Alpha Test Analysis of WIPL-DP

Saad Tabet[1], Laurie Joiner[2], and Milissa Benincasa[3]

1. NAVAIR, Radar and Antenna Systems Division, Patuxent River, MD, saad.tabet@navy.mil
2. University of Alabama in Huntsville, Dept. of ECE, Huntsville, AL, ljoiner@ece.uah.edu
3. Black River Systems Company, Utica, NY, benincas@brsc.com

Abstract: WIPL-D (WIre PLate Dielectric) has become an increasingly popular Method of Moments (MoM) code used in computational electromagnetics (CEM) modeling. WIPL-D was chosen for parallelization under the Common High Performance Computing Software Support Initiative (CHSSI) program of the High Performance Computing Modernization Office (HPCMO). Hence, the new code was given the name WIPL-DP, where “P” stands for Parallelized.

Any computer code chosen for parallelization under the CHSSI program must undergo four rigorous phases of testing: Software Acceptance Test, Alpha Test (AT), Beta Test (BT), and Initial Operational Test and Evaluation (IOT&E). WIPL-DP is currently undergoing those tests, with the Alpha Test recently completed and reported on in this effort. The Beta Test and the IOT&E are to be completed by 30 September 2004.

WIPL-DP is a parallelized C/C++ version of the original FORTRAN 77 WIPL-D code. During the Alpha Test period, WIPL-DP was successfully parallelized for frequency. It also received optimal performance rating for the following Critical Technical Parameters (CTPs): scalability; portability; and correctness, stability, and accuracy. The chosen test case for the Alpha Test was a modified version of the “Human Head Adjacent to a Cellular Phone” (DEMO-531) problem available under the tutorial sub-directory in the PC version of the WIPL-D software. The Alpha Test was performed on two distinct High Performance Computing (HPC) platforms, Tempest and Huinalu, both at the Maui HPC Center.

1. Background on WIPL-DP CHSSI Effort

1.1 CHSSI Program: The CHSSI program is an initiative that is part of the High Performance Computing Modernization Program (HPCMP). The HPCMP program was started in 1992 to provide high performance computing resources to Department of Defense (DoD) laboratories. The HPCMP provides high performance computing services, high-speed network communications, and computation science expertise. CHSSI is one of the programs under the HPCMP. The objective of the CHSSI initiative is to provide portable, scalable, efficient software codes, algorithms, and tools that can be run on the high performance computing resources available through the HPCMP program. CHSSI consists of ten computational technology areas (CTAs) that were designed to support the collaboration between government, industry and academia.

Under the CHSSI program DoD laboratories can submit proposals for the development of parallel software codes, algorithms, and tools that support a specific CTA. The proposals can be a collaboration of government, industry, and academia. The CHSSI Initiative each year puts out a call for proposals. The proposals are then evaluated and selected based upon specific criteria.

1.2 WIPL-DP Effort: The WIPL-DP effort is currently funded out of the Integrated Modeling and Test Environments (IMT) CHSSI CTA. It is a three-year parallel software tool development that is just starting its third year. The goal of the WIPL-DP effort is to develop a scalable, portable, parallel scene generation tool that will provide tri-service capability to quickly generate scenes of radiating and scattering structures (targets and their surrounding environment) in realistically complex electromagnetic environments. This improvement will move users from accurate modeling of targets only to accurate modeling of targets in their environment. WIPL-DP will be applicable to a wide range of end-user applications in the Army, Navy, Air Force, and Marines. Applications that can benefit from such a tool include: Detection of Targets under Trees, Ship Radar Performance, Strategic Subsurface Target Detection, Land Mine Imaging, and Land Mine Detection just to name a few.

As mentioned above efforts in the CHSSI program must undertake four phases of testing. In the first year of the WIPL-DP effort, SAT needed to be completed. The purpose of SAT was to provide a formal test that provides a starting point for the effort, evaluate/refine the technical approach, identify any risks in the project, and make sure the effort has DoD relevance. SAT was completed in the summer of 2002. The next major milestone test for the second year of the effort was Alpha Test. The purpose of AT is to provide an evaluation of the performance measurements of well-functioning software in the middle phase of development. Alpha testing took place in August 2003. Below is a detailed discussion of how the WIPL-DP ST process was conducted and the performance results obtained.

