Performance Prediction of Stencil Applications based on Machine Learning

Víctor Martínez,

and Philippe O. A. Navaux

Federal University of Rio Grande do Sul – BR

Outline

- **Motivation**
  - Multicore Architectures.
  - Stencil Models.
- **Machine Learning Methodology.**
  - Input/Output vectors.
  - Hardware Counters Behavior.
  - ML Model.
- **Experiments.**
  - Testbed
  - Training and Validation sets.
- **Results.**
- **Conclusions.**

Motivation
Multicore and Many core Architectures

- Complexity:
  - Cache levels.
  - Cores.
  - Sockets.
- Optimization:
  - Non-uniform memory access.
  - Vectorization.
  - Compiler flags.
  - Memory policies.

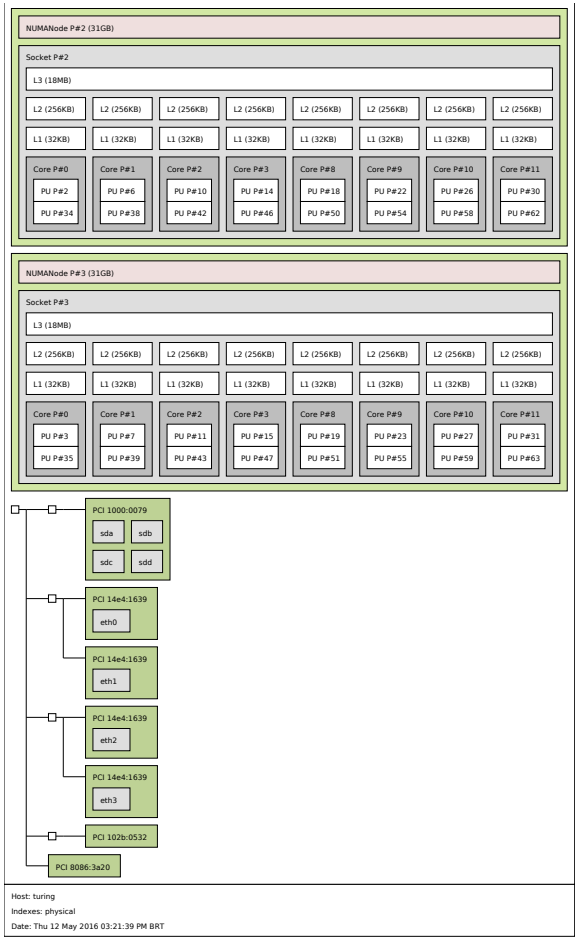

Motivation
Stencil Models

- 7-point Jacobi
  - Heat transference.

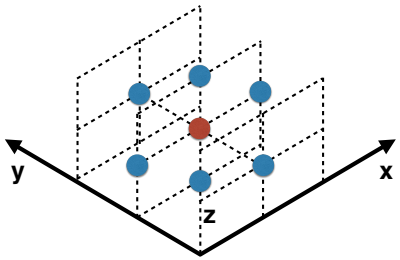

- Seismic
  - Velocity and Stress.

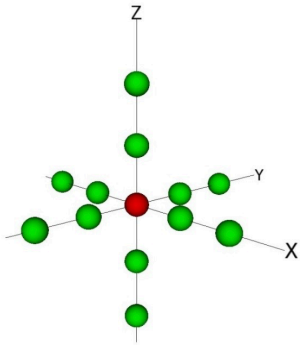

- Acoustic wave
  - Reverse time migration.

$$\frac{1}{V^2}\frac{\partial^2 p}{\partial t^2} = \nabla^2 p$$

$$\frac{1}{V^2}\frac{\partial^2 p}{\partial t^2} = \nabla^2 p - \frac{\nabla \rho}{\rho}\nabla p$$

Machine Learning Methodology
Hardware Counters Behavior

- **Input Vector**
  - Policy Scheduling.
  - Chunk Size.
  - Executed Threads.

- **Hardware Counters**
  - Last Level Cache Misses (L2 & L3)
  - Total Cycles.
  - Data Translation Lookaside Buffer Misses.

