

# AICTE Report

---

## Model Curriculum

for

Undergraduate Degree Courses

in

**Computer Science and Engineering**  
(Engineering & Technology)

Report submitted by the committee constituted by AICTE for model design.  
March 2022

---

# TABLE OF CONTENT

<b>Framework for CSE Curriculum Design</b>	1
Context and Background	1
Graduate Attributes	2
<b>CSE Core Courses</b>	3
Coverage of CS Graduate Attributes by the Core Courses	5
Micro Specializations (and Professional Electives)	5
Recommendations for Online Credits	6
Recommendations for Multiple Exits	6
<b>Syllabus for CSE Core Courses</b>	7
Data Structures and Algorithms	7
Discrete Mathematics	14
Computer Organization & Architecture	19
Advanced Programming	25
Operating Systems	32
Design and Analysis of Algorithms	39
Database Systems	44
Computer Networks	48
Machine Learning	55
Cyber Security	62
<b>Core Courses (Additional)</b>	67
Compiler Construction	67
Theory of Computation	72
<b>Micro Specializations (and Professional Electives)</b>	78
Software Engineering Micro Specialization	79
Machine Learning Micro Specialization	81
Distributed and Cloud Systems Micro Specialization	83
Human-Computer Interaction (HCI) Micro Specialization	85

---

<b>Appendix: Recommendations for Using Online Content in Courses at Colleges and Universities</b>	87
Background	87
Issues	87
Modes of Learning	87
Evaluation	90
Other Issues	90
Online offering of AICTE Model Curriculum Courses by Experts	91
<b>Appendix: Recommendations for Possible Exits for a BTech CSE</b>	93
<b>Committee and Area Experts</b>	94

---

# Framework for CSE Curriculum Design

## Context and Background

Computer Science has changed considerably in the last few years with areas like Machine Learning and Cloud computing becoming much more important. At the same time, the technologies and underlying computing systems have also evolved, considerably easing executing some of the tasks that earlier took much more training and experience. These changes require that teaching of computer science ought to suitably adapt – to reflect the changed nature of the discipline, as well as to update courses with the more recent technology platforms. And this ought to be done while providing some flexibility to the HEIs to address their specific constraints and focus. With these in mind, the committee for model curriculum design for CSE established some of the guiding principles for the exercise:

- The focus of curriculum design is the 4 yr BE/BTech program in CSE and the target audience of the curriculum is the vast majority of universities/institutions, rather than the premier institutes (who design their own curriculum and have the expertise for it.)
- The existing CSE curriculum of AICTE will be the starting point. Suitable enhancements/modifications will be made.
- **Flexibility.** Even within the large non-top tier education system, which is the primary target, some have better infrastructure and capability, so it is desirable to provide a limited degree of flexibility to the HEIs on the curriculum.

To provide this flexibility, for each course, the learning outcomes will be grouped in two – essential, and desired/advanced. The core courses will also be identified as essential and desired/advanced.

This simple framework of grouping outcomes and courses as essential and desired provides limited flexibility to institutions to design their curriculum depending on their capabilities, resources, goals, etc, while still providing guidance for a sound curriculum. The flexibility can be leveraged by institutions to provide different pathways to students, and multiple exits.

- **Early Exits.** The new education policy (NEP) aims to provide multiple exit points to students. To support more than one exits for students, it is important to develop employability skills early, and not take the approach of first focusing on foundations and then develop practical skills. Also, we believe that all but a few students will go for full degree, hence while providing flexibility for early exit, the outcomes of the 4 yr degree should not be compromised. We propose:
  - Eliminate the separation of theory and labs, instead break courses topic or theme-wise. This will allow both basic theory and basic skills to be taught together and build skills early.
  - Introduce discipline courses early, so disciplinary skills can be developed early. For this have one sem common program, and allow some discipline courses in 2<sup>nd</sup> sem – this will also help students in getting internships / part-time jobs in summer. This will require

---

branch change to be done after one semester. (For lateral entry, some make-up courses may be needed.)

With this, and the flexibility provided in the curriculum, an HEI can design suitable exits. We explicitly provide a recommendation for possible exits for a Btech CSE student.

- **Provide guidance for labs/assignments.** A small analysis the committee did on usage of AICTE curriculum indicated that while the courses and topics specified in the curriculum are widely used, HEIs differ considerably in types of assignments/labs/projects they give. Also, this is where there is a perceived deficiency in education – while theory is covered suitable assignments/labs are not give for students to develop skills. To help in this, it was decided that for each course design, experts will also recommend nature of labs and assignments for each module in the course.
- **Pedagogy suggestions.** With the emergence of a host of online resources that can be used to support teaching, there may be different ways to teach a topic. To help teachers in pedagogy, for each module in a course, pedagogy suggestions have also been provided.
- **Multiple pathways.** For supporting multiple pathways within the academic program, we propose to provide for micro specializations through thematic course streams. These can be further enhanced by HEIs with programs like honors for advanced students with more credits or advanced learning outcomes, etc.

## Graduate Attributes

Curriculum of a program is finally a network of credit units – courses (core, disciplinary core, disciplinary elective, open), internships, practice, projects, etc. which help achieve program goals. Program goals can be stated as attributes the students should possess on graduation, i.e. statements about the learning, values, capabilities etc. of graduates. These are called Graduate Attributes (GAs). A program typically has:

- **General GAs:** which are often common across many similar programs (e.g. BTechs) and focus on generalized skills and capabilities in the graduate.
- **Discipline GAs:** are discipline specific attributes, which focus on understanding of different concepts and systems related to the discipline, and on competencies and skills in that discipline.

Together the GAs define the goals of the program. The aim of curriculum design is to evolve a curriculum that can develop in students the stated graduate attributes. While specifying the GAs and designing a curriculum for it, we must keep a basic constraint in mind: a full BTech program has 8 semesters, each with about 5 full courses. GAs should specify only what can be taught and absorbed in this timebox (i.e. we cannot push more simply by adding more.)

Desired Graduate Attributes for the CSE program are given below. The CSE curriculum design will focus more on delivering the discipline GAs, while strengthening the general GAs, where possible. Feedback on these Graduate Attributed was taken from many representatives from industry, as well as from HEIs. GAs should be read by adding this at the start of each: At graduation time, a student should have:

General Graduate Attributes	Discipline Graduate Attributes
G1 Ability to identify a problem, analyze using design thinking techniques, and evolve innovative approaches for solving it.	CS1 Proficiency in writing in at least two dissimilar programming languages programs of modest complexity which are: readable, tested for correctness, efficient, and secure.
G2 Ability to apply mathematical concepts and techniques in problem solving.	CS2 Ability to design and apply appropriate algorithms and data structures for evolving efficient computing based solutions for new problems.
G3 Ability to function effectively in multi-cultural teams to accomplish a common goal.	CS3 Understanding of computing systems at computer architecture, operating systems, and distributed- computing levels, and how they affect the performance of software applications.
G4 Ability to communicate effectively with a wide range of audience.	CS4 Understanding of theoretical foundations, fundamental principles, and limits of computing.
G5 Ability to self-learn and engage in life-long learning and upgrade technical skills	CS5 Ability to analyse large volumes of data employing a variety of techniques for learning, better prediction, decision making, etc
G6 An understanding of professional and ethical responsibility	<b>ADVANCED/OPTIONAL GAs</b> CS6 Ability to design, implement, and evaluate computer based system or application to meet the desired needs using modern tools and methodologies
G7 Ability to undertake small research tasks and projects.	CS7 Ability to develop full stack applications using one commonly used tech-stack and modern tool.
G8 An entrepreneurial mindset for opportunities using technology and innovations.	CS8 Understanding of and ability to use advanced techniques and tools in a few different domain areas (e.g. parallel processing, image processing, IR, ...)
G9 An understanding of impact of solutions on economic, societal, and environment context.	CS9 Exposure to emerging technologies such as Cloud Computing, IoT, etc
G10 Strong emotional intelligence, human and cultural values	

## CSE Core Courses

For the BTech CSE program, students need to do as part of their common 1<sup>st</sup> year curriculum a course on programming (“Introduction to Programming”) and two Math courses. The AICTE 1<sup>st</sup> year curriculum provides for these.

The list of core courses recommended for CSE are given below. As explained above, the core has been split into two – essential and additional. The essential core is what all programs must have. The additional core list those courses which HEIs can choose to make core or an elective, depending on their education approach and resources.

<b>Suggested Year</b>	<b>Course Name</b>	<b>GAs it Contributes to (GAs it supports)</b>
	<b>Professional Core Courses (Essential)</b>	
2 <sup>nd</sup> year	Data structures and Algorithms (DS)	CS1, CS2, CS4
2 <sup>nd</sup> year	Discrete Mathematics (DM)	CS4 (CS2, CS4, CS6)
2 <sup>nd</sup> year	Computer Organization and Architecture (COA)	CS3 (CS4, CS6)
2 <sup>nd</sup> year	Advanced Programming (AP)	CS1, CS3, CS7, CS8 (CS2)
3 <sup>rd</sup> year	Operating Systems (OS)	CS3 (CS6, CS8)
3 <sup>rd</sup> year	Design and Analysis of Algorithms (Algo)	CS1, CS2, CS4
3 <sup>rd</sup> year	Database Systems (DB)	CS3, CS5, CS6 (CS1, CS7, CS9)
3 <sup>rd</sup> /4 <sup>th</sup> year	Computer Networks (NW)	CS3 (CS4, CS6, CS8)
3 <sup>rd</sup> year	Machine Learning (New) (ML)	
3 <sup>rd</sup> year	Security (New)	CS1, CS3, CS6
	<b>Extended Professional Core (Additional)</b>	
3 <sup>rd</sup> /4 <sup>th</sup> year	Compiler Design	
3 <sup>rd</sup> /4 <sup>th</sup> year	Theory of Computation	CS4 (CS2)

Detailed syllabus of each of these courses follows – a separate subsection for each course. A common template has been used to specify the course design. Each course design lists essential learning outcomes for the course, and desired/advanced learning outcomes. As the terms suggest, all HEIs should ensure that the essential learning outcomes are achieved. The desired/advanced learning outcomes are those which those HEIs which have the capabilities and resources to deliver may include in their courses.

## Coverage of CS Graduate Attributes by the Core Courses

This table summarizes which CS GAs different courses contribute to. For each GA, list of main courses is given – these are the courses which directly contribute to the GA. Also mentioned are the courses which support the GA (though perhaps less directly):

CS1 Proficiency in writing in at least two dissimilar programming languages programs of modest complexity which are: readable, tested for correctness, efficient, and secure	Main Courses: DS, AP, Algo, Security Supporting Courses: DB
CS2 Ability to design and apply appropriate algorithms and data structures for evolving efficient computing based solutions for new problems	Main: DS, Algo Supporting: DM, AP, ToC
CS3 Understanding of computing systems at computer architecture, operating systems, and distributed-computing levels, and how they affect the performance of software application	Main: CO/CA, AP, OS, DB, NW, Security Supporting:
CS4 Understanding of theoretical foundations, fundamental principles, and limits of computing	Main: DS, DM, Algo, ToC Supporting: NW, CO/CA
CS5 Ability to analyse large volumes of data employing a variety of techniques for learning, better prediction, decision making,	Main: DB Supporting: DM
<b>ADVANCED/OPTIONAL</b>	Main: AP, DB, Security
CS6 Ability to design, implement, and evaluate computer based system or application to meet the desired needs using modern tools and methodologies	Supporting: DM, CO/CA
CS7 Ability to develop full stack applications using one commonly used tech-stack and modern tools	Main: AP Supporting: OS, NW, Security
CS8 Understanding of and ability to use advanced techniques and tools in a few domain areas (e.g. parallel processing, image processing, IR, ...)	Main: Supporting: DB
CS9 Exposure to emerging technologies such as Cloud Computing, IoT, etc	Main: Supporting: DB, Security, NW

## Micro Specializations (and Professional Electives)

Besides the core courses, programs normally have professional elective courses, which HEIs decide. It is possible to use the electives to provide a limited specialization in some sub-area of CSE to a BTech student. We call these as micro-specializations. These also allow multiple pathways to students, as different students can graduate with different specializations (or not). A micro



---

specialization is a set of a few full or half courses, which build upon the core CSE program. The report gives possible structure of a few micro specializations in:

- Software Engineering
- Machine Learning
- Distributed and Cloud Systems
- Human Computer Interaction (HCI)

More specializations definition may be added (e.g. in Security, High Performance Computing, Algorithms). Micro specializations are optional for HEIs – i.e. they can decide to offer them or not, and if they do, which ones they want to. If they decide not to have these, the list of courses mentioned in the micro specializations can be used as a list of suggested professional electives.

### **Recommendations for Online Credits**

A sub-committee was formed to recommend how online credits may be given by HEIs. The recommendations are given as part of the report as an appendix. The recommendations were accepted by the full steering committee.

### **Recommendations for Multiple Exits**

A sub-committee was formed to recommend possible exits for a BTech CSE student – in line with one of the goals of the NEP. The recommendations are given as part of the report as an appendix. The recommendations were accepted by the full steering committee.

# Syllabus for CSE Core Courses

This section gives the detailed syllabus of all the core courses for the model CSE program. Each course design states the learning outcomes – both essential and desirable – and then the modules in the course and topics in each of the modules. The design also suggests for each module a suitable pedagogy approach and suitable assignments/labs.

## Data Structures and Algorithms

<b>CS301</b>	<b>Data Structure &amp; Algorithms</b>	<b>3L:0T: 4P</b>	<b>5 credits</b>	<b>Pre-Reqs: ESC201</b>
--------------	--	------------------	------------------	-------------------------

**Learning outcomes of the course (i.e. statements on students’ understanding and skills at the end of the course the student shall have):**

### Essential (<=6):

1. Understanding abstract specification of data-structures and their implementation.
2. Understanding time and space complexity of programs and data-structures.
3. Knowledge of basic data-structures, their applications and relative merits.
4. Ability to convert an algorithmic solution to a program using suitable data-structures and analyze the trade-offs involved in terms of time and space complexity.

### Desirable/Advanced (<= 3):

1. Amortized complexity
2. Use of randomization in data-structures

### Detailed contents for Essential Learning Outcomes:

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 1:</b> Introduction and basic terminology (1 week)	(i) Notion of data-structures and algorithms. (ii) $\log n, n, 2^n$ : understanding growth of these functions, and applications (binary search and extensions to similar problems) (ii) Worst-case, average-case time/space	T1: Chapter 3 (i) Explain the interplay between algorithms and data-structures (ii) Explain the meaning of worst and average case (iii) Examples of $O(\ )$ , Motivation behind asymptotic analysis	(i) Worst/average case analysis for small pseudo-codes (ii) Prove/disprove why a function $f(n)$ is $O(g(n))$ (and similarly for $\Omega$ ). (iii) Variations on binary search with applications, recursive and

	<p>complexity and their relative merits.</p> <p>(iii) Asymptotic Notation: <math>O( )</math>, <math>\Omega( )</math></p>	<p>(large <math>n</math> and ignoring constants).</p> <p>(iv) Discuss recurrence relation for binary search.</p>	<p>iterative implementation of binary search with applications to problems.</p>
<p><b>Module 2:</b></p> <p>Abstract Data-types, Arrays, Linked Lists, Stacks, Queues Dictionary ADT, Trees, Binary Trees</p> <p>(2.5 weeks)</p>	<p>(i) Abstract data-type (ADTs): arrays and linked list ADTs.</p> <p>(ii) Stacks, Queues: ADTs and implementations using arrays, linked lists.</p> <p>(iii) Doubly linked lists: ADT and implementation</p> <p>(iv) Dictionary ADT: implementation using array, linked lists, binary search.</p> <p>(v) Tree ADT and examples</p> <p>(vi) Implementation of trees and basic traversal algorithms</p> <p>(vii) Binary trees and inorder traversal</p>	<p>T1: chapter 4, Chapter 6</p> <p>(i) Explain the difference between specification and implementation of ADTs.</p> <p>(ii) Discuss stack implementation using arrays of fixed size and linked list, and relative merits.</p> <p>(iii) Circular array and linked list implementation of queues, and relative merits.</p> <p>(iv) Discuss some applications of stacks: post-fix notation, matching parentheses etc.</p> <p>(v) Dictionary ADT: discuss why neither arrays nor linked list is a good implementation, but sorted arrays are good for searching.</p> <p>(vi) Discuss situations where we need to store hierarchical data.</p> <p>(vii) Discuss pre-order and post-order ,</p>	<p>(i) Implementation of stacks with application to a problem.</p> <p>(ii) Implementation of queues with application to a problem.</p> <p>(i) Implementation of trees with applications for storing and accessing hierarchical data.</p>

		traversals and applications in finding height and similar problems.	
<p><b>Module 3:</b> Priority Queues and Heaps</p> <p>(1 week)</p>	<p>(i) Priority Queue ADT</p> <p>(ii) Definition of heaps</p> <p>(iii) Implementation of Priority Queues using heaps and running time analysis</p> <p>(iv) Implementation of heaps using arrays.</p> <p>(v) Heap-sort</p>	<p>T1: Chapter 7.1-7.3</p> <p>(i) Start with FindMin and Insert, and a simple <math>O(1)</math> time and <math>O(1)</math> space algorithm. Buildup towards DeleteMin.</p> <p>(ii) Discuss algorithm for heaps: inserting an element, modifying an element, and deleting the minimum element.</p> <p>(iii) Discuss why array implementation is more efficient than balanced binary trees</p>	<p>(i) Array implementation of heaps and application to a problem.</p> <p>(ii) k-ary heaps: compare with binary heaps (both in theory and practice)</p>
<p><b>Module 4:</b> Binary Search Trees, AVL Trees, 2-4 trees</p> <p>(3 weeks)</p>	<p>(i) Binary Search Trees: definition and some basic algorithms.</p> <p>(ii) Implementation of Dictionary ADTs using Binary Search trees and running time analysis</p> <p>(iii) AVL trees: height balance condition, rotations, and implementation of dictionary ADT</p> <p>(iv) 2-4 Trees: Multi-way search trees, implementation of dictionary ADT, Informal discussion of extension to B-trees.</p>	<p>T1: Chapter 9.1, 9.2, 9.4</p> <p>(i) Discuss algorithms for finding predecessor or successor, and similar problems in Binary Search Trees.</p> <p>(ii) Explain why height of a Binary Search Tree may not remain <math>O(\log n)</math>.</p> <p>(v) Explain how the height balance condition ensures that the height is <math>O(\log n)</math>, and how rotation changes the structure of a tree.</p> <p>(vi) Explain why rotations in AVL trees</p>	<p>(i) Implementation of AVL trees with search, insert, delete operations and application to a problem. Comparison with unbalanced Binary Search Trees.</p> <p>(ii) Implementation of 2-4 trees with search, insert, delete operations and application to a problem.</p> <p>(iii) Comparison of the two implementations above.</p>

