Technical Lead / Software Architect / System Architect

Education

PHD | 2013 - PRESENT | FACULTY OF ELECTRICAL ENGINEERING, UNIVERSITY OF BELGRADE

- Major: **Software Engineering**
- Area of research: **Quantum Computing**
- Currently on 2nd year.

MASTER | 2005 - 2009 | FACULTY OF PHYSICS, UNIVERSITY OF BELGRADE

- Major: **Nuclear Physics and Elementary Particle Physics**
- Area of research: **Anomalous Triple/Quartic Gradient Gauge Boson Couplings**

BACHELOR | 2000 - 2005 | FACULTY OF PHYSICS, UNIVERSITY OF BELGRADE

- Major: **Applied Physics and Informatics**

Theoretical Foundation

- Object Oriented Analysis and Design, Design Patterns, Enterprise Application Architecture Patterns, UML…
- System Architecture
- Algorithms and Data Structures
- Operating Systems Concepts
- Databases System Concepts, Relational Algebra, Object-Relational Mapping, NoSQL, Big Data
- Computer and Network Security, Cryptography Algorithms
- Computer Architecture
- Computer Networks, TCP/IP, OSI
- Quantum Mechanics, Quantum Computation
- Standard Model of Particle Physics

Skills

SYSTEM ARCHITECTURES

- Serverless, Event Driven, Microservices, Pipe and Filter, N-Tier, Big Compute, Big Data, Workflow Management…

DESIGN PATTERNS

- Object Oriented Design Patterns (GoF), Patterns of Enterprise Application Architecture, Cloud Design Patterns, Twelve Factor App, Python Design Patterns…

PROGRAMMING LANGUAGES

- Java, Python, Scala, JavaScript, TypeScript, C/C++, Objective C, Go, Fortran…

FRAMEWORKS/LIBRARIES

- Python: Django, Flask, Airflow, Tensorflow, SQLAlchemy, Celery, pandas, Fire, pytest, Dogpile, Multithreading, Multiprocessing, Asyncio, SQLAlchemy, Elasticsearch DSL, mypy, zeep, pymongo, motor, alembic…
- Java: Spring Framework (Spring Boot, Spring Security, Spring Data JPA, Spring Data REST, Spring HATEOAS, Spring MVC, Spring REST Docs, Thymeleaf), Play Framework, Java EE (EJB, JPA, JTA, Bean Validation, JAX-RS, JMS, JSP, JSF), Stream API, Guice, Ebean, Maven, Gradle, Collections Framework, Hibernate, JDBC, jUnit, Mockito, JaCoCo, Jersey, RabbitMQ, AspectJ, Swing…
- Scala: Spark, SBT
- JavaScript, TypeScript: AngularJS, React/Flux, jQuery, gulp, grunt, npm…
- C++: Qt, Netburner…

DATABASES

- SQL (DQL, DCL, DML, DDL)
- Relational: MySQL, PostgreSQL, SQLite, H2…
- NoSQL: MongoDB, Neo4j, Cassandra, Firestore, BigTable, DynamoDB, Azure Tables…

BIG DATA / BIG COMPUTE

- Hadoop: HBase, Hive, Pig, Zookeeper…
- Spark: Dataframe API, RDD API
- ElasticSearch: Logstash, Kibana…
- Kafka
- Solr

WORKFLOW ORCHESTRATION

- Apache Airflow

CLOUD SERVICES

- AWS: EC2, ECS, Load Balancing, RDS, ElastiCache, Lambda, SNS, S3, EMR, Glue, Athena, ESS, ECR, SQS, Batch, EFS, CloudWatch..
- Google Cloud Platform**:** Cloud Engine, Cloud Run, Firestore**,** BigQuery, Bigtable, Dataflow, Dataproc, Pub/Sub, Eventarc, Cloud Functions…
- Azure
- OpenShift
- Heroku

INFRASTRUCTURE MANAGEMENT

- Terraform**,** Ansible

CONTAINERS AND ORCHESTRATION

· Docker (Compose), Kubernetes, Amazon Elastic Container Service…

OPERATING SYSTEMS

· Linux Administration (Gentoo, CentOS, Ubuntu, Fedora, SUSE, SLC…), GNU CLI tools, Bash, scripting