2. Alpha Test (AT)

2.1. Test Participants: Two subject matter experts, Dr. Saad Tabet from NAVAIR and Dr. Laurie Joiner from the University of Alabama in Huntsville, performed the WIPL-DP AT. The tests were run independently to ensure the validity of the results. A third member of the team, Mr. Joseph Schneible, also tested the software for ease of use and compatibility, but did not perform the AT.

2.2. Software Test Environment: The AT plan called for showing compatibility on two different systems. The two systems chosen were the Huinalu Linux super cluster and the Tempest IBM super cluster; both are part of the Maui High Performance Computing Center (MHPCC), which are HPCMP high performance computing resources.

Huinalu is a 520 processor IBM Netfinity Linux Super cluster, containing 260 nodes. Each node has two Pentium III 933 MHz processors and 1 GB of local memory (512 MB of memory per CPU). Although the nodes are connected via a high-performance Myrinet switch (with 200 MB/second of sustained bandwidth), only the 100 Mbit Ethernet connections were employed for this test.

Tempest has two partitions under its scheduler. The first partition, P3, has 46 nodes and 736 processors. The nodes are 16-way, 375 MHz Nightnawk-2 nodes, each sharing 8 GB of memory. The IBM Colony Switch, with a bandwidth of 400 MB/second, connects the nodes. The second partition, P4, contains ten 1.3 GHz Regatta nodes. Each node contains 32 CPUs and 32 GB of shared memory. Since one of the main objectives of the AT was to demonstrate the run time execution, all of the Tempest jobs had to run on the same partition. In order to get access to the system under the time constraints of the test, the slower P3 partition was utilized.

2.3. 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 AT. The example was modified for the AT, as a means to make the problem more computationally intensive, as well as, cover the entire cellular communications frequency band (900 – 2400 MHz). The modified DEMO-531 test problem is shown in Figure 1.

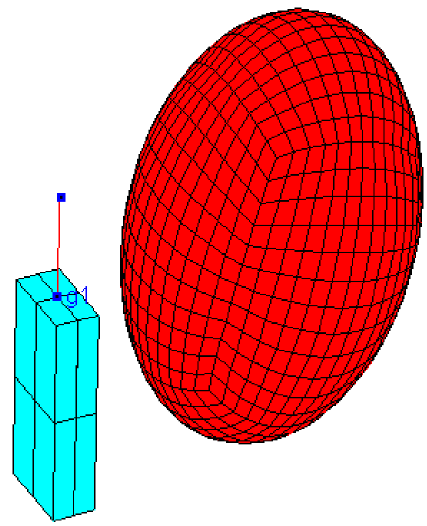

Figure 1. Modified DEMO-531 "Human Head Adjacent to Cellular Phone" Model.

Test cases of 2, 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 to twice the number of processors being used in the analysis.

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

2.4. Test Metrics: The AT 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 32 processors. The minimum threshold is set to a scaled speed-up exceeding 25% of optimum on 16 processors. The scalability CTP is determined by comparing the WIPL-DP runs to the two-frequency baseline case, using the following equation:

$$S = \frac{\mathrm{T(1)}}{\mathrm{T(N)}} \times 100\% \quad \text{where,}$$

S = Scaled speed-up in percent
T(1) = Runtime from running non-parallel WIPL-D C/C++ baseline code

T(N) = Runtime from running same problem using WIPL-DP on N-processors
N = 2, 4, 8, 16, or 32 processors

The portability CTP optimum objective and minimum threshold are one in the same. In both cases, the WIPL-DP code has to run on two HPC platforms (Tempest and Huinalu in this case) producing the same valid results.

