

Noor Nashid

CONTACT Software Analysis and Testing (SALT) Lab
INFORMATION Kaiser 3070, Electrical and Computer Engineering
The University of British Columbia, Vancouver, Canada

Cell: +1 236 335 7047
Email: noor.nashid1@gmail.com
Web: nashid.github.io

EDUCATION

Ph.D. in Software Engineering, (January 2020 - April 2026)

- Thesis: On the Role of Context in AI-Driven Program Repair and Test Generation
- Advisor: Professor Ali Mesbah
- Department of Electrical and Computer Engineering, University of British Columbia (UBC), Vancouver, Canada

Master's by Research in Computer Science, (January 2010-December 2011)

- Thesis: A Framework for HTTP Traffic Acceleration over High Latency Links
- Advisor: Professor Cormac Sreenan
- Department of Computer Science, University College Cork (UCC), Cork, Ireland

Bachelor of Science in Computer Science and Engineering (2002-2007)

- Thesis: Statistical Method Based Bangla Word Sense Disambiguation
- Advisor: Dr. Muhammad Masroor Ali, Professor, Dept. of CSE, BUET
- Department of Computer Science and Engineering (CSE), Bangladesh University of Engineering and Technology (BUET), Dhaka, Bangladesh

JOURNAL PAPERS

Beyond Accuracy: Behavioral Dynamics of Agentic Multi-Hunk Repair, **Noor Nashid**, Daniel Ding, Keheliya Gallaba, Ahmed E Hassan, Ali Mesbah, 2025. (Under Submission)

Katana: Dual Slicing Based Context for Learning Bug Fixes, Mifta Sintaha, **Noor Nashid**, Ali Mesbah, ACM Transactions on Software Engineering and Methodology (TOSEM), 2023.

A controlled experiment of different code representations for learning-based program repair, *Marjane Namavar, ***Noor Nashid**, Ali Mesbah, Empirical Software Engineering, 2022. (***Marjane Namavar and Noor Nashid contributed equally to this work.**)

Client-Side Framework for Automated Evaluation of Mechanisms to Improve HTTP Performance, Paul Davern, **Noor Nashid**, C.J. Sreenan, Ahmed Zahran, *Journal of Networks (JNW)*, Vol 7, No 11, 2012.

HTTPEP: a HTTP Performance Enhancing Proxy for Satellite Systems, Paul Davern, **Noor Nashid**, C.J. Sreenan, Ahmed Zahran, *International Journal of Next Generation Computing*, Vol 2, No 3, November 2011.

CONFERENCE PUBLICATIONS Issue2Test: Generating Reproducing Test Cases from Issue Reports, **Noor Nashid**, Islem Bouzenia, Michael Pradel, Ali Mesbah, International Conference on Software Engineering (ICSE), 2026.

Sijia Gu, **Noor Nashid**, Ali Mesbah. LLM Test Generation via Iterative Hybrid Program Analysis, International Conference on Software Engineering (ICSE), 2026.

Characterizing Multi-Hunk Patches: Divergence, Proximity, and LLM Repair Challenges, **Noor Nashid**, Daniel Ding, Keheliya Gallaba, Ahmed E Hassan, Ali Mesbah, International Conference on Automated Software Engineering (ASE), 2025.

Taha Shabani, **Noor Nashid**, Parsa Alian, Ali Mesbah. Dockerfile Flakiness: Characterization and Repair. IEEE/ACM International Conference on Software Engineering (ICSE), 2025.

Parsa Alian, **Noor Nashid**, Mobina Shahbandeh, Taha Shabani, Ali Mesbah. A Feature-Based Approach to Generating Comprehensive End-to-End Tests. ACM/IEEE International Conference on Software Engineering (ICSE), 2025.

Retrieval-Based Prompt Selection for Code-Related Few-Shot Learning, **Noor Nashid**, Mifta Sintaha, Ali Mesbah, International Conference on Software Engineering (ICSE), 2023.

Embedding Context as Code Dependencies for Neural Program Repair, **Noor Nashid**, Mifta Sintaha, Ali Mesbah, IEEE Conference on Software Testing, Verification and Validation (ICST), 2023.

