

A cluster-based router is a new router architecture that is composed of a cluster of commodity processing nodes interconnected by a high-speed and low-latency network. It inherits packet processing extensibility from the software router, and forwarding performance scalability from clustering.

In this thesis, we describe a prototype cluster-based router, including the design of the cluster-based router architecture and the addressing of critical issues such as the design of a highly efficient communication layer, reduction of operating system overheads, buffer recycling and packet packing. By experimental evaluation, we expose its forwarding capacity scalability and latency variance. We also evaluate and analyze the potential hardware bottlenecks of its commodity processing nodes, and present the correlation between the reception and transmission capabilities of an individual port as well as ports on the same bus. We propose an adaptive scheduling mechanism based on system state information to manage the adverse effect of this correlation on the router performance.

We also investigate internal congestion in the cluster-based router. To manage the internal congestion, we propose two backward explicit congestion notification schemes: a novel queue scheduling method and an optimal utility-based scheme. We show the effectiveness of these schemes either by ns-3 simulation, experimental evaluation, or both. We also analyze the stability of the optimal utility-based BECN internal congestion control scheme through theoretical proof, simulation and experimental evaluation.