· Windows Administration (all versions), Batch scripting

COMPUTER NETWORKS

· OSI, TCP/IP, Security, Design, Routing, Load Balancing, VPN, nginx, Apache HTTP Server,…

CI/CD

· Codeship, Bitbucket Pipelines, CircleCI, Jenkins…

Experience

HEAD OF ENGINEERING / SOFTWARE ARCHITECT / SYSTEM ARCHITECT / DATA ENGINEER @ ACADEMICLABS (JULY 2018 - JUNE 2022)

Technology Stack: Python 3 (multiprocessing, multithreading, asyncio, SQLAlchemy, Elasticsearch DSL, Airflow, flask, mypy, zeep, pymongo, motor, dogpile, alembic, fire, pytest...), Go, Ruby, Docker, Docker Compose, ElasticSearch, Hadoop, Spark, Hive, SOAP, MongoDB, Neo4j, Codeship, Jenkins, GNU/Linux CLI, Bash scripts, Linux, AWS (EMR, Glue, Athena, ESS, ECS, ECR, SQS, Lambda, SNS, Batch, EC2, RDS, ElastiCache, S3, EFS, CloudWatch...), PostgreSQL…

AcademicLabs is unlocking world's R&D capabilities via an online matchmaking platform

Worked on:

- Designed and implemented data pipeline architecture for parallel and distributed ingesting, processing and indexing of hundreds of millions of documents daily.
- Designed and wrote several declarative and reflexive frameworks in Python for parsing and processing of structured and semi structured data
- Designed and wrote framework in Python for efficient and distributed indexing of billions of documents daily to ElasticSearch
- Designed, configured and managed complete project cloud infrastructure on AWS (ECS containers, RDS and other persistence services, ElasticSearch, highly distributed computing using fleet of EC2, load balancing, network security, VPN…)
- Designed and implemented several microservices in Python
- Designed platform and data pipeline relational schemas
- Designed all non relational data models (MongoDB, Neo4j, ElasticSearch…)
- Managed, configured and optimized runtime of persistence services (PostgreSQL, MongoDB, ElasticSearch…)
- Implemented and managed all CI/CD pipelines
- Complete Dockerization of all the microservices and services.
- Various automation Bash scripts.
- Managed VPNs and complete security of the project
- R/D with different big data tools and services (Hadoop, Spark, AWS Big Data services…)

BIG DATA ENGINEER @ SEAVUS (MAY 2018 – JULY 2018)

Technology stack: Hadoop (Hbase, Hive, Pig, Spark, Zookeeper…)

Seavus is software development services provider

Worked on:

- Big data ETL and analytics for big Swiss telecom company

SENIOR SOFTWARE ENGINEER / BIG DATA ENGINEER @ HIGHFIVE (APRIL 2018 – MAY 2018)

Technology stack: GCP Big Data (BigQuery, Cloud Bigtable, Cloud Dataflow, Cloud Dataproc, Cloud Pub/Sub), Elastic Stack (Elasticsearch, Logstash), Hadoop Ecosytstem, Kafka, Spark, Solr, Cassandra…

Highfive is video and web conferencing platform.

Worked on:

- Migrating ETL and data pipelines from GCP Big Data Stack and Elastic Stack to open source solutions.

LEAD SOFTWARE ENGINEER / SOFTWARE ARCHITECT / BIG DATA ENGINEER / DEVOPS @ EPINOMICS (SEPTEMBER 2017 – APRIL 2018)

Technology Stack: Java 8 (Spring Boot, Spring Security, Spring Data for Apache Cassandra, Spring Data Redis, Play Framework, Guice, Ebean, Stream API...), Python 3.5 (SQLAlchemy, pytest, fire, Celery,…), Scala, Spark, Airflow, Bash Scripting, Docker, Docker Compose, AWS (EC2, ECS, Load Balancing, RDS, ElastiCache, S3), Solr, Cassandra, MySQL, SQLite, Redis, Bitbucket Pipelines, Qubole, Jenkins, nginx, Git, Maven, SBT, GNU/Linux CLI, AngularJS, Ansible…

Personalized Health with Epigenomics. Epinomics is revolutionizing personalized medicine through scientific research, deep analytics, and clinical application development.

