

File Systems 1 Github Pages

Contemporary High Performance Computing: From Petascale toward Exascale, Volume 3 focuses on the ecosystems surrounding the world's leading centers for high performance computing (HPC). It covers many of the important factors involved in each ecosystem: computer architectures, software, applications, facilities, and sponsors. This third volume will be a continuation of the two previous volumes, and will include other HPC ecosystems using the same chapter outline: description of a flagship system, major application workloads, facilities, and sponsors. Features: Describes many prominent, international systems in HPC from 2015 through 2017 including each system's hardware and software architecture Covers facilities for each system including power and cooling Presents application workloads for each site Discusses historic and projected trends in technology and applications Includes contributions from leading experts Designed for researchers and students in high performance computing, computational science, and related areas, this book provides a valuable guide to the state-of-the art research, trends, and resources in the world of HPC.

HPC, Big Data, AI Convergence Towards Exascale provides an updated vision on the most advanced computing, storage, and interconnection technologies, that are at basis of convergence among the HPC, Cloud, Big Data, and artificial intelligence (AI) domains. Through the presentation of the solutions devised within recently founded H2020 European projects, this book provides an insight on challenges faced by integrating such technologies and in achieving performance and energy efficiency targets towards the exascale level. Emphasis is given to innovative ways of provisioning and managing resources, as well as monitoring their usage.

Industrial and scientific use cases give to the reader practical examples of the needs for a cross-domain convergence. All the chapters in this book pave the road to new generation of technologies, support their development and, in addition, verify them on real-world problems. The readers will find this book useful because it provides an overview of currently available technologies that fit with the concept of unified Cloud-HPC-Big Data-AI applications and presents examples of their actual use in scientific and industrial applications.

This book shares original innovations, research, and lessons learned regarding teaching and technological perspectives on trust-based learning systems. Both perspectives are crucial to enhancing the e-Assessment process. In the course of the book, diverse areas of the computer sciences (machine learning, biometric recognition, cloud computing, and learning analytics, amongst others) are addressed. In addition, current trends, privacy, ethical issues, technological solutions, and adaptive educational models are described to provide readers with a global view on the state of the art, the latest challenges, and potential solutions in e-Assessment. As such, the book offers a valuable reference guide for industry, educational institutions, researchers, developers, and practitioners seeking to promote e-Assessment processes.

Unlike other books, courses and training that expect an analyst to piece together individual instructions into a cohesive investigation, Investigating Windows Systems provides a walk-through of the analysis process, with descriptions of the thought

process and analysis decisions along the way. Investigating Windows Systems will not address topics which have been covered in other books, but will expect the reader to have some ability to discover the detailed usage of tools and to perform their own research. The focus of this volume is to provide a walk-through of the analysis process, with descriptions of the thought process and the analysis decisions made along the way. A must-have guide for those in the field of digital forensic analysis and incident response. Provides the reader with a detailed walk-through of the analysis process, with decision points along the way, assisting the user in understanding the resulting data Coverage will include malware detection, user activity, and how to set up a testing environment Written at a beginner to intermediate level for anyone engaging in the field of digital forensic analysis and incident response

Building the New Economy

GitHub For Dummies

ICTs for Improving Patients Rehabilitation Research Techniques

Third International Workshop, REHAB 2015, Lisbon, Portugal, October 1-2, 2015, Revised Selected Papers

Linux Kernel Programming

A Non-Technical Guide

Learn WebAssembly

Build Android apps using the popular and efficient Android Studio 3 suite of tools, an integrated development environment (IDE) with which Android developers can now use the Kotlin programming language. With this book, you'll learn the latest and most productive tools in the Android tools ecosystem, ensuring quick Android app development and minimal effort on your part. Along the way, you'll use Android Studio to develop apps tier by tier through practical examples. These examples cover core Android topics such as Activities, Intents, BroadcastReceivers, Services and AsyncTask. Then, you'll learn how to publish your apps and sell them online and in the Google Play store. What You'll Learn Use Android Studio 3 to quickly and confidently build your first Android apps Build an Android user interface using activities and layouts, event handling, images, menus and the action bar Incorporate new elements including fragments Learn how data is persisted Use Kotlin to build apps Who This Book Is For Those who may be new to Android Studio 3 or Android Studio in general. You may or may not be new to Android development in general. Some prior experience with Java is also recommended.

