

Networks for Future Services in a Smart City: Lessons Learned from the Connected OFCity Challenge 2017

Domaniç Lavery, Marco Ruffini, Luca Valcarengi, Naoto Yoshimoto, Thomas Pfeiffer, Dave Hood, Junwen Zhang, Daniel King, Hal Roberts, Rajesh Yadav, Nicola Sambo, Marco Tacca, Silvia Fichera, Franco Tecchia, Marcello Carrozzino, Elaine Wong, Ning Cheng, Yuki Yoshida, Denis Khotimsky, and Jun Shan Wey

ABSTRACT

The drive towards ubiquitous communications has long been encompassed by the concept of a *connected* or *smart* city. The idea that data transfer and real-time data analysis can enhance the quality of life for urban inhabitants is compelling, and one can easily envision the provision of exciting new services and applications that such an information-driven city could provide. The challenge in achieving a truly smart city stems largely from communications technologies – fixed line, wireless, backhaul, and fronthaul – and how these are combined to provide fast, reliable, and secure communications coverage.

Here, we report on the key observations from the *Connected OFCity Challenge* competition, held at the Optical Fiber Communications Conference (OFC) 2017, which addressed the fixed and wireless access network requirements for smart cities. It is shown that, from a technological perspective, future optical networks will be capable of securely supporting extremely low-latency and high bandwidth applications. However, as shown by using a Networked Music Performance as a particularly challenging example application, how readily this is achieved will depend on the interplay between wired and wireless access services.

Introduction

In March of 2016, the Optical Fiber Communications Conference (OFC) hosted the inaugural *Connected OFCity Challenge*. Teams of industry experts and academics shared ideas to enhance the communications capability of a geographically and demographically representative municipal area (the *OCity*) and, in doing so, debated and proposed solutions for some of today's most pressing issues in optical access networks [1].

The second *Connected OFCity Challenge*, held at OFC 2017, asked deeper technical questions about future services in a smart city. This was set against the backdrop of the *OCity* hosting its Septicentennial celebrations, which comprised, inter alia, a *virtual reality* (VR) concert and major sporting events. Contestants were asked which innovations in access networks would be required to support advanced applications (such as remote, interactive audience participation), which are *latency-sensitive*, *bandwidth-intensive*, and/or require high *network reliability and security*.

Each of the four competing multidisciplinary teams (Alive, Fibrus, Vutopia and Terapolis), proposed innovative solutions for meeting these three, in part contradictory, requirements. The proposed solutions are highly relevant to design choices made today by city planners, application designers, and network engineers. In this article, we take the opportunity to share the key findings of the 2017 competition.

When we think about Internet access, we tend to think in terms of the peak number of bits (or bytes) transferred per second. This is how network connections are sold to consumers, and how we quantify the quality of our service. However, much of the time, the true Quality of Service (QoS) has less to do with the volume of data transferred and

more to do with the consistency and latency of that data transfer. Indeed, the QoS required is application specific, and so we were moved to ask: which applications would go beyond the capabilities of today's network, and what changes can we make to the network to enable them?

The goal of the Connected OFCity competition was, therefore, to stimulate the generation and exchange of innovative ideas to transform access networks for future applications.

The most challenging target application presented to the teams in this competition was to provide suitable connectivity using a virtual/augmented reality (VR/AR) application for an orchestra concert rehearsal and live performance, where the musicians were in different geographical locations. How far apart could the musicians be, and how would the network cope with such a requirement? We envisaged for this challenge the simultaneous transmission of remote audio and video to enable a remote presence for musicians, as well as the concept of a virtual conductor to aid synchronization between all connected musicians.

The second, more conventional, target application challenged the teams to transmit extremely high-quality video within a metropolitan area network. This task was concerned with the sporting events and addressed both remote support for training athletes and broadcast VR for audience viewing. Latency, again, plays a role here, but bandwidth is the driving issue.

Finally, the teams were asked to consider a traffic management scenario, where the passage of autonomous vehicles could be managed via remote data centers. This challenge was a test of a simultaneous demonstration of high security/reliability and low latency network architectures – a requirement common to many other smart city applications.

The following sections discuss challenges and solutions for fixed-mobile/cloud access convergence and intelligent optical networks, when considering these three very demanding scenarios.