		<p>restore height balance condition</p> <p>(vii) Explain why 2-4 trees have <math>O(\log n)</math> height and the running time of insert/delete operations.</p> <p>(viii) Discuss how balanced binary tree data-structures can implement a priority queue</p>	
<p><b>Module 5:</b> Hash tables, tries (2 weeks)</p>	<p>(i) Map ADT</p> <p>(ii) Hash Tables and implementation of Map using Hash Tables</p> <p>(iii) Design of hash functions</p> <p>(iv) Collision resolution schemes: chaining, open addressing schemes like linear probing, quadratic probing, double hashing.</p> <p>(v) Applications of Hashing: finding duplicates, set intersection, etc.</p> <p>(vi) Tries: implementation of Map ADT using tries.</p> <p>(vii) Compressed tries and suffix tries.</p>	<p>T1: Chapter 8.1-8.3, Chapter 11.3</p> <p>(i) Explain the difference between Map and Dictionary ADT.</p> <p>(ii) Discuss how hash functions can have non-numeric keys as input.</p> <p>(iii) Discuss the relative merits of hash tables and balanced binary search trees.</p> <p>(iv) Discuss how hashing can be a substitute for sorting in many cases.</p> <p>(v) Explain why tries can be better than balanced binary search trees in some settings.</p> <p>(vi) Explain how compressed tries save space</p> <p>(vii) Discuss real-life applications of tries.</p>	<p>(i) Implementation of hash tables with applications to a problem.</p> <p>(ii) Implementation of tries and applications to a problem.</p>

<p><b>Module 6:</b> Sorting, Selection (1.5 weeks)</p>	<p>(i) Bubble sort, insertion sort, selection sort.</p> <p>(ii) Merge sort and divide and conquer paradigm</p> <p>(iii) Quick sort: average and worst case analysis, randomized quicksort (intuitive explanation)</p> <p>(iv) Selection based on partitioning ideas used in QuickSort.</p>	<p>T1: Chapter 10.1, 10.2, 10.4, 10.5, 10.7</p> <p>(i) Discuss why <math>O(n^2)</math> time algorithms can be useful sometimes (small data, data nearly sorted etc.)</p> <p>(ii) Only the recurrence for merge sort and mention that divide-and-conquer paradigm will be explored more in algorithms course.</p> <p>(ii) Discuss the randomized splitting algorithm for quicksort and selection and explain intuitively the expected running time.</p>	<p>Implementation of sorting algorithms and comparison of running times on large data-sets.</p>
<p><b>Module 7:</b> Graphs, representations and traversal algorithms, applications of BFS, DFS (2.5 weeks)</p>	<p>(i) Graph ADTs and applications</p> <p>(ii) Adjacency list and adjacency matrix representations and relative merits</p> <p>(iii) Basic graph definitions: paths, cycles, trees, spanning trees, connected components, Euler's formula.</p> <p>(iv) Depth First Search Traversal algorithm for directed graphs: classification of edges into forward, back and cross edges. Applications to cycle finding, topological sort in directed acyclic graphs, finding connected components. Running time analysis.</p>	<p>T1: Chapter 12.1-12.4</p> <p>(i) Discuss the wide applicability of graphs including social networks, internet.</p> <p>(ii) Discuss time and space complexity of basic operations using adjacency list and adjacency matrix.</p> <p>(iii) Discuss why trees have <math>n - 1</math> edges.</p> <p>(iv) Discuss how DFS can be thought of as exploration with backtracking. Explain the role of stack in DFS.</p> <p>(v) Explain how BFS can be thought of as traversal along shortest paths and implementation using queues.</p>	<p>(i) Graph implementation using adjacency list and DFS/BFS traversal with applications.</p>

	(v) Breadth first search algorithm: implementation using queues, shortest path tree property. Running time analysis	(vi) Formal proof of why BFS yields a shortest path tree.	
--	---	---	--

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 1:</b> Amortized Complexity (1 week)	(i) Binary counter (ii) Binomial Heaps (iii) Extendible arrays	T1: Chapter 5.1.3, R2: Chapter 17.1  (i) Explain the motivation behind amortized analysis.  (ii) Analyze amortized complexity by explicit calculation of total number of operations after $n$ steps.	(i) Exercises on amortized time complexity (e.g., a queue using two stacks etc.)  (ii) Implementation of extendible arrays.
<b>Module 2:</b> Randomization in data-structures (1 week)	(i) Skip Lists (ii) Randomized quick-sort	T1: Chapter 8.4, 10.2  (i) Explain the idea of expectation of a random variable.  (ii) How does expected running time translate to actual running time  (iii) Recurrence for expected running time randomized quicksort	(i) Exercises on calculation of expectation of a random variable.  (ii) Implementation of skip lists and comparison with AVL trees.

**Suggested text books / Online lectures or tutorials:**

- 
1. “Data Structures and Algorithms in Java”, by Michael T. Goodrich and Roberto Tamassia, John Wiley & Sons; 3rd Edition.
  2. “Data Structures and Algorithms in Python”, by Michael T. Goodrich and Robert Tamassia, Wiley, 1st Edition.
  3. In case any other programming language is used for this course, some other suitable text book may be chosen.

**Suggested reference books / Online resources:**

1. NPTEL video series, Data-structures and Algorithms, Instructor: Naveen Garg.
2. Introduction to Algorithms, 4TH Edition, Thomas H Cormen, Charles E Lieserson, Ronald L Rivest and Clifford Stein, MIT Press/McGraw-Hill.

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly: CS1, CS2, CS4***

***Other discipline GAs to which this course may contribute somewhat:***



## Discrete Mathematics

CS-X	Discrete Mathematics	3L:1T: 0P	4 credits	Pre-Reqs: 0
------	----------------------	-----------	-----------	-------------

### Prerequisites

Familiarity with some mathematical notation, ideas and concepts covered at the pre-college levels.

**remark:**

**Learning Outcomes of the course (i.e. statements on students' understanding and skills at the end of the course the student shall have):**

### Essential (<=6):

1. Understand examples in Computer Science through mathematical terminology and notation.
2. Construct direct, and indirect, proofs of basic theorems.
3. Understand the differences between a mathematical proof, a heuristic, and a conjecture.
4. Learn how to divide a problem, or a proof, into smaller cases.
5. Formulate mathematical claims and be able to construct counterexamples.
6. Apply the knowledge of mathematics to solve real-world problems.

### Desirable/Advanced (<= 3):

1. Identify formal algebraic structures in computer science.
2. Work with probability & statistics in a rigorous ways.
3. Use this course topics to design, and rigorously analyze, real-life algorithms.

### Detailed contents for Essential Learning Outcomes:

Module	Topics	Pedagogy / teaching suggestions	Nature of lab / assignment / practice
<b>Module 1:</b> Set, Relations, Functions. (~1.5 wks)	Operations and Laws of Sets, Cartesian Products, Binary Relation, Partial Ordering Relation, Equivalence Relation, Image of a Set, Sum and Product of Functions, Bijective functions, Inverse and Composite Function, Size of a Set, Finite and infinite Sets, Countable and uncountable Sets, Cantor's	- Chapter 1,3 of T1. - Chapter 2 of T2 - Discuss some examples from computer science: +Uncountability of Reals +Uncomputability by the diagonal argument (ref. Chap.6 of T1).	Make assignments using the books. To test: (1) what was done in the class (2) whether the student can think and apply the concepts.

	diagonal argument and The Power Set theorem.	+Relational database (ref. Chap.3 of T1. Refer R8).	
<b>Module 2:</b> Proof strategies. (1.5 wks)	Proof Methods and Strategies: Forward Proof, Proof by Contradiction, Proof by Contraposition, Proof of Necessity and Sufficiency, Case analysis, Induction.	- Chapter 1 of T1. - Chapter 1 of T2.  - Discuss some examples from computer science:  +Examples like above.  + $\sqrt{2}$ is irrational.  +Prime numbers are infinite (Chap.4.3 of T2).	Make assignments using the books. To test:  (1) what was done in the class  (2) whether the student can think and apply the concepts.
<b>Module 3:</b> Modular Arithmetic. (1.5 wks)	Extended Euclid's Greatest Common Divisor algorithm, The Fundamental Theorem of Arithmetic, Modular arithmetic, Coprimality (or Euler's totient function), Chinese Remainder Theorem.	- Chapter 4 of T2. Chapter 1.20 of T1.  - References T5 & R5.  - Discuss some examples from computer science:  +Diffie-Hellman key-exchange.  +RSA Cryptosystem.	Make assignments using the books. To test:  (1) what was done in the class  (2) whether the student can think and apply the concepts.
<b>Module 4:</b> Combinatorics. (1.5 wks)	Permutation & Combination, Inclusion-Exclusion, Pigeon-hole principle, Generating functions, Recurrence.	- Chapter 2, 8, 9 of T1. Chapter 6 of T2. Chapter 2, 3, 4, 5, 6 of T4. Also refer to R3, R4.  - Discuss some examples from computer science:  +Count binary trees.  +Count matched parentheses.	Make assignments using the books. To test:  (1) what was done in the class  (2) whether the student can think and apply the concepts.

		+Count matrix chain multiplication.	
<p><b>Module 5:</b> Graphs.</p> <p>(1.5 wks)</p>	<p>Connected components, Paths, Cycles, Trees, Hamiltonian/ Eulerian Walks, Coloring, Planarity, Matching.</p>	<p>- Chapter 4, 5 of T1. Chapter 10, 11 of T2.</p> <p>- Discuss some examples from computer science:</p> <p>+Shortest-distance (Chapter 7 of T1).</p> <p>+Minimum spanning tree.</p> <p>+Prefix codes.</p> <p>+Graph isomorphism problem.</p>	<p>Make assignments using the books. To test:</p> <p>(1) what was done in the class</p> <p>(2) whether the student can think and apply the concepts.</p>
<p><b>Module 6:</b> Logic.</p> <p>(1.5 wks)</p>	<p>Languages of Propositional logic and First-order logic, expressing natural language sentences in languages of propositional and first-order logic, expressing natural language predicates in the language of first-order logic. Semantics of First-order logic: interpretation and its use in evaluating a formula.</p> <p>Optional advanced topics if there is extra time:</p> <p>Semantic entailment, Validity and Satisfiability. What is a proof system? Eg. natural deduction or analytical tableau. Notions of Consistency and Completeness of a logic.</p>	<p>- Chapter 1, 11 of T1. Chapter 1, 12 of T2. Also refer to R2.</p> <p>- Discuss some examples from computer science (also, refer T3):</p> <p>+SAT solvers and why these are so useful.</p> <p>+Relational Calculus.</p> <p>+Language for specification like Z.</p> <p>+Hoare logic for program verification.</p>	<p>Make assignments using the books. To test:</p> <p>(1) what was done in the class</p> <p>(2) whether the student can think and apply the concepts.</p>

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 7:</b> Algebra.  (1.5 wks)	Group, Permutation Groups, Cosets, Normal Subgroups, Ring, Field, Finite fields, Fermat's little theorem, Homomorphisms, Isomorphisms.	- Chapter 10 of T1. Also, refer R5, R6.  - Discuss some examples from computer science (ref. T3):  +RSA Cryptosystem.  +Reed-Solomon error-correction encoding.	Make assignments using the books. To test:  (1) what was done in the class  (2) whether the student can think and apply the concepts.
<b>Module 8:</b> Discrete probability.  (1.5 wks)	Topics to be taught from the viewpoint of CS instead of Maths:  Discrete Sample Space, Probability Distribution, Random variables, Expectation, Variance, Bernoulli trials, Conditional probability & independence (Bayes' Theorem).	- Chapter 2 of T1. Chapter 7 of T2.  - Discuss some examples from computer science (also refer R7):  +Randomized algorithms.  +Heuristics.	Make assignments using the books. To test:  (1) what was done in the class  (2) whether the student can think and apply the concepts.

**Suggested text books / Online lectures or tutorials:**

- T1. Liu, C. L., & Mohapatra, D. P. (2008). Elements of Discrete Mathematics. Tata McGraw-Hill.
- T2. Rosen, K. H. (2019). Discrete Mathematics and Its Applications. (8th Edition) ISBN10: 125967651X ISBN13: 9781259676512.
- T3. Huth, M., & Ryan, M. (2004). Logic in Computer Science: Modelling and Reasoning about Systems (2nd ed.). Cambridge University Press.
- T4. Cohen, D. I. A. (1978). Basic techniques of combinatorial theory. John Wiley.

- 
- T5. Niven, I., Zuckerman, H. S., & Montgomery, H. L. (1991). An introduction to the theory of numbers. John Wiley & Sons.

**Suggested reference books / Online resources:**

- R1. Norman L. Biggs, Discrete Mathematics, (2nd ed. 2002), Oxford University Press.
- R2. Smullyan, R. M. (1995). First-order logic. Courier Corporation.
- R3. Bóna, M. (2006). A walk through combinatorics: an introduction to enumeration and graph theory.
- R4. Cameron, P. J. (1994). Combinatorics: topics, techniques, algorithms. Cambridge University Press.
- R5. Shoup, V. (2009). A computational introduction to number theory and algebra. Cambridge University Press.
- R6. Herstein, I. N. (2006). Topics in algebra. John Wiley & Sons.
- R7. Mitzenmacher, M., & Upfal, E. (2017). Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis. Cambridge University Press.
- R8. C. J. Date (2019). Database Design and Relational Theory. Normal Forms and All That Jazz.

**Prepared by:**

1. Somenath Biswas, Computer Science & Engineering Department, IIT Goa (Retd. IIT Kanpur)
2. Partha Mukhopadhyay, Theoretical Computer Science Group, Chennai Mathematical Institute
3. Nitin Saxena, Computer Science & Engineering Department, IIT Kanpur
4. Bhabani P. Sinha, Advanced Computing & Microelectronics Unit, Indian Statistical Institute Kolkata

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly:*** CS4

***Other discipline GAs to which this course may contribute somewhat:*** CS2, CS5, CS6

## Computer Organization & Architecture

CS402	Computer Organization & Architecture	3L:0T: 4P	5 credits	Pre-Reqs: ESC302
-------	--------------------------------------	-----------	-----------	------------------

**Learning Outcomes of the course (i.e. statements on students' understanding and skills at the end of the course the student shall have):**

**Essential (<=6):**

1. The key components of a basic computer.
2. The key components of a CPU and how the instructions are executed.
3. Execution and time taken by instructions in a pipelined processor.
4. The need for memory hierarchy and efficiency achieved due to the use of cache.
5. How the data is stored and input-output is performed in computers.

**Desirable/Advanced (<= 3):**

1. Super-scalar and multi-core architectures for crossing one clock per instruction barrier
2. Cache data coherence related challenges in multi-core processors

**Detailed contents for Essential Learning Outcomes:**

Module (appx dur in wks)	Topics	Pedagogy / teaching suggestions	Nature of lab / assignment / practice
<b>Module 1:</b> Introduction (Lectures 6) (Weeks 2)	Role of abstraction, basic functional units of a computer, Von-Neumann model of computation, A note on Moore's law, Notion of IPC, and performance. Data representation and basic operations.	T1: Chapter 1 R1: Chapter 1(Section 1.7.2) R3: Chapter 1 R2: Chapter 3	
<b>Module 2:</b> Instruction Set Architecture (RISC-V) (Lectures 8/9) (Weeks 3)	CPU registers, instruction format and encoding, addressing modes, instruction set, instruction types, instruction decoding and execution, basic	R1: Chapter 3, 5 R3: Chapter 2, Chapter 5 (Section 5.1 to 5.5) R7: Chapter 3	Lab Modules 1 and 4

	instruction cycle, Reduced Instruction Set Computer (RISC), Complex Instruction Set Computer (CISC), RISC-V instructions; X86 Instruction set.		
<b>Module 3:</b> The Processor (Lectures 6) (Weeks 2)	Revisiting clocking methodology, Amdahl's law, Building a data path and control, single cycle processor, multi-cycle processor, instruction pipelining, Notion of ILP, data and control hazards and their mitigations.	T1: Chapter 1, 4 R3: Chapter 6	Lab Modules 2 and 4
<b>Module 4:</b> Memory hierarchy (Lectures 8/9) (Weeks 3)	SRAM/DRAM, locality of reference, Caching: different indexing mechanisms, Trade-offs related to block size, associativity, and cache size, Processor-cache interactions for a read/write request, basic optimizations like write-through/write-back caches, Average memory access time, Cache replacement policies (LRU), Memory interleaving.	T1: Chapter 5 R3: Chapter 8	Lab Module 2
<b>Module 5:</b> Storage and I/O (Lectures 6) (Weeks 2)	Introduction to magnetic disks (notion of tracks, sectors), flash memory. I/O mapped, and memory mapped I/O. I/O data transfer techniques: programmed I/O,	R5: Chapter 5 (Section 5.4) Chapter 6 (Section 6.1) Chapter 7 (Section 7.1-7.5)	Lab Modules 3 and 5

	Interrupt-driven I/O, and DMA.		
--	--------------------------------	--	--

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 6:</b> Superscalar processors and multicore systems (Lectures 6) (Weeks 2)	Limits of ILP, SMT processors, Introduction to multicore systems and cache coherence issues	Lab Modules 1 and 4	

**Laboratory Modules**

The laboratory component consists of 5 modules out of which three can be chosen based on the overall curriculum and its emphasis. Modules 4 and 5 assume that students have a background in using FPGA kits in their Digital Electronics course (ESC302) and have also been introduced in programming in one of the HDLs (VHD or Verilog).

**Possible three options are:**

1. Module 1 + Module 2 + Module 3  
 Instructions & assembly language + basic performance + advanced performance analysis
2. Architecture + Module 1 + Module 2 + Module 4  
 Instructions & assembly language + basic performance + basic processor design  
 Module 1 + Module 4 + Module 5
3. Instructions & assembly language + basic processor design + I/O and architecture enhancements.

**Detailed contents for Essential Learning Outcomes:**

<b>Module (appx dur in wks)</b>	<b>Topics</b>	<b>Comments</b>
<b>Module 1:</b> (Weeks 4)	Write programs in ARM/RISC V assembly language and test these on an instruction set simulator. Typical examples are given below. Some of these are dependent on I/O facilities provided by the simulator.	Essential component



<p>Objective: Understanding architecture and instructions through assembly programming</p>	<ul style="list-style-type: none"> <li>• Generate some interesting numbers (example - Happy numbers, Autonomic numbers, Hardy-Ramanujan numbers etc.)</li> <li>• Implement a 4-function calculator</li> <li>• Sort an integer array using merge sort (recursive)</li> <li>• Evaluate an arithmetic expression specified as a string (using recursive functions)</li> <li>• Implement a simple game</li> </ul>	
<p><b>Module 2:</b></p> <p>Understanding performance issues related to pipelining and cache using architectural simulator</p> <p>(Weeks 4/5 )</p>	<p>Usage of a instruction pipeline visualization tool like RIPES</p> <p>Write or generate sequence of instructions and observe the overall pipeline stalls with and without data hazards, control hazards, and with/without data forwarding.</p> <p>Rearrange the sequence of instructions or the program so that the pipeline stalls will be minimized.</p>	<p>Optional</p>
<p><b>Module 3:</b></p> <p>Understanding memory access patterns and changing basic cache memory parameters to analyze the impact of standard programs or benchmarks using architectural simulators.</p> <p>(Weeks 3)</p>	<p>Configure the simulator [gem5 is preferred] to operate on the binaries of the benchmark as the input.</p> <p>Run the program and examine the IPC, cache hit rate, number of conflict misses and block replacements.</p> <p>Vary the cache size, block size, and associativity and analyze the metrics and reason the changes observed.</p> <p>Modify the block replacement algorithms and see the impact at cache memory performance</p> <p>Calculate the access time, power and are associated with a given cache configuration.</p>	<p>Optional:</p> <p>Familiarity with tools like GEM5, CACTI and PIN, and access to benchmarks like SPEC, PARSEC, SPLASH.</p> <p>Any other tools/simulators that can support memory pattern analysis is also fine.</p>

	Vary the cache size, block size, and associativity and analyze the metrics and reason the changes observed.	
<p><b>Module 4:</b> (Weeks 4/5)</p> <p>Objective: Understanding computer architecture by designing a CPU on FPGA kit</p>	<p>Design a simple ARM/RISC V processor for a small subset of instructions and implement on FPGA board. This is done in stages as follows.</p> <ul style="list-style-type: none"> <li>• Design CPU for the instruction subset {add, sub, cmp, mov, ldr, str, beq, bne, b}, with each instruction executing in a single cycle. No sequential control required. For each instruction, only a limited set of variants are considered.</li> <li>• Use memory generator to add program and data memories.</li> <li>• Include circuit for single step execution and for displaying signals of interest.</li> <li>• Modify the design to allow multi-cycle execution of instructions.</li> <li>• Enhance the design to include all DP instructions and all variants of the second operand.</li> </ul>	<p>Optional</p> <p>Prerequisite:</p> <p>Use of FPGA kits as well as exposure to VHDL/Verilog</p>
<p><b>Module 5:</b></p> <p>Objective: Understanding computer architecture performance issues as well as I/O through architectural enhancements (on FPGA)</p> <p>(Weeks 4/5)</p>	<p>Extension of the CPU design and I/O programming</p> <ul style="list-style-type: none"> <li>• Enhance the design to include all variants of DT instructions.</li> <li>• Implement multiply group of instructions.</li> <li>• Enhance the design to implement "branch and link" instruction and include full predication.</li> <li>• Include limited exception handling.</li> <li>• Interface 7-segment display and 4x4 keypad.</li> <li>• Demonstrate execution of simple programs.</li> </ul>	<p>Optional</p> <p>Prerequisite:</p> <p>Use of FPGA kits as well as exposure to VHDL/Verilog</p>

---

### **Suggested text books / Online lectures or tutorials:**

T1 “Computer Organization and Design: The Hardware/Software Interface”, [David A. Patterson](#) and [John L. Hennessy](#), 5th Ed, Elsevier.

