

Research Statement

Ta Nguyen Binh Duong
 School of Computing and Information Systems
 Singapore Management University
 Tel: (65) 6828-0158; Email: donta@smu.edu.sg
 Dec 2023

Background

My research strategy is mainly application-driven, i.e., the focus is on developing novel applications to exploit the full potential, and at the same time, to expose any limitations of emerging technologies. My background is firmly rooted in cluster and cloud computing, however, I believe that the decentralized computing model, currently made possible and popularized by blockchain technologies, complements very well and might gradually take over the prevalent cloud-based, centralized computing model in the near future.

Research Areas



From the cloud based computing model...

I have been spending many years (since 2003) working on several kinds of distributed computing applications. The journey started with online gaming supported first by local server clusters; and later via grid computing. Multi-player online gaming is a very demanding application due to its stringent requirements in terms of processing and network latency required for delivering a highly interactive experience. As a result, computing and network resource optimizations have been the central research problems in this area. We first developed new load sharing techniques for local gaming clusters hosting a seamless virtual environment. Our paper on this topic [1] has been well received with 139 citations till date (per Google Scholar). However, the issue of network latency mandated a geographically distributed server architecture leading to the adoption of grid computing and our formulation of a novel server placement problem, which has been published in top tier distributed computing venues, e.g., *IEEE TPDS and IPDPS* [2-4]. Our papers were also **nominated for the Best Paper Award** twice at the flagship conference for simulation research, ACM SIGSIM PADS, in 2006 [5] and 2013 [6]. I received a **Best Paper Award** in 2011 as the main author for a conference paper on virtualized game server migration [7].

Grid computing resource tends to be heterogeneous, which had made it not really suitable for time-sensitive applications like online gaming. More recently, it is cloud computing that has taken the central attention in my research agenda. Public cloud services such as AWS EC2 supply well-curated computing resource, but at a significant cost. Therefore, my focus has been on finding a good trade-off between the quality of services of applications, e.g., gaming and simulation [8-9], machine learning/big data processing [10-12], and their resource provisioning cost/energy usage. I received the **NTU Tan Chin Tuan Exchange Fellowship in Engineering**, 2017, in recognition of my work in this area. In addition, my research has been well supported by competitive funding from

the Singapore Ministry of Education (MOE) Tier-1 Academic Fund (2016, 2018, as PI), IMDA Green Data Centre Programme (2016-2019, as Co-PI), Amazon Web Services (2017, as PI), and most recently the SMU funding for practice research in efficient cloud game streaming (2021, as PI).

The consolidation of data and analytic capabilities into centralized data centers brings to wider attention co-location security issues, in which side-channel attacks could retrieve confidential data from co-located cloud virtual machines. We framed this problem as a security-resource trade-off, in which practical VM placement approaches could be implemented into existing cloud providers. The research has been published in highly reputable cloud computing venues such as IEEE CLOUD (2018) [13], Elsevier FGCS (2019) [14-15], IEEE WETICE (2020, 2021) [16-17]. Our paper on secure cloud infrastructure optimization was selected as **one of the three best papers** at the 10th IEEE Service Oriented Computing and Applications, Japan, 2017 [18].

To the decentralized computing model of the future...

As we have seen, centralized clouds pose challenging problems in terms of energy usage and data privacy. Decentralized computing has been around for a while – but with blockchain it has started to show some potential. I strongly believe that it can be an essential component in the future of computing; as it enables very large-scale cooperation and collaborative decision making:

- 1) There are inherent security features – it is by design that a single malicious participant cannot influence the system in the way he wants without consensus. As a result, there is no explicit trust relationship needed between collaborators – this enables very large scale cooperative networks.
- 2) It is privacy-preserving: you can collaborate without exposing your private data – this is a big advantage for any organizations or individuals. For instance, people have been doing blockchain-based AI in which different stakeholders can jointly develop far better AI models without the need to share their datasets or their decision making logic/processes. This is also what we did with **blockchain-based reinforcement learning agents**, 2019 [19]. This research direction is currently an on-going collaboration with the Air Traffic Management Research Institute.
- 3) The centralized computing model would still be relevant in the meantime – there are just too much high quality data accumulated at big corporations, and it is currently not exactly fast and cheap enough to transport such data around. But network speed and throughput are getting better, at the same time regulations, e.g., GDPR (EU), on data privacy are catching up with the big players – they might have to split into smaller independent but cooperative entities eventually. I believe this will form the basis to push decentralized computing into the mainstream.
- 4) Big, centralized data centers as they are today inefficient in terms of energy consumption. They also represent the bottlenecks of the world's economy – they are obvious targets of cyberwarfare. Micro-data centers at the edge, securely linked through blockchain could be an essential component of the future computing platform.

Hence the current research directions and plan...

Despite the rise of decentralized computing, it should be noted that much of the computing workload today still have to be handled on cloud data centers, especially those that require serious processing powers, storage and network performance. Recently, I have been moving towards a combination of cloud and decentralized computing by exploring and developing various interesting applications in this area. I am currently focusing on the following:

- 1) **Blockchain and AI in education.** It is my long term goal to integrate the recent advances in blockchain and AI to build a decentralized environment for supporting better collaborative, student-centric learning. **On the AI side**, my team has published work on automatic scoring of short text answers in computer science courses, 2019 [20]. Our proposal (for which I am the PI) has been **awarded for competitive funding by the Singapore Ministry of Education** (Tertiary Education Research Fund, 2022-2024). Through years of teaching undergraduates in various Singapore IHLs, I have observed that a key element of learning programming is that students have to practice on their own frequently. In this project, we develop AP-Coach, an Automatic Programming Coaching system based on a combination of AI and software engineering techniques to support students practice coding via formative feedback generation. AP-Coach implements a hybrid pedagogical strategy that embraces both constructionism and deconstructionism learning theories. To achieve this, beyond source code, AP-Coach also emphasizes the analysis of textual content in the form of natural language-like pseudo-code, which is common in learning activities such as code reading, tracing and algorithm designing. This is different from existing approaches that mainly deal with source code or textual content in isolation. Our initial result in this project has been published at IEEE TALE 2022 [25]. Recently, we worked on a complementary direction in this area, which is generating ready-to-use programming exercises for students to practice. Our work has been published at the 31st ICCE 2023 [26] and was awarded the **Best Technical Design Paper Award**.

On the blockchain side, I observe that the teaching and learning process in many university courses follows a centralized model which is fully controlled by course instructors. In particular, the instructors design learning material, deliver lectures/tutorials, conduct assessments, and assist with queries from students. In this centralized model, encouraging participatory learning and co-creation of knowledge by engaging students as contributing partners is a challenging problem. To this end, I propose a **peer-learning platform** in which students could be active contributors and true owners of the learning process. The first iteration of this project focuses on computer science courses, but extensions to other courses are possible. The proposed learning platform leverages blockchain technologies, e.g., permissioned Ethereum-based network and its smart contract model, to decentralize and incentivize student participation in learning resource development and co-creation of knowledge. My team has successfully implemented and evaluated an early prototype system, which has been presented at IEEE TALE 2021 [21]. I am currently working on a project funded by SMU Centre for Teaching Excellence (CTE) to deploy a SMU-wide blockchain network to support peer-learning in more courses.

Moving forward, I plan to integrate research outcomes from both AI and blockchain domains for the aim of better participatory learning in IHLs. For instance, we are exploring the integration of AI-based text scoring and automatic formative feedback generation capabilities into the blockchain based peer-learning platform developed in [21]. This serves two goals: 1) providing automatic assessment of relevant contributions from students which can have a wide range of qualities and contents; and 2) assisting students via formative suggestion/feedback for difficult topics and assessment tasks on the peer-learning platform.

