NTT DATA Mathematical Systems Inc.

Trusted Global Innovator
NTT DATA Group

ParaNUOPT: Parallelization of NUOPT by using UG on cloud computing platform

June 25, 2019

NTT DATA Mathematical Systems Inc. *Yasumi Ishibashi
Zuse Institute Berlin Yuji Shinano

© 2018 NTT DATA Mathematical Systems Inc.

Contents

1. Overview
2. UG and ParaNUOPT
3. HPC on Cloud
4. Computational Experiment
5. Conclusion

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

Contents

1. Overview
2. UG and ParaNUOPT
3. HPC on Cloud
4. Computational Experiment
5. Conclusion

© 2018 NTT DATA Mathematical Systems Inc.

1. Overview

■ Motivation

- Easily get machine resources by cloud computing services
- Easily get distributed B&B solver parallelized by UG
- Easily run distributed B&B solver on a cloud HPC environment

■ Question

- Does UG work efficiently on a cloud HPC environment such as AWS ?
  - Confirm the speedup of ParaNUOPT on AWS
  - Confirm the impact of network performance on speedup

■ Answer

- Under investigation. But,
- 4 times speedup with 8 compute nodes in some problems on AWS
- Super-linear speedup on AWS, not a supercomputer

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA
NTT DATA Mathematical Systems Inc.

Contents

1. Overview
2. UG and ParaNUOPT
3. HPC on Cloud
4. Computational Experiment
5. Conclusion

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

2-1. What is UG ?

- UG (http://ug.zib.de)
  - Framework to parallelize branch and bound solver (Shinano 2010)
- Already parallelize many solvers
  - SCIP (Shinano 2010) -> ParaSCIP
  - Xpress (Shinano 2016) -> ParaXpress
  - PIPS-SBB (Munguía 2017)
  - and so on
- Computational results on a supercomputer are reported
- ParaSCIP and ParaXpress solved open instances from MIPLIB2017
- Details of UG will be given in the following talk

Tuesday, 14:30-16:00 - L249
Software for large-scale optimization II
Configuring ParaXpress to Enhance its Heuristic Performance
Yuji Shinano, Timo Berthold, Lluis-Miquel Munguia

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA
NTT DATA Mathematical Systems Inc.

2-2. Mechanism of UG (1/2)

- Supervisor-Worker coordination mechanism with subtree parallelism (Ralphs+ 2016)
  - Worker solves unexplored nodes of search tree
  - Supervisor coordinates workload, communicates with each workers
    - ✓ Communication is basically one to one

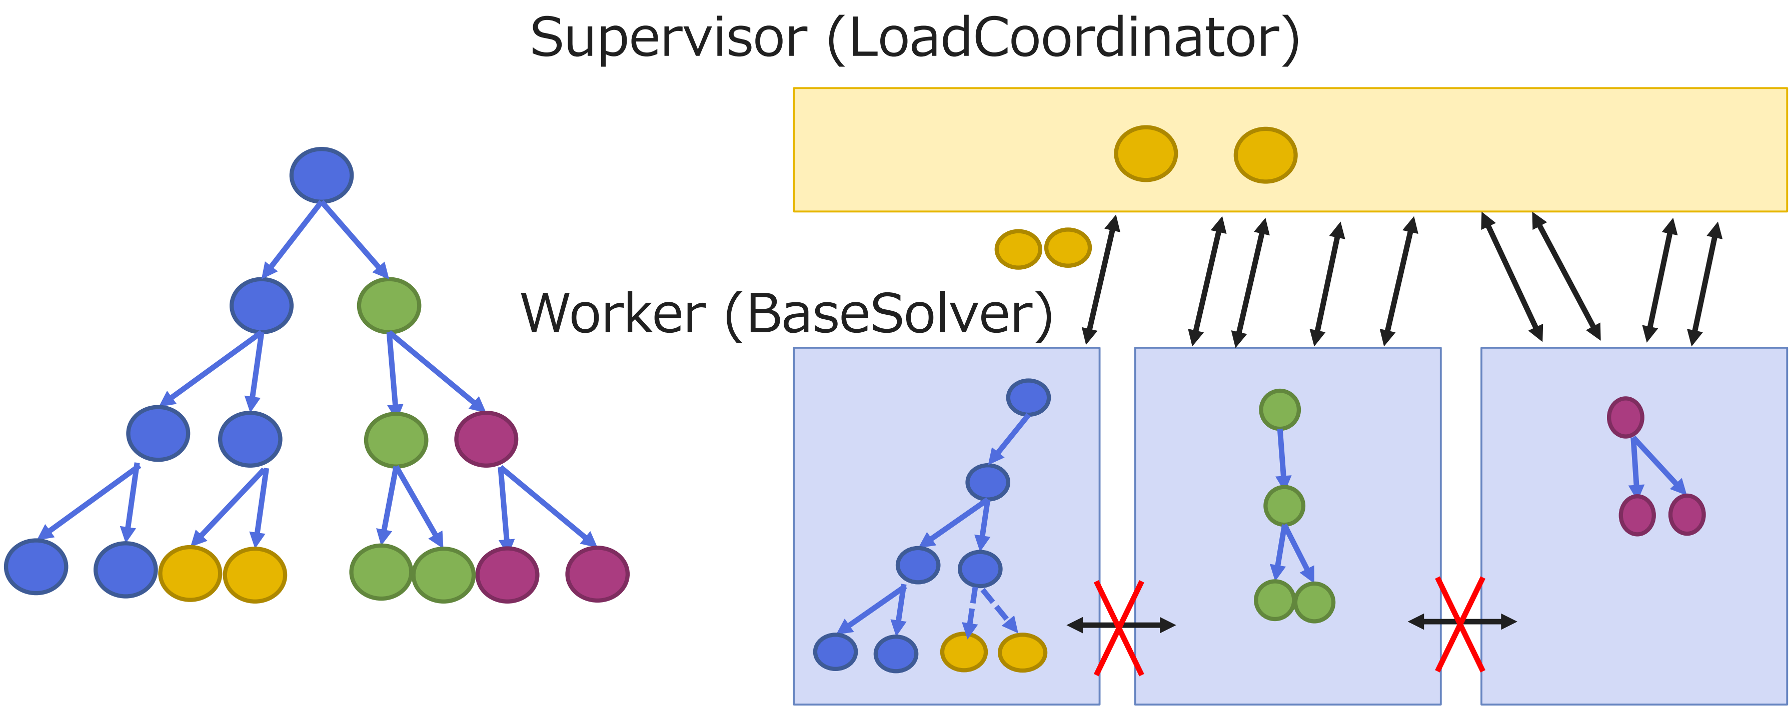

2-2. Mechanism of UG (2/2)