### **Suggested reference books / Online resources:**

R1 “Computer Organisation & Architecture”, [Smruti Ranjan Sarangi](#), McGraw Hill

R2 “Computer System Architecture”, Mano M. Morris, Pearson.

R3 “Computer Organization and Embedded Systems”, 6th Edition by Carl Hamacher, McGraw-Hill Higher Education

R4 “Computer Architecture and Organization”, 3rd Edition by John P. Hayes, WCB/McGraw-Hill

R5 “Computer Organization and Architecture: Designing for Performance”, 10th Edition by William Stallings, Pearson Education.

R6 “Computer System Design and Architecture”, 2nd Edition by Vincent P. Heuring and Harry F. Jordan, Pearson Education.

R7 <http://web.cecs.pdx.edu/~harry/riscv/RISCV-Summary.pdf>

### **Online simulators and tools:**

RIPES: <https://freesoft.dev/program/108505982>

GEM5: [https://www.gem5.org/documentation/learning\\_gem5/introduction/](https://www.gem5.org/documentation/learning_gem5/introduction/)

CACTI: <https://github.com/HewlettPackard/cacti>

PIN: <https://www.intel.com/content/www/us/en/developer/articles/tool/pin-a-binary-instrumentation-tool-downloads.html>

TEJAS: <https://www.cse.iitd.ac.in/~srsarangi/archbooksoft.html>

XILINX (VHDL/Verilog tools): <https://www.xilinx.com/support/university/students.html>

### **Prepared by:**

1. Prof. Anshul Kumar, IIT Delhi
2. Prof. Anupam Basu, IIT Kharagpur/NIT Durgapur
3. Prof. Indranil Sengupta, IIT Kharagpur
4. Prof. Biswabandan Panda, IIT Bombay
5. Prof. John Jose, IIT Guwahati
6. Prof. M. Balakrishnan, IIT Delhi

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly:*** CS3 (and also to some graduate attributes for ECE)

***Other discipline GAs to which this course may contribute somewhat:*** CS4, CS6

## Advanced Programming

CS-X	Advanced Programming	3L:1T: 0P	4 credits	Pre-Reqs: 0
------	----------------------	-----------	-----------	-------------

Familiarity with data structures, and introduction to programming.

**Learning Outcomes of the course (i.e., statements on students' understanding and skills at the end of the course the student shall have):**

**Essential (<=6):**

1. Understanding the build system: IDE, tools for testing, debugging, profiling, and source code management.
2. Students are able to demonstrate proficiency in object-oriented programming .
3. Identify and abstract the programming task involved for a given programming problem.
4. Learning and using language libraries for building large programs.
5. Ability to apply defensive programming techniques (e.g., assertions, exceptions).

**Desirable/Advanced:**

1. Ability to implement basic event-driven programming.
2. Understanding of the fundamentals of parallel programming.
3. Understanding of the basics of cloud computing.

**Detailed contents for Essential Learning Outcomes:**

Module (appx dur in wks)	Topics	Pedagogy/teaching suggestions	Nature of lab / assignment / practice
<b>Module 1:</b> Familiarity with the programming environment  (~1 wks)	Understanding the build system, IDE, debugging, profiling (Eclipse TPTP / gprof / VTune etc.), and source code management	R5;	Familiarity with terminal/command prompt, using git commands and github to pull/commit/push/merge code, writing, compiling and running simple programs, debugging by setting breakpoints

<p><b>Module 2:</b> Basic principles of the object-oriented development process</p> <p>(~1.5 wks)</p>	<p>Introduction to Object-Oriented Paradigm: Data encapsulation, modularity, code reuse, identifying classes, attributes, methods and objects, class relationships</p>	<p>Chapters 2-3, T1;</p> <p>Chapter 10, T2;</p> <p>Chapter 7, R;</p> <p>Chapters 11-12, R1;</p>	<p>Importing pre-written classes using the this keyword, calling and defining methods, writing and instantiating classes, setter/getter methods, instance variables, returning values, debugging using print function, containment and association, scope and parameter passing</p>
<p><b>Module 3:</b> Advanced features of OOP</p> <p>(~ 3 wks)</p>	<p>Interfaces, inheritance, polymorphism, abstract classes, immutability, copying and cloning objects</p>	<p>Chapters 2-3, T1;</p> <p>Chapter 10, T2;</p> <p>Chapter 7, R4;</p> <p>Chapters 11-12, R1;</p>	<p>Parameter polymorphism, method resolution, declared v/s actual type, partially and fully overriding methods, calling superclass constructor from child class constructor, protected fields and methods, using an abstract parent class v/s an interface with default and abstract methods, object equality check, object comparison (Comparable/Comparator interface), Cloneable interface/copy constructor</p>
<p><b>Module 4:</b> Unit testing</p> <p>(~0.5 wks)</p>	<p>Unit testing, developing test suite</p>	<p>R6;</p> <p>R7;</p>	<p>JUnit/Boost.test testing framework, assertion methods, testcase timeout, testing for exceptions, test suite</p>
<p><b>Module 5:</b> Using language APIs</p> <p>(~1 wks)</p>	<p>Language supported libraries for handling advanced data structures</p>	<p>Section 13.2, T2;</p> <p>Section 13.7, R1;</p>	<p>Big-O notation, Java collection framework (or Boost libraries), sorting objects, iterating over objects</p>
<p><b>Module 6:</b> Defensive programming</p> <p>(~1 wks)</p>	<p>Exception handling, assertions</p>	<p>Section 9.4, T2;</p> <p>Chapter 14, R1;</p> <p>Section 2.7, R4;</p>	<p>Exception handling using try/catch block, nesting try/catch blocks, throw and throws keywords, rethrowing exceptions, handling checked exception, user defined exceptions</p>

<b>Module 7:</b> Modeling and Design patterns  (~2 wks)	Basic modeling techniques – eg. Class diagram, sequence diagram, use case diagrams, etc. Introduction to design patterns: iterator, singleton, flyweight, adapter, strategy, template, prototype, factory, façade, decorator, composite, proxy, chain of responsibility, observer, state)	Chapter 5, T1; Chapters 3-9, R3	UML may be used for modeling.
--	---	------------------------------------	-------------------------------

**Suggested text books:**

T1. Grady Booch, Robert A. Maksimchuk, Michael W. Engle, Bobbi J. Young, Jim Conallen, Kelli A. Houston. Object-Oriented Analysis and Design with Applications.

T2. M. Scott. Programming Language Pragmatics. 4th edition.

**Suggested reference books / Online resources:**

- R1. R. Sebesta. Concepts of Programming Languages. 10th edition
- R2. J. Rumbaugh et al. The Unified Modeling Language Reference Manual.
- R3. Erich Gamma, Richard Helm, Ralph Johnson, John Vlissides, and Grady Booch. Design Patterns: Elements of Reusable Object-Oriented Software.
- R4. P. Van Roy and S. Haridi. Concepts, Techniques, and Models of Computer Programming.
- R5. <https://missing.csail.mit.edu/>
- R6. <https://www.baeldung.com/junit>
- R7. <https://www.tutorialspoint.com/junit/index.htm>
- R8. For UML tools, open source tools may be used (e.g. [www.starUML.io](http://www.starUML.io), [argouml.tigris.org/](http://argouml.tigris.org/))

---

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

**Option-1: Ability to implement basic event-driven programming-**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module-1:</b> Introduction to event-driven programming (~1 wks)	Adding UI controls, GUI events, event handling	O1; Chapter 10, O2;	Create GUI forms/application with different GUI components (buttons, text boxes, etc.), responding to user input by creating event handlers, productivity in creating event handlers using anonymous classes and lambda methods
<b>Module-2:</b> Advanced GUI programming (~2 wks)	Creating animations, connection with the database	O1	Implement simple 2-D game that uses animations (timelines, moving shapes), client-server programming (multiplayer game), object serialization/deserialization (save/reload game)

**Suggested text books / Online lectures or tutorials:**

1. Oracle documentation: <https://docs.oracle.com/javase/8/javase-clienttechnologies.htm>
2. P. Van Roy and S. Haridi. Concepts, Techniques, and Models of Computer Programming.

**Option-2: Understanding of the fundamentals of parallel programming**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module-1:</b> Introduction to multithreading (~1 wks)	Ahmdal's law, speedup, parallel efficiency, thread creation	Chapters 2 and 4, T2;	Creating and joining threads, concurrency decomposition, assigning tasks to threads, calculate parallel speedup and efficiency, debugging multithreaded programs
<b>Module-2:</b> Thread pool (~1 wks)	Task parallelism and thread pools (e.g., Java ForkJoin framework, OpenMP/CilkPlus/TBB in C++)	Section 13.2, T1; Section 13.7, R1;	Asynchronous task creation and joining, task parallelism for recursive parallelism, task dependency, controlling task granularity, comparing performance/productivity of explicit multithreading v/s using thread pools
<b>Module-3:</b> Mutual exclusion (~1 wks)	Race conditions, deadlocks, producer-consumer problem	Chapter 13, T1; Chapter 4, T2; Chapter 13, R1; Chapter 8, R2;	Multithreaded program to push/pop items from shared queue - using monitor/Mutex locks, conditional wait and signalling, volatile keyword



---

**Suggested text books / Online lectures or tutorials:**

T1. M. Scott. Programming Language Pragmatics. 4th edition

T2. P. Pacheco. An Introduction to Parallel Programming.

**Suggested reference books / Online resources:**

R1. R. Sebesta. Concepts of Programming Languages. 10th edition

R2. P. Van Roy and S. Haridi. Concepts, Techniques, and Models of Computer Programming.

**Option-3: Understanding of the basics of cloud computing**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 1:</b> Basic concepts plus a distributed programming framework such as MR/Spark (2 weeks)	Basics, Need, Advantages/Benefits, models such as PaaS/SaaS/IaaS/, distributed/cloud computing architectures  Programming frameworks such as Map-Reduce or Spark, how to write programs in these programs, features such as fault-tolerance, using commodity hardware,	T1, Ch. 1, 2, 3  Illustrate the differences between M-R style programming and C-style programming	Set up some pre installed programs and infrastructure to see how to run programs in the cloud.  Programming in Map-Reduce/Spark to study scalability. Typical problems are from matrix multiplication, text processing, graphs , web crawling, and the like
<b>Module 2:</b> (1 week)	Key challenges in cloud/distributed computing: communication vs computation/ problem decomposition/ failures/ etc. and analyze the communication required of basic algorithms such as matrix multiplication	T2, Ch 3.  T3, Ch 1	Redo exercises from week 2 to measure communication time

---

**Suggested text books / Online lectures or tutorials:**

T1: Data-Intensive Text Processing with MapReduce, Jimmy Lin and Chris Dyer, Morgan & Claypool Publishers, 2010.

T2: Parallel Computer Architecture, David Culler, J. P. Singh, and A. Gupta, Elsevier, 1998.

T3. [Distributed Computing: Principles, Algorithms, and Systems](#), A.D. Kshemkalyani, M. Singhal, Cambridge University Press, March 2011.

**Prepared by:**

1. Vivek Kumar, Computer Science & Engineering Department, IIIT Delhi
2. Kishore Kothapalli, Computer Science & Engineering Department, IIIT Hyderabad
3. Swarnendu Biswas, Computer Science & Engineering Department, IIT Kanpur

***The Discipline Graduate Attributes to which this course contributes significantly: CS1, CS3, CS7, CS8***

***Other discipline GAs to which this course may contribute somewhat: CS2***

## Operating Systems

<b>PCC- CS403</b>	<b>Operating Systems</b>	<b>3L:0T: 4P</b>	<b>5 credits</b>	<b>Pre-Reqs: PCC-CS402</b>
-------------------	--------------------------	------------------	------------------	----------------------------

### Prerequisites details:

1. Familiarity with the C/C++ programming language
2. PCC 402: Computer Organization and Computer Architecture
3. CS403 to be scheduled in semester after CS402  
(both courses not to be scheduled in the same semester)

### Learning Outcomes of the course (i.e. statements on students' understanding and skills at the end of the course the student shall have):

#### Essential (<=6):

1. To understand the role, functionality and layering of the systems software components
2. To understand the design and usage of the OS API and OS mechanisms
3. To understand the details of the abstractions and interfaces provided by the OS for program execution and execution requirements --- processes, threads, memory management, files.
4. To understand problems arising due to concurrency and related synchronization based solutions.
5. Hands-on and practical experience with usage of the OS API and basics of OS mechanisms

#### Desirable/Advanced (<= 3):

1. To gain an in-depth understanding of the design and implementation of OS internals via a teaching OS
2. To be able to implement incremental changes to the functionality of a teaching OS

#### Detailed contents for Essential Learning Outcomes:

<b>Module (appx dur in wks)</b>	<b>Topics</b>	<b>Pedagogy / teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 1:</b> Introduction to Operating Systems (0.5 week)	Application requirements  The systems stack and role of OS, resources, abstractions and interfaces  Components overview of an OS  Examples of different types of OSes (RTOS vs. desktop vs. mobile etc.), OS and OS distributions.	T1: Chapter 2 T2: Chapters 1, 2 T3: Chapter 1  R1: Chapter 1 R2: Chapters 1, 2  R4: Chapter 1  R7:Chapter 1 R8: Chapters 1 to 8	1. Usage of tools --- unix shell commands (file commands, ps, ls, top), text editor (nano, vi, gedit, emacs)  2. C programming language refresher --- header files, compilation and linking using gcc, program execution,

		O1: Bash Guide for Beginners	functions, argument passing, structures, pointers, file handling.
<b>Module 2:</b> Computer organization and computer architecture refresher (0.5 week)	Basic organization of hardware components  Role of OS relative to hardware functionality with examples related to the von Neumann architecture	T2: Chapter 1 T3: Chapter 1  R1: Chapter 1 R2: Chapters 1, 2  O1: Bash Guide for Beginners	1. Usage of tools --- gcc, gdb, objdump, shell scripts
<b>Module 3:</b> Processes (1.5 weeks)	Process abstraction --- program vs. process, Process Control Block (PCB)  Design of system calls --- invocation and basic OS handling  Process control system calls --- fork, wait, exec, getpid, getppid and variants  The limited direct execution model	T1: Chapters 4,5,6  T2: Chapters 6, 7  T3: Chapters 7, 8	1. Simple strace usage to showcase different interfaces (stdlib, system call)  2. Tools usage --- ps, pstree, top  3. Usage of process control system calls to identify process identifiers, create process hierarchies, launch new executables, control exit sequence of parent and child processes.  4. Familiarity with files in the /proc/<pid>/ directory
<b>Module 4:</b> Memory management (3 weeks)	Address bus and memory access  Memory view of a process ---- heap, stack, code, data  Process memory usage requirements  The address space abstraction using virtual memory and related system calls (mmap, munmap, sbrk, mprotect)	T1: Chapters 13,14,15,16, 17,18,19, 20  T2: Chapter 9  R8: Chapters 5,6	1. (Virtual) addresses of variables and initialized pointers.  2. Use of malloc() and demonstration of per-process virtual addresses  3. Tools usage --- strace, free, top, htop,

	<p>Address translation mechanisms - -- static mapping, segmentation, paging</p> <p>Page faults, page sharing, read/write permissions, swapping, process vs. OS memory</p> <p>Memory bookkeeping and management --- motivation and mechanisms (process and OS)</p> <p>Case studies --- (i) malloc and (ii) role of OS for program to process</p>		<p>vmstat, /proc/&lt;pid&gt;/maps</p> <p>4. Free memory statistics correlated with malloc(). Number of system calls and malloc() usage.</p> <p>5. Implement a custom memory allocator using system calls</p>
<p><b>Module 5:</b> Process management (1.5 weeks)</p>	<p>The process lifecycle---source code to execution</p> <p>The OS mode of execution --- limited direct execution recap, interrupts, system calls</p> <p>The process context switch mechanism and PCB state</p> <p>Scheduling policies --- set of scheduling metrics, goals and examples (interactive vs. realtime, priority)</p>	<p>T1: Chapters 7,8</p> <p>T2: Chapter 8</p>	<p>1. User mode programs to demonstrate LDE</p> <p>2. Demonstration of process execution interleaving in different orders</p> <p>3. Simulation based analysis of scheduling policies</p> <p>4. Tools usage --- nice, /proc/&lt;pid&gt;/status</p>
<p><b>Module 6:</b> Concurrency and Synchronization (3 weeks)</p>	<p>Motivation --- application, process and OS use cases.</p> <p>Introduction to threads and the pthread API</p> <p>Synchronization primitives --- limitations of software solutions, atomic instructions, test-and-set, spinlocks, mutexes, condition variables, semaphores</p> <p>Introduction to the pthread synchronization API</p> <p>Case studies --- producer-consumer, reader-writers, barriers</p>	<p>T1: Chapters 26, 27, 28, 29, 30, 31, 32</p> <p>T3: Chapters 11, 12</p>	<p>1. Creation of threads using the pthread API and modification of shared variables with and without synchronization</p> <p>2. Using spinlock, mutexes and condition variables to implement semaphores, barriers (using the pthreads API)</p> <p>3. Implement solutions to the producer-consumer, reader-writers problems using the different</p>

