

Computer Science Curricula 2013

Ironman Draft
(Version 0.8)

November 2012

The Joint Task Force on Computing Curricula
Association for Computing Machinery
IEEE-Computer Society

CS2013 Steering Committee

ACM Delegation

Mehran Sahami, *Chair* (Stanford University)

Andrea Danyluk (Williams College)

Sally Fincher (University of Kent)

Kathleen Fisher (Tufts University)

Dan Grossman (University of Washington)

Beth Hawthorne (Union County College)

Randy Katz (UC Berkeley)

Rich LeBlanc (Seattle University)

Dave Reed (Creighton University)

IEEE-CS Delegation

Steve Roach, *Chair* (Univ. of Texas, El Paso)

Ernesto Cuadros-Vargas (Univ. Catolica San Pablo)

Ronald Dodge (US Military Academy)

Robert France (Colorado State University)

Amruth Kumar (Ramapo Coll. of New Jersey)

Brian Robinson (ABB Corporation)

Remzi Seker (Embry-Riddle Aeronautical Univ.)

Alfred Thompson (Microsoft)

Table of Contents

| | |
|---|----|
| Chapter 1: Introduction..... | 5 |
| Charter..... | 6 |
| High-level Themes..... | 6 |
| Knowledge Areas..... | 6 |
| Previous Input..... | 8 |
| Coming Attractions in CS2013..... | 9 |
| Timeline..... | 10 |
| Exemplars of Curricula and Courses..... | 10 |
| Professional Practice..... | 10 |
| Institutional Challenges..... | 11 |
| Opportunities for Involvement..... | 12 |
| References..... | 12 |
| Acknowledgments..... | 13 |
| Chapter 2: Principles..... | 16 |
| Chapter 3: Characteristics of Graduates..... | 19 |
| Chapter 4: Constructing a Complete Curriculum..... | 22 |
| Knowledge Areas are Not Necessarily Courses (and Important Examples Thereof)..... | 23 |
| Tier-1 Core, Tier-2 Core, Elective: What These Terms Mean, What is Required..... | 24 |
| Further Considerations..... | 26 |

| | |
|---|-----|
| Chapter 5: Introduction to the Body of Knowledge..... | 28 |
| Process for Updating the Body of Knowledge | 28 |
| Overview of New Knowledge Areas | 29 |
| How to Read the Body of Knowledge | 31 |
| Appendix A: The Body of Knowledge | 36 |
| Algorithms and Complexity (AL)..... | 36 |
| Architecture and Organization (AR)..... | 43 |
| Computational Science (CN) | 49 |
| Discrete Structures (DS) | 56 |
| Graphics and Visualization (GV)..... | 62 |
| Human-Computer Interaction (HCI)..... | 69 |
| Information Assurance and Security (IAS)..... | 77 |
| Information Management (IM) | 87 |
| Intelligent Systems (IS)..... | 96 |
| Networking and Communication (NC)..... | 106 |
| Operating Systems (OS) | 111 |
| Platform-Based Development (PBD) | 118 |
| Parallel and Distributed Computing (PD)..... | 121 |
| Programming Languages (PL)..... | 132 |
| Software Development Fundamentals (SDF) | 143 |
| Software Engineering (SE) | 148 |
| Systems Fundamentals (SF)..... | 162 |
| Social Issues and Professional Practice (SP) | 169 |

Chapter 1: Introduction

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

Continuing a process that began over 40 years ago with the publication of Curriculum 68 [1], the major professional societies in computing—ACM and IEEE-Computer Society—have sponsored efforts to establish international curricular guidelines for undergraduate programs in computing on roughly a 10-year cycle. As the field of computing has grown and diversified, so too have the curricular recommendations, and there are now curricular volumes for Computer Engineering, Information Systems, Information Technology, and Software Engineering in addition to Computer Science [3]. These volumes are updated regularly with the aim of keeping computing curricula modern and relevant. The last complete Computer Science curricular volume was released in 2001 (CC2001) [2], and an interim review effort concluded in 2008 (CS2008) [4].

This volume, Computer Science Curricula 2013 (CS2013), represents a comprehensive revision. CS2013 redefines the knowledge units in CS, rethinking the essentials necessary for a Computer Science curriculum. It also seeks to identify exemplars of actual courses and programs to provide concrete guidance on curricular structure and development in a variety of institutional contexts.

The development of curricular guidelines for Computer Science is particularly challenging given the rapid evolution and expansion of the field: material dates fast. Moreover, the growing diversity of topics in Computer Science and the increasing integration of computing with other disciplines create additional challenges. Balancing topical growth with the need to keep recommendations realistic and implementable in the context of undergraduate education is particularly difficult. As a result, it is important to engage the broader computer science education community in a dialog to better understand new opportunities, local needs, and to identify successful models of computing curriculum – whether established or novel. One aim of this Strawman report is to provide the basis for such engagement, by providing an early draft of the CS2013 volume that can be scrutinized by members of the computing community with the goal of augmenting and refining the final report.

28 **Charter**

29 The ACM and IEEE-Computer Society chartered the CS2013 effort with the following directive:

30 *To review the Joint ACM and IEEE-CS Computer Science volume of*
31 *Computing Curricula 2001 and the accompanying interim review CS 2008, and*
32 *develop a revised and enhanced version for the year 2013 that will match the*
33 *latest developments in the discipline and have lasting impact.*

34 *The CS2013 task force will seek input from a diverse audience with the goal of*
35 *broadening participation in computer science. The report will seek to be*
36 *international in scope and offer curricular and pedagogical guidance*
37 *applicable to a wide range of institutions. The process of producing the final*
38 *report will include multiple opportunities for public consultation and scrutiny.*

39 Consequently, the CS2013 task force welcomes review of, and comment on, this draft report.

40 **High-level Themes**

41 In developing CS2013, several high-level themes provided an overarching guide for this volume.

42 These themes, which embody and reflect the CS2013 Principles (described in detail in another
43 section of this volume) are:

- 44 • *The “Big Tent” view of CS.* As CS expands to include more cross-disciplinary work and
45 new programs of the form “Computational Biology,” “Computational Engineering,” and
46 “Computational X” are developed, it is important to embrace an outward-looking view
47 that sees CS as a discipline actively seeking to work with and integrate into other
48 disciplines.
- 49 • *Managing the size of the curriculum.* Although the field of Computer Science continues
50 to grow unabated, it is not feasible to proportionately expand the size of the curriculum.
51 As a result, CS2013 seeks to re-evaluate the essential topics in computing to make room
52 for new topics without requiring more total instructional hours than the CS2008
53 guidelines. At the same time, the circumscription of curriculum size promotes more
54 flexible models for curricula without losing the essence of a rigorous CS education.
- 55 • *Actual course exemplars.* CS2001 took on the significant challenge of providing
56 descriptions of six *curriculum models* and forty-seven possible *course descriptions*
57 variously incorporating the knowledge units as defined in that report. While this effort
58 was valiant, in retrospect such course guidance did not seem to have much impact on
59 actual course design. CS2013 plans to take a different approach: to identify and describe
60 existing successful courses and curricula to show how relevant knowledge units are
61 addressed and incorporated in actual programs.

- 62 • *Institutional needs.* CS2013 aims to be applicable in a broad range of geographic and
63 cultural contexts, understanding that curricula exist within specific institutional needs,
64 goals, and resource constraints. As a result, CS2013 allows for explicit flexibility in
65 curricular structure through a tiered set of core topics, where a small set of Core-Tier 1
66 topics are considered essential for all CS programs, but individual programs choose their
67 coverage of Core-Tier 2 topics. This tiered structure is described in more detail in
68 Chapter 4 of this report.

69 **Knowledge Areas**

70 The CS2013 Body of Knowledge is organized into a set of 18 Knowledge Areas (KAs),
71 corresponding to topical areas of study in computing. The Knowledge Areas are:

- 72 • AL - Algorithms and Complexity
- 73 • AR - Architecture and Organization
- 74 • CN - Computational Science
- 75 • DS - Discrete Structures
- 76 • GV - Graphics and Visual Computing
- 77 • HCI - Human-Computer Interaction
- 78 • IAS - Information Assurance and Security
- 79 • IM - Information Management
- 80 • IS - Intelligent Systems
- 81 • NC - Networking and Communications
- 82 • OS - Operating Systems
- 83 • PBD - Platform-based Development
- 84 • PD - Parallel and Distributed Computing
- 85 • PL - Programming Languages
- 86 • SDF - Software Development Fundamentals
- 87 • SE - Software Engineering
- 88 • SF - Systems Fundamentals
- 89 • SP - Social Issues and Professional Issues

90

91 Many of these Knowledge Areas are derived from CC2001/CS2008 but have been revised—in
92 some cases quite significantly—in CS2013; others are new. There are three major causes of KA
93 change: the reorganization of existing KAs, the development of cross-cutting KAs, and the
94 creation of entirely new KAs. Reorganized KAs are a refactoring of existing topics to better
95 reflect coherent units of knowledge as the field of Computer Science has evolved. For example,
96 Software Development Fundamentals is a significant reorganization of the previous
97 Programming Fundamentals KA. Cross-cutting KAs are a refactoring of existing KAs that
98 extract and integrates cross-cutting foundational topics into their own KA rather than duplicating
99 them across many others. Examples include SF-System Fundamentals and IAS-Information
100 Assurance and Security. Finally, new KAs reflect emerging topics in CS that have become
101 sufficiently prevalent to be included in the volume. PBD-Platform-based Development is an
102 example of such a KA. Chapter 5 contains a more comprehensive overview of these changes.

103 **Previous Input**

104 To lay the groundwork for CS2013, we conducted a survey of the usage of the CC2001 and
105 CS2008 volumes. The survey was sent to approximately 1500 Computer Science (and related
106 discipline) Department Chairs and Directors of Undergraduate Studies in the United States and
107 an additional 2000 Department Chairs internationally. We received 201 responses, representing a
108 wide range of institutions (self-identified):

- 109 • research-oriented universities (55%)
- 110 • teaching-oriented universities (17.5%)
- 111 • undergraduate-only colleges (22.5%)
- 112 • community colleges (5%)

113 The institutions also varied considerably in size, with the following distribution:

- 114 • less than 1,000 students (6.5%)
- 115 • 1,000 to 5,000 students (30%)
- 116 • 5,000 to 10,000 students (19%)
- 117 • more than 10,000 students (44.5%)

118 In examining the *usage* of the CC2001/CS2008 reports, survey respondents reported that the
119 Body of Knowledge (i.e., the outline of topics that should appear in undergraduate Computer
120 Science curricula) was the most used aspect. When questioned about new topical areas that
121 should be added to the Body of Knowledge, survey respondents indicated a strong need to add
122 the topics of *Security* as well as *Parallel and Distributed Computing*. Indeed, feedback during
123 the CS2008 review had also indicated the importance of these two areas, but the CS2008 steering
124 committee had felt that creating new KAs was beyond their purview and deferred the
125 development of those areas to the next full curricular report. CS2013 includes these two new
126 KAs (among others): *Information Assurance and Security*, and *Parallel and Distributed*
127 *Computing*.

128 **Coming Attractions in CS2013**

129 The final version of the CS2013 volume is, naturally enough, scheduled for release in 2013.
130 Hence, this Ironman draft is—by design—incomplete. Not only will the final report include
131 revisions of the Body of Knowledge presented here, based on community feedback, it will also
132 include several sections which do not yet exist. Here we provide a timeline for CS2013 efforts
133 and outline some of the “coming attractions” (i.e., additional sections) that are planned for
134 inclusion in future drafts.

135

136 **Timeline**

137 The 2013 curricular guidelines will comprise several sorts of materials: the Body of Knowledge,
138 Exemplars of Curricula and Courses, Professional Practice, and Institutional Challenges. These
139 are being developed in offset phases, starting with the Body of Knowledge.

140 A summary of the CS2013 timeline is as follows:

- Fall 2010: CS2013 chartered and effort begins
- February 2011: CS2013 Principles outlined and Body of Knowledge revision begins
- February 2012: CS2013 Strawman report released
Includes: Body of Knowledge, Characteristics of Graduates
- July 15, 2012: Comment period for Strawman draft closes
- February 2013: CS2013 Ironman report version 1.0 planned for release
Includes: Body of Knowledge, Characteristics of Graduates, Curricula
and Course Exemplars, Professional Practice, Institutional Challenges
- June 2013: Comment period for Ironman draft closes
- Summer 2013: CS2013 Final report planned for release

141

142 **Exemplars of Curricula and Courses**

143 Perhaps the most significant section of the CS2013 final report that is not included in the current
144 Ironman version 0.8 draft is the presentation of actual curricula and courses that embody the
145 topics in the CS2013 Body of Knowledge. The CS2013 Ironman draft will include examples
146 used in practice—from a variety of universities and colleges—to illustrate how topics in the
147 Knowledge Areas may be covered and combined in diverse ways.

148 Importantly, we believe that the identification of such exemplary courses and curricula provides
149 a tremendous opportunity for further community involvement in the development of the CS2013
150 volume. We invite members of the computing community to contribute courses and curricula

151 from their own institutions (or other institutions that they may be familiar with). Those
152 interested in potentially mapping courses/curricula to the CS2013 Body of Knowledge are
153 encouraged to contact members of the CS2013 steering committee for more details.

154 **Professional Practice**

155 The education that undergraduates in Computer Science receive must adequately prepare them
156 for the workforce in a more holistic way than simply conveying technical facts. Indeed, “soft
157 skills” (such as teamwork and communication) and personal attributes (such as identification of
158 opportunity and risk) play a critical role in the workplace. Successfully applying technical
159 knowledge in practice often requires an ability to tolerate ambiguity and work well with others
160 from different backgrounds and disciplines. These overarching considerations are important for
161 promoting successful professional practice in a variety of career paths. We will include
162 suggestions for, and examples of, ways in which curricula encourage the development of such
163 skills, including professional competencies and entrepreneurship, as part of an undergraduate
164 Computer Science program in the CS2013 Ironman version 1.0 draft.

165 **Institutional Challenges**

166 CS departments and programs often face institutional challenges in implementing a curriculum:
167 they may have too few faculty to cover all the knowledge areas, insufficient number of students
168 for a full program, and/or inadequate institutional resource for professional development. This
169 section will identify such challenges and provide suggestions for their amelioration.

170

171 **Opportunities for Involvement**

172 We believe it is essential for endeavours of this kind to engage the broad computing community
173 to review and critique successive drafts. To this end, the development of this Strawman report
174 has already benefited from the input of more than 100 contributors beyond the steering
175 committee. We welcome further community engagement on this effort in multiple ways,
176 including (but not limited to):

- 177 • Comments on the Ironman version 0.8 draft.
- 178 • Contribution of exemplar courses/curricula that are mapped against the Body of
179 Knowledge.
- 180 • Descriptions of pedagogic approaches and instructional designs (both time-tested and
181 novel) that address professional practice within undergraduate curricula.
- 182 • Sharing of institutional challenges, and solutions to them.

183 Comments on all aspects of this report are welcome and encouraged via the CS2013 website:

184 **<http://cs2013.org>**

185

186 **References**

- 187 [1] ACM Curriculum Committee on Computer Science. 1968. Curriculum 68:
188 Recommendations for Academic Programs in Computer Science. *Comm. ACM* 11, 3 (Mar.
189 1968), 151-197.
- 190 [2] ACM/IEEE-CS Joint Task Force on Computing Curricula. 2001. ACM/IEEE Computing
191 Curricula 2001 Final Report. <http://www.acm.org/sigcse/cc2001>.
- 192 [3] ACM/IEEE-CS Joint Task Force for Computer Curricula 2005. Computing Curricula
193 2005: An Overview Report. [http://www.acm.org/education/curric_vols/CC2005-](http://www.acm.org/education/curric_vols/CC2005-March06Final.pdf)
194 [March06Final.pdf](http://www.acm.org/education/curric_vols/CC2005-March06Final.pdf)
- 195 [4] ACM/IEEE-CS Joint Interim Review Task Force. 2008. Computer Science Curriculum
196 2008: An Interim Revision of CS 2001, Report from the Interim Review Task Force.
197 <http://www.acm.org/education/curricula/ComputerScience2008.pdf>

198

199 **Acknowledgments**

200 The CS2013 Strawman report has benefited from the input of many individuals, including: Alex
201 Aiken (Stanford University), Ross Anderson (Cambridge University), Florence Appel (Saint
202 Xavier University), Helen Armstrong (Curtin university), Colin Armstrong (Curtin university),
203 Krste Asanovic (UC Berkeley), Radu F. Babiceanu (University of Arkansas at Little Rock),
204 Mike Barker (Massachusetts Institute of Technology), Michael Barker (Nara Institute of Science
205 and Technology), Paul Beame (University of Washington), Robert Beck (Villanova University),
206 Matt Bishop (University of California, Davis), Alan Blackwell (Cambridge University), Don
207 Blaheta (Longwood University), Olivier Bonaventure (Universite Catholique de Louvain), Roger
208 Boyle (University of Leeds), Clay Breshears (Intel), Bo Brinkman (Miami University), David
209 Broman (Linkoping University), Kim Bruce (Pomona College), Jonathan Buss (University of
210 Waterloo), Netiva Caftori (Northeastern Illinois University, Chicago), Paul Cairns (University of
211 York), Alison Clear (Christchurch Polytechnic Institute of Technology), Curt Clifton (Rose-
212 Hulman and The Omni Group), Yvonne Cody (University of Victoria), Tony Cowling
213 (University of Sheffield), Joyce Currie-Little (Towson University), Ron Cytron (Washington
214 University in St. Louis), Melissa Dark (Purdue University), Janet Davis (Grinnell College),
215 Marie DesJardins (University of Maryland, Baltimore County), Zachary Dodds (Harvey Mudd
216 College), Paul Dourish (University of California, Irvine), Lynette Drevin (North-West
217 Universit), Scot Drysdale (Dartmouth College), Kathi Fisler (Worcester Polytechnic Institute),
218 Susan Fox (Macalester College), Edward Fox (Virginia Tech), Eric Freudenthal (University of
219 Texas El Paso), Stephen Freund (Williams College), Lynn Fatcher (Nelson Mandela
220 Metropolitan University), Greg Gagne (Wesminister College), Dan Garcia (UC Berkeley), Judy
221 Gersting (Indiana University-Purdue University Indianapolis), Yolanda Gil (University of
222 Southern California), Michael Gleicher (University Wisconsin, Madison), Frances Grodzinsky
223 (Sacred Heart University), Anshul Gupta (IBM), Mark Guzdial (Georgia Tech), Brian Hay
224 (University of Alaska, Fairbanks), Brian Henderson-Sellers (University of Technology, Sydney),
225 Matthew Hertz (Canisius College), Tom Hilburn (Embry-Riddle Aeronautical University), Tony
226 Hosking (Purdue University), Johan Jeuring (Utrecht University), Yiming Ji (University of South
227 Carolina Beaufort), Maggie Johnson (Google), Matt Jones (Swansea University), Frans
228 Kaashoek (MIT), Lisa Kaczmarczyk (ACM Education Council), Jennifer Kay (Rowan

229 University), Scott Klemmer (Stanford University), Jim Kurose (University of Massachusetts,
230 Amherst), Doug Lea (SUNY Oswego), Terry Linkletter (Central Washington University), David
231 Lubke (NVIDIA), Bill Manaris (College of Charleston), Samuel Mann (Otago Polytechnic), C.
232 Diane Martin (George Washington University), Andrew McGettrick (University of Strathclyde),
233 Morgan Mcguire (Williams College), Keith Miller (University of Illinois at Springfield),
234 Narayan Murthy (Pace University), Kara Nance (University of Alaska, Fairbanks), Todd Neller
235 (Gettysburg College), Reece Newman (Sinclair Community College), Christine Nickell
236 (Information Assurance Center for Computer Network Operations, CyberSecurity, and
237 Information Assurance), James Noble (Victoria University of Wellington), Peter Norvig
238 (Google), Joseph O'Rourke (Smith College), Jens Palsberg (UCLA), Robert Panoff (Shodor.org),
239 Sushil Prasad (Georgia State University), Michael Quinn (Seattle University), Matt Ratto
240 (University of Toronto), Penny Rheingans (U. Maryland Baltimore County), Carols Rieder
241 (Lucerne University of Applied Sciences), Eric Roberts (Stanford University), Arny Rosenberg
242 (Northeastern and Colorado State University), Ingrid Russell (University of Hartford), Dino
243 Schweitzer (United States Air Force Academy), Michael Scott (University of Rochester), Robert
244 Sedgewick (Princeton University), Helen Sharp (Open Univeristy), Robert Sloan (University of
245 Illinois, Chicago), Ann Sobel (Miami University), Carol Spradling (Northwest Missouri State
246 University), Michelle Strout (Colorado State University), Alan Sussman (University of
247 Maryland, College Park), Blair Taylor (Towson University), Simon Thompson (University of
248 Kent), Johan Vanniekerk (Nelson Mandela Metropolitan University), Christoph von Praun
249 (Georg-Simon-Ohm Hochschule Nürnberg), Rossouw Von Solms (Nelson Mandela
250 Metropolitan University), John Wawrzynek (UC Berkeley), Charles Weems (Univ. of
251 Massachusetts, Amherst), David Wetherall (University of Washington), Michael Wrinn (Intel)

252 Additionally, review of various portions of the Strawman report took part in several venues,
253 including: the 42nd ACM Technical Symposium of the Special Interest Group on Computer
254 Science Education (SIGCSE-11), the 24th IEEE-CS Conference on Software Engineering
255 Education and Training (CSEET-11), the 2011 IEEE Frontiers in Education Conference (FIE-
256 11), the 2011 Federated Computing Research Conference (FCRC-11), the 2nd Symposium on
257 Educational Advances in Artificial Intelligence (EAAI-11), the Conference of ACM Special
258 Interest Group on Data Communication 2011 (SIGCOMM-11), the 2011 IEEE International
259 Joint Conference on Computer, Information, and Systems Sciences and Engineering (CISSE-11),

260 the 2011 Systems, Programming, Languages and Applications: Software for Humanity
261 Conference (SPLASH-11), the 15th Colloquium for Information Systems Security Education, the
262 2011 National Centers of Academic Excellence in IA Education (CAE/IAE) Principles meeting,
263 and the 7th IFIP TC 11.8 World Conference on Information Security Education (WISE).

264 Several more conference special sessions to review and comment on drafts of CS2013 are
265 planned for the coming year, including 43rd ACM Technical Symposium of the Special Interest
266 Group on Computer Science Education (SIGCSE-12), the Special Session of the Special Interest
267 Group on Computers and Society at SIGCSE-12, Computer Research Association Snowbird
268 Conference 2012, and the 2012 IEEE Frontiers in Education Conference (FIE-12), among others.

269 A number of organizations also provided valuable feedback to the CS2013 Strawman effort,
270 including: the ACM Education Board and Council, the IEEE-CS Educational Activities Board,
271 the ACM SIGPLAN Education Board, the ACM Special Interest Group Computers and Society,
272 the Liberal Arts Computer Science Consortium (LACS), and the NSF/IEEE-TCPP Curriculum
273 Initiative on Parallel and Distributed Computing Committee.

Chapter 2: Principles

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

Early in its work, the 2013 Steering Committee agreed upon a set of principles to guide the development of this volume. The principles adopted for CS2013 overlap significantly with the principles adopted for previous curricular efforts, most notably CC2001 and CS2008. As with previous ACM/IEEE curricula volumes, there are a variety of constituencies for CS2013, including individual faculty members and instructors at a wide range of colleges, universities, and technical schools on any of six continents; CS programs and the departments, colleges, and institutions where they are housed; accreditation and certification boards; authors; and researchers. Other constituencies include pre-college preparatory schools and advanced placement curricula as well as graduate programs in computer science.

The principles were developed in consideration of these constituencies, as well as issues related to student outcomes, development of curricula, and the review process. The order of presentation is not intended to imply relative importance.

1. *Computer Science curricula should be designed to provide students with the flexibility to work across many disciplines.* Computing is a broad field that connects to and draws from many disciplines, including mathematics, electrical and systems engineering, psychology, statistics, fine arts, linguistics, and physical and life sciences. Computer Science students should develop the flexibility to work across disciplines.
2. *Computer Science curricula should be designed to prepare graduates for a variety of professions, attracting the full range of talent to the field.* Computer Science impacts nearly every modern endeavour. CS2013 takes a broad view of the field that includes topics such as “computational-x” (e.g., computational finance or computational chemistry) and “x-informatics” (e.g., eco-informatics or bio-informatics). Well-rounded CS graduates will have a balance of theory and application, as described in Chapter 3: Characteristics of Graduates.
3. *CS2013 should provide guidance for the expected level of mastery of topics by graduates.* It should suggest outcomes indicating the intended level of mastery and provide exemplars of fielded curricula covering topics in the Body of Knowledge.

- 28 4. *CS 2013 must provide realistic, adoptable recommendations that provide guidance and*
29 *flexibility, allowing curricular designs that are innovative and track recent developments in*
30 *the field.* The guidelines are intended to provide clear, implementable goals, while also
31 providing the flexibility that programs need in order to respond to a rapidly changing field.
32 CS2013 is intended as guidance, not as a minimal standard against which to evaluate a
33 program.
- 34 5. *The CS2013 guidelines must be relevant to a variety of institutions.* Given the wide range of
35 institutions and programs (including 2-year, 3-year, and 4-year programs; liberal arts,
36 technological, and research institutions; and institutions of every size), it is neither possible
37 nor desirable for these guidelines to dictate curricula for computing. Individual programs will
38 need to evaluate their constraints and environments to construct curricula.
- 39 6. *The size of the essential knowledge must be managed.* While the range of relevant topics has
40 expanded, the size of undergraduate curricula has not. Thus, CS2013 must carefully choose
41 among topics and recommend the essential elements.
- 42 7. *Computer Science curricula should be designed to prepare graduates to succeed in a rapidly*
43 *changing field.* Computer Science is rapidly changing and will continue to change for the
44 foreseeable future. Curricula must prepare students for lifelong learning and must include
45 professional practice (e.g. communication skills, teamwork, ethics) as components of the
46 undergraduate experience. Computer science students must learn to integrate theory and
47 practice, to recognize the importance of abstraction, and to appreciate the value of good
48 engineering design.
- 49 8. *CS2013 should identify the fundamental skills and knowledge that all computer science*
50 *graduates should possess while providing the greatest flexibility in selecting topics.* To this
51 end, we have introduced three levels of knowledge description: Tier-1 Core, Tier-2 Core, and
52 Elective. For a full discussion of Tier-1 Core, Tier-2 Core, and Elective, see Chapter 4:
53 Completing the Curriculum.
- 54 9. *CS2013 should provide the greatest flexibility in organizing topics into courses and*
55 *curricula.* Knowledge areas are not intended to describe specific courses. There are many

56 novel, interesting, and effective ways to combine topics from the Body of Knowledge into
57 courses.

58 10. *The development and review of CS2013 must be broadly based.* The CS2013 Task Force
59 must include participation from many different constituencies including industry,
60 government, and the full range of higher education institutions involved in computer science
61 education. It must take into account relevant feedback from these constituencies.

Chapter 3: Characteristics of Graduates

Graduates of Computer Science programs should have fundamental competency in the areas described by the Body of Knowledge (see Chapter 5), particularly the core topics contained there. However, there are also competences that graduates of CS programs should have that are not explicitly listed in the Body of Knowledge. Professionals in the field typically embody a characteristic style of thinking and problem solving, a style that emerges from the experiences obtained through study of the field and professional practice. Below, we describe the characteristics that we believe should be met at least at an elementary level by graduates of computer science programs. These characteristics will enable their success in the field and further professional development. Some of these characteristics and skills also apply to other fields. They are included here because the development of these skills and characteristics must be explicitly addressed and encouraged by Computer Science programs.

This list is based on a similar list in CC2001 and CS2008. The substantial changes that led to this new version were influenced by responses to a survey conducted by the CS2013 Steering Committee.

At a broad level, the expected characteristics of computer science graduates include the following:

Technical understanding of Computer Science

Graduates should have a mastery of computer science as described by the core of the Body of Knowledge.

Familiarity with common themes and principles

Graduates need understanding of a number of recurring themes, such as abstraction, complexity, and evolutionary change, and a set of general principles, such as sharing a common resource, security, and concurrency. Graduates should recognize that these themes and principles have broad application to the field of computer science and should not consider them as relevant only to the domains in which they were introduced.

28 ***Appreciation of the interplay between theory and practice***

29 A fundamental aspect of computer science is understanding the interplay between theory and practice and
30 the essential links between them. Graduates of a computer science program need to understand how
31 theory and practice influence each other.

32 ***System-level perspective***

33 Graduates of a computer science program need to think at multiple levels of detail and abstraction. This
34 understanding should transcend the implementation details of the various components to encompass an
35 appreciation for the structure of computer systems and the processes involved in their construction and
36 analysis. They need to recognize the context in which a computer system may function, including its
37 interactions with people and the physical world.

38 ***Problem solving skills***

39 Graduates need to understand how to apply the knowledge they have gained to solve real problems, not
40 just write code and move bits. They should be able to design and improve a systems based on a
41 quantitative and qualitative assessment of its functionality, usability and performance. They should realize
42 that there are multiple solutions to a given problem and that selecting among them is not a purely
43 technical activity, as these solutions will have a real impact on people’s lives. Graduates also should be
44 able to communicate their solution to others, including why and how a solution solves the problem and
45 what assumptions were made.

46 ***Project experience***

47 To ensure that graduates can successfully apply the knowledge they have gained, all graduates of
48 computer science programs should have been involved in at least one substantial project. In most cases,
49 this experience will be a software development project, but other experiences are also appropriate in
50 particular circumstances. Such projects should challenge students by being integrative, requiring
51 evaluation of potential solutions, and requiring work on a larger scale than typical course projects.
52 Students should have opportunities to develop their interpersonal communication skills as part of their
53 project experience.

54 ***Commitment to life-long learning***

55 Graduates of a computer science program should realize that the computing field advances at a rapid
56 pace, and graduates must possess a solid foundation that allows and encourages them to maintain relevant
57 skills as the field evolves. Specific languages and technology platforms change over time. Therefore,
58 graduates need to realize that they must continue to learn and adapt their skills throughout their careers.

59 To develop this ability, students should be exposed to multiple programming languages, tools, and
60 technologies as well as the fundamental underlying principles throughout their education.

61 ***Commitment to professional responsibility***

62 Graduates should recognize the social, legal, ethical and cultural issues inherent in the discipline of
63 computing. They must further recognize that social, legal, and ethical standards vary internationally. They
64 should be knowledgeable about the interplay of ethical issues, technical problems, and aesthetic values
65 that play an important part in the development of computing systems. Practitioners must understand their
66 individual and collective responsibility and the possible consequence of failure. They must understand
67 their own limitations as well as the limitations of their tools.

68 ***Communication and organizational skills***

69 Graduates should have the ability to make succinct presentations to a range of audiences about technical
70 problems and their solutions. This may involve face-to-face, written, or electronic communication. They
71 should be prepared to work effectively as members of teams. Graduates should be able to manage their
72 own learning and development, including managing time, priorities, and progress.

73 ***Awareness of the broad applicability of computing***

74 Platforms range from embedded micro-sensors to high-performance clusters and distributed clouds.
75 Computer applications impact nearly every aspect of modern life. Graduates should understand the full
76 range of opportunities available in computing.

77 ***Appreciation of domain-specific knowledge***

78 Graduates should understand that computing interacts with many different domains. Solutions to many
79 problems require both computing skills and domain knowledge. Therefore, graduates need to be able to
80 communicate with, and learn from, experts from different domains throughout their careers.

Chapter 4: Constructing a Complete Curriculum

This chapter provides high-level guidelines on how to use the Body of Knowledge to create an institution's undergraduate curriculum in computer science. It does not propose a particular set of courses or curriculum structure -- that is the role of the (forthcoming) course/curriculum exemplars. Rather, this chapter emphasizes the flexibility that the Body of Knowledge allows in adapting curricula to institutional needs and the continual evolution of the field. In computer-science terms, one can view the Body of Knowledge as a specification of content to cover and a curriculum as an implementation. A large variety of curricula can meet the specification.

The following points are elaborated:

- Knowledge Areas are not intended to be in one-to-one correspondence with particular courses in a curriculum: We expect curricula will have courses incorporating topics from multiple Knowledge Areas.
- Topics are identified as either “core” or “elective” with the core further subdivided into “tier-1” and “tier-2.”
 - A curriculum should include all topics in the tier-1 core and ensure that all students cover this material.
 - A curriculum should include all or almost all topics in the tier-2 core and ensure that all students cover the vast majority of this material.
 - A curriculum should include significant elective material: Covering only “core” topics is insufficient for a complete curriculum.
- Because it is a hierarchical outline, the Body of Knowledge under-emphasizes some key issues that must be considered when constructing a curriculum.

25 **Knowledge Areas are Not Necessarily Courses (and Important** 26 **Examples Thereof)**

27 It is naturally tempting to associate each Knowledge Area with a course. We explicitly
28 discourage this practice in general, even though many curricula will have some courses
29 containing material from only one Knowledge Area or, conversely, all the material from one
30 Knowledge Area in one course. We view the hierarchical structure of the Body of Knowledge as
31 a useful way to group related information, not as a stricture for organizing material into courses.
32 Beyond this general flexibility, in several places we expect many curricula to integrate material
33 from multiple Knowledge Areas, in particular:

- 34 • *Introductory courses:* There are diverse successful approaches to introductory courses in
35 computer science. Many focus on the topics in Software Development Fundamentals
36 together with a subset of the topics in Programming Languages or Software Engineering,
37 while leaving most of the topics in these other Knowledge Areas to advanced courses.
38 But *which* topics from other Knowledge Areas are covered in introductory courses can
39 vary. Some courses use object-oriented programming, others functional programming,
40 others platform-based development (thereby covering topics in the Platform-Based
41 Development Knowledge Area), etc. Conversely, there is no requirement that all
42 Software Development Fundamentals be covered in a first or second course, though in
43 practice most topics will usually be covered in these early courses.
- 44 • *Systems courses:* The topics in the Systems Fundamentals Knowledge Area can be
45 covered in courses designed to cover general systems principles or in courses devoted to
46 particular systems areas such as computer architecture, operating systems, networking, or
47 distributed systems. For example, an Operating Systems course might spend
48 considerable time on topics of more general use, such as low-level programming,
49 concurrency and synchronization, performance measurement, or computer security. Such
50 courses may draw on material in several Knowledge Areas. Certain fundamental systems
51 topics like latency or parallelism will likely arise in many places in a curriculum. While
52 it is important that such topics do arise, preferably in multiple settings, the Body of
53 Knowledge does not specify the particular settings in which to teach such topics.

54 • *Parallel computing*: Among the many changes to the Body of Knowledge compared to
55 previous reports is a new Knowledge Area in Parallel and Distributed Computing. An
56 alternative structure for the Body of Knowledge would place relevant topics in other
57 Knowledge Areas: parallel algorithms with algorithms, programming constructs in
58 software-development focused areas, multi-core design with computer architecture, and
59 so forth. We chose instead to provide guidance on the essential parallelism topics in one
60 place. Some, but not all, curricula will likely have courses dedicated to parallelism, at
61 least in the near term.

62 **Tier-1 Core, Tier-2 Core, Elective: What These Terms Mean, What is** 63 **Required**

64 As described at the beginning of this chapter, computer science curricula should cover all of the
65 core tier-1 topics, all or almost all of the core tier-2 topics, and significant depth in many of the
66 elective topics (i.e., the core is not sufficient for an undergraduate degree in computer science).
67 Here we provide additional perspective on what “tier-1 core,” “tier-2 core”, and “elective” mean,
68 including motivation for these distinctions.

69 ***Motivation for subdividing the core:*** Earlier versions of the ACM/IEEE Computer Science
70 Curricula had only “core” and “elective” with every topic in the former being required. We
71 departed from this strict interpretation of “everything in the core must be taught to every student”
72 for these reasons:

- 73 • It did not sufficiently reflect reality: Many strong computer science curricula were
74 missing at least one hour of core material. It is misleading to suggest that such curricula
75 are outside the definition of an undergraduate degree in computer science.
- 76 • As the field has grown, there is ever-increasing pressure to grow the core and allow
77 students to specialize in areas of interest. Doing so simply becomes impossible within
78 the short time-frame of an undergraduate degree. Providing some flexibility on coverage
79 of core topics enables curricula and students to specialize if they choose to do so.

80 Conversely, we could have allowed for *any* core topic to be skipped provided that the vast
81 majority was part of every student’s education. By retaining a smaller tier-1 core of required

82 material, we provide additional guidance and structure for curriculum designers. In the tier-1
83 core are the topics that are fundamental to the structure of any computer-science program.

