

IEEE-TCPP Parallel and Distributed Computing Curriculum for Computer Science and Engineering Undergraduates - **Version 2 beta**

Sushil K Prasad, Sheikh Ghafoor
Alan Sussman, Charles Weems, and R. Vaidyanathan

CDER Center

Contact: sushil.prasad@utsa.edu

EduHPC-20, Nov 13, 2020

TCPP Curriculum Initiative:

<https://tcpp.cs.gsu.edu/curriculum/>

Outline

- IEEE-TCPP Curriculum – 7 mins
 - History, Opportunities
 - Key Activities and Milestones
 - ACM/IEEE 2013 CS Curriculum Taskforce
 - provided direct link to us for rigorous coverage
 - How formulated? Evaluated? Resources?
- **Announcing Curriculum Version 2-beta** – 18 mins
 - Revision Aspects: Big Data, Energy, Distributed Computing
 - Pervasive topics
 - Areas:
 - Programming
 - Architecture
 - Algorithms

Curriculum Planning Workshops at DC (Feb-10) and at Atlanta (April-10)

- Goals
 - setup mechanism and processes which would provide periodic curricular guidelines
 - employ the mechanism to develop sample curriculums

- Agenda:
 - Review and Scope
 - Formulate Mechanism and Processes
 - Preliminary Curriculum Planning
 - Core Curriculum
 - Introductory and advanced courses
 - Impact Assessment and Evaluation Plan

Main Outcomes

**- Priority:
Core curriculum revision at
undergraduate level**

- Preliminary Core Curriculum Topics

-Sample Intro and Advanced Course Curriculums

Some Participants at the Planning Workshop, Washington DC, Feb 5-6, 2010



3 Curriculum Areas + Cross-Cutting Architecture, Programming, Algorithms

TCCP Curriculum Example

Algorithms Topics		Bloom#	Course	Learning outcome and teaching notes
Algorithmic problems				Algorithmic problems section contains parallel algorithms for certain problems. The important thing here is to emphasize the parallel/distributed aspects of the topic
<i>Communication and Synchronization</i>				Understand (at the pseudo-code level) how certain patterns of communication can be implemented in a parallel/distributed model. Also appreciate the cost of communication in PDC.
	Reduction and Broadcast for communication and synchronization	C	Data Struc/Algo	Understand, for example, how recursive doubling can be used to for all-to-one reduction, and its dual, one-to-all reduction, in $\log(p)$ steps. The same applies to all-to-all broadcast and all-to-all reduction. Recognize that all-to-all broadcast/reduction are synchronizing operations in a distributed (event-driven) environment.
	Parallel Prefix (Scan)	C	Data Struc/Algo	Understand the structure of at least one simple parallel prefix algorithm. One could consider recursive or iterative approaches (such as those of Ladner-Fischer, Kogge-Stone, Brent-Kung)
	Multicast	N		
	Permutation	N		

Early Adopter and Training Programs

- Over 100 institutions worldwide
 - Spring-11: 16 institutions ; Fall'11: 18;
 - Spring-12: 21; Fall-12: 25 institutions, Fall-13: 25 institutions, Fall-14: 25, Fall-15: 13
 - Most from US (4 year to research institutions, one high school)
 - Some from South America, a few from Europe, fewer from Asia (India, China, Indonesia, Singapore), Middle East
- **NSF CyberTraining** Training Workshops - Summer 2018-20
 - **UMass/Maryland; Tennessee Tech**
 - NSF/Intel funded Stipend up to \$1500-5000/proposal
 - *Instructor training + adoption plans*

Edu* Workshop Series

- **EduPar-11** at Alaska, IPDPS-2011
 - Receive feedback from the Adopters
 - Stimulate discussion of curricular and other educational issues.
- EduPar-12 at Shanghai, IPDPS-2012
 - **A regular satellite workshop of IPDPS**
 - EduPar-15 @IPDPS, May, India; EduPar-16, Chicago, EduPar-17 in Orlando; EduPar-18 in Vancouver, EduPar-19 @ IPDPS, Brazil, **EduPar'20 - online**
- **EduHPC** Workshop at SC-13 + BOF at SIGCSE-14
 - EduHPC-14 @ SC-14, Nov – New Orleans; EduHPC-15 in Austin, EduHPC-16, EduHPC-17, EduHPC-18 in Dallas, EduHPC-19 @ SC in Denver
 - **EduHPC-20 @ SC - online**
- **Euro-EduPar Aug 2015**; Euro-EduPar-2016, EEP-2017, EEP-18,
- **EduHiPC 2018 @ HiPC in Bangalore** – for India and the region
 - EduHiPC'19 @ HiPC in Hyderabad Dec'19
- **CFP JPDC Special Issue**

CDER Courseware Website

Upload and Search Course Material

- **Type:**
 - Slides, Syllabus, Tutorial, Video
 - Animation, Article, Award, Blog, Book, Competition
 - Course Template, Course Module, Data
 - Hardware Access, Software/Tools
 - Proposal, Report
- **Courses:**
 - CS1, CS2, Systems, Data Structures and Algorithms, ...

