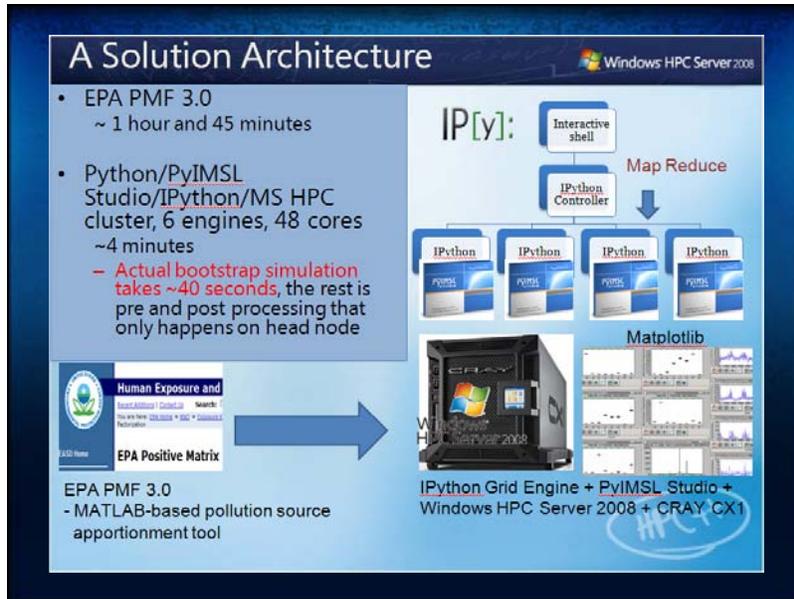
What I love about IPython (besides the interactive shell) is that while it supports multiple ways to do parallel computing, I can use it to quickly and easily parallelize “embarrassingly parallel” problems like I have in this context.

Slide 25



MS HPC Server makes administering cluster easy! Lots of 3rd party support, IPython parallel features dump xml configuration read in by HPC Server. I can remote desktop to each engine and browse the local file system, see a real-time heat map of processor load for each core, manage job schedules, etc, all in an intuitive interface. Great for someone like me who does not know much about the IT issues involved with managing a cluster.

Lots of value in just taking advantage of the cores on one's desktop/laptop, HPC Server not needed. Code works the same, but different IPython profiles would be used to distribute the code on your dual-core laptop vs working with MS HPC Server to work over a cluster.



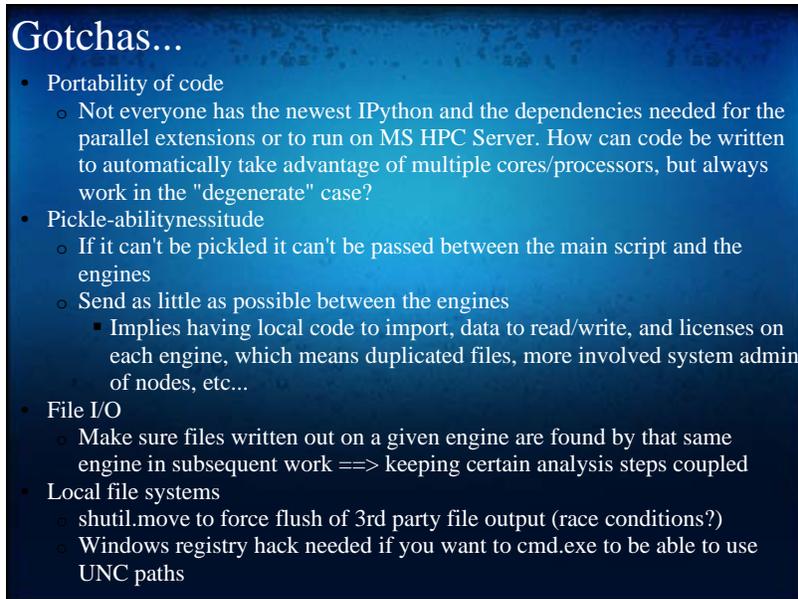
Slide courtesy of Wen-ming Ye, Technical Evangelist (HPC), Microsoft. Wen-ming got me access to the Cray CX1 cluster.

What make parallelizing hard...

There are complex aspects of my application that have nothing to do with cool mathematics...

- Existing application used for a couple of years, not written with parallelization in mind from the start
- Analytics are not just simple calls to pure Python
 - PyIMSL algorithms wrap ctypes objects that sometimes involve C structures (that may contain other complex types), not just simple data types
 - 3rd party Fortran 77 dll called
 - Does it's own file I/O, which is critically important to read, but for which I have little control of (with respect to file names and paths)
- Big time sink is in post-processing of results to set up data for visualization, a whole separate aspect not related to the core analysis

What happens when we move beyond simple examples of computing a matrix product?



Gotchas...

- Portability of code
 - Not everyone has the newest IPython and the dependencies needed for the parallel extensions or to run on MS HPC Server. How can code be written to automatically take advantage of multiple cores/processors, but always work in the "degenerate" case?
- Pickle-ability/nessitude
 - If it can't be pickled it can't be passed between the main script and the engines
 - Send as little as possible between the engines
 - Implies having local code to import, data to read/write, and licenses on each engine, which means duplicated files, more involved system admin of nodes, etc...
- File I/O
 - Make sure files written out on a given engine are found by that same engine in subsequent work ==> keeping certain analysis steps coupled
- Local file systems
 - `shutil.move` to force flush of 3rd party file output (race conditions?)
 - Windows registry hack needed if you want to `cmd.exe` to be able to use UNC paths

Bullet 1: Could be a simple try/except around importing IPython, but...

IPython 0.11 will help this issue, as will time: soon everyone will have multicore machines and be writing parallelized applications from the start.

Gotchas...

- Debugging and diagnostics
 - Sanity checking and introspection can be more involved

```
def parallelBlock(rep):
    ini_file = 'pmf_bootstrap.ini'
    fh = open("NUL", "w")
    subprocess.Popen('pmf2wopt.exe %s %i' \
        % (ini_file, rep), stdout=fh).communicate()
    hostname = socket.gethostname()
    try:
        #Analysis steps. Nothing to do if PMF did not
        #converge for this bootstrap replication...
    except:
        return (-id, hostname, [], [], [], [], [])
    else:
        return (id, hostname, v, w, x, y, z)
```

Whereas the non-parallel code used to just return a tuple of various results, now I have try/except blocks and also return information about the host and process ID, which is helpful for when I have to inspect a given engine's file system say, and track results across various processes.

I have to keep the PMF file I/O and subsequent analysis linked. engine ID is variable "id" in each engine's namespace

I'm happy to talk outside of this
presentation!

josh.hemann@roguewave.com