Introduction to Data Science: Data Analysis and Prediction

Algorithms with R introduces concepts and skills that can help you tackle real-world data analysis challenges. It covers concepts from probability, statistical inference, linear regression, and machine learning. It also helps you develop skills such as R programming, data wrangling, data

visualization, predictive algorithm building, file organization with UNIX/Linux shell, version control with Git and GitHub, and reproducible document preparation. This book is a textbook for a first course in data science. No previous knowledge of R is necessary, although some experience with programming may be helpful. The book is divided into six parts: R, data visualization, statistics with R, data wrangling, machine learning, and productivity tools. Each part has several chapters meant to be presented as one lecture. The author uses motivating case studies that realistically mimic a data scientist's experience. He starts by asking specific questions and answers these through data analysis so concepts are learned as a means to answering the questions. Examples of the case studies included are: US murder rates by state, self-reported student heights, trends in world health and economics, the impact of vaccines on infectious disease rates, the financial crisis of 2007-2008, election forecasting, building a baseball team, image processing of hand-written digits, and movie recommendation systems. The statistical concepts used to answer the case study questions are only briefly introduced, so complementing with a probability and statistics textbook is highly recommended for in-depth understanding of these concepts. If you read and understand the chapters and complete the exercises, you will be prepared to learn the more advanced concepts and skills needed to become an expert.

How to empower people and communities with user-centric data ownership, transparent and accountable algorithms, and secure digital transaction systems. Data is now central to the economy, government, and health systems—so why are data and the AI systems that interpret the data in the hands of so few people? Building the New Economy calls for us to reinvent the ways that data and artificial intelligence are used in civic and government systems. Arguing that we need to think about data as a new type of capital, the authors show that the use of data trusts and distributed ledgers can empower people and communities with user-centric data ownership, transparent and accountable algorithms, machine learning fairness principles and methodologies, and secure digital transaction systems. It's well known that social media generate disinformation and that mobile phone tracking apps threaten privacy. But these same technologies may also enable the creation of more agile systems in which power and decision-making are distributed among stakeholders rather than concentrated in a few hands. Offering both big ideas and detailed blueprints, the authors describe such key building blocks as data cooperatives, tokenized funding mechanisms, and tradecoin architecture. They also discuss

technical issues, including how to build an ecosystem of trusted data, the implementation of digital currencies, and interoperability, and consider the evolution of computational law systems.

Pro Git (Second Edition) is your fully-updated guide to Git and its usage in the modern world. Git has come a long way since it was first developed by Linus Torvalds for Linux kernel development. It has taken the open source world by storm since its inception in 2005, and this book teaches you how to use it like a pro. Effective and well-implemented version control is a necessity for successful web projects, whether large or small. With this book you'll learn how to master the world of distributed version workflow, use the distributed features of Git to the full, and extend Git to meet your every need. Written by Git pros Scott Chacon and Ben Straub, Pro Git (Second Edition) builds on the hugely successful first edition, and is now fully updated for Git version 2.0, as well as including an indispensable chapter on GitHub. It's the best book for all your Git needs.

Technologies and Applications for Big Data Value

15th International Conference, PROPOR 2022, Fortaleza, Brazil, March 21–23, 2022, Proceedings

Information and Communication Technology for Intelligent Systems
A comprehensive guide to kernel internals, writing kernel modules, and kernel synchronization

A Practical Guide to Methods and Tools

Efficient Android App Development

HPC, Big Data, and AI Convergence Towards Exascale

Version control is a widely used practice among software developers. It reduces the risk of changing their software and allows them to manage different configurations and to collaborate with others more efficiently. This is amplified by code sharing platforms such as GitHub or Bitbucket. Most version control systems track files (e.g., Git, Mercurial, and Subversion do), but some programming environments do not operate on files, but on objects instead (many Smalltalk implementations do). Users of such environments want to use version control for their objects anyway. Specialized version control systems, such as the ones available for Smalltalk systems (e.g., ENVY/Developer and Monticello), focus on a small subset of objects that can be versioned. Most of these systems concentrate on the tracking of methods, classes, and configurations of these. Other user-defined and user-built objects are either not eligible for version control at all, tracking them involves complicated workarounds, or a fixed, domain-unspecific serialization format is used that does not equally suit all kinds of objects. Moreover, these version control systems that are specific to a programming environment require their own code sharing platforms; popular, well-established platforms for file-based version control systems cannot be used or

adapter solutions need to be implemented and maintained. To improve the situation for version control of arbitrary objects, a framework for tracking, converting, and storing of objects is presented in this report. It allows editions of objects to be stored in an exchangeable, existing backend version control system. The platforms of the backend version control system can thus be reused. Users and objects have control over how objects are captured for the purpose of version control. Domain-specific requirements can be implemented. The storage format (i.e. the file format, when file-based backend version control systems are used) can also vary from one object to another. Different editions of objects can be compared and sets of changes can be applied to graphs of objects. A generic way for capturing and restoring that supports most kinds of objects is described. It models each object as a collection of slots. Thus, users can begin to track their objects without first having to implement version control supplements for their own kinds of objects. The proposed architecture is evaluated using a prototype implementation that can be used to track objects in Squeak/Smalltalk with Git. The prototype improves the suboptimal standing of user objects with respect to version control described above and also simplifies some version control tasks for classes and methods as well. It also raises new problems, which are discussed in this report as well.

