

# ACTS Workshop

Carlos Rojas  
Interactive Supercomputing

01010110100100111010001000111101011010010010010100101

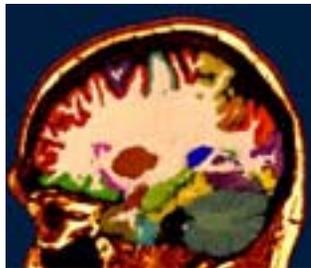


INTERACTIVE  
*supercomputing*

# Our Vision

---

“Bridge the gap between easy-to-use desktop modeling, simulation and development tools on personal computers with the power, scalability and low cost of parallel computer systems, clusters and grids”



INTERACTIVE  
*supercomputing*



# Our Mission

---

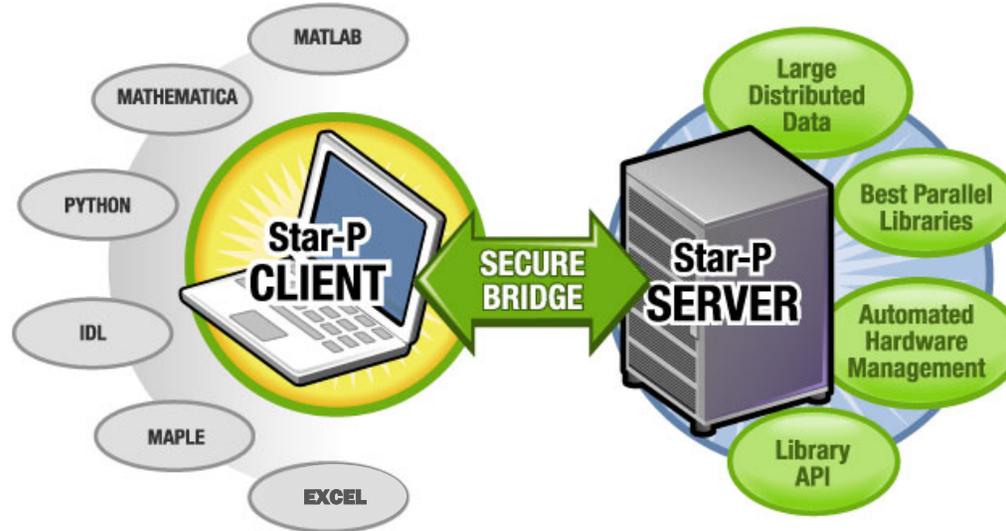
“Establish Star-P as the de facto standard software platform to drive improved productivity in high performance technical computing”



INTERACTIVE  
*supercomputing*



# How Star-P works



- Star-P consists of desktop & server software
- Desktop software – Star-P Client
  - Overloads or intercepts desktop tool functions
  - Connects and communicates to/with server software – securely
- Server software – Star-P Server
  - Manages and directs resources – memory, cpu’s and I/O
  - Contains World Class Libraries for parallel execution
  - User & Session management



# What's the Value

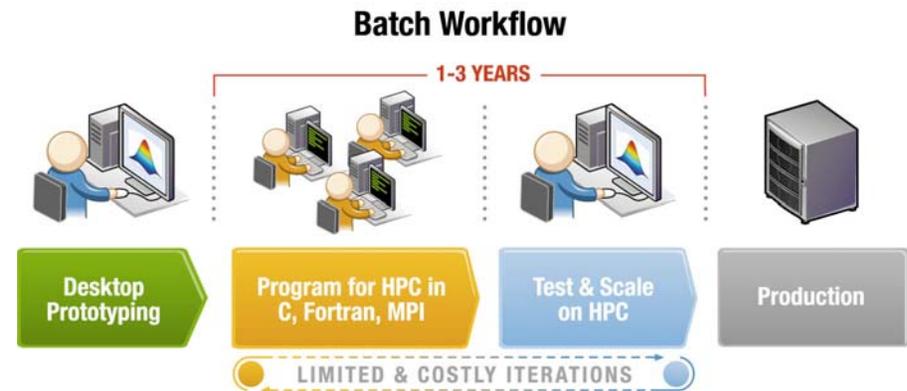
---

- Desktop Users:
  - No change in religion
  - Interactivity
    - On parallel machines
    - For large data
  - No reprogramming
    - No C, Fortran, MPI
  - Reduced run times
    - Not hours or weeks
  - Continued model optimization
- Organization
  - Collapse development cycles
  - Reduce costs
  - Broaden usage
  - Shorten solution time
  - Accelerate research



