

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/154447/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Patros, Panos, Ooi, Melanie, Huang, Victoria, Mayo, Michael, Anderson, Chris, Burroughs, Stephen, Baughman, Matt, Almurshed, Osama, Rana, Omer, Chard, Ryan, Chard, Kyle and Foster, Ian 2023. Rural AI: Serverless-powered federated learning for remote applications. IEEE Internet Computing 27 (2), pp. 28-34. 10.1109/MIC.2022.3202764

Publishers page: <http://dx.doi.org/10.1109/MIC.2022.3202764>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Rural AI: Serverless-Powered Federated Learning for Remote Applications

Panos Patros^{*§}, Melanie Ooi^{*}, Victoria Huang^{*¶}, Michael Mayo^{*}, Chris Anderson^{*}, Stephen Burroughs^{*}, Matt Baughman[‡], Osama Almurshed[‡], Omer Rana[‡], Ryan Chard[‡], Kyle Chard[‡], and Ian Foster[‡]

Abstract—With increasing connectivity to support digital services in urban areas, there is a realization that demand for offering similar capability in rural communities is still limited. To unlock the potential of Artificial Intelligence (AI) within rural economies, we propose Rural AI—the mobilization of serverless computing to enable AI in austere environments. Inspired by problems observed in New Zealand, we analyze major challenges in agrarian communities and define their requirements. We demonstrate a proof-of-concept Rural AI system for cross-field pasture weed detection that illustrates the capabilities serverless computing offers to traditional federated learning.

Index Terms—Serverless computing, cyber-physical infrastructure, computing continuum, federated learning, Rural AI.

I. WHY RURAL AI?

Unequal infrastructure between urban and rural areas—a consequence of disparities in the economics of making such services available—is likely to increase further the gap between those areas’ economic potential [1]. We propose a new discipline of Rural AI, which we define as *the engineering of cyber-physical systems for enabling sovereign, sustainable AI in locations with limited and/or unreliable power/networking infrastructure*. Fig. 1 illustrates an example Rural AI environment for precision agriculture in Aotearoa New Zealand (NZ). Such intelligent systems obtain sparse but valuable sensor data across multiple locations to predict resource needs, adapt to intermittent connectivity, and preserve data privacy. Thus, local software agents need to *collaborate on shared models without sharing data*, e.g., via federated learning [2]. We envision the formation of local cooperatives that benefit from data protection and from aggregation of resources on individual farms.

Rural AI applications require orchestration of computational tasks across the **computing continuum** [3]—a transparent computing fabric unifying cloud, HPC, and the edge—in rural environments. We make the following contributions:

- Introduce the discipline of Rural AI to bridge technical and economic requirements elicited from surveying rural NZ stakeholders.
- Demonstrate an experimental system highlighting the efficacy of federated learning and serverless computing for Rural AI applications.
- Motivate the use of serverless computing and federated learning for future Rural AI applications.

The authors are with: University of Waikato (NZ)^{*}, Cardiff University (UK)[†], University of Chicago and Argonne National Laboratory (USA)[‡], Raygun Performance Monitoring[§], and National Institute of Water and Atmospheric Research (NZ)[¶].











Fig. 1. Rural AI can sustainably support sovereign precision agriculture.

We first present seven NZ use cases and nine Rural AI characteristics we elicited in collaboration with local stakeholders in NZ. Focusing on a Robotic Weeds Removal example, we highlight tradeoffs between computing on the cloud, edge, or a hybrid approach on the computing continuum.

We develop a prototype system building on federated Function as a Service. Our experiments indicate that cloud-only learning exhibits the highest inference accuracy; however, it risks data sovereignty and fault tolerance, and has limited technical feasibility. Local-only (at the edge) learning suffers from low model accuracy, lacking the cooperative federation of Rural AI, which tolerates low node participation but converges slower than cloud-only. Choosing an approach depends on: (i) maturity of computational infrastructure; (ii) complexity of analysis required; (iii) types of end-user communities (farmers, regional planners, local authorities, etc.).