Worked on:

- Production of Spark pipelines using Apache Airflow, Python 3 (SQLAlchemy, pytest...), Docker, Docker Compose, Bitbucket Pipelines, ECS, Qubole, bash scripts and integration with REST API...
- Design and implementation of Scala Apache Spark Dataset pipelines, with integration tests...
- Building the Epinomics Portal infrastructure using EC2 Container Service, Load Balancers, RDS, ElastiCache...
- Building Cassandra microservice using Spring Boot (Spring Security, Spring Data for Apache Cassandra, Spring Data Redis). Dockerization of this service. CI/CD configuration using Bitbucket Pipelines and AWS.
- Dockerization of all the microservices and services.
- CI/CD configuration and deployment to AWS of all the microservices using Bitbucket Pipelines.
- Apache Solr configuration, dockerization, and Schema definition. Development of MySQL and Cassandra import script using Solr and Python scripts.
- Adding new functionalities to Java Play Framework microservices.
- Separation of AngularJS frontend from backend services. Dockerization of frontend with nginx configuration. Adding new functionalities to AngularJS frontend.
- Various automation Bash scripts.
- Various security concerns.

LEAD SOFTWARE ENGINEER / BIG DATA ENGINEER @ FATHOM HEALTH (JUNE 2017 – SEPTEMBER 2017)

Technology stack: Python 3.5 (SQLAlchemy, pandas, pytest, fire, pyinvoke, dogpile, custom built frameworks, multithreading...), TensorFlow, MySQL, PostgreSQL, SQLite, Azure, Docker, Airflow, Git, GNU/Linux CLI, Terraform, Kubernetes, Google protobuf, memcached, Multithreading, CircleCI, Codacy...

Fathom Health is a Deep Learning NLP system applying Deep Learning to medical coding in the US healthcare industry. Medical coding is a natural language processing task to extract symptom, disorder and treatment entities from an unstructured plain-text diagnosis.

Worked on:

- Setting up data pipelines using Apache Airflow, Tensorflow, Python 3 scripts, bash scripts, to train the deep learning model.
- Various statistical and analytical modules in Python 3.
- Moving the persistence layer and data sources to Azure using combination of bash scripting, Python 3 scripts and SQL.
- Database (MySQL, PostgreSQL, SQLite) access parallelization using Python 3 concurrent.
- Various database access caching solutions.

- Various MySQL, PostgreSQL and SQLite creation scripts using the combination of SQL, Python and SQLAlchemy, with strong emphasis on reusable design patterns.
- Dockerization of those scripts.

LEAD SOFTWARE ENGINEER / SOFTWARE ARCHITECT / DEVOPS @ SENTRYS (SEPTEMBER 2016 – MAY 2017)

Technology stack: Java 8 (Spring Boot, Hibernate ORM, Spring Data JPA, Spring Security, Spring Data REST, REST Docs, Spring HATEOAS, Thymeleaf...), AngularJS, MySQL, H2, AWS (EC2, S3, RDS…), Jenkins, Git, REST, Encryption...

Sentrys is a data privacy & data ownership access broker startup.

Worked on:

- Lead software engineering and architect, developing the backend in Spring Framework using Spring Boot, Hibernate ORM, Spring Data JPA, Spring Security, Spring Data REST, REST Docs, Spring HATEOAS…
- Design and the implementation of the REST API.
- Writing more than 500 unit and integration tests.
- Sensitive nature of project's data demanded the use of encryption in the various layers of the application (transportation, persistence etc…)
- Frontend prototyping using Thymeleaf.
- Designed and written big part of AngularJS frontend, especially concentrating on integration with fully secured Spring backend.
- AWS deployment and administration.
- CI/CD pipelines using Jenkins.