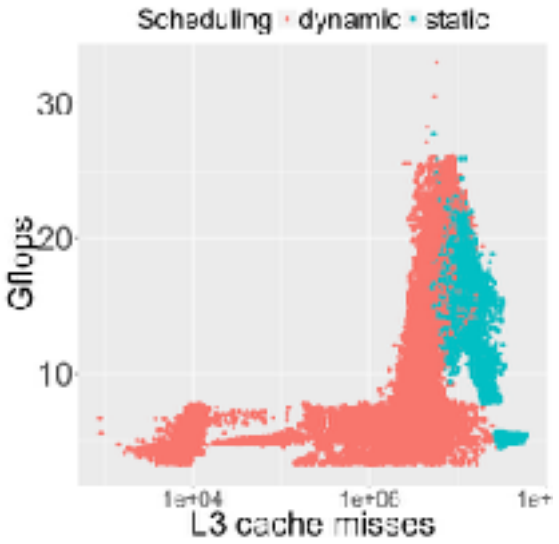

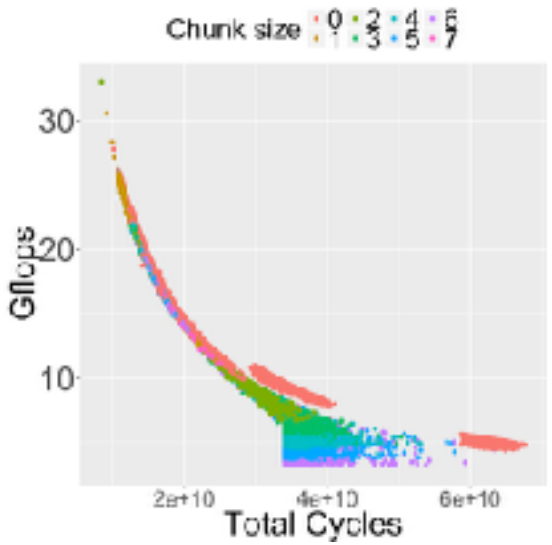

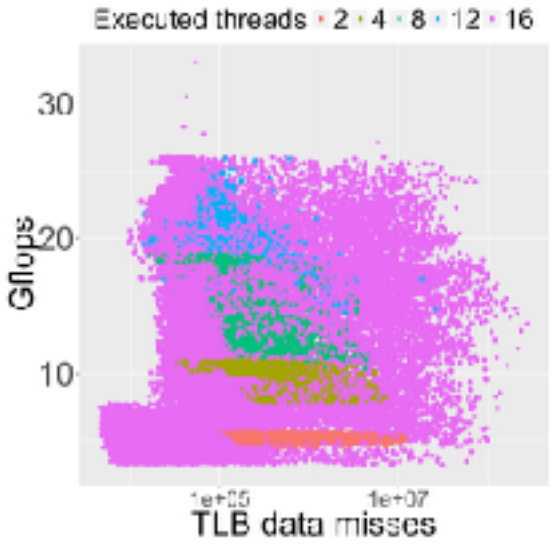

Machine Learning Methodology
Hardware Counters Behavior

- **Input Vector**
  - Policy Scheduling.
  - Chunk Size.
  - Executed Threads.

- **Hardware Counters**
  - Last Level Cache Misses (L2 & L3)
  - Total Cycles.
  - Data Translation Lookaside Buffer Misses.