	Discussion on issues with concurrency —race conditions, deadlocks, order violation.		<p>synchronization primitives</p> <p>4. Develop synchronization solutions for applications that use shared data (e.g., ordering of threads, concurrent hash tables, etc.)</p> <p>5. Optional: Using shared memory and semaphores implement synchronized access to a shared memory area across processes (e.g., a message queue).</p>
<p><b>Module 7:</b> File systems  (2 weeks)</p>	<p>Persistence and the File abstraction</p> <p>Hardware view: Hard disk architecture and its interfacing</p> <p>Process view: System calls for file handling</p> <p>Roles and responsibilities of file system</p> <p>File system design details---file and file system metadata, directory structure, caching optimizations</p> <p>File System case study (the unix file system etc.)</p>	<p>T1: Chapters 36,37,39, 40</p> <p>T2: Chapters 4,5</p> <p>T3: Chapter 3</p>	<p>1. Tools usage --- state, file, du, df, fsck, ...</p> <p>2. Implementation of file utilities (e.g., find, grep, ...) using the system call API.</p> <p>3. Implement a simple file system to handle files on an emulated disk (via a large file) --- file system API, superblock, inode and data block management.</p>

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

Module	Topics	Pedagogy teaching suggestions	Nature of lab / assignment / practice
1. Introduction to setup and usage of a teaching OS	<p>Setup, configuration and usage of a teaching OS</p> <p>Understanding ISA details of the teaching OS</p> <p>Basics of interrupt handling, end-to-end execution of a system call and context switching mechanisms</p>	T4	<p>1. Configuration and setup of a teaching OS and writing of simple user mode programs (e.g., xv6)</p> <p>2. Simple logging style modification to system calls (e.g., #times a system call was invoked, histogram of number of system call invocations, ...)</p> <p>3. Implement new system calls (e.g., read PCB elements to user space of a process) (xv6 based)</p>
2. Basic modifications to the teaching OS functionality saved	<p>OS context of execution --- address space, stack</p> <p>Address translation mechanism in the teaching OS</p> <p>The interrupt handling mechanism (context save and restore, privilege level change, interrupt dispatch)</p> <p>Example usage of locks in the teaching OS</p> <p>File system state in the operating system (file descriptors, file objects, inodes, page cache ...) and in the teaching OS</p>	T4	<p>1. Write a system to call to output per process address space details</p> <p>2. Write a system to call to determine physical address of a virtual address</p> <p>3. Observe the stack pointers, privilege level registers in user and OS modes</p> <p>4. Modifying/profiling behaviour of exception handlers</p> <p>5. Observe process file table entries and file objects across parent and child processes</p>

<p>3. Advanced modifications to teaching OS functionality</p>	<p>Context switch design and implementation in the teaching OS</p> <p>Lazy allocation, page fault management and swapping/demand paging</p> <p>Understand implementation of the synchronization primitives in the teaching OS</p>	<p>T4</p>	<ol style="list-style-type: none"> <li>1. Write a system call to allocate the same physical block to different virtual addresses</li> <li>2. Implement lazy allocation of physical memory to processes</li> <li>3. System call to print saved state of any process</li> <li>4. Write a system call to induce page faults</li> <li>5. Design and implement a shared message queue between processes to be used via the system call interface</li> </ol>
---	---	-----------	--

**Suggested text books / Online lectures or tutorials:**

- T1. Operating Systems: Three Easy Pieces  
 Remzi H. Arpaci-Dusseau and Andrea C. Arpaci-Dusseau LLC  
<https://pages.cs.wisc.edu/~remzi/OSTEP/> (online version)
- T2. Design of the UNIX Operating System  
 Maurice J. Bacchus  
 Pearson Education India; First edition
- T3. Advanced Programming in the UNIX® Environment  
 W. Richard Stevens, Stephen A. Rago  
 Pearson Education India; Third edition
- T4. Xv6, a simple Unix-like teaching operating system  
 Frans Kaashoek, Robert Morris, and Russ Cox  
 [T4-R] <https://github.com/mit-pdos/xv6-riscv> (RISC-V version)  
 [T4-X] <https://github.com/mit-pdos/xv6-public> (x86 version)

**Suggested Online content:**



- 
1. The Linux Documentation Project, [www.tldp.org](http://www.tldp.org)

**Suggested reference books / Online resources:**

- R1. Modern Operating Systems, Andrew S. Tannenbaum and Herbert Bos, Pearson Education India; 4th edd
- R2. Operating System Concepts, Avi Silberschatz, Peter Baer Galvin, Greg Gagne, Wiley India; 9<sup>th</sup>, ed
- R3. Operating System courses offered on NPTEL, <https://nptel.ac.in/>
- R4. Think OS, A Brief Introduction to Operating Systems. Allen B. Downey  
<https://www.greenteapress.com/thinkos/index.html>
- R5. Linux Kernel Development , Robert Love, Pearson Education India; 3<sup>rd</sup> edition
- R6. Operating Systems: Principles and Practice, Thomas Anderson, Michael Dahlin, Recursive Books; 2<sup>nd</sup> Edition, <https://ospp.cs.washington.edu/index.htm>
- R7. Computer Systems: A Programmer's Perspective, Randall E. Bryant, David R.O'Hallaron, Pearson Education India; 3rd edition.
- R8. The C Programming Language, Brian Kernighan, Dennis Ritchie, Pearson Education India; 2nd ed.

**Prepared by:**

1. Purushottam Kulkarni, Department of Computer Science and Engineering (Indian Institute of Technology, Bombay)
2. Chester Rebeiro, Department of Computer Science and Engineering (Indian Institute of Technology, Madras)
3. Debadatta Mishra, Department of Computer Science and Engineering (Indian Institute of Technology, Kanpur)

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly:***  
CS3

***Other discipline GAs to which this course may contribute somewhat: CS6, CS8***

## Design and Analysis of Algorithms

CS404	Design and Analysis of Algorithms	3L:0T: 4P	5 credits	Pre-Reqs: ESC201, CS301
-------	-----------------------------------	-----------	-----------	-------------------------

**Learning Outcomes of the course (i.e. statements on students' understanding and skills at the end of the course the student shall have):**

### Essential (<=6):

1. Analyze the asymptotic performance of algorithms.
2. Establish the correctness of algorithms.
3. Demonstrate familiarity with major algorithms and data structures.
4. Apply important algorithmic design paradigms and methods of analysis.
5. Synthesize efficient algorithms in common engineering design situations.
6. Understanding limits of efficient computation.

### Desirable/Advanced (<= 3):

1. Intuitive understanding of what makes a problem NP-hard.
2. Apply network flow techniques to algorithm design.

### Detailed contents for Essential Learning Outcomes:

Module (approx dur in wks)	Topics	Pedagogy / teaching suggestions	Nature of lab / assignment / practice
<b>Module 1:</b> Applications of Graph Search (~1.5 wks)	(i) Review of BFS/DFS (ii) Checking if an undirected graph is 2-edge connected. (iii) Checking if a directed graph is strongly connected.	T1: Chapter 3, T2: Chapter 3, R1: Chapter 22.  Review breadth first search (BFS) and depth first search (DFS) in undirected and directed graphs; Introduce notion of a bridge and cut-vertex.	Checking if a graph is biconnected

<p><b>Module 2:</b> Greedy algorithms (2 wks)</p>	<p>(i) Introduction to the greedy paradigm</p> <p>(ii) Examples of activity selection, deadline scheduling, fractional knapsack, Kruskal's algorithm for minimum spanning trees, Huffman coding.</p>	<p>T1: Chapter 4, T2: Chapter 5.</p> <p>Exchange arguments are a useful recipe for proving correctness of greedy algorithms. Illustrate these through examples. Show examples where greedy does not give an optimum solution</p>	<p>(i) Prim's algorithm for minimum spanning trees.</p> <p>(ii) Examples where greedy algorithms are not optimal.</p>
<p><b>Module 3:</b> Divide and Conquer (2 wks)</p>	<p>(i) Explain why the divide and conquer paradigm is useful.</p> <p>(ii) Illustrate the paradigm through integer multiplication.</p> <p>(iii) Writing recurrence relations and solving them.</p> <p>(iv) Further examples from geometry – domination number of a set of points, identifying maximal points, closest pair of points.</p> <p>(iv) Linear time algorithm for finding the median.</p> <p>(v) Randomized divide and conquer algorithms: randomized quicksort and selection.</p>	<p>T1: Chapter 5, T2: Chapter 2.</p> <p>Solve recurrences by building the recurrence tree. Motivate choices of parameters in median finding algorithm by showing how recurrence tree changes. Mention expectation of a random variable and how to make sense of running time of randomized algorithms.</p>	<p>(i) Solve some recurrence relations.</p> <p>(ii) Modify discussed algorithms (e.g., dividing into three parts instead of two parts, or two unequal parts, etc.) and analyse using recurrences.</p> <p>(ii) Some elementary exercises on expectation calculation.</p>

<p><b>Module 4:</b> Dynamic Programming and shortest paths (2.5 wks)</p>	<p>(i) Computing Fibonacci numbers and why divide-and-conquer is not a good idea. Idea of storing function calls, tables.</p> <p>(ii) Notion of subproblems and optimal substructure.</p> <p>(iii) Illustration through subset sum, (integer) knapsack, longest increasing subsequence, longest common subsequence, matrix chain multiplication. Dijkstra's algorithm for single-source shortest paths, Bellman-Ford for SSSP with -ve weights, Floyd Warshall for APSP.</p>	<p>T1: Chapter 6, T2: Chapter 4, Chapter 6.</p> <p>Discuss why Dijkstra's algorithm is an example of dynamic programming. Extending Bellman-Ford to APSP and to find negative cycles. Discuss how dynamic programming problems can often be cast as longest paths in acyclic graphs.</p>	<p>Exercises on dynamic programming (textbook problems)</p>
<p><b>Module 5:</b> Network flows (2 wks)</p>	<p>The maximum s-t flow problem in capacitated networks. Ford Fulkerson algorithm or maximum flow. Max-flow min-cut theorem and integrality of maximum flow for integral capacities. Applications of max flow to maximum bipartite matching, max disjoint paths</p>	<p>T1: Chapter 7. 1-7.2, 7.5, 7.6</p> <p>Notion of residual capacities and residual graphs and how this allows us to correct greedy decisions made in earlier steps of FF algorithm</p>	<p>(i) Some simple examples. (ii) Implementation of Ford Fulkerson algorithm.</p>

<p><b>Module 6:</b> Intractability (2wks)</p>	<p>(i) Models of computation, Turing machines, RAM model. Brief discussion on other models of computation e.g. PRAM model, Memory Hierarchy etc.</p> <p>(ii) Notion of polynomial time computation.</p> <p>(iii) Polynomial time reductions. Yes and No instances of decision problems. Decision vs optimization.</p> <p>(iv) NP as a class of problems with Yes certificates which can be efficiently checked.</p> <p>(v) NP-hardness and Cook-Levin theorem (just the statement).</p> <p>(vi) NP-completeness. Examples of Reductions.</p>	<p>T1: Chapter 8, T2: Chapter 8</p> <p>(i) Emphasize how reduction can be used to solve a problem using an algorithm for a different problem.</p> <p>(ii) Emphasize the asymmetry in the definition of NP: No efficiently checkable NO certificates for problems in NP.</p> <p>(iii) Problems which are NP-hard but not in NP.</p> <p>(iv) Examples of poly time reductions. Polytime as a measure of efficiency. Hardness only for the general instance, whereas special instances can be efficiently solvable.</p>	<p>Exercises on reductions, NP-completeness.</p>
---	--	--	--

**Suggested Texts:**

1. Algorithm Design, 1ST Edition, Jon Kleinberg and ÉvaTardos, Pearson.
2. Algorithms, Sanjoy Dasgupta, Christos Papadimitriou, Umesh Vazirani

**Suggested reference books:**

1. Introduction to Algorithms, 4TH Edition, Thomas H Cormen, Charles E Lieserson, Ronald L Rivest and Clifford Stein, MIT Press/McGraw-Hill.
2. Algorithm Design: Foundations, Analysis, and Internet Examples, Second Edition, Michael T Goodrich and Roberto Tamassia, Wiley.

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module1:</b>	Polynomial multiplication using FFT and DFT.	T1: Chapter 5.5, T2: Chapter 2.5  (i) Explain the interplay between the two representations of a polynomial.  (ii) Explain why complex numbers arise naturally here.	Take two polynomials and multiply them by first computing the point-value representation using FFT and then recovering product using inverse FFT
<b>Module2:</b>	Beyond NP-completeness: approximations, exponential algorithms, popular heuristics.	T1: Chapter 10, T2: Chapter 9  (i) Show that TSP can be solved in $2^n \text{poly}(n)$ time (rather than $n!$ Time), and vertex cover in $2^k \text{poly}(n)$ , where $k$ is the size of the optimal vertex cover.  (ii) Give examples of some simple approximation algorithms, e.g., bin packing, vertex cover.  (iii) Popular heuristics for satisfiability.	(i) Implementation of some exponential time algorithms and heuristics and see how they scale with large $n$ .

**The Discipline Graduate Attributes (GAs) to which this course contributes significantly: CS1, CS2, CS4**

---

**Other discipline GAs to which this course may contribute somewhat:**

## Database Systems

<b>Course code: ?</b>	<b>Introduction to Database Systems</b>	<b>3L:0T: 4P</b>	<b>Credits: 5</b>	<b>Pre-Req: Data Structures and Algorithms</b>
-----------------------	---	------------------	-------------------	--

**Learning Outcomes of the course (i.e. statements on students' understanding and skills at the end of the course the student shall have):**

**Essential (<=6):**

1. Ability to design and implement database schema for an application using RDBMS concepts.
2. Ability to write SQL queries for tasks of various complexities.
3. Ability to write an application program that uses a database system as the backend.
4. Understanding of internal working of a DBMS including data storage, indexing, query processing, transaction processing, concurrency control and recovery mechanisms.
5. Awareness of non-relational and parallel/distributed data management systems with a focus on scalability.

**Desirable/Advanced (<= 3):**

[Nil]

**Detailed contents for Essential Learning Outcomes:**

<b>Module (appx dur in wks)</b>	<b>Topics</b>	<b>Pedagogy / teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 1:</b> Introduction (~ 1 wk)	(i) Motivation  (ii) Introduction to Data Models (Relational, Semistructured, ER)		

<b>Module 2:</b> Relational Databases (~ 1.5 wks)	(i) Relational Data Model (ii) Relational Algebra (iii) Relational Calculus or Connection to First Order Logic (Optional)		(i) Simple pen+paper, and using Relax Relational Algebra calculator in browser
<b>Module 3:</b> SQL + interfacing with database (~3 wks)	(i) DDL (ii) Insert/Delete/Update (iii) Simple Queries (select/project/join/ aggregate queries) (iv) Complex queries (With Clause, Nested Subqueries, Views) (v) Programming in a standard language and interfacing with a DB backend		(i) Laboratory exercises where students write SQL queries for various tasks.  Platform can be PostgreSQL preferably, or MySQL. W3Schools/SQLite in web browser can also be used but beware of non-standard SQL features.  (ii) Practice interfacing with a database from a program using connectors like JDBC/ODBC...
<b>Module 4:</b> Big Data (~1 wk)	Key-value Stores and Semi-structured Data, using JSON and MongoDB, or other combinations		(i) Small exercises on MongoDB
<b>Module 5:</b> Database Design (~2 wk)	(i) Introduction to ER model (ii) Mapping from ER to relational model (iii) Functional Dependencies (iv) Normalization (BCNF, Optionally 3NF)		(i) Exercise in ER design for an application starting with natural language description  (ii) Convert ER design to tables  (iii) Pen-and-paper exercises with FDs and normalization
<b>Module 6:</b> Physical Design (~2 wk)	(i) Overview of Physical Storage (Hard Disks, Flash/SSD/RAM), sequential vs random I/O, Reliability via RAID  (ii) Storage Organization (Records, Pages and Files),		(i) Use a B+-tree visualization system to understand how B+-trees work



	Database Buffers, Database Metadata  (iii) Indexing, B+-Trees		
<b>Module 7:</b> Query Processing and Optimization  (~2 wk)	(i) Query Processing: External sort, Joins using nested loops, indexed nested loops  (ii) Overview of Query Optimization: equivalent expressions, and concept of cost-based optimization		(i) Examine query plans for sample queries by using the Explain feature of database systems.  (ii) Small exercises to show benefit of indices.
<b>Module 8:</b> Transaction Processing  (~2 wk)	(i) Concept of transactions and schedules, ACID properties  (ii) Conflict-serializability  (iii) Concurrency control: locks, 2PL, Strict 2PL, optional: isolation levels  (iv) Recovery using undo and redo logs		(i) Pen-and-paper exercises on conflicts, cycles, conflict serializability, recoverability, etc.

**Suggested text books / Online lectures or tutorials:**

1. Database System Concepts, 7th Ed, Silberschatz, Korth and Sudarshan, McGraw-Hill. Indian Edition released 2021
2. Fundamentals of Database Systems, 7th Ed, Elmasri and Navathe, Pearson Pubs, 2017
3. Principles of Database Management, Lemahieu, Broucke and Baesens, Cambridge University Press, 2018

**Suggested reference books / Online resources:**

1. Software
  - a. Relax Relational algebra calculator: <https://dbis-uibk.github.io/relax/landing>
  - b. SQL: PostgreSQL/MySQL/MariaDB, or SQLite in browser
  - c. B+-tree visualization:  
<https://www.cs.usfca.edu/~galles/visualization/BPlusTree.html>
  - d. MongoDB

- 
- e. Various DB systems playground: <https://www.pdbmbook.com/playground>

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly: CS3, CS5, CS6***

***Other discipline GAs to which this course may contribute somewhat: CS1, CS7, CS9***

---

## Computer Networks

CS602	Computer Networks	3L:0T: 2P	4 credits	Pre-Reqs: CS402 Computer Organisation, CS403 Operating Systems
-------	-------------------	-----------	-----------	--

### Pedagogy and Learning Outcomes

This course introduces students to the fundamental principles of computer networks, and to their use in the Internet. Lab assignments cover network programming, network tools, applications, and simulation. Through these assignments, students get a broad understanding by building, operating and tuning components of the Internet. The top-down, application-driven design of the course motivates students with familiar uses and problems of their digital world and enables them to work with real-world applications from early in the course.

### At the end of the course the student shall:

#### **Essential (<=6):**

1. Understand the architecture principles that have enabled the orders of magnitude expansion of the Internet
2. Understand networked applications and their protocols, their installation, operation and performance tuning
3. Understand layering as a means of tackling complexity, layering applied to the Internet
4. Understand protocols as a structured means of reliable communications
5. Be conversant with network programming using the socket API
6. Be familiar with tools for configuring, monitoring and tuning the Internet and hosts

#### **Desirable/Advanced (<= 3):**

1. Understand basics of data centre networks and software-defined networks (SDN)
2. Understand cellular networks, mobility and the impact on applications
3. Understand the interaction between streaming media applications and the underlying network infrastructure.