84 ***On the meaning of tier-1:*** A tier-1 topic should be a required part of every computer-science
85 curriculum for every student. This is not to say that tier-2 or even elective topics should not be,
86 but the tier-1 topics are those with widespread consensus for inclusion. Moreover, at least
87 preliminary treatment of most of these topics typically comes in the first two years of a
88 curriculum, precisely because so much of the field relies on these topics. However, introductory
89 courses need not cover all tier-1 material and will usually draw on tier-2 and elective material as
90 well.

91 ***On the meaning of tier-2:*** Tier-2 topics are generally essential in an undergraduate computer-
92 science degree. Requiring the vast majority of them is a *minimum* expectation, and we
93 encourage institutions to cover all of them for every student. That said, computer science
94 programs can allow students to focus in certain areas in which some tier-2 topics are not
95 required. We also acknowledge that resource constraints, such as a small number of faculty or
96 institutional limits on degree requirements, may make it prohibitively difficult to cover every
97 topic in the core while still providing advanced elective material. **A computer-science
98 curriculum should aim to cover 90-100% of the tier-2 topics for every student, with 80%
99 considered as a minimum.**

100 There is no expectation that tier-1 topics necessarily precede tier-2 topics in a curriculum. In
101 particular, we expect introductory courses will draw on both tier-1 and tier-2 (and possibly
102 elective) material and that some core material will be delayed until later courses.

103 ***On the meaning of elective:*** A program covering only core material would provide
104 insufficient breadth and depth in computer science, but most programs will not cover all the
105 elective material in the Body of Knowledge and certainly few, if any, students will cover all of it
106 within an undergraduate program. Conversely, the Body of Knowledge is by no means
107 exhaustive, and advanced courses may often go beyond the topics and learning outcomes
108 contained in it. Nonetheless, the Body of Knowledge provides a useful guide on material
109 appropriate for a computer-science undergraduate degree, and all students of computer science
110 should deepen their understanding in multiple areas via the elective topics.

111 A curriculum may well require material designated elective in the Body of Knowledge. Many
112 curricula, especially those with a particular focus, will require some elective topics, by virtue of
113 them being covered in required courses.

114 **The size of the core:** The size of the core (tier-1 plus tier-2) is a few hours larger than in
115 previous versions of the computer-science curriculum, but this is counterbalanced by our more
116 flexible treatment of the core. As a result, we are not increasing the number of required courses
117 a curriculum should need. Indeed, a curriculum covering 90% of the tier-2 hours would have the
118 same number of core hours as a curriculum covering the core in the CS2008 volume, and a
119 curriculum covering 80% of the tier-2 hours would have fewer core hours than even a curriculum
120 covering the core in the CC2001 volume (the core grew from 2001 to 2008).

121 **A note on balance:** Computer science is an elegant interplay of theory, software, hardware,
122 and applications. The core in general and the tier-1 core in particular, when viewed in isolation,
123 may seem to focus on programming, discrete structures, and algorithms. This focus results from
124 the fact that these topics typically come early in a curriculum so that advanced courses can use
125 them as pre-requisites. Essential experience with systems and applications can be achieved in
126 more disparate ways using elective material in the Body of Knowledge. Because all curricula
127 will include appropriate elective material, an overall curriculum can and should achieve an
128 appropriate balance.

129 **Further Considerations**

130 As useful as the Body of Knowledge is, it is important to complement it with a thoughtful
131 understanding of cross-cutting themes in a curriculum, the “big ideas” of computer science. In
132 designing a curriculum, it is also valuable to identify curriculum-wide objectives, for which the
133 Principles and the Characteristics of Graduates chapters of this volume should prove useful.

134 In the last few years, two on-going trends have had deep effects on many curricula. First, the
135 continuing growth of computer science has led to many programs organizing their curricula to
136 allow for *intradisciplinary* specialization (using terms such as threads, tracks, vectors, etc.).
137 Second, the importance of computing to almost every other field has increasingly led to the
138 creation of *interdisciplinary* programs (joint majors, double majors, etc.) and incorporating
139 interdisciplinary material into computer-science programs. We applaud both trends and believe

140 a flexible Body of Knowledge, including a flexible core, support them. Conversely, such
141 specialization is not required: Many programs will continue to offer a broad yet thorough
142 coverage of computer science as a distinct and coherent discipline.

Chapter 5: Introduction to the Body of Knowledge

Process for Updating the Body of Knowledge

The CS2013 Steering Committee constituted a subcommittee for each KA, chaired by a member of the Steering Committee, and initially including at least two other members of the Steering Committee. Individual subcommittee Chairs then invited expert members (outside the CS2013 Steering Committee) to join the work of defining and reviewing each KA; drafts of KAs were also presented in various conference panel and special session presentations. The KA subcommittee Chairs (as members of the CS2013 Steering Committee) worked to resolve conflicts, eliminate redundancies and appropriately categorize and cross-reference topics between the various KAs. This year-long process ultimately converged to the draft version of the Body of Knowledge presented here.

As noted in the introduction to this report, we are soliciting continued community feedback which will be considered and incorporated into future drafts of the CS2013 report.

The CS2013 Body of Knowledge is presented as a set of Knowledge Areas (KAs), organized on topical themes rather than by course boundaries. Each KA is further organized into a set of Knowledge Units (KUs), which are summarized in a table at the head of each KA section. We expect that the topics within the KAs will be organized into courses in different ways at different institutions.

Here, we provide background for understanding how to read the Body of Knowledge, and we give an overview of the number of core hours in each KA. We also highlight the KAs that have significant cross-topic components and those that are new to this volume. Chapter 4 presents essential background on how the Body of Knowledge translates into actual curricula.

25 **Overview of New Knowledge Areas**

26 While computer science encompasses technologies that change rapidly over time, it is defined by
27 essential concepts, perspectives, and methodologies that are constant. As a result, much of the
28 core Body of Knowledge remains unchanged from earlier curricular volumes. However, new
29 developments in computing technology and pedagogy mean that some aspects of the core evolve
30 over time, and some of the previous structures and organization may no longer be appropriate for
31 describing the discipline. As a result, CS2013 has modified the organization of the curriculum in
32 various ways, adding some new KAs and restructuring others. We highlight these changes in the
33 remainder of this section.

34 **IAS-Information Assurance and Security**

35 IAS is a new KA in recognition of the world's reliance on information technology and its critical
36 role in computer science education. IAS as a domain is the set of controls and processes, both
37 technical and policy, intended to protect and defend information and information systems. IAS
38 draws together topics that are pervasive throughout other KAs. Topics germane to *only* IAS are
39 presented in depth in this KA, whereas other topics are noted and cross referenced to the KAs
40 that contain them. As such, this KA is prefaced with a detailed table of cross-references to other
41 KAs.

42 **NC-Networking and Communication**

43 CC2001 introduced a KA entitled "Net-Centric Computing" which encompassed a combination
44 of topics including traditional networking, web development, and network security. Given the
45 growth and divergence in these topics since the last report, we renamed and re-factored this KA
46 to focus specifically on topics in networking and communication. Discussions of web
47 applications and mobile device development are now covered in the new PBD-Platform-Based
48 Development KA. Security is covered in the new IAS-Information Assurance and Security KA.

49

50 **PBD-Platform-Based Development**

51 PBD is a new KA that recognizes the increasing use of platform-specific programming
52 environments, both at the introductory level and in upper-level electives. Platforms such as the
53 Web or mobile devices enable students to learn within and about environments constrained by
54 hardware, APIs, and special services (often in cross-disciplinary contexts). These environments
55 are sufficiently different from “general purpose” programming to warrant this new (wholly
56 elective) KA.

57 **PD-Parallel and Distributed Computing**

58 Previous curricular volumes had parallelism topics distributed across disparate KAs as electives.
59 Given the vastly increased importance of parallel and distributed computing, it seemed crucial to
60 identify essential concepts in this area and to promote those topics to the core. To highlight and
61 coordinate this material, CS2013 dedicates a KA to this area. This new KA includes material on
62 programming models, programming pragmatics, algorithms, performance, computer architecture,
63 and distributed systems.

64 **SDF-Software Development Fundamentals**

65 This new KA generalizes introductory programming to focus on the entire software development
66 process, identifying concepts and skills that should be mastered in the first year of a computer
67 science program. As a result of its broad purpose, the SDF KA includes fundamental concepts
68 and skills that could appear in other software-oriented KAs (e.g., programming constructs from
69 Programming Languages, simple algorithm analysis from Algorithms and Complexity, simple
70 development methodologies from Software Engineering). Likewise, each of those KAs will
71 contain more advanced material that builds upon the fundamental concepts and skills in SDF.

72 Compared to previous volumes, key approaches to programming -- including object-oriented
73 programming, functional programming, and event-driven programming -- are kept in one place,
74 namely the PL KA, even though many curricula will cover some of these topics in introductory
75 courses.

76

77 **SF-Systems Fundamentals**

78 In previous curricular volumes, the interacting layers of a typical computing system, from
79 hardware building blocks, to architectural organization, to operating system services, to
80 application execution environments (particularly for parallel execution in a modern view of
81 applications), were presented in independent knowledge units. The new Systems Fundamentals
82 KA presents a unified systems perspective and common conceptual foundation for other KAs
83 (notably Architecture and Organization, Network and Communications, Operating Systems, and
84 Parallel and Distributed Algorithms). An organizational principle is “programming for
85 performance”: what a programmer needs to understand about the underlying system to achieve
86 high performance, particularly in terms of exploiting parallelism.

87

88 **How to Read the Body of Knowledge**

89 **Curricular Hours**

90 Continuing in the tradition of CC2001/CS2008, we define the unit of coverage in the Body of
91 Knowledge in terms of **lecture hours**, as being the sole unit that is understandable in (and
92 transferable to) cross-cultural contexts. An “hour” corresponds to the time required to present the
93 material in a traditional lecture-oriented format; the hour count does not include any additional
94 work that is associated with a lecture (e.g., in self-study, lab classes, assessments, etc.). Indeed,
95 we expect students to spend a significant amount of additional time outside of class developing
96 facility with the material presented in class. As with previous reports, we maintain the principle
97 that the use of a lecture-hour as the unit of measurement does not require or endorse the use of
98 traditional lectures for the presentation of material.

99 The specification of topic hours represents the **minimum** amount of time we expect such
100 coverage to take. Any institution may opt to cover the same material in a longer period of time as
101 warranted by the individual needs of that institution.

102

103 **Courses**

104 Throughout the Body of Knowledge, when we refer to a “course” we mean an institutionally-
105 recognized unit of study. Depending on local circumstance, full-time students will take several
106 “courses” at any one time, typically eight or more per academic year. While “course” is a
107 common term at some institutions, others will use other names, for example “module” or
108 “paper”.

109 **Guidance on Learning Outcomes**

110 Each KU within a KA lists both a set of topics and the learning outcomes students are expected
111 to achieve with respect to the topics specified. Each learning outcome has an associated *level of*
112 *mastery*. In defining different levels we drew from other curriculum approaches, especially
113 Bloom’s Taxonomy, which has been well explored within Computer Science. We did not
114 directly apply Bloom’s levels in part because several of them are driven by pedagogic context,
115 which would introduce too much plurality in a document of this kind; in part because we intend
116 the mastery levels to be indicative and not to impose theoretical constraint on users of this
117 document.

118 There are three levels of mastery, defined as:

- 119 • *Familiarity*: The student understands what a concept is or what it means. This level of
120 mastery concerns a basic awareness of a concept as opposed to expecting real facility
121 with its application. It provides an answer to the question “What do you know about
122 this?”
- 123 • *Usage*: The student is able to use or apply a concept in a concrete way. Using a concept
124 may include, for example, appropriately using a specific concept in a program, use of a
125 particular proof technique, or performing a particular analysis. It provides an answer to
126 the question “What do you know how to do?”
- 127 • *Assessment*: The student is able to consider a concept from multiple viewpoints and/or
128 justify the selection of a particular approach to solve a problem. This level of mastery
129 implies more than using a concept; it involves the ability to select an appropriate
130 approach from understood alternatives. It provides an answer to the question “Why
131 would you do that?”

132 As a concrete, although admittedly simplistic, example of these levels of mastery, we consider
133 the notion of iteration in software development, for example for-loops, while-loops, iterators. At
134 the level of “Familiarity,” a student would be expected to have a definition of the concept of
135 iteration in software development and know why it is a useful technique. In order to show
136 mastery at the “Usage” level, a student should be able to write a program properly using a form
137 of iteration. Understanding iteration at the “Assessment” level would require a student to
138 understand multiple methods for iteration and be able to appropriately select among them for
139 different applications.

140

141 **Core Hours in Knowledge Areas**

142 An overview of the number of core hours (both Tier1 and Tier2) by KA in the CS2013 Body of
 143 Knowledge is provided below (for a discussion of Tier1 and Tier2, see Chapter 4). For
 144 comparison, the number of core hours from both the previous CS2008 and CC2001 reports are
 145 provided as well.

| Knowledge Area | CS2013 | | CS2008 | CC2001 |
|--|------------|------------|------------|------------|
| | Tier1 | Tier2 | Core | Core |
| AL-Algorithms and Complexity | 19 | 9 | 31 | 31 |
| AR-Architecture and Organization | 0 | 16 | 36 | 36 |
| CN-Computational Science | 1 | 0 | 0 | 0 |
| DS-Discrete Structures | 37 | 4 | 43 | 43 |
| GV-Graphics and Visual Computing | 2 | 1 | 3 | 3 |
| HCI-Human-Computer Interaction | 4 | 4 | 8 | 8 |
| IAS-Security and Information Assurance | 2 | 6 | -- | -- |
| IM-Information Management | 1 | 9 | 11 | 10 |
| IS-Intelligent Systems | 0 | 10 | 10 | 10 |
| NC-Networking and Communication | 3 | 7 | 15 | 15 |
| OS-Operating Systems | 4 | 11 | 18 | 18 |
| PBD-Platform-based Development | 0 | 0 | -- | -- |
| PD-Parallel and Distributed Computing | 5 | 10 | -- | -- |
| PL-Programming Languages | 8 | 20 | 21 | 21 |
| SDF-Software Development Fundamentals | 43 | 0 | 47 | 38 |
| SE-Software Engineering | 6 | 21 | 31 | 31 |
| SF-Systems Fundamentals | 18 | 9 | -- | -- |
| SP-Social Issues and Professional Practice | 11 | 5 | 16 | 16 |
| Total Core Hours | 164 | 142 | 290 | 280 |

| | |
|---------------------------------------|--------------|
| All Tier1 + All Tier2 Total | 306 |
| All Tier1 + 90% of Tier2 Total | 291.8 |
| All Tier1 + 80% of Tier2 Total | 277.6 |

146
 147 As seen above, in CS2013 the total Tier1 hours together with the entirety of Tier2 hours slightly
 148 exceeds the total core hours from previous reports. However, it is important to note that the
 149 tiered structure of the core in CS2013 explicitly provides the flexibility for institutions to select

150 topics from Tier2 (to include at least 80%). As a result, it is possible to implement the CS2013
151 guidelines with slightly fewer hours than previous curricular guidelines.

Appendix A: The Body of Knowledge

Algorithms and Complexity (AL)

Algorithms are fundamental to computer science and software engineering. The real-world performance of any software system depends on: (1) the algorithms chosen and (2) the suitability and efficiency of the various layers of implementation. Good algorithm design is therefore crucial for the performance of all software systems. Moreover, the study of algorithms provides insight into the intrinsic nature of the problem as well as possible solution techniques independent of programming language, programming paradigm, computer hardware, or any other implementation aspect.

An important part of computing is the ability to select algorithms appropriate to particular purposes and to apply them, recognizing the possibility that no suitable algorithm may exist. This facility relies on understanding the range of algorithms that address an important set of well-defined problems, recognizing their strengths and weaknesses, and their suitability in particular contexts. Efficiency is a pervasive theme throughout this area.

This knowledge area defines the central concepts and skills required to design, implement, and analyze algorithms for solving problems. Algorithms are essential in all advanced areas of computer science: artificial intelligence, databases, distributed computing, graphics, networking, operating systems, programming languages, security, and so on. Algorithms that have specific utility in each of these are listed in the relevant knowledge areas. Cryptography, for example, appears in the new knowledge area on Information Assurance and Security, while parallel and distributed algorithms appear in PD-Parallel and Distributed Computing.

As with all knowledge areas, the order of topics and their groupings do not necessarily correlate to a specific order of presentation. Different programs will teach the topics in different courses and should do so in the order they believe is most appropriate for their students.

26 **AL. Algorithms and Complexity (19 Core-Tier1 hours, 9 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| AL/Basic Analysis | 2 | 2 | N |
| AL/Algorithmic Strategies | 5 | 1 | N |
| AL/Fundamental Data Structures and Algorithms | 9 | 3 | N |
| AL/Basic Automata, Computability and Complexity | 3 | 3 | N |
| AL/Advanced Computational Complexity | | | Y |
| AL/Advanced Automata Theory and Computability | | | Y |
| AL/Advanced Data Structures, Algorithms, and Analysis | | | Y |

27

28 **AL/Basic Analysis**

29 *[2 Core-Tier1 hours, 2 Core-Tier2 hours]*

30 *Topics:*

31 [Core-Tier1]

- 32 • Differences among best, expected, and worst case behaviors of an algorithm
- 33 • Asymptotic analysis of upper and expected complexity bounds
- 34 • Big O notation: formal definition
- 35 • Complexity classes, such as constant, logarithmic, linear, quadratic, and exponential
- 36 • Empirical measurements of performance
- 37 • Time and space trade-offs in algorithms

38
39 [Core-Tier2]

- 40 • Big O notation: use
- 41 • Little o, big omega and big theta notation
- 42 • Recurrence relations
- 43 • Analysis of iterative and recursive algorithms
- 44 • Some version of a Master Theorem

45

46 *Learning Outcomes:*

47 [Core-Tier1]

- 48 1. Explain what is meant by “best”, “expected”, and “worst” case behavior of an algorithm. [Familiarity]
- 49 2. In the context of specific algorithms, identify the characteristics of data and/or other conditions or assumptions that lead to different behaviors. [Assessment]
- 50 3. Determine informally the time and space complexity of simple algorithms. [Usage]

51

- 52 4. Understand the formal definition of big O. [Familiarity]
 53 5. List and contrast standard complexity classes. [Familiarity]
 54 6. Perform empirical studies to validate hypotheses about runtime stemming from mathematical analysis.
 55 Run algorithms on input of various sizes and compare performance. [Assessment]
 56 7. Give examples that illustrate time-space trade-offs of algorithms. [Familiarity]
 57

58 [Core-Tier2]

- 59 8. Use big O notation formally to give asymptotic upper bounds on time and space complexity of algorithms.
 60 [Usage]
 61 9. Use big O notation formally to give expected case bounds on time complexity of algorithms. [Usage]
 62 10. Explain the use of big omega, big theta, and little o notation to describe the amount of work done by an
 63 algorithm. [Familiarity]
 64 11. Use recurrence relations to determine the time complexity of recursively defined algorithms. [Usage]
 65 12. Solve elementary recurrence relations, e.g., using some form of a Master Theorem. [Usage]
 66

67 **AL/Algorithmic Strategies**

68 *[5 Core-Tier1 hours, 1 Core-Tier2 hours]*

69 An instructor might choose to cover these algorithmic strategies in the context of the algorithms
 70 presented in “Fundamental Data Structures and Algorithms” below. While the total number of
 71 hours for the two knowledge units (18) could be divided differently between them, our sense is
 72 that the 1:2 ratio is reasonable.

73 **Topics:**

74 [Core-Tier1]

- 75 • Brute-force algorithms
- 76 • Greedy algorithms
- 77 • Divide-and-conquer (cross-reference SDF/Algorithms and Design/Problem-solving strategies)
- 78 • Recursive backtracking
- 79 • Dynamic Programming

80
 81 [Core-Tier2]

- 82 • Branch-and-bound
 - 83 • Heuristics
 - 84 • Reduction: transform-and-conquer
- 85

86 **Learning Outcomes:**

87 [Core-Tier1]

- 88
- 89 1. For each of the above strategies, identify a practical example to which it would apply. [Familiarity]
 - 90 2. Have facility mapping pseudocode to implementation, implementing examples of algorithmic strategies
 91 from scratch, and applying them to specific problems. [Usage]
 - 92 3. Use a greedy approach to solve an appropriate problem and determine if the greedy rule chosen leads to an
 93 optimal solution. [Usage, Assessment]
 - 94 4. Use a divide-and-conquer algorithm to solve an appropriate problem. [Usage]
 - 95 5. Use recursive backtracking to solve a problem such as navigating a maze. [Usage]
 - 96 6. Use dynamic programming to solve an appropriate problem. [Usage]
- 97
 98

- 99 [Core-Tier2]
100
101 7. Describe various heuristic problem-solving methods. [Familiarity]
102 8. Use a heuristic approach to solve an appropriate problem. [Usage]
103 9. Describe the trade-offs between brute force and other strategies. [Assessment]
104

105 **AL/Fundamental Data Structures and Algorithms**

106 *[9 Core-Tier1 hours, 3 Core-Tier2 hours]*

107 This knowledge unit builds directly on the foundation provided by Software Development
108 Fundamentals (SDF), particularly the material in SDF/Fundamental Data Structures and
109 SDF/Algorithms and Design.

110 **Topics:**

111 [Core-Tier1]

112 Implementation and use of:

- 113 • Simple numerical algorithms, such as computing the average of a list of numbers, finding the min, max,
114 and mode in a list, approximating the square root of a number, or finding the greatest common divisor
- 115 • Sequential and binary search algorithms
- 116 • Worst case quadratic sorting algorithms (selection, insertion)
- 117 • Worst or average case $O(N \log N)$ sorting algorithms (quicksort, heapsort, mergesort)
- 118 • Hash tables, including strategies for avoiding and resolving collisions
- 119 • Binary search trees
 - 120 ○ Common operations on binary search trees such as select min, max, insert, delete, iterate over tree
- 121 • Graphs and graph algorithms
 - 122 ○ Representations of graphs (e.g., adjacency list, adjacency matrix)
 - 123 ○ Depth- and breadth-first traversals

124
125 [Core-Tier2]

- 126 • Heaps
- 127 • Graphs and graph algorithms
 - 128 ○ Shortest-path algorithms (Dijkstra's and Floyd's algorithms)
 - 129 ○ Minimum spanning tree (Prim's and Kruskal's algorithms)
- 130 • Pattern matching and string/text algorithms (e.g., substring matching, regular expression matching, longest
131 common subsequence algorithms)

132

133 **Learning Outcomes:**

134 [Core-Tier1]

- 135 1. Implement basic numerical algorithms. [Usage]
- 136 2. Implement simple search algorithms and explain the differences in their time complexities. [Usage,
137 Assessment]
- 138 3. Be able to implement common quadratic and $O(N \log N)$ sorting algorithms. [Usage]
- 139 4. Understand the implementation of hash tables, including collision avoidance and resolution. [Familiarity]
- 140 5. Discuss the runtime and memory efficiency of principal algorithms for sorting, searching, and hashing.
141 [Familiarity]

- 142 6. Discuss factors other than computational efficiency that influence the choice of algorithms, such as
 143 programming time, maintainability, and the use of application-specific patterns in the input data.
 144 [Familiarity]
 145 7. Solve problems using fundamental graph algorithms, including depth-first and breadth-first search. [Usage]
 146 8. Demonstrate the ability to evaluate algorithms, to select from a range of possible options, to provide
 147 justification for that selection, and to implement the algorithm in a particular context. [Usage, Assessment]
 148
 149 [Core-Tier2]
 150
 151 9. Understand the heap property and the use of heaps as an implementation of priority queues. [Familiarity]
 152 10. Solve problems using graph algorithms, including single-source and all-pairs shortest paths, and at least
 153 one minimum spanning tree algorithm. [Usage]
 154 11. Be able to implement a string-matching algorithm. [Usage]
 155

156 **AL/Basic Automata Computability and Complexity**

157 *[3 Core-Tier1 hours, 3 Core-Tier2 hours]*

158 *Topics:*

159 [Core-Tier1]

- 160 • Finite-state machines
- 161 • Regular expressions
- 162 • The halting problem
- 163

164 [Core-Tier2]

- 165 • Context-free grammars (cross-reference PL/Syntax Analysis)
- 166 • Introduction to the P and NP classes and the P vs NP problem
- 167 • Introduction to the NP-complete class and exemplary NP-complete problems (e.g., SAT, Knapsack)
- 168

169 *Learning Outcomes:*

170 [Core-Tier1]

- 171 1. Discuss the concept of finite state machines. [Familiarity]
- 172 2. Design a deterministic finite state machine to accept a specified language. [Usage]
- 173 3. Generate a regular expression to represent a specified language. [Usage]
- 174 4. Explain why the halting problem has no algorithmic solution. [Familiarity]
- 175

176 [Core-Tier2]

- 177
- 178 5. Design a context-free grammar to represent a specified language. [Usage]
- 179 6. Define the classes P and NP. [Familiarity]
- 180 7. Explain the significance of NP-completeness. [Familiarity]
- 181

182 **AL/Advanced Computational Complexity**

183 *[Elective]*

184 *Topics:*

- 185 • Review definitions of the classes P and NP; introduce P-space and EXP
- 186 • NP-completeness (Cook's theorem)

- 187 • Classic NP-complete problems
- 188 • Reduction Techniques
- 189

190 **Learning Outcomes:**

- 191 1. Define the classes P and NP. (Also appears in AL/Basic Automata, Computability, and Complexity)
192 [Familiarity]
- 193 2. Define the P-space class and its relation to the EXP class. [Familiarity]
- 194 3. Explain the significance of NP-completeness. (Also appears in AL/Basic Automata, Computability, and
195 Complexity) [Familiarity]
- 196 4. Provide examples of classic NP-complete problems. [Familiarity]
- 197 5. Prove that a problem is NP-complete by reducing a classic known NP-complete problem to it. [Usage]
- 198

199 **AL/Advanced Automata Theory and Computability**

200 **[Elective]**

201 **Topics:**

- 202 • Sets and languages
- 203 • Regular languages
 - 204 ○ Review of deterministic finite automata (DFAs)
 - 205 ○ Nondeterministic finite automata (NFAs)
 - 206 ○ Equivalence of DFAs and NFAs
 - 207 ○ Review of regular expressions; their equivalence to finite automata
 - 208 ○ Closure properties
 - 209 ○ Proving languages non-regular, via the pumping lemma or alternative means
- 210 • Context-free languages
 - 211 ○ Push-down automata (PDAs)
 - 212 ○ Relationship of PDAs and context-free grammars
 - 213 ○ Properties of context-free languages
- 214 • Turing machines, or an equivalent formal model of universal computation
- 215 • Nondeterministic Turing machines
- 216 • Chomsky hierarchy
- 217 • The Church-Turing thesis
- 218 • Computability
- 219 • Rice's Theorem
- 220 • Examples of uncomputable functions
- 221 • Implications of uncomputability
- 222

223 **Learning Outcomes:**

- 224 1. Determine a language's place in the Chomsky hierarchy (regular, context-free, recursively enumerable).
225 [Assessment]
- 226 2. Prove that a language is in a specified class and that it is not in the next lower class. [Assessment]
- 227 3. Convert among equivalently powerful notations for a language, including among DFAs, NFAs, and regular
228 expressions, and between PDAs and CFGs. [Usage]
- 229 4. Explain the Church-Turing thesis and its significance. [Familiarity]
- 230 5. Explain Rice's Theorem and its significance. [Familiarity]
- 231 6. Provide examples of uncomputable functions. [Familiarity]
- 232 7. Prove that a problem is uncomputable by reducing a classic known uncomputable problem to it. [Usage]
- 233

234 **AL/Advanced Data Structures Algorithms and Analysis**

235 *[Elective]*

236 Many programs will want their students to have exposure to more advanced algorithms or
237 methods of analysis. Below is a selection of possible advanced topics that are current and timely
238 but by no means exhaustive.

239 *Topics:*

- 240 • Balanced trees (e.g., AVL trees, red-black trees, splay trees, treaps)
- 241 • Graphs (e.g., topological sort, finding strongly connected components, matching)
- 242 • Advanced data structures (e.g., B-trees, Fibonacci heaps)
- 243 • String-based data structures and algorithms (e.g., suffix arrays, suffix trees, tries)
- 244 • Network flows (e.g., max flow [Ford-Fulkerson algorithm], max flow – min cut, maximum bipartite
245 matching)
- 246 • Linear Programming (e.g., duality, simplex method, interior point algorithms)
- 247 • Number-theoretic algorithms (e.g., modular arithmetic, primality testing, integer factorization)
- 248 • Geometric algorithms (e.g., points, line segments, polygons [properties, intersections], finding convex hull,
249 spatial decomposition, collision detection, geometric search/proximity)
- 250 • Randomized algorithms
- 251 • Approximation algorithms
- 252 • Amortized analysis
- 253 • Probabilistic analysis
- 254 • Online algorithms and competitive analysis
- 255

256 *Learning Outcomes:*

- 257 1. Understand the mapping of real-world problems to algorithmic solutions (e.g., as graph problems, linear
258 programs, etc.) [Usage, Assessment]
- 259 2. Use advanced algorithmic techniques (e.g., randomization, approximation) to solve real problems. [Usage]
- 260 3. Apply advanced analysis techniques (e.g., amortized, probabilistic, etc.) to algorithms. [Usage,
261 Assessment]

1 **Architecture and Organization (AR)**

2 Computing professionals should not regard the computer as just a black box that executes
3 programs by magic. AR-Architecture and Organization builds on SF-Systems Fundamentals to
4 develop a deeper understanding of the hardware environment upon which all of computing is
5 based, and the interface it provides to higher software layers. Students should acquire an
6 understanding and appreciation of a computer system's functional components, their
7 characteristics, performance, and interactions, and, in particular, the challenge of harnessing
8 parallelism to sustain performance improvements now and into the future. Students need to
9 understand computer architecture to develop programs that can achieve high performance
10 through a programmer's awareness of parallelism and latency. In selecting a system to use,
11 students should be able to understand the tradeoff among various components, such as CPU clock
12 speed, cycles per instruction, memory size, and average memory access time.

13 The learning outcomes specified for these topics correspond primarily to the core and are
14 intended to support programs that elect to require only the minimum 16 hours of computer
15 architecture of their students. For programs that want to teach more than the minimum, the same
16 AR topics can be treated at a more advanced level by implementing a two-course sequence. For
17 programs that want to cover the elective topics, those topics can be introduced within a two-
18 course sequence and/or be treated in a more comprehensive way in a third course.

19

20 **AR. Architecture and Organization (0 Core-Tier 1 hours, 16 Core-Tier 2 hours)**

| | Core-Tier 1 hours | Core-Tier 2 Hours | Includes Elective |
|--|-------------------|-------------------|-------------------|
| AR/Digital logic and digital systems | | 3 | N |
| AR/Machine level representation of data | | 3 | N |
| AR/Assembly level machine organization | | 6 | N |
| AR/Memory system organization and architecture | | 3 | N |
| AR/Interfacing and communication | | 1 | N |
| AR/Functional organization | | | Y |
| AR/Multiprocessing and alternative architectures | | | Y |
| AR/Performance enhancements | | | Y |

21

22 **AR/Digital logic and digital systems**

23 *[3 Core-Tier2 hours]*

24 *Topics:*

- 25 • Overview and history of computer architecture
- 26 • Combinational vs. sequential logic/Field programmable gate arrays as a fundamental combinational +
- 27 sequential logic building block
- 28 • Multiple representations/layers of interpretation (hardware is just another layer)
- 29 • Computer-aided design tools that process hardware and architectural representations
- 30 • Register transfer notation/Hardware Description Language (Verilog/VHDL)
- 31 • Physical constraints (gate delays, fan-in, fan-out, energy/power)

32

33 *Learning outcomes:*

- 34 1. Describe the progression of computer technology components from vacuum tubes to VLSI, from
- 35 mainframe computer architectures to the organization of warehouse-scale computers [Familiarity].
- 36 2. Comprehend the trend of modern computer architectures towards multi-core and that parallelism is inherent
- 37 in all hardware systems [Familiarity].
- 38 3. Explain the implications of the “power wall” in terms of further processor performance improvements and
- 39 the drive towards harnessing parallelism [Familiarity].
- 40 4. Articulate that there are many equivalent representations of computer functionality, including logical
- 41 expressions and gates, and be able to use mathematical expressions to describe the functions of simple
- 42 combinational and sequential circuits [Familiarity].
- 43 5. Design the basic building blocks of a computer: arithmetic-logic unit (gate-level), registers (gate-level),
- 44 central processing unit (register transfer-level), memory (register transfer-level) [Usage].
- 45 6. Use CAD tools for capture, synthesis, and simulation to evaluate simple building blocks (e.g., arithmetic-
- 46 logic unit, registers, movement between registers) of a simple computer design [Usage].

- 47 7. Evaluate the functional and timing diagram behavior of a simple processor implemented at the logic circuit
48 level [Assessment].
49

50 **AR/Machine-level representation of data**

51 *[3 Core-Tier2 hours]*

52 *Topics:*

- 53 • Bits, bytes, and words
- 54 • Numeric data representation and number bases
- 55 • Fixed- and floating-point systems
- 56 • Signed and twos-complement representations
- 57 • Representation of non-numeric data (character codes, graphical data)
- 58 • Representation of records and arrays

59
60 *Learning outcomes:*

- 61 1. Explain why everything is data, including instructions, in computers [Familiarity].
- 62 2. Explain the reasons for using alternative formats to represent numerical data [Familiarity].
- 63 3. Describe how negative integers are stored in sign-magnitude and twos-complement representations
64 [Familiarity].
- 65 4. Explain how fixed-length number representations affect accuracy and precision [Familiarity].
- 66 5. Describe the internal representation of non-numeric data, such as characters, strings, records, and arrays
67 [Familiarity].
- 68 6. Convert numerical data from one format to another [Usage].
- 69 7. Write simple programs at the assembly/machine level for string processing and manipulation [Usage].
70

71 **AR/Assembly level machine organization**

72 *[6 Core-Tier2 hours]*

73 *Topics:*

- 74 • Basic organization of the von Neumann machine
- 75 • Control unit; instruction fetch, decode, and execution
- 76 • Instruction sets and types (data manipulation, control, I/O)
- 77 • Assembly/machine language programming
- 78 • Instruction formats
- 79 • Addressing modes
- 80 • Subroutine call and return mechanisms (xref PL/Language Translation and Execution)
- 81 • I/O and interrupts
- 82 • Heap vs. Static vs. Stack vs. Code segments
- 83 • Shared memory multiprocessors/multicore organization
- 84 • Introduction to SIMD vs. MIMD and the Flynn Taxonomy

85
86 *Learning outcomes:*

- 87 1. Explain the organization of the classical von Neumann machine and its major functional units [Familiarity].

- 88 2. Describe how an instruction is executed in a classical von Neumann machine, with extensions for threads,
89 multiprocessor synchronization, and SIMD execution [Familiarity].
90 3. Describe instruction level parallelism and hazards, and how they are managed in typical processor pipelines
91 [Familiarity].
92 4. Summarize how instructions are represented at both the machine level and in the context of a symbolic
93 assembler [Familiarity].
94 5. Demonstrate how to map between high-level language patterns into assembly/machine language notations
95 [Familiarity].
96 6. Explain different instruction formats, such as addresses per instruction and variable length vs. fixed length
97 formats [Familiarity].
98 7. Explain how subroutine calls are handled at the assembly level [Familiarity].
99 8. Explain the basic concepts of interrupts and I/O operations [Familiarity].
100 9. Explain how subroutine calls are handled at the assembly level [Familiarity].
101 10. Write simple assembly language program segments [Usage].
102 11. Show how fundamental high-level programming constructs are implemented at the machine-language level
103 [Usage].
104

105 **AR/Memory system organization and architecture**

106 *[3 Core-Tier2 hours]*

107 [Cross-reference OS/Memory Management--Virtual Machines]

108 **Topics:**

- 109 • Storage systems and their technology
- 110 • Memory hierarchy: importance of temporal and spatial locality
- 111 • Main memory organization and operations
- 112 • Latency, cycle time, bandwidth, and interleaving
- 113 • Cache memories (address mapping, block size, replacement and store policy)
- 114 • Multiprocessor cache consistency/Using the memory system for inter-core synchronization/atomic memory
115 operations
- 116 • Virtual memory (page table, TLB)
- 117 • Fault handling and reliability
- 118 • Error coding, data compression, and data integrity (xref SF/Reliability through Redundancy)

119 **Learning outcomes:**

- 121 1. Identify the main types of memory technology [Familiarity].
122 2. Explain the effect of memory latency on running time [Familiarity].
123 3. Describe how the use of memory hierarchy (cache, virtual memory) is used to reduce the effective memory
124 latency [Familiarity].
125 4. Describe the principles of memory management [Familiarity].
126 5. Explain the workings of a system with virtual memory management [Familiarity].
127 6. Compute Average Memory Access Time under a variety of memory system configurations and workload
128 assumptions [Usage].
129
130

131 **AR/Interfacing and communication**

132 *[1 Core-Tier2 hour]*

133 [Cross-reference OS Knowledge Area for a discussion of the operating system view of
134 input/output processing and management. The focus here is on the hardware mechanisms for
135 supporting device interfacing and processor-to-processor communications.]

