EDITORS: David Abramson, david.abramson@uq.edu.au
Manish Parashar, manish.parashar@utah.edu

DEPARTMENT: CASE STUDIES IN TRANSLATIONAL COMPUTER SCIENCE

# Translational Edge and Cloud Computing to Advance Lake Water Quality Forecasting

Renato J. Figueiredo, *University of Florida, Gainesville, FL, 32611, USA*
Cayelan C. Carey, and R. Quinn Thomas, *Virginia Tech, Blacksburg, VA, 24061, USA*

***In this article, we report on our experiences with interdisciplinary projects at the intersection of freshwater ecology, data science, and computer science. The translational research process has progressively led to the development of distributed systems that apply both edge computing and function-as-a-service (FaaS) cloud computing to support end-to-end water quality forecasting workflows across the edge-to-cloud continuum.***

The environmental sciences are being transformed by an increasing variety, velocity, and volume of data streaming from sensors and Internet-of-Things (IoT) devices, and there is a growing need to access and extract information in real-time to enable transformative applications.[1] In particular, ecological forecasting is poised to significantly increase predictive capacity for effective decision-making with broad societal impact, including improving water quality in lakes and reservoirs.[2,3] Water quality forecasting requires a robust and scalable cyberinfrastructure (CI) with end-to-end workflows encompassing sensors and edge and cloud computing for real-time decision support.[4] However, there are often significant barriers that prevent ecological forecasters from making effective use of available physical CI capacity (e.g., edge and cloud computers).

The effort required in software development, deployment, and resource management makes adopting end-to-end forecasting workflows by the ecology research community challenging. Thus, while new advances in CI (e.g., serverless computing[5]) hold much promise for forecasting applications, they need to be grounded in realistic use cases and deployments to validate their potential. We have approached these challenges through translational research uniquely enabled by multiple years of collaboration at the intersection of two disciplines, which has resulted in the development of a scalable forecasting system now being deployed at lakes around the world.[6,7]


Digital Object Identifier 10.1109/MCSE.2024.3430148
Date of current version 20 November 2024.

## TRANSLATIONAL PILLARS

Translational computer science (TCS) focuses on three pillars[8]: *laboratory, locale, and community*. We extend the analogy by describing these in the context of ecological forecasting. In our experience, the *laboratory* pillar has been multifaceted: It includes the ecosystems where sensors are deployed, a physical laboratory where water samples are analyzed, as well as a virtual laboratory where software is developed for data capture, transfer, and model execution. The *locale* pillar has also been multifaceted: It includes physical CI deployed at/near a lake (e.g., water quality sensors and edge computing gateways) as well as virtual CI (e.g., containers, virtual machines and networks, containers, and cloud storage). The third pillar, *community*, has included early adopters (our research team, which includes students, postdoctoral researchers, technicians, and faculty), managers, and community members. By building upon these pillars, our translational research workflow has led to basic research contributions as well as to the development of new practices—backed by field deployments and data—that are leading state-of-the-art innovation in ecological forecasting, illustrated in Figure 1.

