

1 **Systems Fundamentals (SF)**

2 The underlying hardware and software infrastructure upon which applications are constructed is  
3 collectively described by the term "computer systems." Computer systems broadly span the sub-  
4 disciplines of operating systems, parallel and distributed systems, communications networks, and  
5 computer architecture. Traditionally, these areas are taught in a non-integrated way through  
6 independent courses. However these sub-disciplines increasingly share important common  
7 fundamental concepts within their respective cores. These include computational paradigms,  
8 parallelism, cross-layer communications, state and state transition, resource allocation and  
9 scheduling, and so on. This knowledge area presents an integrative cross-layer view of these  
10 fundamental concepts in a unified albeit simplified fashion, providing a common foundation for  
11 the different specialized mechanisms and policies appropriate to the particular knowledge areas  
12 that it underlies. An organizing principle is "programming for performance": what does a  
13 programmer need to know about the underlying system in order to achieve high performance in  
14 an application being developed.

15

16 **SF. Systems Fundamentals [18 core Tier 1, 9 core Tier 2 hours, 27 total]**

	Core-Tier 1 hours	Core-Tier 2 hours	Includes Electives
<b>SF/Computational Paradigms</b>	3		
<b>SF/Cross-Layer Communications</b>	3		
<b>SF/State-State Transition-State Machines</b>	6		
<b>SF/System Support for Parallelism</b>	3		
<b>SF/Performance</b>	3		
<b>SF/Resource Allocation and Scheduling</b>		2	
<b>SF/Proximity</b>		3	
<b>SF/Virtualization and Isolation</b>		2	
<b>SF/Reliability through Redundancy</b>		2	

17

18

## 19 **SF/Computational Paradigms**

20 *[3 Core-Tier 1 hours]*

21 [Cross-reference PD/parallelism fundamentals: The view presented here is the multiple  
22 representations of a system across layers, from hardware building blocks to application  
23 components, and the parallelism available in each representation; PD/parallelism fundamentals  
24 focuses on the application structuring concepts for parallelism.]

25 **Topics:**

- 26 • A computing system as a layered collection of representations
  - 27 • Basic building blocks and components of a computer (gates, flip-flops, registers, interconnections;  
28 Datapath + Control + Memory)
  - 29 • Hardware as a computational paradigm: Fundamental logic building blocks (logic gates, flip-flops,  
30 counters, registers, PL); Logic expressions, minimization, sum of product forms
  - 31 • Application-level sequential processing: single thread [xref PF/]
  - 32 • Simple application-level parallel processing: request level (web services/client-server/distributed), single  
33 thread per server, multiple threads with multiple servers
  - 34 • Basic concept of pipelining, overlapped processing stages
  - 35 • Basic concept of scaling: going faster vs. handling larger problems
- 36

37 **Learning Outcomes:**

- 38 1. List commonly encountered patterns of how computations are organized [Knowledge].
  - 39 2. Describe the basic building blocks of computers and their role in the historical development of computer  
40 architecture [Knowledge].
  - 41 3. Articulate the differences between single thread vs. multiple thread, single server vs. multiple server  
42 models, motivated by real world examples (e.g., cooking recipes, lines for multiple teller machines, couple  
43 shopping for food, wash-dry-fold, etc.) [Knowledge].
  - 44 4. Articulate the concept of strong vs. weak scaling, i.e., how performance is affected by scale of problem vs.  
45 scale of resources to solve the problem. This can be motivated by the simple, real-world examples  
46 [Knowledge].
  - 47 5. Design and simulate a simple logic circuit using the fundamental building blocks of logic design  
48 [Application].
  - 49 6. Write a simple sequential problem and a simple parallel version of the same program [Application].
  - 50 7. Evaluate performance of simple sequential and parallel versions of a program with different problem sizes,  
51 and be able to describe the speed-ups achieved [Evaluation].
- 52

## 53 **SF/Cross-Layer Communications**

54 *[3 Core-Tier 1 hours]*

55 **Topics:**

- 56 • Programming abstractions, interfaces, use of libraries
  - 57 • Distinction between application and OS services, remote procedure call
  - 58 • Interactions between applications and virtual machines
  - 59 • Reliability
- 60  
61