The three-volume set LNCS 12681-12683 constitutes the proceedings of the 26th International Conference on Database Systems for Advanced Applications, DASFAA 2021, held in Taipei, Taiwan, in April 2021. The total of 156 papers presented in this three-volume set was carefully reviewed and selected from 490 submissions. The topic areas for the selected papers include information retrieval, search and recommendation techniques; RDF, knowledge graphs, semantic web, and knowledge management; and spatial, temporal, sequence, and streaming data management, while the dominant keywords are network, recommendation, graph, learning, and model. These topic areas and keywords shed the light on the direction where the research in DASFAA is moving towards. Due to the Corona pandemic this event was held virtually.

In the first definitive guide on WebAssembly, you 'll learn how you can wield this new technology to break through the current barriers of web development and build an entirely new class of performant applications . Key FeaturesGenerate WebAssembly modules from C and C++ using Emscripten and interact with these modules in the browser Learn how to use WebAssembly outside of the browser and load modules using Node.js Build a high-performance application using C and WebAssembly and port an existing C++ game to WebAssembly using Emscripten Book Description WebAssembly is a brand-new technology that represents a paradigm shift in web development. This book teaches programmers to leverage this technology to write high-performance applications that run in the browser. This book introduces you to powerful WebAssembly concepts to help you write lean and powerful web applications with native performance. You start with the evolution of web programming, the state of things

today, and what can be done with the advent and release of WebAssembly. We take a look at the journey from JavaScript to asm.js to WebAssembly. We then move on to analyze the anatomy of a WebAssembly module and the relationship between binary and text formats, along with the corresponding JavaScript API. Further on, you'll implement all the techniques you've learned to build a high-performance application using C and WebAssembly, and then port an existing game written in C++ to WebAssembly using Emscripten. By the end of this book, you will be well-equipped to create high-performance applications and games for the web using WebAssembly. What you will learn

Learn how WebAssembly came to be and its associated elements (text format, module, and JavaScript API)
Create, load, and debug a WebAssembly module (editor and compiler/toolchain)
Build a high-performance application using C and WebAssembly
Extend WebAssembly 's feature set using Emscripten by porting a game written in C++
Explore upcoming features of WebAssembly, Node.js integration, and alternative compilation methods
Who this book is for
If you are a web developer or C/C++ programmer keen to leverage the powerful technology of WebAssembly to build high-performance web applications, then this book is for you.

This book constitutes the proceedings of the 14th International Conference on Quantitative Evaluation Systems, QEST 2017, held in Berlin, Germany, in September 2017. The 20 full papers and 4 tool papers presented were carefully reviewed and selected from 58 submissions. The papers are organized in topical sections entitled: probabilistic modeling; smart energy systems over the cloud; Petri nets and performance modeling; parametric verification; machine learning and formal methods; tools.

4th International Conference, ICCCS 2018, Haikou, China, June 8-10, 2018, Revised Selected Papers, Part II

Cloud Computing and Security

Mobile Forensics -- the File Format Handbook

ISC High Performance Digital 2021 International Workshops, Frankfurt am Main, Germany, June 24 - July 2, 2021, Revised Selected Papers

Programming Languages and Systems

Data as Capital

Proceedings of the Fifth HPI Cloud Symposium "Operating the Cloud" 2017

Cloud Storage Security: A Practical Guide introduces and discusses the risks associated with cloud-based data storage from a security and privacy perspective. Gain an in-depth understanding of the risks and benefits of cloud storage illustrated using a Use-Case methodology. The authors also provide a checklist that enables the user, as well as the enterprise practitioner to evaluate what security and privacy issues need to be considered when using the cloud to store personal and sensitive information. Describes the history and the evolving nature of cloud storage and security
Explores the threats to privacy and security when using free social media applications that use cloud storage
Covers legal issues and laws that govern privacy, compliance, and legal responsibility for enterprise users
Provides guidelines and a security checklist for selecting a cloud-storage service

provider Includes case studies and best practices for securing data in the cloud Discusses the future of cloud computing

This six volume set LNCS 11063 – 11068 constitutes the thoroughly refereed conference proceedings of the 4th International Conference on Cloud Computing and Security, ICCCS 2018, held in Haikou, China, in June 2018. The 386 full papers of these six volumes were carefully reviewed and selected from 1743 submissions. The papers cover ideas and achievements in the theory and practice of all areas of inventive systems which includes control, artificial intelligence, automation systems, computing systems, electrical and informative systems. The six volumes are arranged according to the subject areas as follows: cloud computing, cloud security, encryption, information hiding, IoT security, multimedia forensics.