136 **Topics:**

- 137 • I/O fundamentals: handshaking, buffering, programmed I/O, interrupt-driven I/O
- 138 • Interrupt structures: vectored and prioritized, interrupt acknowledgment
- 139 • External storage, physical organization, and drives
- 140 • Buses: bus protocols, arbitration, direct-memory access (DMA)
- 141 • Introduction to networks: networks as another layer of access hierarchy
- 142 • Multimedia support
- 143 • RAID architectures

144
145 **Learning outcomes:**

- 146 1. Explain how interrupts are used to implement I/O control and data transfers [Familiarity].
 - 147 2. Identify various types of buses in a computer system [Familiarity].
 - 148 3. Describe data access from a magnetic disk drive [Familiarity].
 - 149 4. Compare common network organizations, such as ethernet/bus, ring, switched vs. routed [Familiarity].
 - 150 5. Identify interfaces needed for multimedia support, from storage, through network, to memory and display
 - 151 [Familiarity].
 - 152 6. Describe the advantages and limitations of RAID architectures [Familiarity].
- 153

154 **AR/Functional organization**

155 *[Elective]*

156 [Note: elective for computer scientist; would be core for computer engineering curriculum]

157 **Topics:**

- 158 • Implementation of simple datapaths, including instruction pipelining, hazard detection and resolution
- 159 • Control unit: hardwired realization vs. microprogrammed realization
- 160 • Instruction pipelining
- 161 • Introduction to instruction-level parallelism (ILP)

162
163 **Learning outcomes:**

- 164 1. Compare alternative implementation of datapaths [Familiarity].
 - 165 2. Discuss the concept of control points and the generation of control signals using hardwired or
 - 166 microprogrammed implementations [Familiarity].
 - 167 3. Explain basic instruction level parallelism using pipelining and the major hazards that may occur
 - 168 [Familiarity].
 - 169 4. Design and implement a complete processor, including datapath and control [Usage].
 - 170 5. Determine, for a given processor and memory system implementation, the average cycles per instruction
 - 171 [Assessment].
- 172

173 **AR/Multiprocessing and alternative architectures**

174 *[Elective]*

175 [Cross-reference PD/Parallel Architecture: The view here is on the hardware implementation of
176 SIMD and MIMD architectures; in PD/Parallel Architecture, it is on the way that algorithms can
177 be matched to the underlying hardware capabilities for these kinds of parallel processing
178 architectures.]

179 *Topics:*

- 180 • Power Law
- 181 • Example SIMD and MIMD instruction sets and architectures
- 182 • Interconnection networks (hypercube, shuffle-exchange, mesh, crossbar)
- 183 • Shared multiprocessor memory systems and memory consistency
- 184 • Multiprocessor cache coherence

185

186 *Learning outcomes:*

- 187 1. Discuss the concept of parallel processing beyond the classical von Neumann model [Familiarity].
- 188 2. Describe alternative architectures such as SIMD and MIMD [Familiarity].
- 189 3. Explain the concept of interconnection networks and characterize different approaches [Familiarity].
- 190 4. Discuss the special concerns that multiprocessing systems present with respect to memory management and
191 describe how these are addressed [Familiarity].
- 192 5. Describe the differences between memory backplane, processor memory interconnect, and remote memory
193 via networks [Familiarity].

194

195 **AR/Performance enhancements**

196 *[Elective]*

197 *Topics:*

- 198 • Superscalar architecture
- 199 • Branch prediction, Speculative execution, Out-of-order execution
- 200 • Prefetching
- 201 • Vector processors and GPUs
- 202 • Hardware support for Multithreading
- 203 • Scalability
- 204 • Alternative architectures, such as VLIW/EPIC, and Accelerators and other kinds of Special-Purpose
205 Processors

206

207 *Learning outcomes:*

- 208 1. Describe superscalar architectures and their advantages [Familiarity].
- 209 2. Explain the concept of branch prediction and its utility [Familiarity].
- 210 3. Characterize the costs and benefits of prefetching [Familiarity].
- 211 4. Explain speculative execution and identify the conditions that justify it [Familiarity].
- 212 5. Discuss the performance advantages that multithreading offered in an architecture along with the factors
213 that make it difficult to derive maximum benefits from this approach [Familiarity].
- 214 6. Describe the relevance of scalability to performance [Familiarity].

1 **Computational Science (CN)**

2 Computational Science is a field of applied computer science, that is, the application of computer
3 science to solve problems across a range of disciplines. According to the book “Introduction to
4 Computational Science”, Shiflet & Shiflet offer the following definition: “the field of
5 computational science combines computer simulation, scientific visualization, mathematical
6 modeling, computer programming and data structures, networking, database design, symbolic
7 computation, and high performance computing with various disciplines.” Computer science,
8 which largely focuses on the theory, design, and implementation of algorithms for manipulating
9 data and information, can trace its roots to the earliest devices used to assist people in
10 computation over four thousand years ago. Various systems were created and used to calculate
11 astronomical positions. Ada Lovelace’s programming achievement was intended to calculate
12 Bernoulli numbers. In the late nineteenth century, mechanical calculators became available, and
13 were immediately put to use by scientists. The needs of scientists and engineers for computation
14 have long driven research and innovation in computing. As computers increase in their problem-
15 solving power, computational science has grown in both breadth and importance. It is a
16 discipline in its own right (President’s Information Technology Advisory Committee, 2005, page
17 13) and is considered to be “one of the five college majors on the rise” (Fischer and Gleen, “5
18 College Majors on the Rise”, The Chronicle of Higher Education, 2009.) An amazing assortment
19 of sub-fields have arisen under the umbrella of Computational Science, including computational
20 biology, computational chemistry, computational mechanics, computational archeology,
21 computational finance, computational sociology and computational forensics.

22 Some fundamental concepts of computational science are germane to every computer scientist,
23 and computational science topics are extremely valuable components of an undergraduate
24 program in computer science. This area offers exposure to many valuable ideas and techniques,
25 including precision of numerical representation, error analysis, numerical techniques, parallel
26 architectures and algorithms, modeling and simulation, information visualization, software
27 engineering, and optimization. At the same time, students who take courses in this area have an
28 opportunity to apply these techniques in a wide range of application areas, such as: molecular
29 and fluid dynamics, celestial mechanics, economics, biology, geology, medicine, and social
30 network analysis. Many of the techniques used in these areas require advanced mathematics such

31 as calculus, differential equations, and linear algebra. The descriptions here assume that students
32 have acquired the needed mathematical background elsewhere.

33 In the computational science community, the terms *run*, *modify*, and *create* are often used to
34 describe levels of understanding. This chapter follows the conventions of other chapters in this
35 volume and uses the terms *familiarity*, *usage*, and *assessment*.

36

37 **CN. Computational Science (1 Core-Tier1 hours, 0 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| CN/Fundamentals | <i>1</i> | | N |
| CN/Modeling and Simulation | | | Y |
| CN/Processing | | | Y |
| CN/Interactive Visualization | | | Y |
| CN/Data, Information, and Knowledge | | | Y |

38

39

40 **CN/Fundamentals**

41 *[1 Core-Tier1 hours]*

42 Abstraction is a fundamental concept in computer science. A principal approach to computing is
43 to abstract the real world, create a model that can be simulated on a machine. The roots of
44 computer science can be traced to this approach, modeling things such as trajectories of artillery
45 shells and the modeling cryptographic protocols, both of which pushed the development of early
46 computing systems in the early and mid-1940's.

47 Modeling and simulation are essential topics for computational science. Any introduction to
48 computational science would either include or presume an introduction to computing. Topics
49 relevant to computational science include fundamental concepts in program construction
50 (SDF/Fundamental Programming Concepts), algorithm design (SDF/Algorithms and Design),
51 program testing (SDF/Development Methods), data representations (AR/Machine
52 Representation of Data), and basic computer architecture (AR/Memory System Organization and
53 Architecture). In addition, a general set of modeling and simulation techniques, data
54 visualization methods, and software testing and evaluation mechanisms are also important CN
55 fundamentals.

56 *Topics:*

- 57 • Models as abstractions of situations
- 58 • Simulations as dynamic modeling
- 59 • Simulation techniques and tools, such as physical simulations, human-in-the-loop guided simulations, and
60 virtual reality.
- 61 • Foundational approaches to validating models
- 62

63 *Learning Outcomes:*

- 64 1. Explain the concept of modeling and the use of abstraction that allows the use of a machine to solve a
65 problem. [Familiarity]
- 66 2. Describe the relationship between modeling and simulation, i.e., thinking of simulation as dynamic
67 modeling. [Familiarity]
- 68 3. Create a simple, formal mathematical model of a real-world situation and use that model in a simulation.
69 [Usage]
- 70 4. Differentiate among the different types of simulations, including physical simulations, human-guided
71 simulations, and virtual reality. [Familiarity]
- 72 5. Describe several approaches to validating models. [Familiarity]
- 73

74 **CN/Modeling and Simulation**

75 *[Elective]*

76 *Topics:*

- 77 • Purpose of modeling and simulation including optimization; supporting decision making, forecasting,
78 safety considerations; for training and education.
- 79 • Tradeoffs including performance, accuracy, validity, and complexity.
- 80 • The simulation process; identification of key characteristics or behaviors, simplifying assumptions;
81 validation of outcomes.

- 82 • Model building: use of mathematical formula or equation, graphs, constraints; methodologies and
- 83 techniques; use of time stepping for dynamic systems.
- 84 • Formal models and modeling techniques: mathematical descriptions involving simplifying assumptions
- 85 and avoiding detail. The descriptions use fundamental mathematical concepts such as set and function.
- 86 Random numbers. Examples of techniques including:
- 87 ○ Monte Carlo methods
- 88 ○ Stochastic processes
- 89 ○ Queuing theory
- 90 ○ Petri nets and colored Petri nets
- 91 ○ Graph structures such as directed graphs, trees, networks
- 92 ○ Games, game theory, the modeling of things using game theory
- 93 ○ Linear programming and its extensions
- 94 ○ Dynamic programming
- 95 ○ Differential equations: ODE, PDE
- 96 ○ Non-linear techniques
- 97 ○ State spaces and transitions
- 98 • Assessing and evaluating models and simulations in a variety of contexts; verification and validation of
- 99 models and simulations.
- 100 • Important application areas including health care and diagnostics, economics and finance, city and urban
- 101 planning, science, and engineering.
- 102 • Software in support of simulation and modeling; packages, languages.
- 103

104 ***Learning Outcomes:***

- 105 1. Explain and give examples of the benefits of simulation and modeling in a range of important application
- 106 areas. [Familiarity]
- 107 2. Demonstrate the ability to apply the techniques of modeling and simulation to a range of problem areas.
- 108 [Usage]
- 109 3. Explain the constructs and concepts of a particular modeling approach. [Familiarity]
- 110 4. Explain the difference between validation and verification of a model; demonstrate the difference with
- 111 specific examples¹. [Assessment]
- 112 5. Verify and validate the results of a simulation. [Assessment]
- 113 6. Evaluate a simulation, highlighting the benefits and the drawbacks. [Assessment]
- 114 7. Choose an appropriate modeling approach for a given problem or situation. [Assessment]
- 115 8. Compare results from different simulations of the same situation and explain any differences. [Assessment]
- 116 9. Infer the behavior of a system from the results of a simulation of the system. [Assessment]
- 117 10. Extend or adapt an existing model to a new situation. [Assessment]
- 118

119

¹ *Verification* means that the computations of the model are correct. If we claim to compute total time, for example, the computation actually does that. *Validation* asks whether the model matches the real situation.

120 CN/Processing

121 *[Elective]*

122 The processing topic area includes numerous topics from other knowledge areas. Specifically,
123 coverage of processing should include a discussion of hardware architectures, including parallel
124 systems, memory hierarchies, and interconnections among processors. These are covered in
125 AR/Interfacing and Communication, AR/Multiprocessing and Alternative Architectures,
126 AR/Performance Enhancements.

127 *Topics:*

- 128 • Fundamental programming concepts:
 - 129 ○ The concept of an algorithm consisting of a finite number of well-defined steps, each of which
 - 130 completes in a finite amount of time, as does the entire process.
 - 131 ○ Examples of well-known algorithms such as sorting and searching.
 - 132 ○ The concept of analysis as understanding what the problem is really asking, how a problem can be
 - 133 approached using an algorithm, and how information is represented so that a machine can process
 - 134 it.
 - 135 ○ The development or identification of a workflow.
 - 136 ○ The process of converting an algorithm to machine-executable code.
 - 137 ○ Software processes including lifecycle models, requirements, design, implementation, verification
 - 138 and maintenance.
 - 139 ○ Machine representation of data computer arithmetic, and numerical methods, specifically
 - 140 sequential and parallel architectures and computations.
- 141 • Fundamental properties of parallel and distributed computation:
 - 142 ○ Bandwidth.
 - 143 ○ Latency.
 - 144 ○ Scalability.
 - 145 ○ Granularity.
 - 146 ○ Parallelism including task, data, and event parallelism.
 - 147 ○ Parallel architectures including processor architectures, memory and caching.
 - 148 ○ Parallel programming paradigms including threading, message passing, event driven techniques,
 - 149 parallel software architectures, and MapReduce.
 - 150 ○ Grid computing.
 - 151 ○ The impact of architecture on computational time.
 - 152 ○ Total time to science curve for parallelism: continuum of things.
- 153 • Computing costs, e.g., the cost of re-computing a value vs. the cost of storing and lookup.
- 154

155 *Learning Outcomes:*

- 156 1. Explain the characteristics and defining properties of algorithms and how they relate to machine
- 157 processing. [Familiarity]
- 158 2. Analyze simple problem statements to identify relevant information and select appropriate processing to
- 159 solve the problem. [Assessment]
- 160 3. Identify or sketch a workflow for an existing computational process such as the creation of a graph based
- 161 on experimental data. [Familiarity]
- 162 4. Describe the process of converting an algorithm to machine-executable code. [Familiarity]
- 163 5. Summarize the phases of software development and compare several common lifecycle models.
- 164 [Familiarity]
- 165 6. Explain how data is represented in a machine. Compare representations of integers to floating point
- 166 numbers. Describe underflow, overflow, round off, and truncation errors in data representations.
- 167 [Familiarity]

- 168 7. Apply standard numerical algorithms to solve ODEs and PDEs. Use computing systems to solve systems of
169 equations. [Usage]
170 8. Describe the basic properties of bandwidth, latency, scalability and granularity. [Familiarity]
171 9. Describe the levels of parallelism including task, data, and event parallelism. [Familiarity]
172 10. Compare and contrast parallel programming paradigms recognizing the strengths and weaknesses of each.
173 [Assessment]
174 11. Identify the issues impacting correctness and efficiency of a computation. [Familiarity]
175 12. Design, code, test and debug programs for a parallel computation. [Usage]
176

177 **CN/Interactive Visualization**

178 *[Elective]*

179 This sub-area is related to modeling and simulation. Most topics are discussed in detail in other
180 knowledge areas in this document. There are many ways to present data and information,
181 including immersion, realism, variable perspectives; haptics and heads-up displays, sonification,
182 and gesture mapping.

183 Interactive visualization in general requires understanding of human perception (GV/Basics);
184 graphics pipelines, geometric representations and data structures (GV/Fundamental Concepts);
185 2D and 3D rendering, surface and volume rendering (GV/Rendering, GV/Modeling, and
186 GV/Advanced Rendering); and the use of APIs for developing user interfaces using standard
187 input components such as menus, sliders, and buttons; and standard output components for data
188 display, including charts, graphs, tables, and histograms (HCI/GUI Construction, HCI/GUI
189 Programming).

190 *Topics:*

- 191 • Principles of data visualization.
- 192 • Graphing and visualization algorithms.
- 193 • Image processing techniques.
- 194 • Scalability concerns.
- 195

196 *Learning Outcomes:*

- 197 1. Compare common computer interface mechanisms with respect to ease-of-use, learnability, and cost.
198 [Assessment]
- 199 2. Use standard APIs and tools to create visual displays of data, including graphs, charts, tables, and
200 histograms. [Usage]
- 201 3. Describe several approaches to using a computer as a means for interacting with and processing data.
202 [Familiarity]
- 203 4. Extract useful information from a dataset. [Assessment]
- 204 5. Analyze and select visualization techniques for specific problems. [Assessment]
- 205 6. Describe issues related to scaling data analysis from small to large data sets. [Familiarity]
206

207

208 **CN/Data, Information, and Knowledge**

209 *[Elective]*

210 Many topics are discussed in detail in other knowledge areas in this document, specifically
211 Information Management (IM/Information Management Concepts, IM/Database Systems, and
212 IM/Data Modeling), Algorithms and Complexity (AL/Basic Analysis, AL/Fundamental Data
213 Structures and Algorithms), and Software Development Fundamentals (SDF/Fundamental
214 Programming Concepts, SDF/Development Methods).

215 *Topics:*

- 216 • Content management models, frameworks, systems, design methods (as in IM. Information Management).
- 217 • Digital representations of content including numbers, text, images (e.g., raster and vector), video (e.g.,
218 QuickTime, MPEG2, MPEG4), audio (e.g., written score, MIDI, sampled digitized sound track) and
219 animations; complex/composite/aggregate objects; FRBR.
- 220 • Digital content creation/capture and preservation, including digitization, sampling, compression,
221 conversion, transformation/translation, migration/emulation, crawling, harvesting.
- 222 • Content structure / management, including digital libraries and static/dynamic/stream aspects for:
 - 223 ○ Data: data structures, databases.
 - 224 ○ Information: document collections, multimedia pools, hyperbases (hypertext, hypermedia),
225 catalogs, repositories.
 - 226 ○ Knowledge: ontologies, triple stores, semantic networks, rules.
- 227 • Processing and pattern recognition, including indexing, searching (including: queries and query languages;
228 central / federated / P2P), retrieving, clustering, classifying/categorizing, analyzing/mining/extracting,
229 rendering, reporting, handling transactions.
- 230 • User / society support for presentation and interaction, including browse, search, filter, route, visualize,
231 share, collaborate, rate, annotate, personalize, recommend.
- 232 • Modeling, design, logical and physical implementation, using relevant systems/software.
- 233

234 *Learning Outcomes:*

- 235 1. Identify all of the data, information, and knowledge elements and related organizations, for a computational
236 science application. [Assessment]
- 237 2. Describe how to represent data and information for processing. [Familiarity]
- 238 3. Describe typical user requirements regarding that data, information, and knowledge. [Familiarity]
- 239 4. Select a suitable system or software implementation to manage data, information, and knowledge.
240 [Assessment]
- 241 5. List and describe the reports, transactions, and other processing needed for a computational science
242 application. [Familiarity]
- 243 6. Compare and contrast database management, information retrieval, and digital library systems with regard
244 to handling typical computational science applications. [Assessment]
- 245 7. Design a digital library for some computational science users / societies, with appropriate content and
246 services. [Usage]

1 **Discrete Structures (DS)**

2 Discrete structures are foundational material for computer science. By foundational we mean that
3 relatively few computer scientists will be working primarily on discrete structures, but that many
4 other areas of computer science require the ability to work with concepts from discrete
5 structures. Discrete structures include important material from such areas as set theory, logic,
6 graph theory, and probability theory.

7 The material in discrete structures is pervasive in the areas of data structures and algorithms but
8 appears elsewhere in computer science as well. For example, an ability to create and understand
9 a proof—either a formal symbolic proof or a less formal but still mathematically rigorous
10 argument—is important in virtually every area of computer science, including (to name just a
11 few) formal specification, verification, databases, and cryptography. Graph theory concepts are
12 used in networks, operating systems, and compilers. Set theory concepts are used in software
13 engineering and in databases. Probability theory is used in intelligent systems, networking, and a
14 number of computing applications.

15 Given that discrete structures serves as a foundation for many other areas in computing, it is
16 worth noting that the boundary between discrete structures and other areas, particularly
17 Algorithms and Complexity, Software Development Fundamentals, Programming Languages,
18 and Intelligent Systems, may not always be crisp. Indeed, different institutions may choose to
19 organize the courses in which they cover this material in very different ways. Some institutions
20 may cover these topics in one or two focused courses with titles like "discrete structures" or
21 "discrete mathematics", whereas others may integrate these topics in courses on programming,
22 algorithms, and/or artificial intelligence. Combinations of these approaches are also prevalent
23 (e.g., covering many of these topics in a single focused introductory course and covering the
24 remaining topics in more advanced topical courses).

25

26 **DS. Discrete Structures (37 Core-Tier1 hours, 4 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| DS/Sets, Relations, and Functions | 4 | | N |
| DS/Basic Logic | 9 | | N |
| DS/Proof Techniques | 10 | 1 | N |
| DS/Basics of Counting | 5 | | N |
| DS/Graphs and Trees | 3 | 1 | N |
| DS/Discrete Probability | 6 | 2 | N |

27

28 **DS/Sets, Relations, and Functions**

29 *[4 Core-Tier1 hours]*

30 *Topics:*

31 [Core-Tier1]

- 32
- 33 • Sets
 - 34 ○ Venn diagrams
 - 35 ○ Union, intersection, complement
 - 36 ○ Cartesian product
 - 37 ○ Power sets
 - 38 ○ Cardinality of finite sets
 - 39 • Relations
 - 40 ○ Reflexivity, symmetry, transitivity
 - 41 ○ Equivalence relations, partial orders
 - 42 • Functions
 - 43 ○ Surjections, injections, bijections
 - 44 ○ Inverses
 - 45 ○ Composition

46 *Learning Outcomes:*

47 [Core-Tier1]

- 48
- 49 1. Explain with examples the basic terminology of functions, relations, and sets. [Familiarity]
 - 50 2. Perform the operations associated with sets, functions, and relations. [Usage]
 - 51 3. Relate practical examples to the appropriate set, function, or relation model, and interpret the associated operations and terminology in context. [Assessment]
- 52

53

54 **DS/Basic Logic**

55 *[9 Core-Tier1 hours]*

56 **Topics:**

57 [Core-Tier1]

- 58 • Propositional logic (cross-reference: Propositional logic is also reviewed in IS/Knowledge Based Reasoning)
- 59
- 60 • Logical connectives
- 61 • Truth tables
- 62 • Normal forms (conjunctive and disjunctive)
- 63 • Validity of well-formed formula
- 64 • Propositional inference rules (concepts of modus ponens and modus tollens)
- 65 • Predicate logic
- 66 ◦ Universal and existential quantification
- 67 • Limitations of propositional and predicate logic (e.g., expressiveness issues)
- 68

69 **Learning Outcomes:**

70 [Core-Tier1]

- 71 1. Convert logical statements from informal language to propositional and predicate logic expressions. [Usage]
- 72
- 73 2. Apply formal methods of symbolic propositional and predicate logic, such as calculating validity of formulae and computing normal forms. [Usage]
- 74
- 75 3. Use the rules of inference to construct proofs in propositional and predicate logic. [Usage]
- 76 4. Describe how symbolic logic can be used to model real-life situations or applications, including those arising in computing contexts such as software analysis (e.g., program correctness), database queries, and algorithms. [Usage]
- 77
- 78
- 79 5. Apply formal logic proofs and/or informal, but rigorous, logical reasoning to real problems, such as predicting the behavior of software or solving problems such as puzzles. [Usage]
- 80
- 81 6. Describe the strengths and limitations of propositional and predicate logic. [Familiarity]
- 82

83 **DS/Proof Techniques**

84 *[10 Core-Tier1 hours, 1 Core-Tier2 hour]*

85 **Topics:**

86 [Core-Tier1]

- 87 • Notions of implication, equivalence, converse, inverse, contrapositive, negation, and contradiction
- 88 • The structure of mathematical proofs
- 89 • Direct proofs
- 90 • Disproving by counterexample
- 91 • Proof by contradiction
- 92 • Induction over natural numbers
- 93 • Structural induction
- 94 • Weak and strong induction (i.e., First and Second Principle of Induction)
- 95 • Recursive mathematical definitions
- 96

97 [Core-Tier2]

- 98 • Well orderings
- 99

100 **Learning Outcomes:**

101 [Core-Tier1]

- 102 1. Identify the proof technique used in a given proof. [Familiarity]
- 103 2. Outline the basic structure of each proof technique described in this unit. [Usage]
- 104 3. Apply each of the proof techniques correctly in the construction of a sound argument. [Usage]
- 105 4. Determine which type of proof is best for a given problem. [Assessment]
- 106 5. Explain the parallels between ideas of mathematical and/or structural induction to recursion and recursively defined structures. [Assessment]
- 107 6. Explain the relationship between weak and strong induction and give examples of the appropriate use of each. [Assessment]

110 [Core-Tier2]

- 112 7. State the well-ordering principle and its relationship to mathematical induction. [Familiarity]
- 113

114 **DS/Basics of Counting**

115 *[5 Core-Tier1 hours]*

116 **Topics:**

117 [Core-Tier1]

- 118 • Counting arguments
 - 119 ○ Set cardinality and counting
 - 120 ○ Sum and product rule
 - 121 ○ Inclusion-exclusion principle
 - 122 ○ Arithmetic and geometric progressions
- 123 • The pigeonhole principle
- 124 • Permutations and combinations
 - 125 ○ Basic definitions
 - 126 ○ Pascal's identity
 - 127 ○ The binomial theorem
- 128 • Solving recurrence relations (cross-reference: AL/Basic Analysis)
 - 129 ○ An example of a simple recurrence relation, such as Fibonacci numbers
 - 130 ○ Other examples, showing a variety of solutions
- 131 • Basic modular arithmetic
- 132

133 **Learning Outcomes:**

134 [Core-Tier1]

- 135 1. Apply counting arguments, including sum and product rules, inclusion-exclusion principle and arithmetic/geometric progressions. [Usage]
- 136 2. Apply the pigeonhole principle in the context of a formal proof. [Usage]
- 137 3. Compute permutations and combinations of a set, and interpret the meaning in the context of the particular application. [Usage]
- 138
- 139

- 140 4. Map real-world applications to appropriate counting formalisms, such as determining the number of ways
 141 to arrange people around a table, subject to constraints on the seating arrangement, or the number of ways
 142 to determine certain hands in cards (e.g., a full house). [Usage]
 143 5. Solve a variety of basic recurrence relations. [Usage]
 144 6. Analyze a problem to determine underlying recurrence relations. [Usage]
 145 7. Perform computations involving modular arithmetic. [Usage]
 146

147 **DS/Graphs and Trees**

148 *[3 Core-Tier1 hours, 1 Core-Tier2 hour]*

149 (cross-reference: AL/Fundamental Data Structures and Algorithms, especially with relation to
 150 graph traversal strategies)

151 **Topics:**

152 [Core-Tier1]

- 153 • Trees
 - 154 ○ Properties
 - 155 ○ Traversal strategies
- 156 • Undirected graphs
- 157 • Directed graphs
- 158 • Weighted graphs

159

160 [Core-Tier2]

- 161 • Spanning trees/forests
- 162 • Graph isomorphism

163

164 **Learning Outcomes:**

165 [Core-Tier1]

- 166 1. Illustrate by example the basic terminology of graph theory, and some of the properties and special cases of
 167 each type of graph/tree. [Familiarity]
- 168 2. Demonstrate different traversal methods for trees and graphs, including pre, post, and in-order traversal of
 169 trees. [Usage]
- 170 3. Model *a variety of* real-world problems in computer science using appropriate forms of graphs and trees,
 171 such as representing a network topology or the organization of a hierarchical file system. [Usage]
- 172 4. Show how concepts from graphs and trees appear in data structures, algorithms, proof techniques
 173 (structural induction), and counting. [Usage]

174

175 [Core-Tier2]

- 176 5. Explain how to construct a spanning tree of a graph. [Usage]
- 177 6. Determine if two graphs are isomorphic. [Usage]

178

179 **DS/Discrete Probability**

180 *[6 Core-Tier1 hours, 2 Core-Tier2 hour]*

181 (Cross-reference IS/Basic Knowledge Representation and Reasoning, which includes a review of
182 basic probability)

183 **Topics:**

184 [Core-Tier1]

- 185 • Finite probability space, events
- 186 • Axioms of probability and probability measures
- 187 • Conditional probability, Bayes' theorem
- 188 • Independence
- 189 • Integer random variables (Bernoulli, binomial)
- 190 • Expectation, including Linearity of Expectation

191

192 [Core-Tier2]

- 193 • Variance
- 194 • Conditional Independence

195

196 **Learning Outcomes:**

197 [Core-Tier1]

- 198 1. Calculate probabilities of events and expectations of random variables for elementary problems such as
199 games of chance. [Usage]
- 200 2. Differentiate between dependent and independent events. [Usage]
- 201 3. Identify a case of the binomial distribution and compute a probability using that distribution. [Usage]
- 202 4. Make a probabilistic inference in a real-world problem using Bayes' theorem to determine the probability
203 of a hypothesis given evidence. [Usage]
- 204 5. Apply the tools of probability to solve problems such as the average case analysis of algorithms or
205 analyzing hashing. [Usage]

206

207 [Core-Tier2]

- 208 6. Compute the variance for a given probability distribution. [Usage]
- 209 7. Explain how events that are independent can be conditionally dependent (and vice-versa). Identify real-
210 world examples of such cases. [Usage]

1 **Graphics and Visualization (GV)**

2 *Computer graphics* is the term commonly used to describe the computer generation and
3 manipulation of images. It is the science of enabling visual communication through computation.
4 Its uses include cartoons, film special effects, video games, medical imaging, engineering, as
5 well as scientific, information, and knowledge visualization. Traditionally, graphics at the
6 undergraduate level has focused on rendering, linear algebra, and phenomenological approaches.
7 More recently, the focus has begun to include physics, numerical integration, scalability, and
8 special-purpose hardware, In order for students to become adept at the use and generation of
9 computer graphics, many implementation-specific issues must be addressed, such as file formats,
10 hardware interfaces, and application program interfaces. These issues change rapidly, and the
11 description that follows attempts to avoid being overly prescriptive about them. The area
12 encompassed by Graphics and Visual Computing (GV) is divided into several interrelated fields:

- 13 • **Fundamentals:** Computer graphics depends on an understanding of how humans use
14 vision to perceive information and how information can be rendered on a display device.
15 Every computer scientist should have some understanding of where and how graphics can
16 be appropriately applied and the fundamental processes involved in display rendering.
- 17 • **Modeling:** Information to be displayed must be encoded in computer memory in some
18 form, often in the form of a mathematical specification of shape and form.
- 19 • **Rendering:** Rendering is the process of displaying the information contained in a model.
- 20 • **Animation:** Animation is the rendering in a manner that makes images appear to move
21 and the synthesis or acquisition of the time variations of models.
- 22 • **Visualization.** The field of visualization seeks to determine and present underlying
23 correlated structures and relationships in data sets from a wide variety of application
24 areas. The prime objective of the presentation should be to communicate the information
25 in a dataset so as to enhance understanding
- 26 • **Computational Geometry:** Computational Geometry is the study of algorithms that are
27 stated in terms of geometry.

29 Graphics and Visualization is related to machine vision and image processing (in the Intelligent
30 Systems KA) and algorithms such as computational geometry, which can be found in the
31 Algorithms and Complexity KA. Topics in virtual reality can be found in the Human Computer
32 Interaction KA.

33 This description assumes students are familiar with fundamental concepts of data representation,
34 abstraction, and program implementation.

35

36 **GV. Graphics and Visualization (2 Core-Tier1 hours, 1 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--------------------------------|------------------|------------------|--------------------|
| GV/Fundamental Concepts | 2 | 1 | N |
| GV/Basic Rendering | | | Y |
| GV/Geometric Modeling | | | Y |
| GV/Advanced Rendering | | | Y |
| GV/Computer Animation | | | Y |
| GV/Visualization | | | Y |

37

38

39 **GV/Fundamental Concepts**

40 *[2 Core-Tier1 and 1 Core-Tier2 hours]*

41 For nearly every computer scientist and software developer, an understanding of how humans
42 interact with machines is essential. While these topics may be covered in a standard
43 undergraduate graphics course, they may also be covered in introductory computer science and
44 programming courses. Part of our motivation for including immediate and retained modes is that
45 these modes are analogous to polling vs. event driven programming. This is a fundamental
46 question in computer science: Is there a button object, or is there just the display of a button on
47 the screen? Note that most of the outcomes in this section are at the knowledge level, and many
48 of these topics are revisited in greater depth in later sections.

49 **Topics:**

50 [Core-Tier1]

- 51 • Applications of computer graphics: including user interfaces, game engines, cad, visualization, virtual
52 reality.
- 53 • Digitization of analog data and the limits of human perception, e.g., pixels for visual display, dots for laser
54 printers, and samples for audio (HCI Foundations)
- 55 • Use of standard graphics APIs for the construction of UIs and display of standard image formats (see HCI
56 GUI construction).
- 57 • Standard image formats, including lossless and lossy formats

58
59 [Core-Tier2]

- 60 • Additive and subtractive color models (CMYK and RGB) and why these provide a range of colors
- 61 • Tradeoffs between storing data and re-computing data as embodied by vector and raster representations of
62 images
- 63 • Animation as a sequence of still images
- 64 • Double buffering.

65 66 **Learning Outcomes:**

67 [Core-Tier1]

- 68 1. Identify common uses of computer graphics. [Familiarity]
- 69 2. Explain in general terms how analog signals can be reasonably represented by discrete samples, for
70 example, how images can be represented by pixels. [Familiarity]
- 71 3. Construct a simple user interface using a standard graphics API. [Usage]
- 72 4. Describe the differences between lossy and lossless image compression techniques, for example as
73 reflected in common graphics image file formats such as JPG, PNG, and GIF. [Familiarity]

74
75 [Core-Tier2]

- 76 5. Describe color models and their use in graphics display devices. [Familiarity]
- 77 6. Describe the tradeoffs between storing information vs. storing enough information to reproduce the
78 information, as in the difference between vector and raster rendering. [Familiarity]
- 79 7. Describe the basic process of producing continuous motion from a sequence of discrete frames (sometimes
80 called “flicker fusion”). [Familiarity]
- 81 8. Describe how double-buffering can remove flicker from animation. [Familiarity]

82

83 **GV/Basic Rendering**

84 *[Elective]*

85 This section describes basic rendering and fundamental graphics techniques that nearly every
86 undergraduate course in graphics will cover and that is essential for further study in graphics.
87 Sampling and anti-aliasing is related to the effect of digitization and appears in other areas of
88 computing, for example, in audio sampling.

