

## Multi-dimensional Torus networks for current and next generation HPC systems: APEnet+ status and perspectives.

*Thursday, 12 February 2015 17:50 (20 minutes)*

The Apenet family interconnect cards have been out for more than a decade, pursuing the legacy of the APE custom massively parallel computing machines and bringing few concepts (3D-torus network, high bandwidth, low latency) to commodity clusters.

Over the years we have been added key features to our custom network, like Remote-DMA programming paradigm or, lately, NVidia peer-to-peer capability for tightly coupling GPUs with our network card.

The most strenuous efforts were made on keeping our communication interface IP cores (both for the host side - on PCIe interface, and on remote link side) up to date with current state-of-the-art technology. Apenet+ is the actual production-class interconnect card equipping Quong, the 16-node CPU-GPU hybrid system deployed.

Exploring and developing innovative ideas in various research fields, such as fault tolerance, fast address translation or embedded processors, has brought interesting results and significant impact on our system performances. A comprehensive view of these topics will be given in this talk, along with the results achieved at this point. Perspectives of future work will also be given, covering topics such as integrating next-generation ARM System-on-Chips, collective communication optimizations and other planned hardware and software enhancements.

This work has been funded by the European FET FP7 project EURETILE (grant 247846) and by the MIUR project SUMA

**Presenter:** AMMENDOLA, Roberto (ROMA2)