62 **Learning Outcomes:**

- 63 1. Describe how computing systems are constructed of layers upon layers, based on separation of concerns,  
64 with well-defined interfaces, hiding details of low layers from the higher layers [knowledge].  
65 2. Recognize that hardware, VM, OS, application are just additional layers of interpretation/processing  
66 [knowledge].  
67 3. Describe the mechanisms of how errors are detected, signaled back, and handled through the layers  
68 [knowledge].  
69 4. Construct a simple program using methods of layering, error detection and recovery, and reflection of error  
70 status across layers [application].  
71 5. Find bugs in a layered program by using tools for program tracing, single stepping, and debugging  
72 [evaluation].  
73

74 **SF/State-State Transition-State Machines**

75 *[6 Core-Tier 1 hours]*

76 [Cross-reference AL/Basic Computability and Complexity, OS/state and state diagrams,  
77 NC/protocols]

78 **Topics:**

- 79 • Digital vs. analog/discrete vs. continuous systems  
80 • Simple logic gates, logical expressions, Boolean logic simplification  
81 • Clocks, state, sequencing  
82 • Combinational Logic, Sequential Logic, Registers, Memories  
83 • Computers and Network Protocols as examples of State Machines  
84

85 **Learning Outcomes:**

- 86 1. Describe computations as a system with a known set of configurations, and a byproduct of the computation  
87 is to transition from one unique configuration (state) to another (state) [Knowledge].  
88 2. Recognize the distinction between systems whose output is only a function of their input (Combinational)  
89 and those with memory or history (Sequential) [Knowledge].  
90 3. Describe a computer as a state machine that interprets machine instructions [Knowledge].  
91 4. Explain how a program or network protocol can also be expressed as a state machine, and that alternative  
92 representations for the same computation can exist [Knowledge].  
93 5. Develop state machine descriptions for simple problem statement solutions (e.g., traffic light sequencing,  
94 pattern recognizers) [Application].  
95 6. Derive time-series behavior of a state machine from its state machine representation [Evaluation].  
96

97 **SF/System Support for Parallelism**

98 *[3 Core-Tier1 hours]*

99 [Cross-reference: PD/Parallelism Fundamentals]

100 **Topics:**

- 101 • Execution and runtime models that distinguish Sequential vs. Parallel processing  
102 • System organizations that support Request and Task parallelism and other parallel processing paradigms,  
103 such as Client-Server/Web Services, Thread parallelism(Fork-Join), and Pipelining  
104 • Multicore architectures and hardware support for parallelism  
105

106 **Learning Outcomes:**

- 107  
108  
109  
110  
111  
112  
113  
114  
115  
116  
117  
118  
119
1. For a given program, distinguish between its sequential and parallel execution, and the performance implications thereof [knowledge].
  2. Demonstrate on an execution time line that parallel events and operations can take place simultaneously (i.e., at the same time). Explain how work can be performed in less elapsed time if this can be exploited [knowledge].
  3. Explain other uses of parallelism, such as for reliability/redundancy of execution [knowledge].
  4. Define the differences between the concepts of Instruction Parallelism, Data Parallelism, Thread Parallelism/Multitasking, Task/Request Parallelism.
  5. Write a simple parallel program in more than one paradigm so as to be able to compare and contrast ease of expression and performance in solving a given problem [application].
  6. Use performance tools to measure speed-up achieved by parallel programs in terms of both problem size and number of resources [evaluation].

## 120 **SF/Performance**

121 *[3 Core-Tier 1 hours]*

122 [Cross-reference PD/Parallel Performance]

### 123 *Topics:*

- 124  
125  
126  
127  
128  
129  
130  
131
- Figures of performance merit (e.g., speed of execution, energy consumption, bandwidth vs. latency, resource cost)
  - Benchmarks (e.g., SPEC) and measurement methods
  - CPI equation (Execution time = # of instructions \* cycles/instruction \* time/cycle) as tool for understanding tradeoffs in the design of instruction sets, processor pipelines, and memory system organizations.
  - Amdahl's Law: the part of the computation that cannot be sped up limits the effect of the parts that can

### 132 *Learning Outcomes:*

- 133  
134  
135  
136  
137  
138  
139
1. Explain how the components of system architecture contribute to improving its performance [Knowledge].
  2. Describe Amdahl's law and its implications for parallel system speed-up when limited by sequential portions, e.g., in processing pipelines [Knowledge].
  3. Benchmark a parallel program with different data sets in order to iteratively improve its performance [Application].
  4. Use software tools to profile and measure program performance [Evaluation].

