

Beta Test Analysis of WIPL-DP

Saad N. Tabet*, NAVAIR, Patuxent River, MD, saad.tabet@navy.mil
Christopher Card, Black River Systems Company, Utica, NY, card@brsc.com

Abstract: During the Beta Test (BT) phase, WIPL-D [1] was parallelized for matrix fill/solution, to supplement the frequency parallelization developed in the Alpha Test (AT) phase [2]. The new WIPL-DP code was required to run on three distinct HPC platforms. The BT phase was the final testing phase of WIPL-DP, since the IOT&E was eliminated from the requirements. WIPL-DP was successfully parallelized for frequency and matrix fill/solution. The revised code received threshold level (or better) performance rating for all Critical Technical Parameters (CTPs) tested. The chosen test case for the BT phase was a further modification to the version used in the AT phase of the “Human Head Adjacent to a Cellular Phone” (DEMO-531) problem.

This WIPL-DP effort was funded under the Common High Performance Computing Software Support Initiative (CHSSI) as described in [2]. The goal of this initiative was to provide efficient, scalable, portable software codes, algorithms, tools, models and simulations that can run on a variety of DOD High Performance Computing platforms that can be used by scientists and engineers to solve computing problems.

Keywords: Parallelization, WIPL-D, WIPL-DP, Alpha Test, Beta Test, Scalability, Portability, Correctness, Stability, Accuracy, CHSSI.

Beta Test (BT)

1. Test Participants: For the WIPL-DP BT, only one subject matter expert (SME), Dr. Saad Tabet, NAVAIR was used. However, an independent group of expert users were assigned to test the WIPL-DP software using their own test cases, i.e., day-to-day problems of interest to them and their respective commands. The names of the users team members are available upon request.

2. Software Test Environment: The BT plan called for showing compatibility on three distinct HPC platforms. The three HPCMP high performance computing resource systems chosen were the Huinalu Linux super cluster and the Tempest IBM super cluster at the Maui High Performance Computing Center (MHPCC), and the Compaq SC-45 at the Aeronautical Systems Center Major Shared Resource Center (ASC MSRC).

The specifications for the Huinalu and Tempest systems are described in detail in [2]. The Compaq SC-45 machine is an SMP system with four CPU’s per node. Each CPU is a 1 GHz EV6.8 Alpha processor and contains a 64 KB primary instruction cache, and an 8 MB on-board cache. The SC-45 is partitioned into two separate systems. The partition used for BT contained 128 nodes, with 4 GB of memory per node.

3. Problem Under Test: A brief description of the test problem is provided. DEMO-531, “Human Head Adjacent to a Cellular Phone”, an example in the “Tutorial” sub-directory of the professional version of WIPL-D, was used as the foundation for the BT. Initially, the example was modified for the AT as a means to make the problem more

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 01 JAN 2005		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Beta Test Analysis of WIPL-DP				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NAVAIR, Patuxent River, MD				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM001846, Applied Computational Electromagnetics Society 2005 Journal, Newsletter, and Conference., The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

computationally intensive, as well as, cover the entire cellular communications frequency band (900 – 2400 MHz). For the BT, the model used in the AT phase was then further modified to make it even more computationally intensive, as well as, better conform to the actual shape of a cellular phone. The modified DEMO-531 test problem used in the BT phase is shown in Figure 1.

Test cases of 4, 8, 16, 32, and 64 frequencies, all bounded by the 900 – 2400 MHz range, were run. The test cases were set up such that the number of frequencies was set equal to the number of processors being used in the analysis.

Moreover, for comparison purposes, one-, two- and four-frequency “baseline” cases were run (on Tempest, Huinalu, and Compaq) using the originally converted non-parallelized C/C++ WIPL-D code. The one-frequency baseline results were used in the analysis of some of the test metrics.

In addition to the frequency parallelization employed during the AT phase, matrix fill/solution was added to the parallelization process of WIPL-DP during the BT phase. However, during the BT phase, the parallelization of WIPL-DP was a hybrid one, i.e., more than one form of parallelization being applied. WIPL-DP was parallelized for frequency and matrix fill/solution.

