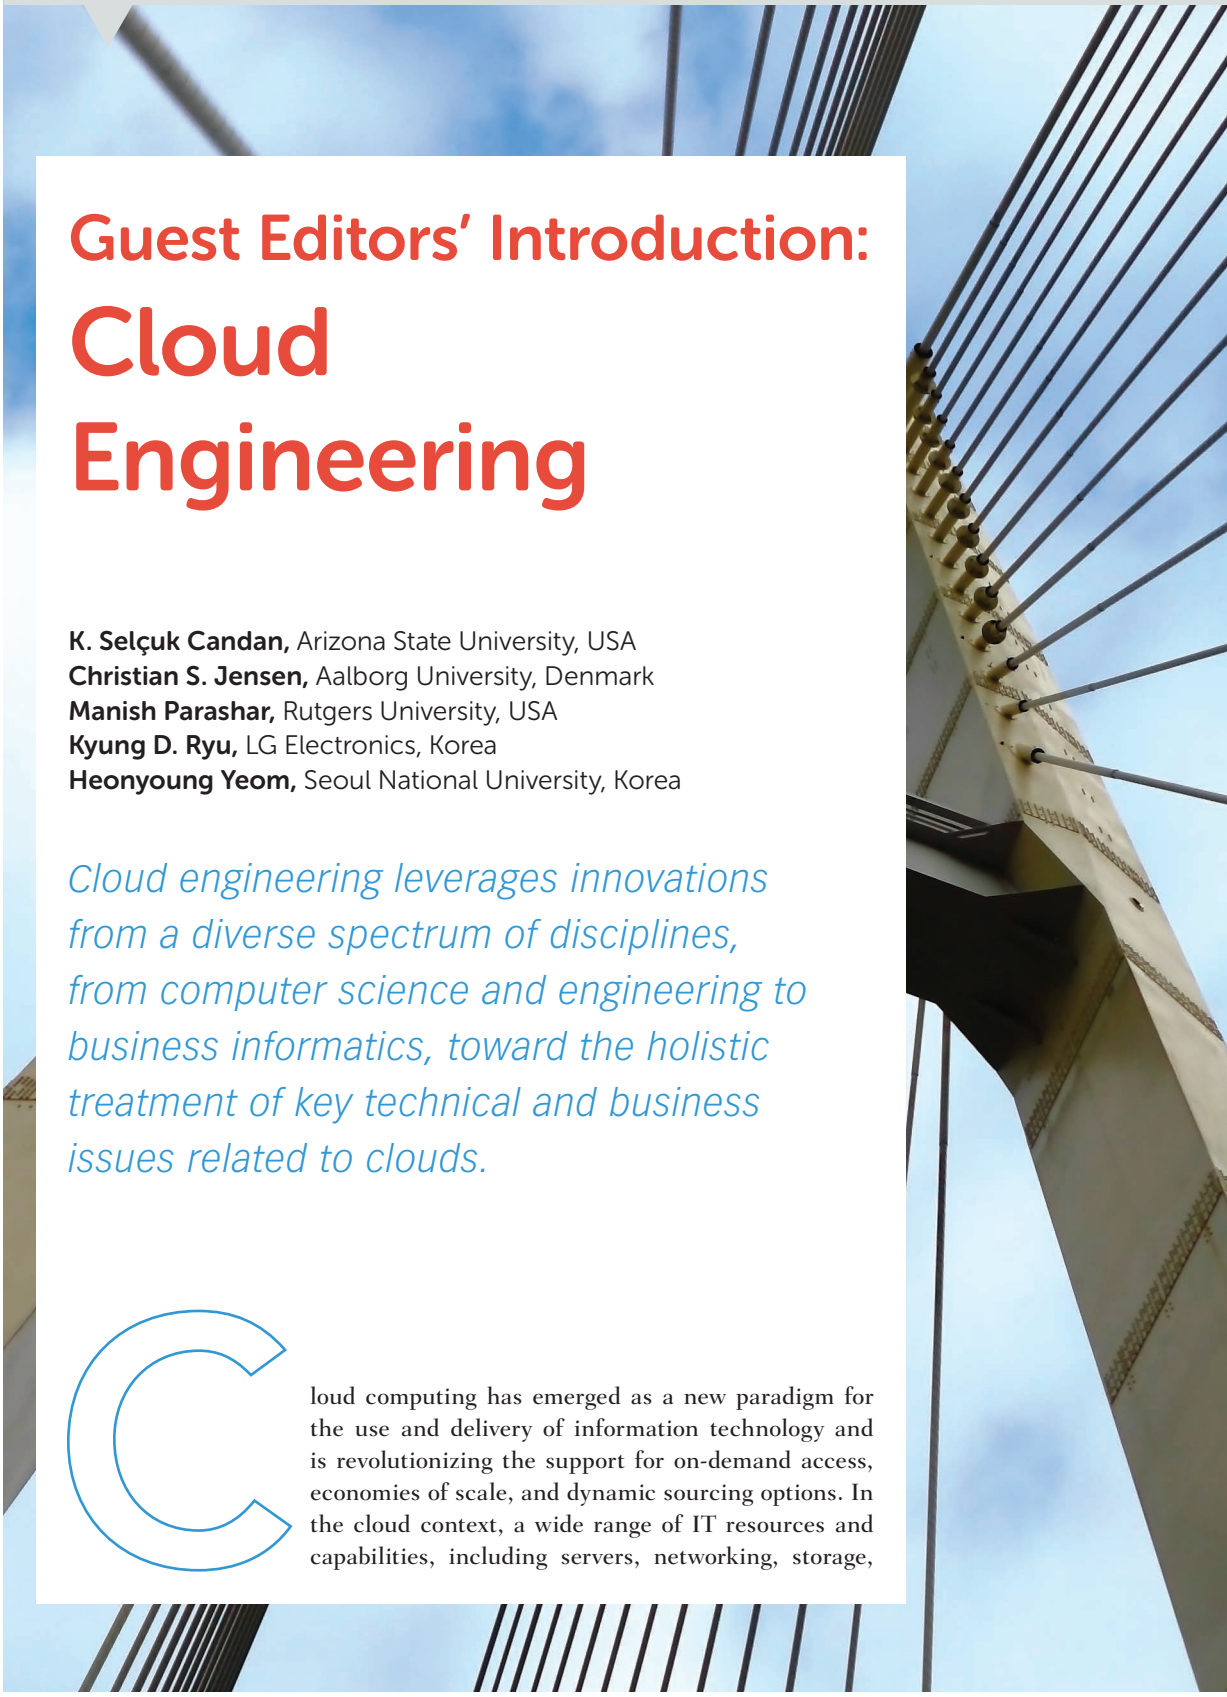

Guest Editors' Introduction: Cloud Engineering

K. Selçuk Candan, Arizona State University, USA
Christian S. Jensen, Aalborg University, Denmark
Manish Parashar, Rutgers University, USA
Kyung D. Ryu, LG Electronics, Korea
Heonyoung Yeom, Seoul National University, Korea

Cloud engineering leverages innovations from a diverse spectrum of disciplines, from computer science and engineering to business informatics, toward the holistic treatment of key technical and business issues related to clouds.

Cloud computing has emerged as a new paradigm for the use and delivery of information technology and is revolutionizing the support for on-demand access, economies of scale, and dynamic sourcing options. In the cloud context, a wide range of IT resources and capabilities, including servers, networking, storage,

 2325-6095/15/$31.00 © 2015 IEEE

RESEARCH DIRECTIONS

In addition to the articles included in this special issue, we highlight some of the other works considered for inclusion in the special issue that show interesting directions for cloud engineering. Two sets of authors to briefly summarize their work.

Security Assurance

"Next-Generation Cloud Security Assurance," by Marco Anisetti, Claudio A. Ardagna, and Ernesto Damiani, Universita degli Studi di Milano, Italy; Antonio Maña, University of Malaga; and George Spanoudakis, City University of London

Cloud users and providers have little or no evidence of the behavior of the cloud services and processes they rely on, or of how their data and applications are secured once deployed in the cloud. Security-assurance solutions based on audit, certification, and compliance have been proposed to evaluate and communicate the security of cloud-based systems. Unfortunately, most existing techniques rely on complex manual approaches and have limited applicability at large scale. Next-generation assurance techniques should support trusted collection and communication of verifiable evidence on cloud service/process behavior based on

Cont. next page

middleware, data, security, applications, and business processes, are being made available as services enabled for rapid provisioning, flexible pricing, elastic scaling, and resilience.

