Guest Editors' Introduction

Pervasive Computing at the Edge

Paramvir (Victor) Bahl
Microsoft Research

Ramón Cáceres
Google

Nigel Davies
Lancaster University

Roy Want
Google

Today, the infrastructure needed to support pervasive computing and the Internet of Things (IoT) is unparalleled as entirely new classes of applications and systems emerge. For example, pervasive systems designed to augment human cognition with tasks such as face recognition must operate at "superhuman speeds," delivering insights to help with human decision-making within very strict and narrow time limits. Similarly, the emergence of pervasive video analytics demands processing of very large volumes of video data in near-real-time. In general, the field of pervasive computing is rapidly changing in the face of major advances in sensing, data processing techniques, and wearable computing. The ever increasing data rates of high-speed networks also factor into the design tradeoff to decide if computing should be local, or remote. The design landscape is multifaceted and complex.

Edge computing (for example, cloudlets and fog computing) offers the promise of supporting the demanding requirements of the next generation of pervasive computing and IoT systems—providing devices with access to nearby processing and storage while simultaneously offering new solutions to issues such as user privacy and resilience. In this new world, research challenges abound: How should these systems be architected? What new applications are enabled? What are the challenges in deploying edge technology? And, how can they be secured against cyberattacks?

In 2009, *IEEE Pervasive Computing* published one of the first papers on Edge Computing ("The Case for VM-Based Cloudlets in Mobile Computing") and in this special issue, we return to the topic to explore recent developments in the field.

Our first article "OpenRTiST: End-to-End Benchmarking for Edge Computing" considers the challenge of comparing the performance of different edge-computing solutions. This is crucial to the future adoption of edge technologies as the approach requires deploying significant amounts of computing infrastructure—and the costs and benefits of these needs to be clearly articulated. The authors introduce OpenRTiST, an open-source application that is highly demanding in terms of computation, bandwidth, and latency requirements and thus is able to provide an ideal test for edge computing solutions.

Our second article "Edge Computing for Legacy Applications" introduces a new use case for edge computing, namely supporting legacy

Digital Object Identifier 10.1109/MPRV.2020.3032205
Date of current version 9 November 2020.

Published by the IEEE Computer Society

applications. The focus is on unmodified legacy content-authoring applications (e.g., Microsoft Office) originally written for personal computers. Drawing on a rich history of supporting VM mobility in mobile and cloud systems, the article describes a new capability termed edge-based virtual desktop infrastructure (EdgeVDI) that allows mobile users to instantiate virtual desktop interfaces on nearby machines—enabling access to legacy applications while benefiting from the performance benefits of a cloudlet infrastructure.

These two articles highlight the significant progress that has been made in the field of edge computing since its inception. The discussion is no longer about the validity of the edge-computing paradigm and the search for a killer application—rather researchers are now focused on how to engineer and evaluate such systems and their potential to support new use cases that further emphasize the benefits of pervasive computing at the edge.

Paramvir (Victor) Bahl is currently a Technical Fellow and Chief Technology Officer of Azure for Operators at Microsoft, Redmond, WA, USA. He received the Ph.D. degree from the University of Massachusetts, Amherst, MA, USA. He is a IEEE and ACM Fellow. Contact him at bahl@microsoft.com.

Ramón Cáceres is currently a Software Engineer with Google, USA. He received the Ph.D. degree from the University of California, Berkeley, CA, USA. Contact him at caceres@google.com.

Nigel Davies is currently a Professor with Lancaster University, Lancaster, U.K., where he also received the Ph.D. degree. Contact him at n.a.davies@lancaster.ac.uk.

Roy Want is currently a Research Scientist with Google, USA. He received the Ph.D. degree from Cambridge University, Cambridge, U.K. Contact him at roywant@google.com.