- Supervisor-Worker coordination mechanism with subtree parallelism (Ralphs+ 2016)
  - Worker solves unexplored nodes of search tree
  - Supervisor coordinates workload, communicates with each workers
    - ✓ Communication is basically one to one

Message	Frequency	Size
Worker status	regularly	small (60 bytes)
Primal bound	when needed	double (8 bytes)
Incumbent solution	when needed	depend on the number of vars.
Unexplored node (task)	when needed	depend on the number of vars.

- Worker status is the only message communicated regularly
  - Expect that network performance does not affect the speedup of UG

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

2-3. What is NUOPT and ParaNUOPT ?

- NUOPT
  - Commercial mathematical optimization solver
- ParaNUOPT
  - NUOPT parallelized by UG for research

Numerical Optimizer

- ParaNUOPT first solved the following open instances from MIPLIB2017
  - gen-ip016 (in 71498 seconds, on PC cluster with 19 cores)
  - rococoC11-010100 (in 32368 seconds, on PC cluster with 9 cores)
- Does UG work also efficiently on a cloud HPC cluster ?
  - If this answer is "Yes",
    anyone who does not have a supercomputer can use it

Contents

1. Overview

2. UG and ParaNUOPT

3. HPC on Cloud

4. Computational Experiment

5. Conclusion

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

3-1. What is Cloud Computing Service?

- Cloud computing service
  - Provide computing resources/services via the internet

- Major cloud venders
  - Amazon (AWS : https://aws.amazon.com)
    - ✓ market leader
  - Microsoft (Azure : https://azure.microsoft.com)
  - Google (GCP : https://cloud.google.com)

- Cloud venders provide various virtual machines
  - Let's see the virtual machines of AWS for HPC

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA
NTT DATA Mathematical Systems Inc.

3-2. Virtual Machines for HPC (AWS)

- Purpose : Computing
- CPU : Intel Xeon Platinum 3.0 GHz
- Network Bandwidth : 10 Gbps
- Price (Tokyo) : Spot price gives about 70% discount

VM (EC2 Instance)	CPU	Memory (GiB)	Network Bandwidth (Gbps)	On-demand Price ($/hour)		Spot Price ($/hour)
c5.xlarge	2	8	10	0.214		0.0736
c5.2xlarge	4	16	10	0.428		0.1364
c5.4xlarge	8	32	10	0.856		0.2661
c5.9xlarge	18	72	10	1.926	-70%	0.5998
c5.18xlarge	36	144	10	3.852	→	1.1976

© 2018 NTT DATA Mathematical Systems Inc.

3-3. How to create HPC on AWS

- ParallelCluster (http://aws-parallelcluster.readthedocs.io)
  - Create flexible HPC cluster with a single command
    - ✓ When there are no jobs, compute nodes will be shutdown
      - ➢ save money

- Other cloud vendors also provide tools to create HPC cluster
  - What cloud vendors provide the best HPC cluster ?

© 2018 NTT DATA Mathematical Systems Inc.

3-4. HPC comparison on the cloud

- Mohammad and Timur compared cloud HPC cluster by using High Performance LINPACK (2018); Microsoft Azure

Vendor	VM	Cores	Freq. (GHz)	RAM (Gb)	Network.
Azure	H16r	18	3.2	112	Infiniband 54Gbps
AWS	c4.8xlarge	18	2.9	60	Ethernet 10Gbps