II. RURAL AI DESIGN AND REQUIREMENTS ANALYSIS

We show seven Rural AI examples in Fig. 2—all funded projects of national interest in NZ. Their requirements can be broadly categorized to the nine shown in Table I. Referring to these examples, it is clear that cooperative federation, which involves the exchange and sharing of information among common stakeholders, is needed to improve performance of machine learning models based on data from many clients.

Robotic Weeds Removal 	Robots scan for produce, pests, disease. Data- and IoT-driven decisions sent to robots that actuate. GreenTech Robotics offer an autonomous weeding robot ^a .
Wireless Telecom Ops 	Local and regional usage analyzed to improve QoS. Provides the basis for other Rural AI applications. Lightwire offer rural broadband in NZ's Waikato region ^b .
Forester Biomonitoring 	Foresters equipped with smart vests to monitor health and emit alerts for themselves and nearby workers. Hakituri are designing a forestry smart vest ^c .
Robotic Pollination 	Robot-fleets detect, precision-spray flowers. Flowering windows are 1-2 weeks: rapid adaptation by federating knowledge from nearby farms. RoboticsPlus offer a broad pollen-delivery system ^d .
Active Biosecurity 	To protect NZ birds from invasive mammals, Cacophony deploy smart traps that collectively and adaptively learn to activate only for invasive specimens ^e .
Renewable Microgrids & Industrial Digital Twins 	Local AI to manage, exchange and store of renewable energy. Ahuora decarbonize industrial sites with renewable sources in their vicinity, using local adaptive digital twins ^f .
Environmental Monitoring 	Science missions maintain equipment and manually transfer data. A smart sensor grid de-risks operations. Waikato Environmental Research Institute conduct frequent missions in caves and estuaries ^g .

^agreentechrobotics.com/products/weedspider ^bwww.lightwire.co.nz
^csdb.cms.waikato.ac.nz/research-projects
^dwww.roboticsplus.co.nz/quaddusterecacophony.org.nz
^fahuora.co.nz ^gwww.waikato.ac.nz/eri

Fig. 2. Rural AI Applications: Seven NZ Examples

Increasingly mission-critical autonomous systems must be able to self-navigate safely even when geo-location services are down: agricultural robots must continue performing AI tasks at brownout quality and upload local updated models upon reconnection.

Given these observations, we define Rural AI requirements as follows:

Data sovereignty: NZ has strict data sovereignty laws, and in particular, those of the nation's indigenous peoples. For instance, Māori data sovereignty means that Māori must have control over their data. This applies to information about native plant species, agriculture, health, etc. Collected data must be stored on the edge and deleted as soon as possible. Only trained models that obfuscate information should be communicated. This enables group knowledge sharing between competitors or consumers while preserving the privacy of information for individuals. Additionally, recent literature

TABLE I
THE RURAL AI REQUIREMENTS MUST BE MET CONCURRENTLY

Adaptive AI Functionality	Cyber-physical and mission-critical applications and the underlying platform need to handle concept drifts, uncertainty, instrument variability, unpredicted changes in the system and environmental fluctuations.
Ruggedness and Resilience	Deployment may take place on remote/ isolated areas with unreliable power, limited networking infrastructure and adverse environmental conditions.
Cooperative Federation	Systems that are locality-correlated share nearby experience to improve performance.
Big data	Risky and cost-prohibitive to securely move data manually. Training of models needs to take place incrementally to avoid saturating resources.
Security and Data Sovereignty	Guarantees that must keep sensitive raw data from being shared to the cloud—economic sovereignty is crucial for industrial Rural AI.
Compute and Network Heterogeneity	Varying embedded computing infrastructure, such as GPGPUs and FPGAs. Varying networking requirements, such as WiFi, LoRaWAN, satellite and 5G/6G.
Power- and Energy-Awareness	Unreliable energy availability due to renewable source fluctuations and power-competition between domain and computing/ networking devices.
Human User and Developer Barriers	Easy-to-use systems by non-technical people. Software engineers to easily develop, test and deploy Rural AI solutions.
Sustainability and Cost Effectiveness	Rural AI systems should be cost effective to develop and operate vs. cloud. Local renewable power sources and executing AI on the edge.

demonstrates significant skepticism in the agricultural community [4] and, as such, it is important to respect those sensitivities to ensure wide-ranging applications.

