

# Deployment of WRF Model Using Bootable Cluster CD (BCCD) on Heterogeneous PC Systems

Shing Yoh and Braden Ward  
Department of Geology & Meteorology  
Kean University, Union, NJ 07083

## Abstract

Using Message-Passing Interface (MPI) implementation, the Weather Research and Forecasting (WRF) Model is suitable for high performance computing on cluster machines. At an education institution like Kean University, there are many student computers running Microsoft Windows as their operating system. These networked computers consist of a variety of processor speeds and are unused during overnight hours and weekends. These PCs are ideal for building a parallel computing cluster environment for the WRF Model.

In order to build a distributed computing environment using diverse PCs (running Windows, Linux, and MAC) with no interruption to the underlying operating system on each computer, the approach of using Bootable Cluster CD (BCCD) is adopted for this study. Each computer will be rebooted using the BCCD (v2.2.1c) to a Linux environment. With the mounting of a Network File System (NFS) to each of the computer for software storage (MPICH library with PGI F90/F95 compiler, netcdf library, and WRF system) and saving output, we have successfully run the WRF Model using 1 to 16 homogeneous and heterogeneous PC systems.

In this paper, we will describe the parallel performance of WRF V2.1.1 (both ideal case and real case) using the distributed computing environment created by BCCD. Comparison will also be made with running the same WRF Model on a dedicated 16-node Linux cluster built by the Computational Science Group at Kean University. The strength and weakness of using BCCD in doing WRF research will be discussed.

## 1. Introduction

The Weather Research and Forecasting (WRF) Model is a development of collaboration between numbers of institutions, providing a forecasting model that can be tailored for the user. The model has a growing user base, with advances in physics, numeric, and data assimilation from the research community.

The WRF software is portable over a variety of platforms using shared, distributed-memory and hybrid parallel architectures, using both vector and scalar processor types (Michalakes et al., 2004). Using Message-Passing Interface (MPI) implementation, the WRF Model is highly suitable for high performance computing on dedicated cluster machines. WRF has been

successfully run on variety of homogeneous massive parallel architectures (Michalakes et al., 2004) and on heterogeneous systems (Huang and Abernethy, 2005). However, building and using a dedicated distributed system to run the WRF Model can be expensive and may involve levels of difficulty beyond WRF users that are not experts in computer science. This project explores increasing the portability and efficiency of running the WRF Model in common, heterogeneous computing environments using 2 to 16 personal computers without extensive knowledge in computer science and parallel computing.

## 2. Computing Environment

Kean University, as such with many institutions, is host to numerous student computers of varying types. The specific personal computers used for this study are sixteen old 448 MHz Pentium 3 desktops with 256 MB of memory (referred as p3 in this paper), and sixteen 2.8 GHz Pentium 4

---

\* *Corresponding author address* : Dr. Shing Yoh, Department of Geology & Meteorology, Kean University, Union, NJ 07083.  
E-mail : syoh@kean.edu

desktops with 2.0 GB of memory (referred as p4 in this paper), all manufactured by Dell. These computers, which are all on a network and sub-netted using 100 Mbps switches, are normally running Microsoft Windows and are used by students for Internet access and schoolwork. These student computers are an ideal resource for running the WRF Model, as they remain idle overnight, and are widely unused over weekends. In addition, for comparison purposes, identical model runs will be performed on a dedicated 16-node Linux cluster (diskless except the master node), built by the Computational Science Group at Kean University using 3.0 GHz Pentium 4 processors with 1 GB of memory per node, linked by 1 Gbps switches.

The key goal of this project is to use the common desktop machines with a non-intrusive, fast setup process in order to realistically operate them as a cluster for WRF Model runs. This ability exists through the use of Linux bootable CD. The Bootable Cluster CD (BCCD) v2.2.1c developed by Paul Gray of the University of Northern Iowa was chosen for this project. The BCCD was created to facilitate instruction of parallel computing aspects and paradigms (<http://bccd.cs.uni.edu/>). It features a 2.4 Linux kernel, booting directly off of the CD, and returning the machine to its original state following shutdown.