A Networked Music Performance

Latency considerations

Perhaps surprisingly, music represents a particular challenge for both the optical and wireless communication infrastructures, due to the acute ability of humans to perceive asynchrony in ostensibly simultaneous audio. In contrast to, say, a Voice Over IP (VOIP) conversation, where participants would rarely require concurrent speech, musicians greatly rely on timing cues from other musicians to make decisions about when, and what, to play. Different types of musical interactions are possible (Fig. 1), as detailed in [2]. Among these, the most challenging is the remote, or networked, synchronous music performance (also known as Remote Jamming), where low latency auditory feedback between musicians is essential.




Fig. 1 Types of musical interaction (presented by Team FIBRUS) from Á. Barbosa [2]

In this scenario, musicians will be capable of rehearsing remotely only if located in places no farther than approximately 20-50 ms in latency [3]. In the OFCity Challenge, Team FIBRUS tested this scenario in a transatlantic Remote Jamming session between Pisa Italy and Dallas Texas US, as depicted in Fig. 2. They showed that this time cannot be estimated simply by distance (Pisa-Dallas, about 9,000km – 45 ms), but should, instead, be based on the network, OSI layer three, latency (e.g., Pisa-Dallas 126 ms round trip time) as reported in Fig. 2 (inset).




Fig. 2 Remote Jamming test scenario (presented by Team FIBRUS). Although there is a 45 ms physical latency between Scuola Superiore Sant’Anna, Pisa (Italy) and Dallas, Texas (USA), the *traceroute* (inset) shows a 40% latency overhead.

Based on this analysis, any network solution would require a maximum differential distance between performers of <4,000 km, that the number of traversed nodes is minimized (i.e., reducing the store and forward delay) and that the access network latency

is generally kept low (i.e., for the optical access network, using XGS-PON or NG-PON2 with short DBA time – 1-2 ms – and few multiple access wireless technologies).

To further understand how network latency could affect an orchestra music performance, Team Alive studied the dynamics of sound propagation delays across a concert stage, which can span over 20 m in length. Considering a speed of sound of 340 m/s, the propagation time across the stage can easily exceed 60 ms. Orchestras are typically divided into instrument sections (e.g., strings, brass, woodwind, percussion) each spanning only a few meters, suggesting that the propagation delay within a section should remain below 15 ms for a 5m distance, while across sections up to 60 ms might be tolerated for a distance of 20 m.

An interesting observation occurs if we extrapolate such values to network propagation latency; as signals travel in optical fiber at speed of approximately 200,000 kilometers per second, the concert stage distances above would dilate respectively to 3,000 and 12,000 km. However, in addition to propagation, we need to consider the network latency due to packet processing (e.g., header processing, switching and protocol delays), packet serialization, packet queuing at switches/routers, and, more importantly, the delays introduced by the application (e.g., coding/decoding and compression for both video and audio signals). The application delays, indeed, dominate when queuing delays are negligible, with the typical lowest values for audio/video jitter buffering of 5 ms, audio codec latency of 10 ms, and (hardware) video codec latency of 20 ms. Considering that audio and video coding could be executed in parallel, and allowing 5 ms for packetization, protocol overheads and hardware switching delays, the minimum overall delay will be around 30ms. Since this does not satisfy the latency constraints for intra-sectional players, the players will be affected by noticeable transmission delays. If such effects are tolerable, then this could work provided the intra-sectional distances are kept to a minimum (say within 100 km, adding just another half a millisecond). Another option is instead to use uncompressed video, thus eliminating all coding delays. This would potentially satisfy the latency requirements even over longer distances, but would require ultra-high capacity solutions, particularly in the optical fiber backhaul (e.g., considering high definition streams, a sustained capacity above $N \times 1$ Gb/s for each player, when N players are rehearsing together).

Regarding the inter-sectional latency, assuming a software defined network (SDN) control plane can reserve capacity end-to-end across the network with priority queuing QoS, the inter-sectional distance could allow up to 4,000 km where we have allowed an additional 10 ms of networking delays across several potential routing hops. If the end-to-end path can be achieved through transparent optical connections, without intermediate packet processing, such distances can increase further due to reduced packet processing delays and jitter buffering at the application.

Networking solutions

Although in the competition itself, each team presented a unique architectural solution, here we focus on the two most radically different proposals.

The first proposal, presented by Team Alive, suggests the use of multi-wavelength PONs as a solution to provide ubiquitous multi-tenant connectivity to any type of service. These services scale from residential access, up to fiber-enabled dense mobile cells (i.e., mobile back/fronthaul). The overall architecture is shown in Fig. 3.