- 2) **Agent-based Air Traffic Flow Management:** ATFM plays an important role in Air Traffic Control (ATC) systems, due to its significant impact on the efficiency and safety of air transportation. For improving local ATFM implementations, recent R&D interests emphasize the importance of **collaborative decision making** towards realizing global ATFM systems, e.g. the EU Single European Sky ATM Research (SESAR). Such collaborative decision making process is challenging to implement in ASEAN, for instance. In this project, we have successfully developed a **decentralized ATFM system prototype** on top of the permissioned blockchain Hyperledger Fabric. The system incorporates reinforcement learning techniques to carry out local optimization for a particular Flight Information Region

(FIR), and leverages Hyperledger's smart contracts to aggregate local results into global optimization strategies, 2019 [19].

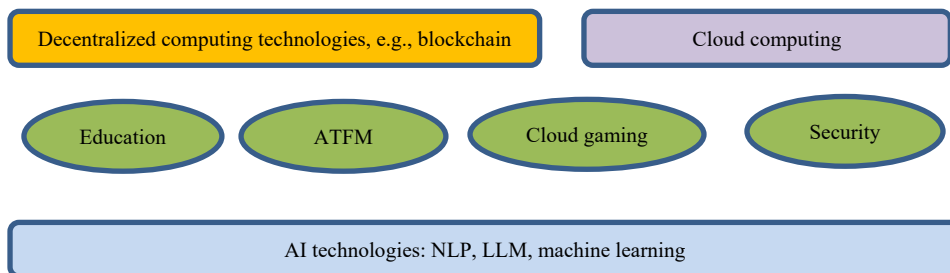
For this research, I have obtained funding as the PI from the Civil Aviation Authority of Singapore (2019), via a collaboration with the Air Traffic Management Institute (ATMRI). In the on-going collaboration, we plan to implement and deploy a distributed agent platform which can allow decentralized and effective coordination between different ATFM stakeholders in an ASEAN ATFM system.

- 3) **Decentralized cloud gaming.** Cloud gaming enables people playing graphically intensive games without the need of upgrading their hardware. At the same time, game developers can support more platforms easily and cost-effectively. It is difficult to realize cloud gaming as it requires minimal latency in server-side processing, rendering, and streaming, which are expensive in terms of resource requirements, e.g., powerful GPU servers and network bandwidth. In our [2021 cloud computing paper](#) [17], I highlighted the key issue with the current cloud game streaming model, e.g., Google Stadia, Amazon Luna, etc., which is basically the centralization of resource and game content. This led to much under-utilization, high cost, and then failure for several previous industry players, e.g., OnLive, Gaikai, etc.

I am advocating a decentralized model in which game content creators can host their own cloud solution, which can scale thanks to public cloud resource. I see the need of a reference, open-source cloud gaming system so that: 1) small, local game developers could quickly evaluate and roll their own cloud gaming solutions easily; and 2) realistic evaluations of cost saving approaches can be done. As a result, practical research and development in cloud gaming could be further advanced. To this end, we have developed and deployed the first prototype of [CloudNPlay](#), a microservices based reference cloud gaming system which is easy to extend, deploy and operate. We plan to open-source CloudNPlay so that any developers can bring their games to the cloud if needed. The development of CloudNPlay is on-going, and is supported by a grant for practical research from SMU School of Computing and Information Systems (2020-2021).

- 4) **AI-based security analysis of malware.** More recently, I have started to look at AI and its application in malware and security analysis. Our work on Android malware and security analysis has been published in [APSEC 2020, 2021 and 2022](#) [22, 23, 24]. I am currently looking at ways to conduct these analyses more efficiently, e.g., reducing the turn-around time in model training and inference, using cloud servers and the decentralized computing infrastructure, e.g., edge devices.

Part of this research direction involves the analysis of drone behaviors using machine learning techniques. For instance, when a criminal/malicious incident happens, one must analyze anomalies in drone behaviors so that those responsible for the incident are identified and held accountable. I am currently a Co-PI of a recently awarded project on auditing drone behaviors, funded by the National Satellite of Excellence - Mobile Systems Security and Cloud Security (2021, S\$594,220).