This open access book constitutes the refereed proceedings of the 6th Asian Supercomputing Conference, SCFA 2020, which was planned to be held in February 2020, but unfortunately, the physical conference was cancelled due to the COVID-19 pandemic. The 8 full papers presented in this book were carefully reviewed and selected from 22 submissions. They cover a range of topics including file systems, memory hierarchy, HPC cloud platform, container image configuration workflow, large-scale applications, and scheduling.

Covering research topics from system software such as programming languages, compilers, runtime systems, operating systems, communication middleware, and large-scale file systems, as well as application development support software and big-data processing software, this book presents cutting-edge software technologies for extreme scale computing. The findings presented here will provide researchers in these fields with important insights for the further development of exascale computing technologies. This book grew out of the post-peta CREST research project funded by the Japan Science and Technology Agency, the goal of which was to establish software technologies for exploring extreme performance computing beyond petascale computing. The respective were contributed by 14 research teams involved in the project. In addition to advanced technologies for large-scale numerical computation, the project addressed the technologies required for big data and graph processing, the complexity of memory hierarchy, and the power problem. Mapping the direction of future high-performance computing was also a central priority.

Investigating Windows Systems

How to take over any company in the world

26th International Conference, DASFAA 2021, Taipei, Taiwan, April 11 – 14, 2021, Proceedings, Part I

A Practical Guide

Data Analysis and Prediction Algorithms with R

Creating Blogs with Jekyll

Learn Android Studio 3 with Kotlin

Software is eating the world, and GitHub is where software is built. GitHub is also a powerful way for people to collaborate on text-based documents, from contracts to screenplays to legislation. With this introductory guide, you'll learn how to use GitHub to manage and collaborate with developers, designers and other business professionals more effectively. Topics include project transparency, collaboration tools, the basics of Git version control management and how to make changes yourself - without having to bother your development team.

This book provides the users with quick and easy data acquisition, processing, storage and product generation services. It describes the entire life cycle of

remote sensing data and builds an entire high performance remote sensing data processing system framework. It also develops a series of remote sensing data management and processing standards. Features: Covers remote sensing cloud computing Covers remote sensing data integration across distributed data centers Covers cloud storage based remote sensing data share service Covers high performance remote sensing data processing Covers distributed remote sensing products analysis

This book constitutes the refereed post-conference proceedings of 9 workshops held at the 35th International ISC High Performance 2021 Conference, in Frankfurt, Germany, in June-July 2021: Second International Workshop on the Application of Machine Learning Techniques to Computational Fluid Dynamics and Solid Mechanics Simulations and Analysis; HPC-IODC: HPC I/O in the Data Center Workshop; Compiler-assisted Correctness Checking and Performance Optimization for HPC; Machine Learning on HPC Systems; 4th International Workshop on Interoperability of Supercomputing and Cloud Technologies; 2nd International Workshop on Monitoring and Operational Data Analytics; 16th Workshop on Virtualization in High-Performance Cloud Computing; Deep Learning on Supercomputers; 5th International Workshop on In Situ Visualization. The 35 papers included in this volume were carefully reviewed and selected. They cover all aspects of research, development, and application of large-scale, high performance experimental and commercial systems. Topics include high-performance computing (HPC), computer architecture and hardware, programming models, system software, performance analysis and modeling, compiler analysis and optimization techniques, software sustainability, scientific applications, deep learning.

This book gathers the carefully reviewed proceedings of the 19th International Conference on Systems Science, presenting recent research findings in the areas of Artificial Intelligence, Machine Learning, Communication/Networking and Information Technology, Control Theory, Decision Support, Image Processing and Computer Vision, Optimization Techniques, Pattern Recognition, Robotics, Service Science, Web-based Services, Uncertain Systems and Transportation Systems. The International Conference on Systems Science was held in Wroclaw, Poland from September 7 to 9, 2016, and addressed a range of topics, including systems theory, control theory, machine learning, artificial intelligence, signal processing, communication and information technologies, transportation systems, multi-robotic systems and uncertain systems, as well as their applications. The aim of the conference is to provide a platform for communication between young and established researchers and practitioners, fostering future joint research in systems science.