89

90 **Topics:**

- 91 • Rendering in nature, i.e., the emission and scattering of light and its relation to numerical integration.
- 92 • Forward and backward rendering (i.e., ray-casting and rasterization).
- 93 • Polygonal representation.
- 94 • Basic radiometry, similar triangles, and projection model.
- 95 • Affine and coordinate system transformations.
- 96 • Ray tracing.
- 97 • Visibility and occlusion, including solutions to this problem such as depth buffering, Paiter’s algorithm,
98 and ray tracing.
- 99 • The forward and backward rendering equation.
- 100 • Simple triangle rasterization.
- 101 • Rendering with a shader-based API.
- 102 • Texture mapping, including minification and magnification (e.g., trilinear MIP-mapping).
- 103 • Application of spatial data structures to rendering.
- 104 • Sampling and anti-aliasing.
- 105 • Scene graphs and the graphics pipeline.
- 106

107 **Learning Outcomes:**

- 108 1. Discuss the light transport problem and its relation to numerical integration i.e., light is emitted, scatters
109 around the scene, and is measured by the eye; the form is an integral equation without analytic solution, but
110 we can approach it as numerical integration. [Familiarity]
- 111 2. Describe the basic graphics pipeline and how forward and backward rendering factor in this. [Familiarity]
- 112 3. Model simple graphics images. [Usage]
- 113 4. Derive linear perspective from similar triangles by converting points (x, y, z) to points (x/z, y/z, 1). [Usage]
- 114 5. Obtain 2-dimensional and 3-dimensional points by applying affine transformations. [Usage]
- 115 6. Apply 3-dimensional coordinate system and the changes required to extend 2D transformation operations to
116 handle transformations in 3D. [Usage]
- 117 7. Contrast forward and backward rendering. [Assessment]
- 118 8. Explain the concept and applications of texture mapping, sampling, and anti-aliasing. [Familiarity]
- 119 9. Explain the ray tracing – rasterization duality for the visibility problem. [Familiarity]
- 120 10. Implement simple procedures that perform transformation and clipping operations on simple 2-dimensional
121 images. [Usage]
- 122 11. Implement a simple real-time renderer using a rasterization API (e.g., OpenGL) using vertex buffers and
123 shaders. [Usage]
- 124 12. Compare and contrast the different rendering techniques. [Assessment]
- 125 13. Compute space requirements based on resolution and color coding. [Assessment]
- 126 14. Compute time requirements based on refresh rates, rasterization techniques. [Assessment]
- 127

128

129 **GV/Geometric Modeling**

130 *[Elective]*

131 *Topics:*

- 132 • Basic geometric operations such as intersection calculation and proximity tests
- 133 • Volumes, voxels, and point-based representations.
- 134 • Parametric polynomial curves and surfaces.
- 135 • Implicit representation of curves and surfaces.
- 136 • Approximation techniques such as polynomial curves, Bezier curves, spline curves and surfaces, and non-
- 137 uniform rational basis (NURB) spines, and level set method.
- 138 • Surface representation techniques including tessellation, mesh representation, mesh fairing, and mesh
- 139 generation techniques such as Delaunay triangulation, marching cubes, .
- 140 • Spatial subdivision techniques.
- 141 • Procedural models such as fractals, generative modeling, and L-systems.
- 142 • Graftals, cross referenced with programming languages (grammars to generated pictures).
- 143 • Elastically deformable and freeform deformable models.
- 144 • Subdivision surfaces.
- 145 • Multiresolution modeling.
- 146 • Reconstruction.
- 147 • Constructive Solid Geometry (CSG) representation.
- 148

149 *Learning Outcomes:*

- 150 1. Represent curves and surfaces using both implicit and parametric forms. [Usage]
- 151 2. Create simple polyhedral models by surface tessellation. [Usage]
- 152 3. Implement such algorithms as
- 153 4. Generate a mesh representation from an implicit surface. [Usage]
- 154 5. Generate a fractal model or terrain using a procedural method. [Usage]
- 155 6. Generate a mesh from data points acquired with a laser scanner. [Usage]
- 156 7. Construct CSG models from simple primitives, such as cubes and quadric surfaces. [Usage]
- 157 8. Contrast modeling approaches with respect to space and time complexity and quality of image.
- 158 [Assessment]
- 159

160 **GV/Advanced Rendering**

161 *[Elective]*

162 *Topics:*

- 163 • Solutions and approximations to the rendering equation, for example:
 - 164 ○ Distribution ray tracing and path tracing
 - 165 ○ Photon mapping
 - 166 ○ Bidirectional path tracing
 - 167 ○ Reyes (micropolygon) rendering
 - 168 ○ Metropolis light transport
- 169 • Considering the dimensions of time (motion blur), lens position (focus), and continuous frequency (color).
- 170 • Shadow mapping.
- 171 • Occlusion culling.
- 172 • Bidirectional Scattering Distribution function (BSDF) theory and microfacets.
- 173 • Subsurface scattering.

- 174 • Area light sources.
- 175 • Hierarchical depth buffering.
- 176 • The Light Field, image-based rendering.
- 177 • Non-photorealistic rendering.
- 178 • GPU architecture.
- 179 • Human visual systems including adaptation to light, sensitivity to noise, and flicker fusion.
- 180

181 **Learning Outcomes:**

- 182 1. Demonstrate how an algorithm estimates a solution to the rendering equation. [Assessment]
- 183 2. Prove the properties of a rendering algorithm, e.g., complete, consistent, and/or unbiased. [Assessment]
- 184 3. Analyze the bandwidth and computation demands of a simple algorithm. [Assessment]
- 185 4. Implement a non-trivial shading algorithm (e.g., toon shading, cascaded shadow maps) under a rasterization API. [Usage]
- 186 5. Discuss how a particular artistic technique might be implemented in a renderer. [Familiarity]
- 187 6. Explain how to recognize the graphics techniques used to create a particular image. [Familiarity]
- 188 7. Implement any of the specified graphics techniques using a primitive graphics system at the individual pixel level. [Usage]
- 189 8. Implement a ray tracer for scenes using a simple (e.g., Phong's) BRDF plus reflection and refraction.
- 190 [Usage]
- 191
- 192
- 193

194 **GV/Computer Animation**

195 **[Elective]**

196 **Topics:**

- 197 • Forward and inverse kinematics.
- 198 • Collision detection and response
- 199 • Procedural animation using noise, rules (boids/crowds), and particle systems.
- 200 • Skinning algorithms.
- 201 • Physics based motions including rigid body dynamics, physical particle systems, mass-spring networks for
- 202 cloth and flesh and hair.
- 203 • Key-frame animation.
- 204 • Splines.
- 205 • Data structures for rotations, such as quaternions.
- 206 • Camera animation.
- 207 • Motion capture.
- 208

209 **Learning Outcomes:**

- 210 1. Compute the location and orientation of model parts using an forward kinematic approach. [Usage]
- 211 2. Compute the orientation of articulated parts of a model from a location and orientation using an inverse
- 212 kinematic approach. [Usage]
- 213 3. Describe the tradeoffs in different representations of rotations. [Assessment]
- 214 4. Implement the spline interpolation method for producing in-between positions and orientations. [Usage]
- 215 5. Implement algorithms for physical modeling of particle dynamics using simple Newtonian mechanics, for
- 216 example Witkin & Kass, snakes and worms, symplectic Euler, Stormer/Verlet, or midpoint Euler methods.
- 217 [Usage]
- 218 6. Describe the tradeoffs in different approaches to ODE integration for particle modeling. [Assessment]
- 219 7. Discuss the basic ideas behind some methods for fluid dynamics for modeling ballistic trajectories, for
- 220 example for splashes, dust, fire, or smoke. [Familiarity]

221 8. Use common animation software to construct simple organic forms using metaball and skeleton. [Usage]
222

223 **GV/Visualization**

224 *[Elective]*

225 Visualization has strong ties to Human Computer Interaction as well as Computational Science.
226 Readers should refer to the HCI and CN KAs for additional topics related to user population and
227 interface evaluations.

228 **Topics:**

- 229 • Visualization of 2D/3D scalar fields: color mapping, isosurfaces.
- 230 • Direct volume data rendering: ray-casting, transfer functions, segmentation.
- 231 • Visualization of:
 - 232 ○ Vector fields and flow data
 - 233 ○ Time-varying data
 - 234 ○ High-dimensional data: dimension reduction, parallel coordinates,
 - 235 ○ Non-spatial data: multi-variate, tree/graph structured, text
- 236 • Perceptual and cognitive foundations that drive visual abstractions.
- 237 • Visualization design.
- 238 • Evaluation of visualization methods.
- 239 • Applications of visualization.
- 240

241 **Learning Outcomes:**

- 242 1. Describe the basic algorithms for scalar and vector visualization. [Familiarity]
- 243 2. Describe the tradeoffs of algorithms in terms of accuracy and performance. [Assessment]
- 244 3. Propose a suitable visualization design for a particular combination of data characteristics and application
245 tasks. [Assessment]
- 246 4. Discuss the effectiveness of a given visualization for a particular task. [Assessment]
- 247 5. Design a process to evaluate the utility of a visualization algorithm or system. [Assessment]
- 248 6. Recognize a variety of applications of visualization including representations of scientific, medical, and
249 mathematical data; flow visualization; and spatial analysis. [Familiarity]
- 250

1 **Human-Computer Interaction (HCI)**

2 Human-computer interaction (HCI) is concerned with designing interactions between human
3 activities and the computational systems that support them, with constructing interfaces to afford
4 those interactions, and with the study of major phenomena surrounding them.

5 Interaction between users and computational artifacts occurs at an interface which includes both
6 software and hardware. Thus interface design impacts the software life-cycle in that it should
7 occur early; the design and implementation of core functionality can influence the user interface
8 – for better or worse.

9 Because it deals with people as well as computational systems, as a knowledge area HCI
10 demands the consideration of cultural, social, organizational, cognitive and perceptual issues.

11 Consequently it draws on a variety of disciplinary traditions, including psychology, ergonomics,
12 computer science, graphic and product design, anthropology and engineering.

13 **HCI: Human Computer Interaction (4 Core-Tier1 hours, 4 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|---------------------|---------------------|-----------------------|
| HCI/Foundations | 4 | | N |
| HCI/Designing Interaction | | 4 | N |
| HCI/Programming Interactive Systems | | | Y |
| HCI/User-Centered Design & Testing | | | Y |
| HCI/Design for Non-Mouse Interfaces | | | Y |
| HCI/Collaboration & Communication | | | Y |
| HCI/Statistical Methods for HCI | | | Y |
| HCI/Human Factors & Security | | | Y |
| HCI/Design-Oriented HCI | | | Y |
| HCI/Mixed, Augmented and Virtual Reality | | | Y |

14

15

16 HCI/Foundations

17 [4 Core-Tier1 hours]

18 **Motivation:** For end-users, the interface *is* the system. So design in this domain must be
19 interaction-focused and human-centered. Students need a different repertoire of techniques to
20 address this than is provided elsewhere in the curriculum.

21 **Topics:**

- 22 • Contexts for HCI (anything with a user interface: webpage, business applications, mobile applications,
23 games, etc.)
- 24 • Processes for user-centered development: early focus on users, empirical testing, iterative design.
- 25 • Different measures for evaluation: utility, efficiency, learnability, user satisfaction.
- 26 • Physical capabilities that inform interaction design: colour perception, ergonomics
- 27 • Cognitive models that inform interaction design: attention, perception and recognition, movement, and
28 memory. Gulfs of expectation and execution.
- 29 • Social models that inform interaction design: culture, communication, networks and organizations.
- 30 • Principles of good design and good designers; engineering tradeoffs
- 31 • Accessibility: interfaces for differently-abled populations (e.g blind, motion-impaired)
- 32 • Interfaces for differently-aged population groups (e.g. children, 80+)
- 33

34 **Learning Outcomes:**

35 Students should be able to:

- 36 1. Discuss why human-centered software development is important [Familiarity]
- 37 2. Summarize the basic precepts of psychological and social interaction [Familiarity]
- 38 3. Develop and use a conceptual vocabulary for analyzing human interaction with software: affordance,
39 conceptual model, feedback, and so forth [Usage]
- 40 4. Define a user-centered design process that explicitly recognizes that the user is not like the developer or her
41 acquaintances [Usage]
- 42 5. Create and conduct a simple usability test for an existing software application [Assessment]
- 43

44 HCI/Designing Interaction

45 [4 Core-Tier2 hours]

46 **Motivation:** CS students need a minimal set of well-established methods and tools to bring to
47 interface construction.

48 **Topics:**

- 49 • Principles of graphical user interfaces (GUIs).
- 50 • Elements of visual design (layout, color, fonts, labelling)
- 51 • Task analysis
- 52 • Paper prototyping
- 53 • Keystroke-level evaluation
- 54 • Help & documentation
- 55 • Handling human/system failure
- 56 • User interface standards
- 57

58

59 **Learning Outcomes:**

60 Students should be able to apply the principles of HCI foundations to:

- 61 1. Create a simple application, together with help & documentation, that supports a graphical user interface [Usage]
- 62 [Usage]
- 63 2. Conduct a quantitative evaluation and discuss/report the results [Usage]
- 64 3. Discuss at least one national or international user interface design standard [Assessment]

65

66 **HCI/Programming Interactive Systems**

67 **[Elective]**

68 **Motivation:** To take a user-experience-centered view of software development and then cover
69 approaches and technologies to make that happen.

70 **Topics:**

- 71 • Software Architecture Patterns: Model-View controller; command objects, online, offline, [*cross reference*
- 72 *SE/Software Design*]
- 73 • Interaction Design Patterns: visual hierarchy, navigational distance
- 74 • Event management and user interaction
- 75 • Geometry management [*cross reference GV/Geometric Modeling*]
- 76 • Choosing interaction styles and interaction techniques
- 77 • Presenting information: navigation, representation, manipulation
- 78 • Interface animation techniques (scene graphs, etc)
- 79 • Widget classes and libraries
- 80 • Modern GUI libraries (e.g. iOS, Android, JavaFX) GUI builders and UI programming environments [*cross*
- 81 *reference to PBD/Mobile Platforms*]
- 82 • Declarative Interface Specification: Stylesheets and DOMs
- 83 • Data-driven applications (database-backed web pages)
- 84 • Cross-platform design
- 85 • Design for resource-constrained devices (e.g. small, mobile devices)

86

87 **Learning Outcomes:**

88 Students should be able to apply the principles of HCI foundations to:

- 89 1. Understand there are common approaches to design problems, and be able to explain the importance of
- 90 Model-View controller to interface programming [Familiarity]
- 91 2. Create an application with a modern graphical user interface [Usage]
- 92 3. Identify commonalities and differences in UIs across different platforms [Usage]
- 93 4. Explain and use GUI programming concepts: event handling, constraint-based layout management, etc
- 94 [Assessment]

95

96 HCI/User-Centered Design and Testing

97 *[Elective]*

98 **Motivation:** An exploration of techniques to ensure that end-users are fully considered at all
99 stages of the design process, from inception to implementation.

100 **Topics:**

- 101 • Approaches to, and characteristics of, the design process
- 102 • Functionality and usability requirements [*cross reference to Software Engineering*]
- 103 • Techniques for gathering requirements: interviews, surveys, ethnographic & contextual enquiry [*cross*
104 *reference to Software Engineering*]
- 105 • Techniques and tools for analysis & presentation of requirements: reports, personas
- 106 • Prototyping techniques and tools: sketching, storyboards, low-fidelity prototyping, wireframes
- 107 • Evaluation without users, using both qualitative and quantitative techniques: walkthroughs, GOMS, expert-
108 based analysis, heuristics, guidelines, and standards
- 109 • Evaluation with users: observation, think-aloud, interview, survey, experiment.
- 110 • Challenges to effective evaluation: sampling, generalization.
- 111 • Reporting the results of evaluations
- 112 • Internationalisation, designing for users from other cultures, cross-cultural evaluation [*cross reference to*
113 *Software Engineering*]

114

115 **Learning Outcomes:**

116 Students should be able to apply the principles of HCI foundations to:

- 117 1. Understand how user-centred design complements other software process models [Familiarity]
- 118 2. Use lo-fi prototyping techniques to gather, and report, user responses [Usage]
- 119 3. Choose appropriate methods to support the development of a specific UI [Assessment]
- 120 4. Use a variety of techniques to evaluate a given UI [Assessment]
- 121 5. Describe the constraints and benefits of different evaluative methods [Assessment]

122

123 HCI/Design for Non-Mouse Interfaces

124 *[Elective]*

125 **Motivation:** As technologies evolve, new interaction styles are made possible. This knowledge
126 unit should be considered extensible, to track emergent technology.

127 **Topics:**

- 128 • Choosing interaction styles and interaction techniques
- 129 • Representing information to users: navigation, representation, manipulation
- 130 • Approaches to design, implementation and evaluation of non-mouse interaction
 - 131 ○ Touch and multi-touch interfaces
 - 132 ○ New Windows (iPhone, Android)
 - 133 ○ Speech recognition and natural language processing [*cross reference to Intelligent Systems*]
 - 134 ○ Wearable and tangible interfaces
 - 135 ○ Persuasive interaction and emotion
 - 136 ○ Ubiquitous and context-aware (UbiComp)

- 137 ○ Bayesian inference (e.g. predictive text, guided pointing)
138 ○ Ambient/peripheral display and interaction
139

140 **Learning Outcomes:**

- 141 Students should be able to apply the principles of HCI foundations to:
142 1. Describe when non-mouse interfaces are appropriate [Familiarity]
143 2. Understand the interaction possibilities beyond mouse-and-pointer interfaces [Usage]
144 3. Discuss the advantages (and disadvantages) of non-mouse interfaces [Assessment]
145

146 **HCI/Collaboration and Communication**

147 **[Elective]**

148 **Motivation:** Computer interfaces not only support users in achieving their individual goals but
149 also in their interaction with others, whether that is task-focussed (work or gaming) or task-
150 unfocussed (social networking).

151 **Topics:**

- 152 • Asynchronous group communication: e-mail, forums, facebook
153 • Synchronous group communication: chat rooms, conferencing, online games
154 • Online communities
155 • Software characters and intelligent agents, virtual worlds and avatars (cross-reference IS/Agents)
156 • Social psychology
157 • Social networking
158 • Social computing
159

160 **Learning Outcomes:**

- 161 Students should be able to apply the principles of HCI foundations to:
162 1. Describe the difference between synchronous and asynchronous communication [Familiarity]
163 2. Compare the HCI issues in individual interaction with group interaction [Usage]
164 3. Discuss several issues of social concern raised by collaborative software [Assessment]
165 4. Discuss the HCI issues in software that embodies human intention [Assessment]
166

167 **HCI/Statistical Methods for HCI**

168 **[Elective]**

169 **Motivation:** Much HCI work depends on the proper use, understanding and application of
170 statistics. This knowledge is often held by students who join the field from psychology, but less
171 common in students with a CS background.

172 **Topics:**

- 173 • t-tests
174 • ANOVA
175 • randomization (non-parametric) testing, within v. between-subjects design

- 176 • calculating effect size
- 177 • exploratory data analysis
- 178 • presenting statistical data
- 179 • using statistical data
- 180 • using qualitative and quantitative results together

181
182 **Learning Outcomes:**

- 183 Students should be able to apply the principles of HCI foundations to:
- 184 1. Explain basic statistical concepts and their areas of application [Familiarity]
 - 185 2. Extract and articulate the statistical arguments used in papers which report [Usage]
- 186

187 **HCI/Human Factors and Security**

188 *[Elective]*

189 **Motivation:** Effective interface design requires basic knowledge of security psychology. Many
 190 attacks do not have a technological basis, but exploit human propensities and vulnerabilities.
 191 “Only amateurs attack machines; professionals target people” (Bruce Schneier)

192 **Topics:**

- 193 • Applied psychology and security policies
- 194 • Security economics
- 195 • Regulatory environments – responsibility, liability and self-determination
- 196 • Organizational vulnerabilities and threats
- 197 • Usability design and security (cross reference to IAS)
- 198 • Pretext, impersonation and fraud. Phishing and spear phishing (cross reference to IAS)
- 199 • Trust, privacy and deception
- 200 • Biometric authentication (camera, voice)
- 201 • Identity management

202
203 **Learning Outcomes:**

- 204 Students should be able to apply the principles of HCI foundations to:
- 205 1. Explain the concepts of phishing and spear phishing, and how to recognize them [Familiarity]
 - 206 2. Explain the concept of identity management and its importance [Familiarity]
 - 207 3. Describe the issues of trust in interface design with an example of a high and low trust system [Usage]
 - 208 4. Design a user interface for a security mechanism [Assessment]
 - 209 5. Analyze a security policy and/or procedures to show where they consider, or fail to consider, human factors
 - 210 [Assessment]
- 211

212 HCI/Design-Oriented HCI

213 *[Elective]*

214 **Motivation:** Some curricula will want to emphasise an understanding of the norms and values of
215 HCI work itself as emerging from, and deployed within specific historical, disciplinary and
216 cultural contexts.

217 **Topics:**

- 218 • Intellectual styles and perspectives to technology and its interfaces
- 219 • Consideration of HCI as a design discipline:
 - 220 ○ Sketching
 - 221 ○ Participatory design
- 222 • Critically reflective HCI
 - 223 ○ Critical technical practice
 - 224 ○ Technologies for political activism
 - 225 ○ Philosophy of user experience
 - 226 ○ Ethnography and ethnomethodology
- 227 • Indicative domains of application
 - 228 ○ Sustainability
 - 229 ○ Arts-informed computing

230

231 **Learning Objectives**

232 Students should be able to apply the principles of HCI foundations to:

- 233 1. Detail the processes of design appropriate to specific design orientations [Familiarity]
- 234 2. Apply a variety of design methods to a given problem [Usage]
- 235 3. Understand HCI as a design-oriented discipline. [Assessment]

236

237 HCI/Mixed, Augmented and Virtual Reality

238 *[Elective]*

239 **Motivation:** A detailed consideration of the interface components required for the creation and
240 development of immersive environments, especially games.

241 **Topics:**

- 242 • Output
 - 243 ○ Sound
 - 244 ○ Stereoscopic display
 - 245 ○ Force feedback simulation, haptic devices
- 246 • User input
 - 247 ○ Viewer and object tracking
 - 248 ○ Pose and gesture recognition
 - 249 ○ Accelerometers
 - 250 ○ Fiducial markers
 - 251 ○ User interface issues
- 252 • Physical modelling and rendering
 - 253 ○ Physical simulation: collision detection & response, animation

- 254 ○ Visibility computation
- 255 ○ Time-critical rendering, multiple levels of details (LOD)
- 256 ● System architectures
- 257 ○ Game engines
- 258 ○ Mobile augmented reality
- 259 ○ Flight simulators
- 260 ○ CAVEs
- 261 ○ Medical imaging
- 262 ● Networking
- 263 ○ p2p, client-server, dead reckoning, encryption, synchronization
- 264 ○ Distributed collaboration
- 265

Learning Objectives:

- 267 1. Describe the optical model realized by a computer graphics system to synthesize stereoscopic view
268 [Familiarity]
- 269 2. Describe the principles of different viewer tracking technologies [Familiarity]
- 270 3. Describe the differences between geometry- and image-based virtual reality [Familiarity]
- 271 4. Describe the issues of user action synchronization and data consistency in a networked environment
272 [Familiarity]
- 273 5. Determine the basic requirements on interface, hardware, and software configurations of a VR system for a
274 specified application [Usage]
- 275 6. To be aware of the range of possibilities for games engines, including their potential and their limitations
276 [Assessment]

1 **Information Assurance and Security (IAS)**

2 In CS2013, the Information Assurance and Security KA is added to the Body of Knowledge in
3 recognition of the world’s reliance on information technology and its critical role in computer
4 science education. Information assurance and security as a domain is the set of controls and
5 processes both technical and policy intended to protect and defend information and information
6 systems by ensuring their availability, integrity, authentication, and confidentiality and providing
7 for non-repudiation. The concept of assurance also carries an attestation that current and past
8 processes and data are valid. Both assurance and security concepts are needed to ensure a
9 complete perspective. Information assurance and security education, then, includes all efforts to
10 prepare a workforce with the needed knowledge, skills, and abilities to protect our information
11 systems and attest to the assurance of the past and current state of processes and data. The
12 Information Assurance and Security KA is unique among the set of KA’s presented here given
13 the manner in which the topics are pervasive throughout other Knowledge Areas. The topics
14 germane to only IAS are presented in depth in the IAS section; other topics are noted and cross
15 referenced in the IAS KA, with the details presented in the KA in which they are tightly
16 integrated.

17 The aim of this KA is two-fold. First, the KA defines the core (tier1and tier2) and the elective
18 components that depict topics that are part of an undergraduate computer science curriculum.
19 Secondly (and almost more importantly), we document the pervasive presence of IAS within a
20 computer science undergraduate curriculum.

21 The IAS KA is shown in two groups; (1) concepts that are, at the first order, germane to
22 Information Assurance and Security and (2) IAS topics that are integrated into other KA’s. For
23 completeness, the total distribution of hours is summarized in the table below.

24

| | Core-Tier1 hours | Core-Tier2 hours | Elective Topics |
|--------------------------------------|-------------------------|-------------------------|------------------------|
| IAS | 2 | 6 | Y |
| IAS distributed in other KA’s | 23 | 46 | Y |

25

26 **IAS. Information Assurance and Security (2 Core-Tier1 hours, 6 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| IAS/Fundamental Concepts | 1 | 2 | N |
| IAS/Network Security | 1 | 4 | N |
| IAS/Cryptography | | | Y |
| IAS/Risk Management | | | Y |
| IAS/Security Policy and Governance | | | Y |
| IAS / Digital Forensics | | | Y |
| IAS / Security Architecture and Systems Administration | | | Y |
| IAS/Secure Software Design and Engineering | | | Y |

27

28

29 **IAS. Information Assurance and Security (distributed) (23 Core-Tier1 hours, 46**
 30 **Core-Tier2 hours)**

| Knowledge Area and Topic | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|-------------------------------------|------------------|------------------|--------------------|
| SDF/Development Methods | 9 | | |
| SE/Software Processes | 1 | | |
| SE/Software Project Management | | 1* | |
| SE/Tools and Environments | | 1* | |
| SE/Software Construction | | 2 | Y |
| SE/Software Verification Validation | | 3 | Y |
| PL/Functional Programming | | 2 | |
| PL/Type Systems | 1 | 4 | |

| | | | |
|--|------------|------------|----------|
| PL/Language Translation And Execution | 1 | 3 | |
| SF/Virtualization and Isolation | | 2* | |
| SF/Reliability through Redundancy | | 2 | |
| PD/Parallelism Fundamentals | 1* | | Y |
| PD/Communication and Coordination | 1 | 3 | |
| OS/ Overview of OS | 1* | | |
| OS/OS Principles | 1* | | |
| OS/Concurrency | | 3 | |
| OS/Scheduling and Dispatch | | 3 | |
| OS/Memory Management | | 1* | |
| OS/Security and Protection | | 2 | |
| OS/Virtual Machines | | | Y |
| OS/Device Management | | | Y |
| OS/File Systems | | | Y |
| OS/Real Time and Embedded Systems | | | Y |
| OS/Fault Tolerance | | | Y |
| OS/System Performance Evaluation | | | Y |
| NC/Introduction | 1.5 | | |
| NC/Networked Applications | 1.5 | | |
| NC/Reliable Data Delivery | | 2 | |
| NC/Routing and Forwarding | | 1.5 | |

| | | | |
|---|----------|------------|----------|
| NC/Local Area Networks | | 1.5 | |
| NC/Resource Allocation | | 1 | |
| NC/Mobility | | 1 | |
| PBD/Web Platforms | | | Y |
| PBD/Mobile Platforms | | | Y |
| PBD/Industrial Platforms | | | Y |
| IM/Information Management Concepts | | 2 | |
| IM/Transaction Processing | | | Y |
| IM/Distributed Databases | | | Y |
| SP/Professional Ethics | 2 | 1 | |
| SP/Intellectual Property | 2 | | |
| SP/Security Policies, Laws and Computer Crimes | | | Y |
| HCI/Human Factors and Security | | | Y |
| IS/Reasoning Under Uncertainty | | | Y |

31 * Indicates not all hours in the KU are classified as cross referenced to IAS

32

33

34 **IAS/Fundamental Concepts**

35 *[1 Core-Tier1 hours, 2 Core-Tier2 hours]*

36 **Topics:**

37 [Core-Tier1]

- 38 • Nature of the Threats (e.g. natural, intentional, accidental)
- 39 • Definition and need for Information Assurance.
- 40 • Basic Information Assurance Concepts that should be recognized. (Confidentiality, Integrity, Availability)

41
42 [Core-Tier2]

- 43 • Industry, Government, and Cultural Guidelines, Standards, and Differences including topics such as
- 44 HIPAA, ISO 27002, Safe Harbor, and data protection laws.
- 45 • Legal, Ethical, and Social Issues (cross-reference SP)
- 46 • Threats and Vulnerabilities.
- 47 • Motivation of Attackers.
- 48 • Protection Mechanisms.
- 49 • Incident Response.

50

51 **Learning outcomes:**

- 52 1. Describe the types of threats to data and information systems [Familiarity]
- 53 2. Describe why processes and data need protection [Familiarity]
- 54 3. Describe the context in which Confidentiality, Integrity, and Availability are important to given processes
- 55 or data. [Familiarity]
- 56 4. Describe the significant national/international level laws affecting the obligation for the protection of data.
- 57 [Familiarity]
- 58 5. Describe the impact of ethics and social issues in information assurance and security. [Familiarity]
- 59 6. Describe the major vulnerabilities present in systems today and the types of attacks. [Familiarity]
- 60 7. Define the fundamental motivations for intentional malicious exploitation of vulnerabilities. [Familiarity]
- 61 8. Define the protection mechanisms that can be used to detect or mitigate malicious activity in information
- 62 systems. [Familiarity]
- 63 9. Define an incident and evaluate the roles and actions taken in response to an incident. [Usage]

64

65 **IAS/Network Security**

66 *[1 Core-Tier1 hours, 4 Core-Tier2 hours]*

67 Discussion of network security relies on previous understanding on fundamental concepts of
68 networking, including protocols, such as TCP/IP, and network architecture/organization (xref
69 NC/Network Communication).

70 **Topics:**

71 [Core-Tier1]

- 72 • Network attack type including denial of service, flooding, sniffing and traffic redirection, and message
- 73 integrity attacks.
- 74 • Use of cryptography for network security.

75

76

77 [Core-Tier2]

- 78 • Protections mechanisms for communication protocols including: (xref NC/Introduction, Networked
79 Applications)
80 • Defense Mechanisms /Countermeasures. (Intrusion Detection, Firewalls, Detection of malware, IPSec,
81 Virtual Private Networks, Network Address Translation.)
82 • Network Auditing.
83

84 **Learning outcomes:**

- 85 1. Identify the common type of network attacks describe how the attack can occur. [Familiarity]
86 2. Describe the architecture for public and private key cryptography and how PKI supports network security.
87 [Usage]
88 3. Describe the appropriate technical controls that can be implemented in the OSI model to support security.
89 [Familiarity]
90 4. Describe the components and their application in the security of networked communications (for example,
91 describe the different impacts of an access control list in a firewall to the use of network access translation.
92 [Familiarity]
93 5. Discuss what information could be found in enterprise systems and network devices to aid in identifying
94 both the presence of a threat and where physically the risk exists [Familiarity]
95

96 **IAS/ Cryptography**

97 **[Elective]**

98 **Topics:**

- 99 • The Basic Cryptography Terminology covers notions pertaining to the different (communication) partners,
100 secure/unsecure channel, attackers and their capabilities, encryption, decryption, keys and their
101 characteristics, signatures, etc.
102 • Cipher types: Caesar cipher, affine cipher, etc. together with typical attack methods such as frequency
103 analysis, etc.
104 • Mathematical Preliminaries; include topics in linear algebra, number theory, probability theory, and
105 statistics. (Discrete Structures)
106 • Cryptographic Primitives include encryption (stream ciphers, block ciphers public key encryption), digital
107 signatures, message authentication codes, and hash functions.
108 • Cryptanalysis covers the state-of-the-art methods including differential cryptanalysis, linear cryptanalysis,
109 factoring, solving discrete logarithm problem, lattice based methods, etc.
110 • Cryptographic Algorithm Design covers principles that govern the design of the various cryptographic
111 primitives, especially block ciphers and hash functions. (Algorithms and Complexity - Hash functions)
112 • The treatment of Common Protocols includes (but should not be limited to) current protocols such as RSA,
113 DES, DSA, AES, ElGamal, MD5, SHA-1, Diffie-Hellman Key exchange, identification and authentication
114 protocols, secret sharing, multi-party computation, etc.
115 • Public Key Infrastructure deals with challenges, opportunities, local infrastructures, and national
116 infrastructure.
117

118 **Learning outcomes:**

- 119 1. Describe the purpose of Cryptography and list ways it is used in data communications. [Familiarity]
120 2. Define the following terms: Cipher, Cryptanalysis, Cryptographic Algorithm, and Cryptology and describe
121 the two basic methods (ciphers) for transforming plain text in cipher text. [Familiarity]
122 3. Discuss the importance of prime numbers in cryptography and explain their use in cryptographic
123 algorithms. [Familiarity]

- 124 4. Discuss the different cryptographic primitives and the work function of each. [Familiarity]
125 5. Describe how the advances in cryptography have made it possible to keep pace with advances in computing
126 power. [Familiarity]
127 6. Discuss the impact of algorithm design and complexity with respect to the work function of a given
128 cryptographic algorithm. [Familiarity]
129 7. Describe the current algorithms used to support various communication security protocols. [Familiarity]
130 8. List the security vulnerabilities of the PKI infrastructure. [Familiarity]
131

132 **IAS/Risk Management**

133 *[Elective]*

134 Risk Analysis involves identifying the assets, probable threats, vulnerabilities and control measures to discern risk
135 levels and likelihoods. It can be applied to a program, organization, sector, etc. Knowledge in this area includes
136 knowing different risk analysis models and methods, their strengths and benefits and the appropriateness of the
137 different methods and models given the situation. This includes periodic reassessment. (cross-reference SE/Software
138 Project Management)

139 *Topics:*

- 140 • Risk acceptance and risk aversion for organizations.
- 141 • Cost/Benefit Analysis used to weigh private and/or public costs versus benefits and can be applied to
142 security policies, investments, programs, tools, deployments, etc.
- 143 • Asset Management minimizes the life cost of assets and includes critical factors such as risk or business
144 continuity.
- 145 • Continuity Planning.
- 146 • Disaster Recovery.
- 147 • Security Auditing.
- 148

149 *Learning outcomes:*

- 150 1. Describe organizational considerations with respect to managing risk and how is risk exposure
151 communicated? [Familiarity]
- 152 2. Describe the methods used to conduct a cost/benefit analysis for risk mitigation. [Familiarity]
- 153 3. Describe an asset and how is the worth established. [Familiarity]
- 154 4. Describe the controls and safeguards an organization may implements to ensure delivery of critical services
155 and ensure survival. [Familiarity]
- 156 5. Critique the trade-off considerations given the value of an asset and the cost of the security controls to
157 mitigate loss/damage/destruction. [Assessment]
- 158 6. Describe the objective of a security audit and how security controls are assessed. [Familiarity]
159

160 **IAS/Security Policy and Governance**

161 *[Elective]*

162 *Topics:*

- 163 • Standards and best practices for organizational security policies.
- 164 • Strategies for creating security policies.
- 165 • Compliance and Enforcement of policies, standards, regulations, and laws.
- 166 • Formal models such as Bell-LaPadula, Biba and Clark-Wilson.
- 167 • Policy as related to Risk Aversion.
- 168

169 **Learning outcomes:**

- 170 1. Describe the function of a security policy in an organization. [Familiarity]
- 171 2. Describe the organizational considerations and challenges when creating and implementing a security
- 172 policy. [Familiarity]
- 173 3. Describe the role of an organization in relation to legal and regulatory compliance in the enforcement of a
- 174 security policy and governance plan. [Familiarity]
- 175 4. Critique formal models such as Bell-LaPadula, Biba and Clark-Wilson and the role in security policy and
- 176 governance. [Familiarity]
- 177 5. Describe the impact of risk aversion on the development and implementation of an organization's security
- 178 policy and governance plan. [Familiarity]
- 179

180 **IAS/ Digital Forensics**

181 **[Elective]**

182 **Topics:**

- 183 • Basic Principles and methodologies for digital forensics.
- 184 • Rules of Evidence – general concepts and differences between jurisdictions and Chain of Custody.
- 185 • Search and Seizure of evidence, e.g., computers, including search warrant issues.
- 186 • Digital Evidence methods and standards.
- 187 • Techniques and standards for Preservation of Data.
- 188 • Legal and Reporting Issues including working as an expert witness.
- 189 • OS/File System Forensics
- 190 • Application Forensics
- 191 • Network Forensics
- 192 • Mobile Device Forensics
- 193 • Computer/network/system attacks.
- 194

195 **Learning outcomes:**

- 196 1. Describe what is a Digital Investigation is, the sources of digital evidence, and the responsibilities for the
- 197 involved parties. [Familiarity]
- 198 2. Describe the legal requirements for use if seized data. [Familiarity]
- 199 3. Describe the process of evidence seizure from the time when the requirement was identified to the
- 200 disposition of the data. [Familiarity]
- 201 4. Describe how data collection is accomplished and the proper storage of the original and forensics copy.
- 202 [Familiarity]
- 203 5. Conduct a data collection on a harddrive. [Usage]
- 204 6. Describe a person's responsibility and liability while testifying as a forensics examiner. [Familiarity]
- 205 7. Describe the file system structure for a given device (NTFA, MFS, iNode, HFS...) and recover data based
- 206 on a given search term from an imaged system. [Usage]
- 207 8. Describe how an application can be evaluated to determine if it is the intended application (pre-install),
- 208 evaluate the application at run-time, and review any error/status logs for unexpected activity. [Usage]
- 209 9. Capture and interpret network traffic. [Usage]
- 210 10. Discuss the challenges associated with mobile device forensics. [Familiarity]
- 211 11. Evaluate a system (network, computer, or application) for the presence of malware or malicious activity.
- 212 [Assessment]
- 213

214

215 IAS/Security Architecture and Systems Administration

216 *[Elective]*

217 *Topics:*

- 218 • Considerations for architecting a secure computing system; for example the Saltzer and Schroeder security
219 principles.
- 220 • Access Control Basic Principles.
- 221 • Physical and information system access controls.
- 222 • Usability of systems: including the difficulty for humans to deal with security (e.g., remembering PINs),
223 social engineering, phishing, and other similar attacks.
- 224 • Analyzing and identifying Threats and Vulnerabilities
- 225 • Multi-level/Multi-lateral Security
- 226 • Supervisory Control and Data Acquisition (SCADA)
- 227

228 *Learning outcomes:*

- 229 1. Describe the security principles that should be considered to secure a computing system. [Familiarity]
- 230 2. Describe the function of an access control and it's integration into an enterprise. [Familiarity]
- 231 3. Describe the considerations for usability and social acceptance of security controls. [Familiarity]
- 232 4. Identify where data exists in a networked environment, what tools can be used to review the data, and how
233 to analyze the data for evidence of a risk. [Familiarity]
- 234 5. Define "Defense in Depth" and how security controls can compliment or interfere with each other.
235 [Familiarity]
- 236 6. Describe the nature of SCADA systems and the security considerations in designing and protecting them.
237 [Familiarity]
- 238

239 IAS/Secure Software Design and Engineering

240 *[Elective]*

241 Fundamentals of secure coding practices covered in other knowledge areas, including
242 SDF/SE/PL/SF.

243 *Topics:*

- 244 • Building security into the Software Development Lifecycle (cross-reference SE/ Software Processes)
- 245 • Secure Design Principles and Patterns (Saltzer and Schroeder, etc)
- 246 • Secure Software Specification and Requirements deals with specifying what the program should and should
247 not do, which can be done either using a requirements document or using formal methods.
- 248 • Secure Coding techniques to minimize vulnerabilities in code, such as data validation, memory handling,
249 and crypto implementation (cross-reference SE/Software Construction)
- 250 • Secure Testing is the process of testing that security requirements are met (including Static and Dynamic
251 analysis).
- 252

253 *Learning outcomes:*

- 254 1. Describe the requirements for integrating security into the SDL. [Familiarity]
- 255 2. Apply the concepts of the Design Principles for Protection Mechanisms (e.g. Saltzer and Schroeder), the
256 Principles for Software Security (Viega and McGraw), and the Principles for Secure Design (Morrie
257 Gasser) on a software development project [Usage]

- 258
259
260
261
262
3. Develop specifications for a software development effort that fully specify functional requirements and identifies the expected execution paths. [Usage]
 4. Describe software development best practices for minimizing vulnerabilities in programming code. [Familiarity]
 5. Conduct a security verification and assessment (static and dynamic) of a software application [Usage]

1 **Information Management (IM)**

2 Information Management (IM) is primarily concerned with the capture, digitization,
3 representation, organization, transformation, and presentation of information; algorithms for
4 efficient and effective access and updating of stored information, data modeling and abstraction,
5 and physical file storage techniques. The student needs to be able to develop conceptual and
6 physical data models, determine what IM methods and techniques are appropriate for a given
7 problem, and be able to select and implement an appropriate IM solution that addresses relevant
8 design concerns including scalability, accessibility and usability.