**Detailed contents for Essential Learning Outcomes:**

<b>Module (approx duration in wks)</b>	<b>Topics</b>	<b>Pedagogy / teaching suggestions</b>	<b>Lab / assignment / practice</b>
<p><b>Module 1:</b> Introduction to the Internet (1 week)</p>	<p>Overview of how the Internet works. Understand at a high level what happens when we browse a website. Understand basic terminology like browser, web server, URL, domain name, IP address, packets. (1 hour)</p> <p>Overview of the design principles of the Internet: packet switching vs circuit switching, store-and-forward networks, layering for modularity. Introduction to the various layers in the Internet. (1 hour)</p> <p>Introduction to performance metrics like end-to-end throughput, delay, jitter and drop rates in a network. Statement of Little's Law. How performance is measured. (1 hour)</p>	<p><b>Kurose &amp; Ross:</b> Chap 1 (sec. 1.1-1.5). <b>Raj Jain:</b> Sec. 3.2 - 3.4</p> <p>Top-down sequence is recommended. Performance metrics and measurement concepts are introduced in Module 1 and used in all modules (lectures and labs) as appropriate.</p>	<p>- Use Linux tools like ifconfig, dig, ethtool, route, netstat, nslookup, and ip to understand the networking configuration of the computer that the student is working on.</p> <p>- Use Wireshark to capture packets when browsing the Internet. Examine the structure of packets: the various layers, protocols, headers, payload.</p> <p>This resource is useful for many of the lab assignments: <a href="https://gaia.cs.umass.edu/kurose_ross/about.php">https://gaia.cs.umass.edu/kurose_ross/about.php</a></p>
<p><b>Module 2:</b> Application layer (2 weeks)</p>	<ul style="list-style-type: none"> <li>- Internet names, how DNS works. (1 hour)</li> <li>- Application layer protocols: HTTP, SMTP, SNMP, web applications. (3 hours)</li> <li>- Peer-to-peer applications. P2P file distribution. (1 hour)</li> <li>- Audio and video streaming. Challenges of streaming over best effort IP. (1 hour)</li> </ul>	<p><b>Kurose and Ross:</b> Chap 2 (sec. 2.1-2.5), Sec. 7.1.3</p>	<ul style="list-style-type: none"> <li>- Install and configure some network applications, eg. Apache, Bind (DNS), etc.</li> <li>- Understand various header fields and their usage in different application layer protocols using Wireshark packet capture.</li> </ul>

Module (approx duration in wks)	Topics	Pedagogy / teaching suggestions	Lab / assignment / practice
<p><b>Module 3:</b></p> <p>Linux Network Programming (0.3 week)</p>	<p>- Introduction to socket programming in Linux. Understand how to build a simple client-server application using TCP/UDP sockets. (1 hour)</p>	<p><b>Kurose and Ross:</b> Sec 2.7.</p>	<p>- Socket programming: write a simple client-server program using TCP and UDP sockets.</p> <p><b>Optional:</b> Modify server to handle multiple clients concurrently.</p>
<p><b>Module 4:</b></p> <p>Transport Layer (2.3 weeks)</p>	<p>- Importance of the transport layer; end-to-end principle. Transport layer protocols: basics of TCP and UDP, process-to-process delivery, multiplexing, port numbers, header structure. (2 hours)</p> <p>- Reliable transmission of packets over an unreliable network: sequence numbers, ACKs, timeout, retransmissions. Stop and wait, and sliding window. (2 hours)</p> <p>- TCP connection setup and teardown. (1 hour)</p> <p>- Flow control and congestion control at the transport layer. Differences between the two. Overview of TCP congestion control: Slow start and reaction to timeouts (2 hours)</p> <p>- TCP congestion control: Slow start; congestion avoidance using loss-based and delay-based control. (1 hour)</p>	<p><b>Kurose &amp; Ross:</b> Chap. 3 (sec. 3.1-3.6)</p> <p><b>Lab:</b> Optional -- <a href="https://www.isi.edu/nsnam/ns/tutorial/nsscript1.html#second">https://www.isi.edu/nsnam/ns/tutorial/nsscript1.html#second</a></p> <p>Validation tests and demos: <a href="https://www.isi.edu/nsnam/ns/ns-tests.html">https://www.isi.edu/nsnam/ns/ns-tests.html</a></p>	<p>- Measure TCP throughput between two hosts in a network using tools like iperf. Modify TCP configuration parameters. Use the <b>tc</b> Linux utility or similar to control bandwidth, delay, loss. Observe impact on measured throughput.</p> <p>- Experiment with multiple applications running concurrently to generate congestion.</p> <p><b>Optional:</b> Observe the behaviour of congestion control protocols in ns-2/ns-3, change various network parameters and observe evolution of the TCP congestion window.</p>

Module (approx duration in wks)	Topics	Pedagogy / teaching suggestions	Lab / assignment / practice
<p><b>Module 5:</b> The IP Layer (2 wks)</p>	<p><b>[A] Network architecture and performance</b></p> <ul style="list-style-type: none"> <li>- Network topology; Router architecture: queueing and switching. (2 hour)</li> <li>- Performance evaluation of a network link: traffic characteristics, performance measures, Kendall's notation. (1 hour)</li> </ul> <p><b>[B] IP Protocol</b></p> <ul style="list-style-type: none"> <li>- Need for an Internet address, and its design. Hierarchical IP addressing, IPv4 and IPv6, structure of IP datagram, IP forwarding.</li> <li>- NATs, security attacks and defences: DMZ, firewalls. (3 hours)</li> </ul>	<p><b>Kurose &amp; Ross:</b> Chap. 1.4; Chap. 4 (Sec. 4.1, 4.3, 4.4)</p> <p><b>S. Bose:</b> Chap. 1; Chap. 2.3</p>	<ul style="list-style-type: none"> <li>- Use tools like ping and traceroute to explore various Internet paths to popular servers.</li> <li>- Use web-based tools like the <b>whois</b> utility to query Internet registries, and understand which IP addresses are allocated to the student's network. Find out which are the major ISPs, and which is the ISP of the student's network.</li> </ul>
<p><b>Module 6:</b> Routing protocols and Internet architecture (2 wks)</p>	<ul style="list-style-type: none"> <li>- Routing protocols: Link state routing. Distance vector routing: count-to-infinity, routing convergence. (3 hours)</li> <li>- Understand the structure of the Internet: end-user organizations and ISPs. Understand the difference between intra-domain (OSPF) and inter-domain (BGP) routing. Intra-domain routing: OSPF. (3 hours)</li> </ul>	<p><b>Kurose &amp; Ross:</b> Chap 4 (Sec. 4.5 - 4.6)</p>	<ul style="list-style-type: none"> <li>- Configure a simple mesh network using computers in the lab, or using Mininet. Setup static routes to conform to the desired mesh topology.</li> <li>- Use NS-2/NS-3 to simulate a mesh of at least 4 nodes and 3 links to evaluate performance under various conditions</li> </ul>

Module (approx duration in wks)	Topics	Pedagogy / teaching suggestions	Lab / assignment / practice
<b>Module 7:</b> Data Link Layer (2 wks)	<ul style="list-style-type: none"> <li>- Mechanisms for error detection/recovery: Parity checks, CRC (1 hour)</li> <li>- Medium access protocols: Polling vs. contention-based: TDM, Aloha, CSMA/CD (3 hours)</li> <li>- Switched LANs: L2 addressing and ARP, Ethernet frame structure, learning switches. (2 hours)</li> </ul>	<b>Kurose &amp; Ross:</b> Chap. 5 (Sec. 5.1 - 5.4)	<ul style="list-style-type: none"> <li>- Use Linux network tools like ethtool to observe and analyze link layer packet statistics and errors.</li> <li><b>Optional:</b> Use NS-2/NS-3 to simulate medium access protocols. Observe contention, collisions and packet loss in medium access protocols. Observe the working of error detection/recovery mechanisms.</li> </ul>
<b>Module 8:</b> Wireless Networks (1 wk)	<ul style="list-style-type: none"> <li>- Wireless physical layer: signal-to-noise ratio, bit error rate, modulation, multipath, interference (1 hour)</li> <li>- Wireless LANs: 802.11 architecture (access points, SSID, channels, beacons, scanning, association), 802.11 CSMA-CA protocol; summary of 802.11 variants (3 hours)</li> </ul>	<b>Kurose &amp; Ross:</b> Chap 6 (Sec. 6.1, 6.2, 6.3)	<ul style="list-style-type: none"> <li>- Use cellphone to measure WiFi signal strength (RSS) at various places in the campus. Draw a contour map with access points and RSS levels. Correlate with upload/download speed using tools like Measurement Lab speedtest.</li> <li><b>Optional:</b> Understand the behavior of WiFi using NS-2/NS-3.</li> </ul>

---

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module (approx duration in wks)</b>	<b>Topics</b>	<b>Pedagogy/ teaching suggestions</b>	<b>Lab / assignment / practice</b>
<b>Module D1:</b> SDN and data centre networking	Data centre network design and topology. Transport protocols optimized for data centres. Introduction to software defined networking.	<b>Kurose and Ross:</b> Chap. 5.6 Refs [5], [6]	Simulate transport protocols optimized for data centres in NS-2/NS-3.
<b>Module D2:</b> Mobile data networks	Cellular Internet access. Mobility management and handovers. Impact of mobility on higher layers.	<b>Kurose and Ross:</b> Chap. 6.4-6.8	Use cellphone to measure cellular signal strength (RSS) at various places in the campus. Draw a contour map with cellphone towers and RSS levels. Correlate with upload/download speed using tools like Measurement Lab speedtest.
<b>Module D3:</b> Streaming media protocols and services	Multimedia applications. Audio and video streaming over UDP, HTTP. Adaptive streaming. Voice over IP. Recovering from packet loss and jitter. QoS. Protocols for real time applications.	<b>Kurose and Ross:</b> Chap. 7	Implement a streaming audio/video server using open-source software.

**Suggested text books / Online lectures or tutorials:**

1. J.F. Kurose and K.F. Ross, *Computer networking: a top-down approach*, 6<sup>th</sup> ed, Pearson, 2017. (6<sup>th</sup> ed. is low-cost Indian edition. 7<sup>th</sup> ed. is high-cost, may be used if available.)

**Suggested reference books / Online resources:**

1. R. Jain, *The art of computer systems performance analysis*, Wiley India, 1991.
2. S.K. Bose, *An Introduction to Queueing Systems*, Springer Science+Business Media New York, 2012.
3. A.S. Tanenbaum and D.J. Wetherall, *Computer Networks*, 5<sup>th</sup> ed., Pearson, 2013.
4. Larry Peterson and Bruce Davie, *Computer Networks: A Systems Approach*, 6<sup>th</sup> Ed., available at <https://book.systemsapproach.org/>
  - N. Feamster, J. Rexford, Ellen Zegura, "The Road to SDN", *ACM Queue*, 2013.
  - Alizadeh et al., "Data Center TCP", *ACM SIGCOMM 2010*.



- 
- *The Network Simulator - ns-2*, <https://www.isi.edu/nsnam/ns/>
  - E. Altman and T. Jimenez, *NS2 Simulator for Beginners*, 2003.
  - *ns-3 network simulator*. <https://www.nsnam.org/>

**Prepared by** (alphabetical order):

1. Timothy A Gonsalves, School of Computing & Electrical Engg, IIT Mandi
2. BN Jain, Department of Computer Science & Engg, IIIT Delhi
3. Vinay Ribeiro, Department of Computer Science & Engg, IIT Bombay
4. PM Sreelakshmi, School of Computing & Electrical Engg, IIT Mandi
5. Mythili Vutukuru, Department of Computer Science & Engg, IIT Bombay

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly:*** CS3

***Other discipline GAs to which this course may contribute somewhat:*** CS4, CS6, CS8

---

## Machine

## Learning

CSXXX	Machine Learning	3L:1T: 0P	4 credits	Pre-Reqs: Engg. mathematics (LinAlg, Prob), Algorithms
-------	------------------	-----------	-----------	--

### Learning Outcomes of the course (i.e., statements on students' understanding and skills at the end of the course the student shall have):

#### Essential (<=6):

1. Understanding popular ML algorithms with their associated mathematical foundations.
2. Capability to implement basic algorithms using basic (python) libraries. Have hands-on experience in applying ML to problems encountered in various domains. Have exposure to high level ML libraries or frameworks such as TF, pytorch.
3. Make aware of the role of data in the future of computing and solving real-world problems.
4. Helping them connect/map real-world problems to the appropriate ML algorithm(s) to solve them
5. Appreciate the mathematical background behind the popular ML algorithms
6. Have awareness about the importance of core CS principles such as algorithmic thinking and systems design in ML

#### Desirable/Advanced (<= 3):

1. Have a solid mathematical understanding of the popular ML algorithms
2. Preparedness to use state of the art machine learning algorithms in formulating and solving new problems.
3. Capability to train (or solve optimization problems) ML models with applications in real world use cases.

#### Detailed contents for Essential Learning Outcomes:

Module (appx duration of 3 weeks or 9-12 Hrs)	Topics	Pedagogy / teaching suggestions	Nature of lab / assignment / practice
Introduction to ML (~2wk)	(i) Motivation and role of machine learning in computer science and problem-solving  (ii) Representation (features), linear	(i) Connect machine learning to the broader theme of Computer Science  (ii) Expose broad canvas of machine learning; brief history and importance	(i)Experiments/notebooks/code that refresh Python, programming frameworks used for the course  (ii) Experiments/Code that allows students to appreciate mathematics and data

	<p>transformations, Appreciate linear transformations and matrix vector operations in the context of data and representation.</p> <p>(iii) Problem formulations (classification and regression).</p> <p>(iv) Appreciate the probability distributions in the context of data, Prior probabilities and Bayes Rule.</p> <p>(v) Introduce paradigms of Learning (primarily supervised and unsupervised. Also a brief overview of others)</p>	<p>(iii) Role of data, Connection to the knowledge /experience in learning.</p> <p>(iv) Show successful examples of machine learning in Industry/working</p> <p>(v) Motivate students by Showing how ML and Data driven solutions could help in our day to day problems around (interdisciplinary such as agriculture, healthcare, education, living etc.)</p> <p>(vi) Refresh the basic mathematical notions that students may know (vectors, matrices, probabilities, etc.) with examples in ML</p> <p>(vii) Make students aware of relevant topics like “what is learnable”? And “what are the disadvantages of data driven solutions”.</p>	<p>manipulation. Appreciate (a) Features, Representation of the data/real-world phenomena (b) mathematical operations or transformations that manipulate the data (c) plot/visualise the data distributions (say in 2D) (d) eigen values, eigen vectors, rank of matrices.</p> <p>(iii) Lab/Experiments that appreciate the problem of Classification and problem of Regression</p> <p>(iv) Lab/Experiments that appreciates the notions related to “Training” and “Testing” by considering algorithms like decision trees, nearest neighbour as black boxes.</p>
<p>Fundamentals of ML (~3 wks)</p>	<p>(i) PCA and Dimensionality Reduction,</p> <p>(ii) Nearest Neighbours and KNN.</p> <p>(iii) Linear Regression</p> <p>(iv) Decision Tree Classifiers</p> <p>(iv) Notion of Generalization and</p>	<p>(i) Focus on mathematical and algorithmic precise description of the content.</p> <p>(ii) Insights into these algorithms, why? When? What are the limitations? Why multiple algorithms exists for a specific problem</p> <p>(iii) Insights into the notion of generalization. Challenges for generalization. Assumptions to make</p>	<p>(i) Dimensionality Reduction using PCA and its applications in (a) removing irrelevant features (b) compression /compaction (c) efficient ML pipeline</p> <p>(ii) Experiment related to Nearest neighbour classifier, (a) visualize the decision boundaries (b) appreciate the role of hyperparameter K. Role of validation data in choice of hyper parameters</p>

	<p>concern of Overfitting</p> <p>(v) Notion of Training, Validation and Testing; Connect to generalisation and overfitting.</p>	<p>(iv) Practical insights and tips on avoiding overfitting.</p>	<p>(iii) Decision Tree as a classifier and see the overfitting with “deep” trees. How the overfitting can be controlled by seeing validation performance during the training.</p>
<p>Selected Algorithms (~4 wks)</p>	<p>(i) Ensembling and RF</p> <p>(ii) Linear SVM,</p> <p>(iii) K Means,</p> <p>(iv) Logistic Regression</p> <p>(v) Naive Bayes</p>	<p>(i) Make students appreciate the role of optimization in machine learning. Challenges in optimization and why we are sometimes happy with sub-optimal solutions. How assumptions make the algorithms simple/tractable.</p> <p>(ii) Make students appreciate the role of uncertainty in data and machine learning problems/solutions. Give probabilistic insights into Loss functions (em MSE, cross entropy)</p> <p>(iii) Introduce iterative algorithms, convergence, role of initialization etc. in a class of ML solutions.</p> <p>(iv) Connect the geometric view of Margin (eg, linear SVM) and Probabilistic View of Margin (Logistic Regression) and the need of Generalization</p>	<p>(i) Experiments related to K-Means, by varying in “K”, “initialization”. How the “analysis of the algorithm” can be seen in the lab (eg. change of objective across iterations). Try multiple datasets. Appreciate that “unsupervised discovery” makes sense in the problem under consideration.</p> <p>(ii) An experiment that demonstrates how SVM can yield a solution better than a simple linear separating solution. Appreciate the role of support vectors. Appreciate how SVMs extend to problems even if data is not linearly separable.</p> <p>(iii) An experiment that makes students appreciate the utility of naive Bayes classifier in practice (say designing a text classifier).</p>
<p>Neural Network Learning (~4 wks)</p>	<p>(i) Role of Loss Functions and Optimization,</p>	<p>(i) Appreciate (a) the neuron model (b) the neural network and its utility in modelling and</p>	<p>(i) Experiment that exposes the GD and BP in simple neural networks. Show the</p>

	<p>(ii) Gradient Descent and Perceptron/Delta Learning,</p> <p>(iii) MLP,</p> <p>(iv) Backpropagation</p> <p>(v) MLP for Classification and Regression,</p> <p>(vi) Regularisation, Early Stopping</p> <p>(vii) Introduction to Deep Learning</p> <p>(viii) CNNs</p>	<p>solving the problem. Connect to the biological motivations and parallelism.</p> <p>(ii) Expose the simple elegant optimization scheme of gradient descent with associated mathematical rigour and insights.</p> <p>(iii) Expose the practical issues in extending GD to multiple layers and how the backpropagation algorithm efficiently computes the gradients.</p> <p>(iv) Expose the practical challenges in training a neural network (such as non-convexity, initialization, size of data, number of parameters) and how they are taken care of in the practical implementations of today.</p> <p>(v) Appreciate the need for empirical skills in training neural networks.</p>	<p>learning process (graphs) and performances.</p> <p>(ii) Experiment that use a modern library and implementation of a deep neural network, expose computational graphs, expose the generalized way of appreciating BP as a learning algorithm in Deep Neural Networks</p> <p>(iii) Experiment that uses a popular CNN architecture for practical application (say image classification).</p> <p>(iv) Experiments that strengthen the empirical skills in training with (a) initializations (b) update strategies (c) regularisation (d) multi fold validation on a small/medium size deep neural network that can be trained in 5 minutes.</p>
--	--	--	--

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
Key Concepts from ML	Kernels (with SVM), Bayesian Methods, Generative Methods, HMM, EM, PAC learning	Focus on mathematical and analytical skills. Expose the intuition behind these algorithms. Introduce analysis of machine learning using a PAC model.	Python notebooks that demonstrate the use of these algorithms on public datasets
Deep Learning Architectures	Popular CNN Architectures, RNNs, GANS and Generative Models,	Introduce popular architectures, models, and the use of it in various settings.	(i) Use of popular architectures for pretrained features and transfer learning  (ii) Use of RNNs in learning “language models” in large text corpus (charRNN)  (iii) Capability and practical challenged in working with GANS
Training Today's Neural Networks	Advances in Backpropagation and Optimization for Neural Networks  Adversarial Learning	Appreciate the challenges in large nonconvex optimization and how many of today's design choices have helped.	(i) See how (a) initialization (b) momentum (c) update rules have helped in getting better minima/soln.  (ii) Experience how regularisation helps in avoiding overfitting and getting better solutions

---

### **Notes on Exercises/Labs/Homeworks:**

1. Use iPython/Jupyter notebooks to hand out assignments; such notebooks allow embedding instructions/videos/etc, making it easy for instructors, students as well as TAs
2. It may be good to have both theory and programming components in the assignment/homework component, to allow students to appreciate and learn both aspects of machine learning
3. Based on target audience, have a healthy mix of using an off-the-shelf machine learning library (e.g. using sklearn's decision tree function) and writing an algorithm from scratch (e.g. coding up decision tree from scratch)
4. Consider having a Kaggle-style hackathon in the course, where students get used to competitive machine learning

### **Notes on Sequencing Lectures:**

1. Consider your target audience (undergrad vs grad; 2nd-year vs 3rd/4th-year students;etc) in deciding the prerequisites that the students may fulfil, and also the topics that you may want to cover.
2. It may be worthwhile visiting publicly available course outlines to verify and confirm that no important topics are missed out from both fundamental and contemporary perspectives (it is important to keep both these considerations in mind – to raise the spark of fundamental curiosity and also to provide utilitarian value)..
3. Sequencing can be done considering the target audience in mind – one option may be to start with simple algorithms and gradually move towards mathematically involved ones.
4. Sequencing can also be modified based on whether there are any other machine learning or AI or related courses in the curriculum/elective list.