Deep infrastructure virtualization, both in the central office (CO) [4] and in the PON [5] allows virtual network operators (VNOs) to deliver assured QoS even for time-critical applications. This allows cost-effective, high-capacity and low latency connectivity giving

the ability to adopt cloud-RAN to maximize resource sharing in high-density cell deployment.

In this architecture, it is envisaged that 5G services will add complexity to COs, for example to support data storage and processing for smart city applications, while meeting strict latency constraints. The proposal is to use fronthaul-based mobile edge computing, which adds edge processing to a locally centralized node covering a number of Radio Access Technologies (RATs) [6] (detailed in relation to autonomous vehicles in Fig. 6). The Radio Units (RU) connect to the edge node through relatively short fronthaul links, while the link from the converged node to the CO operates at lower rates; dynamically switching between MAC (media access control) and PHY (physical layer) splits depending on functional requirements and resource availability. Besides reducing the complexity of each RRU, this solution provides minimal communications latency, facilitating coordination of cells with enhanced aggregated processing.




Fig. 3: Overall architecture of the Team Alive solution with the typical city (OFCity) shown, center. Legend: CO: central office; DC: data center; RN: remote node; SDN: software defined network; QoE: Quality of Experience.

The second solution, created by Team VUTOPIA, focused on provision of low latency communications by placing at the center of their network architecture the stringent low latency physical layer requirements of 5G wireless (specified as 100 μ s one-way latency in the fronthaul link in eCPRI [7]). They proposed a feasible concept of physical network architecture, termed *all-spectrum access*, in which resources spanning an extremely wide frequency range, from optical to millimeter and terahertz waves, can be used dynamically and flexibly. This network architecture consists of not only optical fiber access, but also all-spectrum fully-meshed wireless access using millimeter and terahertz wave as shown in Fig. 4. The feature of this architecture is to reduce latency and respond in real-time.

Firstly, the optical metro-area network configuration is a widely flat-ring based on time-sensitive Ethernet switches connected to the optical line terminal (OLT) and radio base band unit (BBU). Secondly, OLTs, multi-RAT BBUs, and edge servers are virtually co-

located to efficiently transfer data with low latency. And thirdly, the optical feeder fibers covering the full-mesh wireless fronthaul are transformed from a conventional point-to-point architecture to small ring architecture with high resiliency, to maintain a time-sensitive response. To further develop this delay-sensitive network, SDN-orchestration and Artificial Intelligence (AI) networking technologies can be used to build a development environment for the application platform and all-spectrum resource virtualization function.




Fig. 4: Team Vutoptia's solution based on a physical layer platform including an all-spectrum testbed. Legend: NW: network; MW: micro wave; MMW: millimeter wave.

Implications for Realtime Augmented and Virtual Reality

Today, we already have access to sufficient network capacity to support 4K, or even 8K, video streaming, sometimes in stereoscopic 3D. This is a testament to the remarkable developments in source coding (compression), which has enabled support for high quality video streaming, requiring just a few tens of megabits per second. The implication for the recent developments in VR/AR video is that, so long as a video stream can be buffered to avoid encoding latency in the video stream, the data rate requirements are comparable to conventional video.

As a very simple test case of commercially available VR streaming, during the OFCity challenge event, a cloud platform was used to live stream a 360° video between a laboratory in London to the audience and event participants in Los Angeles. The latency (including the network, transcoding and compression) was around 14 seconds. The videos in the live streaming demo can be viewed in the OFCity YouTube Channel [8].

So, what can be done in latency-sensitive scenarios? At the software level, advanced transcoding and compression would incur a significant delay and, consequently, low latency compression (or raw data transmission) would be used to accommodate the higher future bandwidth required by the VR/AR stream.

In the context of the smart city challenge, Team Alive proposed that providing such advanced digital media support for the sporting events would require a mix of mobile communications technologies, in conjunction with fiber front/backhaul, due to the

competing requirements of different end users. On the one hand, they aimed to provide an enhanced audience experience, using media technologies such as AR with a 360° field of view. On the other hand, they also aimed to provide the training teams with realtime feedback information to improve athletes' performance.

A specific example was given for a diving competition, where a number of accelerometers could be worn on the critical parts of an athlete's body to be able to record their movements so that an exact replica of the diving performance can be recreated through post-processing to provide athletes and training teams with critical feedback.