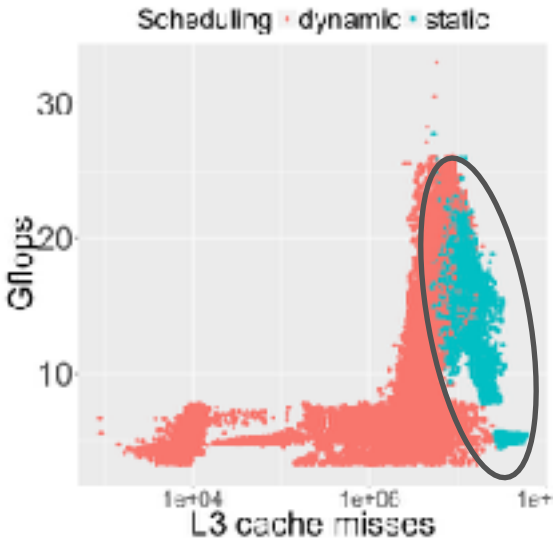

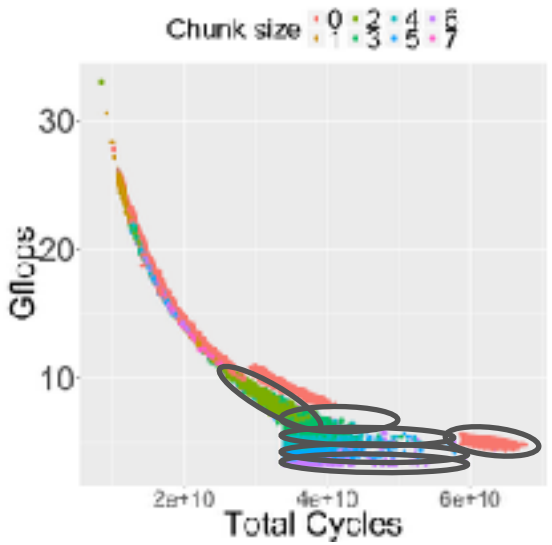

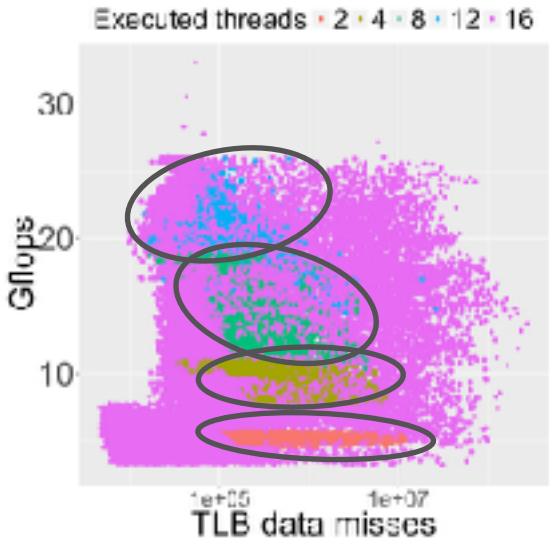

Machine Learning Methodology
Hardware Counters Behavior

- **Input Vector**
  - Policy Scheduling.
  - Chunk Size.
  - Executed Threads.

- **Hardware Counters**
  - Last Level Cache Misses (L2 & L3)
  - Total Cycles.
  - Data Translation Lookaside Buffer Misses.

