TCPP Curriculum Initiative: Integrating Parallel and Distributed Computing in Early Computing Classes

Sushil K Prasad, UT San Antonio
Alan Sussman, U. Maryland
Chip Weems, UMass
Sheikh Ghafoor, TTU
R. Vaidyanathan, LSU

EduPar-24, May 27, San Francisco
https://tcpp.cs.gsu.edu/curriculum/

Public Feedback on TCPP Curriculum & Contact
sushil.prasad@gmail.com

Sponsors

TCPP Curriculum Initiative

What should every Computer Science and Engineering Student know about Parallel and Distributed Computing (PDC)?
https://tcpp.cs.gsu.edu/curriculum/

- **Areas:** Programming, Architecture, & Algorithms
  - Version 1 – 2012
- **New Aspects:** Big Data, Energy, Distributed Computing, Pervasive topics
  - Version-2-beta released 2020
  - URL

- **Companion Activities:**
  - CE-oriented TCPP Curriculum
  - NSF Project on CS1/CS2 Exemplars
    - Recruiting Testing teams
    - Apply by March 31st: **URL**

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

Main Outcomes

- Priority:
Core curriculum revision at undergraduate level

- Preliminary Core Curriculum Topics

- Sample Intro and Advanced Course Curriculums

CDER - Center for Parallel and Distributed Computing Curriculum Development and Education Resources - Timeline