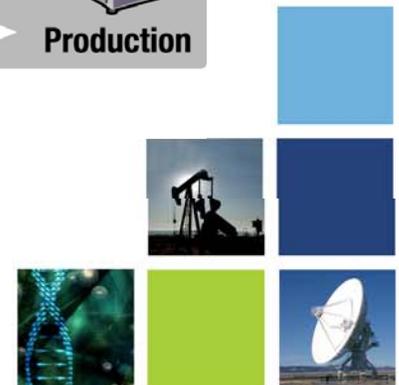
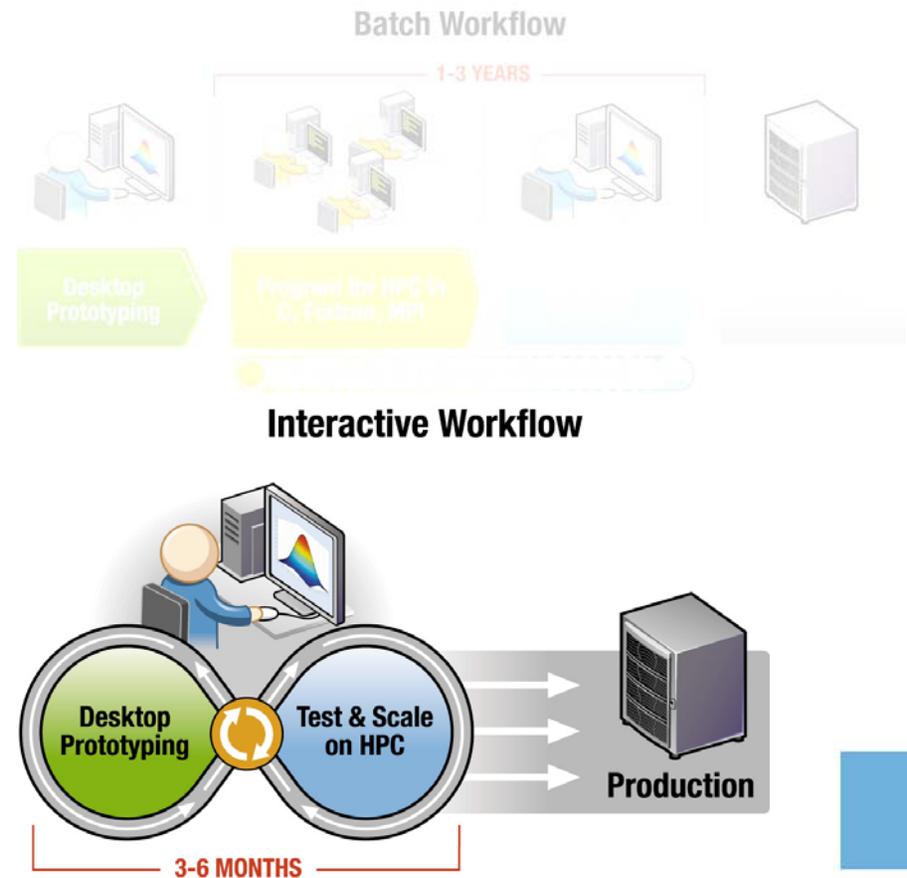
# Parallel Development Takes Too Long

- Months or years are spent porting from desktops to parallel systems
- No interactively on parallel machines from desktop
- Little ability to iterate
- Long compute times for batch runs; hours-days
- Analyst's ability to optimize the model is limited



# Interactive workflow solves the problem

- No re-programming
- Interactive desktop environment
- Rapid iteration, refinement
- 50% reduction in time-to-solution

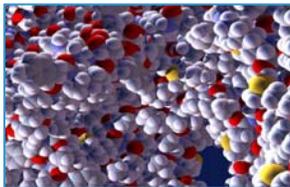


# What Markets & Applications benefit

---

- Market focus

- Government
  - Defense
  - Intelligence
  - Civilian
- Bio-Pharma-Medical
- Financial Services
- Manufacturing
- Academic/Research