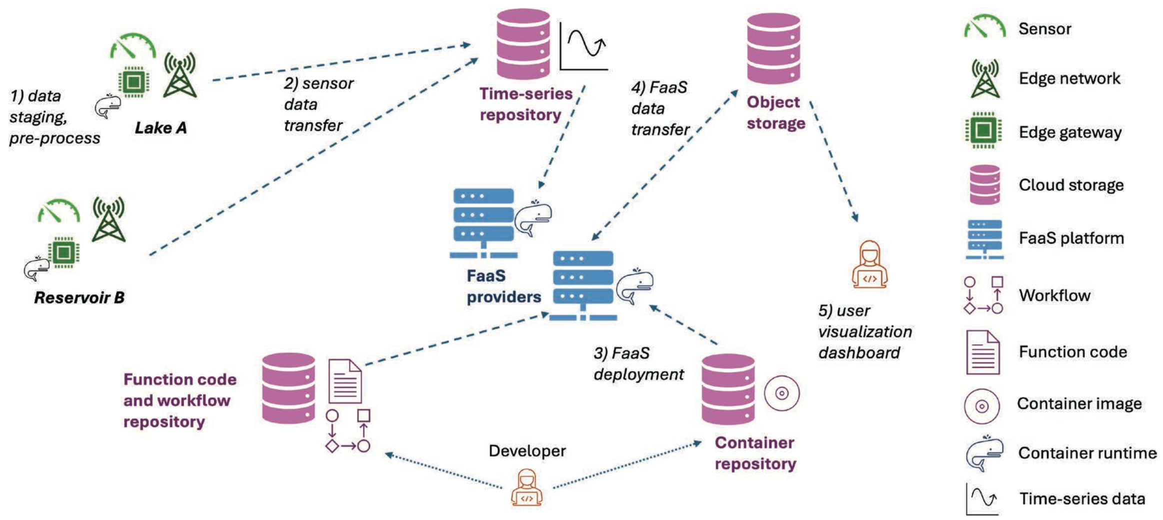


**FIGURE 1.** Overview of water quality forecasting workflows resulting from the translational process. 1) Sensor data are collected and preprocessed at the edge, then 2) time-series data are transferred to Git repositories and made available to 3) event-driven functions deployed as containers by FaaS cloud providers. FaaS functions use S3 storage for 4) data transfers throughout the workflow and for 5) visualization.

## TRANSLATIONAL PROCESS

A question we have been asked many times when describing our collaboration is: How did it all start? The seeds were planted during in-person meetings of the Pacific Rim Applications and Grid Middleware Assembly (PRAGMA) and Global Lake Ecological Observatory Network ()GLEON networks starting in 2014 that provided forums for interdisciplinary cross-pollination. The opportunities for social and technical interactions created in these grassroots meetings have been crucial to interdisciplinary exchange as well as the building of mutual trust.[9] Nurturing these connections requires willingness to listen with curiosity to learn new concepts, terminology, and jargon associated with research questions with which one is initially unfamiliar.

### Laboratory: Basic Research

Throughout our projects, basic research questions have emerged from both computer systems and ecology domains. The computer systems/CI research questions focused on the development of virtualization applied to edge and cloud computing to reduce the complexity associated with the deployment of end-to-end forecasting workflows. Because these workflows involve both edge computing (in situ data collection and staging) and cloud computing (forecasting models and data assimilation), a core research question was: How can virtual network systems suitable for both edge and cloud computing environments leverage software-defined networking (SDN)? In particular, we were constrained by our locale at the physical location of a freshwater ecosystem, where edge devices have limited compute and storage capacity, have limited power, and use commodity cell phone Internet with private, nonroutable Internet addresses subject to network address translation (NAT). Addressing this question led to development of EdgeVPN,[10,11] a novel SDN-based overlay virtual network integrating NAT-constrained edge and cloud devices end-to-end.

Another core computer systems research question was: How can serverless computing be harnessed to support end-to-end ecological forecasting workflows, while presenting an accessible interface for users and developers in this domain? This challenge led to the development of Function-as-a-Service for R (FaaSr), a novel middleware supporting serverless, cross-platform workflows written in R.[12]

For the ecologists, research used the end-to-end forecasting system to quantify fundamental controls of ecosystem predictability,[6] determine the relative forecastability of different freshwater variables,[13] and identify the dominant sources of uncertainty in ecosystem forecasts.[14] In addition to basic ecology questions, the team has also conducted "experiments" with different configurations of the freshwater forecasting system to answer methodological questions on the optimal frequency of data assimilation[15] and model structure.[16]

