Nimrod/EnFuzion

Colin Enticott
Slavisa Garic
Tom Peachey

Monash University

Today's time table

- 9:30 – 11:00 Introduction to Cluster Computing
  - eScience
  - History of Nimrod/EnFuzion
  - Clusters
  - Parameter sweeps
  - EnFuzion Demonstration of EnFuzion running under Windows
- 11:30 – 1:00 Grid Computing
- 2:00 – 3:00 Distributed Optimization
- 3:30 – 4:30 eScience Projects

The Nimrod team

- http://www.csse.monash.edu.au/nimrod

Project head	David Abramson
Nimrod/G core technology	Slavisa Garic
Scheduler & APST Interface	Shahaan Ayyub
Portal and Web Services	Colin Enticott
Active Sheets	Paul Roe, Gavin Cheuk & Slavisa Garic
Applications	Colin Enticott and Tom Peachey
Nimrod/O core technology	Tom Peachey & Andrew Lewis

eScience

- What is eScience?
- Categories
  - Data storage, informational services, metadata etc.
  - Real time data
  - Data transformations
    - Visualisation
  - Computational experiments
    - modelling

Modelling

- Computational experiment
  - Needs CPU power
- Can simulate real world experiments

This is how it all began

Air pollution modelling circa 1990

- Want to control Ozone
  - What happens if we reduce NOx?
  - What happens if we reduce ROC?