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 4% (accuracy of 96% or higher). The minimum threshold relaxes the optimum objective maximum percent error to no worse than 5% (accuracy of 95% or higher). The maximum percent error (ε_{max}) is determined using the following equation:

$$\varepsilon_{max} = \frac{\max[\mathrm{abs}(C_{val} - P_{val})]}{\max[\mathrm{abs}(C_{val})]} \times 100\% \quad \text{where,}$$

C_{val} = Commercial WIPL-D generated value
P_{val} = Parallelized WIPL-D generated value

2.5. Management of the AT: The AT test was conducted in a very systematic fashion. After the Test Plan was approved by HPCMO, the AT SMEs began their independent testing process. Each SME conducted the same series of tests as specified by the Test Plan, and performed the CTP evaluations based on their independently run cases.

The two- and four-frequency baseline cases were run on both Tempest and Huinalu using the non-parallelized C/C++ WIPL-D code, converted from the FORTRAN 77 original code. The two-frequency case was required to determine the scalability CTP results of the parallel code. The four-frequency case was used to ensure that the program is functional, and runs in approximately twice the time of the two-frequency case, since no parallelization was employed.

Utilizing the Windows commercial version of WIPL-D, the modified DEMO-531 model was run for 2, 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 AT was to run WIPL-DP on two distinct HPC platforms. Tempest and Huinalu, two separate clusters in the Maui HPC system, were used. Cases of 2, 4, 8, 16, 32, and 64 frequencies utilizing 1, 2, 4, 8, 16, and 32 nodes, respectively, were run on each platform. The results from these cases, when compared to the baselines, determine whether the AT was a success or not.

2.6. Results and Conclusions: Using MATLAB scripts and other intermediate scripts developed by Black River Systems Company (BRSC), WIPL-DP CTPs were compared to their baseline counterparts. Scalability CTP (speed-up) test results are shown in Figure 2 and Table 1.

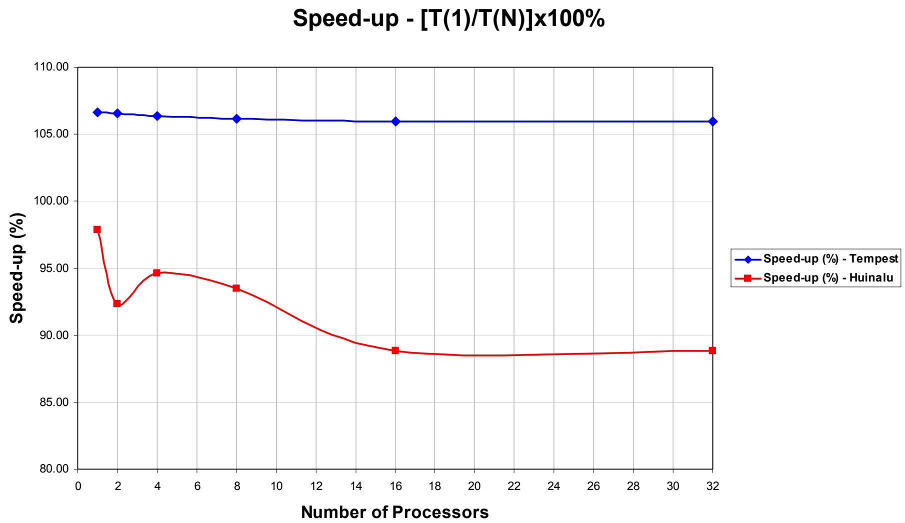

Figure 2. Percent Speed-up Versus Number of Processors for Tempest and Huinalu.

Single Processor Baseline		
Frequencies	Time (s) - Tempest	Time (s) - Huinalu
2	12478	5793
4	24903	11937

Parallel					
Frequencies	Processors	Time (s) - Tempest	Time (s) - Huinalu	Speed-up (%) - Tempest	Speed-up (%) - Huinalu
2	1	11703	5920	106.622233615	97.854729730
4	2	11708	6276	106.576699693	92.304015296
8	4	11731	6123	106.367743585	94.610485056
16	8	11755	6195	106.150574224	93.510895884
32	16	11775	6519	105.970276008	88.863322595
64	32	11776	6519	105.961277174	88.863322595

Table 1. Baseline and Parallel Cases Runtime and Speed-up Data for Tempest and Huinalu.

