Evaluation of Tsunami Inundation Simulation using Vector-Scalar Hybrid MPI on SX-Aurora TSUBASA

Akihiro Musa, Takashi Soga
Tohoku University
NEC Corporation
Sendai, Japan
musa@tohoku.ac.jp, soga_takashi@nec.com

Takashi Abe
International Research Institute of Disaster Science, Tohoku University
Sendai, Japan
t-abe@irides.tohoku.ac.jp

Masayuki Sato
Graduate School of Information Sciences
Tohoku University
Sendai, Japan
masa@tohoku.ac.jp

Kazuhiko Komatsu
Cyberscience Center
Tohoku University
Sendai, Japan
komatsu@tohoku.ac.jp

Shunichi Koshimura
International Research Institute of Disaster Science, Tohoku University
Sendai, Japan
koshimura@irides.tohoku.ac.jp

Hiroaki Kobayashi
Graduate School of Information Sciences
Tohoku University
Sendai, Japan
koba@tohoku.ac.jp

***Abstract*—A real-time tsunami inundation forecast system has been developed since the 2011 Great East Japan earthquake occurred. Reducing the processing time and downsizing the system have been required for the forecast system. To this end, we develop a vector-scalar hybrid MPI code of a tsunami inundation simulation on NEC vector supercomputer SX-Aurora TSUBASA. The vector-scalar hybrid MPI code can run on both Xeon processors and SX vector processors as a single program. Experimental results show that the execution time of the vector-scalar hybrid MPI code is less than five minutes. As a result, the forecast system on a workstation type of SX-Aurora TSUBASA can deliver the prediction information within 17 minutes including the fault estimation and visualization in addition to the simulation.**

***Index Terms*—tsunami, MPI, SX-Aurora TSUBASA, performance, evaluation**

I. Introduction

In this poster, we discuss our vector-scalar hybrid MPI code for a tsunami inundation simulation on SX-Aurora TSUBASA to construct a real-time tsunami inundation system, which contains a *tsunami source modeling* process, a *tsunami inundation simulation*, and a *mapping and delivery* process. The tsunami inundation simulation solves non-linear shallow water equations by using the staggered leap-frog 2D finite difference method [2]. Figure 1 shows its computational domain and model size in the case of the Kochi prefecture region in Japan, on which devastating tsunamis are predicted. The coastal length is about 720 km. The computational domain is divided into four sub-domains with different grid-sizes. The simulation code is parallelized by the domain decomposition method. In this evaluation, the domain data shown in Fig. 1 are used.

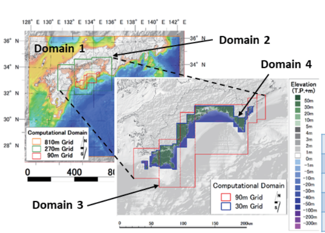

Domain	Grid size (meters)	Number of grid point
1	810	2.01×10^6
2	270	1.70×10^6
3	90	1.81×10^6
4	30	4.98×10^6

Area	1244 Km x 826 Km
Total grid points	1.05×10^7
Time interval (sec.)	0.5
Period (hours)	6
Flop count	2.29×10^{14}

Fig. 1. Computational domains and model size.

SX-Aurora TSUBASA [1] consists of a vector host (VH) and card-type vector engines (VEs). VH is a standard x86 Linux workstation and contains two Intel Xeon processors that is mainly used to control VEs. One VE contains a vector processor that is mainly used for the execution of the tsunami inundation simulation. This poster presents a hybrid implementation that uses not only vector processors but also Xeon processors to further accelerate the tsunami inundation simulation. By using a Vector-Scalar hybrid MPI, both processors can be used as a single system image.

II. Implementation of a hybrid MPI parallelization on SX-Aurora TSUBASA

Applying the vector-scalar hybrid MPI parallelization to the tsunami inundation simulation for SX-Aurora TSUBASA requires that computing loads of the tsunami inundation sim-

TABLE I
Specifications of SX-Aurora TSUBASA.

		VH	VE
Model Number		Xeon Gold 6148	20B
CPU	Number of cores	20	8
	Theoretical performance	1.4 Tflop/s	2.43 Tflop/s
	Memory bandwidth	128 GB/s	1.35 TB/s
System	Number of CPUs	2	1
	Memory capacity	192 GB	48 GB

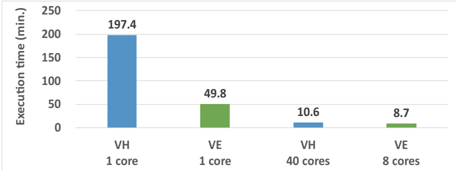

Fig. 2. Execution times of full cores and a single core on VH and VE.