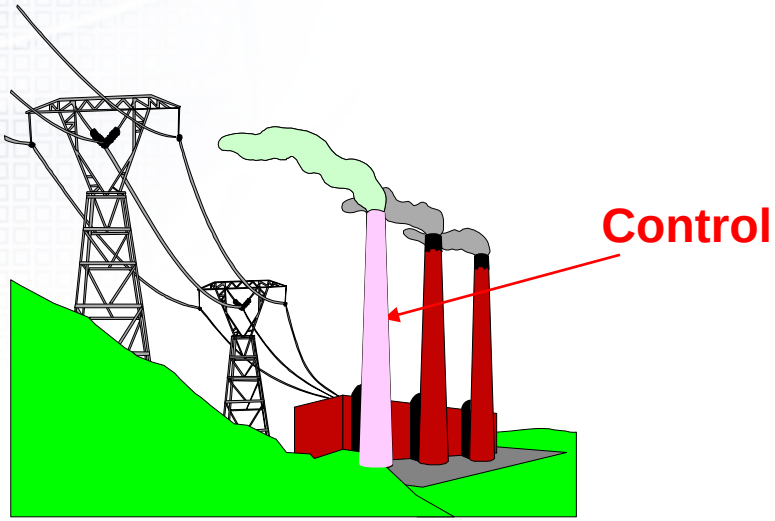

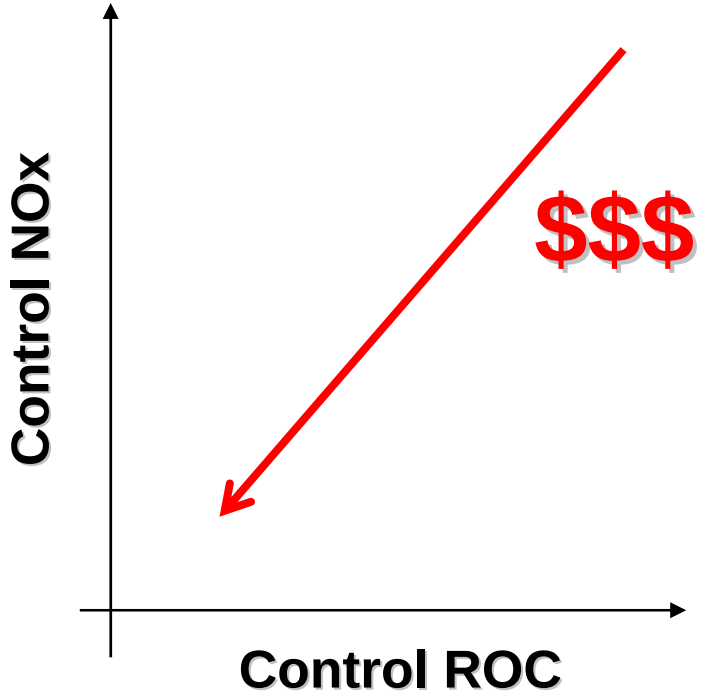

But, Ozone chemistry is non-linear

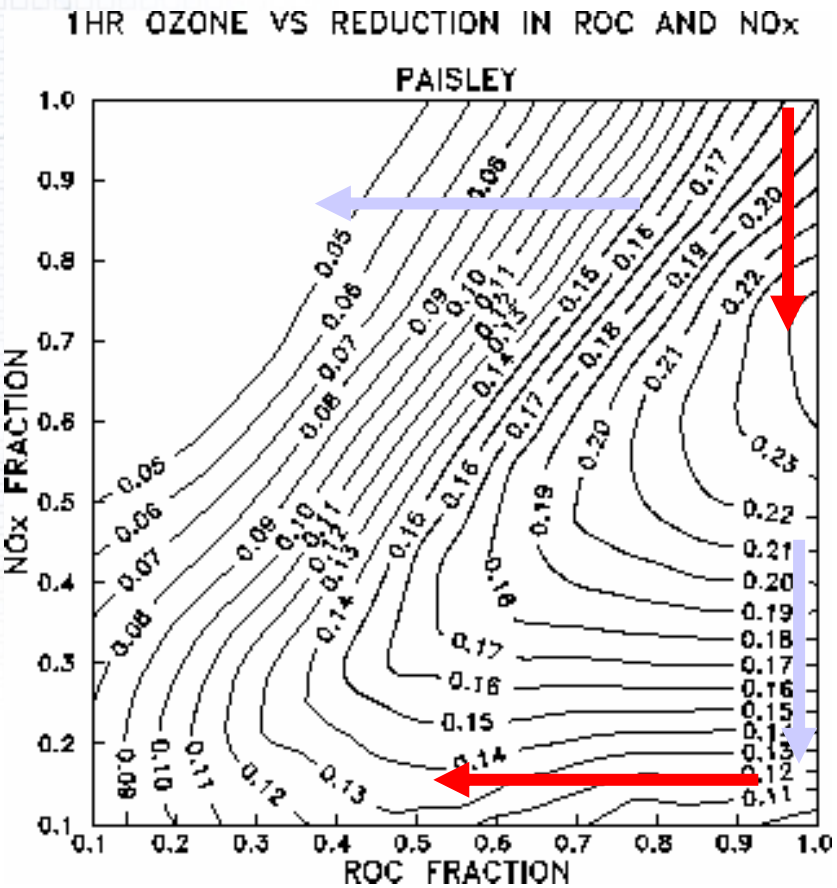

Decreasing NOx increases Ozone

Decreasing ROC decreases Ozone

Decreasing NOx decreases Ozone

Decreasing ROC increases Ozone

Parametric Execution

- Study the behaviour of some of the output variables against a range of different input scenarios.
  - For example, what is the expected Ozone output given the NOx and ROC levels?
- Computations are uncoupled (file transfer)
  - The result of one parameter set does do affect the results of another parameter set
- More realistic simulations
  - Increasing the number of values of NOx and ROC to explore will produce a higher resolution result

Multiple Runs

- to validate the model
  - Comparison with real world data
- to explore the “parameter space”
  - What is the expected behaviour given new parameters?
- for non-deterministic models to average over an ensemble
- to find parameters that optimize some result
  - After lunch

Cross product

- For each value of NOx
  - For each value of ROC
    - What is the Ozone output level?

	0.1	0.2	0.3	0.4	0.5	…
0.1	Job	Job	Job	Job	Job	
0.2	Job	Job	Job	Job	Job	
0.3	Job	Job	Job	Job	Job	
0.4	Job	Job	Job	Job	Job	
0.5	Job	Job	Job	Job	Job	
…						

Cluster computing

- Where can we run these jobs?

- Idle workstations
  - Lab computers
  - Using EnFuzion
  - EnFuzion will wait until the computer is idle before it starts a job

- Dedicated computational resources
  - Resources that allow only remote access to the computers
  - These resources have been set up for the sole purpose of running computational heavy experiments.

Cluster computing - queuing

- What happens if the demands increase?

