Towards Service and Application-Aware Ultrafast Networking

It is expected that the next generation Internet should be able to support demanding new services and applications that require minimal latency and high bandwidth, for example in healthcare, manufacturing and transport, and which cannot be supported by the current infrastructure. Here we present work which is currently taking place at University College London and addresses a number of areas that can collectively contribute towards this direction, and which are the following.

First, on Big Packet Protocol with prioritized payload dropping for congestion avoidance. Second, on the use of machine learning supported networking, which includes neural-assisted routing and zero-queuing with network-centric (as opposed to current application-centric) congestion control. Third, on service-routing with service names encoded in packet headers and very low latency resolution required at the edge without the use of the DNS. Fourth, on a management plane driven by declarative high-level intents which are automatically refined to low-level configuration commands and which is supported by smart, adaptive and scalable line-speed traffic monitoring. And fifth, on high-precision new IP that includes meta-information in headers to support programmability w.r.t. to the way packets are processed. An overview of all these research activities is given, presenting their state and challenges.