Success Story

Intel® Academic Program for oneAPI

Training Multifaceted Programmers for a Heterogeneous World

At the University of Massachusetts Lowell, Professor Yan Luo is extending the reach of the Electrical & Computer Engineering curriculum through affiliation with the Intel® Academic Program for oneAPI. Using technologies and tools from the program, Luo's students are getting the training they need to program future generations of heterogeneous hardware, such as novel cyber-physical systems.

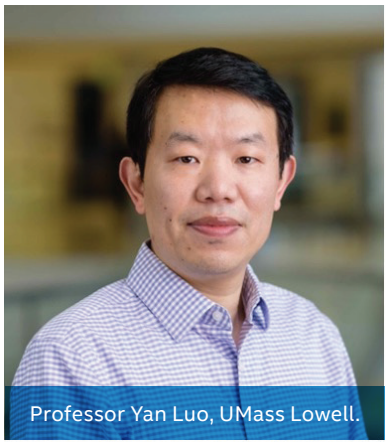

Professor Yan Luo, UMass Lowell.

One of the greatest technology challenges and opportunities of our time is the addition of intelligence to systems in new ways. This can include development of the internet of things, for example, or improving instrumentation and telemetry on network equipment. Broadly, this work consists of gathering novel data and analyzing it to yield the insights that can provide the greatest value possible. This involves integration of low-level sensor and other machine sources with high-level computing mechanisms such as AI and analytics.

At the University of Massachusetts Lowell, Electrical & Computer Engineering Professor Yan Luo applies state-of-the-art technologies to embedded computing and network processing, advancing these capabilities. His lab provides the link between computers and the rest of the world with cyber-physical systems such as acoustic sensors for detecting damage on wind turbines and pathogen sensors for wastewater and aquaculture. Another thrust of his work is with network-oriented workloads such as virtualized 5G radio access network functions and AI-driven network traffic classification.

Programming for heterogeneous architectures is an important aspect of both this research and of Dr. Luo's classroom teaching. To help support this requirement, he developed a course several years ago to teach OpenCL programming on FPGAs. That course set the stage for adopting Intel® Dev Cloud for instruction, as well as a more comprehensive approach to teaching heterogeneous computing. oneAPI provides a critical element of the approach, which Dr. Luo maximizes the value of through his involvement with the Intel Academic Program for oneAPI Educator Program.

Bridging the Instructional Gap to Heterogeneous Programming

In contrast to the Computer Science department, the Electrical & Computer Engineering curriculum at UMass Lowell (as elsewhere) does not traditionally focus on high-level languages such as Java and Python. Instead, its programming approach is embedded closer to the hardware, such as with VHDL and Verilog. Likewise, the Computer Science curriculum doesn't systematically address programming for GPUs and FPGAs.

Both departments have therefore lacked reach into programming for the mixed hardware architectures that are becoming increasingly important for real-world usages. Dr. Luo anticipates that mastery of those skills will unleash tremendous opportunities for the next generation of engineers. He notes that what has been missing is a high-level programming language that can help bridge the gap between general computing and coding for mixed resources such as CPUs, GPUs and FPGAs. oneAPI provides that unified approach, enabling a single code base to traverse all those resources.

"The curriculum modules that I'm authoring will introduce heterogeneous hardware architectures represented by Intel CPUs, GPUs and FPGAs. They will incorporate concrete coding examples to demonstrate how to design data-parallel algorithms and port existing applications. These concepts are of high importance to undergraduate and graduate students, which represent the future workforce."

- Professor Yan Luo, Department of Electrical & Computer Engineering, UMass Lowell

oneAPI: Programming for the New Era of Heterogeneous Computing

Hardware accelerators are critical for maximizing throughput and energy efficiency while driving down workload latency and cost on commercial-off-the-shelf (COTS) servers. Developers have used performance engines such as GPUs and FPGAs to supplement the CPU for years, although proprietary programming models such as CUDA have limited the reach of those efforts.

oneAPI changes all that, with a single, open model for code that can execute on CPU cores as well as various hardware accelerators. Intel oneAPI toolkits provide best-in-class compilers, performance libraries, frameworks, and analysis and debug tools, so developers can code once and run anywhere, from the largest supercomputers to compute nodes on the distributed edge.

Streamlining Adoption of oneAPI in Curriculum and Beyond

Transitioning coursework at the pace of technology is a constant challenge for educators, which is near the practical heart of the Intel Academic Program for oneAPI. Aside from requirements placed by the institution itself, instructors must create course materials, which can be complex and time consuming within the resource constraints of the academic calendar. The Program simplifies the process, with sample code, tools and the classroom-ready, maintenance-free Intel Dev Cloud.

Dr. Luo has availed himself and his students of these resources, and he continues to develop the role of the Program's teaching kits in his coursework. This adoption complements his participation in the Intel Software Innovator Program and other Intel points of contact to increase the depth of connection between UMass and Intel. Looking ahead, Dr. Luo is finding ways to push capabilities for heterogeneous programming with oneAPI out to larger groups of developers. By exploring the potential for coursework in the public domain such as Coursera, he hopes to scale on-campus research and instruction to the broader community and ecosystem to drive greater reach for intelligence in the world around us.

Learn More:

Intel® Academic Program for oneAPI

0623/OT/MESH/PDF 353921-001US