9 We also note that IM is related to fundamental information security concepts that are described
10 in the Information Assurance and Security (IAS) topic area, *IAS/Fundamental Concepts*.

11 **IM. Information Management (1 Core-Tier1 hour; 9 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|---|------------------|------------------|--------------------|
| IM/Information Management Concepts | 1 | 2 | N |
| IM/Database Systems | | 3 | Y |
| IM/Data Modeling | | 4 | N |
| IM/Indexing | | | Y |
| IM/Relational Databases | | | Y |
| IM/Query Languages | | | Y |
| IM/Transaction Processing | | | Y |
| IM/Distributed Databases | | | Y |
| IM/Physical Database Design | | | Y |
| IM/Data Mining | | | Y |
| IM/Information Storage And Retrieval | | | Y |
| IM/MultiMedia Systems | | | Y |

12

13 **IM. Information Management-related topics (distributed) (1 Core-Tier1 hour, 2**
14 **Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|----------------------------------|------------------|------------------|--------------------|
| IAS/Fundamental Concepts* | 1 | 2 | N |

15 * See Information Assurance and Security Knowledge Area for a description of this topic area.

16

17 **IM/Information Management Concepts**

18 *[1 Core-Tier1 hour; 2 Core-Tier2 hours]*

19 **Topics:**

20 [Core-Tier1]

- 21 • Information systems as socio-technical systems
- 22 • Basic information storage and retrieval (IS&R) concepts
- 23 • Information capture and representation
- 24 • Supporting human needs: Searching, retrieving, linking, browsing, navigating

26 [Core-Tier2]

- 27 • Information management applications
- 28 • Declarative and navigational queries, use of links
- 29 • Analysis and indexing
- 30 • Quality issues: Reliability, scalability, efficiency, and effectiveness

31
32 **Learning Outcomes:**

33 [Core-Tier1]

- 34 1. Describe how humans gain access to information and data to support their needs [Familiarity]
- 35 2. Understand advantages and disadvantages of central organizational control over data [Assessment]
- 36 3. Identify the careers/roles associated with information management (e.g., database administrator, data
37 modeler, application developer, end-user) [Familiarity]
- 38 4. Compare and contrast information with data and knowledge [Assessment]
- 39 5. Demonstrate uses of explicitly stored metadata/schema associated with data [Usage]
- 40 6. Identify issues of data persistence for an organization [Familiarity]

41
42 [Core-Tier2]

- 43 7. Critique/defend a small- to medium-size information application with regard to its satisfying real user
44 information needs [Assessment]
- 45 8. Explain uses of declarative queries [Familiarity]
- 46 9. Give a declarative version for a navigational query [Familiarity]
- 47 10. Describe several technical solutions to the problems related to information privacy, integrity, security, and
48 preservation [Familiarity]
- 49 11. Explain measures of efficiency (throughput, response time) and effectiveness (recall, precision)
50 [Familiarity]
- 51 12. Describe approaches that scale up to globally networked systems [Familiarity]
- 52 13. Identify vulnerabilities and failure scenarios in common forms of information systems [Usage]

53

54 **IM/Database Systems**

55 *[3 Core-Tier2 hours]*

56 **Topics:**

57 [Core-Tier2]

- 58 • Approaches to and evolution of database systems

- 59 • Components of database systems
- 60 • Design of core DBMS functions (e.g., query mechanisms, transaction management, buffer management,
- 61 access methods)
- 62 • Database architecture and data independence
- 63 • Use of a declarative query language
- 64 • Systems supporting structured and/or stream content

65
66 [Elective]

- 67 • Approaches for managing large volumes of data (e.g., noSQL database systems, use of MapReduce).

68
69 **Learning Outcomes:**

70 [Core-Tier2]

- 71 1. Explain the characteristics that distinguish the database approach from the traditional approach of
- 72 programming with data files [Familiarity]
- 73 2. Understand the most common designs for core database system components including the query optimizer,
- 74 query executor, storage manager, access methods, and transaction processor. [Familiarity]
- 75 3. Cite the basic goals, functions, models, components, applications, and social impact of database systems
- 76 [Familiarity]
- 77 4. Describe the components of a database system and give examples of their use [Familiarity]
- 78 5. Identify major DBMS functions and describe their role in a database system [Familiarity]
- 79 6. Explain the concept of data independence and its importance in a database system [Familiarity]
- 80 7. Use a declarative query language to elicit information from a database [Usage]
- 81 8. Describe how various types of content cover the notions of structure and/or of stream (sequence), e.g.,
- 82 documents, multimedia, tables [Familiarity]

83
84 [Elective]

- 85 9. Describe major approaches to storing and processing large volumes of data [Familiarity]

86

87 **IM/Data Modeling**

88 **[4 Core-Tier2 hours]**

89 **Topics:**

90 [Core-Tier2]

- 91 • Data modeling
- 92 • Conceptual models (e.g., entity-relationship, UML diagrams)
- 93 • Spreadsheet models
- 94 • Relational data models
- 95 • Object-oriented models
- 96 • Semi-structured data model (expressed using DTD or XML Schema, for example)

97

98

99 **Learning Outcomes:**

100 [Core-Tier2]

- 101 1. Categorize data models based on the types of concepts that they provide to describe the database structure
102 and their usage, for example, use of conceptual, spreadsheet, physical, and representational data models
103 [Assessment]
- 104 2. Describe the modeling concepts and notation of widely used modeling notation (e.g., ERD notation, and
105 UML), including their use in data modeling [Familiarity]
- 106 3. Define the fundamental terminology used in the relational data model [Familiarity]
- 107 4. Describe the basic principles of the relational data model [Familiarity]
- 108 5. Apply the modeling concepts and notation of the relational data model [Usage]
- 109 6. Describe the main concepts of the OO model such as object identity, type constructors, encapsulation,
110 inheritance, polymorphism, and versioning [Familiarity]
- 111 7. Describe the differences between relational and semi-structured data models [Assessment]
- 112 8. Give a semi-structured equivalent (e.g., in DTD or XML Schema) for a given relational schema [Usage]
- 113

114 **IM/Indexing**

115 [*Elective*]

116 *Topics:*

- 117 • The impact of indices on query performance
- 118 • The basic structure of an index
- 119 • Keeping a buffer of data in memory
- 120 • Creating indexes with SQL
- 121 • Indexing text
- 122 • Indexing the web (how search engines work)
- 123

124 **Learning Outcomes:**

- 125 1. Generate an index file for a collection of resources [Usage]
- 126 2. Explain the role of an inverted index in locating a document in a collection [Familiarity]
- 127 3. Explain how stemming and stop words affect indexing [Familiarity]
- 128 4. Identify appropriate indices for given relational schema and query set [Usage]
- 129 5. Estimate time to retrieve information, when indices are used compared to when they are not used [Usage]
- 130

131 **IM/Relational Databases**

132 [*Elective*]

133 *Topics:*

134 Elective

- 135 • Mapping conceptual schema to a relational schema
- 136 • Entity and referential integrity
- 137 • Relational algebra and relational calculus
- 138 • Relational Database design
- 139 • Functional dependency

- 140 • Decomposition of a schema; lossless-join and dependency-preservation properties of a decomposition
- 141 • Candidate keys, superkeys, and closure of a set of attributes
- 142 • Normal forms (BCNF)
- 143 • Multi-valued dependency (4NF)
- 144 • Join dependency (PJNF, 5NF)
- 145 • Representation theory
- 146

147 **Learning Outcomes:**

- 148 1. Prepare a relational schema from a conceptual model developed using the entity- relationship model
- 149 [Usage]
- 150 2. Explain and demonstrate the concepts of entity integrity constraint and referential integrity constraint
- 151 (including definition of the concept of a foreign key) [Usage]
- 152 3. Demonstrate use of the relational algebra operations from mathematical set theory (union, intersection,
- 153 difference, and Cartesian product) and the relational algebra operations developed specifically for relational
- 154 databases (select (restrict), project, join, and division) [Usage]
- 155 4. Demonstrate queries in the relational algebra [Usage]
- 156 5. Demonstrate queries in the tuple relational calculus [Usage]
- 157 6. Determine the functional dependency between two or more attributes that are a subset of a relation
- 158 [Assessment]
- 159 7. Connect constraints expressed as primary key and foreign key, with functional dependencies [Usage]
- 160 8. Compute the closure of a set of attributes under given functional dependencies [Usage]
- 161 9. Determine whether or not a set of attributes form a superkey and/or candidate key for a relation with given
- 162 functional dependencies [Assessment]
- 163 10. Evaluate a proposed decomposition, to say whether or not it has lossless-join and dependency-preservation
- 164 [Assessment]
- 165 11. Describe what is meant by BCNF, PJNF, 5NF [Familiarity]
- 166 12. Explain the impact of normalization on the efficiency of database operations especially query optimization
- 167 [Familiarity]
- 168 13. Describe what is a multi-valued dependency and what type of constraints it specifies [Familiarity]
- 169

170 **IM/Query Languages**

171 *[Elective]*

172 **Topics:**

- 173 • Overview of database languages
- 174 • SQL (data definition, query formulation, update sublanguage, constraints, integrity)
- 175 • Selections
- 176 • Projections
- 177 • Select-project-join
- 178 • Aggregates and group-by
- 179 • Subqueries
- 180 • QBE and 4th-generation environments
- 181 • Different ways to invoke non-procedural queries in conventional languages
- 182 • Introduction to other major query languages (e.g., XPATH, SPARQL)
- 183 • Stored procedures
- 184
- 185

186 **Learning Outcomes:**

- 187 1. Create a relational database schema in SQL that incorporates key, entity integrity, and referential integrity
188 constraints [Usage]
- 189 2. Demonstrate data definition in SQL and retrieving information from a database using the SQL SELECT
190 statement [Usage]
- 191 3. Evaluate a set of query processing strategies and select the optimal strategy [Assessment]
- 192 4. Create a non-procedural query by filling in templates of relations to construct an example of the desired
193 query result [Usage]
- 194 5. Embed object-oriented queries into a stand-alone language such as C++ or Java (e.g., SELECT
195 Col.Method() FROM Object) [Usage]
- 196 6. Write a stored procedure that deals with parameters and has some control flow, to provide a given
197 functionality [Usage]
- 198

199 **IM/Transaction Processing**

200 **[Elective]**

201 **Topics:**

- 202 • Transactions
- 203 • Failure and recovery
- 204 • Concurrency control
- 205 • Interaction of transaction management with storage, especially buffering
- 206

207 **Learning Outcomes:**

- 208 1. Create a transaction by embedding SQL into an application program [Usage]
- 209 2. Explain the concept of implicit commits [Familiarity]
- 210 3. Describe the issues specific to efficient transaction execution [Familiarity]
- 211 4. Explain when and why rollback is needed and how logging assures proper rollback [Assessment]
- 212 5. Explain the effect of different isolation levels on the concurrency control mechanisms [Assessment]
- 213 6. Choose the proper isolation level for implementing a specified transaction protocol [Assessment]
- 214 7. Identify appropriate transaction boundaries in application programs [Assessment]
- 215

216 **IM/Distributed Databases**

217 **[Elective]**

218 **Topics:**

- 219 • Distributed DBMS
 - 220 ○ Distributed data storage
 - 221 ○ Distributed query processing
 - 222 ○ Distributed transaction model
 - 223 ○ Homogeneous and heterogeneous solutions
 - 224 ○ Client-server distributed databases (cross-reference SF/Computational Paradigms)
- 225 • Parallel DBMS
 - 226 ○ Parallel DBMS architectures: shared memory, shared disk, shared nothing;
 - 227 ○ Speedup and scale-up, e.g., use of the MapReduce processing model (cross-reference
 - 228 CN/Processing, PD/Parallel Decomposition)
 - 229 ○ Data replication and weak consistency models

230
231 **Learning Outcomes:**

- 232 1. Explain the techniques used for data fragmentation, replication, and allocation during the distributed
233 database design process [Familiarity]
234 2. Evaluate simple strategies for executing a distributed query to select the strategy that minimizes the amount
235 of data transfer [Assessment]
236 3. Explain how the two-phase commit protocol is used to deal with committing a transaction that accesses
237 databases stored on multiple nodes [Familiarity]
238 4. Describe distributed concurrency control based on the distinguished copy techniques and the voting method
239 [Familiarity]
240 5. Describe the three levels of software in the client-server model [Familiarity]
241

242 **IM/Physical Database Design**

243 **[Elective]**

244 **Topics:**

- 245 • Storage and file structure
246 • Indexed files
247 • Hashed files
248 • Signature files
249 • B-trees
250 • Files with dense index
251 • Files with variable length records
252 • Database efficiency and tuning
253

254 **Learning Outcomes:**

- 255 1. Explain the concepts of records, record types, and files, as well as the different techniques for placing file
256 records on disk [Familiarity]
257 2. Give examples of the application of primary, secondary, and clustering indexes [Familiarity]
258 3. Distinguish between a non-dense index and a dense index [Assessment]
259 4. Implement dynamic multilevel indexes using B-trees [Usage]
260 5. Explain the theory and application of internal and external hashing techniques [Familiarity]
261 6. Use hashing to facilitate dynamic file expansion [Usage]
262 7. Describe the relationships among hashing, compression, and efficient database searches [Familiarity]
263 8. Evaluate costs and benefits of various hashing schemes [Assessment]
264 9. Explain how physical database design affects database transaction efficiency [Familiarity]
265

266 **IM/Data Mining**

267 **[Elective]**

268 **Topics:**

- 269 • The usefulness of data mining
270 • Data mining algorithms
271 • Associative and sequential patterns
272 • Data clustering
273 • Market basket analysis

- 274 • Data cleaning
275 • Data visualization
276

277 **Learning Outcomes:**

- 278 1. Compare and contrast different conceptions of data mining as evidenced in both research and application
279 [Assessment]
280 2. Explain the role of finding associations in commercial market basket data [Familiarity]
281 3. Characterize the kinds of patterns that can be discovered by association rule mining [Assessment]
282 4. Describe how to extend a relational system to find patterns using association rules [Familiarity]
283 5. Evaluate methodological issues underlying the effective application of data mining [Assessment]
284 6. Identify and characterize sources of noise, redundancy, and outliers in presented data [Assessment]
285 7. Identify mechanisms (on-line aggregation, anytime behavior, interactive visualization) to close the loop in
286 the data mining process [Familiarity]
287 8. Describe why the various close-the-loop processes improve the effectiveness of data mining [Familiarity]
288

289 **IM/Information Storage and Retrieval**

290 **[Elective]**

291 **Topics:**

- 292 • Characters, strings, coding, text
293 • Documents, electronic publishing, markup, and markup languages
294 • Tries, inverted files, PAT trees, signature files, indexing
295 • Morphological analysis, stemming, phrases, stop lists
296 • Term frequency distributions, uncertainty, fuzziness, weighting
297 • Vector space, probabilistic, logical, and advanced models
298 • Information needs, relevance, evaluation, effectiveness
299 • Thesauri, ontologies, classification and categorization, metadata
300 • Bibliographic information, bibliometrics, citations
301 • Routing and (community) filtering
302 • Search and search strategy, multimedia search, information seeking behavior, user modeling, feedback
303 • Information summarization and visualization
304 • Integration of citation, keyword, classification scheme, and other terms
305 • Protocols and systems (including Z39.50, OPACs, WWW engines, research systems)
306 • Digital libraries
307 • Digitization, storage, interchange, digital objects, composites, and packages
308 • Metadata, cataloging, author submission
309 • Naming, repositories, archives
310 • Spaces (conceptual, geographical, 2/3D, VR)
311 • Architectures (agents, buses, wrappers/mediators), interoperability
312 • Services (searching, linking, browsing, and so forth)
313 • Intellectual property rights management, privacy, and protection (watermarking)
314 • Archiving and preservation, integrity
315

316

317 **Learning Outcomes:**

- 318 1. Explain basic information storage and retrieval concepts [Familiarity]
- 319 2. Describe what issues are specific to efficient information retrieval [Familiarity]
- 320 3. Give applications of alternative search strategies and explain why the particular search strategy is
321 appropriate for the application [Assessment]
- 322 4. Perform Internet-based research [Usage]
- 323 5. Design and implement a small to medium size information storage and retrieval system, or digital library
324 [Usage]
- 325 6. Describe some of the technical solutions to the problems related to archiving and preserving information in
326 a digital library [Familiarity]

327

328 **IM/Multimedia Systems**

329 *[Elective]*

330 **Topics:**

- 331 • Input and output devices (scanners, digital camera, touch-screens, voice-activated, MIDI keyboards,
332 synthesizers), device drivers, control signals and protocols, DSPs
- 333 • Standards (audio, music, graphics, image, telephony, video, TV), including storage standards (Magnet
334 Optical disk, CD-ROM, DVD)
- 335 • Applications, media editors, authoring systems, and authoring
- 336 • Streams/structures, capture/represent/transform, spaces/domains, compression/coding
- 337 • Content-based analysis, indexing, and retrieval of audio, images, animation, and video
- 338 • Presentation, rendering, synchronization, multi-modal integration/interfaces
- 339 • Real-time delivery, quality of service (including performance), capacity planning, audio/video
340 conferencing, video-on-demand
- 341

342 **Learning Outcomes:**

- 343 1. Describe the media and supporting devices commonly associated with multimedia information and systems
344 [Familiarity]
- 345 2. Explain basic multimedia presentation concepts [Familiarity]
- 346 3. Demonstrate the use of content-based information analysis in a multimedia information system [Usage]
- 347 4. Critique multimedia presentations in terms of their appropriate use of audio, video, graphics, color, and
348 other information presentation concepts [Assessment]
- 349 5. Implement a multimedia application using a commercial authoring system [Usage]
- 350 6. For each of several media or multimedia standards, describe in non-technical language what the standard
351 calls for, and explain how aspects of human perception might be sensitive to the limitations of that standard
352 [Familiarity]
- 353 7. Describe the characteristics of a computer system (including identification of support tools and appropriate
354 standards) that has to host the implementation of one of a range of possible multimedia applications
355 [Familiarity]

1 **Intelligent Systems (IS)**

2 Artificial intelligence (AI) is the study of solutions for problems that are difficult or impractical
3 to solve with traditional methods. It is used pervasively in support of everyday applications such
4 as email, word-processing and search, as well as in the design and analysis of autonomous agents
5 that perceive their environment and interact rationally with the environment.

6 The solutions rely on a broad set of general and specialized knowledge representation
7 schemes, problem solving mechanisms and learning techniques. They deal with sensing (e.g.,
8 speech recognition, natural language understanding, computer vision), problem-solving (e.g.,
9 search, planning), and acting (e.g., robotics) and the architectures needed to support them (e.g.,
10 agents, multi-agents). The study of Artificial Intelligence prepares the student to determine when
11 an AI approach is appropriate for a given problem, identify the appropriate representation and
12 reasoning mechanism, and implement and evaluate it.

13 **IS. Intelligent Systems (10 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|---|---------------------|---------------------|-----------------------|
| IS/Fundamental Issues | | 1 | Y |
| IS/Basic Search Strategies | | 4 | N |
| IS/Basic Knowledge Representation and Reasoning | | 3 | N |
| IS/Basic Machine Learning | | 2 | N |
| IS/Advanced Search | | | Y |
| IS/Advanced Representation and Reasoning | | | Y |
| IS/Reasoning Under Uncertainty | | | Y |
| IS/Agents | | | Y |
| IS/Natural Language Processing | | | Y |
| IS/Advanced Machine Learning | | | Y |
| IS/Robotics | | | Y |
| IS/Perception and Computer Vision | | | Y |

14

15

16 **IS/Fundamental Issues**

17 *[1 Core-Tier2 hours]*

18 **Topics:**

- 19 • Overview of AI problems, Examples of successful recent AI applications
- 20 • What is intelligent behavior?
 - 21 ○ The Turing test
 - 22 ○ Rational versus non-rational reasoning
- 23 • Nature of environments
 - 24 ○ Fully versus partially observable
 - 25 ○ Single versus multi-agent
 - 26 ○ Deterministic versus stochastic
 - 27 ○ Static versus dynamic
 - 28 ○ Discrete versus continuous
- 29 • Nature of agents
 - 30 ○ Autonomous versus semi-autonomous
 - 31 ○ Reflexive, goal-based, and utility-based
 - 32 ○ The importance of perception and environmental interactions
- 33 • Philosophical and ethical issues [elective]
- 34

35 **Learning Outcomes:**

- 36 1. Describe Turing test and the “Chinese Room” thought experiment. [Familiarity]
- 37 2. Differentiate between the concepts of optimal reasoning/behavior and human-like reasoning/behavior.
38 [Familiarity]
- 39 3. Describe a given problem domain using the characteristics of the environments in which intelligent systems
40 must function. [Assessment]
- 41

42 **IS/Basic Search Strategies**

43 *[4 Core-Tier2 hours]*

44 (Cross-reference AL/Basic Analysis, AL/Algorithmic Strategies, AL/Fundamental Data
45 Structures and Algorithms)

46 **Topics:**

- 47 • Problem spaces (states, goals and operators), problem solving by search
- 48 • Factored representation (factoring state into variables)
- 49 • Uninformed search (breadth-first, depth-first, depth-first with iterative deepening)
- 50 • Heuristics and informed search (hill-climbing, generic best-first, A*)
- 51 • Space and time efficiency of search
- 52 • Two-player games (Introduction to minimax search)
- 53 • Constraint satisfaction (backtracking and local search methods)
- 54

55

56 **Learning Outcomes:**

- 57 1. Formulate an efficient problem space for a problem expressed in natural language (e.g., English) in terms
58 of initial and goal states, and operators. [Usage]
59 2. Describe the role of heuristics and describe the trade-offs among completeness, optimality, time
60 complexity, and space complexity. [Familiarity]
61 3. Describe the problem of combinatorial explosion of search space and its consequences. [Familiarity]
62 4. Select and implement an appropriate uninformed search algorithm for a problem, and characterize its time
63 and space complexities. [Assessment, Usage]
64 5. Select and implement an appropriate informed search algorithm for a problem by designing the necessary
65 heuristic evaluation function. [Assessment, Usage]
66 6. Evaluate whether a heuristic for a given problem is admissible/can guarantee optimal solution.
67 [Assessment]
68 7. Formulate a problem specified in natural language (e.g., English) as a constraint satisfaction problem and
69 implement it using a chronological backtracking algorithm or stochastic local search. [Usage]
70 8. Compare and contrast basic search issues with game playing issues [Familiarity]
71

72 **IS/Basic Knowledge Representation and Reasoning**

73 *[3 Core-Tier2 hours]*

74 **Topics:**

- 75 • Review of propositional and predicate logic (cross-reference DS/Basic Logic)
76 • Resolution and theorem proving (propositional logic only)
77 • Forward chaining, backward chaining
78 • Review of probabilistic reasoning, Bayes theorem (cross-reference with DS/Discrete Probability)
79

80 **Learning Outcomes:**

- 81 1. Translate a natural language (e.g., English) sentence into predicate logic statement. [Usage]
82 2. Convert a logic statement into clause form. [Usage]
83 3. Apply resolution to a set of logic statements to answer a query. [Usage]
84 4. Apply Bayes theorem to determine conditional probabilities in a problem. [Usage]
85

86 **IS/Basic Machine Learning**

87 *[2 Core-Tier2 hours]*

88 **Topics:**

- 89 • Definition and examples of broad variety of machine learning tasks, including classification
90 • Inductive learning
91 • Simple statistical-based learning such as Naive Bayesian Classifier, Decision trees
92 • Define over-fitting problem
93 • Measuring classifier accuracy
94

95 **Learning Outcomes:**

- 96 1. List the differences among the three main styles of learning: supervised, reinforcement, and unsupervised.
97 [Familiarity]
98 2. Identify examples of classification tasks, including the available input features and output to be predicted.
99 [Familiarity]
100 3. Explain the difference between inductive and deductive learning. [Familiarity]

- 101 4. Apply the simple statistical learning algorithm such as Naive Bayesian Classifier to a classification task and
102 measure the classifier's accuracy. [Usage]
103

104 **IS/Advanced Search**

105 *[Elective]*

106 *Topics:*

- 107 • Constructing search trees, dynamic search space, combinatorial explosion of search space
- 108 • Stochastic search
 - 109 ○ Simulated annealing
 - 110 ○ Genetic algorithms
 - 111 ○ Monte-Carlo tree search
- 112 • Implementation of A* search, Beam search
- 113 • Minimax Search, Alpha-beta pruning
- 114 • Expectimax search (MDP-solving) and chance nodes
- 115

116 *Learning Outcomes:*

- 117 1. Design and implement a genetic algorithm solution to a problem. [Usage]
- 118 2. Design and implement a simulated annealing schedule to avoid local minima in a problem. [Usage]
- 119 3. Design and implement A*/beam search to solve a problem. [Usage]
- 120 4. Apply minimax search with alpha-beta pruning to prune search space in a two-player game. [Usage]
- 121 5. Compare and contrast genetic algorithms with classic search techniques. [Assessment]
- 122 6. Compare and contrast various heuristic searches vis-a-vis applicability to a given problem. [Assessment]
- 123

124 **IS/Advanced Representation and Reasoning**

125 *[Elective]*

126 *Topics:*

- 127 • Knowledge representation issues
 - 128 ○ Description logics
 - 129 ○ Ontology engineering
- 130 • Non-monotonic reasoning (e.g., non-classical logics, default reasoning, etc.)
- 131 • Argumentation
- 132 • Reasoning about action and change (e.g., situation and event calculus)
- 133 • Temporal and spatial reasoning
- 134 • Rule-based Expert Systems
- 135 • Model-based and Case-based reasoning
- 136 • Planning:
 - 137 ○ Partial and totally ordered planning
 - 138 ○ Plan graphs
 - 139 ○ Hierarchical planning
 - 140 ○ Planning and execution including conditional planning and continuous planning
 - 141 ○ Mobile agent/Multi-agent planning
- 142

143

144 **Learning Outcomes:**

- 145 1. Compare and contrast the most common models used for structured knowledge representation, highlighting
146 their strengths and weaknesses. [Assessment]
- 147 2. Identify the components of non-monotonic reasoning and its usefulness as a representational mechanisms
148 for belief systems. [Familiarity]
- 149 3. Compare and contrast the basic techniques for representing uncertainty. [Familiarity, Assessment]
- 150 4. Compare and contrast the basic techniques for qualitative representation. [Familiarity, Assessment]
- 151 5. Apply situation and event calculus to problems of action and change. [Usage]
- 152 6. Explain the distinction between temporal and spatial reasoning, and how they interrelate. [Familiarity,
153 Assessment]
- 154 7. Explain the difference between rule-based, case-based and model-based reasoning techniques. [Familiarity,
155 Assessment]
- 156 8. Define the concept of a planning system and how they differ from classical search techniques. [Familiarity,
157 Assessment]
- 158 9. Describe the differences between planning as search, operator-based planning, and propositional planning,
159 providing examples of domains where each is most applicable. [Familiarity, Assessment]
- 160 10. Explain the distinction between monotonic and non-monotonic inference. [Familiarity]
- 161

162 **IS/Reasoning Under Uncertainty**

163 **[Elective]**

164 **Topics:**

- 165 • Review of basic probability (cross-reference DS/Discrete Probability)
- 166 • Random variables and probability distributions
 - 167 ○ Axioms of probability
 - 168 ○ Probabilistic inference
 - 169 ○ Bayes' Rule
- 170 • Conditional Independence
- 171 • Knowledge representations
 - 172 ○ Bayesian Networks
 - 173 ▪ Exact inference and its complexity
 - 174 ▪ Randomized sampling (Monte Carlo) methods (e.g. Gibbs sampling)
 - 175 ○ Markov Networks
 - 176 ○ Relational probability models
 - 177 ○ Hidden Markov Models
- 178 • Decision Theory
 - 179 ○ Preferences and utility functions
 - 180 ○ Maximizing expected utility
- 181

182 **Learning Outcomes:**

- 183 1. Apply Bayes' rule to determine the probability of a hypothesis given evidence. [Usage]
- 184 2. Explain how conditional independence assertions allow for greater efficiency of probabilistic systems.
185 [Assessment]
- 186 3. Identify examples of knowledge representations for reasoning under uncertainty. [Familiarity]
- 187 4. State the complexity of exact inference. Identify methods for approximate inference. [Familiarity]
- 188 5. Design and implement at least one knowledge representation for reasoning under uncertainty. [Usage]
- 189 6. Describe the complexities of temporal probabilistic reasoning. [Familiarity]
- 190 7. Explain the complexities of temporal probabilistic reasoning. [Assessment]
- 191 8. Design and implement an HMM as one example of a temporal probabilistic system. [Usage]

- 192 9. Describe the relationship between preferences and utility functions. [Familiarity]
193 10. Explain how utility functions and probabilistic reasoning can be combined to make rational decisions.
194 [Assessment]
195

196 **IS/Agents**

197 *[Elective]*

198 (Cross-reference HCI/Collaboration and Communication)

199 *Topics:*

- 200 • Definitions of agents
 - 201 • Agent architectures (e.g., reactive, layered, cognitive, etc.)
 - 202 • Agent theory
 - 203 • Rationality, Game Theory
 - 204 ○ Decision-theoretic agents
 - 205 ○ Markov decision processes (MDP)
 - 206 • Software agents, personal assistants, and information access
 - 207 ○ Collaborative agents
 - 208 ○ Information-gathering agents
 - 209 ○ Believable agents (synthetic characters, modeling emotions in agents)
 - 210 • Learning agents
 - 211 • Multi-agent systems
 - 212 ○ Collaborating agents
 - 213 ○ Agent teams
 - 214 ○ Competitive agents (e.g., auctions, voting)
 - 215 ○ Swarm systems and biologically inspired models
- 216

217 *Learning Outcomes:*

- 218 1. List the defining characteristics of an intelligent agent. [Familiarity]
- 219 2. Characterize and contrast the standard agent architectures. [Assessment]
- 220 3. Describe the applications of agent theory to domains such as software agents, personal assistants, and
221 believable agents. [Familiarity]
- 222 4. Describe the primary paradigms used by learning agents. [Familiarity]
- 223 5. Demonstrate using appropriate examples how multi-agent systems support agent interaction. [Usage]
- 224

225

226 **IS/Natural Language Processing**

227 *[Elective]*

228 (Cross-reference HCI/Design for Non-Mouse Interfaces)

229 *Topics:*

- 230 • Deterministic and stochastic grammars
- 231 • Parsing algorithms
 - 232 ○ CFGs and chart parsers (e.g. CYK)
 - 233 ○ Probabilistic CFGs and weighted CYK
- 234 • Representing meaning / Semantics
 - 235 ○ Logic-based knowledge representations
 - 236 ○ Semantic roles
 - 237 ○ Temporal representations
 - 238 ○ Beliefs, desires, and intentions
- 239 • Corpus-based methods
- 240 • N-grams and HMMs
- 241 • Smoothing and backoff
- 242 • Examples of use: POS tagging and morphology
- 243 • Information retrieval (Cross-reference IM/Information Storage and Retrieval)
 - 244 ○ Vector space model
 - 245 ■ TF & IDF
 - 246 ○ Precision and recall
- 247 • Information extraction
- 248 • Language translation
- 249 • Text classification, categorization
 - 250 ○ Bag of words model

251

252 *Learning Outcomes:*

- 253 1. Define and contrast deterministic and stochastic grammars, providing examples to show the adequacy of
254 each. [Assessment]
- 255 2. Simulate, apply, or implement classic and stochastic algorithms for parsing natural language. [Usage]
- 256 3. Identify the challenges of representing meaning. [Familiarity]
- 257 4. List the advantages of using standard corpora. Identify examples of current corpora for a variety of NLP
258 tasks. [Familiarity]
- 259 5. Identify techniques for information retrieval, language translation, and text classification. [Familiarity]

260

261

262 **IS/Advanced Machine Learning**

263 *[Elective]*

264 **Topics:**

- 265 • Definition and examples of broad variety of machine learning tasks
- 266 • General statistical-based learning, parameter estimation (maximum likelihood)
- 267 • Inductive logic programming (ILP)
- 268 • Supervised learning
 - 269 ○ Learning decision trees
 - 270 ○ Learning neural networks
 - 271 ○ Support vector machines (SVMs)
- 272 • Ensembles
- 273 • Nearest-neighbor algorithms
- 274 • Unsupervised Learning and clustering
 - 275 ○ EM
 - 276 ○ K-means
 - 277 ○ Self-organizing maps
- 278 • Semi-supervised learning
- 279 • Learning graphical models (Cross-reference IS/Reasoning under Uncertainty)
- 280 • Performance evaluation (such as cross-validation, area under ROC curve)
- 281 • Learning theory
- 282 • The problem of overfitting, the curse of dimensionality
- 283 • Reinforcement learning
 - 284 ○ Exploration vs. exploitation trade-off
 - 285 ○ Markov decision processes
 - 286 ○ Value and policy iteration
- 287 • Application of Machine Learning algorithms to Data Mining (Cross-reference IM/Data Mining)
- 288

289 **Learning Outcomes:**

- 290 1. Explain the differences among the three main styles of learning: supervised, reinforcement, and
- 291 unsupervised. [Familiarity]
- 292 2. Implement simple algorithms for supervised learning, reinforcement learning, and unsupervised learning.
- 293 [Usage]
- 294 3. Determine which of the three learning styles is appropriate to a particular problem domain. [Usage]
- 295 4. Compare and contrast each of the following techniques, providing examples of when each strategy is
- 296 superior: decision trees, neural networks, and belief networks. [Assessment]
- 297 5. Evaluate the performance of a simple learning system on a real-world dataset. [Assessment]
- 298 6. Characterize the state of the art in learning theory, including its achievements and its shortcomings.
- 299 [Familiarity]
- 300 7. Explain the problem of overfitting, along with techniques for detecting and managing the problem. [Usage]
- 301

302

303 **IS/Robotics**

304 *[Elective]*

305 *Topics:*

- 306 • Overview: problems and progress
 - 307 ○ State-of-the-art robot systems, including their sensors and an overview of their sensor processing
 - 308 ○ Robot control architectures, e.g., deliberative vs. reactive control and Braitenberg vehicles
 - 309 ○ World modeling and world models
 - 310 ○ Inherent uncertainty in sensing and in control
- 311 • Configuration space and environmental maps
- 312 • Interpreting uncertain sensor data
- 313 • Localizing and mapping
- 314 • Navigation and control
- 315 • Motion planning
- 316 • Multiple-robot coordination
- 317

318 *Learning Outcomes:*

- 319 1. List capabilities and limitations of today's state-of-the-art robot systems, including their sensors and the
320 crucial sensor processing that informs those systems. [Familiarity]
- 321 2. Integrate sensors, actuators, and software into a robot designed to undertake some task. [Usage]
- 322 3. Program a robot to accomplish simple tasks using deliberative, reactive, and/or hybrid control architectures.
323 [Usage]
- 324 4. Implement fundamental motion planning algorithms within a robot configuration space. [Usage]
- 325 5. Characterize the uncertainties associated with common robot sensors and actuators; articulate strategies for
326 mitigating these uncertainties. [Familiarity]
- 327 6. List the differences among robots' representations of their external environment, including their strengths
328 and shortcomings. [Familiarity]
- 329 7. Compare and contrast at least three strategies for robot navigation within known and/or unknown
330 environments, including their strengths and shortcomings. [Assessment]
- 331 8. Describe at least one approach for coordinating the actions and sensing of several robots to accomplish a
332 single task. [Familiarity]
- 333

334 **IS/Perception and Computer Vision**

335 *[Elective]*

336 *Topics:*

- 337 • Computer vision
 - 338 ○ Image acquisition, representation, processing and properties
 - 339 ○ Shape representation, object recognition and segmentation
 - 340 ○ Motion analysis
- 341 • Audio and speech recognition
- 342 • Modularity in recognition
- 343 • Approaches to pattern recognition [overlapping with machine learning]
 - 344 ○ Classification algorithms and measures of classification quality
 - 345 ○ Statistical techniques
- 346

347 **Learning Outcomes:**

- 348 1. Summarize the importance of image and object recognition in AI and indicate several significant
349 applications of this technology. [Familiarity]
- 350 2. List at least three image-segmentation approaches, such as thresholding, edge-based and region-based
351 algorithms, along with their defining characteristics, strengths, and weaknesses. [Familiarity]
- 352 3. Implement 2d object recognition based on contour- and/or region-based shape representations. [Usage]
- 353 4. Distinguish the goals of sound-recognition, speech-recognition, and speaker-recognition and identify how
354 the raw audio signal will be handled differently in each of these cases. [Familiarity]
- 355 5. Provide at least two examples of a transformation of a data source from one sensory domain to another,
356 e.g., tactile data interpreted as single-band 2d images. [Familiarity]
- 357 6. Implement a feature-extraction algorithm on real data, e.g., an edge or corner detector for images or vectors
358 of Fourier coefficients describing a short slice of audio signal. [Usage]
- 359 7. Implement an algorithm combining features into higher-level percepts, e.g., a contour or polygon from
360 visual primitives or phoneme hypotheses from an audio signal. [Usage]
- 361 8. Implement a classification algorithm that segments input percepts into output categories and quantitatively
362 evaluates the resulting classification. [Usage]
- 363 9. Evaluate the performance of the underlying feature-extraction, relative to at least one alternative possible
364 approach (whether implemented or not) in its contribution to the classification task (8), above.
365 [Assessment]
- 366 10. Describe at least three classification approaches, their prerequisites for applicability, their strengths, and
367 their shortcomings. [Familiarity]
- 368

1 **Networking and Communication (NC)**

2 The Internet and computer networks are now ubiquitous and a growing number of computing
3 activities strongly depend on the correct operation of the underlying network. Networks, both
4 fixed and mobile, are a key part of today's and tomorrow's computing environment. Many
5 computing applications that are used today would not be possible without networks. This
6 dependency on the underlying network is likely to increase in the future.