Selected Research Awards and Recognitions

Best Technical Design Paper Award, at the 31st International Conference on Computers in Education, 2023. The research was about LLM-based programming exercise generation.

Selected for the Best Paper Session at the 10th IEEE Service Oriented Computing and Applications, Japan, 2017. The research was about secure virtual machine placement in IaaS cloud infrastructures.

NTU Tan Chin Tuan Exchange Fellowship in Engineering (Outbound), 2017. Awarded to young, full-time faculty to visit and conduct research in a world-renowned institution of choice.

Best Paper Nominations, ACM SIGSIM PADS, in 2006 and 2013. The research work was about optimizing computation and network performance for distributed simulation and gaming applications.

Best Paper Award (as the main author) at the 2nd ICST/CREATE-NET Simulation and Online Gaming Workshop (DISIO), Spain, 2011. The research was about virtual machine migration techniques to support latency-sensitive distributed applications such as cloud gaming. Link: <https://pads.cs.unibo.it/doku.php?id=disio2011:bestpapers>

Selected Research Grants Obtained

1. "AP-Coach: AI-based formative feedback generation to improve student learning outcomes in introductory programming courses", Ministry of Education Tertiary Research Fund, PI (Project Level): Duong Ta, 2022, S\$200,741.
2. "CloudNPlay: An Open-Source Cloud-Native Gaming System", SCIS Seed Funding for Practice Research, SMU SCIS PI (Project Level): Don TA, 2021, S\$41,478.
3. "ADrone: Auditing Drone Behaviours for Accountability of Criminal/Malicious Activities", National Satellite of Excellence - Mobile Systems Security and Cloud Security, Co-PI (Project Level): Duong Ta, 2021, S\$594,220.
4. "LearningChain: decentralizing teaching and learning via blockchain technology", funded by NTU Edex Grants, (S\$34,059). As PI.
5. "MLAnywhere: Distributed Machine Learning on Heterogeneous Public Computing Services", MOE Academic Research Fund Tier 1 (S\$30,000), 2018. As PI.
6. "BlockAgent: Transforming Regional Air Traffic Flow Management via Blockchain based Decentralized Multi-Agent Systems", funded by Air Traffic Management Research Institute, Singapore (S\$120,000). As PI.
7. "FogGrid: Transforming Microgrid Operations via Blockchain and Fog Computing in Singapore", funded by Singapore Energy Market Authority (S\$2,607,120), 2018-2021. As Co-PI (S\$150,000 allocated).
8. "Toward joint IT-thermal optimization to improve energy efficiency for high-ambient temperature data centre in the tropics via learning-based algorithms", Info-communications Media Development Authority of Singapore – Green Data Centre Programme (S\$1,191,600), 2016-2019. As Co-PI (\$150,000 allocated).
9. "Towards truly on-demand resources for latency-sensitive services and applications on public clouds", MOE Academic Research Fund Tier 1 (S\$98,505), 2016-2018. As PI.
10. "QoS-aware cloud resource provisioning", Amazon Web Services Research Grant (US\$6,800), 2016-2017. As PI.
11. "SEFA: a secure system for executing untrusted code on clouds, to support interactive self-learning of computer programming via fast, automatic program evaluation", NTU Edex Grants (\$18,230), 2016-2017. As PI.