### **Suggested text books / Online lectures or tutorials:**

1. Marc Peter Deisenroth, A. Aldo Faisal, Cheng Soon Ong, Mathematics for Machine Learning, Cambridge University Press (23 April 2020)
2. Tom M. Mitchell- Machine Learning - McGraw Hill Education, International Edition
3. Aurélien Géron Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, O'Reilly Media, Inc. 2nd Edition

### **Reference Books:**

1. Ian Goodfellow, Yoshoua Bengio, and Aaron Courville Deep Learning MIT Press Ltd, Illustrated edition
2. Christopher M. Bishop Pattern Recognition and Machine Learning - Springer, 2nd edition
3. Trevor Hastie, Robert Tibshirani, and Jerome Friedman - The Elements of Statistical Learning: Data Mining, Inference, and Prediction - Springer, 2nd edition

---

**Prepared****by:**

1. Mitesh Khapra, IIT Madras
2. Vineeth N Balasubramanian, IIT Hyderabad
3. P. K. Biswas, IIT Kharagpur
4. Piyush Rai, IIT Kanpur
5. Preethi Jyothi, IIT Bombay
6. Pabitra Mitra, IIT Kharagpur
7. Chetan Arora, IIT Delhi
8. P J Narayanan, IIIT Hyderabad
9. C. V. Jawahar, IIIT Hyderabad



---

## Cyber Security

CSXXX	Introductory Cyber Security	3L:0T: 4P	5 credits	<b>Pre-Reqs:</b> <b>Introduction to Programming, Data Structures &amp; Algorithms, Operating Systems, Networking</b>
-------	-----------------------------	-----------	-----------	---

### Learning Outcomes of the course (i.e. statements on students' understanding and skills at the end of the course the student shall have):

Having completed this course, a student should be able to:

#### Essential (<=6):

1. Understand the importance of cyber security (data confidentiality, Integrity, and Availability) and various recent attacks on important digital systems such as banking, e-commerce systems, e-governance systems etc.
2. Understand basic cryptography concepts – symmetric vs asymmetric cryptography, Public Key Crypto Infrastructure (PKI), Symmetric Ciphers, Hashing, Digital Signatures.
3. Understand methods and tools for authentication, authorization, privilege, and their needs in securing an organization's IT system.
4. Understand the common vulnerabilities in applications, web applications, network, and the Internet Infrastructure.
5. Understand the methods and tools for Intrusion Detection (network and host intrusion detection) and perimeter security (firewall).
6. Understand basic malware functions and indicators of compromise.

#### Desirable/Advanced (<= 3):

1. Understand basic mobile application security issues and android platform architecture for securing app execution .
2. Understand wireless LAN security issues.

**Detailed contents for Essential Learning Outcomes:**

<b>Module (appx dur in wks)</b>	<b>Topics</b>	<b>Pedagogy / teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<p><b>Module 1:</b> Introduction and basic terminology (~1 wks)</p>	<p>Cyber Security and CIA Triad, basic cyber threats to CIA, cyber-attack surfaces, recent cyber-security incidents and their high-level analysis</p>	<p>Example Driven Lectures with examples drawn from most recent incidents</p>	<p>None</p>
<p><b>Module 2:</b> Basic Cryptography (1 - 2 wks)</p>	<p>Role of Cryptography in ensuring confidentiality for data at rest, data in motion, and data in process.</p> <p>Symmetric and Asymmetric Cryptography, their needs as complementary of each other, some basic symmetric and asymmetric algorithm outlines (RSA, DH, DES, AES)</p> <p>Role of cryptography in data integrity, non-repudiation</p> <p>Hashing and Digital Signature and some example hash function outlines (MD5, SHA-256), understanding digital signature and its role.</p> <p>Digital Certificate and PKI.</p> <p>Importance of the role of a proper Pseudo Random Number Generator</p>	<p>Provide good intuition than nitty gritty of algorithms, or going through ciphers in details. The intuition behind PKI, Digital Certificate.</p>	<p>Using library functions to use RSA, AES, SHA-256 and show the result of encryption, Hashing etc.</p> <p>Taking apart a digital certificate and show the various components and their significance</p>
<p><b>Module 3:</b> Authentication, Authorization and Privilege (1 week)</p>	<p>Importance of strong Authentication, distinction between authorization and authentication, importance of authorization, access control, Mandatory and Discretionary Access control, role based</p>	<p>Intuition of distinguishing between authorization from authentication, access control lists, MAC vs</p>	<p>Lab on 2 factor authentication,</p> <p>Lab on privilege escalation example</p>

	authorization, privilege and privilege escalation	DAC with examples, importance of distinct privileges, principle of least privilege, show example of privilege escalation	
<b>Module 4:</b> Application Security (4-5 weeks)	<p>Basic application vulnerabilities (Buffer overflow, Integer Overflow, format string vulnerability) , Basic mitigations of buffer overflow – platform bases, compiler based, secure programming practice</p> <p>Web Client Security, Same Origin Principle, DOM, Java Script Vulnerability, Cookies and Cookie attributes Secure, httponly, Concept of session and session ID, Session hijacking vulnerability, http vs. https and SSL/TLS and version issue</p> <p>Web Server Security – XSS, CSRF, SQL Injection, Command Injection concepts, examples of each and mitigation techniques</p> <p>Vulnerabilities in DNS, Routing and IP protocols especially in IPv4 and suggested remedies with DNSSEC, S-BGP, and IPSec</p>	Intuition of why these vulnerabilities happen, and how various mitigation techniques have been developed, why the mitigation techniques are not enough and can be escaped by determined attackers, and examples from real attacks	<p>Lab1: Buffer overflow, integer overflow and format string vulnerability testing in vulnerable applications</p> <p>Lab 2: DVWA based command injection. SQL injection, XSS and CSRF</p>
<b>Module 5:</b> Perimeter protection and Intrusion Detection (2 weeks)	<p>Host Intrusion Detection techniques, what are the indicators to look for and how an SIEM tool can consolidate such indicators into a management console</p> <p>Network Intrusion Detection – signature based vs. behavior based, Snort</p>	Intuition about indicators that could indicate a host as compromised, and how multiple hosts/endpoint s can be monitored	<p>Lab1: Students are asked to install Wazuh and monitor a host</p> <p>Lab 2: Students are asked to install snort and monitor a</p>

	Firewall vs. Intrusion Detection tool, Firewall rules and customization techniques	Intuition behind signature vs. behavior based network monitoring and detection of intrusion  Hands on with Snort installation	network on their local network
<b>Module 6:</b> Basic Malware Analysis (1 week)	Various malware classes and their characteristics  Difference between static analysis and dynamic analysis  Signature vs. behavioral detection techniques	Intuition about various malware classes and their modus operandi, discuss why static and dynamic analysis complement each other, what kind of information is obtainable from static vs. dynamic analysis, demonstrate some static analysis tools	Lab 1: US static analysis tools to find how an executable can be analyzed.

**Detailed Contents for Desirable Learning Outcomes (optional, <= 3 modules):**

<b>Module</b>	<b>Topics</b>	<b>Pedagogy teaching suggestions</b>	<b>Nature of lab / assignment / practice</b>
<b>Module 7:</b> Mobile Application Security	Basic mobile attack surface and the ideas of permissions, and their abuse	Provide intuition on mobile malware and how they work, give example of mobile malware attacks,	None

(1 week)	Execution model of mobile apps in Android (Sandboxing) and communication	provide intuition of execution model of Android and demonstrate Mandatory Access Control idea in action, SE Linux being part of Android	
<b>Module 8:</b> WLAN Security (1 week)	Some common ways WLAN are compromised including weak cipher such as WEP, evil twin attack, unauthorized access point based attacks (rogue wlan) etc	Provide students idea about how to look for signs of these rogue wlangs, evil twins, public wifi etc.	None

**Suggested text books / Online lectures or tutorials:**

1. Ross J. Anderson, Security Engineering, Third Edition, Wiley, Nov 2020

**Suggested reference books / Online resources:**

1. The Web Application Hacker's Handbook: Finding and Exploiting Security Flaws 2nd Edition by D Stuttard and M Pinto
2. Cryptography and Network Security by William Stallings.
3. The Hacker Playbook: Practical Guide To Penetration Testing (vol. 1 and 2) by Peter Kim.

**Prepared by:** Sandeep K. Shukla (IITK), Vinod Ganapathy (IISc), Sambuddho Chakravarty (IIITD), S. Venkatesan (IIITA).

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly: CS1, CS3, CS6***

***Other discipline GAs to which this course may contribute somewhat:***

---

## Core Courses (Additional)

There are two additional courses, which are not part of the essential core but can be part of an extended-core. In other words, they may be offered as core courses, or as electives that are offered regularly – the choice is for the HEI to decide.

### Compiler Construction

CSXXX	Compiler Construction	3L:0T: 4P	5 credits	Pre-Reqs: Data Structures
-------	-----------------------	-----------	-----------	---------------------------

**Prerequisite:** C/C++/Java programming language. Data structures and algorithms. Automata theory.

#### Learning Outcomes, Essential (<=6):

1. To understand the role, functionality and structure of program translation and interpretation in software development.
2. To understand the difference between abstraction levels of a high level language and a machine language.
3. To understand the role of a sequence of intermediate representations in lowering the level of abstractions in the process of language translation.
4. To get a firsthand experience of a practical application of elegant data structures, algorithms, and other core CS concepts such as automata theory.
5. To make effective use of tools such as lex and yacc.
6. To become a much better programmer by appreciating all that happens behind the scenes in making an HLL program run.

#### Desirable/Advanced (<= 3):

1. To understand some of the critical aspects of machine code generation.
2. To understand the issues in efficient code generation.

#### Syllabus for Essential Learning Outcomes:

Module (appx dur in wks)	Topics	Pedagogy / teaching suggestions	Nature of lab / assignment / practice
<b>Module 1:</b> Introduction to Compilers (1 week)	Comparing abstractions of a high level language and a low level language; compilation as a series of steps for lowering the abstraction level through stepwise refinement; phases of compilation;	Sections 1.1 and 1.2 from the textbook, online resources [2] and [3].	Viewing the intermediate representations and the final assembly code generated by GCC/LLVM, relating them to the input program

	bootstrapping; cross-compilation.		Programming Assignment #A0.
<b>Module 2:</b> Lexical Analysis (2 weeks)	The role of lexical analysis; Token, lexemes, and token codes; Regular Expressions (RE) to represent tokens, Deterministic finite automata (DFA), Traversing a DFA for recognising tokens; Generating a lexical analyzer using lex/Flex.	Sections 3.1, 3.2, 3.4.4, 3.5, 3.6, 3.8.1, 3.8.3 from the textbook.	Writing lex specifications and generating tokens for a given language,  Programming Assignment #A1.
<b>Module 3:</b> Syntax Analysis (3 weeks)	Context Free Grammars (CFG), Concept of parsing, sentences and sentential forms, leftmost and rightmost derivations, parse trees, ambiguous grammars; Overview of top-down and bottom-up parsing;  Option1: Introduction to shift reduce parsing; viable prefixes and valid items, Constructing LR(0) sets of items; Constructing SLR parsing tables; Generating a parser using a parser generator such as Yacc/Bison.  Option 2: Top-down parsing, Left factoring, Elimination of left-recursion, predictive parsing, recursive descent parsing, LL(1) parsing. Generating a parser using a parser generator such as ANTLR, JavaCC, etc.	This module is driven by the chosen parser generator. If Yacc/Bison is chosen, the module should cover bottom up parsing. If Antlr or JavaCC is used, the module should cover top down parsing.  Sections 4.1, 4.2, 4.3 from the textbook.  Option1: Sections 4.5, 4.6 from the textbook.  Option2: Section 4.4 from the textbook.	Writing yacc specifications, generating a parser by using the scanner generated in module 2.  Precise error reporting using yytext and yylينو.  Using a command line switch to optionally print token details.  Programming assignment #A2.
<b>Module 4:</b> Semantic Analysis (1 weeks)	The need of semantic analysis; abstract syntax trees for expressions, assignment statements and control flow statements;	Sections 5.1, 5.2, 5.3, 5.4, and 6.1 from the textbook.	Writing a type checker to ensure that a syntactically correct MMC program is type-safe.

	attribute evaluation, syntax directed translation schemes (SDTS);		Programming Assignment #A3.
<b>Module 5:</b> Applications of Semantic Analysis (3 weeks)	Applications of SDTS for (a) declaration processing and type checking, (b) generating three-address code	Sections 6.2, 6.3, 6.4, 6.5.1, 6.5.2, 6.9 from the textbook.	Write a translator to translate a type-checked MMC program to equivalent three-address code.  Programming Assignment #A4.
<b>Module 6:</b> Run time support (1 week)	Parameter passing by value, reference, and name; activation records, stack and static allocation of activation records; translating a function call, allocating offsets to variables, generating code for function prologue, function epilogue, call sequence, and return sequence.	Sections 1.6.6, 7.1, 7.2.2, 7.2.3 from the textbook, Chapter 6.1 from the reference book and online material [4].	Not Applicable
<b>Module 7:</b> Introduction to Code Optimization (1 week)	Control flow graphs; Local optimizations (common subexpression, copy propagation, dead code elimination).	The focus of this module is on knowing the optimizations and not on the techniques of performing them.  The desirable part would need an additional one week.  Section 9.1 from the textbook..	Not Applicable
<b>Module 8:</b> Code Generation (2 weeks)	Generating assembly code from three address codes using simple register allocation and instruction selection.	By simple register allocation, we mean that all the values of temporaries are held in registers across three address code statements but	



		not those of source variables.  Sections 8.1, 8.2, 8.6 in the textbook.	
--	--	---	--

### Syllabus for Desirable/Advanced topics.

Module (appx dur in wks)	Topics	Pedagogy / teaching suggestions	Nature of lab / assignment / practice
<b>Module 2:</b> Lexical Analysis (2 weeks)	4-arrays representation, observing the data structures in the scanner generated by lex.	Section 3.9.8 from the textbook.	
<b>Module 3:</b> Syntax Analysis (3 weeks)	LR(1) and LALR(1) parsing and the option not chosen in the default offering	Section 4.7 from the textbook.	Error recovery using the error token in yacc and emitting meaningful error messages.
<b>Module 7:</b> Introduction to Code Optimization (1 week)	Global optimization (constant propagation, common subexpression elimination, copy propagation, dead code elimination, strength reduction)	Sections 9.2.5, 9.2.6, 9.4.1, 9.4.2, 9.4.3, 9.4.6, in the textbook.	
<b>Module 8:</b> Code Generation (2 weeks)	Register allocation using graph colouring, Optimal code generation for expression trees, Sethi Ullman algorithm, Aho Johnson algorithm.	Sections 8.8.1, 8.8.4, 8.10, 8.11 in the textbook.	Write a translator to translate code in 3-address-code form to assembly code.  Programming Assignment #A5

### Lab/Assignment Details

- A0: Write 5 simple test-cases in MMC and then inspect the generated code.
- A1: Write a lexer to recognize valid tokens.
- A2: Write a parser to parse the given input MMC program.
- A3: Write a type-checker for a syntactically correct input MMC program.

---

A4: Write a Translator that takes a type-checked MMC program and generates equivalent IR code in TACoC format; TACoC is a subset of MMC such that code is in a form similar to three address code. Details of TACoC is given towards the end of this section.

A5: Generate MIPS code. Use the SPIM simulator to run the code.

### **Suggested Pedagogy for the lab/assignments:**

1. The instructor is expected to prepare a set of test-cases for assignments A1-A5.
2. For A1, A2, and A3 – the test cases should include some test-cases which throw lexical, syntactical and type errors, respectively.
3. For input test-case design the expected output and write a script to verify the output generated by the student assignment.
4. For each assignment, the student should use a script to build their code (using a Makefile/ant build script/shell script). The instructor must create some sample scripts and teach the same to the students.
5. The instructor should create a reference implementation for A1-A5.
6. The implementation can be done in Flex+Bison or Antlr or JavaCC.

### **Details of MMC and TACoC**

MMC is a simple variant of C, and has many simplifications. This along with the simplified nature of TACoC ensures that the compiler can be written within a short span of time. Further details about this can be obtained from:

<http://www.cse.iitm.ac.in/~krishna/aicte-compiler-design/Lab-Syllabus.docx>

### **Texts and Other Material**

1. [Text book] Aho, Lam, Sethi, and Ullman. Compilers: Principles, Techniques, and Tools. 2/e, Addison-Wesley, 2006.
2. [Reference book] Andrew Appel and Jens Palsberg. Modern Compiler Implementation in Java. 2/e, Cambridge University Press, 2002.
3. [https://en.wikipedia.org/wiki/Cross\\_compiler](https://en.wikipedia.org/wiki/Cross_compiler).
4. [https://en.wikipedia.org/wiki/Bootstrapping\\_\(compilers\)](https://en.wikipedia.org/wiki/Bootstrapping_(compilers))
5. [https://en.wikipedia.org/wiki/Function\\_prologue\\_and\\_epilogue](https://en.wikipedia.org/wiki/Function_prologue_and_epilogue)

## Theory of Computation

<b>CSXXX</b>	<b>Theory of Computation</b>	<b>3L:1T: 0P</b>	<b>4 credits</b>	<b>Pre-Reqs: Discrete Maths, Data Structures</b>
--------------	------------------------------	------------------	------------------	--

### Essential Learning Objectives:

1. Understand models and abstractions: automata as a basic model of computation
2. Link between languages, automata, and decision problems.
3. How to build new models from old ones: product, union, closure properties.
4. Argue about limitations of computational models.
5. Understand algebraic formalisms of languages such as regular expressions, context-free grammar.
6. Understand algorithms and computability through the lens of Turing machines.
7. Existence of unsolvable problems and what that means.
8. Relations between the various computational models.