Three external cameras, and one mounted on the diver, could be used to provide an audience with a 3D subjective diving experience. The compressed data rate of each camera, which provide up to 8K resolution at up to 120 frames per second, is on the order of 50 Mb/s and connectivity to the sport center hubs could not be provided using current WiFi standards, but could be provided through the multi-gigabit/s WiFi technology specified in IEEE 802.11ax, which is due for general release in 2019.

The team also discussed the requirements for a separate sporting event in this competition, the cycling, where AR glasses would be used to display relevant physiological parameters and feedback from the team and doctors. The main challenge here is to provide immersive VR experience to the public, across an area spanning over a hundred kilometers.

The proposed solution can be seen in the bottom left part of Fig. 3. Cameras are placed on all cyclists' bikes and on a number of drones, which carry out multiple tasks. They provide multi-hop connectivity to cyclist cameras, receiving their signal over 802.11ax technology and relaying it to 5G base stations over 60 GHz MIMO channels. In addition, they provide aerial camera views, as they follow the cyclists along the course. Cyclists, drones and team cars are all part of the same ad-hoc mobile infrastructure as a self-organizing network to deliver video streams towards a central datacenter. As the race moves along its track, drones will be delivering their large aggregated data streams to different base stations. The agile optical network can dynamically increase the capacity offered to those base stations, as the race progresses, to support higher wireless bandwidth and higher number of MIMO channels.

Team FIBRUS noted that athletes should be capable of training remotely in an immersive environment where they would experience 360° VR with 4K-equivalent resolution and sensory feedback, e.g., uphill or downhill slope, water/wind resistance. This setup would require bandwidth on the order of 1 Gbps – 5 Gbps per VR stream [9]. 10G PON, multiple 10G PON, and NG-PON2 systems are capable of providing the desired, scalable bandwidth to event locations. Advanced VR solutions could be implemented to reduce the required bandwidth with the potential trade-off of more stringent latency requirements. For example, only the 90° athlete view could be transmitted with the rest of it stored at edge nodes for fast transmission in the case of sudden movements. This solution does increase the need for lower latency. The VR experience for mobile devices would be encoded at lower rates to allow for transmission over 4G/5G wireless connectivity.

Finally, Team VUTOPIA noted that, for a VR audience experience at a large-scale event, such as this sporting competition, there would be a temporary broadband resource requirement at a local location with high population density (i.e., the event attendees). In this case, in addition to the aforementioned solutions, a 10 Gbps MMW/THz link (from the all-spectrum testbed) would be useful because it can provide high bandwidth locally, and can be easily installed and removed.

Implications for Reliable Autonomous Vehicular Transport

The final challenge for the teams in this competition was to probe the response of their network designs when handling latency-critical traffic for autonomous vehicle management in the city.

Team FIBRUS made the point that, for OFCity to use autonomous vehicles [10], they should be truly autonomous, without the need for *any communications network*, much less a high availability, low latency one. This is necessary since any car/car, car/curb or car-datacenter-car communication must be wireless and, as such, network availability cannot be guaranteed. That is not to say autonomous cars cannot *benefit* from such a system.

For example, as depicted in Fig. 5, sensors installed on stoplights can signal to cars, equipped with low latency wireless receivers, potential risks (e.g., a human driven car running a red light). In this case, the connectivity must be wireless and latency may be reduced as the communication is local and does not need to connect with low latency to the cloud. Ultra-reliability is not needed as only intersections with such hazards will need them at first.

In another example, for autonomous cars to take advantage of real time, detailed local maps providing a pervasive infrastructure is needed. Ultimately, cars will become a mobile living room that will need all the communication capabilities of a home, only it must be wireless based. In OFCity, 5G base stations transmitting sensor, map and entertainment data may be installed on stoplights or light poles fed by high capacity terrestrial access networks such as 10G PON to enhance the capabilities of autonomous vehicles. This scenario is not atypical of many smart city applications, where sensor data is aggregated and used in situ (see Fig. 6). The stringent and diverse application requirements must first be met in the wireless domain, but the optical network supports these self-managed regions by providing a high capacity backhaul channel. Interestingly, in this model, the optical fiber communications network, at the physical layer, can simply make use of existing standards.




Fig. 5: Sensor-equipped stoplights at a four-way junction. Sensing that vehicles A and B are both approaching this junction enables the appropriate signal (red/green) to be set to avoid collision, and to manage traffic flow.