### From Laboratory to Locale

The need for interdisciplinary collaboration and practical water quality forecasting led the team to address these research questions beyond the confines of individual laboratories. Starting from a foundation of trust and shared knowledge from technical exchanges and team-building social interactions, our process of translational research was based on iterative cycles of codesign and coproduction to develop prototypes suitable for deployment in practice. Key aspects of these cycles have been as follows: 1) field visits and practical demonstrations helped team members build intuition and a sense of possibilities (and challenges) in cross-domain work. For example, in reservoir field visits, the CI experts learned about environmental sensors and data they produce, as well as the practical constraints on power, compute, and networking of edge devices in the field. Conversely, the ecology and forecasting domain experts learned about virtualization, containers, and cloud computing by exploring these systems in hands-on tutorials that were specifically developed for the ecologists on the team, which gave an appreciation of how these software approaches work, as well as their limitations. 2) A culture of inclusion of different perspectives and voices across the spectrum of team members (undergraduate and graduate students, technicians, post-docs, and faculty) in all-hands meetings. This culture enabled all members to participate, which is critical for grounding team decisions. For example, all-hands team discussions helped guide: The configuration of in situ edge gateway hardware, software, and networking; use of different cloud storage modalities (Git for reliably transferring time-series sensor data "deltas" from the edge; S3 for FaaS data); and the choice of GitHub Actions and Docker containers as a primary target for serverless deployment. All decisions were made to respect the needs of the field crew maintaining equipment at a lake, the ecologists who use data and develop forecasts, and the CI research and development team. 3) A focus on usability in which technical decisions were weighed not only with respect to performance, but also the effort needed to develop, deploy, and manage systems, which has significant implications for whether a given approach can be adopted. As a first step, the focus on usability has helped ensure that designs are effectively usable within our own team, while keeping in mind the longer term goal of broad applicability and usability.

Throughout iterative codesign, a recurring goal was to accomplish tangible milestones and deliver usable prototypes—aiming for simplicity in the beginning and incremental progress in each iteration. This ethos has allowed the team to achieve concrete research goals early and then build on experiences and lessons learned to improve and generalize from prototypes to scalable systems that can be widely adopted by the community. One concrete example has been the development of the core forecasting open source software module of our project, FLARE.[14] The first iteration of FLARE benefited from, and brought the team together, around the foundation of using containers for reproducible deployment. A second iteration added support for natively accessing S3 cloud storage from the containers. The next iteration introduced the use of an open source platform (Apache OpenWhisk) for FaaS deployment. A subsequent iteration integrated GitHub Actions, as well as the use of Apache Arrow for efficient data access. The next iteration led to the design of FaaSr, an open source software package that generalizes the FLARE approach to FaaS-based execution of workflows with functions written in R and data stored in S3/Arrow across multiple platforms.[12]

## Engaging Community

While joint participation early on in interdisciplinary research meetings helped establish a shared team-science foundation, once the translational process began to move from laboratory to locale, engagement from different communities enabled interactions that were invaluable in helping shape our activities, as follows:

- *Managers*: Our team has met regularly with the managers that oversee our focal reservoirs to receive iterative feedback on water forecast design and delivery. For example, following manager request, daily reservoir forecasts were emailed to the managers at a set time in the morning to coincide with a staff shift change at the water treatment plant.
- *Public*: We hosted workshops in which we shared draft forecast visualizations with watershed community members to understand how they might use forecasts to guide their decision-making about lake use. This feedback has shaped visualization prototypes developed in collaboration with watershed association staff.
- *Ecology and forecasting*: Through collaboration with researchers around the world, our end-to-end forecasting system has now been adapted and deployed at other lake sites outside of our oversight to generate automated forecasts. For example, we have collaborated with National Ecological Observatory Network (NEON) scientists to remotely deploy FLARE without disrupting their existing data streams.[6] These NEON lakes forecasts contributed to a community-developed forecasting competition focused on examining different modeling approaches for forecasting water quality.[17] Similarly, we have trained visiting students to deploy FLARE at lakes in Australia and Ireland.

- *Cyberinfrastructure*: We have hosted demonstrations and collaborations in the computer systems community to share techniques learned from EdgeVPN to other domains, including smart cities and federated machine learning, as well as FaaS-based scientific workflows.

### Feeding Back Into Research

The field deployments and engagement with the broader community have provided context to both build new capacity and open further lines of research inquiry. For example, initial FLARE development focused solely on water temperature in reservoirs, but new deployments required FLARE to support additional water quality variables (e.g., dissolved oxygen[13]). Incorporating additional water quality variables in our end-to-end forecasting system increased our memory requirements beyond available GitHub Action resources. Consequently, we needed the capacity to change computation resources in response to a particular FLARE configuration, motivating the development of FaaSr.