### Advanced/Desirable Learning Outcomes

Nil

<b>Module (appx dur in wks)</b>	<b>Topics</b>	<b>Teaching Suggestions</b>	<b>Learning outcomes</b>
<b>Module 1:</b> Finite Automaton (4-5 wks)	<p>-Why automata theory?</p> <p>-Alphabets, formal languages, and problems.</p> <p>-What are regular languages and automata models for them: Deterministic Finite automaton, Formal argument of correctness, Regular languages</p> <p>-Properties of regular languages-Closure, properties, product construction</p> <p>-Limitations of Automata Non-regularity, Pumping Lemma</p>	<p>-Sec 1.1 of T2</p> <p>-Sec 1.2, 1.5 of T2</p> <p>-Sec 1.1 of T1</p> <p>-Sec 1.1 of T1</p> <p>-Sec 1.4 of T1</p> <p>-Sec 1.2 of T1</p>	<p>F. Familiarity with notations.</p> <p>U. Give examples of languages, regular languages.</p> <p>U. Design finite automata, both deterministic and nondeterministic for a given language.</p> <p>R. Write formal proof of correctness of a DFA</p> <p>U. Give examples of non-regular languages and prove that language is non-regular using pumping lemma</p>

	<p>-Non-deterministic Finite Automaton, Subset construction, Equivalence with DFAs.</p> <p>-Regular expressions. Equivalence with regular languages.</p> <p>-Algorithms for regular languages, Minimization and its algorithm.</p> <p>-(suggested) Myhill-Nerode relations, Characterization of regular languages</p>	<p>-Sec 1.3 of T1</p> <p>-Sec 4.3, 4.4 of T2</p> <p>-Lecture 15,16 of R1</p> <p>+applications of automata to text search and NLP</p> <p>+applications of regular expressions for text search in UNIX.</p> <p><b>Advanced Topics:</b></p> <p>- 2DFAs, Equivalence with DFAs using Myhill-Nerode Relations (Lecture 17,18 of R1)</p>	<p>F. Understand the difference between determinism and nondeterminism</p> <p>U. Use closure properties to show non-regularity</p> <p>U. Design regular expressions</p> <p>U. Use the minimization algorithm to minimize a given DFA</p> <p>U. (suggested) Apply Myhill-Nerode Theorem to show that a language is regular or non-regular</p>
--	---	--	--

<p><b>Module 2:</b> Grammars, Context-free Languages and machine models.  (4-5 wks)</p>	<p>-Grammars and the motivation from language theory.</p> <p>-Context-free grammars, closure properties. Chomsky Normal Form for CFGs.</p> <p>-PDAs. Empty-stack vs Final state acceptance conditions. Equivalence of PDAs and CFGs.</p> <p>-Limitations of PDA computation, non context-free language. Pumping Lemma for CFLs.</p> <p>-Deterministic CFLs and PDAs.</p> <p>-(suggested) CYK Algorithm for parsing of CFLs.</p>	<p>-Sec 2.1 of T1</p> <p>-Sec 2.1 of T1</p> <p>-Sec 2.2 of T1</p> <p>-Sec 2.3 of T1</p> <p>-Sec 2.4 of T1</p> <p>-Sec 7.4 of T2</p> <p>+applications to parsers and compilers.</p> <p><b>Advanced Topics:</b></p> <p>- Ogden's Lemma.</p>	<p>U. Design CFGs and PDAs for CFLs</p> <p>R. Prove correctness of CFGs</p> <p>F. Understand that regular languages are a subset of CFLs.</p> <p>R. Prove equivalence of CFGs and PDAs</p> <p>U. Argue a language is non-CFL using pumping lemma</p> <p>F. Familiarity with DPDAs</p> <p>U.(suggested) Construction of DPDAs</p> <p>U. (suggested) Parsing using CYK algorithm</p>
---	---	---	--

<p><b>Module 3:</b> Turing machines and Computability, (4-5 wks)</p>	<p>-Modeling computation using Turing Machines. Equivalent models. Church Turing Hypothesis.</p> <p>-Decidability and Turing recognizability (i.e., recursive and recursively enumerable). Closure properties.</p> <p>-Undecidability by diagonalization.</p> <p>-Reductions to show undecidability. Examples of reductions.</p> <p>-Resource bounded Turing machines &amp; Intro to Complexity. Basic complexity classes. Time bounded classes: P, NP, EXP.</p> <p>-(suggested) Post's correspondence problem and other undecidable problems</p> <p>-(suggested) Polytime reductions, NP-completeness, Cook-Levin Theorem without proof</p>	<p>-Sec 3.1, 3.2, 3.3 of T1</p> <p>-Sec 4.1 of T1</p> <p>-Sec 4.2 of T1</p> <p>-Sec 9.3 of T2. Sec 5.1, 5.3 of T1.</p> <p>-Sec 7.1 of T1</p> <p>-Sec 5.2 of T1</p> <p>-Sec 7.3, 7.4, 7.5</p> <p><b>Advanced Topics:</b></p> <p>- Rice's Theorem</p> <p>- Space bounded computations and complexity, PSPACE</p>	<p>F. Understand relation between the various classes such as decidable, Turing recognizable., co-Turing recognizable.</p> <p>F. Give examples of decidable languages, undecidable languages, Turing recognizable languages.</p> <p>U. Prove a language is undecidable by reduction from a known undecidable problem</p> <p>F. Relation between basic complexity classes</p> <p>F. (suggested) Scenarios in which the reductions are used</p> <p>R. (suggested) Proving languages are NP-complete using reductions</p>
--	--	--	--

---

**Notations:**

- Topic Categorization:  
*Compulsory* - Topics that should be covered.  
*Suggested* - Optional topics that the instructor can choose from given availability of time.  
*Advanced* - Advanced topics in each module that an instructor can teach depending on the interest of the class.
- + indicates applications that could be mentioned in the class.
- Learning Outcome Categorization:  
*Familiarity* - Student should be able to identify and comprehend *what* the topic is about. This corresponds to the cognitive levels of *knowledge* and *comprehension* of Bloom's taxonomy (see e.g., [https://en.wikipedia.org/wiki/Bloom's\\_taxonomy](https://en.wikipedia.org/wiki/Bloom's_taxonomy)).  
*Usability* - Student should be able to understand *how* a particular idea/topic can be used, to solve problems, design examples, etc. This corresponds to the cognitive levels of *application* and *synthesis* of Bloom's taxonomy.  
*Reasoning* - Student should have a deeper understanding of a particular concept and *why* it works. This corresponds to the cognitive levels of *analysis* and *synthesis* of Bloom's taxonomy.

**Nature of lab / assignment / practice / tutorial:**

1. Make assignments using the books. To test:  
- what was done in the class  
- whether the student can think and apply the concepts.
2. Tutorials: Weekly problem-solving sessions.

**Suggested textbooks:**

1. Introduction to the Theory of Computation, 3rd edition. Michael Sipser, Cengage Publications (Low-cost Indian edition available).
2. Introduction to Automata, Theory, Languages and Computation. Third Edition. John Hopcroft, Rajeev Motwani, Jeffrey D. Ullmann, Pearson Publications (Low-cost Indian edition available).

**Additional Reference Material:**

1. Automata and Computability, Dexter C. Kozen. Part of the Undergraduate Texts in Computer Science book series (UTCS), Springer.
2. Elements of the Theory of Computation, 2nd edition. Harry Lewis, Christos Papadimitriou, Prentice Hall.

---

**Prepared by:** Jayalal Sarma, IIT Madras. S Akshay, IIT Bombay. Raghunath Tewari, IIT Kanpur.

***The Discipline Graduate Attributes (GAs) to which this course contributes significantly:*** CS4

***Other discipline GAs to which this course may contribute somewhat:*** CS2



---

## Professional Electives and Micro Specializations

Besides the core courses, programs normally have professional elective courses. Each HEI decides the electives it can or wishes to offer. We suggest that these electives be included:

- The key courses of some of the micro specializations (discussed below)
- A course on Advanced Mobile Communications

A *micro-specialization* is to provide a limited specialization in some sub-area of CSE, by offering suitable electives. The goal of the micro specialization is to provide deeper understanding and skill development in that area, and can provide multiple pathways to students, as different students can graduate with different specializations (or not). The areas in which micro specialization are offered should be aligned to industry careers or higher studies.

A micro specialization for CSE is defined as follows:

- It has a core course as the head (starting) course for the micro specialization
- It has a clearly defined goal, and learning outcomes for the goal
- It can have 2 +/- 0.5 additional courses (besides the head course) in the sub-area aligned to the goal. These courses can be full course (4-credits) or half-course (2 credit), and can be taken as electives by students (or extra credits.)

Based on the discussions and inputs, a few desired micro-specializations were identified. For a subset of these, a possible design of the micro-specialization is provided below. It should be pointed out that the list of courses in a specialization is illustrative – an Institution can replace or add a course aligned to the micro specialization goal. Institutions can also define a set of courses for a micro specialization and require that a subset be taken, with perhaps one being compulsory.

It should be added that HEIs are completely free to decide whether to offer micro specializations or not, and if they decide to offer, which areas to provide the specialization in. How the micro specialization is to be reflected in a student's records/certificates is also to be decided entirely by HEIs based on their policies and practices.

Note that the list of additional courses mentioned in the micro specializations also provide a list of suggested professional electives. However, there can be electives which are not a part of any specialization.

---

## Software Engineering Micro Specialization

**Goal:** The aim of this specialization is to provide an understanding of the importance and role of software engineering, an understanding of the important sub-areas of software engineering (SE), and to provide a deeper understanding of some of them.

### Learning outcomes of the specialization:

1. Understand the role and importance of software engineering in developing industrial strength software, and the important tasks involved in engineering such software.
2. Ability to apply proper SE practices to develop in a team a working software system to solve some users' problem.
3. Understand and apply concepts in some sub-areas of SE like testing, maintenance, open source software, model based development, requirements engineering, etc.

**Base Course for this specialization (core course):** Advance Programming

### Courses (Electives) Proposed for the Specialization:

(The courses for a micro specialization can be full course (4-credits) or half-course (2 credit). The total number of courses suggested is 2 +/- 0.5.) For this specialization, it is suggested that the first course should be required, followed by 2 half courses in some sub-areas of SE.

Course Name (full / half)	Purpose/Goal	Topics
Software Engineering (full)	To explain the iterative software development process and different aspects of it (e.g. short cycles, agile approaches, test-driven-development, etc), and to apply the concepts to a team project to develop software	Iterative software development process, basic project planning for such a process, requirements including user interface, architecture and design, coding using modern IDEs, testing, integration and deployment
Software testing (half)	Go deeper in software testing, including some modern tools	Unit testing and the test frameworks, test case design (incl. black box and white box), non-functional testing, test automation and tools, related metrics (e.g. coverage, performance, e.g.)
Open Source Software (half)	To expose to students open source software practices and ecosystem, and to expose them to using them	OSS evolution, current situation, general OSS practices, common OSS platforms and how to use them, small project on some OSS platform (e.g. GitHub)

---

Software Maintenance (half)	To explain the software maintenance cycle and different aspects of it, and develop capabilities in them	Defect cycle and bug management, regression testing, change management, refactoring, software evolution, related metrics, etc.
-----------------------------	---	--

**Prepared by:** Pankaj Jalote (IIIT-Delhi), Raghu Reddy (IIIT Hyderabad), Vinay Kulkarni (TCS), Meenakshi Dsouza (IIIT Bangalore)

---

## Machine Learning Micro Specialization

**Goal:** The aim of this specialization is to provide an understanding of the importance of machine learning in computer science, expose the advances in machine learning, and also to demonstrate the practical role of machine learning in related domains in artificial intelligence. This specialization also aims at providing deeper understanding of some of the sub areas, deeper analytical concepts, and the recent advances.

### Learning outcomes of the specialization:

1. Appreciate the role of machine learning in computer science and data-driven problem solving.
2. Appreciate the role of recent and advanced machine learning algorithms and formulations in solving problems of practical importance. Expose students to principled ways of practising machine learning in many unstructured real world problem settings.
3. Ability to apply principles of ML to develop practical solutions to problems seen around in industry, society and research, focusing on problem formulation, solution design, implementation and experiencing the empirical design process.
4. Providing practical experience in use of popular machine learning libraries and implementations. Get the students trained in the empirical science behind the design of machine learning based solutions.

### Base Course for this specialization (core course):

1. Machine Learning

### Courses (Electives) Proposed for the Specialization:

*Note: The courses for the specialization can be full course (4-credits) or half-course (2 credit). The total number of courses suggested should be 2 +/- 0.5.*

Course Name (full / half)	Purpose/Goal	Topics
Advanced Machine Learning (full)	<ul style="list-style-type: none"><li>- To introduce and be familiar with the popular deep learning architectures.</li><li>- To expose the popular problem formulations beyond simple supervised learning used in situations of data scarcity, domain/distribution shift etc. using the popular ML/DL algorithms.</li></ul>	<ul style="list-style-type: none"><li>- CNNs, RNNs, Auto Encoders, Loss Functions and Training.</li><li>- Transfer learning, domain adaptation, semi-supervised learning, active learning, self-supervised learning, incremental learning, few-shot learning</li><li>Laboratory experiments to appreciate the utility</li></ul>

Computer Vision or Natural Language Processing (Full)	<ul style="list-style-type: none"> <li>- To introduce a sub-area of Artificial Intelligence (perception) where machine learning formulations get extensively used.</li> <li>- Appreciate how ML has influenced the problems in this area, and how the type of data/problems in this area demanded newer solutions in ML.</li> </ul>	<ul style="list-style-type: none"> <li>- Formulations based on classification, regression, structured prediction in sub-tasks of CV/NLP.</li> <li>- Role of data and learning in Feature/Embedding.</li> <li>- Laboratory experiments to appreciate the utility</li> <li>- Encourage Course Project</li> </ul>
Programming for Machine Learning (half; Lab Course)	<ul style="list-style-type: none"> <li>- To expose students to popular ML libraries and frameworks such as PyTorch and Tensorflow). To develop empirical skills in design and implementation of ML solutions</li> <li>- To strengthen the hands-on skill of the student.</li> </ul>	<ul style="list-style-type: none"> <li>- Programming in popular libraries, Use of Cloud, GPUs, Debugging/Visualization. Working on large datasets.</li> <li>- Programming intensive. Use of popular (in industry, research) frameworks and tools expected.</li> </ul>

**Notes:**

1. To strengthen the micro specialization, students may be encouraged to do the BTech Project (or Internship or a summer project) in this area if the regulations (credit system) allow.
2. Second course may be adapted to the strength of the department as an "Applied Machine Learning Course " in an area of expertise in the department. However, the focus is on the use of ML formulations in this course.

**Prepared by:** Prof. C. V. Jawahar (IIIT Hyderabad), Prof. P. J. Narayanan (IIIT Hyderabad)

---

## Distributed and Cloud Systems Micro Specialization

**Goal:** The aim this specialization is to provide: an understanding of the importance and role of distributed computing and data systems, an understanding of the important sub-areas of these systems such as cloud computing and big data platforms, a deeper understanding of their concepts, and the practical aspects of building applications for such systems. The aim of this specialization is to introduce important topics and themes of distributed systems and cloud computing, both from a theoretical and a practical perspective.

### Learning outcomes of the specialization:

1. Understand the need for and the models of distributed systems, and ability to access remote applications and data over the network.
2. Understand key concepts for building scalable, reliable and consistent distributed applications.
3. Understand the design, implementation, deployment and use of distributed computing and data platforms.
4. Understand the models of Cloud computing, enabling technologies and how to build Cloud native applications

### Base Course for this specialization (core course):

Computer Networks and/or Operating Systems

### Courses (Electives) Proposed for the Specialization:

The courses for the specialization can be full course (4-credits) or half-course (2 credit). The total number of courses suggested should be 2 +/- 0.5.

Course Name (full / half)	Purpose/Goal	Topics
Principles of Distributed Systems (Half or Full)+	Understand the need for and the models of distributed systems, and ability to access remote applications and data over the network.  Understand key concepts for building scalable, reliable and consistent distributed applications.	Need for distributed systems, models for coordination among multiple machines over the network. Models of distributed systems, remote invocation and remote storage.  Concepts of performance, scalability, reliability, consistency and correctness.
Building Cloud and Big Data Applications (Full)	Understand the models of Cloud computing, enabling technologies and how to build Cloud native applications.	Designing distributed applications/algorithms, Data-Intensive Computing and Data-Oriented Programming. Distributed execution and runtimes. Streaming data management and processing.

	Understand the design and use of distributed computing and data platforms.	Remote and Distributed file systems, Key Value Storage and NoSQL Columnar Store.  Resource abstraction and Service oriented architecture, Computing Abstractions on the Cloud, and Building Cloud-native applications.
Cloud Systems Engineering (Half or Full)+	Understand the models of Cloud computing, enabling technologies and how to build Cloud native applications.	Cloud services models, Data center architecture and management. Virtualization and Containerization. Big data computing platforms, REST based web services.
Mini-Project/Hands-on Cloud Computing/Big Data (Half)*	Understand the design, implementation, deployment and use of cloud systems to help build distributed, scalable and reliable applications.	Hands-on mini programming project that makes use of the concepts and technologies learnt in the theory courses.

**Notes:**

1. Either separate or as a lab-component of one of the other theory courses
2. The course can be offered as half or full course depending on faculty availability

**Pedagogical References**

1. Cristina L. Abad, Eduardo Ortiz-Holguin, and Edwin F. Boza. 2021. Have We Reached Consensus? An Analysis of Distributed Systems Syllabi. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE '21)*. Association for Computing Machinery, New York, NY, USA, 1082–1088. DOI:<https://doi.org/10.1145/3408877.3432409>
2. Joshua Adams, Brian Hainey, Laurie White, Derek Foster, Narine Hall, Mark Hills, Sara Hooshangi, Karthik Kuber, Sajid Nazir, Majd Sakr, Lee Stott, and Carmen Taglienti. 2020. Cloud Computing Curriculum: Developing Exemplar Modules for General Course Inclusion. In *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education (ITiCSE-WGR '20)*. Association for Computing Machinery, New York, NY, USA, 151–172. DOI:<https://doi.org/10.1145/3437800.3439206>  
DOI:<https://doi.org/10.1145/3304221.3325536>

**Prepared by:** Yogesh Simmhan (IISc Bangalore), Purushottam Kulkarni (IIT Bombay)

---

## Human Computer Interaction (HCI) Micro Specialization

**Goal:** The aim of this specialization is to provide an understanding of the importance and role of user-centered design and associated techniques.

### Learning outcomes of the specialization:

1. Understand the user-centered design process and associated techniques to use it for designing and developing computing based solutions
2. Develop an ability to start with an observable problem in real-life settings and to develop a technology-led solution for the same.
3. To develop soft skills such as empathy for societal problems, the ability to think from another perspective, learning to work in a group
4. To learn methods related to observation, interviewing, problem identification, ideation, prototyping, and evaluation.
5. To apply the design process in a hands-on scenario.

**Base Course for this specialization (core course):** Advanced Programming

### Courses (Electives) Proposed for the Specialization:

Course Name (full / half)	Purpose/Goal	Topics
Introduction to HCI (full)	To provide an understanding of the importance and role of user-centered design and techniques like observation, interviewing, problem identification, ideation, prototyping, and evaluation.	Contextual Inquiry, Interviews, Surveys, Focus Groups, sketching, low-fidelity prototyping, usability evaluation
HCI semester project (half)	Give an opportunity to apply the concepts learnt in the theory course in a hands-on project	Students will work on a project from real-life-like settings should be chosen and a group of students should be guided through the design process
Interactive Systems (half)	To develop a new interaction technique to solve a specific problem in HCI	Using any of the modern-day technologies, like, Virtual Reality, Augmented Reality, Speech technologies, Mobile programming, Drones, Haptic systems, Eye-



---

		trackers, etc. students will develop and test new interaction modalities.
--	--	---

**Notes:**

1. If there is faculty to offer further courses, then one/two other half courses can be added to this specialization – e.g. HCI for Development (HCI4D), Inclusive Design, Learning Technologies,...