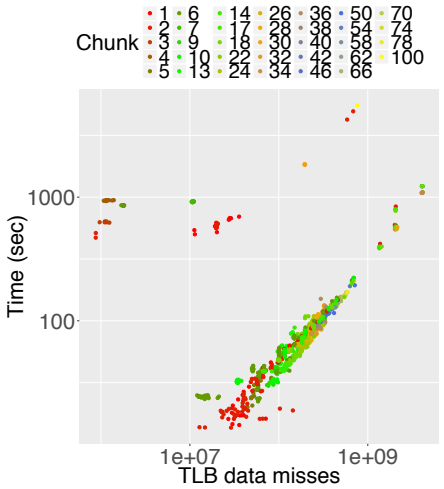

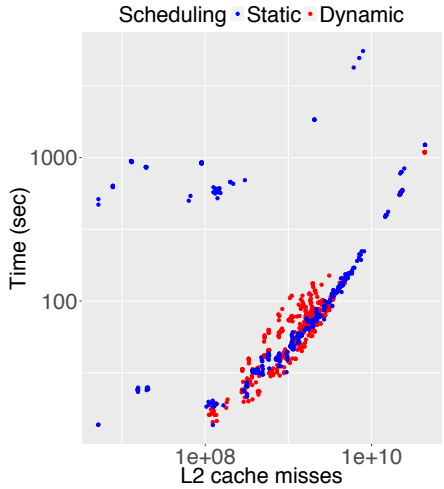

Machine Learning Methodology
Supported Vector Machines Model

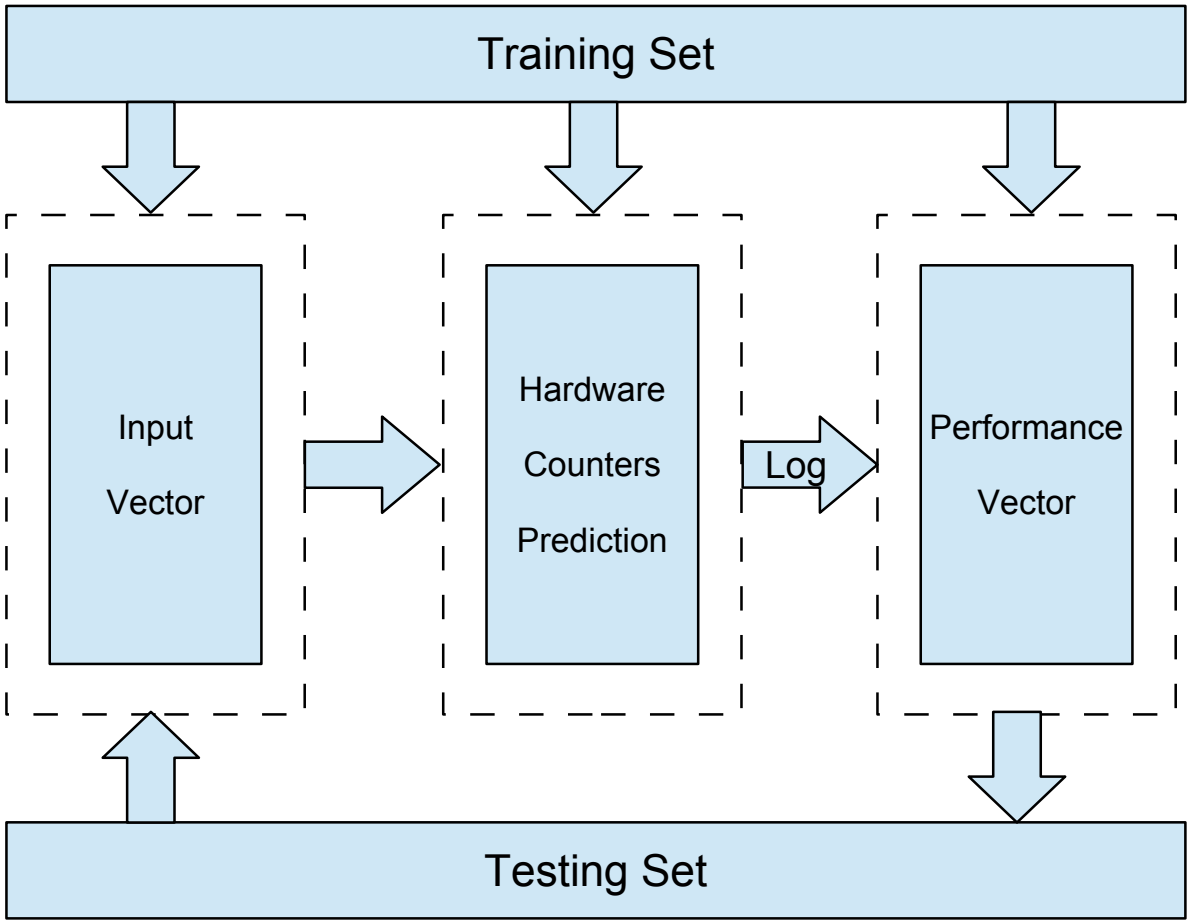

Experiments
Testbed

	Node 1	Node 2	Manycore
Processor	Xeon E5-2650	Xeon E5-4650	Xeon Phi 750
Clock (GHz)	2.0	2.7	1.40
Cores	8	8	68
Sockets	2	4	1
Threads	16	32	272
LL cache size (MB)	20 (L3)	20 (L3)	34 (L2)