### Supporting the Translational Process

Support for our projects has come primarily from the U.S. National Science Foundation, through programs that foster interdisciplinary collaboration, e.g., Smart and Connected Communities (SCC), Cyberinfrastructure for Capacity in Biological Research (CIBR), Software Infrastructure for Sustained Innovation (SSE/CSSI), and Rules of Life (RoL). This funding has been critical for supporting early career trainees and ensuring continuity over time, and we believe that these types of programs are necessary for successful translational, interdisciplinary research. Funding from competitive grants is not sufficient, however, as there is also a significant time and resource investment to build and maintain shared capacity. Two of the major roadblocks to translational computer science[8] are as follows: 1) fewer funding opportunities and 2) typical publication venues may not value translational outputs in the same way as basic research. This loops back to the beginning of our process: While tangible outcomes of funding and publications are crucial, the intangible team science aspects of trust and shared commitment are key to sustaining the research team over time.

## IMPACT

Impacts include the training of students and early career researchers through our interdisciplinary collaboration, the creation of many valuable environmental datasets and computer systems, and the development of automated daily forecasts for managers and other community members. If a water manager receives a forecast that indicates that water quality impairment is likely to happen in the upcoming week, they can act today to preempt or mitigate the impairment by, e.g., implementing management interventions (e.g., algaecide application), changing water treatment processes, or altering staff schedules. Impact also includes open source software[7,11,12] catalyzed by laboratory research and hardened by practical deployments.

## LESSONS LEARNED AND RECOMMENDATIONS

There are several effective ways of fostering research at the interface of ecology and computer science.[9] A key lesson we learned is that translational, interdisciplinary research builds upon a foundation of shared language and knowledge, mutually developed goals, trust, patience, and open communication. This trust, in turn, requires a long-term time investment, continuous interactions, and representation across different career stages (undergraduate and graduate students, technicians, postdoctoral researchers, and faculty).

## CONCLUSION

While initiating these engagements can seem daunting, there are several tangible actions that computer science researchers interested in translational research can take to start. First, step out of disciplinary comfort zones and proactively attend conferences and research coordination network meetings where there can be opportunities to meet, socialize, and begin to appreciate the CI and systems challenges faced by scientists in different domains in which you are interested. This awareness helps both to build rapport and trust and identify collaborative opportunities. Second, conceptualize and codesign what could become the first concrete implementation of a system that applies an interesting CI approach to solve a well-defined application problem. This codevelopment helps identify common fertile ground to build shared knowledge, and further solidifies trust. Third, seek funding opportunities that support taking the kernel of an idea and a prototype to solve a larger, more general problem. Funding, even at small scale, is key to providing resources needed for nurturing a team longer-term. Fourth, consider ways in which addressing specific challenges of a target user community can feed back into new research questions in computer science. This approach can help sustain the collaboration by creating opportunities for new projects and funding across disciplines. Finally, an interdisciplinary culture of training is needed—it takes time, but the goal is to develop the next generation of leaders at the intersection of computer science and application domains that will further enhance and develop translational computer science research.

## ACKNOWLEDGMENTS

This work was supported by the U.S. National Science Foundation (NSF) under Grant EF-2318861, Grant EF-2318862, Grant OAC-2311123, Grant OAC-2311124, Grant OAC-2004441, Grant OAC-2004323, Grant DBI-1933016, Grant DBI-1933102, and Grant CNS-1737424. We thank the many students, postdocs, technicians, and collaborators who have enabled and supported the research described in this article.