## 140 **SF/Resource Allocation and Scheduling**

141 *[2 Core-Tier 2 hours]*

### 142 *Topics:*

- 143  
144  
145  
146
- Kinds of resources: processor share, memory, disk, net bandwidth
  - Kinds of scheduling: first-come, priority
  - Advantages of fair scheduling, preemptive scheduling

### 147 *Learning Outcomes:*

- 148  
149  
150  
151  
152
1. Define how finite computer resources (e.g., processor share, memory, storage and network bandwidth) are managed by their careful allocation to existing entities [Knowledge].
  2. Describe the scheduling algorithms by which resources are allocated to competing entities, and the figures of merit by which these algorithms are evaluated, such as fairness [Knowledge].
  3. Implement simple scheduling algorithms [Application].

153 4. Measure figures of merit of different scheduler implementations [Evaluation].  
154

## 155 **SF/Proximity**

156 *[3 Core-Tier 2 hours]*

157 [Cross-reference: AR/Memory Management, OS/VM/Virtual Memory]

158 **Topics:**

- 159 • Speed of light and computers (one foot per nanosecond vs. one GHz clocks)
- 160 • Latencies in computer systems: memory vs. disk latencies vs. across the network memory
- 161 • Caches, spatial and temporal locality, in processors and systems
- 162 • Elementary introduction into the processor memory hierarchy: registers and multi-level caches, and the  
163 formula for average memory access time  
164

165 **Learning Outcomes:**

- 166 1. Explain the importance of locality in determining performance [Knowledge].
- 167 2. Describe why things that are close in space take less time to access [Knowledge].
- 168 3. Calculate average memory access time and describe the tradeoffs in memory hierarchy performance in  
169 terms of capacity, miss/hit rate, and access time [Evaluation].  
170

## 171 **SF/Virtualization and Isolation**

172 *[2 Core-Tier 2 hours]*

173 **Topics:**

- 174 • Rationale for protection and predictable performance
- 175 • Levels of indirection, illustrated by virtual memory for managing physical memory resources
- 176 • Methods for implementing virtual memory and virtual machines  
177

178 **Learning Outcomes:**

- 179 1. Explain why it is important to isolate and protect the execution of individual programs and environments  
180 that share common underlying resources, including the processor, memory, storage, and network access  
181 [Knowledge].
- 182 2. Describe how the concept of indirection can create the illusion of a dedicated machine and its resources  
183 even when physically shared among multiple programs and environments [Knowledge].
- 184 3. Measure the performance of two application instances running on separate virtual machines, and determine  
185 the effect of performance isolation [Evaluation].  
186

187

188 **SF/Reliability through Redundancy**

189 *[2 Core-Tier 2 hours]*

190 **Topics:**

- 191 • Distinction between bugs and faults, and how they arise in hardware vs. software
- 192 • How errors increase the longer the distance between the communicating entities; the end-to-end principle
- 193 as it applies to systems and networks
- 194 • Redundancy through check and retry
- 195 • Redundancy through redundant encoding (error correcting codes, CRC/Cyclic Redundancy Codes,
- 196 FEC/Forward Error Correction)
- 197 • Duplication/mirroring/replicas

198  
199 **Learning Outcomes:**

- 200 1. Explain the distinction between program errors, system errors, and hardware faults (e.g., bad memory) and
- 201 exceptions (e.g., attempt to divide by zero) [Knowledge].
- 202 2. Articulate the distinction between detecting, handling, and recovering from faults, and the methods for their
- 203 implementation [Knowledge].
- 204 3. Describe the role of error correcting codes in providing error checking and correction techniques in
- 205 memories, storage, and networks [Knowledge].
- 206 4. Apply simple algorithms for exploiting redundant information for the purposes of data correction
- 207 [Application].
- 208 5. Compare different error detection and correction methods for their data overhead, implementation
- 209 complexity, and relative execution time for encoding, detecting, and correcting errors [Evaluation].
- 210