A critical step in using the bootable CD to create a cluster for WRF model runs is the mounting of a permanent file system. For our purpose, we mount a Network File System (NFS) from the department Dell PowerEdge 4400 server. This is a critical step in creation of the processing cluster, as each machine requires access to the same WRF input and output files. It must also be ensured that all machines are in sync with the same system time. The setup process is basic, requiring 2 minutes of user time for setup per machine. This setup time can be greatly reduced if the BCCD is modified and reconfigured specifically for running parallel WRF model.

Once the system setup is completed, there is a one time installation of netcdf library (v3.6.0), MPICH (v.1.2.7) with PGI F90/F95 compiler (the BCCD comes with MPICH without F90 compiler), and WRF Model system (v2.1.1) on the network drive. Following the loading process of the BCCD on a group of machines, the desired systems can be isolated and selected for model runs. The compiled codes, data and output are all saved on the network drive though NFS.

### 3. Results

To best show the benefits and results from running the WRF in this parallel computing environment, two model runs are performed: the ideal supercell (em\_quarter\_ss in the WRF test cases) and a real case 12 hour forecast starting from 00UTC February 11<sup>th</sup>, 2006 centered over New Jersey. The real case is configured with grid resolution close to 15 km (100 x 100 x 31 grid points). The initial and boundary conditions are generated from the 00Z NAM run from NCEP using WRFSI.


All the model runs for this project were executed in student computer labs when no students were using the computers in order to minimize network traffic. The runtimes are captured using the Linux "time" command when the run was performed, as well as scripting a system timestamp before and after the model run. The average timestep is based on the WRF Model output, ignoring the time used to write intermediate output files. All of the timing intervals are used in the average with use of Microsoft Excel.

Table 1 below gives ideal supercell case runtimes and times for a model timestep for the older Pentium 3 (p3), the newer Pentium 4 (p4), and the dedicated cluster (c16) machines. Figure 1 shows the decrease in runtime of the ideal supercell case as the number of processors increase. Except for the slower p3 case, the ideal supercell does not depict as great of parallel efficiency. The faster p4 cluster and the c16

cluster shows less than 2 times faster in runtime when the WRF model is benchmarked using 16 processors instead of 1. However, the slower p3 cluster shows a corresponding 5 times faster when 16 processors are used.

**Table 1. Mean wall time and timesteps for supercell ideal runs**

<i>Machine</i>	<i>runtime</i> (minutes)	<i>timestep</i> (seconds)
<b>Single processor p3:</b>	72.4	7.16
<b>p4:</b>	11.6	1.14
<b>c16:</b>	8.7	0.86
<b>16 processors p3:</b>	14.5	1.36
<b>p4:</b>	6.3	0.58
<b>c16:</b>	3.7	0.33




**Figure 1. Mean wall times vs. number of processors for supercell ideal runs. The time is given in minutes along the vertical axis, the number of processors used is the horizontal axis.**

Table 2 outlines the performance of the real case simulation with the Pentium 4 (p4) machines, compared to the Linux cluster (c16) machine. An increase of processor from 1 to 16 reduces the runtime by about 75% in both cases. Figure 2 shows a decrease in runtime of the real case as the number of processors increases. Similar to Figure 1, the most dramatic improvement in the performance is when the number of processors is increased from 1 to 4 to 8. Above 8 processors, the rate of improvement slows down, and hits the limitations of other components of the PCs. The real case is also performed with the Pentium 3 (p3) machines. Their performance shows the same dramatic improvement as

with the ideal case. However runtimes from only 6 to 16 processors are shown here due to long runtimes when 4 or fewer machines are used. The runtimes do surprisingly mirror those of the faster machines when 8 or more are used, which suggests that the further decreases in runtime are limited by PC components other than the processor speed and memory.