Federated Learning: Initial training is undertaken on edge devices with local data before aggregating models within fog/cloud nodes. Aggregated models can be transferred back to the edge devices for inference/actuation.

Federated Cyberinfrastructure: In-field devices could comprise the first layer of the computing continuum. These devices (e.g., pesticide spraying robots) could then decide (a) when to perform the next round of federated learning by submitting recently captured data, (b) which field-side unit(s) to communicate with, and (c) at what granularity to relay their data.

Scheduling strategy: Self-adaptive methods to determine where and when to distribute workloads are needed to optimize deployment objectives and provide guarantees under uncertainty. Decision making must address conflicting metrics, such as time-cost, geographic proximity and computing/ bandwidth/ energy availability tradeoffs.

III. SERVERLESS ARCHITECTURE FOR RURAL AI

We developed a prototype Rural AI platform based on the requirements in the previous section. Our architecture uses funcX [5], a federated Function as-a-Service (FaaS) platform that enables function execution on remote compute endpoints. FuncX supports reliable execution across heterogeneous computing resources—ideal for heterogeneous computational resources of rural areas. Our prototype (Fig. 3) comprises one node as the federation server and three edge cluster nodes, each with 4 VMs, which we then used for our use case.

Serverless computing is a computing paradigm pioneered by cloud providers [6]. Its aim is to allow software developers to focus on the development side without having to manage

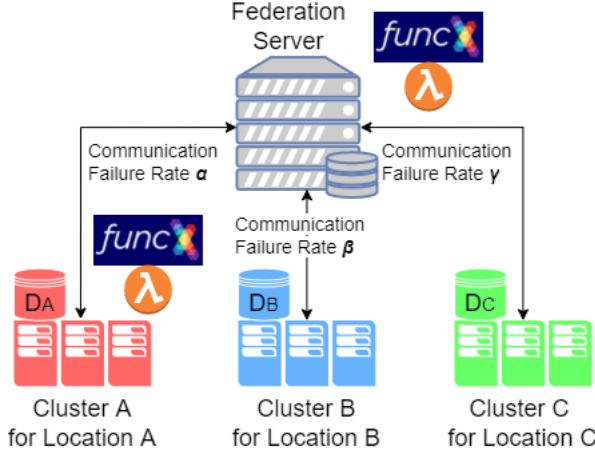


Fig. 3. Federated FaaS Architecture using funcX

physical or virtual servers. For example, in the function as a service (FaaS) model, developers define a programming function; they may then invoke that function via an API by passing the function ID and input arguments to the FaaS platform. The FaaS platform is responsible for provisioning computing resources, executing the function, and relaying results to the user.

Others have explored the use of serverless models for IoT scenarios [7]: cloud providers, such as AWS Greengrass and Microsoft Azure IoT Edge, provide methods to bridge between public clouds and edge devices. However, to the best of our knowledge, such methods have not been used to deliver with federated learning applications in rural settings.

A. Weed Detection Use Case

We evaluated the fitness of our prototype Rural AI platform for weeds detection—a representative application with generalizable observations. Pastures require analysis of highly localized datasets that benefit from cooperative training (Fig. 4).

We used hyperspectral pasture images from three different sites [8] labelled with four classes—three different weeds and a background class.

