

A Single-Kernel Approach to OS Specialization and Node Resource Partitioning for Exascale Computing

Judicael A. Zounmevo[#], Kamil Iskra[#], Kazutomo Yoshii[#], Roberto Gioiosa^{\$},

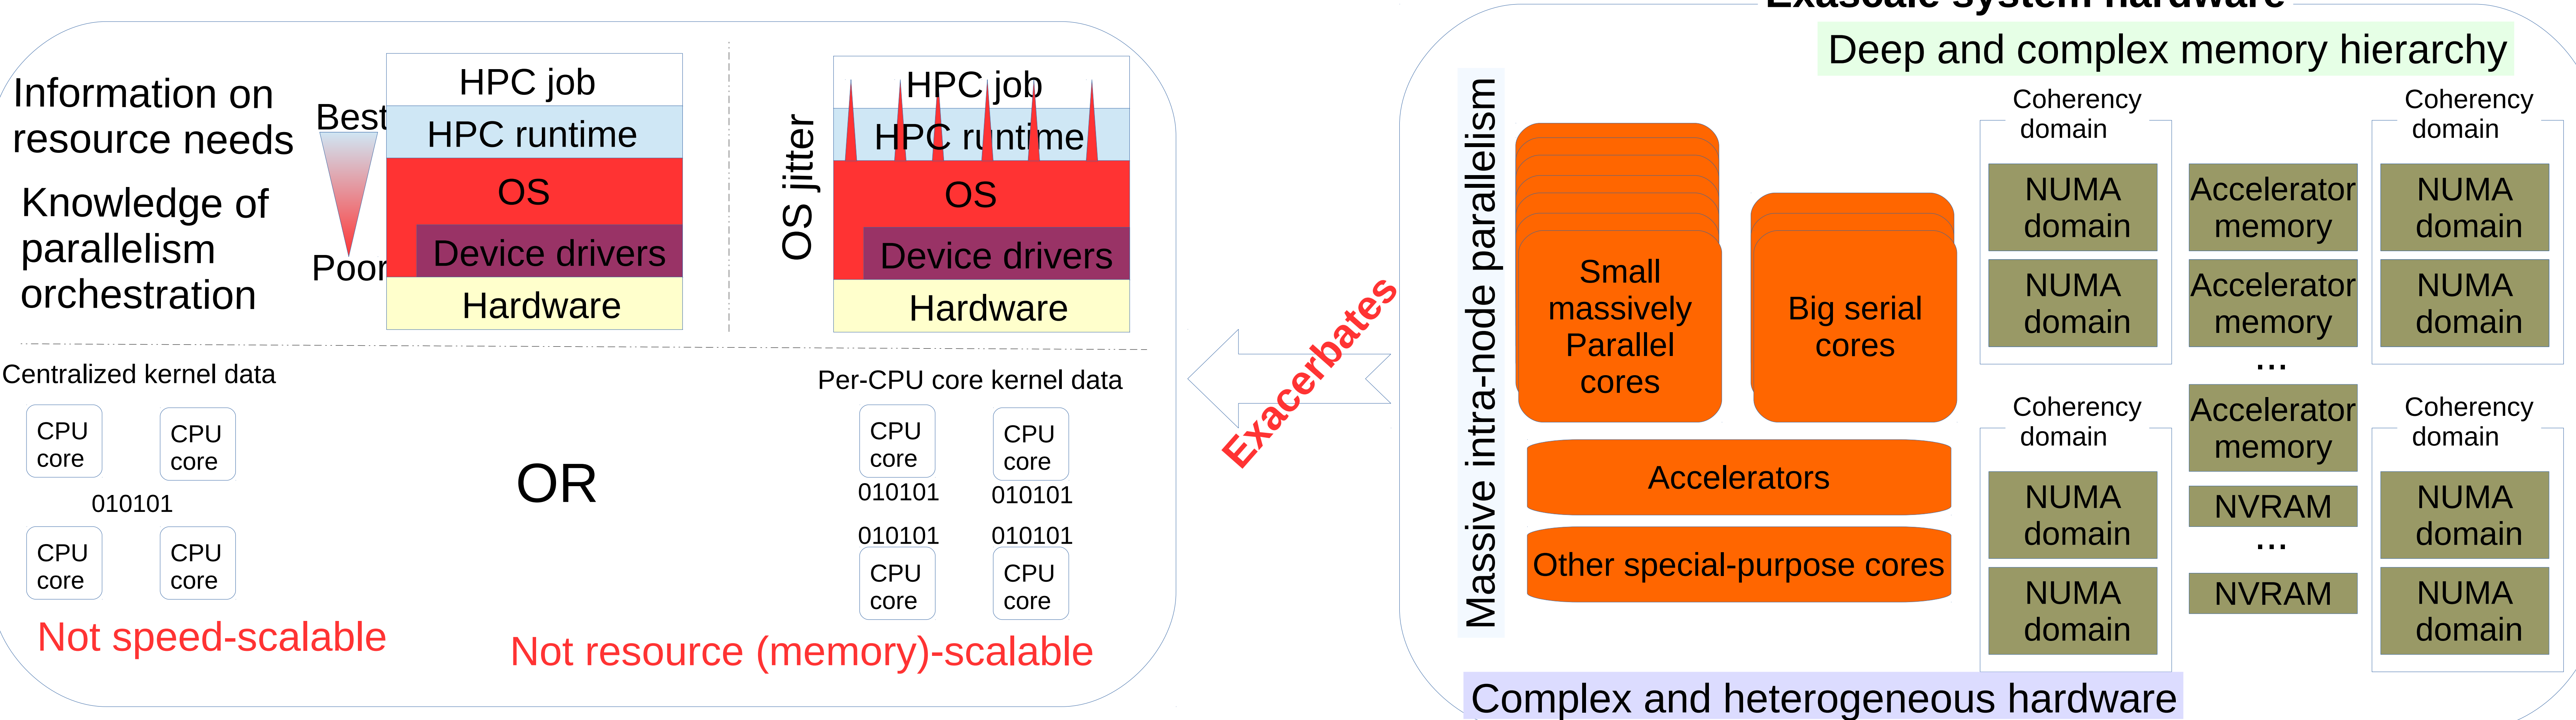
Brian C. Van Essen^{*}, Maya B. Gokhale^{*}, Edgar A. Leon^{*}