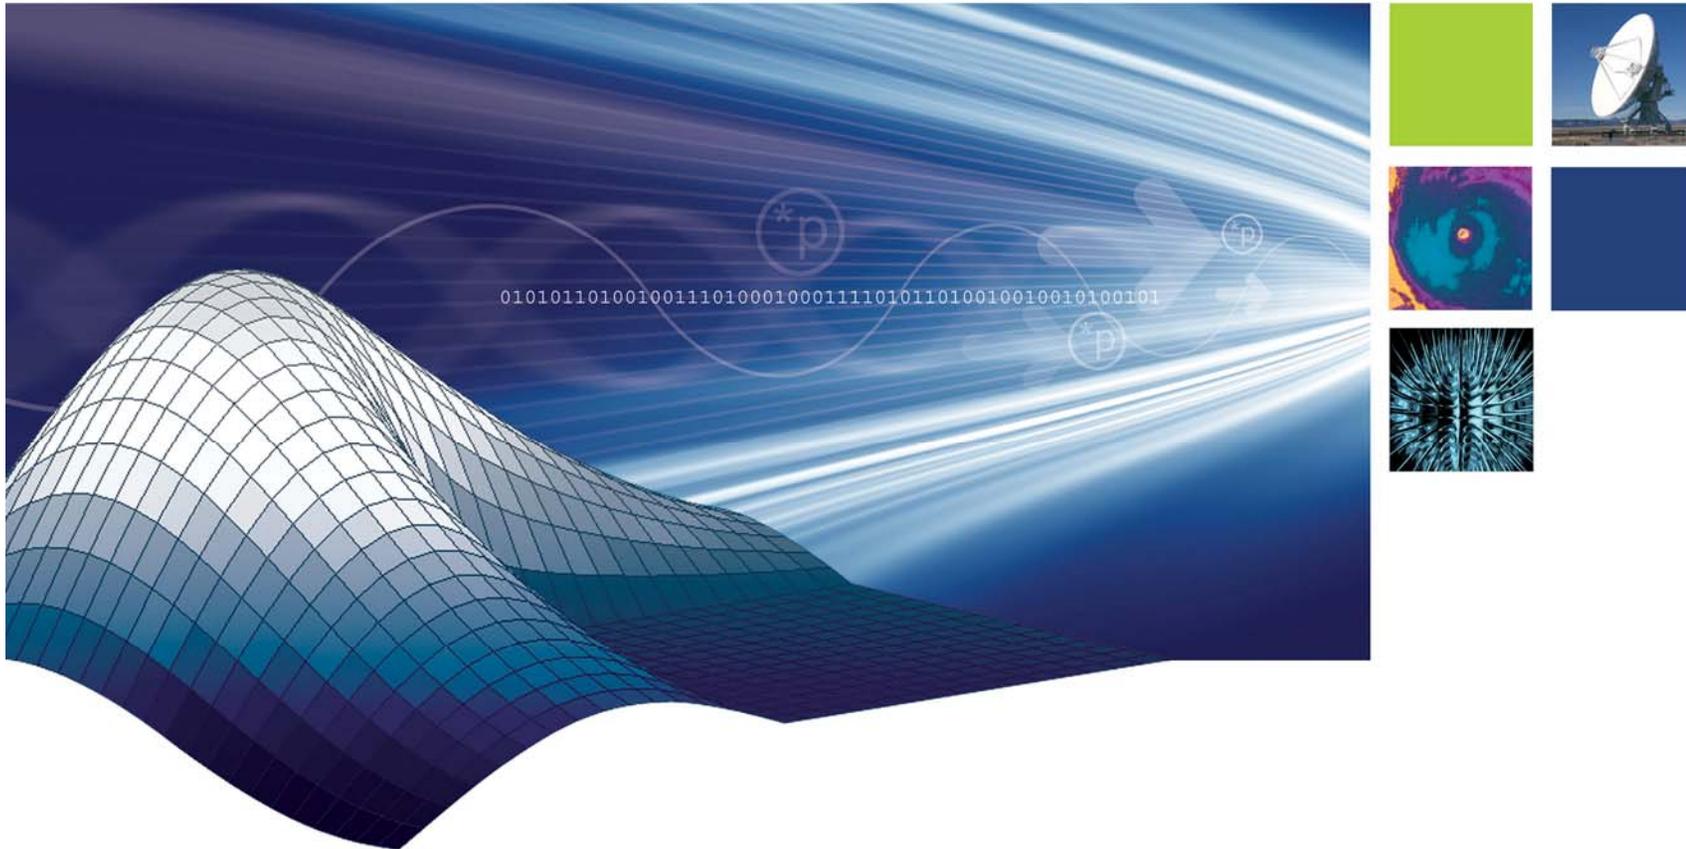
INTERACTIVE  
*supercomputing*

- Applications

- Image Processing
- Signal Processing
- Financial modeling & Analysis
- Graph theory and analysis
- Statistical analysis
- Molecular modeling & simulation
- Finite Element analysis