Dataset: The dataset contains 104,544 labelled 148-channel hypervoxels drawn from three 900nm–1700nm infrared spectrum images, one image per site; 60,072 hypervoxels were sampled from Site A, 30,240 from Site B, and 6,232 from Site C. Sites A and B data were balanced across all four classes; Site C data contained one weed species and the background class. We randomly split the dataset into training and testing with an 80 : 20 ratio. Using linear discriminant analysis in a non-federated (centralised) setting as a baseline yields 91% test accuracy.

Approach: Based on initial experiments and a non-exhaustive architecture search, we selected the following ANN architecture: Two 1D convolutional layers with max-pooling after each, flattened and fed into a dense classifier; ReLU activations for the convolutions and softmax for the classifier.

We implemented a federated averaging [9] algorithm that divides the model training process into T iterations. A random

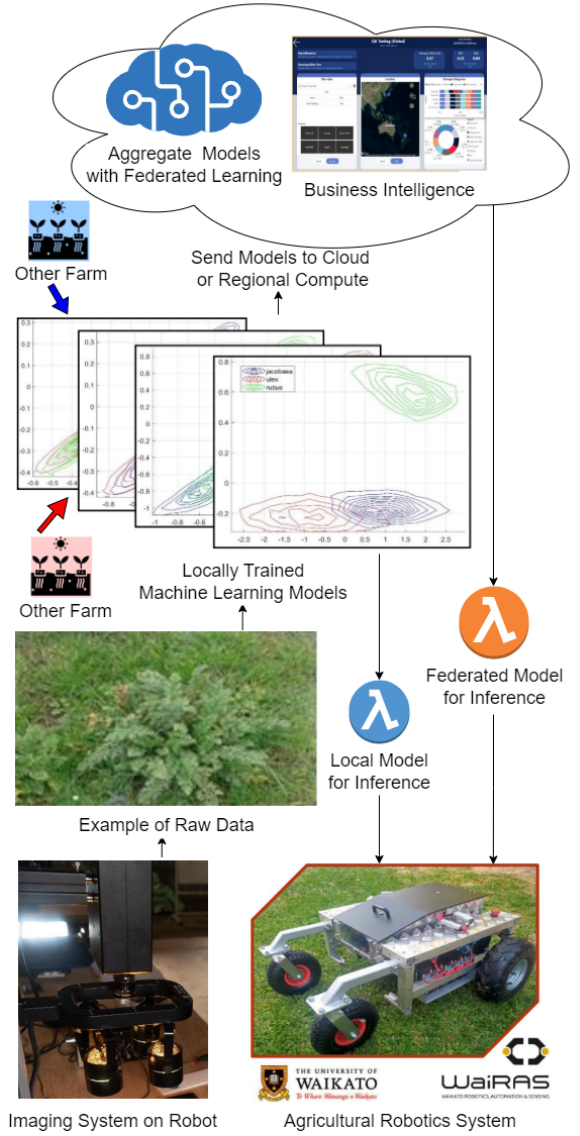


Fig. 4. Exemplar Rural AI Application: Near-infrared hyperspectral imaging system to distinguish weeds from produce [8].

subset of local models is selected and E epochs of training are performed on their local data with minibatch size B .

Both hyperparameters E and B must be optimized: we experimented with $E \in \{1, 5, 10\}$ and $B \in \{10, 50\}$ and found that $E = 1$ and $B = 50$ performed optimally. When each local training iteration is complete, the local sites return their models to the federation server for model aggregation. In the original algorithm [9], a key parameter is N —the proportion of local models uploaded to the cloud server per iteration. In a rural setting, this would correspond to the fraction of local nodes available at any one time. Therefore, we also evaluate different N values ranging from 100% to 25%. After aggregation, each local model is replaced with the downloaded new model.