Experiments
Optimization, Training and Validation Sets

Optimization	Parameters	Total of Configurations		
		Node 1	Node 2	Manycore
Number of thtreads	1	8	12	272
Scheduling policy	1	2	2	2
Chunk size	1	32	32	272
Total for Naive	**3**	**512**	**768**	**147,968**

Stencil	Set	Node 1	Node 2	Manycore
7-point Jacobi	Training	44	38	-
	Testing	11	10	-
	Total	*55*	*48*	-
Seismic wave	Training	211	237	-
	Testing	53	60	-
	Total	*264*	*297*	-
Acoustic	Training	-	-	808
	Testing	-	-	203
	Total	-	-	*1,011*

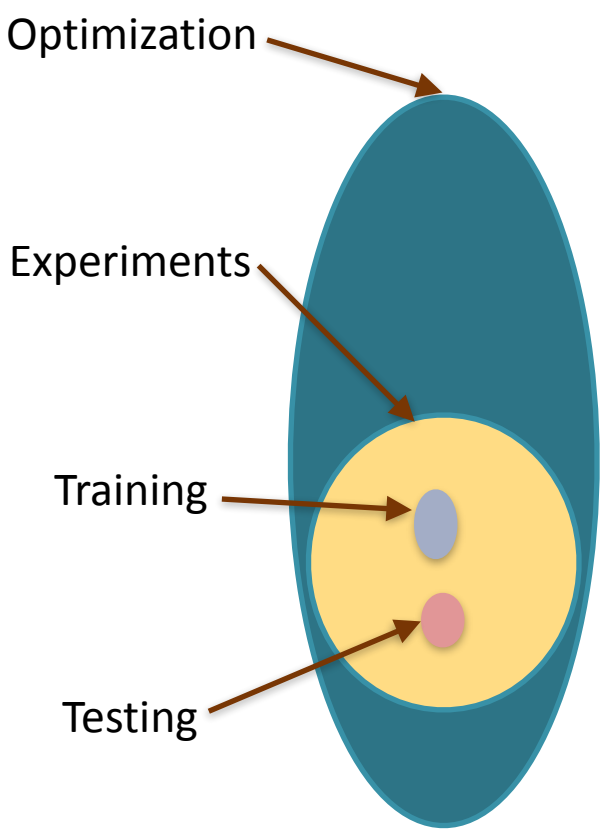

Regression Model Accuracy

- Testing set

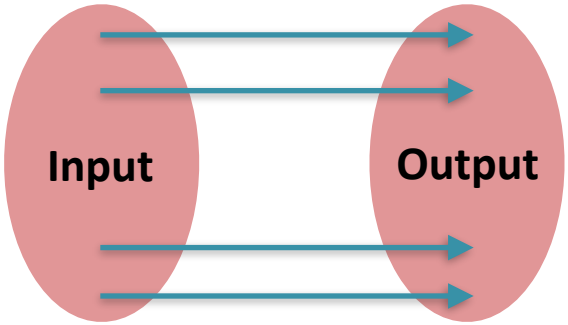

Stencil		Node 1	Node 2	Manycore
7-point Jacobi	*RMSE*	0.66	2.51	-
	R-Square	0.98	0.86	-
Seismic wave	*RMSE*	*13.53*	*212.28*	-
	R-Square	*0.99*	*0.62*	-
Acoustic	*RMSE*	-	-	154.04
	R-Square	-	-	0.94

Conclusions

- We introduced a general predictive performance modeling strategy for geophysical numerical kernel on multi and many-core architectures:
  - Performance of stencil can be predicted with a high accuracy (up to 99%).
  - Model can be integrated into an auto-tuning framework to find the best performance.

- Next step: a model based on unsupervised ML algorithms.

Performance Prediction of Stencil Applications based on Machine Learning

Thanks!

Acknowledgment

Research has received funding from the EU H2020 Programme and from MCTI/RNP-Brazil under the HPC4E Project, grant agreement n° 689772.

Victor Martinez, and Philippe O. A. Navaux

Federal University of Rio Grande do Sul – BR