- Idle workstations
  - EnFuzion will run one job per node.
  - First in first served

- Dedicated computational resources
  - Have quotas and fair sharing policies

Cluster computing – more power

- What happens if I need more computers?
- Use Nimrod/G
  - After morning tea

Nimrod & EnFuzion

Nimrod History

- Project History
  - Initial Cluster version 1994
  - Nimrod/G 1997
  - EnFuzion (Axceleon) 1997
  - Nimrod/O 1999

Nimrod Goals

- Goals
  - Supports parametric execution
  - Execute programs
  - Varying parameters
  - Simple scatter/gather
  - Make parallel computing easy for parametric problems

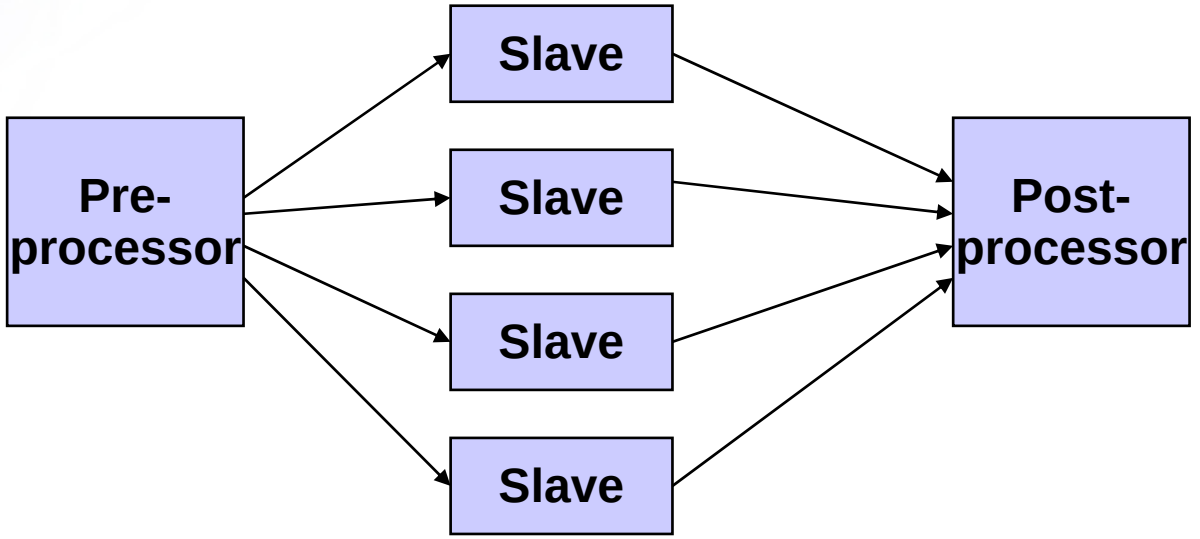

Nimrod Execution Architecture

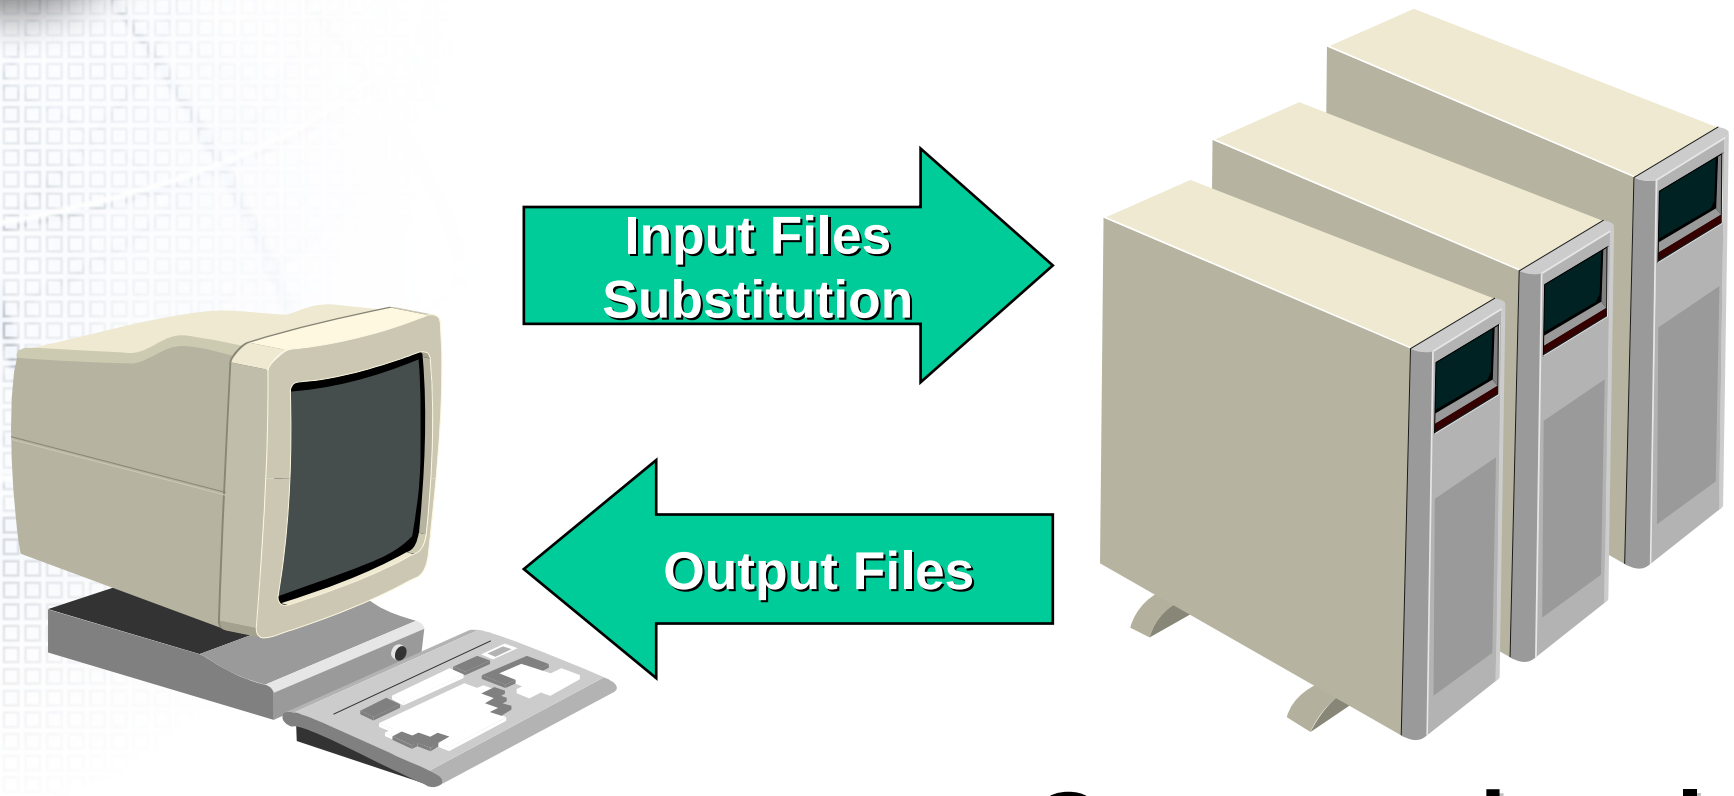

Root Machine

Computational Nodes

How does a user develop an application using EnFuzion?