B. Experimental Results

We demonstrate the benefit of using federated learning vs. building independent models based on locally captured image

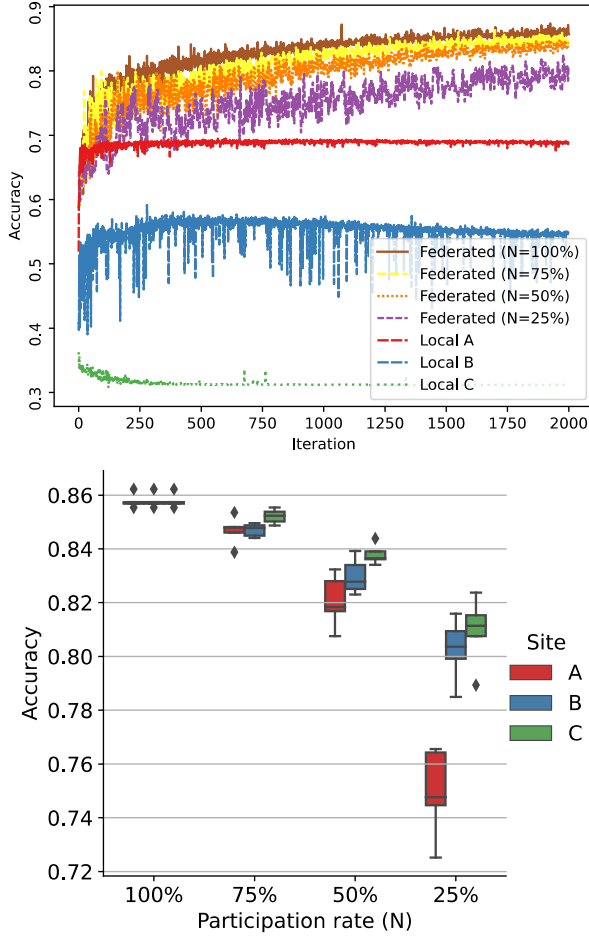


Fig. 5. Top: Accuracy of models trained using local-only and federated (Rural AI) approaches over 2,000 iterations. N , the fraction of local nodes available at any one time. Bottom: Model accuracy with different participation rates N of different local nodes.

data. As illustrated in Fig. 5, three local models (A, B, C) are developed on different sized datasets: accuracy can vary from 30% (Local C) to 70% (Local A).

Compared to cloud-only (98% accuracy) and local-only (30%-70%), our full-participation federated model ($N = 100\%$) reached an accuracy of 85% after 2,000 iterations—but all federated models exhibit slower convergence rate.

Although cloud-only achieves higher accuracy by training a global model using *all* the data from *all* three sites, it requires: (i) data to be shared with the datacenter—incurring additional latency, bandwidth and communication time during the data transfer process; (ii) violates sovereignty; (iii) suffers from sustainability, scalability and synchronization issues.

The federated models' accuracy exhibit higher variance with smaller N values due to the reduced model sharing. Nevertheless, all federated models show consistently increasing accuracy. Even with low local-node availability ($N=25\%$), the federated model outperformed the best local model, achieving 75%. Hence, we hypothesize that *nodes can adaptively, voluntarily forgo participating in training rounds to save energy without penalizing the federated models' performance.*

To test our hypothesis, we simulate different nodes volun-

tarily participating at different rates for each training round. As shown in Fig. 5, sites with more samples have a higher impact on accuracy if they only partially join the training. Nevertheless, the federated model still outperformed the best local model. Meanwhile, for sites with fewer samples (e.g., B and C), they can significantly benefit in terms of model accuracy and energy savings even when only participating in 25% of training rounds.

IV. THE FUTURE OF RURAL AI

We outline in Table I key research challenges in Rural AI, considering constituent technologies and research areas.

Agriculture: To support *cooperating farms*, self-adaptive distributed systems are needed to connect computing resources with those on the farm and progressively moving to the cloud. To support interoperability between farm systems, new specification languages, metamodels and models-at-runtime are needed. With an increase in capability must also follow an increase in usability. A possible future solution is to integrate serverless federated learning with new edge-based farm management architectures [10].