# Product Overview





- Introduction
- Star-P Usage
  - Modes of operation
- Star-P Server
- Architecture
  - Application Programming Interface (API)
- System Requirements



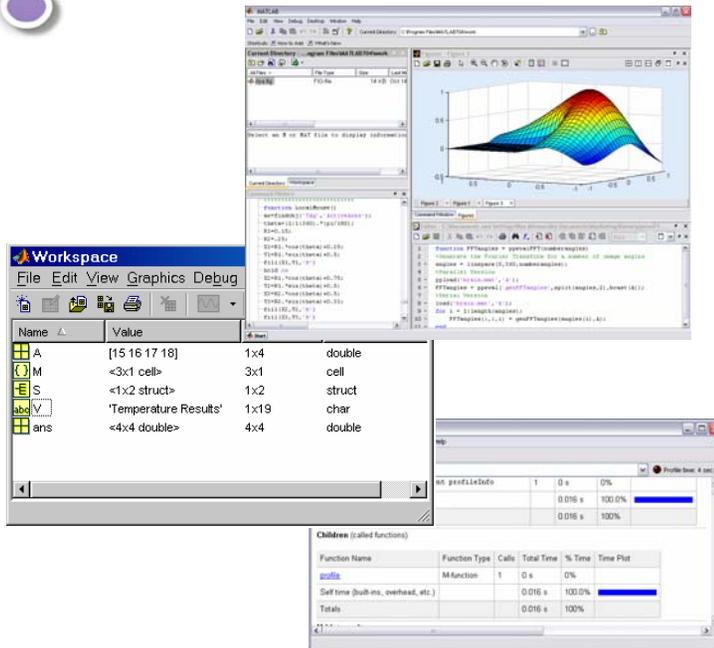
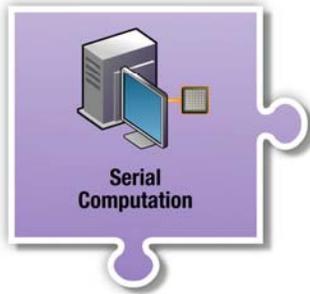
# Star-P Introduction

---

- Intended for codes, models, or algorithms that are no longer solvable/efficient on desktop systems
- User still works in Matlab
  - 100's of commands and functions operate in parallel
  - Some new constructs – (pp, dd, \*p, etc)
  - Serial codes are not good parallel ones - necessarily
- Computations sent to parallel server
  - Tasks and data are spread across processors or nodes
  - World class libraries perform parallel operation
  - Manages inter-processor communication
  - Manages flow and memory storage of large data



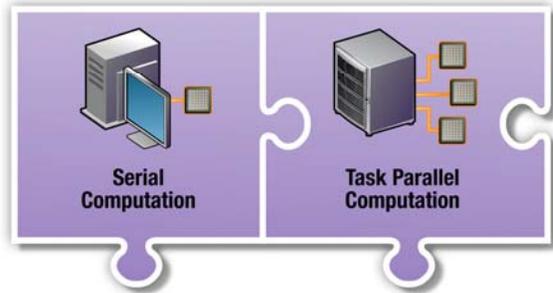
# Star-P usage – Serial operations



- Use Matlab
  - File Editor
  - Profiler
  - Debugger
  - Array Editor
  - Desktop
  - Visualization
  - Small Calculations
- Small operations ~ 100K string
- Computations taking less the .5 seconds