Towards a Framework for Evaluating HTTP/TCP Performance, Paul Davern, **Noor Nashid**, Ahmed Zahran, C.J. Sreenan *2011 International Symposium on Performance Evaluation of Computer and Telecommunication Systems (SPECTS 2011)*, The Hague, June 2011.

HTTP acceleration over high latency links, Paul Davern, **Noor Nashid**, Ahmed Zahran, C.J. Sreenan *Proc. of Fourth IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, Paris, Feb 2011.

LICENSING AGREEMENTS Methods to Optimise Internet Data Service for Satellite Systems. Licensee: Altobridge Ltd, 2011.

Software Architecture for Optimised Delivery of Web Traffic over Satellite. Licensee: Altobridge Ltd, 2012.

PATENTS *A Communication System, Patent Number: US9560468B2*, Granted January 2017.

GRANTS AND FUNDING

- Led proposal development for three awarded research grants at University of British Columbia (PI: Prof. Ali Mesbah): AWS Generative AI Research Grant, NSERC Discovery Grant (DG), and NSERC Research Tools and Instruments grant (RTI). As PhD candidate, contributed primary technical content and research design for the proposals.

AWARDS

- ACM SIGSOFT Distinguished Paper Award at ASE 2025.
- UCC Research Commercialisation Award (Inventor Award) for *Methods to Optimise Internet Data Services for Satellite Systems*, April 2011. Research licensed to Altobridge Ltd.
- UCC Research Commercialisation Award (Inventor Award) for *Software Architecture for Optimised Delivery of Web Traffic Over Satellite*, August 2010. Research licensed to Altobridge Ltd.
- Shortlisted for the University College Cork's Invention of the Year Award, 2012 for *Software enabling accelerated data services for satellite systems at lower costs and increased speed and quality project*.
- ACM SIGSOFT CAPS travel grant
- UBC four-year doctoral fellowship (4YF) award.
- IBM Extreme Blue premier internship award, June 2011.
- Full Scholarship for Master's By Research program under Backhaul Optimisation in Heterogeneous Wireless Access Networks (BACOPT) project at Mobile & Internet Systems Laboratory (MISL), UCC, Cork, Ireland. (January 2010 - December 2011)
- University Admission Test Scholarship: Awarded for securing 47th position in the admission test, comprising more than 5000 applicants.

SERVICE **Program Committee Member**

- Mining Software Repositories (MSR) Industry Track 2026, 2025
- SANER Tool Demo Track 2026
- International Conference on Software Engineering (ICSE) Artifact Evaluation Track 2026, 2025

- 3rd ACM International Conference on AI-powered Software (AIware) 2026
- Foundations of Software Engineering (FSE) Ideas, Visions and Reflections 2026
- ASE Tools and Datasets Track 2026
- Reliable and Secure AI for Software Engineering (ReSAISE) 2025, 2024, 2023
- Open Science for Foundation Models (SCI-FM), ICLR 2025
- Deep Learning for Code (DL4C), ICLR 2025
- Internetwork 2026, 2025
- DeepTest 2025, 2024
- Mining Software Repositories (MSR) 2025, Junior PC
- International Conference on Technical Debt (TechDebt) 2025, Junior PC
- ASE Industry Showcase 2024
- Forge 2024
- ICSE 2025, 2024, Co-reviewer
- LLM4Code 2024, Co-reviewer
- ISSTA 2023, 2020, Co-reviewer
- FSE 2021, 2020, Co-reviewer
- ICSE 2021, 2020, Co-reviewer

Journal Reviewer

- Empirical Software Engineering (EMSE), 2025-present
- ACM Transactions on Software Engineering and Methodology (TOSEM), 2024-present
- Information and Software Technology, 2025-present
- Journal of Systems & Software, 2024-present
- Science of Computer Programming, 2025-present
- Journal of Software - Evolution and Process, 2024-present
- Future Generation Computer Systems, 2026-present
- Computer Applications in Engineering Education, 2026-present

TEACHING EXPERIENCE

Lead Teaching Assistant, University of British Columbia

- **CPEN 522 (Software Verification and Testing)**: Designed lab assignments from scratch, introducing a large project with multiple components: front end, back end, REST API layer, and an AI component. Emphasized realistic, large-scale project work to enhance student learning beyond isolated exercises. Handled grading, office hours, and all aspects of course management.
- **CPEN 422 (Software Testing and Analysis)**: Coordinated the TA team for a course with nearly 100 students, managed project grading, and guided student teams through the full development lifecycle.