The higher Rmax is better

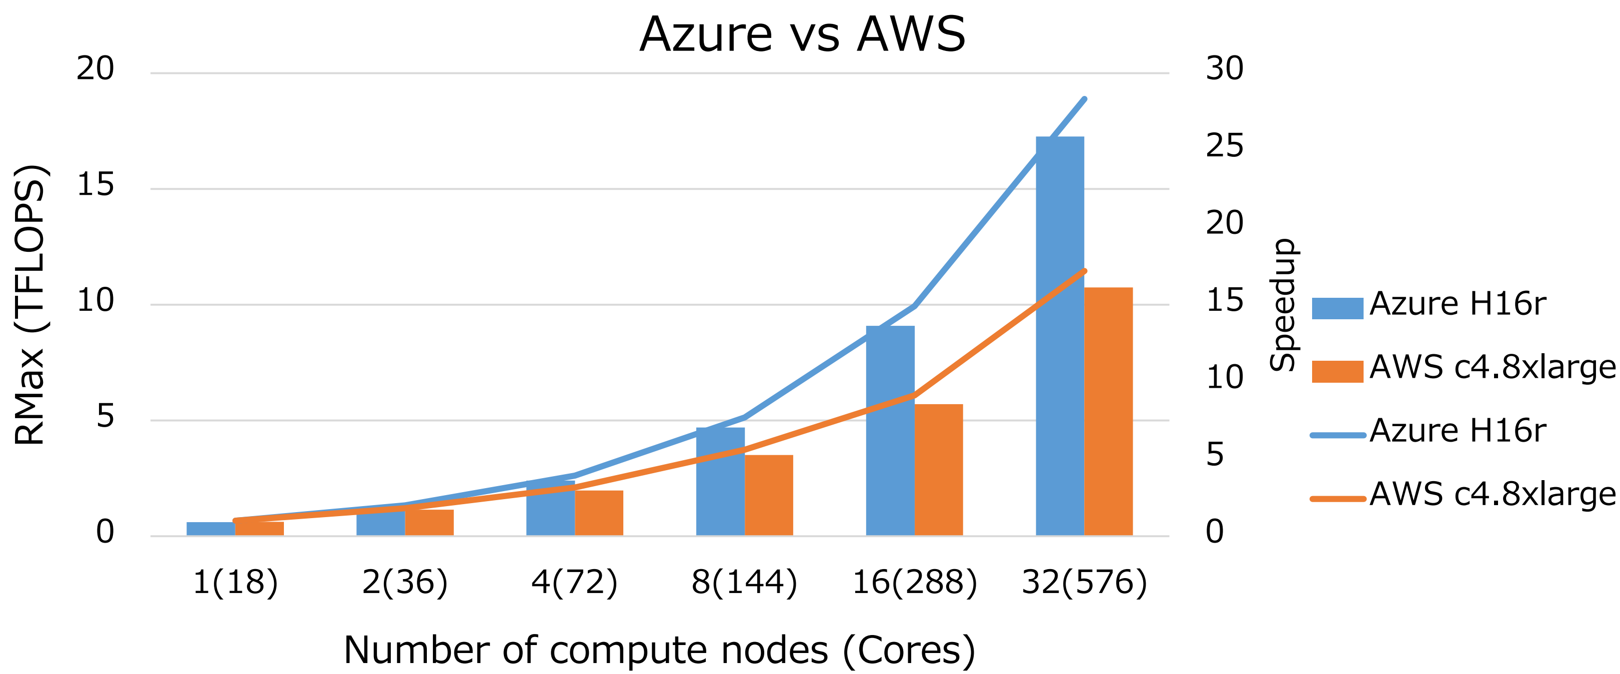

© 2018 NTT DATA Mathematical Systems Inc.

3-5. Question (1/2)

- Communication-intensive application will not be accelerated on AWS
  - ✓ Because AWS does not provide Infiniband.

- HPC cluster on Azure is the best for High Performance LINPACK

- **Question**
  - Does UG work efficiently on a cloud HPC environment such as AWS ?
  - When we use UG, what cloud vendors is the best in terms of cost and performance ?

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA
NTT DATA Mathematical Systems Inc.

3-5. Question (2/2)

- Communication-intensive application will not be accelerated on AWS
  - ✓ Because AWS does not provide Infiniband.

- HPC cluster on Azure is the best for High Performance LINPACK

- **Question**
  - Does UG work efficiently on a cloud HPC environment such as AWS ?
    - Confirm the speedup of ParaNUOPT on AWS
    - Confirm the impact of network performance on speedup
  - When we use UG, what cloud vendors is the best in terms of cost and performance ?

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA
NTT DATA Mathematical Systems Inc.

3-6. Performance of Virtual Machine on AWS

NB = Network Bandwidth

VM (EC2)	Freq. (GHz)	Cores	Memory (GiB)	NB (Single)	NB (Total)	Network Latency
c4.8xlarge	2.9	18	60	?	10 Gbps	?
c5.18xlarge	3.0	limit to 18	144	?	25 Gbps	?

Memory Bandwidth	Copy (GB/s)	Scale (GB/s)	Add (GB/s)	Triad (GB/s)
c4.8xlarge	58.16	56.02	61.04	62.94
c5.18xlarge	129.45	120.94	135.66	132.75

by STREAM (version 5.10)

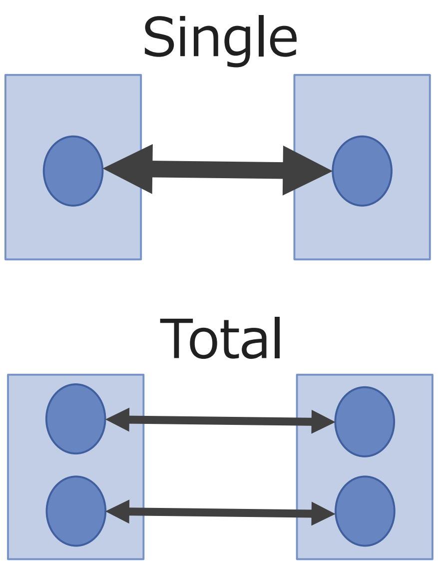

3-7. Measure Network Bandwidth/Latency (1/2)