# Star-P usage – Task parallel

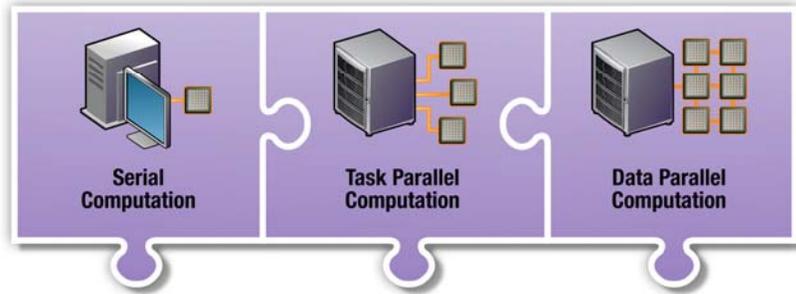


- Data size < 100MB
- Execution time > .5 second
- Code separable in time
- Embarrassingly parallel apps
- Incorporate Star-P's **ppeval**

```
1  %Generate the Fourier Transform on 10 degree spacing
2  angles = linspace(0,360,37);
3  %Serial Version
4  load('brain.mat','A');
5  for i = 1:length(angles);
6      FFTangles(:, :, i) = genFFTangles(angles(i), A);
7  end
8
9
10
```



# Star-P usage – “Data parallel”



```
% explicitly parallel with *p
n=10000*p

% implicitly parallel
A = rand(n, n);

% implicitly parallel
x = randn(n, 1);

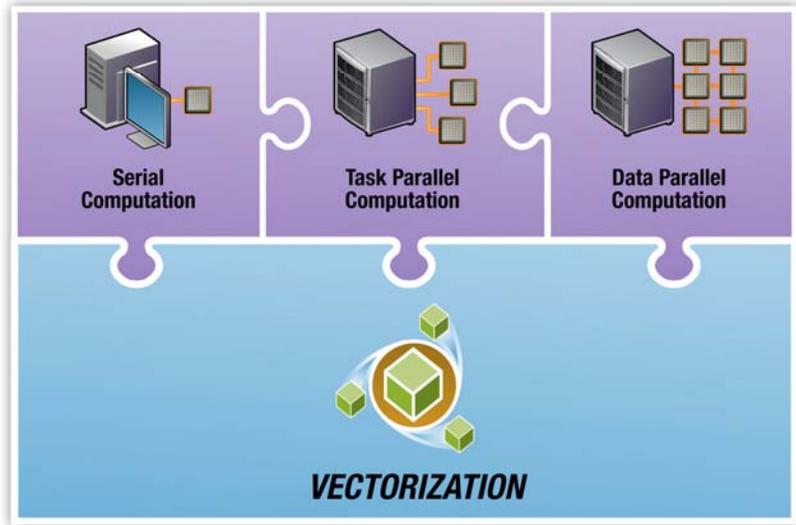
% implicitly parallel
y = zeros(size(x));

while norm(x-y) / norm(x) > 1e-11
    y = x;
    x = A*x;
    x = x / norm(x);
end;
```

- Data sizes >100MB
- Execution time > .5 second
- Data not separable
- Operations on vectors and matrices
- Incorporate \*P
  - Global parallelism
  - Variables become parallel
  - Propagation occurs
    - Results are parallel
    - Functions performed on parallel data



# Star-P usage –\*p w-vectorization



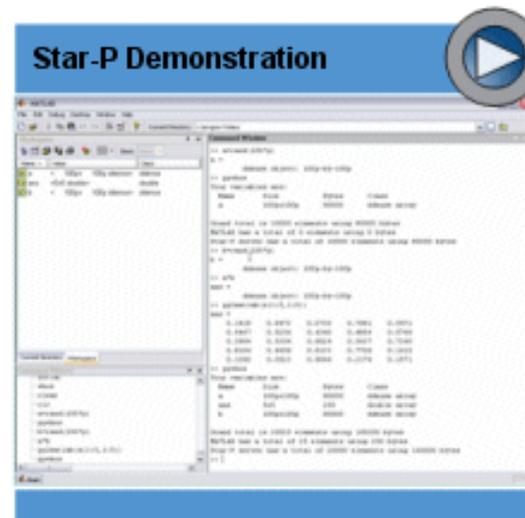
- Data sizes >100MB
- Execution time > .5 second
- Data not separable
- Operations not on vectors and/or matrices
- Vectorization needed
- Incorporate \*P



# Star-P Usage



- Straight serial Matlab
- Task parallel jobs
- Large data problems
- Combine them all