So, regardless of the feasibility of autonomous vehicle control in practice, these requirements are a common need in smart cities, which typically exploit some form of supervisory control for city monitoring and management. Team VUTOPIA proposed that, based on their technical innovation of an all-spectrum access, many of these smart functions can be realized at the center of the city to provide high-capacity and latency-sensitive communications. For example, by exploiting aggregated and anonymous geolocation data from mobile phones, connected vehicles and other sensors around the

city, a comprehensive picture of crowd densities, travel patterns, and group profiles of traveling crowds, can be established. This enables real-time monitoring and crowd management, allowing transport operators to deploy public transport resources as required. The data sets can also be exploited in predictive analysis for future events, allowing potential issues to be highlighted before they occur.

However, VUTOPIA's innovation comes into its own in the context of autonomous vehicle management, where Vehicle-to-Vehicle and Vehicle-to-Curb Infrastructure (V2I) communications can carry sensor and control information. This allows a vehicle to dynamically sense of its immediate vicinity, share this information and likewise receive information that allows anticipation of traffic conditions occurring for miles ahead, with obvious benefits for road safety applications. The V2I interface naturally fits with the all-spectrum testbed, which would enable wider area communication of sensor data as and when required. To reduce latency and respond in real-time by pre-empting an incident, cloudlets with control servers would be installed along key routes at the RRH/RRU/Base stations, allowing sensor and control information to be directly routed within the cells, instead of transporting data through the mobile and optical network via the cloud.




Fig. 6: A simple, converged fiber/wireless edge node, which facilitates realtime processing of live traffic data.

SUMMARY

In this second edition of the *Connected OFCity Challenge*, four multidisciplinary teams competed for the most innovative solution for future services in a mid-size city, which called for achieving the objectives of high bandwidth, low latency, high security and reliability.

In the networked music performance scenario, a transatlantic Remote Jamming session demonstrated how the network latency affected the musicians' experience. In an orchestra performance, up to 120 ms might be tolerated for a distance of 20 m between musicians. Applying these findings to a remote music performance and considering network and application delays, the distance between musicians could be up to 4,000 km. Innovations in optical networking technologies play an important role to provide efficient backhauling and support the high bandwidth requirement.

Two different architectures were proposed in this competition. One approach uses multi-wavelength PONs to provide connectivity from residential to back/fronthaul of dense mobile cells. Another approach is the *all-spectrum access*, consisting of optical fiber and fully-meshed wireless access. The teams agreed that infrastructure virtualization and software defined network control plane capable of reserving capacity with priority QoS are instrumental in building a delay-sensitive network.

The teams were asked to include virtual reality/augmented reality video streaming in their solutions. A mix of mobile technologies in conjunction with optical fiber backhaul/fronthaul would be required to capture and deliver their videos. As a proof-of-concept demonstration during the event, a 360° video from London was live-streamed to the audience in Los Angeles. The latency was around 14 seconds.

Finally, the teams validated their network designs by considering latency-critical traffic for autonomous vehicle management. In some cases, autonomous vehicles must be wireless and information elaborated locally (e.g., for accident prevention) without the need to access remote servers. The stringent and diverse application requirements must first be met in the wireless domain, complemented by the optical network to provide a high capacity backhaul channel.

We originally set out in this competition to brainstorm novel ideas for optical access technologies; however, interestingly, although not surprisingly, all teams proposed that realistic solutions must rely on a close cooperation of fixed and wireless networks using optical and radio access technologies. As a takeaway for the optical communications industry, innovations to accelerate the convergence of these access networks will be critical for future smart city services.

Note: The Connected OFCity Challenge will return at OFC 2018 in San Diego, taking on the task of communications infrastructure deployment in the developing world.

ACKNOWLEDGEMENT

The authors wish to extend their thanks to the 2017 Connected OFCity Challenge judging panel, including Dominique Dagenais, Salim Jabr, Julie Kunstler, Inder Monga, Kazuhide Nakajima, Rod Tucker, Junichi Kani, and Peter Vetter, for their invaluable input to the discussion and debate. We also would like to thank Team Terapolis (Dimitra Simeonidou, Sergi Figuerola, Harald Hass, and Stephen Hilton) for the engaging discussion in the competition.

