Siddhartha

FPGA/Software Engineer

CONTACT INFORMATION

sidmontu@gmail.com | +65 9857-4171

https://sidmontu.github.io | github.com/sidmontu

INTERESTS — FPGAs, Machine Learning Engineering, HPC, RF Communications

PROFESSIONAL SUMMARY — Experienced mid-career software engineer with a FPGA-focused research background looking for a new challenge. Especially interested in projects at the crossroads of hardware acceleration (FPGAs/ASICs) and modern deep learning applications. Referees available on request.

Work Experience

Machine Learning Engineer (Contract), SAP Singapore — Jun 2021 to Present

- Building enterprise search solutions for internal stakeholders using Elasticsearch/Opensearch stacks.
- Implementing DNN-based semantic-similarity algorithms for information retrieval and ranking of results.
- MLOps tooling: setting up labeling, data cleaning, and evaluation pipelines for efficient data-centric ML implementations.
- Tools/frameworks used (non-exhaustive): [Elastic|Open]search, Docker, FastAPI, Kubernetes, Poetry, Kibana, etc.

CTO/Founder, Inpact Technologies (inPact.ai), Singapore — May 2020 to Sep 2021

- As CTO, planned and executed in-house product development across all aspects of machine learning, frontend, backend, and deployment. Hired freelancers and interns to achieve alpha and **beta** product release milestones within 8 months.
- Frontend using **React-Redux** + **SASS** stack; backend using **serverless framework**; deployment on AWS (EC2, DynamoDB, Cognito, S3, etc); ML models trained using PyTorch and **spaCy**, REST API deployment through **cortex**.
- Trained/fine-tuned DNNs (e.g. residual CNNs, BERT) for NLP tasks like text classification and named entity recognition.
- Raised S$75K from Entrepreneur First, an international VC firm funded by some of the top investors in the world.

Posdoctoral Research Associate, University of Sydney, Australia — Oct 2017 to Dec 2019

- Research & development on a "High-Speed Machine Learning for RF Communications" project commissioned by the Australian Defense Agency, leading to a peer-reviewed publication at **MILCOM 2018**. PI: **Prof. Philip Leong**.
- Contribution to other FPGA-based research projects in the lab on topics such as **on-chip training**, **logic cell architecture design**, and **automatic modulation classification**.
- Extensive hands-on experience with technologies such as Vivado[HLS], RFSoC, PYNQ, Tensorflow, etc during stint at the lab. Undertook sysadmin duties to manage lab resources, and advised students on their final-year research projects.

Education & Certifications

Nanyang Technological University (NTU), Singapore — Jan 2013 to Feb 2019

PhD, Computer Science & Engineering, Advisor: **Prof. Nachiket Kapre**
Dissertation: **Dataflow Optimized Overlays for FPGAs**

- Developed a 16x16 mesh token dataflow overlay architecture finely-tuned for efficient mapping on Arria 10 FPGAs.
- Exploited statically-extracted instruction-criticality information to enable out-of-order execution inside each processing element; improved performance by up to 2.4× over existing in-order processor architectures.
- Verilog for hardware implementation; verilator for testbenches and cycle-accurate simulations; C++ compiler to optimize and map dataflow graphs to instructions inside each PE; python/R/shell scripting for data wrangling tasks.

Teaching Assistant: Programmable System on Chip (CE4054), Embedded Software Development (CE4052)

Imperial College London, United Kingdom — Oct 2009 to Jun 2012

Bachelors of Engineering (BEng), Electrical & Electronics Engineering

Machine Learning Engineering for Production Specialization, DeepLearning.AI, Coursera — Ongoing

Data Science Specialization, John Hopkins University, Coursera — Apr 2016