- OpenMPI 4.0.1
- OSU MICRO BENCHMARKS 5.6.1
- Logarithmic scale

Single 1-to-1

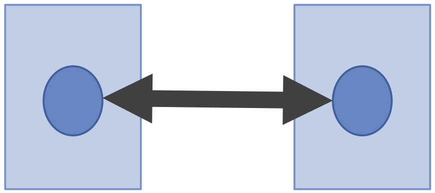

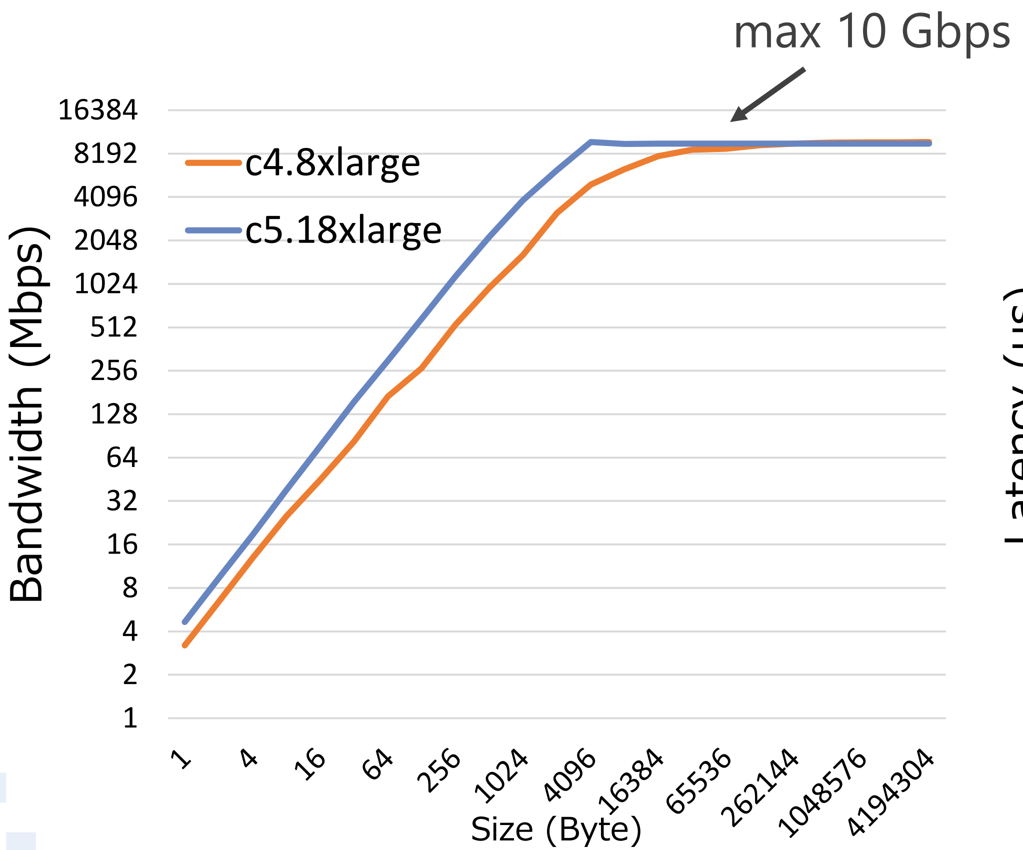

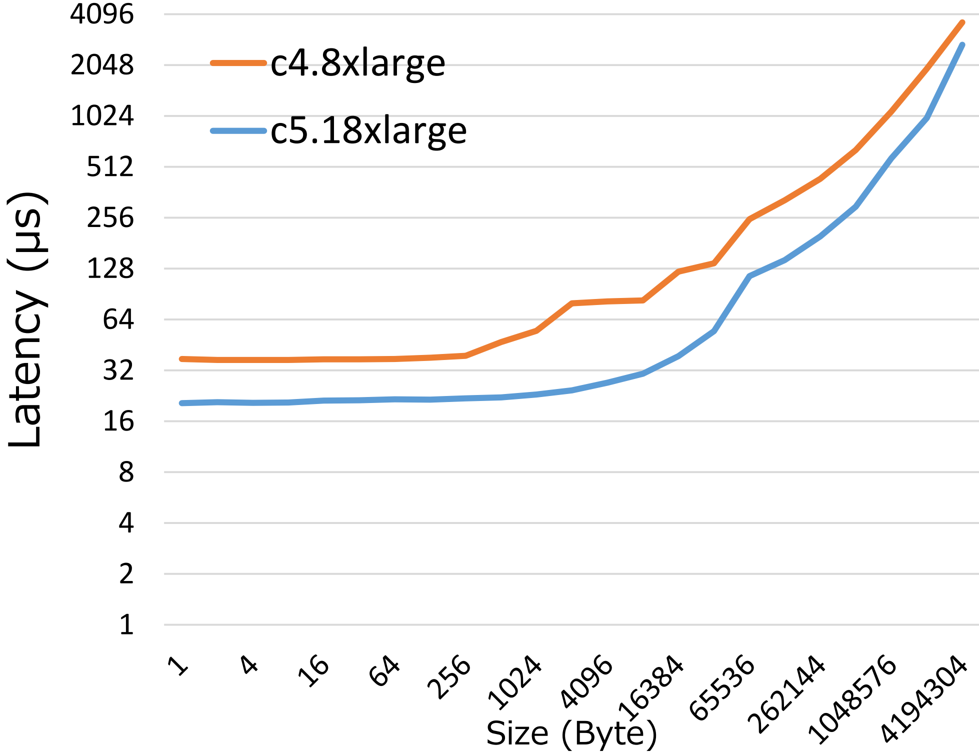

© 2018 NTT DATA Mathematical Systems Inc.

3-7. Measure Network Bandwidth/Latency (2/2)

- Latency of c4.8xlarge is 1.7 times slower than c5.18xlarge
- Bandwidth of c4.8xlarge is about 0.6 times smaller than c5.18xlarge for messages of 8192 bytes or less