Teaching Assistant, University College Cork

- **CS1106 (Introduction to Relational Databases)**: Supported undergraduate students in learning database fundamentals, SQL query design, and database modeling concepts.
- **CS2505 (Network Computing)**: Assisted students with understanding network protocols, distributed systems, and practical network programming.

PROFESSIONAL EXPERIENCE

Before academia, I held industry positions across diverse domains, building developer tools for application performance management and working in software security and e-commerce. I served as a founding engineer on the UNDP national voter ID platform for Bangladesh, which served over 100 million citizens.

Electrical And Computer Engineering, University of British Columbia

(February 2020 - Ongoing)

- I am currently working as a Research Assistant to conduct research in the field of deep learning for software engineering under the supervision of *Prof. Ali Mesbah*. The overarching goal of my work is to investigate how to incorporate contextual information into different learning-based source code processing tasks.

Tech used: Large language model (LLM), Program analysis techniques, Neural Machine Translation (NMT), Graph Neural Network (GNN), TensorFlow, Natural Language Processing (NLP), Java, Python, Amazon Bedrock - Generative AI, Fine-tuning, LangChain, LiteLLM, ChromaDB, Claude Code, OpenAI Codex, Gemini CLI, Amazon Q Developer, Tree-sitter, Docker, Python, Java

- **Characterizing Multi-Hunk Patches: Divergence, Proximity, and LLM Repair Challenges**

(ASE 2025): Multi-hunk bugs, where fixes span disjoint regions of code, are common in practice, yet remain underrepresented in automated repair. Existing techniques and benchmarks pre-dominantly target single-hunk scenarios, overlooking the added complexity of coordinating semantically related changes across the codebase. In this work, we characterize HUNK4J, a dataset of multi-hunk patches derived from 372 real-world defects. We propose hunk divergence, a metric that quantifies the variation among edits in a patch by capturing lexical, structural, and file-level differences, while incorporating the number of hunks involved. We further define spatial proximity, a classification that models how hunks are spatially distributed across the program hierarchy. Our empirical study spanning six LLMs reveals that model success rates decline with increased divergence and spatial dispersion. Notably, when using the LLM alone, no model succeeds in the most dispersed Fragment class. These findings highlight a critical gap in LLM capabilities and motivate divergence-aware repair strategies.