4. Test Metrics: The BT had to meet or exceed several test metrics, known as Critical Technical Parameters (CTPs). The CTPs are: scalability; portability; and correctness, stability, and accuracy. Each CTP had to meet an optimum objective and a minimum threshold.

The scalability CTP optimum objective is set to a scaled speed-up exceeding 80% of optimum on 64 processors. The minimum threshold is set to a scaled speed-up exceeding 25% of optimum on 32 processors. The scalability CTP is determined by comparing the WIPL-DP runs to the one-frequency baseline case, using the scaled speed-up in percent (S) given in [2].

The portability CTP optimum objective required that WIPL-DP runs on three HPC platforms (Tempest, Huinalu, and Compaq in this case) producing very similar and valid results. The threshold (i.e., minimum) objective required that WIPL-DP runs on two HPC platforms.

The correctness, stability, and accuracy CTP optimum objective is for WIPL-DP to produce results that match the commercial WIPL-D results, value for value, with a maximum percent error of no worse than 2% (accuracy of 98% or higher). The minimum threshold relaxes the optimum objective maximum percent error to no worse than 3% (accuracy of 97% or higher). The maximum percent error (ϵ_{\max}) equation is given in [2].

5. BT Management: As soon as the BT Plan was approved by HPCMO, the SME started the testing process. The one-, two-, and four-frequency baseline cases were run on all three HPC systems using the non-parallelized C/C++ WIPL-D code. These cases were run to prove that the processing time scaled proportionally to the number of frequencies used. In each successive case the processing time was doubled since the numbers of frequencies were doubled. Unlike in [2], the single processor non-parallelized case could not be used as a comparison baseline for the

parallel code due to the use of a different matrix solver. The parallelized version of the code used a parallel matrix solver that proved to be more efficient for the single processor case than that of the original solver used in WIPL-D. Thus the baseline timing that was used came from running the parallelized code with a single processor and frequency, allowing for a true speedup test as described below.

Utilizing the Windows commercial version of WIPL-D, the modified DEMO-531 model was run for 4, 8, 16, 32, and 64 frequencies. These runs were necessary to determine the accuracy CTP results of WIPL-DP. The PC results were treated as theoretical values, since the commercial WIPL-D code has been well validated over its years of existence.

The next stage in conducting the BT was to run WIPL-DP on the three distinct HPC platforms; Tempest, Huinalu, and Compaq. Cases of 4, 8, 16, 32, and 64 frequencies utilizing 4, 8, 16, 32, and 64 nodes, respectively, were run on each platform. Also, a single processor single frequency case was run on each system. These runs were used as a baseline to calculate the speed-up, for reasons described above. The results from these cases, when compared to the single processor baseline for each machine, determined whether the BT was a success or not.

6. Results and Conclusions: WIPL-DP CTPs were compared to their baseline counterparts. Scalability CTP (i.e., speed-up) test results for Tempest, Huinalu, and Compaq are shown in Figure 2. Figure 2 shows that the scalability CTP measure was established. The worst speed-up achieved was over 70% (Compaq for 64 processors), which is within 10% of the optimum objective of 80% for 64 processors (green line.) However, it far exceeded the threshold requirement of 25% for 32 processors (red line.)

The portability CTP was successfully achieved, since WIPL-DP ran quite successfully on three distinct HPC platforms (Tempest, Huinalu, and Compaq.) Moreover, similar results were achieved on all three HPC systems.

Accuracy CTP results are shown in Table 1. The results in Table 1 show that the accuracy CTP was established. The maximum error recorded in all the compared cases was less than 0.083%, registered in the “.ra1” file of the 4-frequency hybrid case on all three HPC systems. This maximum error was more than an order of magnitude below the optimal objective set for the test (i.e., less than 2%.)