**Table 2. Mean wall time and timesteps for real case runs**

<i>Machine</i>	<i>runtime</i> (minutes)	<i>timestep</i> (seconds)
<b>Single processor p4:</b>	105.7	3.88
<b>c16:</b>	82.1	3.42
<b>16 processors p4:</b>	25.3	0.97
<b>c16:</b>	17.5	0.73



**Figure 2. Mean wall time vs. number of processors for real case runs. The time is given in minutes along the vertical axis; the number of processors used is the horizontal axis**

#### 4. Discussion and Summary

This study shows that the WRF Model can be run across heterogeneous computer systems, and done so with a low level of difficulty and setup time. As apparent with the ideal and real case runs, the Pentium 4 cluster built from common desktop computers through use of the BCCD closely mirror the results of the dedicated Linux cluster. The times of the p4 cluster fall behind that of the c16 cluster machines, due to the slower processor speed (2.8 GHz vs. 3.0 GHz) and switches (100 Mbps vs. 1 Gbps). As depicted in figure 2, p4 desktops

run an average difference of 22% behind the c16 runs.

With the much older, Pentium 3 systems, amiable results are shown. Although the runs with smaller processor numbers lag behind by a significant amount, 8 or more of them together perform the ideal case runs an average of only 51% slower than the p4 cluster machines. This is good result, as the p4 machines processor speed is over 3 times that of the p3 machines, as well as 8 times the amount of memory. Although these old computers perform faster in large numbers, they should not be mixed with the p4 machines to run the model. Experiments were performed using four slow p3 machines with four fast p4 machines running the ideal supercell case. With this configuration it takes 14.8 minutes to complete, which is close to the 16.2 minutes taken when all eight are p3 systems. However, the same model run only takes 8.1 minutes with four p4 systems. Adding a slow machine will slow down the cluster; the slowest system is therefore a limiting factor. Care should be taken when selecting heterogeneous systems for parallel computing. Mixing fast and slow processors seems to be feasible only when WRF and MPI can assign different weights to different processors in parallel computing.

The BCCD is an excellent tool to utilize for implementation of the WRF Model. In principle, the whole WRF model system can be distributed on a bootable CD format such that users can simply boot the CD on any number of computer systems, mount a permanent file system on each system, and then compile (hopefully without the need of any propriety FORTRAN compiler), and run the model. With the popularity of Networked Attached Storage (NAS) and networked personal computers, the idea of running the WRF Model on parallel clusters with no interruption to the underlying operating system on each computer is very appealing to many users, such as undergraduate students, researchers

or operational forecasters with limited resources.

Relating to this project, even with dated resources one could improve model runtime performance efficiently in parallel. WRF Model can easily be deployed for parallel computing on any number of heterogeneous systems using BCCD, but the most effective deployment would be from 4 to 8 computers with similar specifications.

## 5. Acknowledgements

We would like to thank Dr. George Chang from the Computer Science Department and Computational Science Group at Kean University for building the Linux Cluster and providing technical support. Dr. Paul Gray of the University of Northern Iowa for developing the BCCD and making the software freely available. The project is supported by the Kean University Office of Research and Sponsored Programs and the Department of Geology of Meteorology.

## 6. References

- Huang, W. and J. Abernethy, 2004 : Examination of the Weather Research and Forecast (WRF) Model on Heterogeneous Distributed Systems. *5<sup>th</sup> WRF User Workshop*, June 22-25, 2004, Boulder, CO.
- Michalakes, J., J. Dudhia, D. Gill, T. Henderson, J. Klemp, W. Skamarock, and W. Wang, 2004: The Weather Research and Forecast Model: Software Architecture and Performance. *Proceedings of the 11<sup>th</sup> ECMWF Workshop on the Use of High Performance Computing In Meteorology*, 25-29 October 2004, Reading U.K.