

ECE/ME/EMA/CS 759 High Performance Computing for Engineering Applications
Fall 2017 – Syllabus [**subject to change**]

Date	Title	HW Assigned	Reading Assignments Other Observations
09/06 [L01]	Syllabus related issues. Course overview.	HW01 (due 09/13): C programming related	Job Scheduling/Execution on Euler. For beginners: Linux command line online tutorial.
09/08 [L02]	From Code to Instructions. The FDX Cycle. Pipelining. Measuring Computer Performance. Superscalar architectures		Read Chapter 5 of Brian W. Kernighan and Dennis M. Ritchie “The C Programming Language” book . Read gdb tutorial .
09/11 [L03]	Memory Aspects. Caches. Virtual Memory		Read the git & GitHub tutorial . Read the CMake tutorial .
09/13 [L04]	The Walls to Sequential Computing. Moore’s Law. Parallel Computing. Flynn’s Taxonomy. Amdahl’s Law.	HW02 (due 09/20): C programming related	Read the Manferdelli and Amdahl articles
09/15 [L05]	Vectorization. Code Optimization Aspects		
09/18 [L06]	Multi-core Parallel Computing with OpenMP. Parallel Regions. Work Sharing		
09/20 [L07]	OpenMP Work Sharing (Tasks)	HW03 (due 09/27): Vectorization related	Chapter 12, from Agner Fog's optimization tutorial .
09/22 [L08]	Parallel Programming with OpenMP: Variable Scoping		
09/25 [L09]	OpenMP NUMA Aspects. Multi-level parallelism in OpenMP		
09/27 [L10]	OpenMP Hands-on	HW04 (due 10/04): OpenMP related	Knuth paper on premature optimization.
09/29 [L11]	GPU Computing Intro. The CUDA Programming Model.		
10/02 [L12]	CUDA Execution Configuration		
10/04 [L13]	GPU Memory Spaces	HW05 (due 10/11): OpenMP related	GPU computing evolution article of Nickolls & Dally
10/06 [L14]	GPU Scheduling Issues		
10/09 [L15]	Execution Divergence. Control Flow in CUDA. Global Memory Access Patterns and Implications.		
10/11 [L16]	Guest Lecture, AMD Staff	HW06 (due 10/18): GPU/CUDA related	Skim through 2017 GPU Tech Conference (GTC) talk titles . Skim through CUDA Programming Guide .
10/13 [L17]	CUDA Shared Memory Issues. Atomic operations in CUDA.		
10/16 [L18]	Using the CUDA profiler. Example: 1D Stencil Operation		
10/18 [L19]	Tiling as a Programming Pattern in CUDA Example: Vector Reduction in CUDA	HW07 (due 10/25): GPU/CUDA related	CUDA C Best Practices Guide .
10/20 [L20]	CUDA Optimization Issues. Resource Utilization Issues. Parallel Prefix Scan on the GPU. Using Multiple Streams in CUDA.		

10/23 [L21]	Streams, and overlapping data copy with execution.		
10/25 [L22]	GPU Computing with thrust and cub	HW08 (due 11/01): GPU/CUDA related.	Paper on thrust in GPU Gems 4, by Nathan Bell and Jared Hoberock.
10/27 [L23]	GPU Computing: Advanced Features.		
10/30 [L24]	Computing with Supercomputers. MPI Parallel Programming General Introduction, Point-to-Point Communication		
11/01 [L25]	MPI Parallel Programming Point-to-Point communication: Blocking vs. Non-blocking sends	HW09 (due 11/08): GPU/CUDA related.	2005 article of Dongarra et al. for an overview of HPC .
11/03 [L26]	MPI Parallel Programming: MPI Collectives		Two-page Final Project Proposal due at 11:59 PM
11/06 [L27]	MPI Data Types. Class Wrap-up.		
11/08	NO CLASS	HW10 (due 11/15): thrust/cub related	
11/10	NO CLASS		
11/13	NO CLASS		
11/15	NO CLASS	HW11 (due 11/22): MPI related	
11/17	NO CLASS		
11/20	NO CLASS		
11/22	NO CLASS	HW12 (due 12/04): MPI-related	
11/24	THANKSGIVING RECESS		
11/27	NO CLASS		
11/29	NO CLASS		
12/01	NO CLASS		
12/04	NO CLASS		
12/06	NO CLASS		
12/08	NO CLASS		
12/11	NO CLASS		
12/13	NO CLASS		

Final Project Presentations time slots available during finals week. Scheduled through doodle.

Final Project due on Monday, Dec. 21 at 11:59 PM (submitted via GitHub)

Time for Final Exam: TBA.