- **Issue2Test: Generating Reproducing Test Cases from Issue Reports (Under Submission):** Automated tools for solving GitHub issues are receiving significant attention by both researchers and practitioners, e.g., in the form of foundation models and LLM-based agents prompted with issues. A crucial step toward successfully solving an issue is creating a test case that accurately reproduces the issue. Such a test case can guide the search for an appropriate patch and help validate whether the patch matches the issue’s intent. However, existing techniques for issue reproduction show only moderate success. This paper presents Issue2Test, an LLM-based technique for automatically generating a reproducing test case for a given issue report. Unlike automated regression test generators, which aim at creating passing tests, our approach aims at a test that fails, and that fails specifically for the reason described in the issue. To this end, Issue2Test performs three steps: (1) understand the issue and gather context (e.g., related files and project-specific guidelines) relevant for reproducing it; (2) generate a candidate test case; and (3) iteratively refine the test case based on compilation and runtime feedback until it fails and the failure aligns with the problem described in the issue. We evaluate Issue2Test on the SWT-bench-lite dataset, where it successfully reproduces 30.4% of the issues, achieving a 16.3
- **LLM Test Generation via Iterative Hybrid Program Analysis (ICSE 2026)** Automating unit test generation remains a significant challenge, particularly for complex methods in real-world projects. While Large Language Models (LLMs) have made strides in code generation, they struggle to achieve high branch coverage due to their limited ability to reason about intricate control flow structures. To address this limitation, we introduce *Panta*, a technique that emulates the iterative process human developers follow when analyzing code and constructing test cases. *Panta* integrates static control flow analysis and dynamic code coverage analysis to systematically guide LLMs in identifying uncovered execution paths and generating better test cases. By incorporating an iterative feedback-driven mechanism, our technique continuously refines test generation based on static and dynamic path coverage insights, ensuring more comprehensive and effective testing. Our empirical evaluation, conducted on classes with high cyclomatic complexity from open-source projects, demonstrates that *Panta* achieves 26% higher line coverage and 23% higher branch coverage compared to the state-of-the-art.
- **Dockerfile Flakiness: Characterization and Repair (ICSE 2025):** Dockerfile flakiness-unpredictable temporal build failures caused by external dependencies and evolving environments-undermines deployment reliability and increases debugging overhead. Unlike traditional Dockerfile issues, flakiness occurs without modifications to the Dockerfile itself, complicating its resolution. In this work, we present the first comprehensive study of Dockerfile flakiness, featuring a nine-month analysis of 8,132 Dockerized projects, revealing that around 10% exhibit flaky behavior. We propose a taxonomy categorizing common flakiness causes, including dependency errors and server connectivity issues. Existing tools fail to effectively address these challenges due to their reliance on pre-defined rules and limited generalizability. To overcome these limitations, we introduce *FLAKIDOCK*, a novel repair framework combining static and dynamic analysis, similarity retrieval, and an iterative feedback loop powered by Large Language Models (LLMs). Our evaluation demonstrates that *FLAKIDOCK* achieves a repair accuracy of 73.55%, significantly surpassing state-of-the-art tools and baselines.
- **Curing Hallucinations with Context for Local API Completion:** Large language models have made substantial progress in addressing diverse code-related tasks. However, the adoption of language models is hindered by hallucinations, where models generate incorrect or inaccurate output due to the lack of real-world, domain-specific information, such as for intra-repository API calls. We introduce a novel technique to mitigate hallucinations by leveraging global and local contextual information within a code repository for API completion tasks. Our approach is tailored to refine code completion tasks, with a focus on optimizing local API completions. We examine import statements during API completion to derive insights into local APIs, drawing from their method signatures. Following that, we analyze the inline variables and correlate them with the appropriate imported modules, thereby allowing our approach to rank

the most contextually relevant suggestions from the available local APIs within those imports. We further employ a contextual pruning step where irrelevant contextual information is filtered out, retaining only the pertinent details relevant to local APIs around the code under edit. The implications of our findings are substantial for developers, indicating that our method can be applied in multilingual environments without the need for language-specific training or fine-tuning. This allows for the easy adoption of our approach across a wide array of large language models (LLMs) with minimal effort.

- **Retrieval-Based Prompt Selection for Code-Related Few-Shot Learning (ICSE 2023)**: Large language models trained on massive code corpora can generalize to new tasks without the need for task-specific fine-tuning. In few-shot learning, these models take as input a prompt, composed of natural language instructions, a few instances of task demonstration, and a query and generate an output. However, the creation of an effective prompt for code-related tasks in few-shot learning has received little attention. We present a technique for prompt creation that automatically retrieves code demonstrations similar to the developer task, based on embedding or frequency analysis. We apply our approach, CEDAR, to two different programming languages, statically and dynamically typed, and two different tasks, namely, test assertion generation and program repair. For each task, we compare CEDAR with state-of-the-art task-specific and fine-tuned models. The empirical results show that, with only a few relevant code demonstrations, our prompt creation technique is effective in both tasks, with an accuracy of 76% and 52% for exact matches in test assertion generation and program repair tasks, respectively. For assertion generation, CEDAR outperforms existing task-specific and fine-tuned models by 333% and 11%, respectively. For program repair, CEDAR yields 189% better accuracy than task-specific models and is competitive with recent fine-tuned models. These findings have practical implications for practitioners, as CEDAR could potentially be applied to multilingual and multitask settings without task or language-specific training with minimal examples and effort
- **Embedding Context as Code Dependencies for Neural Program Repair (TOSEM 2023, ICST 2023)**: Contextual information plays a vital role for software developers when understanding and fixing a bug. Consequently, deep learning-based program repair techniques leverage context for bug fixes. However, existing techniques treat context in an arbitrary manner, by extracting code in close proximity to the buggy statement within the enclosing file, class, or method, without any analysis to find actual relations with the bug. To reduce noise, they use a predefined maximum limit on the number of tokens to be used as context. We present a program slicing-based approach, in which instead of arbitrarily including code as context, we analyze statements that have a control or data dependency on the buggy statement. We propose a novel concept, which leverages the context of both buggy and fixed versions of the code to capture relevant repair ingredients. We are the first to apply slicing-based context to the learning-based program repair task.
- **A Controlled Experiment of Different Code Representations for Learning-Based Program Repair (EMSE 2022)**: Training a deep learning model on source code has gained significant traction recently. Since such models reason about vectors of numbers, source code needs to be converted to a code representation before vectorization. Numerous approaches have been proposed to represent source code, from sequences of tokens to abstract syntax trees. However, there is no systematic study to understand the effect of code representation on learning performance. Through a controlled experiment, we examine the impact of various code representations on model accuracy and usefulness in deep learning-based program repair. We train 21 different generative models that suggest fixes for name-based bugs, including 14 different homogeneous code representations, four mixed representations for the buggy and fixed code, and three different embeddings. We assess if fix suggestions produced by the model in various code representations are automatically patchable, meaning they can be transformed to a valid code that is ready to be applied to the buggy code to fix it. We also conducted a developer study to qualitatively evaluate the usefulness of inferred fixes in different code representations.

