

The book was found

Raspberry Pi Super Cluster



Synopsis

As a Raspberry Pi enthusiast have you ever considered increasing their performance with parallel computing? Discover just how easy it can be with the right help - this guide takes you through the process from start to finish.

Overview Learn about parallel computing by building your own system using Raspberry Pi

Build a two-node parallel computing cluster Integrate Raspberry Pi with Hadoop to build your own super cluster

In Detail A cluster is a type of parallel/distributed processing system which consists of a collection of interconnected stand-alone computers cooperatively working together. Using Raspberry Pi computers, you can build a two-node parallel computing cluster which enhances performance and availability. This practical, example-oriented guide will teach you how to set up the hardware and operating systems of multiple Raspberry Pi computers to create your own cluster. It will then navigate you through how to install the necessary software to write your own programs such as Hadoop and MPICH before moving on to cover topics such as MapReduce.

Throughout this book, you will explore the technology with the help of practical examples and tutorials to help you learn quickly and efficiently. Starting from a pile of hardware, with this book, you will be guided through exciting tutorials that will help you turn your hardware into your own super-computing cluster. You'll start out by learning how to set up your Raspberry Pi cluster's hardware. Following this, you will be taken through how to install the operating system, and you will also be given a taste of what parallel computing is about. With your Raspberry Pi cluster successfully set up, you will then install software such as MPI and Hadoop. Having reviewed some examples and written some programs that explore these two technologies, you will then wrap up with some fun ancillary projects. Finally, you will be provided with useful links to help take your projects to the next step.

What you will learn from this book

- Discover how to set up the hardware to build your parallel computing cluster
- Set up your Raspberry Pi computers and install an operating system
- Network your two Raspberry Pi's together
- Gain an understanding of MPI through practical examples
- Learn how to work with MPICH to write parallel applications
- Install Hadoop and experiment with processing text files
- Get acquainted with MapReduce, the paradigm at the heart of Hadoop

Approach This book follows a step-by-step, tutorial-based approach which will teach you how to develop your own super cluster using Raspberry Pi computers quickly and efficiently. Who this book is written for

Raspberry Pi Super Cluster is an introductory guide for those interested in experimenting with parallel computing at home. Aimed at Raspberry Pi enthusiasts, this book is a primer for getting your first cluster up and running. Basic knowledge of C or Java would be helpful but no prior knowledge of parallel computing is necessary.

Book Information

Paperback: 126 pages

Publisher: Packt Publishing (November 20, 2013)

Language: English

ISBN-10: 1783286199

ISBN-13: 978-1783286195

Product Dimensions: 7.5 x 0.3 x 9.2 inches

Shipping Weight: 10.7 ounces (View shipping rates and policies)

Average Customer Review: 3.9 out of 5 stars [See all reviews](#) (12 customer reviews)

Best Sellers Rank: #1,047,849 in Books (See Top 100 in Books) [#34 in Books > Computers & Technology > Programming > Languages & Tools > Fortran](#) [#80 in Books > Computers & Technology > Hardware & DIY > Mainframes & Minicomputers](#) [#195 in Books > Computers & Technology > Hardware & DIY > Peripherals](#)

Customer Reviews

To be honest and straightforward I expected more from a book with title like Raspberry Pi Super Cluster. The author Andrew K. Denis has a very clear vision on the subject (like in his previous book Raspberry Pi Home Automation with Arduino, which I liked a lot). He's done his best to deliver an exhaustive set-up while being concise at the same time, but it seems to me, this clearly is the wrong format for a book on the given topic. Now having this book at hand, I finally got the chance to answer many of the questions I had about clustering, and how it can be applied to a set of Raspberry Pis. The first impression is that it is very well structured and gradual. Lets see, the first two chapters are short introductions to parallel computing (background history and the contemporary systems) and the initial set-up respectively. They're short and to the point. And that's the way it should be - it is presumed that if you're going parallel, then you're somewhat advanced tinkerer already. Actually the second chapter is pretty abundant in details on how to install the operating system, the required software and tools. I skimmed through it, because I already had the two Raspberry Pi units pretty well equipped with what was needed. The next chapter is the first encounter with a parallel software in the face of MPICH - one of the oldest and most widely adopted implementations of the MPI (Message Passing Interface) implementations, which is designed for applications written in C, C++ or Fortran. In this chapter we also come to one tricky part - setting up of the second (equally applicable to third, fourth, and so on) Raspberry Pi unit. It is tricky because it's a continuation of the set-up started in the second chapter and must be followed strictly. Especially the part with the RSA keys exchange. If you get just one thing wrong, you may have to start all over (like myself). The