Electronic Imaging & the Visual Arts EVA 2019 Florence

High Performance Computing

Cloud Computing in Remote Sensing

14th International Conference, QEST 2017, Berlin, Germany, September 5-7, 2017, Proceedings

Contemporary High Performance Computing

Proceedings of ICTIS 2020, Volume 1

Designing Data-Intensive Applications

Learn how to write high-quality kernel module code, solve common Linux kernel programming issues, and understand the fundamentals of Linux kernel internals

Key Features Discover how to write kernel code using the Loadable Kernel Module framework Explore industry-grade techniques to perform efficient memory allocation and data synchronization within the kernel Understand the essentials of key internals topics such as kernel architecture, memory management, CPU scheduling, and kernel synchronization

Book Description Linux Kernel Programming is a comprehensive introduction for those new to Linux kernel and module development. This easy-to-follow guide will have you up and running with writing kernel code in next-to-no time. This book uses the latest 5.4 Long-Term Support (LTS) Linux kernel, which will be maintained from November 2019 through to December 2025. By working with the 5.4 LTS kernel throughout the book, you can be confident that your knowledge will continue to be valid for years to come. You'll start the journey by learning how to build the kernel from the source. Next, you'll write your first kernel module using the powerful Loadable Kernel Module (LKM) framework. The following chapters will cover key kernel internals topics including Linux kernel architecture, memory management, and CPU scheduling. During the course of this book, you'll delve into the fairly complex topic of concurrency within the kernel, understand the issues it can cause, and learn how they can be addressed with various locking technologies (mutexes, spinlocks, atomic, and refcount operators). You'll also benefit from more advanced material on cache effects, a primer on lock-free techniques within the kernel, deadlock avoidance (with lockdep), and kernel lock debugging techniques. By the end of this kernel book, you'll have a detailed understanding of the fundamentals of writing Linux kernel module code for real-world projects and products. What you will learn

Write high-quality modular kernel code (LKM framework) for 5.x kernels

Configure and build a kernel from source

Explore the Linux kernel architecture

Get to grips with key internals regarding memory management within the kernel

Understand and work with various dynamic kernel memory alloc/dealloc APIs

Discover key internals aspects regarding CPU scheduling within the kernel

Gain an understanding of kernel concurrency issues

Find out how to work with key kernel synchronization primitives

Who this book is for This book is for Linux programmers beginning to find their way with Linux kernel development. If you're a Linux kernel and driver developer looking to overcome frequent and common kernel development issues, or understand kernel internals, you'll find plenty of useful information. You'll need a solid foundation of Linux CLI and C programming before you can jump in.

The Art of Network Penetration Testing is a guide to simulating an internal security breach. You ' ll take on the role of the attacker and work through every stage of a professional pentest, from information gathering to seizing control of a system and owning the network. Summary

Penetration testing is about more than just getting through a perimeter firewall. The biggest security threats are inside the network, where attackers can rampage through sensitive data by exploiting weak access controls and poorly patched software. Designed for up-and-coming security professionals, The Art of Network Penetration Testing teaches you how to take over an enterprise network from the inside. It lays out every stage of an internal security assessment step-by-step, showing you how to identify weaknesses before a malicious invader can do real damage. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the technology

Penetration testers uncover security gaps by attacking networks exactly like malicious intruders do. To become a world-class pentester, you need to master offensive security concepts, leverage a proven methodology, and practice, practice, practice. Th is book delivers insights from security expert Royce Davis, along with a virtual testing environment you can use to hone your skills. About the book

The Art of Network Penetration Testing is a guide to simulating an internal security breach. You ' ll take on the role of the attacker and work through every stage of a professional

pentest, from information gathering to seizing control of a system and owning the network. As you brute force passwords, exploit unpatched services, and elevate network level privileges, you'll learn where the weaknesses are—and how to take advantage of them. What's inside Set up a virtual pentest lab Exploit Windows and Linux network vulnerabilities Establish persistent re-entry to compromised targets Detail your findings in an engagement report About the reader For tech professionals. No security experience required. About the author Royce Davis has orchestrated hundreds of penetration tests, helping to secure many of the largest companies in the world. Table of Contents 1 Network Penetration Testing PHASE 1 - INFORMATION GATHERING 2 Discovering network hosts 3 Discovering network services 4 Discovering network vulnerabilities PHASE 2 - FOCUSED PENETRATION 5 Attacking vulnerable web services 6 Attacking vulnerable database services 7 Attacking unpatched services PHASE 3 - POST-EXPLOITATION AND PRIVILEGE ESCALATION 8 Windows post-exploitation 9 Linux or UNIX post-exploitation 10 Controlling the entire network PHASE 4 - DOCUMENTATION 11 Post-engagement cleanup 12 Writing a solid pentest deliverable

This two-volume book constitutes the post-conference proceedings of the 5th International Conference on Advances in Computing and Data Sciences, ICACDS 2021, held in Nashik, India, in April 2021.* The 103 full papers were carefully reviewed and selected from 781 submissions. Part II is devoted to data sciences, organizing principles, medical technologies, computational linguistics etc. *The conference was held virtually due to the COVID-19 pandemic.

This is the first monograph dedicated to this interdisciplinary research area, combining the views of music, computer science, education, creativity studies, psychology, and engineering. The contributions include introductions to ubiquitous music research, featuring theory, applications, and technological development, and descriptions of permanent community initiatives such as virtual forums, multi-institutional research projects, and collaborative publications. The book will be of value to researchers and educators in all domains engaged with creativity, computing, music, and digital arts.