Merchant Technology, Jet Technology Engineering, Walmart Labs (August 2019 - February 2020)

- I am working as a Lead Software Engineer at Walmart Labs, which is the tech innovation arm of Walmart. This initiative is part of Walmart's continued investment in the technology landscape to bolster the online services domain aggressively. I worked in the dynamic pricing initiative, a greenfield project to build an event-driven elastic highly scalable system. At Walmart Labs, the typical load can range from 90 million monthly visitors, loads up to 15,000 requests per second, and more than 20 million products. My primary responsibility is to design the overall architecture and to develop the system so that it can satisfy 99.99% uptime.

Tech used: Scala, Java, Microsoft Azure, Google Cloud Platform, Apache Flink, Apache Kafka, PostgreSQL, Cassandra

Search Engine Optimisation (Team Fire), Hudson's Bay Company (HBC) (November 2017 - June 2019)

- Working as a Senior Team Member in the HBC Tech in Dublin. Hudson's Bay Company is one of the fastest-growing department store retailers in the world. Founded in 1670, HBC is the oldest company in North America. I initially started to work on the Gilt Groupe, which was a flash sales site. Currently, my responsibility is to work on the SEO/SEM initiative across HBC banner. HBC's leading banners include Hudson's Bay, Lord & Taylor, Saks Fifth Avenue, and Saks OFF 5TH. I have been primarily developing tools and services which can facilitate HBC in the SEO/SEM front across all the banners.

Tech used: Java, Scala, Python, Amazon Web Services (AWS), Docker, Kubernetes, GoCD, Shell Scripting, Oracle Commerce, Akamai

Senior Software Engineer

- I augmented product feed publisher microservice for the Gilt banner, which is to facilitate affiliate marketing. This microservice was interfacing with Gilt inventory and was publishing product feed according to the stakeholders need.
- I have been working to optimize the URL generation for the Product Detail Page (PDP). Prior to this initiative, PDP URL was not SEO friendly and to mitigate that new PDP URL is generated with the brand name, product short description and produce code with the goal to increase the organic web traffic. This has been a highly visible project in the organisation, and my work resulted in a revenue increase in the range of \$5M.
- I have developed a Sitemap Generation Service which builds and publishes Sitemap for Search Engines in a periodic interval.

Legacy sitemap service was deployed in mutable server infrastructure, and the deployment process was manual and error-prone. I took the initiative to build consensus across the team and proposed a microservice for building Sitemap in an immutable server infrastructure. I have also improved the existing sitemap generation process and alerting infrastructure around it. This project was highly successful, which resulted in a significant improvement in search index coverage (97% from 15%) for PDP pages.

