Co-Designed SDIs for Computing and Physical Complexes Using Integrated Analytics

Nagi Rao, Oak Ridge National Laboratory
raons@ornl.gov

Expanding Reach of Networks and Virtualization: The capabilities of on-the-fly remote visualization and analysis of complex cyber-physical datasets, and monitoring and steering of supercomputations and large physical instruments, seem increasingly realizable with advances in Software Defined Infrastructures (SDI) and virtualization technologies. The underlying data flows in these cases, for example, involving large computing systems and physical facilities, are expected to generate large volumes of data, which are to be transported over wide-area networks. These flows must be composed from complex and disparate component systems, which require a concerted optimization of the local storage-to-host data and execution paths to match the edge and long-haul network connections. Currently, such data flows are mostly manually composed, configured, and optimized, which require teams with multiple areas of expertise and significant setup times. SDI frameworks combined with virtualization of disparate components offer opportunities to provide these flows *directly and transparently* to users and applications entirely through software, They not only make the process agile and fast but also enable sophisticated dynamic end-to-end performance optimizations. The underlying solutions must be developed and tested to "software-enable" all components, and dynamically solve the complex control and optimization problems.

Co-Designed SDIs: Under the concept of *Co-Design*, the network technologies will be developed in concert with other systems for realizing the data flows across the complexes of computing, storage, and physical devices. This approach is in sharp contrast with conventional approaches that separately design the network infrastructures. We propose that comprehensive software suites be designed, developed and tested to include control plane modules and controllers for site flows, and orchestration systems be devised to realize end-to-end flows using site and multi-domain long-haul software defined exchanges, The orchestration software will be designed to support both on-demand and advanced scheduling of custom, dynamic data flows that help complex workflows involving networked computing and physical facilities. These co-designs will be build upon the emerging Open Networking technologies (e.g., OpenFlow, OpenSM), and also will develop novel control plane extensions to encompass other components, including processing nodes, cross-connects, I/O systems, storage networks, and file systems. These end-to-end network flows will be seamlessly composed and realized entirely in software to achieve unprecedented speeds and performances through dynamic configuration and optimal resource utilization.

Testbed with Integrated Analytics: Development and testing of the required co-design technologies will be supported by powerful testbeds that include diverse components in addition to network devices and connections. Furthermore, we propose that they be augmented with analytics capabilities both in terms of analytical models and analysis tools as an integral part of the infrastructure. In addition to hardware, they include emulators, simulators and analytical modules that implement capabilities for which hardware is yet to be designed and built. Rigorous measurement analysis methods will be devised for performance assessment and optimization through a structured experimentation regimen. The data processing, storage and analysis tools will be available as a part of testbed infrastructure, unlike conventional network testbeds.

Biography: Nagi Rao is a Corporate Fellow at Oak Ridge National Laboratory. He is a Fellow of IEEE, and published more than 350 technical journal and conference papers. He has been involved in the design and implementation of high-performance networking solutions for Department of Energy (DOE) and Department of Defense, over past two decades. He co-designed and implemented the Ultra Science Net (USN), including installing equipment and writing control plane codes; it is an R&D network with a nation-scale footprint, and is dynamically provisioned using Linux scripts, often by (non-network) scientists, since 2004. He also developed analytical models for transport and protocols, and structured network experiments and analysis methods. He is Lead PI of multi-year DOE SDN-SF project that develops SDN solutions for science flows involving supercomputers, science instruments and 100G networks.