Statement of Research

H.M.N. Dilum Bandara

After starting with networking research, I transitioned to distributed systems and data engineering before focusing on software architecture and security over the past six years. I have recently focused on leveraging blockchains and Machine Learning (ML) to address data sharing, privacy, and smart contract engineering challenges. Across my 18-year journey in research, I have embraced a multidisciplinary approach, applying my expertise in areas such as supply chains; digital finance; environmental, social, and governance (ESG); fleet management; and weather monitoring.

Research Projects and Supervision – CSIRO's Data61 (2019–Present)

My role at Data61 involves delivering and leading research projects, engaging the industry, and co-supervising a PostDoc and seven PhD students from USyd, UNSW, RMIT, UoN, MQ, and UoM. Although much of our R&D endeavours remain confidential, my contributions have led to the ICSE 2023 Showcase of our ICSA 2021 paper and recognition through six Data61 awards. The impact of my blockchain, software architecture, and security research extends to several domains. This includes designing architectures for the digital uplift of water trade services in the Murray-Darling Basin, creating a base-level prototype for Australian Cotton Supply Chain Traceability, evaluating architectures for a potential Australian Central Bank Digital Currency (CBDC), conducting a blockchain security review of a key Australian financial market infrastructure, and providing expert evaluations to DISR. Moreover, my research has had educational impacts, as evidenced by the development and delivery of modules on Software Architecture for Blockchain Applications at UNSW and Digital Finance Technology at the Digital Finance CRC.

My ongoing research primarily focuses on preserving commercial sensitivity and trusted data sharing for supply chain traceability and ESG disclosure using blockchains. A pioneering solution has arisen from a research challenge we encountered during the Australian Cotton Supply Chain Traceability project, tackling the design problem of multi-ledger systems for industry ecosystems. I devised a method to model the data access needs of an industry ecosystem and an algorithm to identify optimal shared ledger combinations and allocate data elements to each ledger based on those needs. I have also identified design patterns for blockchain data migration, payment methods, and upgradable smart contracts, alongside developing a decision model for selecting patterns in the design of blockchain-based applications.

I also delve into privacy-preserving and trustworthy distributed ML. My PhD students and I have developed a model verification technique for hierarchical federated learning, a hybrid federated-split learning framework to bolster data privacy, an edge-learning framework for multi-task learning in IoT, and a framework for managing data and model provenance while ensuring copyright compliance in ML. Another project used ML to detect rug pulls and Ponzi scams in public blockchains early. Moreover, I am currently adopting Large Language Models (LLM) to promote ethical and responsible ML Operations (MLOps).

Research Supervision – University of Moratuwa (2013–2018)

I supervised or co-supervised two PhD, eight MSc by research, and 34 MSc and MBA in IT thesis students, often collaborating with industry partners or students from the industry. These efforts led to securing six grants from industry and university sources, as well as recognition through three conference paper and poster awards and four university research accolades.

My thesis students and I developed a real-time vehicular data acquisition and analytics platform, utilising cloud computing and ML, to monitor driver behaviour to support Usage-Based Insurance applications. This endeavour led to the co-founding of a startup with students and the launch of two mobile apps. Also, I secured three industry and university grants to predict fuel consumption and enhance scheduling and driver allocation within a vehicle fleet.

I collaborated with the Centre for Urban Water (CUrW), Sri Lanka to design and develop a containerised, public cloud platform for its weather and flood forecasting models. My MSc by research students designed weather data integration, assimilation, and decision support systems supporting CUrW's multimodal weather data. My attention also turned to addressing the needs of cloud providers and developers, resulting in the development of proactive, workload, resource,

policy, and cost-aware auto-scaling solutions for IaaS and PaaS clouds. These solutions incorporated an ML-based ensemble workload prediction model, cost model, and smart killing.

Another line of research involved Complex Event Processing (CEP), a crucial element in real-time IoT stream analytics and fraud detection in e-commerce. Collaborating closely with WSO2 Inc.—a leader in web services—we developed techniques to enhance their Siddhi CEP engine's performance by leveraging GPUs and multi-node scaling. Students and I also automated CEP query generation and developed a CEP engine tailored for Arduino-based edge devices in IoT.

One of my PhD students focused on software development methodologies, particularly Robotic Process Automation (RPA) and Agile practices within a multinational IT services company. Through empirical studies and case analyses, we introduced a new software process, Raban, for RPA projects. We also identified a set of software metrics for Agile projects, incorporated design thinking to align with customer expectations in Agile projects, and proposed Agile-based strategies to address software testing challenges.

Graduate Research – Colorado State University (2007–2012)

In my Master's thesis, I addressed challenges related to data transfer and energy efficiency in a wireless sensor network for monitoring subsurface chemical plumes. By integrating Virtual Sensor Networks (VSNs), clustering, and key pre-distribution, I developed a large-scale, power-efficient, and secure VSN design to enhance plume monitoring.

During my PhD, I focused on Distributed Collaborative Adaptive Sensing (DCAS) systems, where timely data processing among resource and data-rich sensors like radars is crucial. I proposed leveraging Peer-to-Peer (P2P) communities, caching, and Named Data Networking (NDN) to enable nodes with varying capabilities to collaborate effectively in complex sensing, communicating, and computing tasks. After identifying limitations in four real-world resource discovery solutions, I developed a novel multi-attribute resource and query-aware, P2P-based resource discovery solution that is scalable, efficient, and load-balanced. Moreover, I applied NDN for DCAS systems to create a distributed weather data fusion system supporting multiple users, applications, and sensors.

Future Directions

As identified through my industry engagements, a key challenge to supply chain traceability and ESG disclosure is the conflict between commercial sensitivity and transparency needs. Businesses are reluctant to share data beyond their immediate upstream and downstream partners in the supply chain. A related challenge arises in MLOps when base models are fine-tuned, transfer-learned, unlearned, or verified across multiple organisations. This challenge could be mitigated using methods such as data segregation and computations on encrypted data to varying extents. For instance, the technique I introduced to segregate data between upstream and downstream partners has been incorporated into the draft ISO/TC 298 WG3 on Traceability, Packaging, and Labeling for Rare Earths standards. Recently, I outlined a high-level design for a Fully Homomorphic Encryption (FHE) based method to calculate and verify product recycling claims. Building on this, I plan to design a framework or protocol to tackle the significant computational overhead and complex FHE configurations, using blockchain-based supply chains and ESG assurance platforms as use cases.

I advocate for addressing the conflict between commercial sensitivity and transparency by embracing a paradigm where only verifiable attestations regarding product or ML model provenance/traceability and ESG claims are shared throughout the supply chain, without divulging any underlying data. This allows businesses to publicly endorse attestation over data, offering the capability to substantiate their claims when necessary. While Zero-Knowledge Proofs (ZKP), fraud proofs, and Verifiable Credentials (VCs) offer a promising foundation, further research is imperative to adopt them effectively across diverse product or software supply chains and MLOps ecosystems with varying levels of commercial sensitivity and multiple ledgers.

I am keen to explore Computer Science (CS) education research, particularly in the areas of educational technologies to promote active learning, micro-credentials, and the use of generative AI to ensure equitable opportunities for first-year students and those from diverse economic and cultural backgrounds. My extensive experience with qualitative research methodologies—evidenced by seven of my publications using surveys, interviews, and case studies—combined with my Skills for Research teaching in the MBA in IT program and familiarity with ethics and privacy approval processes at CSIRO and UNSW, provides a strong foundation for pursuing research in CS education.