PhD research fellowship on Optical Internet Archtectures

Although Internet traffic growth has sensibly decreased over the past few years, traffic forecast show that annual growth of about 40% remains a plausible scenario over the next five years. Especially if we consider that service providers are now deploying FTTx solutions (many of which are FTTH), which have the potential to increase customer peak bandwidth by well over two order of magnitude compared to current offers on copper access, such traffic forecast seem to represent a lower bound.

Traffic growth has put the core network under pressure to provide more and more bandwidth at lower and lower cost. Traffic forecast show that this trend will continue in the foreseeable future. Over the past decade growth on Internet related technologies and services leveraged on the exponential growth of available bandwidth, mainly supported by EDFA and WDM technologies at the optical layer and by extensive silicon chip integration in the electronics. Such technologies seem however to have reached saturation. We do not currently have competitive technologies able to exploit the additional bandwidth in the fiber. Similarly, scalability issues in dynamic heat dissipation in silicon technology have sensibly slowed down the capacity growth of network electronic equipment, such as IP routers.

Power consumption has also become a major concern in networking and more in general for the IT sector. Although IT related power consumption is currently between 1 and 2 % of the total, it is growing three times faster than the rest. The concern generated by such growth has pushed much of current research on networks towards reduction of power consumption. Probably the most relevant example is the Greentouch consortium, which includes among the top vendors, operators and universities in the world, and aims at reducing power consumption in networks by 1000 times within the next 5 years. In addition, networking technology could help further reduce global power consumption in other areas such as transport, by making ideas (i.e., bits) travel instead of people.

Although new techniques are being developed such as multi-level modulations, OFDM, coherent transmission, that aim at increasing the data rate that can be carried over a single fiber, it is believed that their ability in reducing the cost-per-bandwidth ratio will not be comparable to that allowed by EDFA and WDM technologies over the past decade. Although network capacity will keep growing over the foreseeable future, its growth will not benefit from the same economies of scale we have seen in the past, and bandwidth will be increasingly perceived as a scarce resource. Such trend will continue unless significant innovation will emerge in the optical transmission (e.g., with new transmission media) and in the electronic processing domain (e.g., for switching and routing functions). The quest for reduction in bandwidth cost will thus need to be tackled at multiple levels.

Besides looking for technological advances at the physical layer, means for a more efficient use of the scarce bandwidth resources are also required. For example achieving quality of service through massive static over-provisioning of network resources is not sustainable. Similarly, current use of routing traffic almost exclusively through IP routers is not sustainable. These technologies do not seem scalable to support the Internet evolution in the foreseeable future, where high peak rates will be required to offer quality of experience to a multitude of to bandwidth-hungry applications, such as HD video on demand and thin client computing.

We believe that dynamic and intelligent bandwidth allocation will be a key feature to enable efficient bandwidth usage and reduce the pressure on the network core. Such efficiency will be also reflected in reduced power consumption. Both access and core architectures will need to be redesigned to be more flexible, self-aware and self-managing, bandwidth and power efficient. We intend to further investigate a set of topics that will lead to the design and implementation of the smart network of the future. Such topics are highly interconnected and include:

- Dynamic traffic aggregation to support transparent optical networks and efficient multi-layer traffic switching (i.e., switching traffic at the lowest/most efficient layer)
- Intelligent, self-aware networks, capable of reconfiguring links depending on application service requirements and network state
- Dynamic traffic switching at multiple layers: fiber, waveband, wavelength, packet (layer2), service (layer-3), as needed.
- Peer-to-peer based content delivery to help reducing the load on core networks (where peer-to-peer is used as a network strategy rather than as an individual user application)

Applications are invited for funded Ph.D. positions on the topics above, in the CTVR Telecommunications Research Centre, School of Computer Science and Statistics, Trinity College Dublin.

The successful candidate will have an excellent academic record (i.e., first class honors and ideally a postgraduate qualification, e.g. M.Sc.) in Computer Science, Electronic engineering or related discipline. She/he will be highly motivated, with strong communication skills and excellent knowledge of C/C++ and Python programming languages.

The successful candidates will have knowledge and experience in one or more of the following areas:

- Routing and switching protocols (IP, Ethernet, GMPLS, BGP, OSPF)
- Optical switched networks
- Practical experience in testbed implementations
- Peer to peer networks

The Ph.D. Position is for 4 years duration, and fully funded by either SFI or IRCSET framework. It includes payment of academic fees and a tax-exempt annual stipend.

For application please send an updated CV with at least two references and motivation letter to:

Marco.ruffini@cs.tcd.ie