# The Star-P Server



# Star-P Server

---

- Library Management
  - ISC and World class parallel libraries
  - Libraries from commercial providers
    - IMSL, NAG, etc
  - Other commercial libraries and proprietary via Star-P API/SDK
- Data Management
  - Efficient handling of distributed data
  - Distributed Sparse Matrix operations
  - Distributed Dense Matrix operations
  - Data decomposition
    - Row or Column distributed



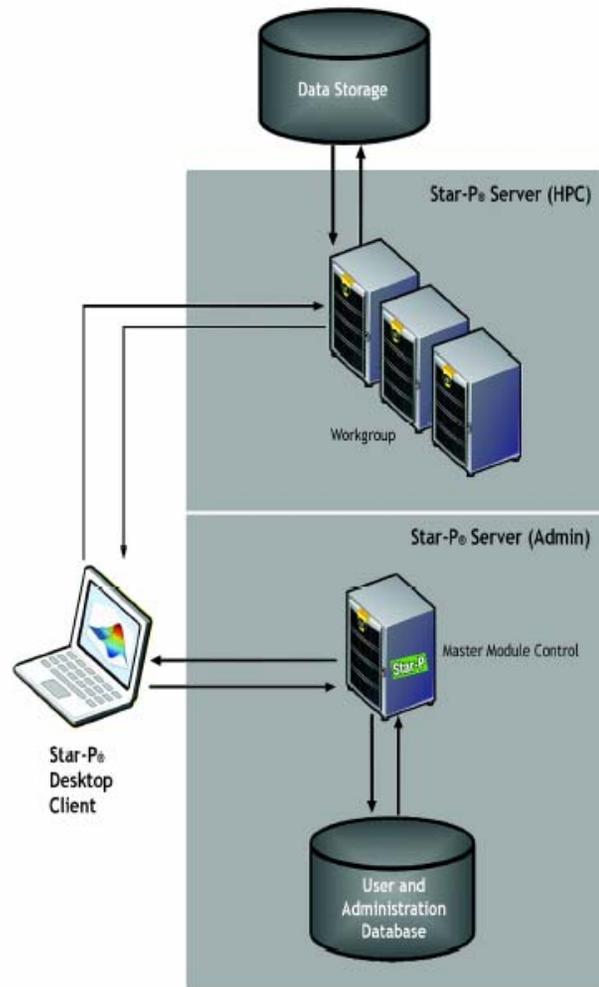
# Star-P Server

---

- Interactive Computing
  - Assignment of CPU's to sessions
  - Schedules memory and CPU allocations
  - Commercial grade server – error recovery
  - Comprehensive System Administration
  - Standard ports and protocols for communication
  - Support for multiple clients/users
  - Support for multiple servers
  - Workgroup pooling for cpu resources
  - Not a compiler or simulator
  - Customizable security



# Star-P Architecture - Logical



- Client
- Workgroup server(s)
- Master Control Module
- User & Admin database
- Data Storage

A screenshot of the Star-P Application web interface in Mozilla Firefox. The page title is 'Star-P Application - Mozilla Firefox'. The interface includes a navigation menu on the left with options like 'Manage Users', 'Manage Groups', 'Manage Sessions', 'Manage HPCs', 'Manage Applications', 'View Workgroups', 'Install Windows', 'Client Version: trunk-3515', and 'Install Linux Client Version: trunk-3515'. The main content area is titled 'Administrative' and 'Listing HPC Sessions'. It features a table with the following data:

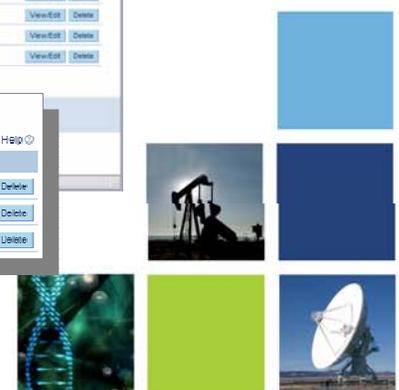
Session Name	Num CPUs	Exclusive	Max Mbytes	
Default	2	No	64000	<a href="#">View/Edit</a> <a href="#">Delete</a>
Kong 16	16	No	64000	<a href="#">View/Edit</a> <a href="#">Delete</a>
Kong 8	8	No	64000	<a href="#">View/Edit</a> <a href="#">Delete</a>
Shared Session for altix2	4	No	8000	<a href="#">View/Edit</a> <a href="#">Delete</a>
altix-4	4	No	8000	<a href="#">View/Edit</a> <a href="#">Delete</a>

