Software-Defined Edge Cloud for Smart City

Kaikai Liu, Jerry Gao *, Ke Xu†
*Computer Engineering, San Jose State University, San Jose, CA 95192
†Dell Research, Santa Clara, CA 95054

Introduction. As the capital of the Silicon Valley with over a million residents, San Jose faces multiple challenges in engaging the communities and enabling multiple smart applications, for example, city life, energy and environment, traffic flows, safety and emergency response. How to enable all these smart applications in an efficient manner has become a key issue not only for San Jose but for many city authorities.

Different applications proposed by companies and domain specialists have different data and facility needs, and some of these requirements are critical in terms of residents' privacy and security. For example, multiple companies are applying for installing smart sensors in the street light. City authorities need to enforce policy, and evaluate all the applications. It can be rather slow and labor-intensive for the administrative processing. Another option for city planners is via the build-it-yourself route by selecting state-of-the-art solutions from the portfolio of various wireless standards, hardware and software vendors. This active approach helps the city to gain the agility, and probably enjoy lower Capex along the way, but at the expense of more internal technical resources.

To become effective in this new smart cities initiative, city planners must move to the next generation of automation that understands the requirements of different applications and responds to those requirements according to defined policies, i.e., launch groups of isolated multi-tenant applications on demand. This setting minimizes the deployment cost, and enables computing and smart applications at the wireless network edge. The research challenges are tri-fold:

(1) Software-defined Cloud Orchestration and Policy-based Automation. Different from existing cloud infrastructure in data centers, citywide infrastructure contains thousands of Geo-distributed wireless micro-basestations deployed in residential areas, streets, or parks. Enabling multi-tenant applications in each micro-basestation with high reliability is unachievable by humans or existing solutions. To provide a robust environment where developers can remotely deploy and debug their applications in the cloud of micro-basestations, the orchestration layer with the remote management dashboard is a must-have to enable policy-based automation. The requirements described by business users will be systematically validated through defined policies in the dashboard. The infrastructure will not only automatically control the configuration of the computing, storage, and networking, but also virtualize heterogeneous sensing modules. This software-defined cloud environment will simplify the smart city infrastructure through open standards and resource sharing, as well as responsive to shifting requirements adaptively.

(2) Edge Computing and Virtualization. Low latency analytics and real-time response are especially important for smart city applications, for example, security enforcement, video surveillance, traffic monitoring. The dominant approach of aggregating all the data to the datacenter stresses communication links, and inflates the timeliness of analytics. Moving much of the processing to the locations where the event is happening facilitates real time response and low communication overhead. At the same time, the computing and sensing modules available in a micro-basestation will be virtualized as a resource pool that integrates a cluster of application containers to provide agility, responsiveness, and less overhead than traditional hypervisor-based virtualization. To operate, maintain, and secure this edge cloud network, researchers must grapple with multiple vendor-specific computing and sensing modules to implement complex high-level container management policies. Despite many previous proposals to make the VMs in datacenter easier to manage, many approaches are not fit for this edge cloud scenario or only amount to stop-gap solutions because of the underlying highly distributed and low-complex infrastructure.

(3) Resilient Network. Different from datacenter networks where the cable/fiber connection is more reliable, edge cloud infrastructure involves distributed heterogeneous gateways and unreliable wireless links. As unpredictable disasters and attacks increase, we need a resilient network design for the edge cloud infrastructure that avoids any single point of failure and keeps residents connected to vital city services. The deployed smart application should easily re-route via the software defined infrastructure, and collaborate with nearby peers with improved disaster preparedness and response. The adaptive monitoring software will determine when hardware is likely to fail, when resources will exceed capacity, or where the attack is happening, and finally deliver the agility and flexibility needed to support smart applications that enhance the livability, workability and sustainability of cities.

The Team. Kaikai Liu, Jerry Gao, Robert Morelos-Zaragoza, and Harry Li from San Jose State University (SJSU), work with Khanh Duy Russo (Director of Strategic Partnerships, Office of Mayor Sam Liccardo) from the city of San Jose, for smart cities initiatives. Jerry Gao is the director of the Silicon Valley Excellence Research Center for smart city in the MetroLab network. Ke Xu (Research Staff Member), Rajesh Narayanan (Distinguished Engineer), and David Shen (Business Strategy Director), from Dell Research have a strong expertise in datacenter networking, cloud architecture, services and management.