## REFERENCES

1. S. S. Farley, A. Dawson, S. J. Goring, and J. W. Williams, "Situating ecology as a big-data science: Current advances, challenges, and solutions," *BioScience*, vol. 68, no. 8, pp. 563–576, 2018, doi: 10.1093/biosci/biy068.
2. M. E. Lofton, D. W. Howard, R. Q. Thomas, and C. C. Carey, "Progress and opportunities in advancing near-term forecasting of freshwater quality," *Global Change Biol.*, vol. 29, no. 7, pp. 1691–1714, 2023, doi: 10.1111/gcb.16590.
3. A. S. L. Lewis et al., "Increased adoption of best practices in ecological forecasting enables comparisons of forecastability," *Ecol. Appl.*, vol. 32, no. 2, 2022, Art. no. e2500, doi: 10.1002/eap.2500.
4. V. Daneshmand et al., "Edge-to-cloud virtualized cyberinfrastructure for near real-time water quality forecasting in lakes and reservoirs," in *Proc. 17th IEEE Int. Conf. eSci. (eScience)*, Innsbruck, Austria, 2021, pp. 138–148, doi: 10.1109/eScience51609.2021.00024.
5. M. Malawski and B. Balis, "Serverless computing for scientific applications," *IEEE Internet Comput.*, vol. 26, no. 4, pp. 53–58, Jul./Aug. 2022, doi: 10.1109/MIC.2022.3168810.
6. R. Q. Thomas et al., "Near-term forecasts of NEON lakes reveal gradients of environmental predictability across the US," *Front. Ecol. Environ.*, vol. 21, no. 5, pp. 220–226, Jun. 2023, doi: 10.1002/fee.2623.
7. "Forecasting lake and reservoir ecosystems." FLARE. Accessed: Jul. 2024. [Online]. Available: https://flare-forecast.org
8. D. Abramson and M. Parashar, "Translational research in computer science," *Computer*, vol. 52, no. 9, pp. 16–23, 2019, doi: 10.1109/MC.2019.2925650.
9. C. C. Carey et al., "Enhancing collaboration between ecologists and computer scientists: Lessons learned and recommendations forward," *Ecosphere*, vol. 10, no. 5, 2019, Art. no. e02753, doi: 10.1002/ecs2.2753.
10. K. Subratie, S. Aditya, and R. J. Figueiredo, "EdgeVPN: Self-organizing layer-2 virtual edge networks," *Future Gener. Comput. Syst.*, vol. 140, pp. 104–116, Mar. 2023, doi: 10.1016/j.future.2022.10.007.
11. "Open-source VPN for edge computing." EdgeVPN. Accessed: Jul. 2024. [Online]. Available: https://edgevpn.io
12. "Function-as-a-service package for R." FaaSr. Accessed: Jul. 2024. [Online]. Available: https://faasr.io
13. W. M. Woelmer, R. Q. Thomas, F. Olsson, B. G. Steele, K. C. Weathers, and C. C. Carey. "Process-based forecasts of lake water temperature and dissolved oxygen outperform null models, with variability over time and depth." SSRN. Accessed: Jul. 2024. [Online]. Available: http://dx.doi.org/10.2139/ssrn.4699835
14. R. Q. Thomas, R. J. Figueiredo, V. Daneshmand, B. J. Bookout, L. K. Puckett, and C. C. Carey, "A near-term iterative forecasting system successfully predicts reservoir hydrodynamics and partitions uncertainty in real time," *Water Resour. Res.*, vol. 56, no. 11, 2020, Art. no. e2019WR026138, doi: 10.1029/2019WR026138.
15. H. L. Wander, R. Q. Thomas, T. N. Moore, M. E. Lofton, A. Breef-Pilz, and C. C. Carey, "Data assimilation experiments inform monitoring needs for near-term ecological forecasts in a eutrophic reservoir," *Ecosphere*, vol. 15, no. 2, Feb. 2024, Art. no. e4752, doi: 10.1002/ecs2.4752.
16. F. Olsson, T. N. Moore, C. C. Carey, A. Breef-Pilz, and R. Q. Thomas, "A multi-model ensemble of baseline and process-based models improves the predictive skill of near-term lake forecasts," *Water Resour. Res.*, vol. 60, no. 3, 2024, Art. no. e2023WR035901, doi: 10.1029/2023WR035901.
17. R. Q. Thomas et al., "The NEON ecological forecasting challenge," *Front. Ecol. Environ.*, vol. 21, no. 3, pp. 112–113, 2023, doi: 10.1002/fee.2616.

**RENATO J. FIGUEIREDO** is a professor in the Electrical and Computer Engineering Department, University of Florida, Gainesville, FL, 32611, USA. Contact him at renatof@ufl.edu.

**CAYELAN C. CAREY** is a professor in the Department of Biological Sciences and codirector of the Virginia Tech Center for Ecosystem Forecasting, Virginia Tech, Blacksburg, VA, 24061, USA. Contact her at cayelan@vt.edu.

**R. QUINN THOMAS** is an associate professor in the Department of Forest Resources and Environmental Conservation and codirector of the Virginia Tech Center for Ecosystem Forecasting, Virginia Tech, Blacksburg, VA, 24061, USA. Contact him at rqthomas@vt.edu.