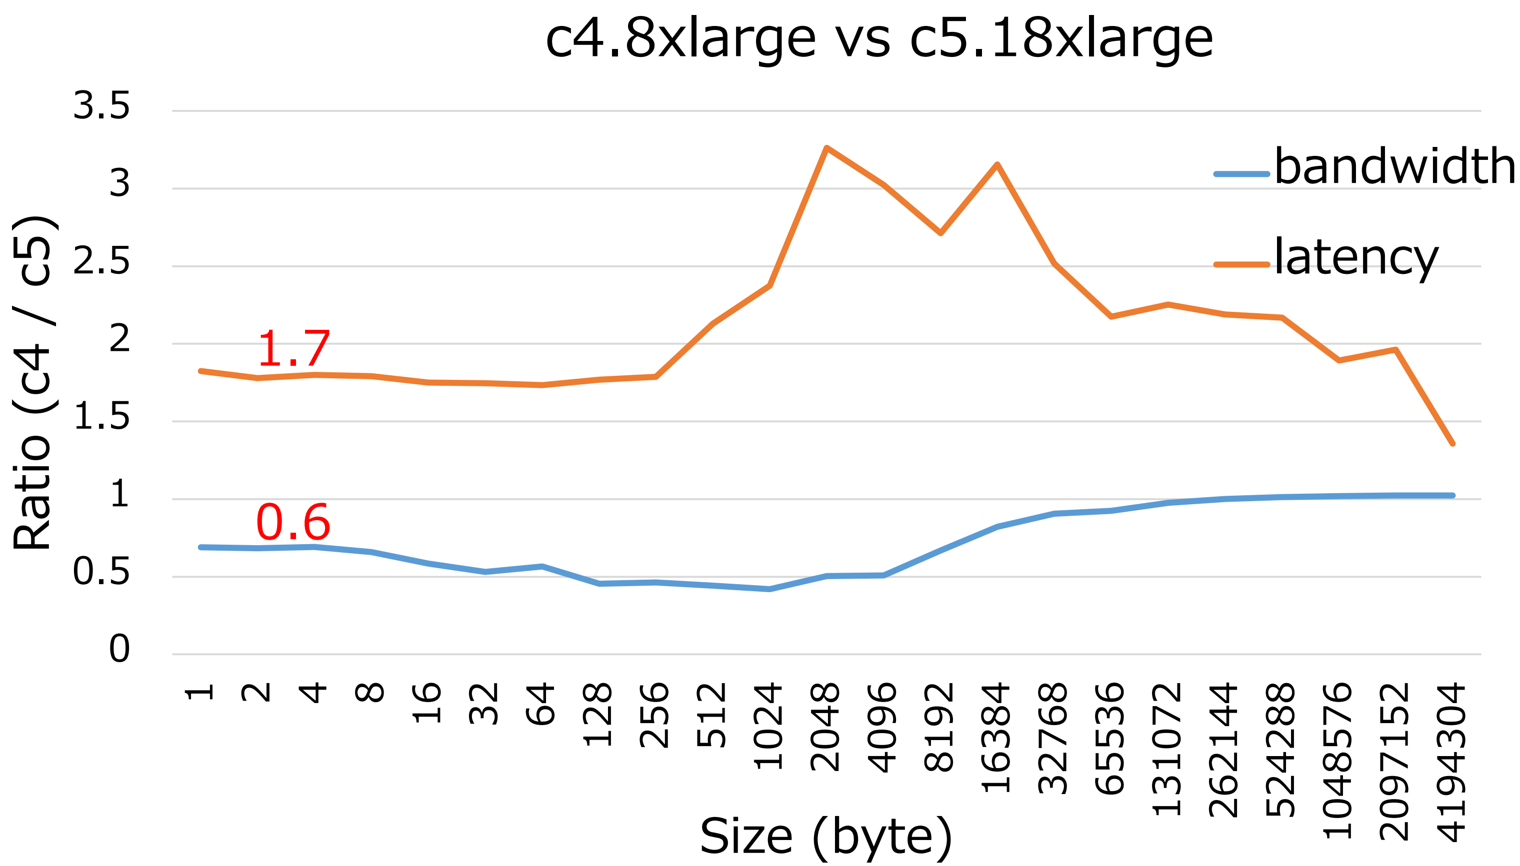

3-8. Performance of Virtual Machine on AWS

NB = Network Bandwidth

VM (EC2)	Freq. (GHz)	Cores	Memory (GiB)	NB (Single)	NB (Total)	Ratio of Network Latency
c4.8xlarge	2.9	18	60	max 10 Gbps	10 Gbps	slower 1.7x
c5.18xlarge	3.0	limit to 18	144	max 10 Gbps	25 Gbps	1.0

Memory Bandwidth	Copy (GB/s)	Scale (GB/s)	Add (GB/s)	Triad (GB/s)
c4.8xlarge	58.16	56.02	61.04	62.94
c5.18xlarge	129.45	120.94	135.66	132.75

by STREAM (version 5.10)

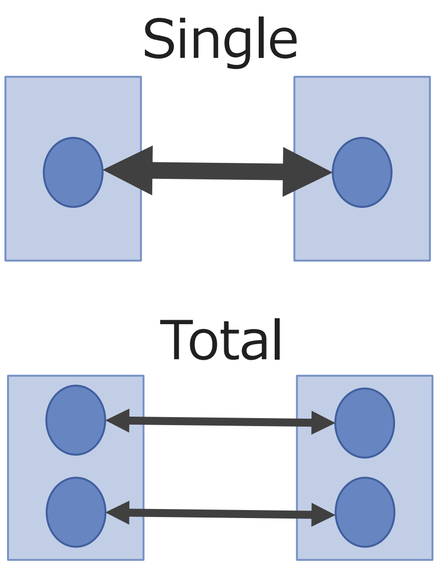

Contents

1. Overview
2. UG and ParaNUOPT
3. HPC on Cloud
4. Computational Experiment
5. Conclusion

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

4-1. Configuration of Computational Experiment

- Use c4.8xlarge and c5.18xlarge
  - Confirm the speedup of ParaNUOPT on AWS
  - Confirm the impact of network performance on the speedup
- Use three problems as MIP

Problem	LIB	Variables	Constraints	Nonzeros
chr20a (*)	QAPLIB	800	441	16,440
fastxgemm-n2r6s0t2	MIPLIB2017	784	5,998	19,376
nu25-pr12	MIPLIB2017	5,868	2,313	17,712

(*) QAP is linearized by Kauffmann and Broeckx formulation (1978)

- Run ParaNUOPT 10 times to give an average value
  - ✓ Because ParaNUOPT is nondeterministic
- Turn off racing of UG so that parameters will not change dynamically

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA
NTT DATA Mathematical Systems Inc.