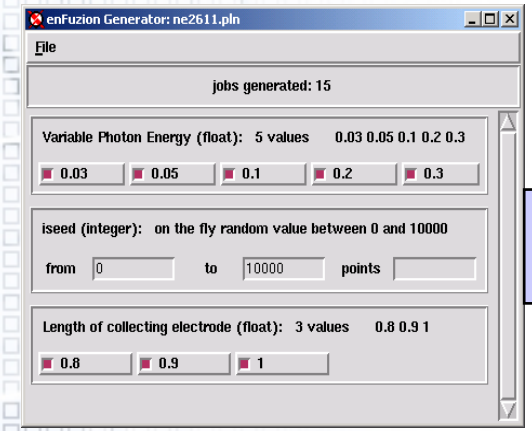

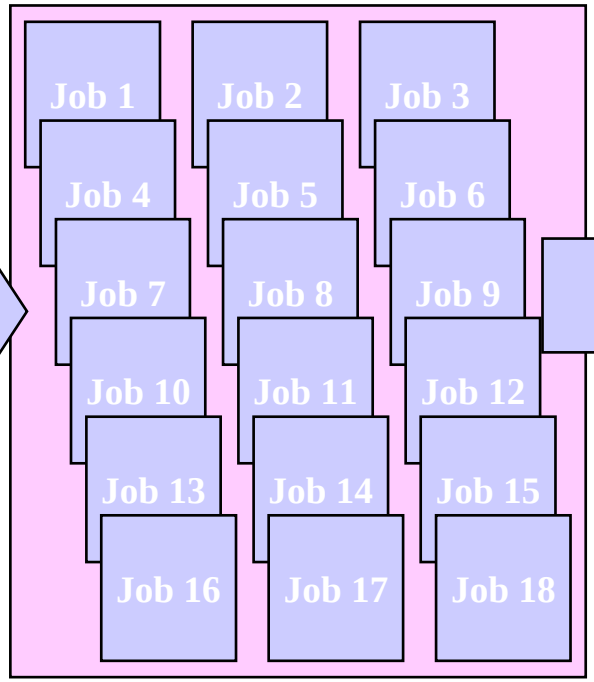

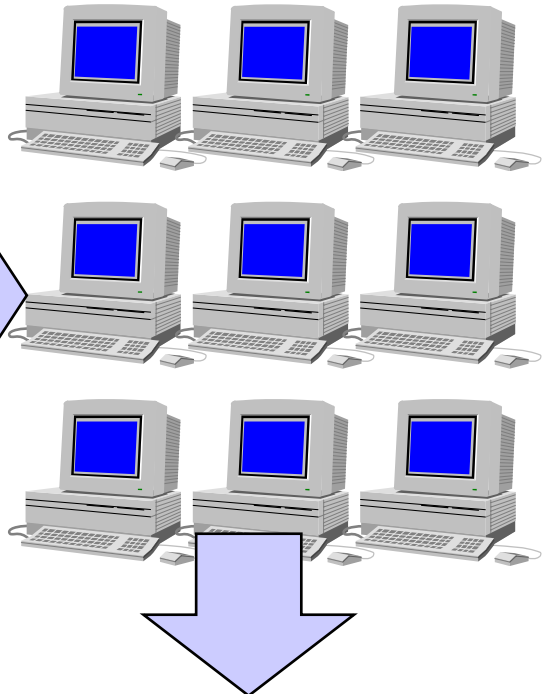

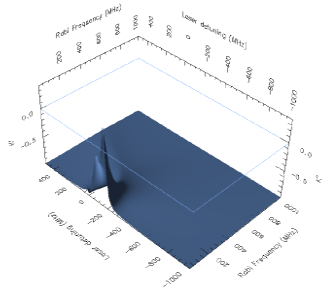

K' (Rabi Frequency, Laser Detuning)
Rubidium I=5/2 isotope : Zero Doppler width

Description of Parameters PLAN FILE

Plan file for HLmeteorite

```
parameter input_seed integer random from 1 to 1000000 points 2000;

task main
    copy p.x node:p.x
    copy projectile.input.sub node:projectile.input.sub
    copy projectile.x node:projectile.x
    substitute projectile.input.sub projectile.input
    node:execute ./p.x
    copy node:distance.out distance.out.$input_seed
    copy node:temperature.out temperature.out.$input_seed
    copy node:error.out error.out.$input_seed
    copy node:disk.out disk.out.$input_seed
    copy node:density.out density.out.$input_seed
    copy node:initial_particle.out initial_particle.out.$input_seed
    copy node:ejected_particle.out ejected_particle.out.$input_seed
    copy node:stopped_particle.out stopped_particle.out.$input_seed
endtask
```

How is the experiment coordinated?