Handbook of Research on Big Data Storage and Visualization Techniques

Advances in Computing and Data Sciences

6th Asian Conference, SCFA 2020, Singapore, February 24-27, 2020, Proceedings

Cloud Storage Security

The Big Ideas Behind Reliable, Scalable, and Maintainable Systems

Introducing GitHub

Advances in Systems Science

Learn to create your own blog using the Jekyll static site generator. You'll start with a simple template, add new features to it, automate any maintenance, attach social sharing, and begin writing. By the end of *Creating Blogs with Jekyll*, you will be able to create custom blogs with Jekyll, update the content with ease, and reach out to your readers with minimal effort. Because you've built your blog yourself, you'll know exactly how each component works, and you won't be dependent on an admin panel to maintain it. *Creating Blogs with Jekyll* equips you with the knowledge to create an elegantly designed blog and scale it to capture more readers. Recapture the magic of writing by creating great content and use an easy workflow in Jekyll to maintain it for blogging. Do new things and write about them in style with Jekyll. Takes you through building a fully functional blog from scratch using Jekyll Provides a fun way to work on a side-project and integrate cutting edge web technologies Teaches you how to update and maintain your awesome blog Jekyll is a simple, secure and very low maintenance blog engine that converts naturally written content in markdown into a beautiful and minimal blog. It allows you to focus on content creation and expressing yourself instead of spending all your time updating the plugins and maintaining the database. Jekyll does not rely on a database as a backend so

your blog will be far more secure and reliable than any traditional blogging engines such as WordPress. We live in a day and age where short attention spans make it very difficult to expose a reader to interesting content. What better way to capture a reader's attention and retain viewers by captivating them by your own unique style and taste? Jekyll allows the content to shine with minimal distractions and a greater focus on the content and easy sharing of the content. What You'll Learn Choose a base theme appropriate for your style and development Integrate various web technologies that will work well together and enhance your blog Automate social sharing components and comments workflow Make adjustments to themes, views and styles of blog posts Update any of the modular components of the blog and integrate new technologies Implement Jekyll and deploying static websites for future projects Who This Book Is For The developer who is ready to move beyond the complexities of maintaining a content management system by creating their own unique blog in their own style. It's for the project manager tired of spending all their time editing their blog on the admin panel and updating the content management system. Creating Blogs with Jekyll is an excellent choice for new developers to start blogging because of the simplicity of Jekyll's theming layer and writing workflow. It's an excellent choice for the web developer wanting to build their blog from scratch and expand their knowledge of higher level web technologies.

Both Traditional Students and Working Professionals Acquire the Skills to Analyze Social Problems. *Big Data and Social Science: A Practical Guide to Methods and Tools* shows how to apply data science to real-world problems in both research and the practice. The book provides practical guidance on combining methods and tools from computer science, statistics, and social science. This concrete approach is illustrated throughout using an important national problem, the quantitative study of innovation. The text draws on the expertise of prominent leaders in statistics, the social sciences, data science, and computer science to teach students how to use modern social science research principles as well as the best analytical and computational tools. It uses a real-world challenge to introduce how these tools are used to identify and capture appropriate data, apply data science models and tools to that data, and recognize and respond to data errors and limitations. For more information, including sample chapters and news, please visit the author's website.

This book gathers papers addressing state-of-the-art research in all areas of information and communication technologies and their applications in intelligent computing, cloud storage, data mining and software analysis. It presents the outcomes of the Fourth International Conference on Information and Communication Technology for Intelligent Systems, which was held in Ahmedabad, India. Divided into two volumes, the book discusses the fundamentals of various data analysis techniques and algorithms, making it a valuable resource for researchers and practitioners alike.

This book constitutes the thoroughly refereed proceedings of the Third International Workshop on ICTs for Improving Patients Rehabilitation Research Techniques, REHAB 2015, held in Lisbon, Portugal, in October 2015. The 15 revised full papers presented were carefully reviewed and selected from 59 submissions. The papers explore how technology can contribute toward smarter and effective rehabilitation methods.