Selected Publications

- [1] Ta Duong, S. Zhou: A dynamic load sharing algorithm for massively multiplayer online games. ICON 2003: 131-136.
- [2] Ta Duong, S. Zhou: Efficient client-to-server assignments for distributed virtual environments. IPDPS 2006.
- [3] Ta Duong, S. Zhou: A two-phase approach to interactivity enhancement for large-scale distributed virtual environments. Computer Networks 51(14): 4131-4152 (2007).
- [4] Ta Duong, S. Zhou, et al.: Interactivity Constrained Server Provisioning in Large-Scale Distributed Virtual Environments. IEEE Trans. Parallel Distrib. Syst. 23(2): 304-312 (2012).
- [5] Ta Duong, S. Zhou and H. Shen: Greedy Algorithms for Client Assignment in Large-Scale Distributed Virtual Environments, ACM SIGSIM PADS, 2006.
- [6] Z. Li, X. Li, Ta Duong, W. Cai and S. Turner: Accelerating Optimistic HLA-based Simulations in Virtual Execution Environments, ACM SIGSIM PADS, 2013.
- [7] Ta Duong, S. Zhou, et al.: A Virtualization-based Approach for Zone Migration in Distributed Virtual Environments, 2nd Distributed Simulation & Online Gaming Workshop, Barcelona, Spain, 2011.
- [8] Ta Duong, et al.: QoS-Aware Revenue-Cost Optimization for Latency-Sensitive Services in IaaS Clouds. DS-RT 2012.
- [9] Ta Duong, et al.: A Framework for Dynamic Resource Provisioning and Adaptation in IaaS Clouds. CloudCom 2011.
- [10] P. Sun, Y. Wen, Ta Duong, and X. Xiao: GraphMP: I/O-Efficient Big Graph Analytics on a Single Commodity Machine. IEEE Trans. Big Data 6(4): 816-829 (2020)
- [11] Ta Duong, "FC2: cloud-based cluster provisioning for distributed machine learning", Cluster Computing, Springer, 2019.
- [12] P. Sun, Y. Wen, Ta Duong, and H. Xie, "MetaFlow: A Scalable Metadata Lookup Service for Distributed File Systems in Data Centers", IEEE Transactions on Big Data 4(2): pp. 203-216 (2018).
- [13] Ta Duong, N. Pimpalkar: Handling Co-Resident Attacks: A Case for Cost-Efficient Dedicated Resource Provisioning. IEEE CLOUD 2018.
- [14] A. Agarwal, Ta Duong: Secure Virtual Machine Placement in Cloud Data Centers, Elsevier Future Generation Computer Systems, 2019.
- [15] A. Agarwal, Ta Duong: Co- Location Resistant Virtual Machine Placement in Cloud Data Centers. ICPADS 2018.
- [16] L. Vu, Ta Duong: Group Instance: Flexible Co-Location Resistant Virtual Machine Placement in IaaS Clouds. IEEE WETICE 2020.
- [17] A. Wibowo, Ta Duong: CloudNPlay: Resource optimization for a cloud-native gaming system, IEEE WETICE 2021.
- [18] V. Natu, Ta Duong: Secure Virtual Machine Placement in Infrastructure Cloud Services. IEEE SOCA 2017.
- [19] Ta Duong, K. Todi, Hong-Linh Truong: Decentralizing Air Traffic Flow Management with Blockchain based Reinforcement Learning, IEEE Industrial Informatics, 2019.
- [20] A. Prabhudesai, Ta Duong: Automatic Short Answer Grading using Siamese Bidirectional LSTM Based Regression. IEEE TALE 2019: 1-6
- [21] Ta Duong, Joel Yang: EtherLearn: Decentralizing Learning via Blockchain, IEEE TALE, 2021.
- [22] Lwin Khin Shar, Ta Duong, et al.: SmartFuzz: An Automated Smart Fuzzing Approach for Testing SmartThings Apps, APSEC 2020.
- [23] Lwin Khin Shar, Ta Duong, David Lo: Empirical evaluation of minority oversampling techniques in the context of Android malware detection, APSEC 2021.
- [24] Lwin Khin Shar, Minn Wei, Ta Duong: DronLomaly: runtime detection of anomalous drone behaviors via log analysis and deep learning, APSEC 2022
- [25] Ta Duong, Lwin Khin Shar, Venky Shankararaman: AP-Coach: formative feedback generation for learning introductory programming concepts, IEEE TALE 2022.
- [26] Ta Duong, Phuc Nguyen, Swapna Gottipati: ExGen: Ready-To-Use exercise generation in introductory programming courses, ICCE 2023.