**Prepared by:** Pushpendra Singh (IIT-Delhi), Anirudha Joshi (IIT Bombay)

---

# Appendix: Recommendations for Using Online Content in Courses at Colleges and Universities

## Background

A lot of online high quality content is available today either free or at a low cost. Besides the government supported NPTEL, we have companies like Coursera and EdX who aggregate courses from several universities (and even private commercial organizations) and offer them to students all over the world. There are other companies like Unacademy who offer courses designed and developed by them. Further, a lot of companies have online content available to students. These include IBM, Microsoft, Amazon, Infosys, Google, LinkedIn, Cisco, and so on.

On the other hand, most of the Computer Science (and related) departments face serious shortage of faculty, particularly in areas where there is a significant demand in the industry.

So, on one hand, we have quality content available for free or low cost, and on the other hand, we don't have faculty to teach such courses. The natural solution is to find ways to use online content for the courses in the curriculum (with credits). This way, either a knowledgeable faculty can "teach" a much larger class, or a faculty member with inadequate background in the topic can still "teach" the course better than what s/he would have done without the support of such online content.

## Issues

The online content has been around for several years. The regulatory bodies have also been encouraging use of such content (particularly, NPTEL). And yet, the online content has not been integrated with the curriculum in most colleges. The pandemic has allowed people to take a fresh look at online content and the mental barriers to using such content in the curriculum have been breached. At this time, it is felt that a lot of colleges would want to use this content. However, there are two primary academic issues that need to be addressed (besides logistics, financial, and HR issues). These are:

1. How do we decide what material to use. This has two sub-issues. One, what content would be equivalent to the content that is mentioned in our curriculum (course mapping). Note this is a challenge because there may not be a single module which covers all parts of the curriculum. So we may need to select more than one module. Two, given the plethora of content, which content is of reasonable quality.
2. How would a college do evaluation of students in order to assign marks/grades.

## Modes of Learning

There are several ways of using the online content in the curriculum.

The simplest mode (Mode 1) is to use online content as additional reference material. In this mode, the normal teaching is anyway being done and the faculty is referring to online content in the same way he/she would refer to a book. This mode requires no change or suggestions and hence is outside the scope of this document.

The next mode (Mode 2) is the flipped classroom model where the students go through the online content (including writing programs, small quizzes, etc.), and the local faculty takes discussion sessions and does all the evaluation. This mode still requires a knowledgeable faculty member to be the instructor. However, given that we can now reduce the contact hours for the students and faculty, the same faculty can handle a much larger class (or multiple sections in case there is an upper limit on the size of the class).

The next mode (Mode 3) is that the students are studying only through the online mode and there is very little interaction with the faculty at the college. May be there can be some sessions once in a while, but mostly, the role of a local faculty is only to handle evaluation (exams, assignments, projects, etc.). This partially addresses the issue of lack of faculty in certain areas since the expertise required for handling evaluation is arguably lesser than the expertise required to teach the course.

The last mode (Mode 4) is where the online provider does everything, including evaluations. We don't need any faculty member at the college to offer this course. We only need to decide how to translate the evaluation done by an outsider to an equivalent grade/marks on the college transcript.

The table below illustrate the four models in brief.

<b>S.No.</b>	<b>Model</b>	<b>Local Part</b>	<b>Online Part</b>	<b>Guidelines</b>
1	Reference Mode	All lectures, homeworks, evaluations	Additional reference material	No additional recommendations.  This requires no additional details.
2	Flipped Mode	Discussion sessions, evaluations	Lectures, Homeworks, Quizzes, Project options,	Requires local faculty with some knowledge in the subject area  Large sections/class sizes can be handled
3	Screen Mode	Evaluations	All lectures, Homeworks, Projects	Place reasonable limit on the number of credits that can be earned in this mode  Use in cases where there is faculty shortage in critical areas
4	Fully online	--	All lectures, evaluations, Projects, Tests	Place reasonable limit on the number of credits that can be earned in this mode. Advised not to repeat this model for the same course more than two times.

---

				Local faculty also to register and go through the course. Official score/grade to be suitably added to the student transcript
--	--	--	--	--

## Course Mapping

Each course in the curriculum has course outcomes and the content defined by the university. Typically, we may have some course outcomes which are important ones, and may be some outcome which is desirable or optional. Similarly, the curriculum may also define the rough duration of each topic. Sometimes the curriculum will also include the kind of projects or assignments that the student shall be asked to do in the course.

The key to course mapping is to realize that any two individuals defining a course will have some differences and we must be flexible to accept differences to some extent. The committee trying to do course mapping should have an understanding of what are the important components (in all three: outcomes, topics, and projects) and should ensure that the online content meets all the important requirements. But looking for identical course would be usually futile.

Also, what is noticed is that typical online courses are often available in smaller modules than a typical 4-credit course in our curriculum. And therefore, one may need to consider more than one online courses together to be equivalent to one course in the college. It may also be noted that just like some small aspects of the course may not be present in the modules chosen, there may be some aspects of the online modules which were not part of the college course. This amount of flexibility should be acceptable to the college. Typically, if the online content covers 80% of the college course, it may be accepted.

While autonomous colleges and universities do have this flexibility, the affiliated colleges may not have the flexibility of not teaching even 20% of the content. In case of affiliated colleges, one will have to either be stricter in course mapping, or find a way to cover the gaps through a local faculty or a visiting faculty. It is assumed that it is easier to find a visiting faculty for a small part of the course and hence it is still a useful mode.

It may also be noted here that for Core courses, the overlap needs to be significantly higher while for the elective courses, the overlap could be relaxed somewhat. It is because the core courses typically are pre-requisites for other courses. Also, core courses have been defined to be such because it is assumed that that content is more important for the graduates than what is taught in electives.

However, course mapping is still not an easy thing to do. It requires an understanding of important versus less important components, and quality of content among the plethora of options available. It may require someone to go through the content patiently.

And hence it is recommended that for standard courses recommended in the AICTE model curriculum, a course mapping may be suggested for the benefit of colleges and universities. A few example course mappings are, therefore, attached with this report.

---

## Evaluation

In Modes 1, 2, and 3, the complete evaluation is local, and hence there is no issue. In Mode 4, we need to consider an external evaluation and use that internally. This is a challenge. There is a difference in how to handle this in a university versus an affiliating college. In a university, a simple way of handling this would be to assign a Pass/Fail grade to the student. In case of an affiliating college, where only the internal marks need to be forwarded to the university, one could consider the external evaluation since it impacts only 30% of the marks, and the university will anyway have its own exam of 70% marks.

There are other models like normalizing the external evaluation to the college's internal policy or average distribution of marks in other courses.

It is assumed that in Mode 4, there is a formal way of communicating performance of the student by the provider of the online courses since colleges and universities will not accept screen shots, emails, etc.

## Other Issues

**Financial:** If there is a fee to be paid for online content, the college should have a policy on that. Typically, if the savings due to reduced faculty requirements are significant, then the fees for the online provider may be reimbursed by the college.

**Teaching Load:** Faculty member supporting the course whether by taking a few discussion sessions (in Mode 2) or by evaluating the students (in Mode 3) is still putting in substantial effort in managing the course, and an appropriate credit should be given to the faculty member when his/her teaching load is computed.

**Training the teacher:** When a course is being done in Mode 3, the local faculty member should also be expected to register for the course and go through the course (with load being appropriately counted). After a faculty member has gone through the course in two academic sessions, s/he would be well prepared to teach the course in the class in a much better way. Even if the course is being offered in Mode 4, there is no harm in asking a faculty member to register for the course and go through it. Some responsible person in the college would know the level of the course and what exactly students have done, and again, after two such sessions, the faculty member would be well prepared to teach the course. Hence this mode will also lead to better training of the teachers.

**Limits on Credits:** The committee believes that there should be a limit on the number of credits students can earn through online courses. In case of Mode 4, where even the evaluation is done by the online course provider, the proposed limit is 8 credits only. It is felt by the committee that the evaluation by online providers is still not fully trustworthy. As the technology for online exams or the processes for evaluation by online providers improve over a period of time, this limit may be increased. In case of Mode 3 where the content is delivered online but evaluation is local, the limit can be high. For Mode 3 and Mode 4 combined, the limit can be what the regulatory bodies like UGC have announced for online courses, which is currently 40% of the total credits. In Modes 1 and 2 where the online content is really the reference material, there is no need for any limit.

---

Another constraint the committee would want the colleges to consider is that in a sequence of courses in one stream of Computer Science, at least one course should be in class. For example, if we consider the sequence of systems courses – Operating Systems, Databases, Networks, Architecture, at least one course should be in class. This is to ensure that if there were some gaps in online courses, the faculty in the face to face class can try to cover that to some extent.

**Faculty Incentive:** There is a need to provide some incentive to faculty members who would manage the course that is being taught in the online mode. If a course is being taught in Mode 2, the load on the faculty is only marginally less than the load of teaching an in person course. So the full teaching load should be considered for the faculty. In Mode 3, the load is much less, and in Mode 4, the load is only that we are asking the faculty to also go through the course along with the students. In these two modes, the college may consider this as reduced load. However, their learning the course may be treated as equivalent to having done a Faculty Development Program when it comes to their appraisal and promotions.

**Consideration in NBA Accreditation:** One of the prime reason why online courses haven't become popular with colleges is that they must recruit faculty with a certain faculty to student ratio for accreditation and ranking. And once they have recruited faculty, one would always want the faculty to teach and not keep them under-loaded. If one can consider online courses as equivalent to faculty strength while deciding faculty-to-student ratio, then colleges would be attracted to online courses. A typical faculty member teaches about 100 students in a semester (across 2-3 courses). If 100 students do a course in Mode 4, we may consider this as equivalent to having one additional full time equivalent (FTE) faculty member on the rolls of the college for that semester. Similarly, if 200 students do a course in Mode 3, we may consider this as equivalent to having one additional FTE faculty member for that semester.

## **Online offering of AICTE Model Curriculum Courses by Experts**

It is known that a significant number of colleges do not have adequately qualified faculty to teach the courses in the model curriculum properly. This proposal tries to alleviate this challenge by utilizing our newly acquired comfort level with online education. Essentially, the suggestion is that AICTE incentivizes top faculty to offer the AICTE model courses online, and publishes a calendar of such courses well in advance, so colleges/HEIs can use them for teaching the AICTE curriculum.

The following would be the salient features of the proposed scheme

1. All the identified core courses and later on some key electives of the AICTE curriculum may offered regularly (maybe every semester) with a predefined calendar enabling colleges to plan using them for their students.
2. For each course AICTE identifies through a process a set of "Distinguished National Technical Teachers" who are subject experts and who also have experience teaching the course at the undergraduate level.
3. These "Distinguished Teachers" are offered financial reward for offering one course online once every 2 to 3 years. The reward must be substantial to make it prestigious as well as attractive. Suitable funding is also made available to the instructor for engaging TAs for help sessions as well as laboratory exercises.
4. AICTE can also recognize industry professionals in the panel of distinguished faculty as they can make these courses more interesting and industry oriented. They could teach jointly

---

with academics appointed as Distinguished Teachers and handle certain components of the course. The Distinguished Teachers will be encouraged to actively invite guest faculty from industry and/or professional associations (e.g. ACM) to give lectures in these courses to make them more interesting and industry oriented.

5. There is a registration procedure for colleges (not individual students) to register for making this course available to their students. They should also commit to assigning an instructor for coordinating classes including infrastructure, conducting assessments, help in grading etc.
6. Colleges may be asked to pay a reasonable amount as registration charges so that the scheme is self-sustaining.
7. This will be an AICTE scheme for teaching its model curriculum courses by experts. It can be managed by AICTE directly, or it can offer it to other institutions/consortium to manage them. Any platform (e.g. NPTEL) can be used to deliver these courses.

This approach differs from the currently available online courses from NPTEL and other platforms substantially.

1. Syllabus for each course is the approved AICTE syllabus
2. The courses would be available against a set calendar each semester making it possible for colleges to rely on them
3. It empowers the college rather than individual students to make use of the online course in a “mixed” mode and thus addresses the deficiency of expertise in specific areas in that college
4. With suitable incentive (both prestige and financial) top experts may be motivated to offer these courses.

**Prepared by:** Dheeraj Sanghi (JKLU), Kishore Kothapalli (IIITH), Ashalata Nayak (MAHE-MIT), Suchismita Roy (NIT Durgapur), RBV Subramanyam (NITW), Divya Bansal (PEC), M. Balakrishnan (IIT Delhi)

---

## Appendix: Recommendations for Possible Exits for a BTech CSE

### Preamble

NEP2020 suggests that a student should have multiple exits. This note suggests a possibility for BTech CSE students. It should be emphasized that it is a choice which a student may wish to take due to his/her financial/family/.... situation and needs, and that it should not be considered as a failure option. We suggest two exits, and flexible reentry options.

### Certificate in Computer Science

A student should be able to get a certificate if he/she completes:

1. 50% of the credits for BTech
2. 50% of CSE program core courses (some specific courses may be specified, so the leaving student has decent skills)
3. Institution may specify some CGPA requirements

### BSc in Computer Science

A student should be able to get a BSc if he/she completes:

1. 75% of the credits for BTech, and at least 3 years in the program
2. 100% of CSE core program courses
3. Institution may specify some CGPA requirements

With BSc, the student is eligible for entry into programs which take BSc as an eligibility criteria.

### Reentry to complete the program

A student exiting with a certificate or BSc should be entitled to reenroll in the program. It is suggested that all credits will be transferred, if the student enrolls back within a limited period (suggested: 3 years) of exiting. In case a student enrolls after that, then the transfer of credits should be examined by looking at the change in the curriculum from what the student did.

### Completion Possibility in other Institutions

It will be desirable for HEIs to make it possible for a student to earn a certificate/BSc in one institution and complete the degree program in another. This will enhance the value of certificates and BSc and encourage competition among HEIs.

If these exit options are accepted for multiple BTech programs, it is suggested that AICTE actively communicate these to the industry and other bodies, so they recognize these and accept them as bona-fide credentials for the purposes of recruitment and/or eligibility for admission to programs, appearing in competitive examinations, etc.

**Prepared by:** Prof. M. Balakrishnan (IIT Delhi), Prof. Anshul Kumar (IIT Delhi), Prof. Pankaj Jalote (IIIT-Delhi)



## Committee and Area Experts

### Experts appointed by AICTE:

- Pankaj Jalote, Distinguished Professor and founding Director, IIIT-Delhi (Chair)
- Manoj Singh Gaur, Director IIT Jammu
- Nutan Limaye, IIT Bombay
- Ramkumar, Pro Vice-Chancellor at Krea University
- Dheeraj Sanghi, Vice-Chancellor, JK Lakshmipat University, Jaipur
- Amit Aggarwal, NASSCOM
- Sukumar Nandi, Professor, IIT Gauhati
- Suchismita Roy, Professor, NIT Durgapur
- Ashalatha Nayak, Professor, Manipal Institute of Technology
- RBV Subramanyam, Professor, NIT Warangal
- Sanjiva Prasad, Professor, IIT Delhi
- Venkatesh R, TCS Pune
- Viraj Kumar, ACM India Education Committee
- Vishram Thatte, Amazon India
- Vinnie Jauhari, Microsoft India
- R Latha, IBM India
- Gaurav Aggarwal, Google India
- Vinayaka Ram Gururajan, TCS
- Thirumala and Sundar K S, Infosys
- P.B. Kotur, Wipro
- Ishvinder Singh, Cisco Systems, Inc.
- Rahul Suresh Ghali, Accenture

### Other Experts in the Committee:

- Kishore Kothapalli, Professor, IIIT Hyderabad
- Sudeshna Sarkar, Professor, IIT Karaghpur

### Expert Committees for Different Courses, Micro specializations, Credits for online, ...

Course	Experts in the Committee
Data Structures	Madhavan Mukund (CMI), Manindra Agrawal (IIT Kanpur), Naveen Garg (IIT Delhi), Amit Kumar (IIT Delhi), Venkatesh Raman (IMSC)
Discrete Mathematics	Nitin Saxena (IIT Kanpur), Somenath Biswas (IIT Goa), Partha Mukhopadhyay (CMI), Bhabani P Sinha (ISI Calcutta)
Algorithm Design and Analysis	Madhavan Mukund (CMI), Manindra Agrawal (IIT Kanpur), Naveen Garg (IIT Delhi), Amit Kumar (IIT Delhi), Venkatesh Raman (IMSC)
Computer Architecture/Organization	M. Balakrishnan (IIT Delhi), John Jose (IIT Gauwhati), Biswadip Panda (IIT Bombay), Anupam Basu (IIT Kgp), Indranil Sengupta (IIT Kgp), Anshul Kumar (IIT Delhi)
Advanced Programming	Kishore Kothapalli (IIIT Hyderabad), Vivek Kumar (IIIT Delhi), Swarnendu Biswas (IIT Kanpur)
Operating Systems	Purushottam (Puru) Kulkarni (IIT Bombay), Debadatta Mishra (IIT Kanpur), Chester Rebeiro (IIT Madras)
Databases	Arnab Bhattacharya (IIT Kanput), Sreenivasa Kumar (IIT Madras), Vikram Goel (IIIT-Delhi), S Sudarshan (IIT Bombay)

Computer Networks	BN Jain (IIIT Delhi), Timothy Gonsalves (IIT Mandi), Vinay Ribero (IITB), Mythili Vutukury (IIT Bombay), PM Sreelakshmi (IIT Mandi)
Machine Learning	CV Jawahar (IIIT Hyderabad), PJ Narayanan (IIIT Hyderabad), Vineeth Balasubramanian (IIT Hyderabad), P. K. Biswas (IIT Kgp), Mitesh Khapra (IIT Madras), Piyush Rai (IIT Kanpur), Preethi Jyothi (IIT Bombay), Pabitra Mitra (IIT Kgp), Chetan Arora (IIT Delhi)
Security	Sandeep Shukla (IIT Kanpur), Sambuddho Chakravarty (IIIT Delhi), S.Venkatesan (IIIT Allahabad), Vinod Ganapathy (IISc)
Theory of Computing	Raghunath Tewari (IIT Kanpur), Jayalal Sarma (IIT Madras), S Akshay (IIT Bombay)
Compilers	Uday P. Khedker (IIT Bombay), V. Krishna Nandivada (IIT Madras), Dibyapran Sanyal (nvidia)
Software Engineering Micro specialization	Pankaj Jalote (IIIT-Delhi), Raghu Reddy (IIIT Hyderabad), Vinay Kulkarni (TCS), Meenakshi Dsouza (IIIT Bangalore)
Cyber Security Micro specialization	C. V. Jawahar (IIIT Hyderabad), P. J. Narayanan (IIIT Hyderabad)
Distributed and Cloud Micro specialization	Yogesh Simmhan (IISc Bangalore), Purushottam Kulkarni (IIT Bombay)
HCI micro specialization	Pushpendra Singh (IIIT-Delhi), Anirudha Joshi (IIT Bombay)
Subcommittee for recommendations for online offerings	Dheeraj Sanghi (JKLU), Kishore Kothapalli (IIITH), Ashalata Nayak (MAHE-MIT), Suchismita Roy (NIT Durgapur), RBV Subramanyam (NITW), Divya Bansal (PEC), M. Balakrishnan (IIT Delhi)
Subcommittee for multiple exit options	Prof. M. Balakrishnan (IIT Delhi), Prof. Anshul Kumar (IIT Delhi), Prof. Pankaj Jalote (IIIT-Delhi)