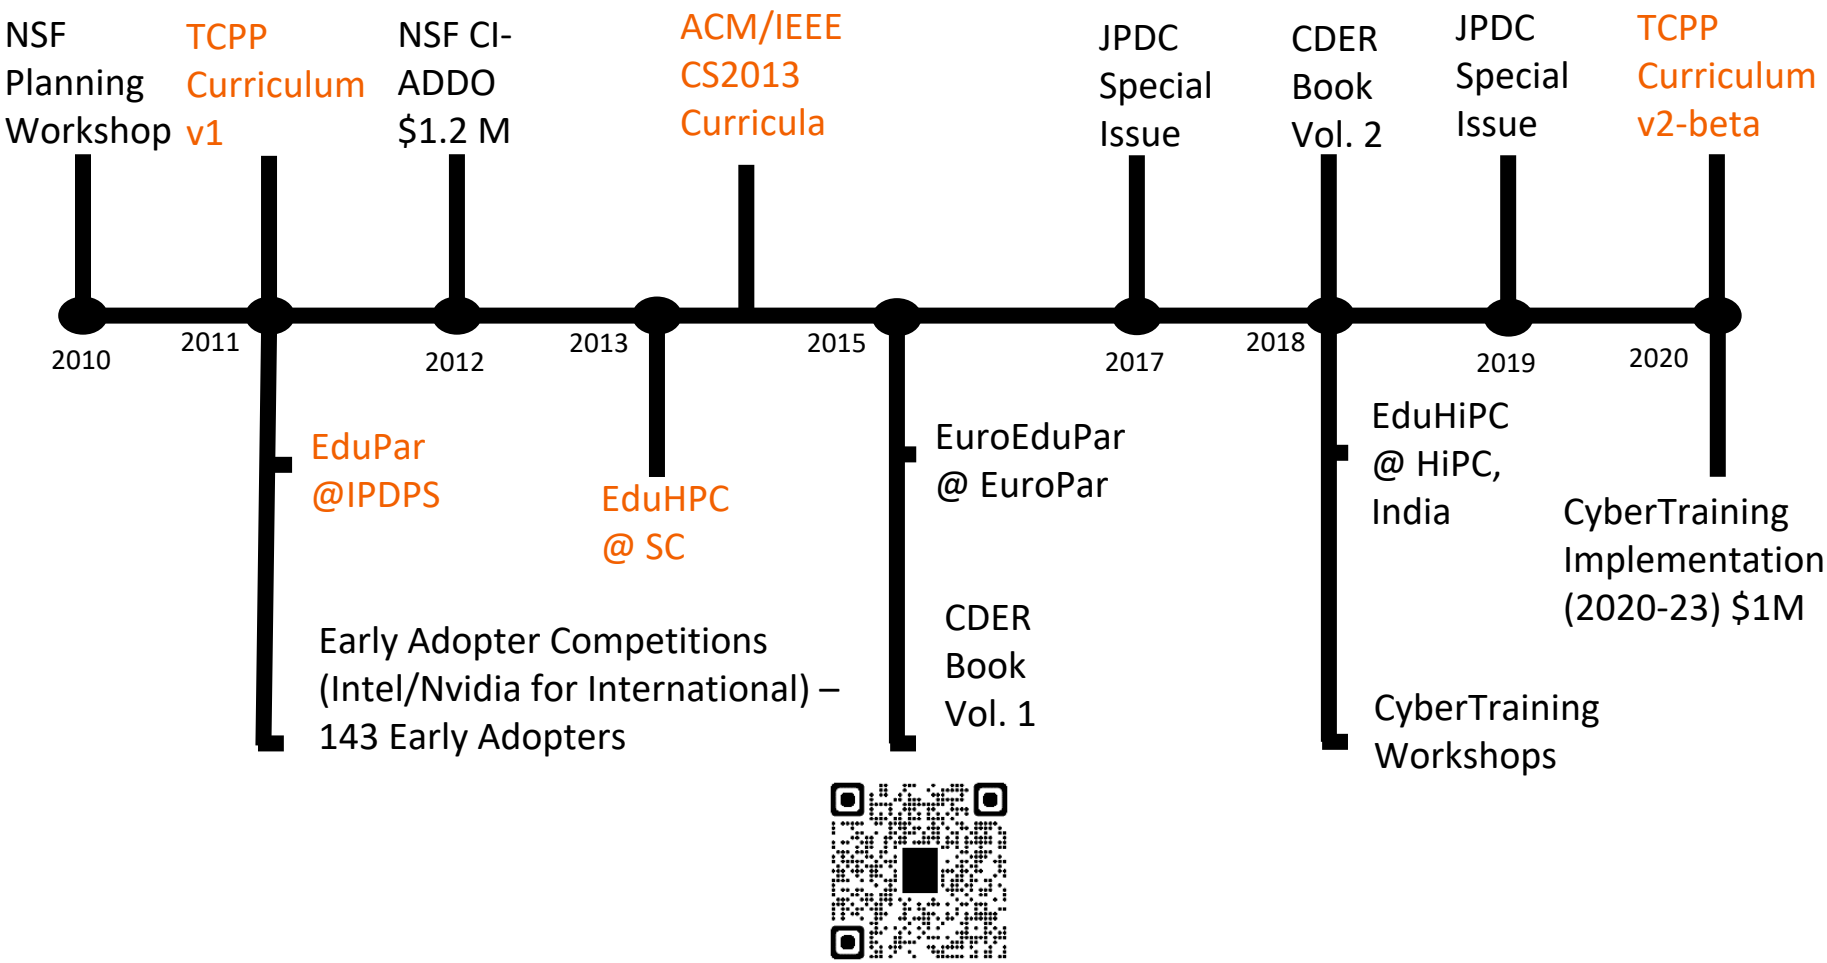

3 Curriculum Areas + Cross-Cutting
Architecture, Programming, Algorithms

TCPP 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
Communication and Synchronization			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 200 early adopter and trainee 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 Workshops – Weeklong in Summer 2018-24
  - **UMass,** LSU; UMass/Maryland; Tennessee Tech
  - NSF funded stipend up to $5K/proposal
  - *Instructor training + adoption plans*
- Additional Training workshops
  - **SIGCSE** 2023, 2024
  - **HiPC** 2022, 2023, **2024**

Edu* Workshop Series

- **EduPar-11** at 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 India; EduPar-16, Chicago, EduPar-17 Orlando; EduPar-18 Vancouver, EduPar-19 Brazil, EduPar'20, EduPar21 & EduPar22 - online – May 30, EduPar23 FL, **EduPar'24, San Francisco, May 27**
- **EduHPC at SC-13** + BOF at SIGCSE-14
  - EduHPC-14, New Orleans; EduHPC-15 Austin, EduHPC-16, EduHPC-17, EduHPC-18 Dallas, EduHPC-19, EduHPC-20 - online**,** EduHPC-21 @ SC – hybrid, EduHPC-22, EduHPC-23, **EduHPC'24 – Atlanta.**
- **Euro-EduPar** Aug 2015; Euro-EduPar-2016, EEP-2017, EEP-18
- **EduHiPC 2018 @ HiPC**– for India and the region
  - EduHiPC'19 Hyderabad, EduHiPC'21 – online, EduHiPC'22, EduHiPC'23 in Goa**, EduHiPC'24 in Bangalore**

Additional CDER Resources

- CDER Book series:
  - Vol 1: Topics in Parallel and Distributed Computing
    - Introducing Concurrency in Undergraduate Courses, *Morgan Kaufman*
  - Vol 2: Topics in Parallel and Distributed Computing
    - Enhancing the Undergraduate Curriculum: Performance, Concurrency, and Programming on Modern Platforms, *Springer*
  - Free Pre-Print Version on CDER site (**50K+ downloads**)
  - Plan for 3rd and 4th Volume – Experience of Adopters; Version 2 topics
    - Exemplars + Resources on courses and topics
- JPDC Special Issue - Keeping up with Technology: Teaching Parallel, Distributed and High-Performance Computing (2017, 19, 21)
  - **CFP -** Paper Submission deadline: May 31, 2024, URL