SENIOR SOFTWARE ENGINEER / DATA ENGINEER / DEVOPS @ EHUB CONNECTED (MARCH 2015 – AUGUST 2016)

Technology Stack: Java (Java SE, Java EE, Spring, Hibernate, Android), Python (httplib, pandas, IPython, matplotlib…), C++, AWS (S3, DynamoDB, RDS, Glacier), Linux, Bash Scripting, MySQL, Jenkins, Git...

eHub is a vehicle IoT device with abilities to capture OBD (on-board diagnostic) data, audio/video, geolocations and various other parameters…

Worked on:

- Development of REST microservices using Java EE and Spring. GlassFish deployment and administration.
- Python scripts for various data analysis and testing tasks.
- Development of Android USB accessory functionality.
- Design solutions and code review of C++ firmware for IoT device.
- Experimentation with various AWS persistence options.
- CI/CD pipelines using Jenkins.

SOFTWARE ENGINEER / DEVOPS / NETWORK ENGINEER @ ETAG SISTEMI, BELGRADE (FEBRUARY 2010 – MARCH 2015)

Technology Stack: Java (Java SE, Java EE), Android, iOS (Objective C), React.js/Flux, Linux, Bash Scripting, Batch Scripting, Virtualization, TCP/IP, MySQL, MSSQL, Jenkins, Apache, TFS, SVN...

Worked on:

- Development and maintenance of POS and bookkeeping software for petrol stations.
- Development of various applications, microservices and scripts using Java (Java SE, Java EE), Android, iOS, C/C++, JavaScript, encryption…
- Development of embedded software for dispenser controllers using C++ Netburner toolkit.
- Installation and administration of various tools and services: Apache HTTP server, Tomcat, Glassfish, MySQL Server, MSSQL Server, SoftEther VPN, Redmine, Jenkins, Nessus, MediaWiki, TWiki, Team Foundation Server, Bonobo Git Server, Subversion Server, IIS, cPanel hosting, various automatic backup solutions, distributed virtual machines…
- Installations and administration of various OS: Windows Servers, CentOS, RHEL, Gentoo, Ubuntu…
- Design, installation, administration and security of various computer networks.
- Various Jenkins automation tasks.

RESEARCH ASSISTANT @ INSTITUT LAUE-LANGEVIN, GRENOBLE (OCTOBER 2007 – NOVEMBER 2007)

- Main area of research: **Anomalous Triple Gradient Gauge Boson Couplings**

RESEARCH ASSISTANT / SOFTWARE ENGINEER @ CERN (EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH), GENEVA (FEBRUARY 2007 – JULY 2007)

Technology stack: Java SE, Swing, Bash scripting, GNU/Linux CLI, TCP/IP, Multithreading, Eclipse...

- Main area of research: **Anomalous Quartic Gradient Gauge Boson Couplings**
- Development of multithreaded client/server TCP/IP application in Java SE/Swing for the ATLAS Detector Control Room that facilitated coordination and synchronization between detector operators.
- Statistical analysis and generation of big volumes of data using the simulator of proton collisions inside of Large Hadron Colider's Atlas Detector. Data was being generated on LHC Computing Grid.
- Python, Fortran and C++ scripts for statistical analysis and processing of the big data.

RESEARCH ASSISTANT/ SOFTWARE ENGINEER / SYSTEM ADMINISTRATOR @ INSTITUTE OF PHYSICS, BELGRADE (NOVEMBER 2005 – DECEMBER 2009)

Technology Stack: Java SE, Python, C++, Swing, Fortran, Bash scripting, GNU/Linux CLI, Virtualization, Eclipse, LaTeX...

- Main area of research: **Anomalous Triple Gradient Gauge Boson Couplings**
- Statistical analysis and generation of big volumes of data using the simulator of proton collisions inside of Large Hadron Colider's Atlas Detector. Data was being generated either on LHC Computing Grid, or Institute's own supercomputer.
- Python, Fortran and C++ scripts for statistical analysis and processing of the big data.
- Automation of the workflow for the researchers in group (submitting jobs on CERN Grid, data processing, data analysis, generation of reports) using combination of Python and Bash scripts.
- Code review and debugging of collisions simulator written in C++.
- Administration of various Linux systems, software tools and administration of the part of institute's computer network.

Languages

SERBIAN– NATIVE

ENGLISH – C1

SPANISH – C1 (INSTITUTO CERVANTES)