A screenshot of the 'Listing hpcs' web interface. It displays a table with the following data:

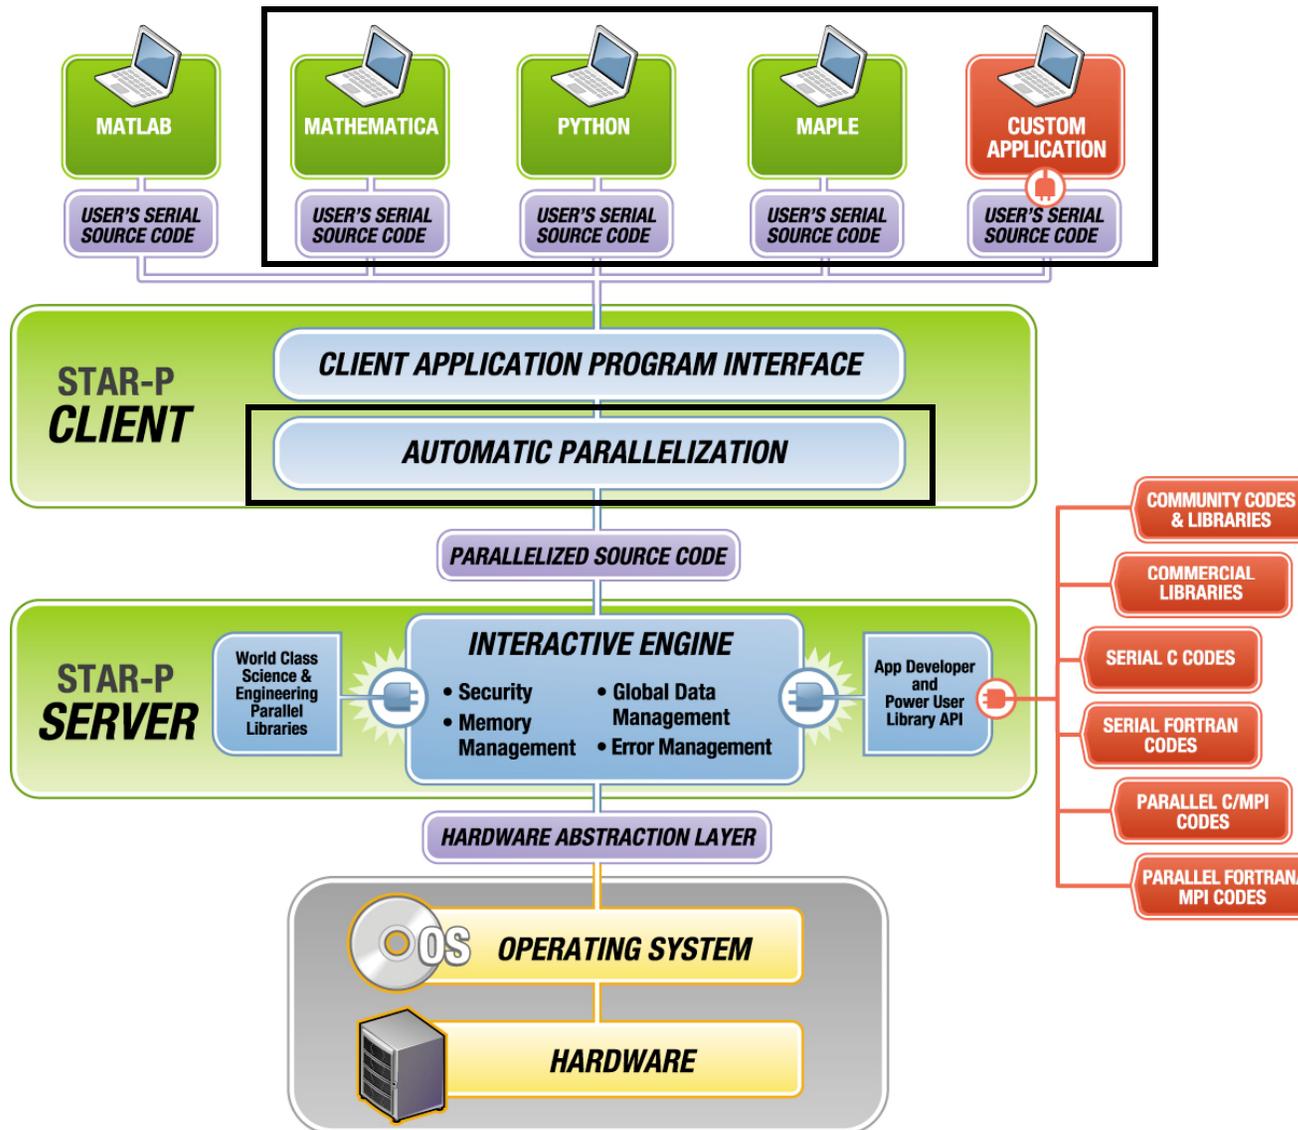
Name	Ip address	Num cpus	Mbyte size	Configuration	
altix	10.0.1.54	0	0192	SGIAItix	<a href="#">View/Edit</a> <a href="#">Delete</a>
				SGIAItix	<a href="#">View/Edit</a> <a href="#">Delete</a>
				SGIAItix	<a href="#">View/Edit</a> <a href="#">Delete</a>