REFERENCES

- [1] Rod Tucker *et al.*, “Connected OFCity: Technology Innovations for a Smart City Project [Invited],” *J. Opt. Commun. Netw.* vol. 9, no. 2, pp. A245-A255, 2017.
- [2] Á. Barbosa, “Displaced soundscapes: A survey of network systems for music and sonic art creation”, *Leonardo Music Journal*, pp. 53-59, 2003
- [3] Z. Kurtisi, X. Gu, and L. Wolf, “Enabling network-centric music performance in wide-area networks,” *Communications of the ACM*, Vol. 49, no. 11, pp. 52-54 Nov. 2006.
- [4] Peterson *et al.*, “Central office re-architected as a data center,” *IEEE Communications Magazine*, vol. 54, no. 10, pp. 96-101, Oct. 2016.
- [5] A. Elrasad *et al.*, “Virtual Dynamic Bandwidth Allocation Enabling True PON Multi-Tenancy,” Paper M3I.3 in proceedings of Optical Fiber Communication Conference, March 2017.
- [6] Fiber-Wireless Integration and Networking (FiWIN), NSF-Center, Georgia Tech, <http://www.fiwin.org/> [Accessed on 11th January 2018]
- [7] eCPRI Transport Network D0.1 (2017-08-30)
- [8] *Connected OFCity Challenge* YouTube channel <http://goo.gl/h3seKq> [Accessed on 11th January 2018]

[9] "Whitepaper on the VR-Oriented Bearer Network Requirement," Huawei Technologies Co., Ltd., <https://goo.gl/2Bkc9d> [Accessed on 11th January 2018]

[10] J. Levinson *et al.*, "Towards fully autonomous driving: Systems and algorithms," in Proceedings of IEEE Intelligent Vehicles Symposium (IV), pp. 163-168, 2011.

Biographies

Domanic Lavery

Domanic Lavery joined the Optical Networks Group at University College London (UCL) in 2009, and currently holds the position of Lecturer and Royal Academy of Engineering Research Fellow. His primary research interest is the use of digital coherent receivers for optical networks at all scales. Between 2014 and 2016 Dr. Lavery was a member of the Optical Access Networks Subcommittee at OFC, and he co-organised the Connected OFCity Challenge workshop in 2016.

Marco Ruffini

Prof. Ruffini is academic and principal investigator of Trinity College Dublin. He is part of the CONNECT telecommunications research centre, where he leads the Optical Network Architectures laboratory. His main research areas are converged access-metro and fixed-mobile network architectures, and includes topics such as access network virtualisation, multi-tenancy and SDN control planes. Prof. Ruffini has to date raised research funding in excess of € 4M, and has authored over 100 international publications, and 10 patents.

Luca Valcarenghi

Luca Valcarenghi has been Associate Professor at the Scuola Superiore Sant'Anna of Pisa, Italy, since 2014. He published more than 100 papers in International Journals and Conference Proceedings and actively participated in the TPC of several IEEE conferences, such as Globecom and ICC. Dr. Valcarenghi received a Fulbright Research Scholar Fellowship in 2009 and a JSPS "Invitation Fellowship Program for Research in Japan (Long Term)" in 2013.

Naoto Yoshimoto

Naoto Yoshimoto joined NTT Laboratories in 1988, and engaged in the research and development of high-speed optical transmission components and broadband optical access network systems. He has also engaged in the planning of next-generation optical-wireless converged access network based on advanced PON technologies. From 2014, he is a professor of Chitose Institute of Science and Technology in Hokkaido, Japan. He has recently been devoting effort to advanced research of virtualized broadband access networks for smart cities.

Thomas Pfeiffer

Thomas Pfeiffer has been leading the optical access research in Nokia Bell Labs for more than a decade. By the end of 2017 he took on a new role as Principal Architect for Optical Access and Fixed Mobile Convergence in Bell Labs. Building on a profound background in ultrafast optical and optoelectronic research he evolved during his industrial career from physical layer and component oriented towards systems and network oriented work.

Dave Hood

Dave Hood chaired the ONF Architecture group during its development of the SDN architecture and SDN-NFV documents. Other contributions related to SDN include aspects of information modeling, network management, and signaling. Earlier, Dave worked in PON access networks, where he was active in standards organizations, and wrote a book on G-PON.

Junwen Zhang

Dr. Junwen Zhang received his Ph.D degree from Fudan University in 2014. From 2012 to 2015, he studied and worked at Georgia Tech. Dr. Zhang joined ZTE(Tx)