Artificial Intelligence: At the time of writing, there are 27 ISO/IEC standards on AI under development to address data quality, management, applications, etc. Adaptive AI algorithms will be required in planning (e.g., robotic automation of orchards [11]), natural language processing (e.g., ease of control of complex farming systems), and heuristic optimization (e.g., developing plans for robotic harvesters). More generally, there is no guarantee that algorithms developed and evaluated in one rural region will work effectively in another region due to the likelihood of the data from the new region not following the same distribution as that used to train and build the AI systems. Climate change presents further complications, and can be viewed as an potential instance of concept drift. Both systems will necessitate extensive and continuous model testing well beyond initial deployment.

Computing continuum: The increase in compute capabilities is mirrored readily by networking infrastructure. While not always the case, as networking increases with compute, adaptive task placement across different devices becomes more pressing as the resources are no longer obfuscated by network bottlenecks. With Rural AI, we envision the ultimate computing continuum—seamlessly combining resources at every level (including local renewables) to dynamically, adaptively, and autonomously execute tasks where and when required.

Serverless computing: As our prototype shows, serverless computing [12] provides a promising direction for delivering a Rural AI platform and offers an intuitive programming model to develop and execute AI applications. Extending FaaS to federated environments is necessary to support Rural AI, enabling functions to be dispatched to different devices in the computing continuum based on resource availability, performance, data privacy, and location and scheduling requirements.

V. CONCLUSION

Using federated learning to support independent model construction and subsequent aggregation provides significant benefits for Rural AI. Using this approach, a robot can capture

data in its proximity using on-board sensors, develop a model based on this data, and share this with field-side units to carry out model “stitching”. In an environment with limited network and computational resource availability, this approach provides greater flexibility to sustainably support decision making and actuation (e.g., weed spraying). Thus, considering the tradeoffs between local-only (edge computing) vs. cloud-only (cloud computing), Rural AI architectures offer a valuable middle-ground, hybrid solution that can achieve better performance, fault tolerance and data sovereignty without sacrificing accuracy or increasing costs and power/energy requirements.

ACKNOWLEDGMENT

This work was supported by the Royal Society of NZ Te Apārangī through a Rutherford Discovery Fellowship and by Te Ipu o Te Mahara AI Institute. We thank Dr. Simpkin from Callaghan Innovation for assistance with the hyperspectral imaging system.