- **NSF/TCPP Topic/Subtopic Classification:**

ALGORITHMS

Parallel and Distributed Models and Complexity

Algorithmic Paradigms

Divide & conquer (parallel aspects)

Algorithmic problems

ARCHITECTURE

PROGRAMMING

CROSS-CUTTING

- [open](#) - Work in Progress

CDER Book Project

- Lack of suitable textbooks to integrate PDC topics into the core courses
 - CS1, CS2, Systems, and Data Structures and Algorithms
 - **Part 1 - For instructors:** Basic Concepts and References on what and how to teach
 - **Part 2: For students:** Supplemental teaching material for core courses
 - 9 chapters
 - 40K chapter downloads – [free downloads](#)
- 2nd Volume – Published Nov'18
- **Vol 3 (CFP upcoming)** – Early Adopter course and topic exemplars and accompanying resources

Curriculum Version II Activities

	Areas	Architecture	Algorithms	Programming
New Aspects	Area Lead/ Aspect Lead	Chip Weems	Anshul Gupta	Alan Sussman
Exemplars	Sushil Prasad	Karen Karavanic, Eric Freudenthal	Erik Saule, Duane Merrill, David Bunde	David Brown, Eric Freudenthal
Distributed	Vaidyanathan Ramachandran	Vaidyanathan Ramachandran, Manish Parashar	Vaidyanathan Ramachandran, Costas Busch, Denis Trystram	Alan Sussman, Chi Shen
Big Data	Trilce Estrada	Craig Stunkel	Cynthia Phillips,	Debzani Deb
Energy	Krishna Kant, Craig Stunkel	Craig Stunkel, Karen Karavanic	Denis Trystram	John Dougherty
Pervasive	Sheikh Ghafoor	Craig Stunkel, Eric Freudenthal	Robert Robey, Martina Barnas	Sheikh Gafoor, Eric Freudenthal

- Timeline:
 - Version-1.9 announced @ EduPar'19
 - Revised based on expert reviews
 - **Version-2-beta released @ EduHPC'20**
 - **Public Feedback:** sushil.prasad@utsa.edu
 - **Companion Activities:**
 - Exemplars
 - CE-oriented TCPP Curriculum
- **New:** NSF Institute Planning Grant => 4 planning workshops
 - SC'19
 - SIGCSE'20 - online
 - July 27, 2020 - online
 - **NSF – Spring'21**

Pervasive and Emerging Topics

- Craig Stunkel IBM
- Randy Bryant CMU
- Martina Barnas Indiana University
- Eric Freudenthal University of Texas El Paso
- Bob Robey LANL
- **Sheikh Ghafoor Tennessee Tech**

Major Update

- Crosscutting view of the 2012 curriculum is modified (no crosscutting table in the 2020 curriculum)
 - Pervasive Concepts
 - Fundamental in nature
 - Likely to remain unchanged for long time
 - Emerging Topics
 - Significant current or emerging interest
 - Continually evolving as topics mature and newer topics with time
 - Often Better suited for advanced courses but may be introduced in lower level courses in a limited way.
 - Applications of many of these topics will be familiar to students

Pervasive Concepts

- Broad themes that are conceptual underpinning of PDC
- Appears at different depth throughout the curriculum
- Transcend PDC and often computing
- Should be taught at the core
 - May be through unplugged activities
 - least at “C” level in terms of Bloom classification.
- These should be further reinforced in architecture, programming, and algorithms
- Four topics have been identified
 - Asynchrony
 - Concurrency and Sequential Dependency
 - Locality
 - Performance

Pervasive Concepts

Topics	Learning Outcome and Teaching Suggestion
Asynchrony	1) Understand cause and effect of Asynchrony and how to ensure that computational correctness. 2) Understand asynchrony is the characteristics of modern system and understand asynchrony in PDC context. 3) Can utilize a standard coordination strategy to prevent incorrect operation due to uncontrolled concurrency (race conditions).
Concurrency & Dependency	1) Understand concurrency is an algorithmic property and difference between concurrency and parallelism. 2) Understand sequential dependency limit degree of concurrency and hence parallelism. 3) Identify sequential dependency in algorithms.