Developing, operating, and maintaining cloud computing systems and cloud-based applications and services are challenging conventional technical wisdom and business practices. Consequently, fully reaping the benefits of cloud computing calls for holistic treatment of key technical and business issues (including commercialization, standardization, and governance) as well as for "cloud engineering" methodologies that draw upon innovations from a diverse spectrum, from computer science and engineering to business informatics.

This special issue focused on cloud engineering seeks to provide a compilation of high-quality papers by researchers and practitioners involved in the development of cloud infrastructure and applications. We invited original, high-quality contributions, foundational as well as applied, describing fully developed or ongoing work relating to all aspects of cloud engineering. We especially targeted the best papers from the 2015 IEEE International Conference on Cloud Engineering (IC2E 2015), which was held in Tempe, Arizona on 9–12 March 2015.

The Articles

The special issue features three articles that highlight the challenges and state of the art in cloud engineering (see the sidebar for a description of other work considered for the special issue).

In "Scalable Attestation: A Step toward Secure and Trusted Clouds," Stefan Berger, Kenneth Goldman, Dimitrios Pendarakis, David Safford, Enriquillo Valdez, and Mimi Zohar address security challenges in cloud

Cont. from previous page
hybrid assessments, using static and dynamic proofs and monitoring. In addition, they shouldn't assume that a trusted party (such as a certification authority) will be available throughout the evaluation process, as this can't be guaranteed in the cloud. Nor should they assume that cloud providers are trusted entities. Rather, they should include attestation protocols to ensure the integrity of critical code for assurance evaluation and of the collected data, using hardware techniques such as trusted platform modules. Finally, the development of applications and services in this new paradigm should be supported by proper engineering processes and tools.

Big Data Processing

"Distributed SSD Caching for Big-Data Systems," by Ming Zhao, Arizona State University; Michel Angelo Roger, Florida Power and Light; and Yiqi Xu, Florida International University

Existing big data systems rely on traditional hard-disk drive (HDD) based storage to provide the volume required by big data applications. But these applications' velocity—the speed of storing and processing data—and variety—the types of data and their processing methods—are also growing rapidly. HDDs alone can't satisfy the increasingly challenging and diverse I/O demands. Emerging solid-state drive (SSD) based storage offers excellent I/O performance in terms of both throughput and latency. However, SSDs have much smaller capacity than HDDs and are much more expensive per unit size, making it difficult to cost effectively provision enough volumes for big data storage. SSDs also wear out due to writes, leading to data durability issues if they're used for permanent storage. Therefore, instead of replacing HDDs, SSD storage needs to be incorporated strategically into existing big data systems to satisfy volume, velocity, and variety requirements. To achieve this objective, the researchers have developed BigCache, an SSD-based distributed caching layer that allows seamless integration of SSDs with existing HDD-based big data systems and enables transparent acceleration of different types of data accesses, thereby exploiting the performance of SSDs and the capacity of HDDs.

computing architectures. Specifically, they look at the ability to monitor and verify the integrity of sensitive resources, such as physical server firmware, hypervisors, guest operating systems, and key applications. They present a scalable attestation method that combines secure boot and trusted boot technologies and extends them into the host and its programs, and into the guest's operating system and workloads, to both detect and prevent integrity attacks.

In "RADical Strategies for Engineering Web-Scale Cloud Solutions," Rohit Ranchal, Ajay Mohindra, Justin G. Manweiler, and Bharat Bhargava address the problem of the lack of clear guidelines for designing and deploying cloud solutions that can seamlessly operate and handle Web-scale workloads. They review industry best practices and identify principles for operating Web-scale cloud solutions by deriving design patterns that enable each principle in cloud solutions.

In "Introspecting for RSA Key Material to Assist Intrusion Detection," John Saxon, Behzad Bordbar, and Keith Harrison note that the increasing use of Transport Layer Security (TLS), a protocol used to encrypt data between two communicating parties, has the potential to complicate the work of intrusion detection systems that try to protect cloud architectures and services. To tackle this challenge, the authors present an efficient virtual machine introspection method based on a structured walk through a virtual machine's memory to find and acquire RSA keys. Their method has the potential to help cloud providers decrypt malicious traffic for analysis to support intrusion detection and protect both themselves and their users.

We hope you'll find this special issue interesting.