REFERENCES

- [1] C. Richardson, “Wind River Brandvoice: Three changes that will make rural 5G access a reality. the dream of a workforce empowered to thrive no matter where it lives,” Feb 2021. [Online]. Available: <https://www.forbes.com/sites/windriver/2021/02/22/three-changes-that-will-make-rural-5g-access-a-reality-the-dream-of-a-workforce-empowered-to-thrive-no-matter-where-it-lives/>
- [2] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, “A survey on federated learning,” *Knowledge-Based Systems*, vol. 216, p. 106775, 2021.
- [3] R. Kumar, M. Baughman, R. Chard, Z. Li, Y. Babuji, I. Foster, and K. Chard, “Coding the computing continuum: Fluid function execution in heterogeneous computing environments,” in *IEEE International Parallel and Distributed Processing Symposium Workshops*, 2021, pp. 66–75.
- [4] L. Wiseman, J. Sanderson, A. Zhang, and E. Jakku, “Farmers and their data: An examination of farmers’ reluctance to share their data through the lens of the laws impacting smart farming,” *NJAS-Wageningen Journal of Life Sciences*, vol. 90, p. 100301, 2019.
- [5] R. Chard, Y. Babuji, Z. Li, T. Skluzacek, A. Woodard, B. Blaiszik, I. Foster, and K. Chard, “FuncX: A federated function serving fabric for science,” in *29th International Symposium on High-Performance Parallel and Distributed Computing*, 2020, pp. 65–76.
- [6] J. Schleier-Smith, V. Sreekanti, A. Khandelwal, J. Carreira, N. J. Yadwadkar, R. A. Popa, J. E. Gonzalez, I. Stoica, and D. A. Patterson, “What serverless computing is and should become: The next phase of cloud computing,” *Communications of the ACM*, vol. 64, no. 5, pp. 76–84, 2021.
- [7] R. Wolski, C. Krintz, F. Bakir, G. George, and W.-T. Lin, “Cspot: Portable, multi-scale functions-as-a-service for iot,” in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, 2019, p. 236–249.
- [8] W. Holmes, M. P.-L. Ooi, M. Look, Y. C. Kuang, R. Simpkin, D. Blanchon, and S. Demidenko, “Proximal near-infrared spectral reflectance characterisation of weeds species in New Zealand pasture,” in *IEEE International Instrumentation and Measurement Technology Conference*, 2019, pp. 1–6.
- [9] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *20th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 54, 2017, pp. 1273–1282.
- [10] A. Jandl, P. Frangoudis, and S. Dustdar, “Edge-based autonomous management of vertical farms,” *IEEE Internet Computing*, vol. 26, no. 1, pp. 68–75, 2022.
- [11] R. Verbiest, K. Ruysen, T. Vanwalleghem, E. Demeester, and K. Kellens, “Automation and robotics in the cultivation of pome fruit: Where do we stand today?” *Journal of Field Robotics*, vol. 38, no. 4, pp. 513–531, 2021.
- [12] S. Kounev *et al.*, “Toward a definition for serverless computing,” *Dagsuhl Reports*, 2021.

BIOGRAPHIES

P. Patros (PhD UNB, CA): Principal Engineer at Raygun, developing front- and back-end performance monitoring and crash reporting tools. Previously, Senior Lecturer, University of Waikato. Chartered Professional Engineer (CPEng) and IEEE Member. Email: patros.panos@gmail.com.

M. Ooi (MIEEEE’05-SMIEEEE’12, FIET, C.Eng. PhD Monash, AUS): Rutherford Discovery Fellowship recipient, (NZ), Assoc. Professor in Mechatronics Engineering, University of Waikato and Adjunct Professor in Computer Science, Sunway University. Chartered Engineer (UK). Email: melanie.ooi@waikato.ac.nz

V. Huang (PhD VUoW, NZ): Data Scientist at National Institute of Water and Atmospheric Research, NZ. Previously, Research Fellow at University of Waikato, NZ. Email: victoria.huang@niwa.co.nz

M. Mayo (PhD Cant, NZ): Senior Lecturer and member of the AI Institute, University of Waikato. Email: michael.mayo@waikato.ac.nz

C. Anderson (MSc UoW, NZ) and S. Burroughs (MSc UoW, NZ): PhD candidates, University of Waikato. Emails: zsmyna@gmail.com, spb23@students.waikato.ac.nz

M. Baughman (MSc UChicago, USA): PhD candidate, University of Chicago, USA. Email: mbaughman@uchicago.edu

O. Almurshed (MSc UoE, UK): PhD Candidate, Cardiff University, Wales, UK. Email: AlmurshedO@cardiff.ac.uk

O. Rana (PhD Imperial College, UK): Professor of Performance Engineering at the School of Computer Science & Informatics, Cardiff University, Wales, UK. Email: ranaof@cardiff.ac.uk

R. Chard (PhD VUoW, NZ): Researcher working with Argonne National Laboratory and University of Chicago, USA. Email: rchard@anl.gov

K. Chard (PhD VUoW, NZ): Research Associate Professor, University of Chicago with a joint appointment at Argonne National Laboratory, USA. IEEE and ACM Member. Email: chard@uchicago.edu

I. Foster (PhD Imperial College, UK): Argonne National Laboratory and University of Chicago, USA. Email: foster@anl.gov