4-2. Result (chr20a)

- 800 variables, 441 constraints, 16,440 nonzeros

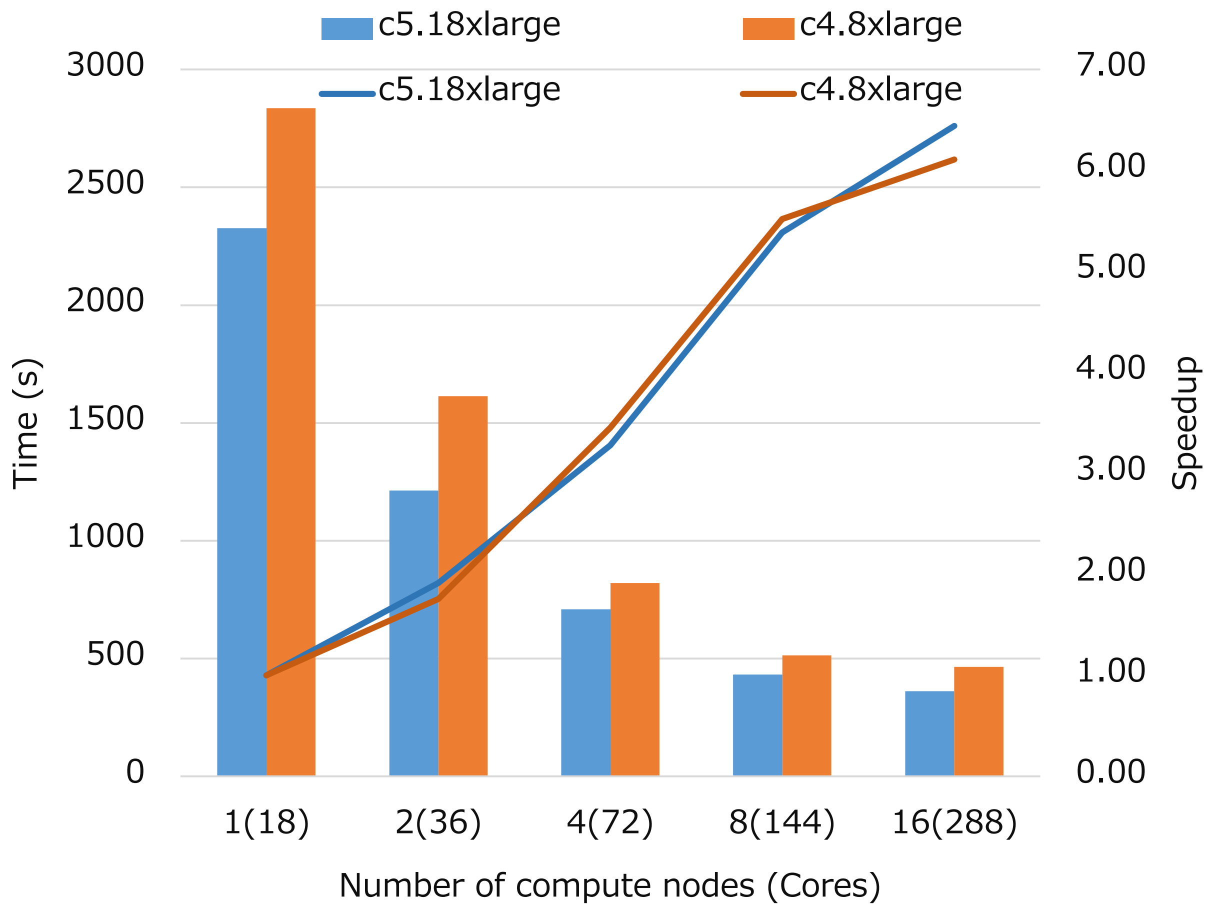

- The slight difference in speedup with 16 compute nodes, but the speedup with up to 8 compute nodes is the same
- chr20a is so easy that the speedup may saturate with 16 compute nodes

© 2018 NTT DATA Mathematical Systems Inc.

4-3. Result (fastxgemm-n2r6s0t2)