good news is that the procedure is short and not as much obscure, as one can imagine for a set-up concerning security matters. Once the berries are prepared correctly, the only thing you'll care further on would be the parallel frameworks and your applications. At chapter four, after we've calculated the number Pi with a small MPICH application written in C, we finally arrive at one of the most popular representatives of the modern trends in parallel software - Apache Hadoop. Its installation is quick, but the configuration is a bit detailed, especially when you take into account that most of the things have to be done at least twice. Here I met the biggest downside of the book - the lack of any trouble shooting for the situations when you get stuck. Although I followed every step verbatim, there were errors logged on the console, for which there was no help around. Fortunately the messages are somewhat self explanatory, so with little deduction one can get to the next step fairly easy. Another disadvantage, if I may call it that, is Hadoop's version. I don't know when the book was written, but when it says "Download the latest version", on the project's site you get version 2.2.0, while for the book this is version 1.2.1. This wouldn't be much of a problem, if Hadoop's architecture hasn't been changed significantly. So if you prefer the latest, the instructions in the chapter are of no use for you. If the author had a good reason to stick to the older branch of the software, this reason remains obscure to the reader. There are few lesser inaccuracies like wrong documentation URL, not whole scp commands, and a sense of text that was a bit too rushed. Now having the framework for parallel computing already set, it is time to test it with an application. Since Hadoop is written in Java (as a typical Apache project), its main target implementation language is Java. The test application is counting some words from an input file and is not particularly interesting, but gives a simple and comprehensible introduction to the MapReduce concept. More interesting is the Monte Carlo algorithm's approach described in the sixth chapter. The good thing is that it is compared side by side with analogous C program for MPI. This actually is the culmination and the essence of the book. For further investigations of the concepts and ways to apply the parallelism in practice, help is available online. Many resources are given in the appendix. The last chapter is quite handy in general and beyond the scope of the book. The instructions for booting the Raspberry Pi with an external USB HDD as an auxiliary data storage seem very useful. The building of LEGO case for the cluster, and the suggestions for alternative energy sources give interesting views to Raspberry Pi on their own. All in all setting up a cluster from Raspberry Pi units is shown to be not so complex as expected. Only the correct set of steps should be followed, and followed strictly at times. If not giving a hint to certain project, this book at least puts you in a firm starting position on the road to parallelism.

I have a number of Raspberry Pi's sitting around from various projects I have experimented with in the past. I came across this book on building a cluster using two Raspberry Pi's. The book guides you through MPI (Message Passing Interface) and Hadoop and provides novices with a good guide for getting started. As well as software development aspects there are some fun little side projects like building your own stackable case using Lego's. There is also an interesting introduction to super computing to give you a background to the field. Overall a great introduction to parallel computing aimed at the novice.

I could describe this book as a really nice introduction to create your own Raspberry Pi cluster, consider this an offline compiled tutorial to do it yourself with some code that would give some insights to start your own apps. Little Java or C++ knowledge is necessary, but a lot of self enthusiasm for have it working. It is easy to read and it will walk you through step by step, but in order to bring all what it is exposed within, it is necessary to have the hardware in place (at least two Raspberry Pi). I liked that the read is smooth, well explained and has a lot of e-references to extends on the contents provided, on the other hand, It would be a better book with more images and, for really tech people, more code to play with. This is the main reason to skip the 5 stars. I read it in a day, but if you follow the guidelines step by step, it can take you at least a whole weekend to have everything up and running. Overall a good read for your sunday night,

I've seen several YouTube videos showing clusters of Raspberry Pi's and, as I have a few Raspberry Pi's, thought I would have a go at building my cluster. Looking for information on how to do it, I came across this book which provides a great introduction into parallel computing and also provided very clear, concise instructions to set up a cluster of 2 Raspberry Pi's. I love reading the back story to the things I learn and the first chapter provides a really good history to parallel/distributed computing. Setting up the Raspberry Pi's is covered in the second chapter and is done well and without going into the finite detail of setting up a Pi that is covered (and repeated!) in so many other places. In the third chapter, the process to setting up MPI on the Raspberry Pi's is covered which enables the Raspberry Pi's to be connected in a parallel computing environment. This is then taken forward into chapter 4 and where we set up Hadoop and MapReduce. Hadoop enables distributed applications to be written and MapReduce is intended to enable systems to process large datasets. In setting up MPI, Hadoop and MapReduce, simple applications are written but the book then brings this all together by writing an application to calculate pi using Hadoop, and then the same in MPI to compare the two technologies. Finally, the book provides some very useful

information on how to take things further. Overall, I thoroughly enjoyed reading and using this book to take me into an area of computing I've never delved into. I thought the book was concise, easy to read and the examples were clear and easy to follow and I'll certainly be keen to read Andrew K. Dennis' future books.

[Download to continue reading...](#)

Raspberry Pi 3: 2016 Raspberry Pi 3 User Guide (Raspberry Pi, Raspberry Pi 2, Raspberry Pi Programming, Raspberry Pi Projects) (Volume 1) Raspberry Pi Super Cluster Linux Enterprise Cluster: Build a Highly Available Cluster with Commodity Hardware and Free Software Additional Aix Security Tools on IBM Elogo Server Pseries, IBM Rs/6000, and Sp/Cluster Cluster Mystery: Epidemic and the Children of Woburn, Mass. Unsupervised Machine Learning in Python: Master Data Science and Machine Learning with Cluster Analysis, Gaussian Mixture Models, and Principal Components Analysis Super Simple Jewelry: Fun and Easy-To-Make Crafts for Kids (Super Sandcastle: Super Simple Crafts) Kindergarten Super Math Success (Sylvan Super Workbooks) (Math Super Workbooks) Fourth Grade Super Math Success (Sylvan Super Workbooks) (Math Super Workbooks) First Grade Super Math Success (Sylvan Super Workbooks) (Math Super Workbooks) First Grade Super Math Success (Sylvan Super Workbooks) (Sylvan Math Super Workbooks) Kindergarten Super Math Success (Sylvan Super Workbooks) (Sylvan Math Super Workbooks) Third Grade Super Math Success (Sylvan Super Workbooks) (Sylvan Math Super Workbooks) Fifth Grade Super Math Success (Sylvan Super Workbooks) (Sylvan Math Super Workbooks) Raspberry Pi For Kids For Dummies Adventures in Raspberry Pi Getting Started with Raspberry Pi: Electronic Projects with Python, Scratch, and Linux Getting Started With Raspberry Pi: An Introduction to the Fastest-Selling Computer in the World JavaScript Robotics: Building NodeBots with Johnny-Five, Raspberry Pi, Arduino, and BeagleBone (Make) Smart Home Automation with Linux and Raspberry Pi

[Dmca](#)