

# Application Scaling Techniques on HPC Clusters

Mike Mikailov, Nicholas Petrick, Fu-Jyh Luo, Stuart Barkley, Nadia Guimont

Center for Devices and Radiological Health (CDRH), Office of Science and Engineering Laboratories (OSEL), Division of Imaging Diagnostics and Software Reliability (DIDSR)

## Introduction

**Modeling and Simulation** Modeling and Simulation (M&S) are used extensively by FDA scientists<sup>1</sup>. M&S applications can overwhelm even HPC clusters. An FDA simulation project would require 28 years of computations per product on a single workstation. Applying our techniques reduces this run time to only seven days.

**Bioinformatics** FDA's Precision Medicine Initiative relies heavily on the next-generation sequencing technology which is associated with lengthy processing of exponentially growing sequence databases. A BLAST+ job consumes more than 22 days of computations. Applying the techniques reduces this run time to only few hours.

**Deep Neural Networks Testing** run time of 228 reduced to 3 weeks.

**Big Data Analysis** Interactome project, 1,800-year reduced to 6 months.

## Materials & Methods

### Scaling simulations:

- Partition simulation loops to delegate iterations to independent tasks.
- Adapt L'Ecuyer RngStream<sup>2</sup> technique to provide quality random numbers across all independent tasks.
- Use array job facility of the job schedulers to launch the tasks on all available cores and save all partial results.
- Merge all the partial results and summarize.

### Scaling Bioinformatics:

- Partition queries and reference database into M and N subsets. Generate M x N unique combinations of the subsets.
- Use array job to launch M x N tasks, each task processing a unique combination of the subsets and save all partial results.
- Merge all the partial results and summarize.

### Scaling Deep Neural Networks Testing:

- Create a look-up table, mapping contiguous intervals of sub-image IDs (total of L) to whole-slide image IDs.
- Use array job to launch L tasks, each task retrieving (using the look-up table) and processing a sub-image for testing and saving results.

### Big Data Analysis:

- Split big data into P smaller subsets.
- Use array job to launch P tasks, each processing a subset and saving partial results.
- Merge all the partial results.

\* **Note:** Number of array job tasks can exceed the max capacity.

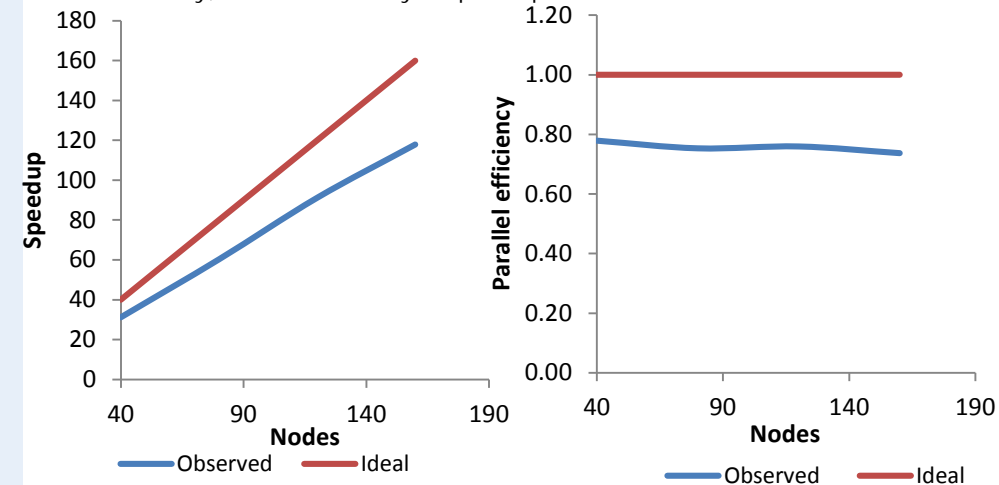
The challenge comes not from how big and fast the cluster can be, but how the HPC applications can be **scaled** to take advantage of all the resources in parallel and scalable manner.



Scan the QR code to access additional information online

## Results

**Speedup and parallel efficiency example using BlastX** . Speedup increases sub-linearly; Parallel efficiency = speed up / number of nodes



## Conclusions

- The scaling techniques presented here are already in use by FDA scientists.
- The techniques enable reduction of the data subset processed by each job task to a size that fits into the memory of the computing nodes where computations are performed. The resulting reduction in disk I/O produces excellent results, enabling substantial drops in run times.
- The described methods use only open source code, adds no hardware cost.

## Acknowledgments

Arifa S. Khan, PhD, CBER; Junshan Qiu, PhD, CDER; Weijie Chen, PhD, CDRH  
Weizhe Li, PhD, CDRH, and many others at FDA.

## References

1. <https://www.nuventra.com/resources/blog/simulation-transforms-regulatory-pathways/>
2. <http://simul.iro.umontreal.ca/>