7 The high-level learning objective of this module can be summarized as follows:

- 8 • Thinking in a networked world. The world is more and more interconnected and the use
9 of networks will continue to increase. Students must understand how the network
10 behaves and the key principles behind the organization and the operation of the computer
11 networks.
- 12 • Continued study. The networking domain is rapidly evolving and a first networking
13 course should be a starting point to other more advanced courses on network design,
14 network management, sensor networks, etc.
- 15 • Principles and practice interact. Networking is real and many of the design choices that
16 involve networks also depend on practical constraints. Students should be exposed to
17 these practical constraints by experimenting with networking, using tools, and writing
18 networked software.

19 There are different ways of organizing a networking course. Some educators prefer a top-down
20 approach, i.e. the course starts from the applications and then explains reliable delivery, routing
21 and forwarding, etc. Other educators prefer a bottom-up approach where the students start with
22 the lower layers and build their understanding of the network, transport and application layers
23 later.

24

25

26 **NC. Networking and Communication (3 Core-Tier1 hours, 7 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|---------------------------|------------------|------------------|--------------------|
| NC/Introduction | 1.5 | | N |
| NC/Networked Applications | 1.5 | | N |
| NC/Reliable Data Delivery | | 2 | N |
| NC/Routing And Forwarding | | 1.5 | N |
| NC/Local Area Networks | | 1.5 | N |
| NC/Resource Allocation | | 1 | N |
| NC/Mobility | | 1 | N |
| NC/Social Networking | | | Y |

27
28

29 **NC/Introduction**

30 *[1.5 Core-Tier1 hours]*

31 (Cross-reference IAS/Network Security, which discusses network security and its applications.)

32 **Topics:**

33 [Core-Tier1]

- 34
- Organization of the Internet (Internet Service Providers, Content Providers, etc.)
 - 35 • Switching techniques (Circuit, packet, etc.)
 - 36 • Physical pieces of a network (hosts, routers, switches, ISPs, wireless, LAN, access point, firewalls, etc.)
 - 37 • Layering principles (encapsulation, multiplexing)
 - 38 • Roles of the different layers (application, transport, network, datalink, physical)
- 39

40 **Learning Outcomes:**

41 [Core-Tier1]

- 42
1. Articulate the organization of the Internet [Familiarity]
 - 43 2. List and define the appropriate network terminology [Familiarity]
 - 44 3. Describe the layered structure of a typical networked architecture [Familiarity]
 - 45 4. Identify the different levels of complexity in a network (edges, core, etc.) [Familiarity]
- 46

47

48 **NC/Networked Applications**

49 *[1.5 Core-Tier1 hours]*

50 **Topics:**

51 [Core-Tier1]

- 52 • Naming and address schemes (DNS, IP addresses, Uniform Resource Identifiers, etc.)
- 53 • Distributed applications (client/server, peer-to-peer, cloud, etc.)
- 54 • HTTP as an application layer protocol
- 55 • Multiplexing with TCP and UDP
- 56 • Socket APIs

57

58 **Learning Outcomes:**

59 [Core-Tier1]

- 60 1. List the differences and the relations between names and addresses in a network [Familiarity]
- 61 2. Define the principles behind naming schemes and resource location [Familiarity]
- 62 3. Implement a simple client-server socket-based application [Usage]

63

64 **NC/Reliable Data Delivery**

65 *[2 Core-Tier2 hours]*

66 This Knowledge Unit is related to SF-Systems Fundamentals. (Cross-reference SF/State-State
67 Transition and SF/Reliability through Redundancy.)

68 **Topics:**

69 [Core-Tier2]

- 70 • Error control (retransmission techniques, timers)
- 71 • Flow control (acknowledgements, sliding window)
- 72 • Performance issues (pipelining)
- 73 • TCP

74

75 **Learning Outcomes:**

76 [Core-Tier2]

- 77 1. Describe the operation of reliable delivery protocols [Familiarity]
- 78 2. List the factors that affect the performance of reliable delivery protocols [Familiarity]
- 79 3. Design and implement a simple reliable protocol [Usage]

80

81 **NC/Routing and Forwarding**

82 *[1.5 Core-Tier2 hours]*

83 **Topics:**

84 [Core-Tier2]

- 85 • Routing versus forwarding

- 86 • Static routing
- 87 • Internet Protocol (IP)
- 88 • Scalability issues (hierarchical addressing)
- 89

90 **Learning Outcomes:**

91 [Core-Tier2]

- 92 1. Describe the organization of the network layer [Familiarity]
- 93 2. Describe how packets are forwarded in an IP networks [Familiarity]
- 94 3. List the scalability benefits of hierarchical addressing [Familiarity]
- 95

96 **NC/Local Area Networks**

97 *[1.5 Core-Tier2 hours]*

98 **Topics:**

99 [Core-Tier2]

- 100 • Multiple Access Problem
- 101 • Common approaches to multiple access (exponential-backoff, time division multiplexing, etc)
- 102 • Local Area Networks
- 103 • Ethernet
- 104 • Switching
- 105

106 **Learning Outcomes:**

107 [Core-Tier2]

- 108 1. Describe how frames are forwarded in an Ethernet network [Familiarity]
- 109 2. Identify the differences between IP and Ethernet [Familiarity]
- 110 3. Describe the steps used in one common approach to the multiple access problem [Familiarity]
- 111 4. Describe the interrelations between IP and Ethernet [Familiarity]
- 112

113 **NC/Resource Allocation**

114 *[1 Core-Tier2 hours]*

115 **Topics:**

116 [Core-Tier2]

- 117 • Need for resource allocation
- 118 • Fixed allocation (TDM, FDM, WDM) versus dynamic allocation
- 119 • End-to-end versus network assisted approaches
- 120 • Fairness
- 121 • Principles of congestion control
- 122 • Approaches to Congestion (Content Distribution Networks, etc)
- 123

124 **Learning Outcomes:**

125 [Core-Tier2]

- 126 1. Describe how resources can be allocated in a network [Familiarity]

- 127 2. Describe the congestion problem in a large network [Familiarity]
128 3. Compare and contrast the fixed and dynamic allocation techniques [Assessment]
129 4. Compare and contrast current approaches to congestion [Assessment]
130

131 **NC/Mobility**

132 *[1 Core-Tier2 hours]*

133 *Topics:*

134 [Core-Tier2]

- 135 • Principles of cellular networks
136 • 802.11 networks
137 • Issues in supporting mobile nodes (home agents)
138

139 *Learning Outcomes:*

140 [Core-Tier2]

- 141 1. Describe the organization of a wireless network [Familiarity]
142 2. Describe how wireless networks support mobile users [Familiarity]
143

144 **NC/Social Networking**

145 *[Elective]*

146 *Topics:*

147 [Elective]

- 148 • Social Networks Overview
149 • Example Social Network Platforms
150 • Structure of social network graphs
151 • Social Network Analysis
152

153 *Learning Outcomes:*

154 [Elective]

- 155 1. Discuss the key principles of social networking [Familiarity]
156 2. Describe how existing social networks operate [Familiarity]
157 3. Construct a social network graph from network data [Usage]
158 4. Analyze a social network to determine who the key people are [Usage]
159 5. Evaluate a given interpretation of a social network question with associated data [Assessment]

1 **Operating Systems (OS)**

2 An operating system defines an abstraction of hardware and manages resource sharing among
3 the computer's users. The topics in this area explain the most basic knowledge of operating
4 systems in the sense of interfacing an operating system to networks, teaching the difference
5 between the kernel and user modes, and developing key approaches to operating system design
6 and implementation. This knowledge area is structured to be complementary to Systems
7 Fundamentals, Networks, Information Assurance, and the Parallel and Distributed Computing
8 knowledge areas. The Systems Fundamentals and Information Assurance knowledge areas are
9 the new ones to include contemporary issues. For example, the Systems Fundamentals includes
10 topics such as performance, virtualization and isolation, and resource allocation and scheduling;
11 Parallel and Distributed Systems knowledge area includes parallelism fundamentals; and
12 Information Assurance includes forensics and security issues in depth. Many courses in
13 Operating Systems will draw material from across these Knowledge Areas.

14 **OS. Operating Systems (4 Core-Tier1 hours; 11 Core Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| OS/Overview of Operating Systems | 2 | | N |
| OS/Operating System Principles | 2 | | N |
| OS/Concurrency | | 3 | N |
| OS/Scheduling and Dispatch | | 3 | N |
| OS/Memory Management | | 3 | N |
| OS/Security and Protection | | 2 | N |
| OS/Virtual Machines | | | Y |
| OS/Device Management | | | Y |
| OS/File Systems | | | Y |
| OS/Real Time and Embedded Systems | | | Y |
| OS/Fault Tolerance | | | Y |
| OS/System Performance Evaluation | | | Y |

16 OS/Overview of Operating Systems

17 [2 Core-Tier1 hours]

18 *Topics:*

- 19 • Role and purpose of the operating system
- 20 • Functionality of a typical operating system
- 21 • Mechanisms to support client-server models, hand-held devices
- 22 • Design issues (efficiency, robustness, flexibility, portability, security, compatibility)
- 23 • Influences of security, networking, multimedia, windows
- 24

25 *Learning Objectives:*

- 26 1. Explain the objectives and functions of modern operating systems [Familiarity].
- 27 2. Analyze the tradeoffs inherent in operating system design [Usage].
- 28 3. Describe the functions of a contemporary operating system with respect to convenience, efficiency, and the
- 29 ability to evolve [Familiarity].
- 30 4. Discuss networked, client-server, distributed operating systems and how they differ from single user
- 31 operating systems [Familiarity].
- 32 5. Identify potential threats to operating systems and the security features design to guard against them
- 33 [Familiarity].
- 34

35 OS/Operating System Principles

36 [2 Core-Tier1 hours]

37 *Topics:*

- 38 • Structuring methods (monolithic, layered, modular, micro-kernel models)
- 39 • Abstractions, processes, and resources
- 40 • Concepts of application program interfaces (APIs)
- 41 • Application needs and the evolution of hardware/software techniques
- 42 • Device organization
- 43 • Interrupts: methods and implementations
- 44 • Concept of user/system state and protection, transition to kernel mode
- 45

46 *Learning Objectives:*

- 47 1. Explain the concept of a logical layer [Familiarity].
- 48 2. Explain the benefits of building abstract layers in hierarchical fashion [Familiarity].
- 49 3. Defend the need for APIs and middleware [Assessment].
- 50 4. Describe how computing resources are used by application software and managed by system software
- 51 [Familiarity].
- 52 5. Contrast kernel and user mode in an operating system [Usage].
- 53 6. Discuss the advantages and disadvantages of using interrupt processing [Familiarity].
- 54 7. Explain the use of a device list and driver I/O queue [Familiarity].
- 55

56

57 **OS/Concurrency**

58 *[3 Core-Tier2 hours]*

59 *Topics:*

- 60 • States and state diagrams (cross reference SF/State-State Transition-State Machines)
- 61 • Structures (ready list, process control blocks, and so forth)
- 62 • Dispatching and context switching
- 63 • The role of interrupts
- 64 • Managing atomic access to OS objects
- 65 • Implementing synchronization primitives
- 66 • Multiprocessor issues (spin-locks, reentrancy) (cross reference SF/Parallelism)

67

68 *Learning Objectives:*

- 69 1. Describe the need for concurrency within the framework of an operating system [Familiarity].
- 70 2. Demonstrate the potential run-time problems arising from the concurrent operation of many separate tasks
- 71 [Usage].
- 72 3. Summarize the range of mechanisms that can be employed at the operating system level to realize
- 73 concurrent systems and describe the benefits of each [Familiarity].
- 74 4. Explain the different states that a task may pass through and the data structures needed to support the
- 75 management of many tasks [Familiarity].
- 76 5. Summarize techniques for achieving synchronization in an operating system (e.g., describe how to
- 77 implement a semaphore using OS primitives) [Familiarity].
- 78 6. Describe reasons for using interrupts, dispatching, and context switching to support concurrency in an
- 79 operating system [Familiarity].
- 80 7. Create state and transition diagrams for simple problem domains [Usage].

81

82 **OS/Scheduling and Dispatch**

83 *[3 Core-Tier2 hours]*

84 *Topics:*

- 85 • Preemptive and non-preemptive scheduling (cross reference SF/Resource Allocation and Scheduling,
- 86 PD/Parallel Performance)
- 87 • Schedulers and policies (cross reference SF/Resource Allocation and Scheduling, PD/Parallel Performance)
- 88 • Processes and threads (cross reference SF/computational paradigms)
- 89 • Deadlines and real-time issues

90

91 *Learning Objectives:*

- 92 1. Compare and contrast the common algorithms used for both preemptive and non-preemptive scheduling of
- 93 tasks in operating systems, such as priority, performance comparison, and fair-share schemes [Usage].
- 94 2. Describe relationships between scheduling algorithms and application domains [Familiarity].
- 95 3. Discuss the types of processor scheduling such as short-term, medium-term, long-term, and I/O
- 96 [Familiarity].
- 97 4. Describe the difference between processes and threads [Usage].
- 98 5. Compare and contrast static and dynamic approaches to real-time scheduling [Usage].
- 99 6. Discuss the need for preemption and deadline scheduling [Familiarity].
- 100 7. Identify ways that the logic embodied in scheduling algorithms are applicable to other domains, such as
- 101 disk I/O, network scheduling, project scheduling, and problems beyond computing [Usage].

102

103 OS/Memory Management

104 *[3 Core-Tier2 hours]*

105 *Topics:*

- 106 • Review of physical memory and memory management hardware
- 107 • Working sets and thrashing
- 108 • Caching

109

110 *Learning Objectives:*

- 111 1. Explain memory hierarchy and cost-performance trade-offs [Familiarity].
- 112 2. Summarize the principles of virtual memory as applied to caching and paging [Familiarity].
- 113 3. Evaluate the trade-offs in terms of memory size (main memory, cache memory, auxiliary memory) and
- 114 processor speed [Assessment].
- 115 4. Defend the different ways of allocating memory to tasks, citing the relative merits of each [Assessment].
- 116 5. Describe the reason for and use of cache memory (performance and proximity, different dimension of how
- 117 caches complicate isolation and VM abstraction) [Familiarity].
- 118 6. Discuss the concept of thrashing, both in terms of the reasons it occurs and the techniques used to recognize
- 119 and manage the problem [Familiarity].

120

121 OS/Security and Protection

122 *[2 Core-Tier2 hours]*

123 *Topics:*

- 124 • Overview of system security
- 125 • Policy/mechanism separation
- 126 • Security methods and devices
- 127 • Protection, access control, and authentication
- 128 • Backups

129

130 *Learning Objectives:*

- 131 1. Defend the need for protection and security in an OS (cross reference IAS/Security Architecture and
- 132 Systems Administration/Investigating Operating Systems Security for various systems) [Assessment].
- 133 2. Summarize the features and limitations of an operating system used to provide protection and security
- 134 (cross reference IAS/Security Architecture and Systems Administration) [Familiarity].
- 135 3. Explain the mechanisms available in an OS to control access to resources (cross reference IAS/Security
- 136 Architecture and Systems Administration/Access Control/Configuring systems to operate securely as an IT
- 137 system) [Familiarity].
- 138 4. Carry out simple system administration tasks according to a security policy, for example creating accounts,
- 139 setting permissions, applying patches, and arranging for regular backups (cross reference IAS/Security
- 140 Architecture and Systems Administration) [Usage].

141

142

143 **OS/Virtual Machines**

144 *[Elective]*

145 *Topics:*

- 146 • Types of virtualization (Hardware/Software, OS, Server, Service, Network, etc.)
- 147 • Paging and virtual memory
- 148 • Virtual file systems
- 149 • Virtual file
- 150 • Hypervisors
- 151 • Portable virtualization; emulation vs. isolation
- 152 • Cost of virtualization
- 153

154 *Learning Objectives:*

- 155 1. Explain the concept of virtual memory and how it is realized in hardware and software [Familiarity].
- 156 2. Differentiate emulation and isolation [Familiarity].
- 157 3. Evaluate virtualization trade-offs [Assessment].
- 158 4. Discuss hypervisors and the need for them in conjunction with different types of hypervisors [Usage].
- 159

160 **OS/Device Management**

161 *[Elective]*

162 *Topics:*

- 163 • Characteristics of serial and parallel devices
- 164 • Abstracting device differences
- 165 • Buffering strategies
- 166 • Direct memory access
- 167 • Recovery from failures
- 168

169 *Learning Objectives:*

- 170 1. Explain the key difference between serial and parallel devices and identify the conditions in which each is
- 171 appropriate [Familiarity].
- 172 2. Identify the relationship between the physical hardware and the virtual devices maintained by the operating
- 173 system [Usage].
- 174 3. Explain buffering and describe strategies for implementing it [Familiarity].
- 175 4. Differentiate the mechanisms used in interfacing a range of devices (including hand-held devices,
- 176 networks, multimedia) to a computer and explain the implications of these for the design of an operating
- 177 system [Usage].
- 178 5. Describe the advantages and disadvantages of direct memory access and discuss the circumstances in
- 179 which its use is warranted [Usage].
- 180 6. Identify the requirements for failure recovery [Familiarity].
- 181 7. Implement a simple device driver for a range of possible devices [Usage].
- 182

183

184 **OS/File Systems**

185 *[Elective]*

186 *Topics:*

- 187 • Files: data, metadata, operations, organization, buffering, sequential, nonsequential
- 188 • Directories: contents and structure
- 189 • File systems: partitioning, mount/unmount, virtual file systems
- 190 • Standard implementation techniques
- 191 • Memory-mapped files
- 192 • Special-purpose file systems
- 193 • Naming, searching, access, backups
- 194 • Journaling and log-structured file systems
- 195

196 *Learning Objectives:*

- 197 1. Summarize the full range of considerations in the design of file systems [Familiarity].
- 198 2. Compare and contrast different approaches to file organization, recognizing the strengths and weaknesses
- 199 of each [Usage].
- 200 3. Summarize how hardware developments have led to changes in the priorities for the design and the
- 201 management of file systems [Familiarity].
- 202 4. Summarize the use of journaling and how log-structured file systems enhance fault tolerance [Familiarity].
- 203

204 **OS/Real Time and Embedded Systems**

205 *[Elective]*

206 *Topics:*

- 207 • Process and task scheduling
- 208 • Memory/disk management requirements in a real-time environment
- 209 • Failures, risks, and recovery
- 210 • Special concerns in real-time systems
- 211

212 *Learning Objectives:*

- 213 1. Describe what makes a system a real-time system [Familiarity].
- 214 2. Explain the presence of and describe the characteristics of latency in real-time systems [Familiarity].
- 215 3. Summarize special concerns that real-time systems present and how these concerns are addressed
- 216 [Familiarity].
- 217

218 **OS/Fault Tolerance**

219 *[Elective]*

220 *Topics:*

- 221 • Fundamental concepts: reliable and available systems (cross reference SF/Reliability through Redundancy)
- 222 • Spatial and temporal redundancy (cross reference SF/Reliability through Redundancy)
- 223 • Methods used to implement fault tolerance

- 224 • Examples of OS mechanisms for detection, recovery, restart to implement fault tolerance, use of these
225 techniques for the OS's own services
226

227 ***Learning Objectives:***

- 228 1. Explain the relevance of the terms fault tolerance, reliability, and availability [Familiarity].
229 2. Outline the range of methods for implementing fault tolerance in an operating system [Familiarity].
230 3. Explain how an operating system can continue functioning after a fault occurs [Familiarity].
231

232 **OS/System Performance Evaluation**

233 ***[Elective]***

234 ***Topics:***

- 235 • Why system performance needs to be evaluated (cross reference SF/Performance/Figures of performance
236 merit)
237 • What is to be evaluated (cross reference SF/Performance/Figures of performance merit)
238 • Policies for caching, paging, scheduling, memory management, security, and so forth
239 • Evaluation models: deterministic, analytic, simulation, or implementation-specific
240 • How to collect evaluation data (profiling and tracing mechanisms)
241

242 ***Learning Objectives:***

- 243 1. Describe the performance measurements used to determine how a system performs [Familiarity].
244 2. Explain the main evaluation models used to evaluate a system [Familiarity].

1 **Platform-Based Development (PBD)**

2 Platform-based development is concerned with the design and development of software
3 applications that reside on specific software platforms. In contrast to general purpose
4 programming, platform-based development takes into account platform-specific constraints. For
5 instance web programming, multimedia development, mobile computing, app development, and
6 robotics are examples of relevant platforms which provide specific services/APIs/hardware
7 which constrain development. Such platforms are characterized by the use of specialized APIs,
8 distinct delivery/update mechanisms, and being abstracted away from the machine level.
9 Platform-based development may be applied over a wide breadth of ecosystems.

10 While we recognize that some platforms (e.g., web development) are prominent, we are also
11 cognizant of the fact that no particular platform should be specified as a requirement in the
12 CS2013 curricular guidelines. Consequently, this Knowledge Area highlights many of the
13 platforms which have become popular, without including any such platform in the core
14 curriculum. We note that the general skill of developing with respect to an API or a constrained
15 environment is covered in other Knowledge Areas, such as SDF-Software Development
16 Fundamentals. Platform-based development further emphasizes such general skills within the
17 context of particular platforms.

18

19 **PBD. Platform-Based Development (Elective)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--------------------------|------------------|------------------|--------------------|
| PBD/Introduction | | | Y |
| PBD/Web Platforms | | | Y |
| PBD/Mobile Platforms | | | Y |
| PBD/Industrial Platforms | | | Y |
| PBD/Game Platforms | | | Y |

20

21

22 **PBD/Introduction**

23 *[Elective]*

24 This unit describes the fundamental differences that Platform-Based Development has over
25 traditional software development.

26 *Topics:*

- 27 • Overview of platforms (Web, Mobile, Game, Industrial etc)
- 28 • Programming via platform-specific APIs
- 29 • Overview of Platform Languages (Objective C, HTML5, etc)
- 30 • Programming under platform constraints

31

32 *Learning Outcomes:*

33 [Elective]

- 34 1. Describe how platform-based development differs from general purpose programming [Familiarity]
- 35 2. List characteristics of platform languages [Familiarity]
- 36 3. Write and execute a simple platform-based program [Usage]
- 37 4. List the advantages and disadvantages of programming with platform constraints [Familiarity]

38

39 **PBD/Web Platforms**

40 *[Elective]*

41 *Topics:*

- 42 • Web programming languages (HTML5, Java Script, PHP, CSS, etc.)
- 43 • Web platform constraints
- 44 • Software as a Service (SaaS)
- 45 • Web standards

46

47 *Learning Outcomes:*

48 [Elective]

- 49 1. Design and Implement a simple web application [Usage]
- 50 2. Describe the constraints that the web puts on developers [Familiarity]
- 51 3. Compare and contrast web programming with general purpose programming [Assessment]
- 52 4. Describe the differences between Software-as-a-Service and traditional software products [Familiarity]
- 53 5. Discuss how web standards impact software development [Familiarity]
- 54 6. Review an existing web application against a current web standard [Assessment]

55

56 **PBD/Mobile Platforms**

57 *[Elective]*

58 *Topics:*

- 59 • Mobile Programming Languages (Objective C, Java Script, Java, etc.)
- 60 • Challenges with mobility and wireless communication
- 61 • Location-aware applications

- 62 • Performance / power tradeoffs
- 63 • Mobile platform constraints
- 64 • Emerging Technologies
- 65

66 **Learning Outcomes:**

67 [Elective]

- 68 1. Design and implement a mobile application for a given mobile platform. [Usage]
- 69 2. Discuss the constraints that mobile platforms put on developers [Familiarity]
- 70 3. Discuss the performance vs. power tradeoff [Familiarity]
- 71 4. Compare and Contrast mobile programming with general purpose programming [Assessment]
- 72

73 **PBD/Industrial Platforms**

74 *[Elective]*

75 This knowledge unit is related to IS/Robotics.

76 **Topics:**

- 77 • Types of Industrial Platforms (Mathematic, Robotics, Industrial Controls, etc.)
- 78 • Robotic Software and its Architecture
- 79 • Domain Specific Languages
- 80 • Industrial Platform Constraints
- 81

82 **Learning Outcomes:**

83 [Elective]

- 84 1. Design and implement an industrial application on a given platform (Lego Mindstorms, Matlab, etc.)
- 85 [Usage]
- 86 2. Compare and contrast domain specific languages with general purpose programming languages.
- 87 [Assessment]
- 88 3. Discuss the constraints that a given industrial platforms impose on developers [Familiarity]
- 89

90 **PBD/Game Platforms**

91 *[Elective]*

92 **Topics:**

- 93 • Types of Game Platforms (XBox, Wii, PlayStation, etc)
- 94 • Game Platform Languages (C++, Java, Lua, Python, etc)
- 95 • Game Platform Constraints
- 96

97 **Learning Outcomes:**

98 [Elective]

- 99 1. Design and Implement a simple application on a game platform. [Usage]
- 100 2. Describe the constraints that game platforms impose on developers. [Familiarity]
- 101 3. Compare and contrast game programming with general purpose programming [Assessment]

1 **Parallel and Distributed Computing (PD)**

2 The past decade has brought explosive growth in multiprocessor computing, including multi-core
3 processors and distributed data centers. As a result, parallel and distributed computing has
4 moved from a largely elective topic to become more of a core component of undergraduate
5 computing curricula. Both parallel and distributed computing entail the logically simultaneous
6 execution of multiple processes, whose operations have the potential to interleave in complex
7 ways. Parallel and distributed computing builds on foundations in many areas, including an
8 understanding of fundamental systems concepts such as concurrency and parallel execution,
9 consistency in state/memory manipulation, and latency. Communication and coordination
10 among processes is rooted in the message-passing and shared-memory models of computing and
11 such algorithmic concepts as atomicity, consensus, and conditional waiting. Achieving speedup
12 in practice requires an understanding of parallel algorithms, strategies for problem
13 decomposition, system architecture, detailed implementation strategies, and performance
14 analysis and tuning. Distributed systems highlight the problems of security and fault tolerance,
15 emphasize the maintenance of replicated state, and introduce additional issues that bridge to
16 computer networking.

17 Because parallelism interacts with so many areas of computing, including at least algorithms,
18 languages, systems, networking, and hardware, many curricula will put different parts of the
19 knowledge area in different courses, rather than in a dedicated course. While we acknowledge
20 that computer science is moving in this direction and may reach that point, in 2013 this process is
21 still in flux and we feel it provides more useful guidance to curriculum designers to aggregate the
22 fundamental parallelism topics in one place. Note, however, that the fundamentals of
23 concurrency and mutual exclusion appear in Systems Fundamentals. Many curricula may
24 choose to introduce parallelism and concurrency in the same course (see below for the distinction
25 intended by these terms). Further, we note that the topics and learning outcomes listed below
26 include only brief mentions of purely elective coverage. At the present time, there is too much
27 diversity in topics that share little in common (including for example, parallel scientific
28 computing, process calculi, and non-blocking data structures) to recommend particular topics be
29 covered in elective courses.

30 Because the terminology of parallel and distributed computing varies among communities, we
31 provide here brief descriptions of the intended senses of a few terms. This list is not exhaustive
32 or definitive, but is provided for the sake of clarity:

- 33 • *Parallelism*: Using additional computational resources simultaneously, usually for
34 speedup.
- 35 • *Concurrency*: Efficiently and correctly managing concurrent access to resources.
- 36 • *Activity*: A computation that may proceed concurrently with others; for example a
37 program, process, thread, or active parallel hardware component.
- 38 • *Atomicity*: Rules and properties governing whether an action is observationally
39 indivisible; for example setting all of the bits in a word, transmitting a single packet, or
40 completing a transaction.
- 41 • *Consensus*: Agreement among two or more activities about a given predicate; for
42 example the value of a counter, the owner of a lock, or the termination of a thread.
- 43 • *Consistency*: Rules and properties governing agreement about the values of variables
44 written, or messages produced, by some activities and used by others (thus possibly
45 exhibiting a *data race*); for example, *sequential consistency*, stating that the values of all
46 variables in a shared memory parallel program are equivalent to that of a single program
47 performing some interleaving of the memory accesses of these activities.
- 48 • *Multicast*: A message sent to possibly many recipients, generally without any constraints
49 about whether some recipients receive the message before others. An *event* is a multicast
50 message sent to a designated set of *listeners* or *subscribers*.

51 As multi-processor computing continues to grow in the coming years, so too will the role of
52 parallel and distributed computing in undergraduate computing curricula. In addition to the
53 guidelines presented here, we also direct the interested reader to the document entitled
54 "NSF/TCPP Curriculum Initiative on Parallel and Distributed Computing - Core Topics for
55 Undergraduates", available from the website: <http://www.cs.gsu.edu/~tcpp/curriculum/>.

56 **General cross-referencing note:** Systems Fundamentals also contains an introduction to
57 parallelism (SF/Computational Paradigms, SF/System Support for Parallelism, SF/Performance).

58 The introduction to parallelism in SF complements the one here and there is no ordering
 59 constraint between them. In SF, the idea is to provide a unified view of the system support for
 60 simultaneous execution at multiple levels of abstraction (parallelism is inherent in gates,
 61 processors, operating systems, servers, etc.), whereas here the focus is on a preliminary
 62 understanding of parallelism as a computing primitive and the complications that arise in parallel
 63 and concurrent programming. Given these different perspectives, the hours assigned to each are
 64 not redundant: the layered systems view and the high-level computing concepts are accounted
 65 for separately in terms of the core hours.

66

67 **PD. Parallel and Distributed Computing (5 Core-Tier1 hours, 9 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| PD/Parallelism Fundamentals | 2 | | N |
| PD/Parallel Decomposition | 1 | 3 | N |
| PD/Communication and Coordination | 1 | 3 | Y |
| PD/Parallel Algorithms, Analysis, and Programming | | 3 | Y |
| PD/Parallel Architecture | 1 | 1 | Y |
| PD/Parallel Performance | | | Y |
| PD/Distributed Systems | | | Y |
| PD/Cloud Computing | | | Y |
| PD/Formal Models and Semantics | | | Y |

68

69

70 **PD/Parallelism Fundamentals**

71 *[2 Core-Tier1 hours]*

72 Build upon students' familiarity with the notion of basic parallel execution--a concept addressed
73 in Systems Fundamentals--to delve into the complicating issues that stem from this notion, such
74 as race conditions and liveness.

75 (Cross-reference SF/Computational Paradigms and SF/System Support for Parallelism)

76 *Topics:*

77 [Core-Tier1]

- 78 • Multiple simultaneous computations
- 79 • Goals of parallelism (e.g., throughput) versus concurrency (e.g., controlling access to shared resources)
- 80 • Programming constructs for creating parallelism, communicating, and coordinating
- 81 • Programming errors not found in sequential programming
 - 82 ○ Data races (simultaneous read/write or write/write of shared state)
 - 83 ○ Higher-level races (interleavings violating program intention)
 - 84 ○ Lack of liveness/progress (deadlock, starvation)
 - 85

86 *Learning outcomes:*

87 [Core-Tier1]

- 88 1. Distinguish using computational resources for a faster answer from managing efficient access to a shared
89 resource [Familiarity]
- 90 2. Distinguish multiple sufficient programming constructs for synchronization that may be inter-
91 implementable but have complementary advantages [Familiarity]
- 92 3. Distinguish data races from higher level races [Familiarity]
- 93

94 **PD/Parallel Decomposition**

95 *[1 Core-Tier1 hour, 3 Core-Tier2 hours]*

96 (Cross-reference SF/System Support for Parallelism)

97 *Topics:*

98 [Core-Tier1]

- 99 • Need for communication and coordination/synchronization
- 100 • Independence and partitioning
- 101

102 [Core-Tier2]

- 103 • Basic knowledge of parallel decomposition concepts (cross-reference SF/System Support for Parallelism)
- 104 • Task-based decomposition
 - 105 ○ Implementation strategies such as threads
- 106 • Data-parallel decomposition
 - 107 ○ Strategies such as SIMD and MapReduce

- 108 • Actors and reactive processes (e.g., request handlers)
109

110 **Learning outcomes:**

111 [Core-Tier1]

- 112 1. Explain why synchronization is necessary in a specific parallel program [Usage]
113

114 [Core-Tier2]

- 115 2. Write a correct and scalable parallel algorithm [Usage]
116 3. Parallelize an algorithm by applying task-based decomposition [Usage]
117 4. Parallelize an algorithm by applying data-parallel decomposition [Usage]
118

119 **PD/Communication and Coordination**

120 *[1 Core-Tier1 hour, 3 Core-Tier2 hours]*

121 (Cross-reference OS/Concurrency for mechanism implementation issues.)