Targeted Attack Analysis (TAA), Kaspersky Lab (February 2017 - November 2017)

- Working as a Senior Team Member in the newly established team in Dublin. Dublin office is the first R&D centre of Kaspersky outside of Russia. Here I am developing a new enterprise product offering that goes beyond the traditional SIEM technologies and provides enterprise-wide visibility, integrates a disparate set of data sources (endpoint agents, network agents) and apply different malware detection techniques to reveal suspicious user behavior, patterns and generate alerts that matter. I contributed to the overall framework design of the Targeted Attack Analysis (TAA) platform, which consists of agent server, data ingestion pipeline (streaming, batch), and detector. I took a pivotal role to bridge the gap between data scientists, malware detection team, and engineers and ensured the insights drawn by the data scientists and malware researchers are brought into the mainstream product.

Tech used: Python, Java, Kafka, Hbase, Hadoop, Hive, Spark, PostgreSQL, Shell Scripting

Lead Software Developer

- I integrated the TAA solution with Kaspersky Security Network (KSN), a complex distributed infrastructure dedicated to processing cybersecurity-related data streams from millions of participants worldwide.
- My major contribution is to develop different detection techniques by integrating threat intelligence from KSN cloud to detect the suspicious file, URL, domain by considering different features like application fingerprints/reputation, domain reputation/popularity, URL classification and applied heuristics for detecting malicious behavior.

Search Engine Marketing (SEM), Groupon (May 2015 - February 2017)

- At Groupon, I am exploring technologies for computational marketing to deal with Groupon's massive international customer base. I have taken multiple products from proof of concept to customer-facing in short time. During this time, I have contributed to different domains of search engine marketing, for example, Keyword Knowledge Base, Deal performance service, Keyword insights, Social insights.

Tech used: Java, Dropwizard, MongoDB, Cassandra, Redis, Hive, PostgreSQL, Shell Scripting, Docker, Opswise (a web-based scheduling tool)

Software Development Engineer - 2

- Keyword Knowledge Base (Keyword-KB) was started from scratch which is Groupon's single, authoritative point of all relevant knowledge about search keywords. Its goal is to store every bit of useful information about keywords, their demand, and their organization in one place. Its purpose is to enable automated and human decision-making around keyword-intent in marketing, supply, and product at Groupon. Overall this service combines a set of tools for assisting keyword generation, keyword classification, categorization, keyword statistics management as well as a convenient API layer to access that information.
- The Keyword Insight Tool (KIT) provides users the ability to quickly examine keyword history for what SEM team is actively advertising on. The tool is designed to improve the long latencies business stakeholders have experienced with Hive, where to get performance data for a single keyword could take a few minutes. Instead, the information will be displayed within a few seconds in KIT with Redis cache and Cassandra storage.
- Social Insight project was undertaken as an exploratory task to quantify the social reach of Groupon deals. With that in mind a tool was developed to visualize the social health of a Groupon deal by aggregating metrics from multiple social feeds, for example, Facebook, Twitter and from Groupon's own site and to provide the user with a consolidated view.

IBM Application Performance Management

(January 2012 - April 2015)

- I was part of the IBM Application Performance Management (APM) team from 2012. Since then, I have contributed to different facets of application performance management (APM). I started developing an extension for Hypervisor support within the existing WebSphere application server monitoring solution.

Followed to that I have worked in the next generation SaaS-based APM solutions. I worked as a Lead Developer to create a lightweight monitoring solution for WebSphere Application Server. This lightweight monitoring solution mainly collects JMX based data and provides a seamless way to transition from JMX based monitoring to deep-dive monitoring without end-user involvement. I also lead a team of developers and QA engineers to introduce continuous delivery pipeline. Another exciting project that I contributed to was the development of a high-throughput low-latency data layer, which was a core component of the SaaS architecture.

And then I worked as a Senior Software Engineer on a greenfield SaaS-based Synthetic Monitoring project to enhance the IBM APM SaaS portfolio.

Technologies & tools: Agile, Scrum, Java, Python, Design Patterns, REST, MongoDB, Shell Scripting, build tools (Ant, Maven), Websphere Application Server, different integration tools (Jenkins, Smart Cloud Orchestrator), IBM Bluemix (Cloud Foundry-based PaaS), Docker.