- 784 variables, 5,998 constraints, 19,376 nonzeros

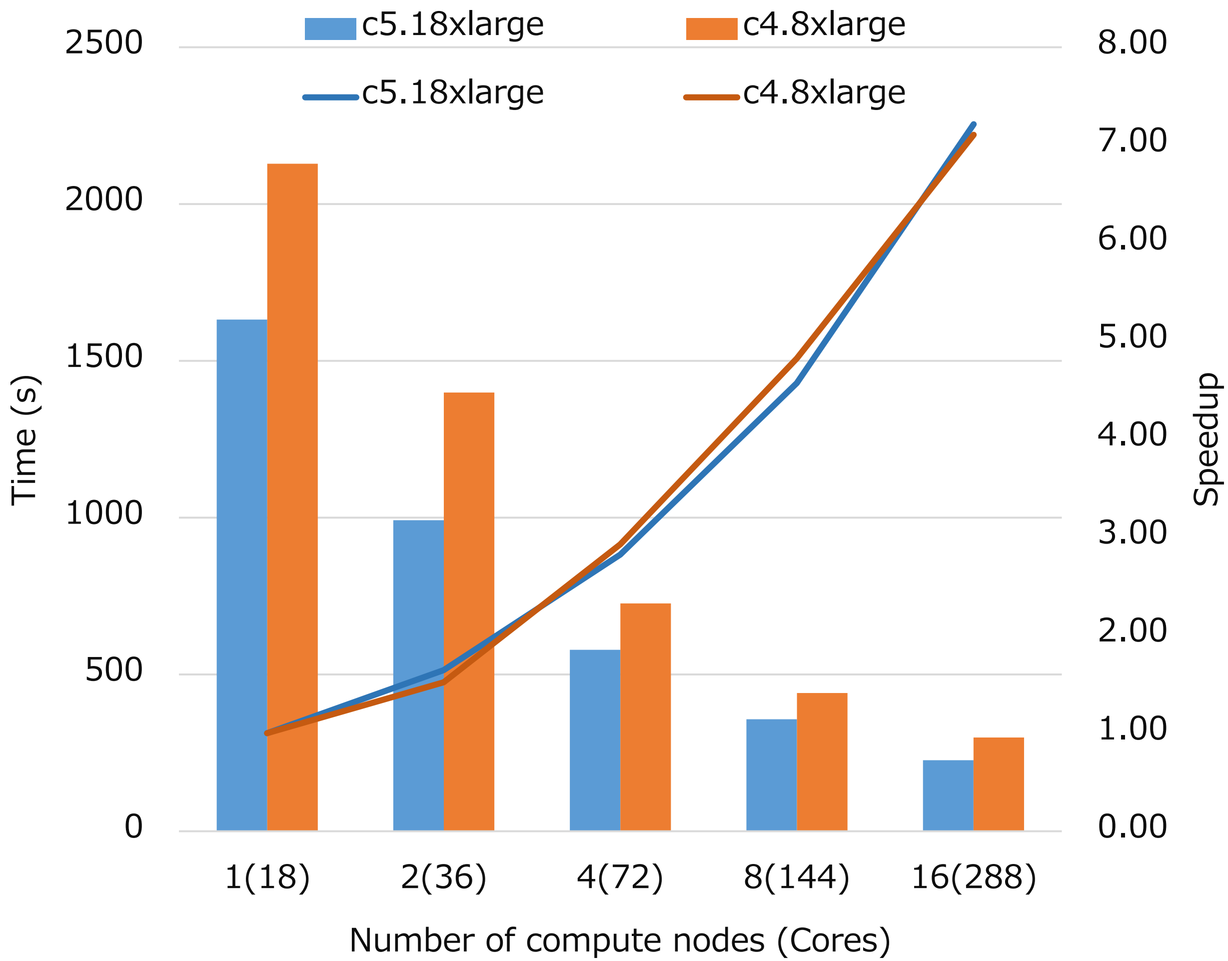

- Can not confirm the difference in speedup

© 2018 NTT DATA Mathematical Systems Inc.

4-4. Result (nu25-pr12)

- 5,868 variables, 2,313 constraints, 17,712 nonzeros

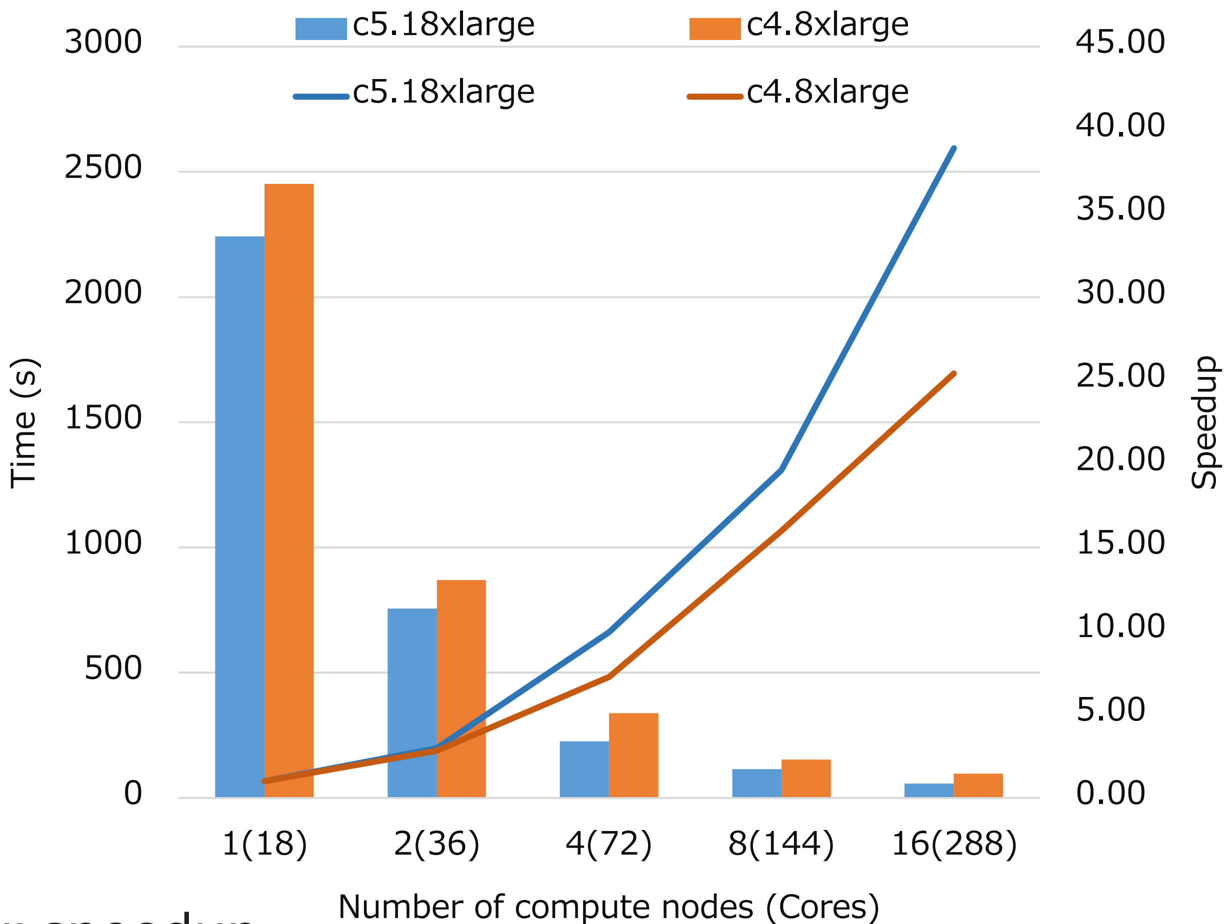

- Super-linear speedup
- The difference in speedup at more than 4 compute nodes

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

4-5. Number of Explored Nodes (nu25-pr12)