Supercomputing Frontiers

Proceedings of the International Conference on Systems Science 2016 (ICSS 2016)

From Petascale toward Exascale, Volume 3

The Art of Network Penetration Testing

Engineering Data-Driven Adaptive Trust-based e-Assessment Systems

Challenges and Infrastructure Solutions

Challenge and Vision

In today's data driven biology, programming knowledge is essential in turning ideas into testable hypothesis. Based on the author's extensive experience, Python for Bioinformatics, Second Edition helps biologists get to grips with the basics of software development. Requiring no prior knowledge of programming-related concepts, the book focuses on the easy-to-use, yet powerful, Python computer language. This new edition is updated throughout to Python 3 and is designed not just to help scientists master the basics, but to do more in less time and in a reproducible way. New developments added in this edition include NoSQL databases, the Anaconda Python distribution, graphical libraries like Bokeh, and the use of Github for collaborative development. Leverage the power of Git to smooth out the development cycle Professional Git takes a professional approach to learning this massively popular software development tool, and provides an up-to-date guide for new users. More than just a development manual, this book helps you get into the Git mindset—extensive discussion of corollaries to traditional systems as well as considerations unique to Git help you draw upon existing skills while looking out—and planning for—the differences. Connected labs and exercises are interspersed at key points to reinforce important concepts and deepen your understanding, and a focus on the practical goes beyond technical tutorials to help you integrate the Git model into your real-world workflow. Git greatly simplifies the software development cycle, enabling users to create, use, and switch between versions as easily as you switch between files. This book shows you how to harness that power and flexibility to streamline your development cycle. Understand the basic Git model and overall workflow Learn the Git versions of common source management concepts and commands Track changes, work with branches, and take advantage of Git's full functionality Avoid trip-ups and missteps common to new users Git works with the most popular software development tools and is used by almost all of the major technology companies. More than 40 percent of software developers use it as their primary source control tool, and that number continues to grow; the ability to work effectively with Git is rapidly approaching must-have status, and Professional Git is the comprehensive

guide you need to get up to speed quickly.

The digital age has presented an exponential growth in the amount of data available to individuals looking to draw conclusions based on given or collected information across industries. Challenges associated with the analysis, security, sharing, storage, and visualization of large and complex data sets continue to plague data scientists and analysts alike as traditional data processing applications struggle to adequately manage big data. The Handbook of Research on Big Data Storage and Visualization Techniques is a critical scholarly resource that explores big data analytics and technologies and their role in developing a broad understanding of issues pertaining to the use of big data in multidisciplinary fields. Featuring coverage on a broad range of topics, such as architecture patterns, programing systems, and computational energy, this publication is geared towards professionals, researchers, and students seeking current research and application topics on the subject.

Data is at the center of many challenges in system design today. Difficult issues need to be figured out, such as scalability, consistency, reliability, efficiency, and maintainability. In addition, we have an overwhelming variety of tools, including relational databases, NoSQL datastores, stream or batch processors, and message brokers. What are the right choices for your application? How do you make sense of all these buzzwords? In this practical and comprehensive guide, author Martin Kleppmann helps you navigate this diverse landscape by examining the pros and cons of various technologies for processing and storing data. Software keeps changing, but the fundamental principles remain the same. With this book, software engineers and architects will learn how to apply those ideas in practice, and how to make full use of data in modern applications. Peer under the hood of the systems you already use, and learn how to use and operate them more effectively. Make informed decisions by identifying the strengths and weaknesses of different tools. Navigate the trade-offs around consistency, scalability, fault tolerance, and complexity. Understand the distributed systems research upon which modern databases are built. Peek behind the scenes of major online services, and learn from their architectures. Common File Formats and File Systems Used in Mobile Devices

Quantitative Evaluation of Systems

Build web applications with native performance using Wasm and C/C++

Professional Git

27th European Symposium on Programming, ESOP 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14–20, 2018, Proceedings

5th International Conference, ICACDS 2021, Nashik, India, April 23–24, 2021, Revised Selected Papers, Part II

Computational Processing of the Portuguese Language

Code collaboratively with GitHub Once you've learned the basics of coding the next step is to start sharing your expertise, learning from other coding pros, or working as a collaborative member of development teams. GitHub is the go-to community for facilitating coding collaboration, and GitHub For Dummies is the next step on your journey as a developer. Written by a GitHub engineer, this book is packed with insight on how GitHub works and how you can use it to become a more effective, efficient, and valuable member of any collaborative programming team. Store and share your work online with GitHub Collaborate with others on your team or across the international coding community Embrace open-source values and processes Establish yourself as a valuable member of the GitHub community From setting up GitHub on your desktop and launching your first project to cloning repositories, finding useful apps on the marketplace, and improving workflow, GitHub For Dummies covers the essentials the novice programmer needs to enhance collaboration and teamwork with this industry-standard tool.

This open access book constitutes the proceedings of the 27th European Symposium on Programming, ESOP 2018, which took place in Thessaloniki, Greece in April 2018, held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018. The 36 papers presented in this volume were carefully reviewed and selected from 114 submissions. The papers are organized in topical sections named: language design; probabilistic programming; types and effects; concurrency; security; program verification; program analysis and automated verification; session types and concurrency; concurrency and distribution; and compiler verification.