Staff Software Engineer, Synthetic Monitoring

(July 2014 - April, 2015)

- I have developed the Synthetic Monitoring solution for extending the IBM Application Performance Management portfolio. End users can record the functional tests for their Web apps using Selenium. Test scripts can be written using Selenese commands or can be recorded easily using the Selenium-IDE. Upon uploading these Selenium Scripts, scripts are played back from multiple points of presence across the globe in a custom interval. Performance metrics are collected using the browser, and real-time and historical dashboard can be analyzed for providing flawless app experience 24X7.
- This challenging project was started from scratch, and I took part in capability analysis of the existing APM vendors, architecture development, UX/UI wireframes, and finally contributing to the software development.
- While I am involved with the overall development of different components, I am leading the data provider component for real-time dynamic dashboard support, which is mainly a REST-based data provider interface. Used tools and technologies - Java, MongoDB, Shell Scripting, build tools (Ant, Maven), Docker, Selenium.

Staff Software Engineer, Application Performance Diagnostics (APD)

(June 2013 - June 2014)

- Development of a high-throughput low-latency data layer which is a fundamental component to gather deep-dive metrics for the APM SaaS platform. Technologies used - Java, Apache Kafka, ZooKeeper, JMS, horizontally scalable mongo setup.

- Worked as a Lead Software Engineer to develop the lightweight monitoring component for WebSphere Application Server, which is one of the major features of the APD project. I have written and assigned stories for each iteration for the lightweight monitoring (DC-less) feature development and assigned them to the fellow developers, QA, ID, and ensured this feature could be finished in time. I have also demonstrated this developed feature from a live testbed and communicated features and functionality to the customers.
- Lead a team of Software Engineers and QA to introduce continuous integration for APD. APD evolved from a legacy product over the years, and there was no complete end-to-end continuous integration. I led the team and gradually improved the awareness of DevOps and Continuous integration, which eventually lead to streamlining the build and release process within the APD organization.
- Working as a mentor for the graduate Software Developers.
- Work with critical clients directly to educate and help with their product issues

Software Engineer, IBM Tivoli Composite Application Manager for Applications Diagnostics (ITCAM-AD) **(June 2012 - May 2013)**

- I have worked on different facets of ITCAM -AD. ITCAM-AD provides Web-based diagnostic view for faster problem isolation and determination through memory trends graphing, metric comparison, heap analysis, and transaction tracing for WebSphere and other J2EE application servers.
- Developed and delivered ITCAM for AD 7.2 support for WAS Hypervisor edition - interactive and silent configuration in ESX environment
- Developed OS agent for IBM Smart Cloud Application Performance Management for Continuous Delivery (DevOps) project
- Script package development for WAS agent with basic and advanced set of parameters for DevOps project

IBM - Tivoli Netcool Wireless Products Engineering

Software Engineer, Tivoli Netcool Performance Manager for Wireless (TNPMW) **(January 2012 - May 2012)**

- Tivoli Netcool Performance Manager for Wireless (TNPMW) provides a real-time view of critical performance metrics to deliver a better overall quality of service by proactively managing the network infrastructure. It is scalable and seamlessly aggregates and correlates data from multi-vendor, multi-technology networks. I am involved in the development of predefined and customized KPIs, web-based graphical user interface, and interactive data reporting.

Software Engineering Internship, IBM Extreme BlueTM program **(June 2011 - August 2011)**

- Extreme Blue is a highly competitive placement program of IBM. We have developed a prototype for an upcoming IBM product portfolio for network service provisioning in Cloud networks. A framework for Network as a Service (NaaS) is proposed to deploy network services independent of the underlying network topology in a repeatable manner. It introduces a template-driven approach and an intelligent orchestration mechanism to provide required network services in a cloud environment in a transactional manner. It allows to provision complex network configurations, integrates a range of Tivoli products, provides real-time provisioning update. This NaaS Layer provides the concept of physical and logical modeling and automates the cloud network provisioning. My work included the integration of several IBM products - IBM Tivoli Netcool Configuration Manager (ITNCM), IBM Tivoli Network Manager (ITNM), Tivoli Netcool/OMNIBus, and Tivoli Service Automation Manager (TSAM).