A screenshot of the 'Listing users' web interface. It displays a table with the following data:

id	Login	Is admin	Is shared	
1	admin	true	false	<a href="#">View/Edit</a> <a href="#">Delete</a>
3	ajenkins	false	false	<a href="#">View/Edit</a> <a href="#">Delete</a>
2	Shared_User	false	true	<a href="#">View/Edit</a> <a href="#">Delete</a>
4	puig	false	false	<a href="#">View/Edit</a> <a href="#">Delete</a>
5	aha	false	false	<a href="#">View/Edit</a> <a href="#">Delete</a>



# Star-P Architecture - Conceptual

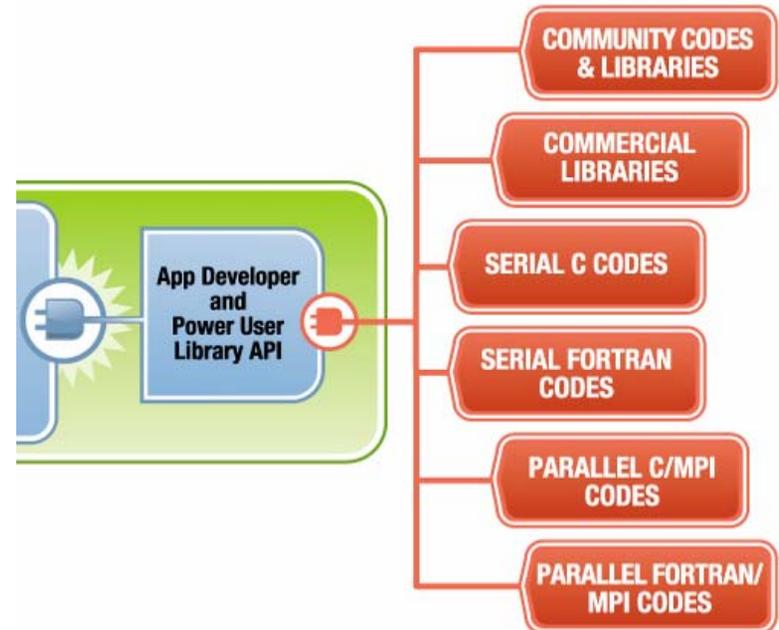


Future



# Plug into Star-P through Server API

- Through Matlab, access:
  - Your own library functions
  - Proprietary solvers
  - Other commercial algorithms
  - Specialized hardware (FPGA's)
- Serial and parallel codes
  - Coarse-grained “multiply effect”
  - Parallel codes



# System requirements & support

---

- Server side

- Itanium or x86-64
- 4+ CPU's (cores)
- 4GB+ memory
- 64-bit Linux
  - SUSE 9.2 or higher
  - RHEL 4.1 or higher
  - Fedora 4 or higher
- Upto 16 nodes
  - Certified
  - Gige, Infiniband or Myrinet

- Client side

- IA32 processors
- O/S
  - Windows XP
  - SUSE 9.3 or higher
  - RH Fedora 3 or higher
- Connect to Server
  - 100 Mbps or higher
  - Gige recommended
- Matlab R14 SP2 or higher



# More demonstrations

---

<http://www.interactivesupercomputing.com/IT/>



# Questions...