122 **Topics:**

123 [Core-Tier1]

- 124 • Shared Memory
125 ○ Consistency, and its role in programming language guarantees for data-race-free programs
126

127 [Core-Tier2]

- 128 • Consistency in shared memory models
129 • Message passing
130 ○ Point-to-point versus multicast (or event-based) messages
131 ○ Blocking versus non-blocking styles for sending and receiving messages
132 ○ Message buffering (cross-reference PF/Fundamental Data Structures/Queues)
133 • Atomicity
134 ○ Specifying and testing atomicity and safety requirements
135 ○ Granularity of atomic accesses and updates, and the use of constructs such as critical sections or
136 transactions to describe them
137 ○ Mutual Exclusion using locks, semaphores, monitors, or related constructs
138 ■ Potential for liveness failures and deadlock (causes, conditions, prevention)
139 ○ Composition
140 ■ Composing larger granularity atomic actions using synchronization
141 ■ Transactions, including optimistic and conservative approaches
142

143 [Elective]

- 144 • Consensus
145 ○ (Cyclic) barriers, counters, or related constructs
146 • Conditional actions
147 ○ Conditional waiting (e.g., using condition variables)
148

149 **Learning outcomes:**

150 [Core-Tier1]

151 1. Use mutual exclusion to avoid a given race condition [Usage]

152

153 [Core-Tier2]

154 2. Give an example of an ordering of accesses among concurrent activities that is not sequentially consistent
155 [Familiarity]

156 3. Give an example of a scenario in which blocking message sends can deadlock [Usage]

157 4. Explain when and why multicast or event-based messaging can be preferable to alternatives [Familiarity]

158 5. Write a program that correctly terminates when all of a set of concurrent tasks have completed [Usage]

159 6. Use a properly synchronized queue to buffer data passed among activities [Usage]

160 7. Explain why checks for preconditions, and actions based on these checks, must share the same unit of
161 atomicity to be effective [Familiarity]

162 8. Write a test program that can reveal a concurrent programming error; for example, missing an update when
163 two activities both try to increment a variable [Usage]

164 9. Describe at least one design technique for avoiding liveness failures in programs using multiple locks or
165 semaphores [Familiarity]

166 10. Describe the relative merits of optimistic versus conservative concurrency control under different rates of
167 contention among updates [Familiarity]

168 11. Give an example of a scenario in which an attempted optimistic update may never complete [Familiarity]

169

170 [Elective]

171 12. Use semaphores or condition variables to block threads until a necessary precondition holds [Usage]

172

173 **PD/Parallel Algorithms, Analysis, and Programming**

174 *[3 Core-Tier2 hours]*

175 **Topics:**

176 [Core-Tier2]

177 • Critical paths, work and span, and the relation to Amdahl's law (cross-reference SF/Performance)

178 • Speed-up and scalability

179 • Naturally (embarrassingly) parallel algorithms

180 • Parallel algorithmic patterns (divide-and-conquer, map and reduce, others)

181 ○ Specific algorithms (e.g., parallel MergeSort)

182

183 [Elective]

184 • Parallel graph algorithms (e.g., parallel shortest path, parallel spanning tree) (cross-reference
185 AL/Algorithmic Strategies/Divide-and-conquer)

186 • Producer-consumer and pipelined algorithms

187

188 **Learning outcomes:**

189 [Core-Tier2]

190 1. Define “critical path”, “work”, and “span” [Familiarity]

- 191 2. Compute the work and span, and determine the critical path with respect to a parallel execution diagram
 192 [Usage]
 193 3. Define “speed-up” and explain the notion of an algorithm’s scalability in this regard [Familiarity]
 194 4. Identify independent tasks in a program that may be parallelized [Usage]
 195 5. Characterize features of a workload that allow or prevent it from being naturally parallelized [Familiarity]
 196 6. Implement a parallel divide-and-conquer (and/or graph algorithm) and empirically measure its performance
 197 relative to its sequential analog [Usage]
 198 7. Decompose a problem (e.g., counting the number of occurrences of some word in a document) via map and
 199 reduce operations [Usage]
 200

201 [Elective]

- 202 8. Provide an example of a problem that fits the producer-consumer paradigm [Familiarity]
 203 9. Give examples of problems where pipelining would be an effective means of parallelization [Familiarity]
 204 10. Identify issues that arise in producer-consumer algorithms and mechanisms that may be used for addressing
 205 them [Familiarity]
 206

207 **PD/Parallel Architecture**

208 *[1 Core-Tier1 hour, 1 Core-Tier2 hour]*

209 The topics listed here are related to knowledge units in the Architecture and Organization area
 210 (AR/Assembly Level Machine Organization and AR/Multiprocessing and Alternative
 211 Architectures). Here, we focus on parallel architecture from the standpoint of applications,
 212 whereas the Architecture and Organization area presents the topic from the hardware
 213 perspective.

214 [Core-Tier1]

- 215 • Multicore processors
- 216 • Shared vs. distributed memory

217
 218 [Core-Tier2]

- 219 • Symmetric multiprocessing (SMP)
- 220 • SIMD, vector processing

221
 222 [Elective]

- 223 • GPU, co-processing
 - 224 • Flynn’s taxonomy
 - 225 • Instruction level support for parallel programming
 - 226 ○ Atomic instructions such as Compare and Set
 - 227 • Memory issues
 - 228 ○ Multiprocessor caches and cache coherence
 - 229 ○ Non-uniform memory access (NUMA)
 - 230 • Topologies
 - 231 ○ Interconnects
 - 232 ○ Clusters
 - 233 ○ Resource sharing (e.g., buses and interconnects)
- 234

235 **Learning outcomes:**

236 [Core-Tier1]

237 1. Explain the differences between shared and distributed memory [Familiarity]

238

239 [Core-Tier2]

240 2. Describe the SMP architecture and note its key features [Familiarity]

241 3. Characterize the kinds of tasks that are a natural match for SIMD machines [Familiarity]

242

243 [Elective]

244 4. Explain the features of each classification in Flynn's taxonomy [Familiarity]

245 5. Describe the challenges in maintaining cache coherence [Familiarity]

246 6. Describe the key features of different distributed system topologies [Familiarity]

247

248 **PD/Parallel Performance**

249 [*Elective*]

250 **Topics:**

251 • Load balancing

252 • Performance measurement

253 • Scheduling and contention (cross-reference OS/Scheduling and Dispatch)

254 • Data management

255 ○ Non-uniform communication costs due to proximity (cross-reference SF/Proximity)

256 ○ Cache effects (e.g., false sharing)

257 ○ Maintaining spatial locality

258 • Impact of composing multiple concurrent components

259 • Power usage and management

260

261 **Learning outcomes:**

262 [Elective]

263 1. Calculate the implications of Amdahl's law for a particular parallel algorithm [Usage]

264 2. Describe how data distribution/layout can affect an algorithm's communication costs [Familiarity]

265 3. Detect and correct a load imbalance [Usage]

266 4. Detect and correct an instance of false sharing [Usage]

267 5. Explain the impact of scheduling on parallel performance [Familiarity]

268 6. Explain performance impacts of data locality [Familiarity]

269 7. Explain the impact and trade-off related to power usage on parallel performance [Familiarity]

270

271

272 **PD/Distributed Systems**

273 *[Elective]*

274 *Topics:*

- 275 • Faults (cross-reference OS/Fault Tolerance)
- 276 ○ Network-based (including partitions) and node-based failures
- 277 ○ Impact on system wide guarantees (e.g., availability)
- 278 • Distributed message sending
- 279 ○ Data conversion and transmission
- 280 ○ Sockets
- 281 ○ Message sequencing
- 282 ○ Buffering, retrying, and dropping messages
- 283 • Distributed system design tradeoffs
- 284 ○ Latency versus throughput
- 285 ○ Consistency, availability, partition tolerance
- 286 • Distributed service design
- 287 ○ Stateful versus stateless protocols and services
- 288 ○ Session (connection-based) designs
- 289 ○ Reactive (IO-triggered) and multithreaded designs
- 290 • Core distributed algorithms
- 291 ○ Election, discovery
- 292

293 *Learning outcomes:*

294 *[Elective]*

- 295 1. Distinguish network faults from other kinds of failures [Familiarity]
- 296 2. Explain why synchronization constructs such as simple locks are not useful in the presence of distributed
- 297 faults [Familiarity]
- 298 3. Give examples of problems for which consensus algorithms such as leader election are required [Usage]
- 299 4. Write a program that performs any required marshalling and conversion into message units, such as
- 300 packets, to communicate interesting data between two hosts [Usage]
- 301 5. Measure the observed throughput and response latency across hosts in a given network [Usage]
- 302 6. Explain why no distributed system can be simultaneously consistent, available, and partition tolerant
- 303 [Familiarity]
- 304 7. Implement a simple server -- for example, a spell checking service [Usage]
- 305 8. Explain the tradeoffs among overhead, scalability, and fault tolerance when choosing a stateful v. stateless
- 306 design for a given service [Familiarity]
- 307 9. Describe the scalability challenges associated with a service growing to accommodate many clients, as well
- 308 as those associated with a service only transiently having many clients [Familiarity]
- 309

310

311 **PD/Cloud Computing**

312 *[Elective]*

313 *Topics:*

- 314 • Internet-Scale computing
 - 315 ○ Task partitioning (cross-reference PD/Parallel Algorithms, Analysis, and Programming)
 - 316 ○ Data access
 - 317 ○ Clusters, grids, and meshes
- 318 • Cloud services
 - 319 ○ Infrastructure as a service
 - 320 ■ Elasticity of resources
 - 321 ■ Platform APIs
 - 322 ○ Software as a service
 - 323 ○ Security
 - 324 ○ Cost management
- 325 • Virtualization (cross-reference SF/Virtualization and Isolation and OS/Virtual Machines)
 - 326 ○ Shared resource management
 - 327 ○ Migration of processes
- 328 • Cloud-based data storage
 - 329 ○ Shared access to weakly consistent data stores
 - 330 ○ Data synchronization
 - 331 ○ Data partitioning
 - 332 ○ Distributed file systems (cross-reference IM/Distributed Databases)
 - 333 ○ Replication

335 *Learning outcomes:*

336 *[Elective]*

- 337 1. Discuss the importance of elasticity and resource management in cloud computing. [Familiarity]
- 338 2. Explain strategies to synchronize a common view of shared data across a collection of devices [Familiarity]
- 339 3. Explain the advantages and disadvantages of using virtualized infrastructure [Familiarity]
- 340 4. Deploy an application that uses cloud infrastructure for computing and/or data resources [Usage]
- 341 5. Appropriately partition an application between a client and resources [Usage]

342

343 **PD/Formal Models and Semantics**

344 *[Elective]*

345 *Topics:*

- 346 • Formal models of processes and message passing, including algebras such as Communicating Sequential Processes (CSP) and pi-calculus
- 347
- 348 • Formal models of parallel computation, including the Parallel Random Access Machine (PRAM) and alternatives such as Bulk Synchronous Parallel (BSP)
- 349
- 350 • Models of (relaxed) shared memory consistency and their relation to programming language specifications
- 351 • Algorithmic correctness criteria including linearizability
- 352 • Models of algorithmic progress, including non-blocking guarantees and fairness

353 • Techniques for specifying and checking correctness properties such as atomicity and freedom from data
354 races
355

356 ***Learning outcomes:***

357 [Elective]

- 358 1. Model a concurrent process using a formal model, such as pi-calculus [Usage]
359 2. Explain the characteristics of a particular formal parallel model [Familiarity]
360 3. Formally model a shared memory system to show if it is consistent [Application]
361 4. Use a model to show progress guarantees in a parallel algorithm [Usage]
362 5. Use formal techniques to show that a parallel algorithm is correct with respect to a safety or liveness
363 property [Usage]
364 6. Decide if a specific execution is linearizable or not [Usage]

1 **Programming Languages (PL)**

2 Programming languages are the medium through which programmers precisely describe
3 concepts, formulate algorithms, and reason about solutions. In the course of a career, a computer
4 scientist will work with many different languages, separately or together. Software developers
5 must understand the programming models underlying different languages, and make informed
6 design choices in languages supporting multiple complementary approaches. Computer
7 scientists will often need to learn new languages and programming constructs, and must
8 understand the principles underlying how programming language features are defined,
9 composed, and implemented. The effective use of programming languages, and appreciation of
10 their limitations, also requires a basic knowledge of programming language translation and static
11 program analysis, as well as run-time components such as memory management.

12

13 **PL. Programming Languages (8 Core-Tier1 hours, 20 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| PL/Object-Oriented Programming | 4 | 6 | N |
| PL/Functional Programming | 3 | 4 | N |
| PL/Event-Driven and Reactive Programming | | 2 | N |
| PL/Basic Type Systems | 1 | 4 | N |
| PL/Program Representation | | 1 | N |
| PL/Language Translation and Execution | | 3 | N |
| PL/Syntax Analysis | | | Y |
| PL/Compiler Semantic Analysis | | | Y |
| PL/Code Generation | | | Y |
| PL/Runtime Systems | | | Y |
| PL/Static Analysis | | | Y |
| PL/Advanced Programming Constructs | | | Y |
| PL/Concurrency and Parallelism | | | Y |
| PL/Type Systems | | | Y |
| PL/Formal Semantics | | | Y |
| PL/Language Pragmatics | | | Y |
| PL/Logic Programming | | | Y |

14
15
16
17
18
19
20
21
22
23
24
25

Note:

- Some topics from one or more of the first three Knowledge Units (*Object-Oriented Programming, Functional Programming, Event-Driven and Reactive Programming*) are likely to be integrated with topics in the Software Development Fundamentals Knowledge Area in a curriculum’s introductory courses. Curricula will differ on which topics are integrated in this fashion and which are delayed until later courses on software development and programming languages.
- Some of the most important core learning outcomes are relevant to object-oriented programming, functional programming, and, in fact, all programming. These learning outcomes are *repeated* in the *Object-Oriented Programming* and *Functional*

26 *Programming* Knowledge Units, with a note to this effect. We do not intend that a
27 curriculum necessarily needs to cover them multiple times, though some will. We repeat
28 them only because they do not naturally fit in only one Knowledge Unit.

29

30 **PL/Object-Oriented Programming**

31 *[4 Core-Tier1 hours, 6 Core-Tier2 hours]*

32 *Topics:*

33 [Core-Tier1]

- 34 • Object-oriented design
 - 35 ○ Decomposition into objects carrying state and having behavior
 - 36 ○ Class-hierarchy design for modeling
- 37 • Definition of classes: fields, methods, and constructors
- 38 • Subclasses, inheritance, and method overriding
- 39 • Dynamic dispatch: definition of method-call

40

41 [Core-Tier2]

- 42 • Subtyping (cross-reference PL/Type Systems)
 - 43 ○ Subtype polymorphism; implicit upcasts in typed languages
 - 44 ○ Notion of behavioral replacement: subtypes acting like supertypes
 - 45 ○ Relationship between subtyping and inheritance
- 46 • Object-oriented idioms for encapsulation
 - 47 ○ Privacy and visibility of class members
 - 48 ○ Interfaces revealing only method signatures
 - 49 ○ Abstract base classes
- 50 • Using collection classes, iterators, and other common library components

51

52 *Learning outcomes:*

53 [Core-Tier1]

- 54 1. Compare and contrast (1) the procedural/functional approach—defining a function for each operation with
55 the function body providing a case for each data variant—and (2) the object-oriented approach—defining a
56 class for each data variant with the class definition providing a method for each operation. Understand
57 both as defining a matrix of operations and variants. [Assessment] *This outcome also appears in*
58 *PL/Functional Programming.*
- 59 2. Use subclassing to design simple class hierarchies that allow code to be reused for distinct subclasses.
60 [Usage]
- 61 3. Correctly reason about control flow in a program using dynamic dispatch. [Usage]

62

63 [Core-Tier2]

- 64 4. Explain the relationship between object-oriented inheritance (code-sharing and overriding) and subtyping
65 (the idea of a subtype being usable in a context that expects the supertype). [Familiarity]
- 66 5. Use multiple encapsulation mechanisms, such as function closures, object-oriented interfaces, and support
67 for abstract datatypes, in multiple programming languages. [Usage] *This outcome also appears in*
68 *PL/Functional Programming.*

- 69 6. Define and use iterators and other operations on aggregates, including operations that take functions as
70 arguments, in multiple programming languages, selecting the most natural idioms for each language.
71 [Usage] *This outcome also appears in PL/Functional Programming.*
72

73 **PL/Functional Programming**

74 *[3 Core-Tier1 hours, 4 Core-Tier2 hours]*

75 *Topics:*

76 [Core-Tier1]

- 77 • Effect-free programming
- 78 ○ Function calls have no side effects, facilitating compositional reasoning
- 79 ○ Variables are immutable, preventing unexpected changes to program data by other code
- 80 ○ Data can be freely aliased or copied without introducing unintended effects from mutation
- 81 • Processing structured data (e.g., trees) via functions with cases for each data variant
- 82 ○ Associated language constructs such as discriminated unions and pattern-matching over them
- 83 ○ Functions defined over compound data in terms of functions applied to the constituent pieces
- 84 • First-class functions (taking, returning, and storing functions)

85

86 [Core-Tier2]

- 87 • Function closures (functions using variables in the enclosing lexical environment)
- 88 ○ Basic meaning and definition -- creating closures at run-time by capturing the environment
- 89 ○ Canonical idioms: call-backs, arguments to iterators, reusable code via function arguments
- 90 ○ Using a closure to encapsulate data in its environment
- 91 ○ Currying and partial application
- 92 • Defining higher-order operations on aggregates, especially map, reduce/fold, and filter

93

94 *Learning outcomes:*

95 [Core-Tier1]

- 96 1. Compare and contrast (1) the procedural/functional approach—defining a function for each operation with
97 the function body providing a case for each data variant—and (2) the object-oriented approach—defining a
98 class for each data variant with the class definition providing a method for each operation. Understand
99 both as defining a matrix of operations and variants. [Assessment] *This outcome also appears in*
100 *PL/Object-Oriented Programming.*
- 101 2. Write basic algorithms that avoid assigning to mutable state or considering reference equality. [Usage]
- 102 3. Write useful functions that take and return other functions. [Usage]

103

104 [Core-Tier2]

- 105 4. Correctly reason about variables and lexical scope in a program using function closures. [Usage]
- 106 5. Use multiple encapsulation mechanisms, such as function closures, object-oriented interfaces, and support
107 for abstract datatypes, in multiple programming languages. [Usage] *This outcome also appears in*
108 *PL/Object-Oriented Programming.*
- 109 6. Define and use iterators and other operations on aggregates, including operations that take functions as
110 arguments, in multiple programming languages, selecting the most natural idioms for each language.
111 [Usage] *This outcome also appears in PL/Object-Oriented Programming.*
112

113 **PL/Event-Driven and Reactive Programming**

114 *[2 Core-Tier2 hours]*

115 This material can stand alone or be integrated with other knowledge units on concurrency,
116 asynchrony, and threading to allow contrasting events with threads.

117 **Topics:**

- 118 • Events and event handlers
- 119 • Canonical uses such as GUIs, mobile devices, robots, servers
- 120 • Using a reactive framework
 - 121 ○ Defining event handlers/listeners
 - 122 ○ Main event loop not under event-handler-writer's control
- 123 • Externally-generated events and program-generated events
- 124 • Separation of model, view, and controller
- 125

126 **Learning outcomes:**

- 127 1. Write event handlers for use in reactive systems, such as GUIs. [Usage]
- 128 2. Explain why an event-driven programming style is natural in domains where programs react to external
- 129 events. [Familiarity]
- 130

131 **PL/Basic Type Systems**

132 *[1 Core-Tier1 hour, 4 Core-Tier2 hours]*

133 The core-tier2 hours would be profitably spent both on the core-tier2 topics and on a less shallow
134 treatment of the core-tier1 topics and learning outcomes.

135 **Topics:**

136 [Core-Tier1]

- 137
- 138 • A type as a set of values together with a set of operations
 - 139 ○ Primitive types (e.g., numbers, Booleans)
 - 140 ○ Compound types built from other types (e.g., records, unions, arrays, lists, functions, references)
- 141 • Association of types to variables, arguments, results, and fields
- 142 • Type safety and errors caused by using values inconsistently with their intended types
- 143 • Goals and limitations of static typing
 - 144 ○ Eliminating some classes of errors without running the program
 - 145 ○ Undecidability means static analysis must conservatively approximate program behavior
- 146

147 [Core-Tier2]

- 148
- 149 • Generic types (parametric polymorphism)
 - 150 ○ Definition
 - 151 ○ Use for generic libraries such as collections
 - 152 ○ Comparison with ad hoc polymorphism (overloading) and subtype polymorphism
- 153 • Complementary benefits of static and dynamic typing
 - 154 ○ Errors early vs. errors late/avoided

- 155 ○ Enforce invariants during code development and code maintenance vs. postpone typing decisions
- 156 while prototyping and conveniently allow flexible coding patterns such as heterogeneous
- 157 collections
- 158 ○ Avoid misuse of code vs. allow more code reuse
- 159 ○ Detect incomplete programs vs. allow incomplete programs to run

160
161 **Learning outcomes:**

162 [Core-Tier1]

- 163 1. For multiple programming languages, identify program properties checked statically and program
- 164 properties checked dynamically. Use this knowledge when writing and debugging programs. [Usage]

165
166 [Core-Tier2]

- 167 2. Define and use program pieces (such as functions, classes, methods) that use generic types. [Usage]
- 168 3. Explain benefits and limitations of static typing. [Familiarity]

169
170 **PL/Program Representation**

171 *[1 Core-Tier2 hour]*

172 **Topics:**

- 173 • Programs that take (other) programs as input such as interpreters, compilers, type-checkers, documentation
- 174 generators, etc.
- 175 • Abstract syntax trees; contrast with concrete syntax
- 176 • Data structures to represent code for execution, translation, or transmission

177
178 **Learning outcomes:**

- 179 1. Write a program to process some representation of code for some purpose, such as an interpreter, an
- 180 expression optimizer, a documentation generator, etc. [Usage]

181
182 **PL/Language Translation and Execution**

183 *[3 Core-Tier2 hours]*

184 **Topics:**

- 185 • Interpretation vs. compilation to native code vs. compilation to portable intermediate representation
- 186 • Language translation pipeline: parsing, optional type-checking, translation, linking, execution
- 187 ○ Execution as native code or within a virtual machine
- 188 ○ Alternatives like dynamic loading and dynamic (or “just-in-time”) code generation
- 189 • Run-time representation of core language constructs such as objects (method tables) and first-class
- 190 functions (closures)
- 191 • Run-time layout of memory: call-stack, heap, static data
- 192 ○ Implementing loops, recursion, and tail calls
- 193 • Memory management
- 194 ○ Manual memory management: allocating, deallocating, and reusing heap memory
- 195 ○ Automated memory management: garbage collection as an automated technique using the notion
- 196 of reachability

197

198 **Learning outcomes:**

- 199 1. Distinguish syntax and parsing from semantics and evaluation. [Familiarity]
200 2. Distinguish a language definition (what constructs mean) from a particular language implementation
201 (compiler vs. interpreter, run-time representation of data objects, etc.). [Familiarity]
202 3. Explain how programming language implementations typically organize memory into global data, text,
203 heap, and stack sections and how features such as recursion and memory management map to this memory
204 model. [Familiarity]
205 4. Reason about memory leaks, dangling-pointer dereferences, and the benefits and limitations of garbage
206 collection. [Usage]
207

208 **PL/Syntax Analysis**

209 *[Elective]*

210 **Topics:**

- 211 • Scanning (lexical analysis) using regular expressions
212 • Parsing strategies including top-down (e.g., recursive descent, Earley parsing, or LL) and bottom-up (e.g.,
213 backtracking or LR) techniques; role of context-free grammars
214 • Generating scanners and parsers from declarative specifications
215

216 **Learning outcomes:**

- 217 1. Use formal grammars to specify the syntax of languages. [Usage]
218 2. Use declarative tools to generate parsers and scanners. [Usage]
219 3. Identify key issues in syntax definitions: ambiguity, associativity, precedence. [Familiarity]
220

221 **PL/Compiler Semantic Analysis**

222 *[Elective]*

223 **Topics:**

- 224 • High-level program representations such as abstract syntax trees
225 • Scope and binding resolution
226 • Type checking
227 • Declarative specifications such as attribute grammars
228

229 **Learning outcomes:**

- 230 1. Implement context-sensitive, source-level static analyses such as type-checkers or resolving identifiers to
231 identify their binding occurrences. [Usage]
232

233

234 **PL/Code Generation**

235 *[Elective]*

236 *Topics:*

- 237 • Instruction selection
- 238 • Procedure calls and method dispatching
- 239 • Register allocation
- 240 • Separate compilation; linking
- 241 • Instruction scheduling
- 242 • Peephole optimization
- 243

244 *Learning outcomes:*

- 245 1. Identify all essential steps for automatically converting source code into assembly or other low-level
246 languages. [Familiarity]
- 247 2. Generate the low-level code for calling functions/methods in modern languages. [Usage]
- 248 3. Discuss opportunities for optimization introduced by naive translation and approaches for achieving
249 optimization. [Familiarity]
- 250

251 **PL/Runtime Systems**

252 *[Elective]*

253 *Topics:*

- 254 • Target-platform characteristics such as registers, instructions, bytecodes
- 255 • Dynamic memory management approaches and techniques: malloc/free, garbage collection (mark-sweep,
256 copying, reference counting), regions (also known as arenas or zones)
- 257 • Data layout for objects and activation records
- 258 • Just-in-time compilation and dynamic recompilation
- 259 • Other features such as class loading, threads, security, etc.
- 260

261 *Learning outcomes:*

- 262 1. Compare the benefits of different memory-management schemes, using concepts such as fragmentation,
263 locality, and memory overhead. [Familiarity]
- 264 2. Discuss benefits and limitations of automatic memory management. [Familiarity]
- 265 3. Identify the services provided by modern language run-time systems. [Familiarity]
- 266 4. Discuss advantages, disadvantages, and difficulties of dynamic recompilation. [Familiarity]
- 267

268

269 **PL/Static Analysis**

270 *[Elective]*

271 *Topics:*

- 272 • Relevant program representations, such as basic blocks, control-flow graphs, def-use chains, static single
273 assignment, etc.
- 274 • Undecidability and consequences for program analysis
- 275 • Flow-insensitive analyses, such as type-checking and scalable pointer and alias analyses
- 276 • Flow-sensitive analyses, such as forward and backward dataflow analyses
- 277 • Path-sensitive analyses, such as software model checking
- 278 • Tools and frameworks for defining analyses
- 279 • Role of static analysis in program optimization
- 280 • Role of static analysis in (partial) verification and bug-finding
- 281

282 *Learning outcomes:*

- 283 1. Define useful static analyses in terms of a conceptual framework such as dataflow analysis. [Usage]
- 284 2. Communicate why an analysis is correct (sound and terminating). [Usage]
- 285 3. Explain why non-trivial sound static analyses must be approximate. [Familiarity]
- 286 4. Distinguish “may” and “must” analyses. [Familiarity]
- 287 5. Explain why potential aliasing limits sound program analysis and how alias analysis can help. [Familiarity]
- 288 6. Use the results of a static analysis for program optimization and/or partial program correctness. [Usage]
- 289

290 **PL/Advanced Programming Constructs**

291 *[Elective]*

292 *Topics:*

- 293 • Lazy evaluation and infinite streams
- 294 • Control Abstractions: Exception Handling, Continuations, Monads
- 295 • Object-oriented abstractions: Multiple inheritance, Mixins, Traits, Multimethods
- 296 • Metaprogramming: Macros, Generative programming, Model-based development
- 297 • Module systems
- 298 • String manipulation via pattern-matching (regular expressions)
- 299 • Dynamic code evaluation (“eval”)
- 300 • Language support for checking assertions, invariants, and pre/post-conditions
- 301

302 *Learning outcomes:*

- 303 1. Use various advanced programming constructs and idioms correctly. [Usage]
- 304 2. Discuss how various advanced programming constructs aim to improve program structure, software
305 quality, and programmer productivity. [Familiarity]
- 306 3. Discuss how various advanced programming constructs interact with the definition and implementation of
307 other language features. [Familiarity]
- 308

309

310 **PL/Concurrency and Parallelism**

311 *[Elective]*

312 Support for concurrency is a fundamental programming-languages issue with rich material in
313 programming language design, language implementation, and language theory. Due to coverage
314 in other Knowledge Areas, this elective Knowledge Unit aims only to complement the material
315 included elsewhere in the body of knowledge. Courses on programming languages are an
316 excellent place to include a general treatment of concurrency including this other material.

317 (Cross-reference: PD-Parallel and Distributed Computing)

318 *Topics:*

- 319 • Constructs for thread-shared variables and shared-memory synchronization
- 320 • Actor models
- 321 • Futures
- 322 • Language support for data parallelism
- 323 • Models for passing messages between sequential processes
- 324 • Effect of memory-consistency models on language semantics and correct code generation

325

326 *Learning outcomes:*

- 327 1. Write correct concurrent programs using multiple programming models. [Usage]
- 328 2. Explain why programming languages do not guarantee sequential consistency in the presence of data races
329 and what programmers must do as a result. [Familiarity]

330

331 **PL/Type Systems**

332 *[Elective]*

333 *Topics:*

- 334 • Compositional type constructors, such as product types (for aggregates), sum types (for unions), function
335 types, quantified types, and recursive types
- 336 • Type checking
- 337 • Type safety as preservation plus progress
- 338 • Type inference
- 339 • Static overloading

340

341 *Learning outcomes:*

- 342 1. Define a type system precisely and compositionally. [Usage]
- 343 2. For various foundational type constructors, identify the values they describe and the invariants they
344 enforce. [Familiarity]
- 345 3. Precisely specify the invariants preserved by a sound type system. [Familiarity]

346

347

348 **PL/Formal Semantics**

349 *[Elective]*

350 *Topics:*

- 351 • Syntax vs. semantics
- 352 • Lambda Calculus
- 353 • Approaches to semantics: Operational, Denotational, Axiomatic
- 354 • Proofs by induction over language semantics
- 355 • Formal definitions and proofs for type systems
- 356 • Parametricity

357
358 *Learning outcomes:*

- 359 1. Give a formal semantics for a small language. [Usage]
- 360 2. Use induction to prove properties of all (or a well-defined subset of) programs in a language. [Usage]
- 361 3. Use language-based techniques to build a formal model of a software system. [Usage]

362

363 **PL/Language Pragmatics**

364 *[Elective]*

365 *Topics:*

- 366 • Principles of language design such as orthogonality
- 367 • Evaluation order, precedence, and associativity
- 368 • Eager vs. delayed evaluation
- 369 • Defining control and iteration constructs
- 370 • External calls and system libraries

371
372 *Learning outcomes:*

- 373 1. Discuss the role of concepts such as orthogonality and well-chosen defaults in language design.
374 [Familiarity]
- 375 2. Use crisp and objective criteria for evaluating language-design decisions. [Usage]

376

377 **PL/Logic Programming**

378 *[Elective]*

379 *Topics:*

- 380 • Clausal representation of data structures and algorithms
- 381 • Unification
- 382 • Backtracking and search

383
384 *Learning outcomes:*

- 385 1. Use a logic language to implement conventional algorithms. [Usage]
- 386 2. Use a logic language to implement algorithms employing implicit search using clauses and relations.
387 [Usage]

1 **Software Development Fundamentals (SDF)**

2 Fluency in the process of software development is a prerequisite to the study of most of
3 computer science. In order to use computers to solve problems effectively, students must be
4 competent at reading and writing programs in multiple programming languages. Beyond
5 programming skills, however, they must be able to design and analyze algorithms, select
6 appropriate paradigms, and utilize modern development and testing tools. This knowledge area
7 brings together those fundamental concepts and skills related to the software development
8 process. As such, it provides a foundation for other software-oriented knowledge areas, most
9 notably Programming Languages, Algorithms and Complexity, and Software Engineering.

10 It is important to note that this knowledge area is distinct from the old Programming
11 Fundamentals knowledge area from CC2001. Whereas that knowledge area focused exclusively
12 on the programming skills required in an introductory computer science course, this new
13 knowledge area is intended to fill a much broader purpose. It focuses on the entire software
14 development process, identifying those concepts and skills that should be mastered in the first
15 year of a computer science program. This includes the design and simple analysis of algorithms,
16 fundamental programming concepts and data structures, and basic software development
17 methods and tools. As a result of its broader purpose, the Software Development Fundamentals
18 knowledge area includes fundamental concepts and skills that could naturally be listed in other
19 software-oriented knowledge areas (e.g., programming constructs from Programming
20 Languages, simple algorithm analysis from Algorithms & Complexity, simple development
21 methodologies from Software Engineering). Likewise, each of these knowledge areas will
22 contain more advanced material that builds upon the fundamental concepts and skills listed here.

23 While broader in scope than the old Programming Fundamentals, this knowledge area still allows
24 for considerable flexibility in the design of first-year curricula. For example, the Fundamental
25 Programming Concepts unit identifies only those concepts that are common to all programming
26 paradigms. It is expected that an instructor would select one or more programming paradigms
27 (e.g., object-oriented programming, functional programming, scripting) to illustrate these
28 programming concepts, and would pull paradigm-specific content from the Programming
29 Languages knowledge area to fill out a course. Likewise, an instructor could choose to

30 emphasize formal analysis (e.g., Big-Oh, computability) or design methodologies (e.g., team
31 projects, software life cycle) early, thus integrating hours from the Programming Languages,
32 Algorithms and Complexity, and/or Software Engineering knowledge areas. Thus, the 43 hours
33 of material in this knowledge area will typically be augmented with core material from one or
34 more of these knowledge areas to form a complete and coherent first-year experience.

35 When considering the hours allocated to each knowledge unit, it should be noted that these hours
36 reflect the minimal amount of classroom coverage needed to introduce the material. Many
37 software development topics will reappear and be reinforced by later topics (e.g., applying
38 iteration constructs when processing lists). In addition, the mastery of concepts and skills from
39 this knowledge area requires a significant amount of software development experience outside of
40 class.

41

42 **SDF. Software Development Fundamentals (43 Core-Tier1 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|---|------------------|------------------|--------------------|
| SDF/Algorithms and Design | 11 | | N |
| SDF/Fundamental Programming Concepts | 10 | | N |
| SDF/Fundamental Data Structures | 12 | | N |
| SDF/Development Methods | 10 | | N |

43

44

45 **SDF/Algorithms and Design**

46 *[11 Core-Tier1 hours]*

47 This unit builds the foundation for core concepts in the Algorithms & Complexity knowledge
48 area, most notably in the Basic Analysis and Algorithmic Strategies units.

49 *Topics:*

- 50 • The concept and properties of algorithms
- 51 ○ Informal comparison of algorithm efficiency (e.g., operation counts)
- 52 • The role of algorithms in the problem-solving process
- 53 • Problem-solving strategies
- 54 ○ Iterative and recursive mathematical functions
- 55 ○ Iterative and recursive traversal of data structures
- 56 ○ Divide-and-conquer strategies
- 57 • Fundamental design concepts and principles
- 58 ○ Abstraction
- 59 ○ Program decomposition
- 60 ○ Encapsulation and information hiding
- 61 ○ Separation of behavior and implementation
- 62

63 *Learning Outcomes:*

- 64 1. Discuss the importance of algorithms in the problem-solving process. [Familiarity]
- 65 2. Discuss how a problem may be solved by multiple algorithms, each with different properties. [Familiarity]
- 66 3. Create algorithms for solving simple problems. [Usage]
- 67 4. Use a programming language to implement, test, and debug algorithms for solving simple problems.
68 [Usage]
- 69 5. Implement, test, and debug simple recursive functions and procedures. [Usage]
- 70 6. Determine whether a recursive or iterative solution is most appropriate for a problem. [Assessment]
- 71 7. Implement a divide-and-conquer algorithm for solving a problem. [Usage]
- 72 8. Apply the techniques of decomposition to break a program into smaller pieces. [Usage]
- 73 9. Identify the data components and behaviors of multiple abstract data types. [Usage]
- 74 10. Implement a coherent abstract data type, with loose coupling between components and behaviors. [Usage]
- 75 11. Identify the relative strengths and weaknesses among multiple designs or implementations for a problem.
76 [Assessment]
- 77

78 **SDF/Fundamental Programming Concepts**

79 *[10 Core-Tier1 hours]*

80 This unit builds the foundation for core concepts in the Programming Languages knowledge
81 area, most notably in the paradigm-specific units: Object-Oriented Programming, Functional
82 Programming, and Event-Driven & Reactive Programming.