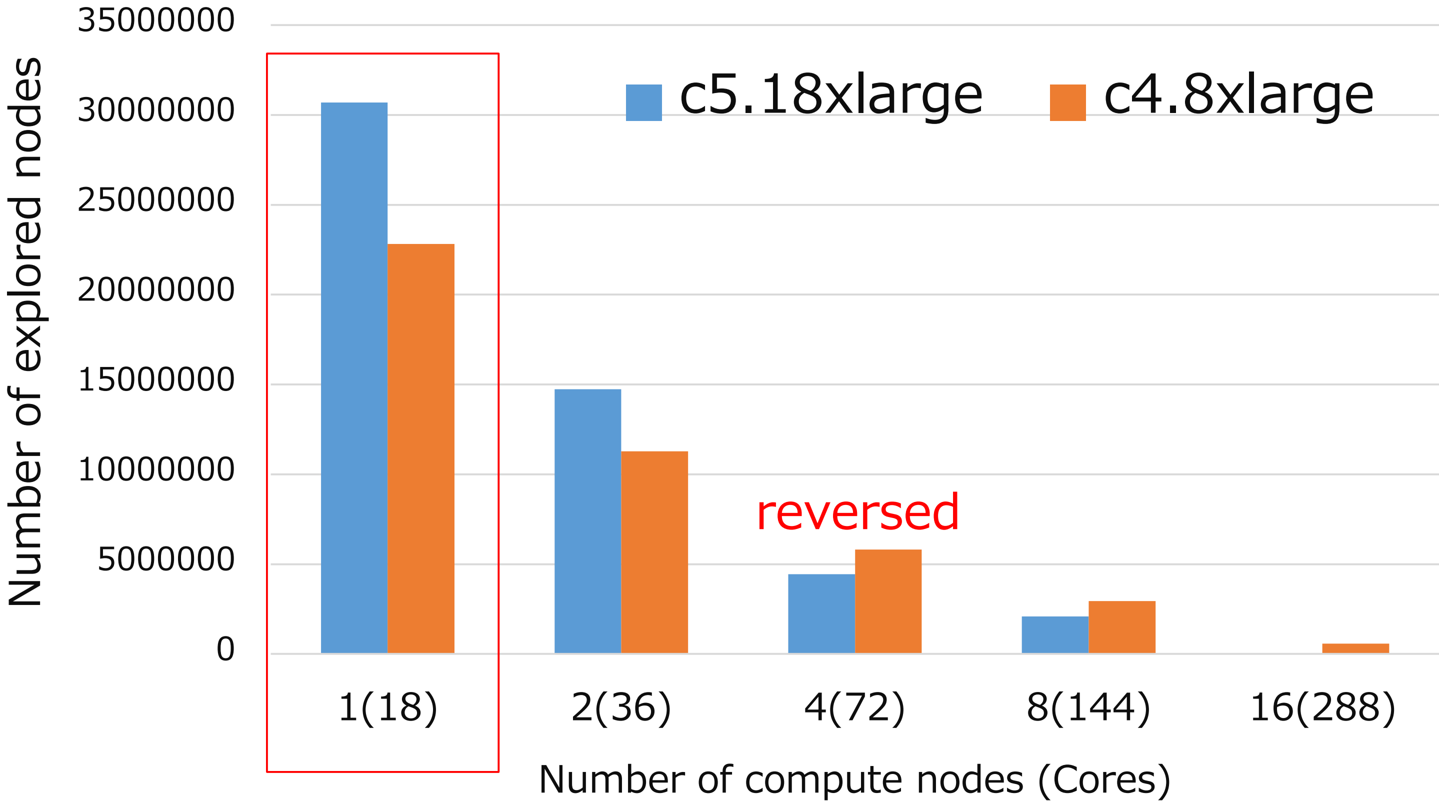

- The number of explored nodes on c5 with 1 comp. node is larger than c4
    - Base of the speedup on c5 is worse than c4
- The number of explored nodes is reversed at 4 compute nodes
    - C5 gets the larger speedup than c4

Contents

1. Overview
2. UG and ParaNUOPT
3. HPC on Cloud
4. Computational Experiment
5. Conclusion

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA Mathematical Systems Inc.

5. Conclusion

- **Question**
  - Does UG work efficiently on a cloud HPC environment such as AWS ?

- **Answer**
  - Under investigation. But,
  - 4 times speedup with 8 compute nodes in some problems on AWS
  - Super-linear speedup on AWS, not a supercomputer

- **Future work**
  - Problems used for this experiment may be too easy
    - ✓ Try large and hard problems
  - Confirm clearly the impact of network performance on the speedup
    - ✓ Compare the performance of ParaNUOPT between Ethernet and Infiniband

© 2018 NTT DATA Mathematical Systems Inc.

References

1. L. Kaufman, *et al*. “An algorithm for the quadratic assignment problem using Bender's decomposition". *European Journal of Operational Research 2.3,* 1978, pp. 207-211.

2. Y. Shinano, *et al*. “ParaSCIP: a parallel extension of SCIP". In *Competence in High Performance Computing 2010*. Springer, 2011, pp. 135-148.

3. Y. Shinano, *et al*. “A first implementation of ParaXpress: Combining internal and external parallelization to solve MIPs on supercomputers". In *International Congress on Mathematical Software.* Springer, 2016, pp. 308-316.

4. L. M. Munguía, *et al*. “Parallel PIPS-SBB: Multi-level parallelism for stochastic mixed-integer programs". *Computational Optimization and Applications*, 2017, pp. 1-27.

5. T. Ralphs, *et al*. “Parallel solvers for mixed integer linear optimization". In *Handbook of parallel constraint reasoning*, Springer, 2018, pp. 283-336.

6. M. Mohammadi, *et al*. “Comparative benchmarking of cloud computing vendors with high performance linpack". In *Proceedings of the 2nd International Conference on High Performance Compilation, Computing and Communications*, 2018, pp. 1-5.

© 2018 NTT DATA Mathematical Systems Inc.

NTT DATA

NTT DATA Mathematical Systems Inc.

Trusted Global Innovator

NTT DATA Group

NTT DATA

© 2018 NTT DATA Mathematical Systems Inc.