Emerging Topics

Topics	Core Bloom Level	Learning Outcome and Teaching Suggestion	Where Covered
Machine Learning	N	Understand that accelerators, parallel matrix decompositions, parallel matrix multiplication, and vectorized operations are fundamental aspects that enable modern machine learning	ParProg, Machine Learning(K)
Mobile Adhoc Network	N	Understand that mobility admits a dynamic topology and a different relationship between "sender" and "receiver." Explore the impact of these differences in the system.	Networking (c)
Quantum Computing	N	Know that quantum computing algorithms can solve some problems (e.g., optimization problems) exponentially faster than classical algorithms via parallelism, and that quantum processors to implement some of those algorithms are under construction, using new parallel programming models	Arch2, ParProg(k)

Programming Topics

Alan Sussman

University of Maryland and NSF

Programming Topics - Overview

- Main goal is to introduce parallel programming topics into intro programming, data structures, and systems classes
 - Secondary goal is to target upper-level classes
- High-level themes include:
 - **Paradigms and notations** – SIMD, shared memory, message passing, client/server, big data stack, threads, tasks, data parallel, etc.
 - **Semantics and correctness** – synchronization, concurrency defects, ...
 - **Memory models** – sequential consistency, weak consistency, ...
 - **Performance and energy** – computation and data decomposition, scheduling/mapping, data layout and locality, tools and metrics
- Most topics at a shallow level (Bloom level C or K) for intro courses, but at a deeper level for upper-level advanced courses (or deferred to upper-level completely, so at N Bloom level for intro courses)

Programming Topics - Updates

- Incorporated new programming topics related to distributed computing (e.g. client/server), big data (e.g., MapReduce), and power/energy
- Eliminated some topics from original guidelines completely, since no longer relevant
- Added small number of other topics missed in original guidelines, or newer ideas (e.g., accelerator programming)
- Revised learning outcomes, Bloom levels for some topics (mainly between C and K)
- Added learning outcomes and Bloom levels for advanced courses

Programming Topics

– Input from Reviewers

- Several suggestions to do a better job on pervasive ideas
 - Across programming areas, algorithms, architecture – e.g., scalability
- Better definitions of terms and acronyms
- More limited energy/power topics – keep them high level
- Some suggestions for updating topics, new topics, eliminating topics no longer relevant
 - Informed discussion for revisions

Architecture Topics

Chip Weems
Univ of Massachusetts

Overview

- Limited coverage in core courses
 - Only K or C Bloom levels
 - Most in Systems course (a few topics in CS1, CS2, such as floating point, atomicity as basis for threaded libraries, event handlers as threads in GUIs)
 - Goal is for all students to be aware of hardware aspects of PDC affecting and providing opportunities for algorithmic problem solving
- Most topics in table are for advanced courses
 - Many already normally covered there, at or beyond suggested bloom levels

Topics Overview

- Classes of Parallelism: data (superscalar, SIMD, dataflow), pipelines (OoO, streams), control (multicore, multithread, heterogeneity), shared memory (snooping, directory), distributed memory (topology, latency)
- Underlying mechanisms (caching, atomicity, consistency, coherence, interrupts/events, handshaking, ID, virtualization)
- Floating point representation (in support of HPC)
- Performance metrics (IPC, benchmarks, network/memory bandwidth, peak performance, sustained performance)
- Power (power/energy, larger scale, embedded, density, static/dynamic, DVFS, heterogeneous cores)
- Scaling (Big data, HPC, fault tolerance, data bound computation, volume, velocity, scale out, cost of data movement)

Algorithms

R. Vaidyanathan, LSU

Algorithms Overview

- Algorithms topics are recommended for coverage along with their sequential counterparts to minimize additional instruction time
- Organized into three sub-areas:
 - ***Parallel/Distributed Models and Complexity:*** Foundational topics aimed at equipping the students with basic knowhow for designing and analyzing parallel algorithms
 - ***Algorithmic Techniques:*** Recurrent themes or constructs that are generally useful in designing a wide variety of parallel algorithms
 - ***Algorithmic Problems:*** Basic problems for which learning both the sequential and parallel algorithms would be considered valuable for almost all CS/CE students

Algorithms: Major update 1

- Model(s) of Choice:
 - Understand concurrency basics without the trappings of real systems (routing, data alignment etc.).
 - Recognize the PRAM as embodying the simplest forms of parallel computation: Embarrassingly parallel problems can be sped up easily just by employing many processors.
 - Recognize how a completely connected network abstracts away from routing details. Recognize the difference between the model(s) and real systems.
 - Such a Model of Choice (MoC) is assumed to be chosen and adopted by the instructor on which PDC concepts would be discussed.

Algorithms: Major update 2

- Deeper Algorithmic Experience:
 - Experience through class instruction and assignment/project the (1) design, (2) analysis, and (3) implementation aspects of at least one parallel or distributed algorithm of choice in detail.
 - Master PDC algorithmic concepts through a detailed exploration, including recognizing how algorithm design reflects the structure of the computational problem(s) and the PDC model/environment.
 - Possible computational problems (to be chosen by instructor(s)) include, but are not limited to, matrix product, map reduce, sorting, search, convolution, a graph algorithm of your choice, etc.
 - We recommend that all CS/CE undergrads have this experience at some point before graduating.

Q&A

TCPP Curriculum Initiative:

<https://tcpp.cs.gsu.edu/curriculum/>

Feedback & Participation: sushil.prasad@utsa.edu