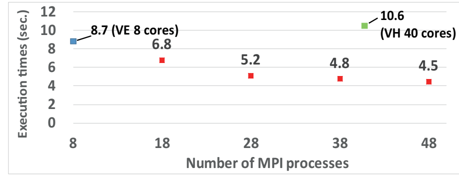

Fig. 4. Execution times of the hybrid MPI.

ulation should be proportionately allocated to VH and VEs depending on their performances and that the MPI communication between VH and VE is not a significant overhead. Therefore, we perform two preliminary evaluations on SX-Aurora TSUBASA. First, the computational performances of the simulation on VH and VE are evaluated, and then, the MPI communication is evaluated. In these evaluations, SX-Aurora TSUBASA with one VH and one VE, called a workstation type, is used. Its specifications are listed in Table I.

Figure 2 shows the execution times of full cores and a single core on VH and VE. The execution time of the full cores (40 cores) on VH is 1.2 times longer than that of the full cores (eight cores) on VE. Moreover, the single core of VE calculates the simulation in one-fourth the execution time of the single core of VH. The difference in performance between VH and VE on the simulation is clarified in this evaluation.

Furthermore, Figure 3 shows the performance of the MPI communication using the *PingPong* benchmark. *VE* indicates that MPI data are transferred between two cores within one VE. *VH* similarly transfers data within one Xeon processor. *Xeon-Xeon* exchanges data between two Xeon processors, and *Xeon-VE* is the case of transfer between one Xeon processor and one VE. The communication time on *Xeon-Xeon* quickly increases with the data size of 64 KB or more, and the times on *Xeon* and *Xeon-VE* also increase with 256 KB or more. Meanwhile, the times on *VE* is nearly constant, 1.0×10^{-6} seconds. In the tsunami inundation simulation using a 30-meter grid size, the maximum data size of MPI communication is about 64 KB. Therefore, the times for MPI communication have little impact on the hybrid MPI code.

Consequently, these preliminary evaluations show that the vector-scalar hybrid MPI parallelization of the tsunami inundation simulation has the potential of increasing the performance. As the implementation of a vector-scalar hybrid MPI parallelization, to balance load among VE and VH, four time more tasks should be assigned to a core of VE. Moreover, the simulation has two patterns of MPI communication: that within a domain and that between neighboring domains. To minimize the MPI communication time, neighboring domains are assigned into the same processors as much as possible.

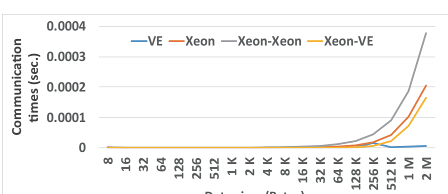

Fig. 3. MPI communication times on the workstation type.

III. Evaluations on Hybrid MPI

We evaluate the performance of the vector-scalar hybrid MPI code on the workstation type. Figure 4 shows the execution times of the code from eight MPI processes (VE 8 cores) to 48 MPI processes (VE 8 cores + VH 40 cores). The hybrid implementation can achieve 4.8 minutes of the execution time in the case of 48 MPI processes, although the scalability of the hybrid MPI code begins to slow down from 28 cores. As a result, the forecast system will be able to distribute the prediction information in less than 17 minutes which includes 7.5 minutes for the *tsunami source modeling* and 5 minutes for the *mapping and delivery* processes, both on VH.

IV. Summary

We have developed a vector-scalar hybrid MPI code for the tsunami inundation simulation on SX-Aurora TSUBASA. The experimental results show that the real-time tsunami inundation forecast system can deliver prediction information of tsunami inundation within 17 minutes. We are constructing a real-time tsunami inundation forecast system for prefectures on the workstation type of SX-Aurora TSUBASA, as shown Fig. 5. This system will provide tsunami prediction information: such as tsunami inundation depth and structural damage.

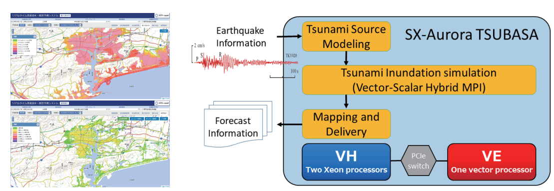

Fig. 5. Left: Samples of information: Inundation depth, Structural damage, Right: Diagram of the real-time tsunami inundation forecast system.

Acknowledgment

This study was partially supported by JSPS KAKENHI Grant Numbers 18K11322 and Ministry of Education, Culture, Sports, Science and Technology of Japan (Next Generation High-Performance Computing Infrastructures and Applications R&D Program).

References

[1] SX-Aurora TSUBASA. https://www.nec.com/en/global/solutions/hpc/

[2] A. Musa, O. Watanabe, H. Matsuoka, H. Hokari, T. Inoue, Y. Murashima, Y. Ohta, R. Hino, S. Koshimura, H. Kobayashi, "Real-time tsunami inundation forecast system for tsunami disaster prevention and mitigation," Journal of Supercomputing, vol. 74, issue 7, pp 3093–3113, 2018, https://doi.org/10.1007/s11227-018-2363-0.