If you're new to GitHub, this concise book shows you just what you need to get started and no more. It's perfect for project and product managers, stakeholders, and other team members who want to collaborate on a development project—whether it's to review and comment on work in progress or to contribute specific changes. It's also great for developers just learning GitHub. GitHub has rapidly become the default platform for software development, but it's also ideal for other text-based documents, from contracts to screenplays. This hands-on book shows you how to use GitHub's web interface to view projects and collaborate effectively with your team. Learn how and why people use GitHub to collaborate View the status of a project—recent changes, outstanding work, and historic changes Create and edit files through GitHub without learning Git Suggest changes to projects you don't have permission to edit directly Use tools like issues, pull requests, and branches to specify and collaborate on changes

Create a new GitHub repository to control who has access to your project

Jedes Jahr lädt das Hasso-Plattner-Institut (HPI) Gäste aus der Industrie und der Wissenschaft zu einem kooperativen und wissenschaftlichen Symposium zum Thema Cloud Computing ein. Unser Ziel ist es, ein Forum für den Austausch von Wissen und Erfahrungen zwischen der Industrie und der Wissenschaft zu bieten. Parallel zur Veranstaltung findet der HPI Future SOC Lab Tag statt, der eine zusätzliche attraktive Umgebung für wissenschaftliche und branchenbezogene Diskussionen bietet. Das Symposium zielt darauf ab, eine Plattform für produktive Interaktionen von innovativen Ideen, Visionen und aufkommenden Technologien im Bereich von Cloud Computing zu bieten. Anlässlich dieses Symposiums fordern wir die Einreichung von Forschungsarbeiten und Erfahrungsberichte. Dieser technische Bericht umfasst eine Zusammenstellung der im Rahmen des fünften HPI Cloud Symposiums "Operating the Cloud" 2017 angenommenen Forschungspapiere. Wir danken den Autoren für spannende Vorträge und Einblicke in ihre aktuelle Arbeit und Forschung. Darüber hinaus freuen wir uns auf weitere interessante Einreichungen für das kommende Symposium im Laufe des Jahres. Every year, the Hasso Plattner Institute (HPI) invites guests from industry and academia to a collaborative scientific workshop on the topic Operating the Cloud. Our goal is to provide a forum for the exchange of knowledge and experience between industry and academia. Co-located with the event is the HPI's Future SOC Lab day, which offers an additional attractive and conducive environment for scientific and industry related discussions. Operating the Cloud aims to be a platform for productive interactions of innovative ideas, visions, and upcoming technologies in the field of cloud operation and administration. In these proceedings, the results of the fifth HPI cloud symposium Operating the Cloud 2017 are published. We thank the authors for exciting presentations and insights into their current work and research. Moreover, we look forward to more interesting submissions for the upcoming symposium in 2018.

Advanced Software Technologies for Post-Peta Scale Computing
Database Systems for Advanced Applications
Big Data and Social Science
The Japanese Post-Peta CREST Research Project
Towards version control in object-based systems
Introduction to Data Science
Pro Git

This open access book summarizes knowledge about several file systems and file formats commonly used in mobile devices. In addition to the fundamental description of the formats, there are hints about the forensic value of possible artefacts, along with an outline of tools that can decode the relevant data. The book is organized into two distinct parts. First, Part I describes several different file systems that are commonly used in mobile devices: APFS is the file system that is used in all modern Apple devices including iPhones, iPads, and even Apple Computers, like the MacBook series. Ext4 is very common in Android devices and is the successor of the Ext2 and Ext3 file systems that were commonly used on Linux-based computers. The Flash-Friendly File System (F2FS) is a Linux system

designed explicitly for NAND Flash memory, common in removable storage devices and mobile devices, which Samsung Electronics developed in 2012. The QNX6 file system is present in Smartphones delivered by Blackberry (e.g. devices that are using Blackberry 10) and modern vehicle infotainment systems that use QNX as their operating system. Second, Part II describes five different file formats that are commonly used on mobile devices: SQLite is nearly omnipresent in mobile devices with an overwhelming majority of all mobile applications storing their data in such databases. The second leading file format in the mobile world are Property Lists, which are predominantly found on Apple devices. Java Serialization is a popular technique for storing object states in the Java programming language. Mobile application (app) developers very often resort to this technique to make their application state persistent. The Realm database format has emerged over recent years as a possible successor to the now ageing SQLite format and has begun to appear as part of some modern applications on mobile devices. Protocol Buffers provide a format for taking compiled data and serializing it by turning it into bytes represented in decimal values, which is a technique commonly used in mobile devices. The aim of this book is to act as a knowledge base and reference guide for digital forensic practitioners who need knowledge about a specific file system or file format. It is also hoped to provide useful insight and knowledge for students or other aspiring professionals who want to work within the field of digital forensics. The book is written with the assumption that the reader will have some existing knowledge and understanding about computers, mobile devices, file systems and file formats.

Ubiquitous Music

Python for Bioinformatics