Department of Computer Science, **University College Cork** **(January 2010-December 2011)**

- I served as a Research Assistant, focusing on the cost optimization of cellular backhaul traffic over satellites. This was part of a research commercialization project in collaboration with Altobridge, and funded by Enterprise Ireland). The integration of new data-centric high-speed radio access technologies within Altobridge's backhaul solutions coupled with the increased availability of content-rich services such as VoIP and Video on Demand are driving factors for advanced research in the area of data optimisation and dynamic link management. For data traffic, the diversity in application characteristics and QoS requirements demand new approaches and thinking to optimise backhaul utilisation. The main objective of this project is to research and develop a novel framework to reduce the bandwidth requirement and thus to optimise capacity utilisation in satellite backhaul. My major responsibility is to investigate potential problems with caching, pre-fetching, compression, and client-side web acceleration techniques. A novel split-proxy web acceleration framework has been developed, which is transparent to the rest of the

network elements. The performance evaluation shows along with compression, TCP optimization, and bundling mechanism reduction of end-user response time up to 27% on average can be attained even in a link which is extremely resource-scarce. As bandwidth or RTT increases, the performance benefits of using the bundling mechanism becomes more apparent.

Grameen Phone Pvt. Ltd (GP), Telenor

System Engineer, Service Network Operations (SNO) / Intelligent Network (IN) (June 2008 - December 2009)

- Design and Implementation of Call Block System (CBS) using existing feature of Ericsson Charging System 3.0. It is an innovative solution that enables GrameenPhone to offer their subscribers a network-based incoming call filtering service.
- Generation of effective statistical reports using shell script and PHP for Ericsson IN, based on the SDP KPI, CCN Counters and CDR files that help to forecast the trend of Call Types, Network congestion and Quality of Service.
- Integration, Dimensioning, and O&M of Ericsson Charging System 3.0 Nodes (SDP, AIR, INS, CCN, MINSAT, VS, IVR, EMA)
- Signaling optimization in prepaid network
- Routing control in Service Switching and Service Control nodes
- Designing, planning, and implementation of SIGTRAN network architecture for IN network
- Extensive knowledge in the area of Fault handling, Network quality & Performance, Procedures, E2E service management.

System Engineer, Operation Support System (OSS) / System Automation(SA) (January 2008 - May 2008)

- GPTTS: This in house trouble ticketing system (TTS) was developed for all sorts of N/W troubleshooting of GP. The major advantage of this TTS solution is it provides integrated workflow management for customer problem and network problem. Based upon event time, fault rectification time, every kind of Service Level Agreement (SLA) are measured. This tool is now being used in every department of GP for individual and group performance evaluation.

United Nations Development Program (UNDP)

Software Engineer(June 2007 - December 2007)

- In 2007, amidst a pivotal period for the Bangladesh Election Commission (BEC), reports surfaced about the electoral rolls containing over ten million *ghost voters*, which undermined the integrity of the electoral process. BEC, aided by the technical expertise of the United Nations Development Program (UNDP), initiated an ambitious and historically significant project: the development of a National ID system in Bangladesh. This project, a first in the nation's history, was characterized by its complexity and the urgent need to meet strict deadlines.

In this high-pressure and challenging scenario, I served as a founding member and Software Engineer on the team. My responsibilities entailed developing the voter ID software, a ground breaking initiative in Bangladesh. This project was not only technically challenging but also pivotal in establishing a foundation for transparent electoral processes within the country.

Tech used: MegaMatcher, VeriFinger, Visual C#, MySQL, MS SQL Server, MS Access and Install-Shield

REFERENCES

Professor Ali Mesbah

PhD Supervisor

Department of Electrical and Computer Engineering

The University of British Columbia

Vancouver, BC, Canada

Email: amesbah@ece.ubc.ca

Web: <https://people.ece.ubc.ca/amesbah/>

Professor Cormac J. Sreenan

Master's Supervisor

School of Computer Science and Information Technology

University College Cork
Cork, Ireland
Email: cormac.sreenan@ucc.ie

Dr. Ahmed H. Zahran
Master's Co-Supervisor
School of Computer Science and Information Technology
University College Cork
Cork, Ireland
Email: a.zahran@cs.ucc.ie