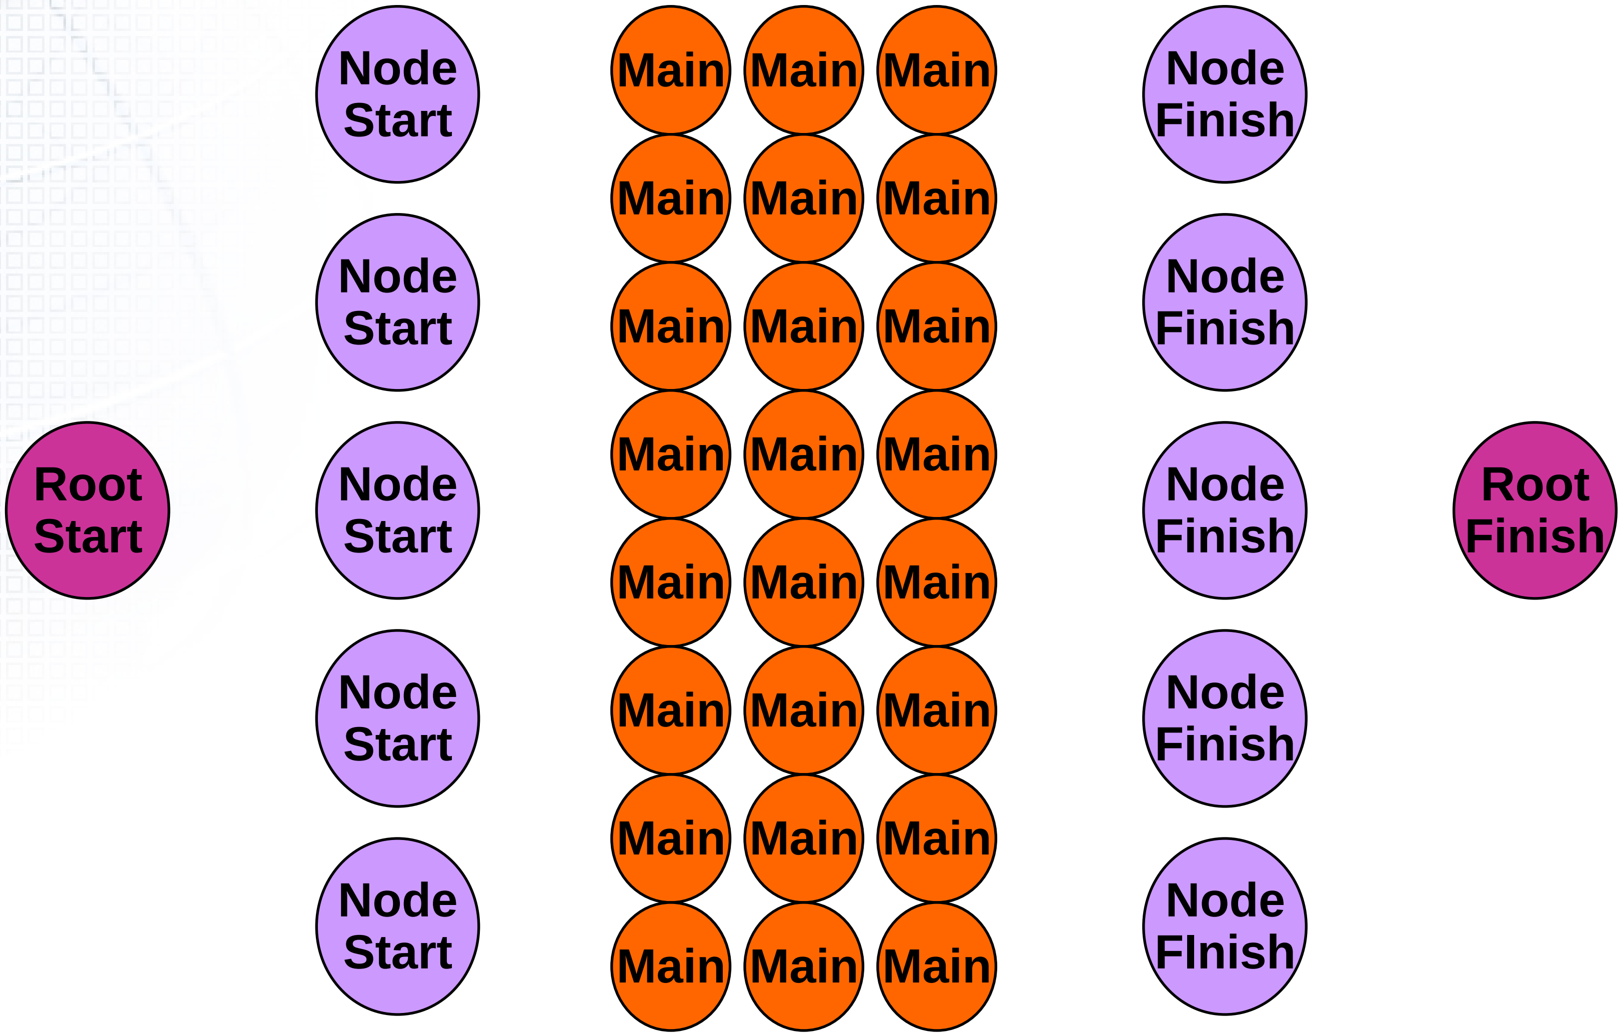

EnFuzion demonstration