[#]Argonne National Laboratory. {jzounmevo@, iskra@mcs., kazutomo@}anl.gov

^{\$}Pacific Northwest National Laboratory. Roberto.Gioiosa@pnnl.gov

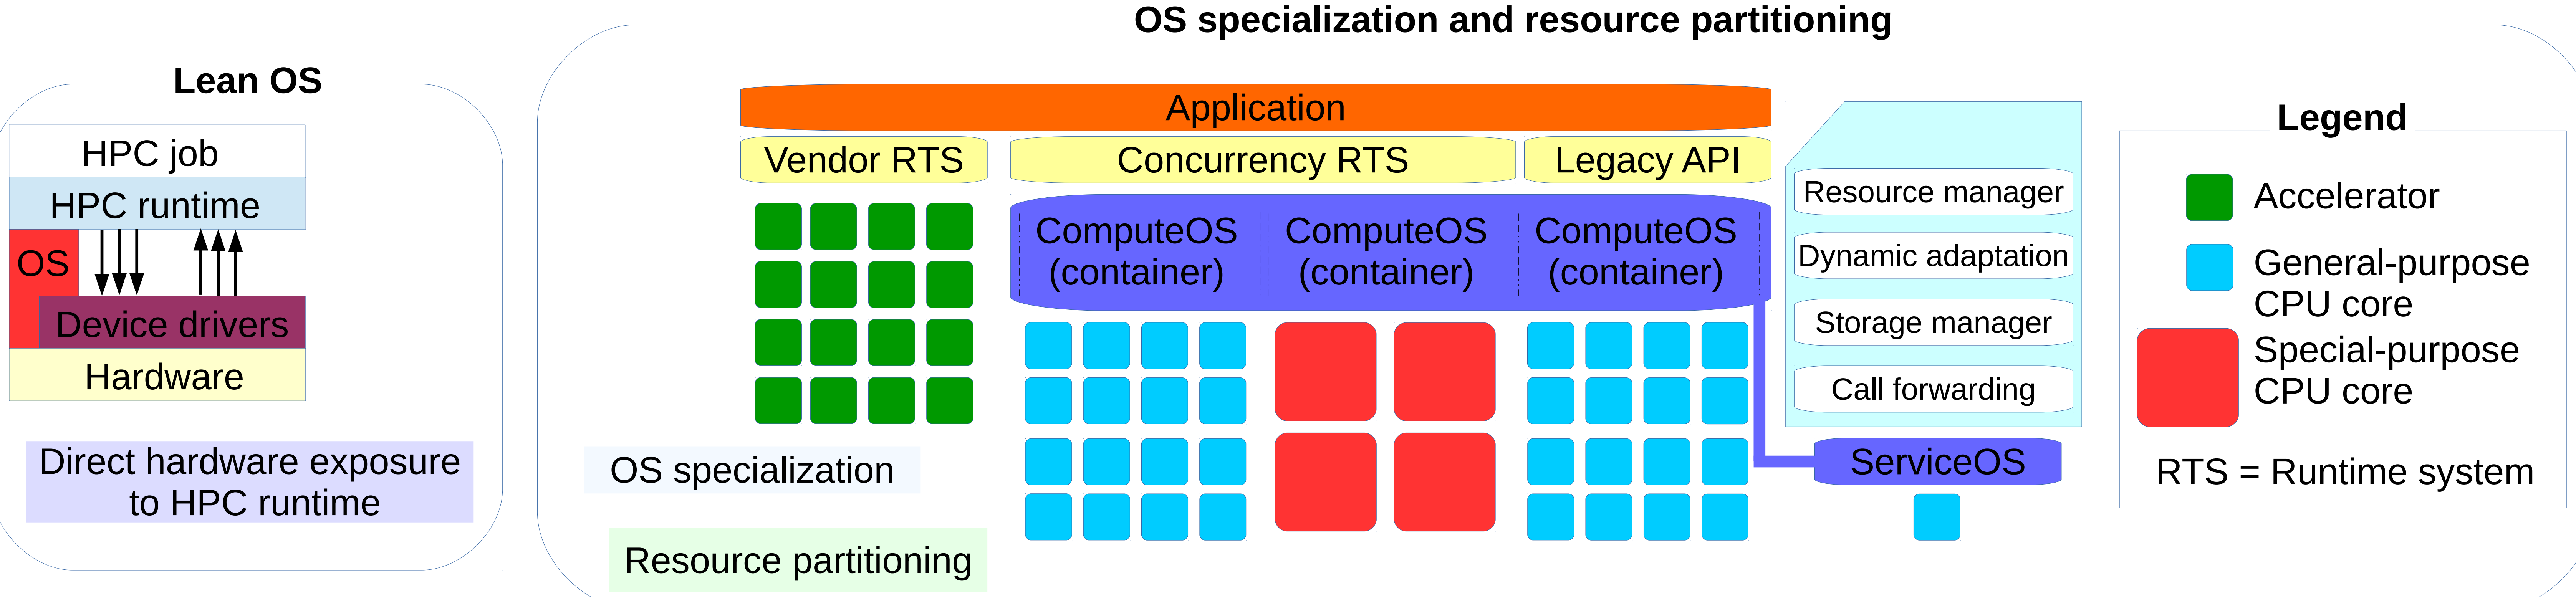
^{*}Lawrence Livermore National Laboratory. {vanessen1, maya, leon}@llnl.gov

General purpose OS may not fulfill exascale performance and scalability needs



... and there is a strong need for latency propagation avoidance

Node OS and runtime view of the Argo exascale OS architecture



Some novelties under development