Those results show that a scalability CTP measure was established. The worst speed-up achieved was over 88% (for Huinalu), far exceeding the optimum objective of 80%. Also, worth noting was the above 105% speed-up for all the Tempest tests performed. Such speed-ups were found to be somewhat unbelievable, and after some investigation by Mr. Christopher Card of BRSC, a minor error in the Tempest setup files was found. A profiling switch was active in the non-parallel C/C++ baseline WIPL-D code, but was inactive in the WIPL-DP runs. The corrected results are shown in Appendix A, where it is obvious that speed-ups exceeding 99% were achieved for all the Tempest cases tested.

The portability CTP was successfully achieved, since WIPL-DP ran quite successfully on two distinct HPC platforms, Tempest and Huinalu. Accuracy CTP results are shown in Table 2.

	% Error for Huinalu - Frequencies					
Output File	2	4	8	16	32	64
ad1	0.001492588	0.001492588	0.001492588	0.001458720	0.001465832	0.006813435
cu1	0.011694553	0.006301755	0.004998560	0.004684291	0.004615556	0.004599328
nf1	0.011429815	0.008199148	0.011874061	0.034300655	0.016161789	0.054710623
ra1	0.035063114	0.079595173	0.025900026	0.067751003	0.068143060	0.068099226

	% Error for Tempest - Frequencies					
Output File	2	4	8	16	32	64
ad1	0.001653883	0.000932868	0.001492588	0.001465981	0.001465832	0.001465819
cu1	0.000000000	0.000630176	0.004998560	0.004684291	0.004615556	0.004599328
nf1	0.005575669	0.009380797	0.010468812	0.034300583	0.016161789	0.034216534
ra1	0.053188366	0.079595173	0.068843553	0.067754391	0.068143060	0.067764651

Accuracy Summary			
Output File	**Max % Error**	**Frequency**	**Machine**
ad1	0.006813435	64	Huinalu
cu1	0.011694553	2	Huinalu
nf1	0.054710623	64	Huinalu
ra1	0.079595173	4	Tempest

Table 2. Baseline and Parallel Cases Runtime and Speed-up Data for Tempest and Huinalu.

Those results show that an accuracy CTP was established. The maximum error recorded in all the compared cases (< 0.08%, Tempest) was more than an order of magnitude below the optimal objective set for the test.

In conclusion, WIPL-DP passed the AT far exceeding the optimum objectives set for the CTPs. With the employment of some recommendations, WIPL-DP will become a more effective CEM tool. Moreover, parallelizing the impedance matrix inversion/solution during the BT phase, in addition to the AT frequency parallelization, shall significantly increase the versatility and abilities of WIPL-DP.

3.0 Future Direction

In the final year of the effort two more critical tests must be passed. The next test is Beta Test. In Beta test the WIPL-DP code will be tested to assure that performance measurements are met for a software code in the later stages of development. BT will evaluate the code and the code documentation to determine if WIPL-DP meets the agreed goals set forth in the beginning of the program. BT is scheduled to take place in April 2004. The final test will be OTTR, which is when the software final acceptance testing takes place. This test will occur in September 2004.

Acknowledgements: The authors would like to recognize Mr. Christopher Card of Black River Systems Company for his dedication to the CHSSI WIPL-DP project. We are also very thankful for the long hours of assistance that he lent the SMEs to make sure that they were able to fully perform their duties during the Alpha Test.

<u>Appendix A</u>: The below results are the corrected Tempest results, turning off the profiling switch.

Frequencies	Processors	Speedup - Tempest
2	1	100.034179270
4	2	99.991458832
8	4	99.795413861
16	8	99.591663122
32	16	99.422505308
64	32	99.414062500