In conclusion, WIPL-DP passed the BT with only the scalability CTP not achieving its optimum objective, and only for the 64-processor case for all three HPCs.

7. Final Comments: A side goal of the WIPL-DP development was to be able to solve up to 100,000 unknowns, a far improvement over the 15,000-unknown limitation existing under the 32-bit Windows environment. However, at this point, the number of solvable unknowns is around 30,000. Late in the project, a problem was identified with high numbers of unknowns. This problem was unfortunately not resolved before the end of the project due to time and financial constraints. Two follow on projects have been proposed that will look into this problem and push the parallel code past the current matrix size limitation.

References

- [1] B.M. Kolundzija, J.S. Ognjanovic, and T.K. Sarkar, "WIPL-D: Electromagnetic Modeling of Composite Metallic and Dielectric Structures – Software and User's Manual," Artech House Publishers, Norwood, MA, 2000.
- [2] S. Tabet, L. Joiner, and M. Benincasa, "Alpha Test Analysis of WIPL-DP," 20th Annual Review of Progress in Applied Computational Electromagnetics, 2004 ACES Conference Proceedings, Syracuse, NY, April 19-23, 2004.

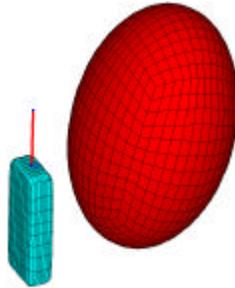


Figure 1. Modified DEMO-531 Model Used During The BT Phase.

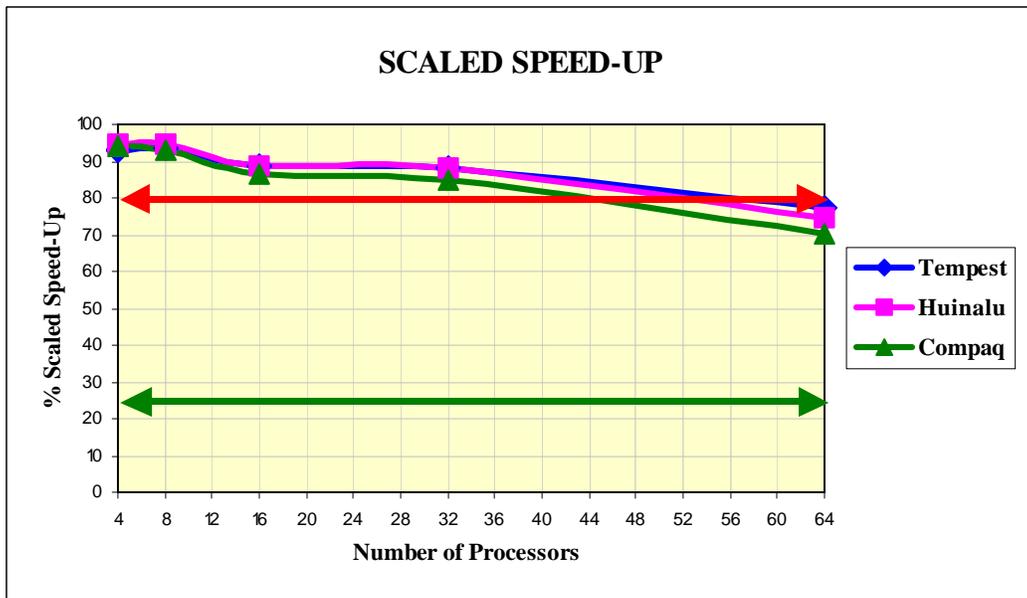


Figure 2. Percent Speed-up Versus Number of Processors.

Table 1. Hybrid WIPL-DP Accuracy CTP Data.

<u>Accuracy Summary</u>			
Output File	Max Error (%)	Frequency Case	HPCC System
ad1	0.0027746	16	Tempest, Huinalu
cu1	0.0048228	16	Tempest, Huinalu, Compaq
nf1	0.0221264	64	Compaq
ra1	0.0823788	4	Tempest, Huinalu, Compaq