CDER Books - Topics in Parallel and Distributed Computing for Undergraduates

50K+ Chapter Downloads

Free Downlo ad

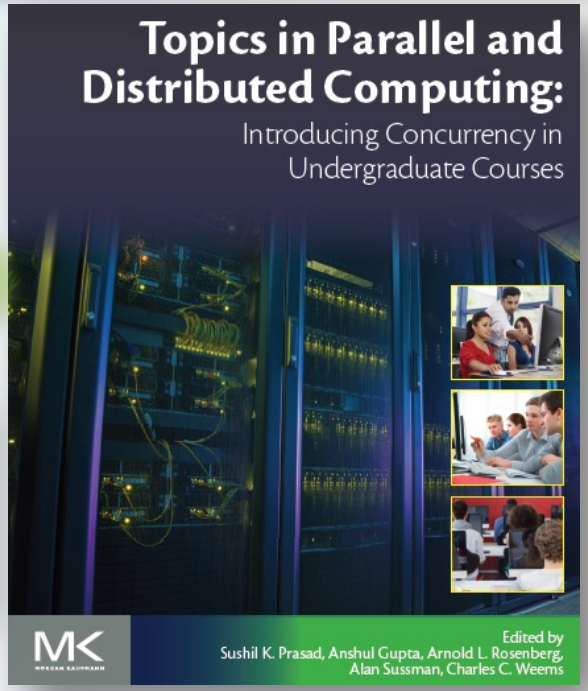

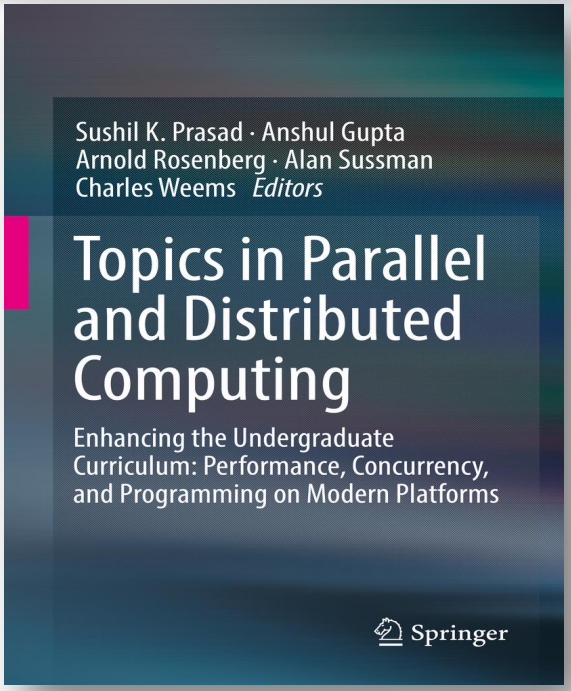

Free Downlo ad

PART 1 FOR INSTRUCTORS
5 Chapters

- ✓ Hands-on Parallelism with no Prerequisites and Little Time Using Scratch
- ✓ Parallelism in Python for Novices
- ✓ Modules for Introducing Threads
- ✓ Introducing Parallel and Distributed Computing Concepts in Digital Logic
- ✓ Networks and MPI for Cluster Computing

PART 2 FOR STUDENTS
4 Chapters

- ✓ Fork-join Parallelism with a Data-Structures Focus
- ✓ Shared-Memory Concurrency Control with a Data-Structures Focus
- ✓ Parallel Computing in a Python-Based Computer Science Course
- ✓ Parallel Programming Illustrated through Conway's Game of Life

PART 1 FOR INSTRUCTORS
5 Chapters

- ✓ What do we need to know about parallel algorithms and their efficient implementation?
- ✓ Models for Teaching Parallel Performance Concepts
- ✓ Scalability in Parallel Processing
- ✓ Energy Efficiency Issues in Computing Systems
- ✓ Scheduling for fault-tolerance: an introduction

PART 2 FOR INSTRUCTORS
5 Chapters

- ✓ MapReduce - The Scalable Distributed Data Processing Solution
- ✓ The Realm of Graphics Processing Unit (GPU) Computation
- ✓ Managing Concurrency in Mobile User Interfaces with Examples in Android
- ✓ Parallel Programming for Integrative GUI Applications

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

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 Merril, 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 Ghafoor, Eric Freudenthal

- **Timeline:**
  - **Version-2-beta released** @ EduHPC'20
    - Public Feedback: sushil.prasad@gmail.com
  - **Companion Activities:**
    - CE-oriented TCPP Curriculum
    - NSF Project on CS1/CS2 Exemplars Development (Oct 23-26)
      - Recruiting Testing teams
      - Apply by March 31st: **URL**
- NSF Institute Planning Grant => 5 planning workshops

1. SC'19, SIGCSE'20, July'20, & Mar'21 - online
2. NSF Reporting Workshop – Oct'21 - See Final Report

Computer Engineering Curriculum

- Current Status
  - PDC Principles
    - Concurrency, Asynchrony and Locality
    - Decomposition and Coordination
    - Performance and Pitfalls
  - CE Courses
    - Intro, Math, Logic, Circuits, Programming, Signals and Communication, , Networks, Embedded Systems, Organization and Architecture, OS, CPS, ML
  - Broad areas
    - Hardware and Architecture
    - Programming and Algorithms
    - Communications and Systems
  - PDC intro ideas

- In Progress
  - Subcommittees in each broad area
  - Within broad areas, mapping topics to courses

- In the future
  - Mapping topics to PDC principles
  - Curriculum Guideline

We need your help