Manoj Raja Rao

CONTACT

manoj@manojrao.com
805-915-9501

TECHNICAL SKILLS

- *Languages*: C, C++, Python, Java, Emacs Lisp, Ruby, Perl, Erlang
- *ML/DL FW*: PyTorch, MXNet, scikit-learn
- *OS Dev*: Linux Kernel Development
- *Tracing*: eBPF, ftrace, systrace, perf
- *Research*: Distributed Systems, Fault Tolerance, Reliability, Message Passing in Clusters
- *Concepts*: Deep Learning, Distributed Systems, Cloud Computing, Computer Architecure, OS
- *Platforms*: Linux, Android, Linux on ARM, EFI, Intel's BIOS, ACPI
- *Dev Tools*: Emacs, Vim, Git, Trace32, GDB, kdb
- *Protocols*: HDMI-CEC, HDCP, I^2C, MHL-CBUS, USB detection, SPI

EDUCATION

UCLA, Los Angeles, USA

MASTER OF SCIENCE, DEPT. OF COMPUTER SCIENCE **2008 – 2010**

- Advisor: Dr. Yuval Tamir
- Research at UCLA's Concurrent Systems Laboratory

RVCE, Bangalore, India

BACHELOR OF ENGG., DEPT. OF COMPUTER SCIENCE&ENGG. **2001 – 2005**

PROFESSIONAL EXPERIENCE

AWS AI, E Palo Alto, USA

Tech Lead, Systems Software for AWS ML and Deep Learning **2019 - Present**

AWS SageMaker Edge Manager: ***2020-2021*** Tech Lead of SageMaker Edge Manager, launched in re:Invent 2020. Designed and Implemented the Deep Learning Runtime and model management for Embedded Devices via AWS SageMaker. Manage, Deploy, and Serve Deep Learning Models via multiple interfaces efficiently implemented in Modern C++ More info here

TorchServe and AWS Personalize ***2019 - 2020*** Lead developer and maintainer of *TorchServe*, the official Deep Learning Model Server for Facebook's popular Deep Learning Framework - PyTorch Link - my whitepaper on AWS Blog Contributed to AWS ML services to enable Predictive Maintenance solution on large scale industrial data. Perform feature extraction on training data. Hands-on with auto-scaling clusters with AWS EMR for massive data processing. Involved in the Link:AWS Personalize core personalization algorithm and inference at AWS scale.

Amazon, Sunnyvale, USA

Senior Software Engineer, Platform Software for Robot **2017 - 2019**

Full stack software development for Robotics project. *Amazon's Robot:* Involved in building low level stack for a robotics project at Amazon. Low-level performance monitoring for sophisticated workloads, board bringups, protoyping PID motor controllers, sensor fusion, compute resource management for multiple CV workloads.

Amazon Lab126, Sunnyvale, USA

Senior Software Engineer, Linux Kernel and Android/FireOS Software **2014 - 2017**

Linux Kernel Development involving embedded systems development for Android/FireOS based devices. Scheduler optimizations, Device Performance, Power Management, Thermal Management, Display Subsystem, Privacy Implementation, Camera, Touch Controllers, Ambient Light Sensors *Amazon's Next Gen Devices:* Involved in product lifecycle from research, prototyping and productizing next generation devices like Alexa, FireTV, Fire Tablets, and more. BSP, board bring-ups, Power and Thermal management, OS-level through UI Performance Engg. Linux Kernel Development involving active embedded systems development for Android/FireOS based devices. The responsibilities include Linux BSP Development for Amazon's next generation devices, Board bringups and prototype development for early stage products, Linux Kernel development involving Amazon specific features, device's performance optimizations for systems with turn-around constraints, experience in early life-cycle product decisions with EE/board designs, experience optimizing device's power consumption under various use cases resulting in long battery life (exceeding product KPIs) and meeting stringent thermal budgets, experience working on firmwares, bootloaders, linux kernel and HALs layers of device software for achieving new board bringup activities for next gen products. Experi-

ence working on time critical activities in factory environment and helping optimize for faster factory diagnostics completion. Expertise with Linux tools such as perf, ftrace, systrace to debug knotty low-level system issues. Experience in all stages of product life-cycle from prototype development through product ship date and beyond.

Qualcomm Innovation Center Inc., San Diego, USA

Senior Software Engineer, Linux Kernel Development for Snapdragon Chipsets **2010 - 2014**

Snapdragon Linux Kernel: Involved in the development of MSM chipsets for Snapdragon's 64-bit CPU architecture. Involved in silicon bring-ups, Linux Kernel Security, Linux Kernel Code Review and Device Tree Code Reviewer prior to upstream process. Device driver Development for various display devices such as Mobile High-Definition Link (MHL), Mobile Display Processor (MDP), Display Serial Interface (DSI). The responsibilites involve active contribution to the Qualcomm MSM open-source tree, development of the device driver for various Linux-based mobile platforms such as Android, Tizen.

UCLA, Los Angeles, USA

Graduate Student Researcher **2008 - 2010**

Fault Tolerance in Distributed Systems: Message Passing Framework for High Availability and Fault Tolerance in Distributed Systems. *HostMon:* A Linux Monitoring System with the ability to monitor and co-ordinate on distributed nodes as well stand-alone systems.

Compiler Construction: Developed a compiler for translating Mini-Java to MIPS assembly in Java as a part of the course I audited at CS Dept., UCLA.

Branch Prediction and Out-of-Order Processors: Analysis of impact of Cache size on performance and Miss-Rates as a part of the course CS251A at CS Dept., UCLA. Analysis and Evaluation of Various types of Branch-Predictors for a low-cost Processor as a part of the course CS251B at CS Dept., UCLA.

UCLA Medical: Research Assistant at the UCLA Medical Center under Dr. Douglas Bell where I implement various types of software ranging from web-based Diabetes study tool, various patient record search and auto-notification applications in Ruby to E-Prescription tools that help physicians to electronically prescribe medicines for patients.

Aylus Networks Pvt Ltd., Bangalore, India

Software Engineer **2006 - 2008**

Developed software modules for the 3G-Telecom Application Server for media share services for 3G users. Developed Service Provisioning System for provisioning users to 3G networks. Developed HA functionality of critical modules.

Huawei Technologies India Pvt Ltd., Bangalore, India

Software Engineer **2005 - 2006**

Worked as Junior Researcher in the R&D for an internal Linux Cluster Middleware for providing Carrier-Grade HA. I was part of the team which applied for patents in Group communication among cluster nodes.

External Links

- TorchServe
- My commits to 3.10 msm Linux Kernel
- My commits to <= 3.4 msm Linux Kernel
- My Technical Blog
- My fledgling podcast series

RELEVANT GRADUATE COURSES

- Distributed Algorithms, Cloud Computing, Operating Systems, Advanced Scalable Systems, Advanced Parallel Systems, Online Algorithms, Advanced System Design, Advanced Computer Architecture, Wireless and Mobile Computing, Cyber Physical Systems

Miscellaneous Recognition

- Patent idea at Huawei Technologies., All India Ranking of 382 among 100000 participants in Entrance Tests., Huawei Certified .C. Programming Specialist.

Membership

- MENSA, FOSS, Computer Society of India