83 *Topics:*

- 84 • Basic syntax and semantics of a higher-level language
- 85 • Variables and primitive data types (e.g., numbers, characters, Booleans)
- 86 • Expressions and assignments
- 87 • Simple I/O including file I/O
- 88 • Conditional and iterative control structures

- 89 • Functions and parameter passing
- 90 • The concept of recursion

91
92 **Learning Outcomes:**

- 93 1. Analyze and explain the behavior of simple programs involving the fundamental programming constructs
94 covered by this unit. [Assessment]
 - 95 2. Identify and describe uses of primitive data types. [Familiarity]
 - 96 3. Write programs that use primitive data types. [Usage]
 - 97 4. Modify and expand short programs that use standard conditional and iterative control structures and
98 functions. [Usage]
 - 99 5. Design, implement, test, and debug a program that uses each of the following fundamental programming
100 constructs: basic computation, simple I/O, standard conditional and iterative structures, the definition of
101 functions, and parameter passing. [Usage]
 - 102 6. Write a program that uses file I/O to provide persistence across multiple executions. [Usage]
 - 103 7. Choose appropriate conditional and iteration constructs for a given programming task. [Assessment]
 - 104 8. Describe the concept of recursion and give examples of its use. [Familiarity]
 - 105 9. Identify the base case and the general case of a recursively-defined problem. [Assessment]
- 106

107 **SDF/Fundamental Data Structures**

108 *[12 Core-Tier1 hours]*

109 This unit builds the foundation for core concepts in the Algorithms & Complexity knowledge
110 area, most notably in the Fundamental Data Structures & Algorithms and Basic Computability &
111 Complexity units.

112 **Topics:**

- 113 • Arrays
 - 114 • Records/structs (heterogeneous aggregates)
 - 115 • Strings and string processing
 - 116 • Abstract data types and their implementation
 - 117 ○ Stacks
 - 118 ○ Queues
 - 119 ○ Priority queues
 - 120 ○ Sets
 - 121 ○ Maps
 - 122 • References and aliasing
 - 123 • Linked lists
 - 124 • Strategies for choosing the appropriate data structure
- 125

126 **Learning Outcomes:**

- 127 1. Discuss the appropriate use of built-in data structures. [Familiarity]
 - 128 2. Describe common applications for each data structure in the topic list. [Familiarity]
 - 129 3. Write programs that use each of the following data structures: arrays, strings, linked lists, stacks, queues,
130 sets, and maps. [Usage]
 - 131 4. Compare alternative implementations of data structures with respect to performance. [Assessment]
 - 132 5. Compare and contrast the costs and benefits of dynamic and static data structure implementations.
133 [Assessment]
 - 134 6. Choose the appropriate data structure for modeling a given problem. [Assessment]
- 135

136

137 **SDF/Development Methods**

138 *[10 Core-Tier1 hours]*

139 This unit builds the foundation for core concepts in the Software Engineering knowledge area,
140 most notably in the Software Processes, Software Design and Software Evolution units.

141 **Topics:**

- 142 • Program comprehension
- 143 • Program correctness
 - 144 ○ Types or errors (syntax, logic, run-time)
 - 145 ○ The concept of a specification
 - 146 ○ Defensive programming (e.g. secure coding, exception handling)
 - 147 ○ Code reviews
 - 148 ○ Testing fundamentals and test-case generation
 - 149 ○ Test-driven development
 - 150 ○ The role and the use of contracts, including pre- and post-conditions
 - 151 ○ Unit testing
- 152 • Simple refactoring
- 153 • Modern programming environments
 - 154 ○ Code search
 - 155 ○ Programming using library components and their APIs
- 156 • Debugging strategies
- 157 • Documentation and program style
- 158

159 **Learning Outcomes:**

- 160 1. Trace the execution of a variety of code segments and write summaries of their computations. [Assessment]
- 161 2. Explain why the creation of correct program components is important in the production of high-quality
162 software. [Familiarity]
- 163 3. Identify common coding errors that lead to insecure programs (e.g., buffer overflows, memory leaks,
164 malicious code) and apply strategies for avoiding such errors. [Usage]
- 165 4. Conduct a personal code review (focused on common coding errors) on a program component using a
166 provided checklist. [Usage]
- 167 5. Contribute to a small-team code review focused on component correctness. [Usage]
- 168 6. Describe how a contract can be used to specify the behavior of a program component. [Familiarity]
- 169 7. Create a unit test plan for a medium-size code segment. [Usage]
- 170 8. Refactor a program by identifying opportunities to apply procedural abstraction. [Usage]
- 171 9. Apply a variety of strategies to the testing and debugging of simple programs. [Usage]
- 172 10. Construct, execute and debug programs using a modern IDE and associated tools such as unit testing tools
173 and visual debuggers. [Usage]
- 174 11. Construct and debug programs using the standard libraries available with a chosen programming language.
175 [Usage]
- 176 12. Analyze the extent to which another programmer's code meets documentation and programming style
177 standards. [Assessment]
- 178 13. Apply consistent documentation and program style standards that contribute to the readability and
179 maintainability of software. [Usage]
- 180

1 **Software Engineering (SE)**

2 In every computing application domain, professionalism, quality, schedule, and cost are critical
3 to producing software systems. Because of this, the elements of software engineering are
4 applicable to developing software in all areas of computing. A wide variety of software
5 engineering practices have been developed and utilized since the need for a discipline of
6 software engineering was first recognized. Many trade-offs between these different practices
7 have also been identified. Practicing software engineers have to select and apply appropriate
8 techniques and practices to a given development effort to maximize value. To learn how to do
9 this, they study the elements of software engineering.

10 Software engineering is the discipline concerned with the application of theory, knowledge, and
11 practice to effectively and efficiently build reliable software systems that satisfy the requirements
12 of customers and users. This discipline is applicable to small, medium, and large-scale systems.
13 It encompasses all phases of the lifecycle of a software system, including requirements
14 elicitation, analysis and specification; design; construction; verification and validation;
15 deployment; and operation and maintenance. Whether small or large, following a traditional
16 disciplined development process, an agile approach, or some other method, software engineering
17 is concerned with the best way to build good software systems.

18 Software engineering uses engineering methods, processes, techniques, and measurements. It
19 benefits from the use of tools for managing software development; analyzing and modeling
20 software artifacts; assessing and controlling quality; and for ensuring a disciplined, controlled
21 approach to software evolution and reuse. The software engineering toolbox has evolved over the
22 years. For instance, the use of contracts, with requires and ensure clauses and class invariants, is
23 one good practice that has become more common. Software development, which can involve an
24 individual developer or a team or teams of developers, requires choosing the most appropriate
25 tools, methods, and approaches for a given development environment.

26

27 Students and instructors need to understand the impacts of specialization on software engineering
28 approaches. For example, specialized systems include:

- 29 • Real time systems
- 30 • Client-server systems
- 31 • Distributed systems
- 32 • Parallel systems
- 33 • Web-based systems
- 34 • High integrity systems
- 35 • Games
- 36 • Mobile computing
- 37 • Domain specific software (e.g., scientific computing or business applications)

38 Issues raised by each of these specialized systems demand specific treatments in each phase of
39 software engineering. Students must become aware of the differences between general software
40 engineering techniques and principles and the techniques and principles needed to address issues
41 specific to specialized systems.

42 An important effect of specialization is that different choices of material may need to be made
43 when teaching applications of software engineering, such as between different process models,
44 different approaches to modeling systems, or different choices of techniques for carrying out any
45 of the key activities. This is reflected in the assignment of core and elective material, with the
46 core topics and learning outcomes focusing on the principles underlying the various choices, and
47 the details of the various alternatives from which the choices have to be made being assigned to
48 the elective material.

49 Another division of the practices of software engineering is between those concerned with the
50 fundamental need to develop systems that implement correctly the functionality that is required
51 for them, and those concerned with other qualities for systems and the trade-offs needed to
52 balance these qualities. This division too is reflected in the assignment of core and elective
53 material, so that topics and learning outcomes concerned with the basic methods for developing

54 such system are assigned to the core, and those that are concerned with other qualities and trade-
55 offs between them are assigned to the elective material.

56 In general, students learn best at the application level much of the material defined in the SE KA
57 by participating in a project. Such projects should require students to work on a team to develop
58 a software system through as much of its lifecycle as is possible. Much of software engineering
59 is devoted to effective communication among team members and stakeholders. Utilizing project
60 teams, projects can be sufficiently challenging to require the use of effective software
61 engineering techniques and that students develop and practice their communication skills. While
62 organizing and running effective projects within the academic framework can be challenging, the
63 best way to learn to apply software engineering theory and knowledge is in the practical
64 environment of a project. The minimum hours specified for some knowledge units in this
65 document may appear insufficient to accomplish associated application-level learning outcomes.
66 It should be understood that these outcomes are to be achieved through project experience that
67 may even occur later in the curriculum than when the topics within the knowledge unit are
68 introduced.

69 Note: The SDF/Development Methods knowledge unit includes 10 Core-Tier1 hours that
70 constitute an introduction to certain aspects of software engineering. The knowledge units,
71 topics and core hour specifications in this document must be understood as assuming previous
72 exposure to the material described in SDF/Development Methods.

73

74 **SE. Software Engineering (6 Core-Tier1 hours; 21 Core-Tier2 hours)**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| SE/Software Processes | 2 | 1 | Y |
| SE/Software Project Management | | 2 | Y |
| SE/Tools and Environments | | 2 | N |
| SE/Requirements Engineering | 1 | 3 | Y |
| SE/Software Design | 3 | 5 | Y |
| SE/Software Construction | | 2 | Y |
| SE/Software Verification and Validation | | 3 | Y |
| SE/Software Evolution | | 2 | Y |
| SE/Formal Methods | | | Y |
| SE/Software Reliability | | 1 | Y |

75

76 **SE/Software Processes**

77 *[2 Core-Tier1 hours; 1 Core-Tier2 hour]*

78 *Topics:*

79 [Core-Tier1]

- 80
- Systems level considerations, i.e., the interaction of software with its intended environment
 - 81 • Introduction to software process models (e.g., waterfall, incremental, agile)
 - 82 ○ Phases of software life-cycles
 - 83 • Programming in the large vs. individual programming

84

85 [Core-Tier2]

- 86 • Applying software process models

87

88 [Elective]

- 89 • Software quality concepts
- 90 • Process improvement
- 91 • Software process capability maturity models
- 92 • Software process measurements

93

94

95 **Learning Outcomes:**

96 [Core-Tier1]

- 97 1. Describe how software can interact with and participate in various systems including information
98 management, embedded, process control, and communications systems. [Familiarity]
99 2. Describe the difference between principles of the waterfall model and models using iterations. [Familiarity]
100 3. Describe the different practices that are key components of various process model. [Familiarity]
101 4. Differentiate among the phases of software development. [Familiarity]
102 5. Describe how programming in the large differs from individual efforts with respect to understanding a large
103 code base, code reading, understanding builds, and understanding context of changes. [Familiarity]
104

105 [Core-Tier2]

- 106 6. Explain the concept of a software life cycle and provide an example, illustrating its phases including the
107 deliverables that are produced. [Familiarity]
108 7. Compare several common process models with respect to their value for development of particular classes
109 of software systems taking into account issues such as requirement stability, size, and non-functional
110 characteristics. [Usage]
111

112 [Elective]

- 113 8. Define software quality and describe the role of quality assurance activities in the software process.
114 [Familiarity]
115 9. Describe the intent and fundamental similarities among process improvement approaches. [Familiarity]
116 10. Compare several process improvement models such as CMM, CMMI, CQI, Plan-Do-Check-Act, or
117 ISO9000. [Familiarity]
118 11. Use a process improvement model such as PSP to assess a development effort and recommend approaches
119 to improvement. [Usage]
120 12. Explain the role of process maturity models in process improvement. [Familiarity]
121 13. Describe several process metrics for assessing and controlling a project. [Familiarity]
122 14. Use project metrics to describe the current state of a project. [Usage]
123

124 **SE/Software Project Management**

125 *[2 Core-Tier2 hours]*

126 **Topics:**

127 [Core-Tier2]

- 128 • Team participation
129 ○ Team processes including responsibilities for tasks, meeting structure, and work schedule
130 ○ Roles and responsibilities in a software team
131 ○ Team conflict resolution
132 ○ Risks associated with virtual teams (communication, perception, structure)
133 • Effort Estimation (at the personal level)
134 • Risk
135 ○ The role of risk in the life cycle
136 ○ Risk categories including security, safety, market, financial, technology, people, quality, structure
137 and process
138

139 [Elective]

- 140 • Team management
141 ○ Team organization and decision-making

- 142 ○ Role identification and assignment
- 143 ○ Individual and team performance assessment
- 144 ● Project management
 - 145 ○ Scheduling and tracking
 - 146 ○ Project management tools
 - 147 ○ Cost/benefit analysis
- 148 ● Software measurement and estimation techniques
- 149 ● Software quality assurance and the role of measurements
- 150 ● Risk
 - 151 ○ Risk identification and management
 - 152 ○ Risk analysis and evaluation
 - 153 ○ Risk tolerance (e.g., risk-adverse, risk-neutral, risk-seeking)
 - 154 ○ Risk planning
- 155 ● System-wide approach to risk including hazards associated with tools

156
157 **Learning Outcomes:**

158 [Core-Tier2]

- 159 1. Identify behaviors that contribute to the effective functioning of a team. [Familiarity]
- 160 2. Create and follow an agenda for a team meeting. [Usage]
- 161 3. Identify and justify necessary roles in a software development team. [Usage]
- 162 4. Understand the sources, hazards, and potential benefits of team conflict. [Usage]
- 163 5. Apply a conflict resolution strategy in a team setting. [Usage]
- 164 6. Use an *ad hoc* method to estimate software development effort (e.g., time) and compare to actual effort required. [Usage]
- 165 7. List several examples of software risks. [Familiarity]
- 166 8. Describe the impact of risk in a software development life cycle. [Familiarity]
- 167 9. Describe different categories of risk in software systems. [Familiarity]

168
169 [Elective]

- 171 10. Identify security risks for a software system. [Usage]
- 172 11. Demonstrate through involvement in a team project the central elements of team building and team management. [Usage]
- 173 12. Identify several possible team organizational structures and team decision-making processes. [Familiarity]
- 174 13. Create a team by identifying appropriate roles and assigning roles to team members. [Usage]
- 175 14. Assess and provide feedback to teams and individuals on their performance in a team setting. [Usage]
- 176 15. Prepare a project plan for a software project that includes estimates of size and effort, a schedule, resource allocation, configuration control, change management, and project risk identification and management. [Usage]
- 177 16. Track the progress of a project using appropriate project metrics. [Usage]
- 178 17. Compare simple software size and cost estimation techniques. [Usage]
- 179 18. Use a project management tool to assist in the assignment and tracking of tasks in a software development project. [Usage]
- 180 19. Describe the impact of risk tolerance on the software development process. [Assessment]
- 181 20. Identify risks and describe approaches to managing risk (avoidance, acceptance, transference, mitigation), and characterize the strengths and shortcomings of each. [Familiarity]
- 182 21. Explain how risk affects decisions in the software development process. [Usage]
- 183 22. Demonstrate a systematic approach to the task of identifying hazards and risks in a particular situation. [Usage]
- 184 23. Apply the basic principles of risk management in a variety of simple scenarios including a security situation. [Usage]
- 185 24. Conduct a cost/benefit analysis for a risk mitigation approach. [Usage]
- 186
- 187
- 188
- 189
- 190
- 191
- 192

193 25. Identify and analyze some of the risks for an entire system that arise from aspects other than the software.
194 [Usage]
195

196 **SE/Tools and Environments**

197 *[2 Core-Tier2 hours]*

198 *Topics:*

199 [Core-Tier2]

- 200 • Software configuration management and version control; release management
- 201 • Requirements analysis and design modeling tools
- 202 • Testing tools including static and dynamic analysis tools
- 203 • Programming environments that automate parts of program construction processes (e.g., automated builds)
- 204 ○ Continuous integration
- 205 • Tool integration concepts and mechanisms

206 *Learning Outcomes:*

207 [Core-Tier2]

- 208 1. Describe the difference between centralized and distributed software configuration management.
209 [Familiarity]
- 210 2. Identify configuration items and use a source code control tool in a small team-based project. [Usage]
- 211 3. Describe the issues that are important in selecting a set of tools for the development of a particular software
212 system, including tools for requirements tracking, design modeling, implementation, build automation, and
213 testing. [Familiarity]
- 214 4. Demonstrate the capability to use software tools in support of the development of a software product of
215 medium size. [Usage]
- 216
- 217
- 218

219 **SE/Requirements Engineering**

220 *[1 Core-Tier1 hour; 3 Core-Tier2 hours]*

221 *Topics:*

222 [Core-Tier1]

- 223 • Properties of requirements including consistency, validity, completeness, and feasibility
- 224 • Describing functional requirements using, for example, use cases or users stories

225

226 [Core-Tier2]

- 227 • Software requirements elicitation
- 228 • Non-functional requirements and their relationship to software quality
- 229 • Describing system data using, for example, class diagrams or entity-relationship diagrams
- 230 • Evaluation and use of requirements specifications

231

232 [Elective]

- 233 • Requirements analysis modeling techniques
- 234 • Acceptability of certainty / uncertainty considerations regarding software / system behavior

- 235 • Prototyping
- 236 • Basic concepts of formal requirements specification
- 237 • Requirements specification
- 238 • Requirements validation
- 239 • Requirements tracing

241 **Learning Outcomes:**

242 [Core-Tier1]

- 243 1. List the key components of a use case or similar description of some behavior that is required for a system
- 244 and discuss their role in the requirements engineering process. [Familiarity]
- 245 2. Interpret a given requirements model for a simple software system. [Familiarity]
- 246 3. Conduct a review of a set of software requirements to determine the quality of the requirements with
- 247 respect to the characteristics of good requirements. [Usage]

248 [Core-Tier2]

- 250 4. Describe the fundamental challenges of and common techniques used for requirements elicitation.
- 251 [Familiarity]
- 252 5. List the key components of a class diagram or similar description of the data that a system is required to
- 253 handle. [Familiarity]
- 254 6. Identify both functional and non-functional requirements in a given requirements specification for a
- 255 software system. [Usage]

256 [Elective]

- 257 7. Apply key elements and common methods for elicitation and analysis to produce a set of software
- 258 requirements for a medium-sized software system. [Usage]
- 259 8. Use a common, non-formal method to model and specify (in the form of a requirements specification
- 260 document) the requirements for a medium-size software system [Usage]
- 261 9. Translate into natural language a software requirements specification (e.g., a software component contract)
- 262 written in a formal specification language. [Usage]
- 263 10. Create a prototype of a software system to mitigate risk in requirements. [Usage]
- 264 11. Differentiate between forward and backward tracing and explain their roles in the requirements validation
- 265 process. [Familiarity]
- 266
- 267

268 **SE/Software Design**

269 *[3 Core-Tier1 hours; 5 Core-Tier2 hours]*

270 **Topics:**

271 [Core-Tier1]

- 272 • Overview of design paradigms
- 273 • System design principles: divide and conquer (architectural design and detailed design), separation of
- 274 concerns, information hiding, coupling and cohesion, re-use of standard structures.
- 275 • Appropriate models of software designs, including structure and behavior.

276 [Core-Tier2]

- 277
- 278 • Design Paradigms such as structured design (top-down functional decomposition), object-oriented analysis
- 279 and design, event driven design, component-level design, data-structured centered, aspect oriented,
- 280 function oriented, service oriented.
- 281 • Relationships between requirements and designs: transformation of models, design of contracts, invariants.

- 282 • Software architecture concepts and standard architectures (e.g. client-server, n-layer, transform centered,
283 pipes-and-filters, etc).
- 284 • Refactoring designs and the use of design patterns.
- 285 • The use of components in design: component selection, design, adaptation and assembly of components,
286 components and patterns, components and objects, (for example, build a GUI using a standard widget set).

287
288 [Elective]

- 289 • Internal design qualities, and models for them: efficiency and performance, redundancy and fault
290 tolerance, traceability of requirements.
- 291 • External design qualities, and models for them: functionality, reliability, performance and efficiency,
292 usability, maintainability, portability.
- 293 • Measurement and analysis of design quality.
- 294 • Tradeoffs between different aspects of quality.
- 295 • Application frameworks.
- 296 • Middleware: the object-oriented paradigm within middleware, object request brokers and marshalling,
297 transaction processing monitors, workflow systems.

299 ***Learning Outcomes:***

300 [Core-Tier1]

- 301 1. Articulate design principles including separation of concerns, information hiding, coupling and cohesion,
302 and encapsulation. [Familiarity]
- 303 2. Use a design paradigm to design a simple software system, and explain how system design principles have
304 been applied in this design. [Usage]
- 305 3. Construct models of the design of a simple software system that are appropriate for the paradigm used to
306 design it. [Usage]
- 307 4. For the design of a simple software system within the context of a single design paradigm, describe the
308 software architecture of that system. [Familiarity]
- 309 5. Within the context of a single design paradigm, describe one or more design patterns that could be
310 applicable to the design of a simple software system. [Familiarity]

311
312 [Core-Tier2]

- 313 6. For a simple system suitable for a given scenario, discuss and select an appropriate design paradigm.
314 [Usage]
- 315 7. Create appropriate models for the structure and behavior of software products from their requirements
316 specifications. [Usage]
- 317 8. Explain the relationships between the requirements for a software product and the designed structure and
318 behavior, in terms of the appropriate models and transformations of them. [Assessment]
- 319 9. Apply simple examples of patterns in a software design. [Usage]
- 320 10. Given a high-level design, identify the software architecture by differentiating among common software
321 architectures such as 3-tier, pipe-and-filter, and client-server. [Familiarity]
- 322 11. Investigate the impact of software architectures selection on the design of a simple system. [Assessment]
- 323 12. Select suitable components for use in the design of a software product. [Usage]
- 324 13. Explain how suitable components might need to be adapted for use in the design of a software product.
325 [Familiarity].
- 326 14. Design a contract for a typical small software component for use in a given system. [Usage]

327
328 [Elective]

- 329 15. Discuss and select appropriate software architecture for a simple system suitable for a given scenario.
330 [Usage]

- 331 16. Apply models for internal and external qualities in designing software components to achieve an acceptable
 332 tradeoff between conflicting quality aspects. [Usage]
 333 17. Analyze a software design from the perspective of a significant internal quality attribute. [Assessment]
 334 18. Analyze a software design from the perspective of a significant external quality attribute. [Assessment]
 335 19. Explain the role of objects in middleware systems and the relationship with components. [Familiarity]
 336 20. Apply component-oriented approaches to the design of a range of software, such as using components for
 337 concurrency and transactions, for reliable communication services, for database interaction including
 338 services for remote query and database management, or for secure communication and access. [Usage]
 339

340 SE/Software Construction

341 *[2 Core-Tier2 hours]*

342 **Topics:**

343 [Core-Tier2]

- 344 • Coding practices: techniques, idioms/patterns, mechanisms for building quality programs
 - 345 ○ Defensive coding practices
 - 346 ○ Secure coding practices
 - 347 ○ Using exception handling mechanisms to make programs more robust, fault-tolerant
- 348 • Coding standards
- 349 • Integration strategies
- 350 • Development context: “green field” vs. existing code base
 - 351 ○ Change impact analysis
 - 352 ○ Change actualization

353
 354 [Elective]

- 355 • Robust And Security Enhanced Programming
 - 356 ○ Defensive programming
 - 357 ○ Principles of secure design and coding:
 - 358 ○ Principle of least privilege
 - 359 ○ Principle of fail-safe defaults
 - 360 ○ Principle of psychological acceptability
- 361 • Potential security problems in programs
 - 362 ○ Buffer and other types of overflows
 - 363 ○ Race conditions
 - 364 ○ Improper initialization, including choice of privileges
 - 365 ○ Checking input
 - 366 ○ Assuming success and correctness
 - 367 ○ Validating assumptions
- 368 • Documenting security considerations in using a program
- 369

370 **Learning Outcomes:**

371 [Core-Tier2]

- 372 1. Describe techniques, coding idioms and mechanisms for implementing designs to achieve desired
 373 properties such as reliability, efficiency, and robustness. [Familiarity]
- 374 2. Build robust code using exception handling mechanisms. [Usage]
- 375 3. Describe secure coding and defensive coding practices. [Familiarity]
- 376 4. Select and use a defined coding standard in a small software project. [Usage]

- 377 5. Compare and contrast integration strategies including top-down, bottom-up, and sandwich integration.
 378 [Familiarity]
 379 6. Describe the process of analyzing and implementing changes to code base developed for a specific project.
 380 [Familiarity]
 381 7. Describe the process of analyzing and implementing changes to a large existing code base. [Familiarity]
 382
 383 [Elective]
 384 8. Rewrite a simple program to remove common vulnerabilities, such as buffer overflows, integer overflows
 385 and race conditions [Usage]
 386 9. State and apply the principles of least privilege and fail-safe defaults. [Familiarity]
 387 10. Write a simple library that performs some non-trivial task and will not terminate the calling program
 388 regardless of how it is called [Usage]
 389

390 **SE/Software Verification Validation**

391 *[3 Core-Tier2 hours]*

392 *Topics:*

393 [Core-Tier2]

- 394 • Verification and validation concepts
- 395 • Inspections, reviews, audits
- 396 • Testing types, including human computer interface, usability, reliability, security, conformance to
 397 specification
- 398 • Testing fundamentals
 - 399 ○ Unit, integration, validation, and system testing
 - 400 ○ Test plan creation and test case generation
 - 401 ○ Black-box and white-box testing techniques
- 402 • Defect tracking
- 403 • Testing parallel and distributed systems

404
 405 [Elective]

- 406 • Static approaches and dynamic approaches to verification
- 407 • Regression testing
- 408 • Test-driven development
- 409 • Validation planning; documentation for validation
- 410 • Object-oriented testing; systems testing
- 411 • Verification and validation of non-code artifacts (documentation, help files, training materials)
- 412 • Fault logging, fault tracking and technical support for such activities
- 413 • Fault estimation and testing termination including defect seeding

414
 415 *Learning Outcomes:*

416 [Core-Tier2]

- 417 1. Distinguish between program validation and verification. [Familiarity]
- 418 2. Describe the role that tools can play in the validation of software. [Familiarity]
- 419 3. Undertake, as part of a team activity, an inspection of a medium-size code segment. [Usage]
- 420 4. Describe and distinguish among the different types and levels of testing (unit, integration, systems, and
 421 acceptance). [Familiarity]

- 422 5. Describe techniques for identifying significant test cases for unit, integration, and system testing.
423 [Familiarity]
424 6. Use a defect tracking tool to manage software defects in a small software project. [Usage]
425 7. Describe the issues and approaches to testing parallel and distributed systems. [Familiarity]
426
427 [Elective]
- 428 8. Create, evaluate, and implement a test plan for a medium-size code segment. [Usage]
429 9. Compare static and dynamic approaches to verification. [Familiarity]
430 10. Discuss the issues involving the testing of object-oriented software. [Usage]
431 11. Describe techniques for the verification and validation of non-code artifacts. [Familiarity]
432 12. Describe approaches for fault estimation. [Familiarity]
433 13. Estimate the number of faults in a small software application based on fault density and fault seeding.
434 [Usage]
435 14. Conduct an inspection or review of software source code for a small or medium sized software project.
436 [Usage]
437

438 **SE/Software Evolution**

439 *[2 Core-Tier2 hour]*

440 *Topics:*

441 [Core-Tier2]

- 442 • Software development in the context of large, pre-existing code bases
 - 443 ○ Software change
 - 444 ○ Concerns and concern location
 - 445 ○ Refactoring
- 446 • Software evolution
- 447 • Characteristics of maintainable software
- 448 • Reengineering systems
- 449 • Software reuse
- 450

451 *Learning Outcomes:*

452 [Core-Tier2]

- 453 1. Identify the principal issues associated with software evolution and explain their impact on the software life
454 cycle. [Familiarity]
- 455 2. Estimate the impact of a change request to an existing product of medium size. [Usage]
- 456 3. Identify weaknesses in a given simple design, and removed them through refactoring. [Usage]
- 457 4. Discuss the challenges of evolving systems in a changing environment. [Familiarity]
- 458 5. Outline the process of regression testing and its role in release management. [Usage]
- 459 6. Discuss the advantages and disadvantages of software reuse. [Familiarity]
- 460

461 **SE/Formal Methods**

462 *[Elective]*

463 The topics listed below have a strong dependency on core material from the Discrete Structures
464 area, particularly knowledge units DS/Functions Relations And Sets, DS/Basic Logic and
465 DS/Proof Techniques.

466 **Topics:**

- 467 • Role of formal specification and analysis techniques in the software development cycle
- 468 • Program assertion languages and analysis approaches (including languages for writing and analyzing pre-
469 and post-conditions, such as OCL, JML)
- 470 • Formal approaches to software modeling and analysis
 - 471 ○ Model checkers
 - 472 ○ Model finders
- 473 • Tools in support of formal methods

474
475 **Learning Outcomes:**

- 476 1. Describe the role formal specification and analysis techniques can play in the development of complex
477 software and compare their use as validation and verification techniques with testing. [Familiarity]
- 478 2. Apply formal specification and analysis techniques to software designs and programs with low complexity.
479 [Usage]
- 480 3. Explain the potential benefits and drawbacks of using formal specification languages. [Familiarity]
- 481 4. Create and evaluate program assertions for a variety of behaviors ranging from simple through complex.
482 [Usage]
- 483 5. Using a common formal specification language, formulate the specification of a simple software system
484 and derive examples of test cases from the specification. [Usage]

485

486 **SE/Software Reliability**

487 **[1 Core-Tier2]**

488 **Topics:**

489 [Core-Tier2]

- 490 • Software reliability engineering concepts
- 491 • Software reliability, system reliability and failure behavior (cross-reference SF9/Reliability Through
492 Redundancy)
- 493 • Fault lifecycle concepts and techniques

494

495 [Elective]

- 496 • Software reliability models
- 497 • Software fault tolerance techniques and models
- 498 • Software reliability engineering practices
- 499 • Measurement-based analysis of software reliability

500

501 **Learning Outcomes:**

502 [Core-Tier2]

- 503 1. Explain the problems that exist in achieving very high levels of reliability. [Familiarity]
- 504 2. Describe how software reliability contributes to system reliability [Familiarity]

- 505 3. List approaches to minimizing faults that can be applied at each stage of the software lifecycle.
506 [Familiarity]
507
508 [Elective]
- 509 4. Compare the characteristics of three different reliability modeling approaches. [Familiarity]
510 5. Demonstrate the ability to apply multiple methods to develop reliability estimates for a software system.
511 [Usage]
512 6. Identify methods that will lead to the realization of a software architecture that achieves a specified
513 reliability level of reliability. [Usage]
514 7. Identify ways to apply redundancy to achieve fault tolerance for a medium-sized application. [Usage]

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 concepts include computational
8 paradigms, parallelism, cross-layer communications, state and state transition, resource
9 allocation and scheduling, and so on. This knowledge area is designed to present an integrative
10 view of these fundamental concepts in a unified albeit simplified fashion, providing a common
11 foundation for the different specialized mechanisms and policies appropriate to the particular
12 domain area.

13

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

| | Core-Tier 1 hours | Core-Tier 2 hours | Includes Electives |
|---|----------------------|----------------------|-----------------------|
| SF/Computational Paradigms | 3 | | N |
| SF/Cross-Layer Communications | 3 | | N |
| SF/State-State Transition-State Machines | 6 | | N |
| SF/Parallelism | 3 | | N |
| SF/Evaluation | 3 | | N |
| SF/Resource Allocation and Scheduling | | 2 | N |
| SF/Proximity | | 3 | N |
| SF/Virtualization and Isolation | | 2 | N |
| SF/Reliability through Redundancy | | 2 | N |
| SF/Quantitative Evaluation | | | Y |

15

16

17 **SF/Computational Paradigms**

18 *[3 Core-Tier 1 hours]*

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

23 **Topics:**

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

34 **Learning Outcomes:**

35 [Core-Tier1]

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

51 **SF/Cross-Layer Communications**

52 [Conceptual presentation here, practical experience in programming these abstractions in PD,
53 NC, OS]

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 • Application-Virtual Machine Interaction
- 59 • Reliability
- 60

61 **Learning Outcomes:**

62 [Core-Tier1]

- 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. This can be motivated by
65 real-world systems, like how a car works, or libraries. [Familiarity]
66 2. Recognize that hardware, VM, OS, application are additional layers of interpretation/processing.
67 [Familiarity]
68 3. Describe the mechanisms of how errors are detected, signaled back, and handled through the layers.
69 [Familiarity]
70 4. Construct a simple program using methods of layering, error detection and recovery, and reflection of error
71 status across layers. [Usage]
72 5. Find bugs in a layered program by using tools for program tracing, single stepping, and debugging. [Usage]
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 [Core-Tier1]

- 87 1. Describe computations as a system with a known set of configurations, and a byproduct of the computation
88 is to transition from one unique configuration (state) to another (state). [Familiarity]
89 2. Recognize the distinction between systems whose output is only a function of their input (Combinational)
90 and those with memory/history (Sequential). [Familiarity]
91 3. Describe a computer as a state machine that interprets machine instructions. [Familiarity]
92 4. Explain how a program or network protocol can also be expressed as a state machine, and that alternative
93 representations for the same computation can exist. [Familiarity]
94 5. Develop state machine descriptions for simple problem statement solutions (e.g., traffic light sequencing,
95 pattern recognizers). [Usage]
96 6. Derive time-series behavior of a state machine from its state machine representation. [Assessment]
97

98 **SF/Parallelism**

99 [3 Core-Tier1 hours]

100 [Cross-reference: PD/Parallelism Fundamentals]

101 **Topics:**

- 102 • Sequential vs. parallel processing

- 103 • Parallel programming (e.g., synchronization for producer-consumer for performance improvement) vs.
- 104 concurrent programming (e.g., mutual exclusion/atomic operations for reactive programs)
- 105 • Request parallelism (e.g., web services) vs. Task parallelism (map-reduce processing)
- 106 • Client-Server/Web Services, Thread (Fork-Join), Pipelining
- 107 • Multicore architectures and hardware support for synchronization

108 **Learning Outcomes:**

109 [Core-Tier1]

- 110 1. For a given program, distinguish between its sequential and parallel execution, and the performance
- 111 implications thereof. [Familiarity]
- 112 2. Demonstrate on an execution time line that parallelism events and operations can take place simultaneously
- 113 (i.e., at the same time). Explain how work can be performed in less elapsed time if this can be exploited.
- 114 [Familiarity]
- 115 3. Explain other uses of parallelism, such as for reliability/redundancy of execution. [Familiarity]
- 116 4. Define the differences between the concepts of Instruction Parallelism, Data Parallelism, Thread
- 117 Parallelism/Multitasking, Task/Request Parallelism. [Familiarity]
- 118 5. Write more than one parallel program (e.g., one simple parallel program in more than one parallel
- 119 programming paradigm; a simple parallel program that manages shared resources through synchronization
- 120 primitives; a simple parallel program that performs simultaneous operation on partitioned data through task
- 121 parallel (e.g., parallel search terms; a simple parallel program that performs step-by-step pipeline
- 122 processing through message passing). [Usage]
- 123 6. Use performance tools to measure speed-up achieved by parallel programs in terms of both problem size
- 124 and number of resources. [Assessment]
- 125
- 126

127 **SF/Evaluation**

128 *[3 Core-Tier 1 hours]*

129 [Cross-reference PD/Parallel Performance]

130 **Topics:**

- 131 • Choosing and understanding performance figures of merit (e.g., speed of execution, energy consumption,
- 132 bandwidth vs. latency, resource cost)
- 133 • Choosing and understanding workloads and representative benchmarks (e.g., SPEC, Dhrystone), and
- 134 methods of collecting and analyzing performance figures of merit
- 135 • CPI equation ($\text{Execution time} = \# \text{ of instructions} * \text{cycles/instruction} * \text{time/cycle}$) as tool for
- 136 understanding tradeoffs in the design of instruction sets, processor pipelines, and memory system
- 137 organizations.
- 138 • Amdahl's Law: the part of the computation that cannot be sped up limits the effect of the parts that can
- 139

140 **Learning Outcomes:**

141 [Core-Tier1]

- 142 1. Explain how the components of system architecture contribute to improving its performance. [Familiarity]
- 143 2. Describe Amdahl's law and discuss its limitations. [Familiarity]
- 144 3. Design and conduct a performance-oriented experiment, e.g., benchmark a parallel program with different
- 145 data sets in order to iteratively improve its performance. [Usage]
- 146 4. Use software tools to profile and measure program performance. [Assessment]
- 147

148 **SF/Resource Allocation and Scheduling**

149 *[2 Core-Tier 2 hours]*

150 *Topics:*

- 151 • Kinds of resources: processor share, memory, disk, net bandwidth
- 152 • Kinds of scheduling: first-come, priority
- 153 • Advantages of fair scheduling, preemptive scheduling

154

155 *Learning Outcomes:*

156 [Core-Tier2]

- 157 1. Define how finite computer resources (e.g., processor share, memory, storage and network bandwidth) are
- 158 managed by their careful allocation to existing entities. [Familiarity]
- 159 2. Describe the scheduling algorithms by which resources are allocated to competing entities, and the figures
- 160 of merit by which these algorithms are evaluated, such as fairness. [Familiarity]
- 161 3. Implement simple schedule algorithms. [Usage]
- 162 4. Measure figures of merit of alternative scheduler implementations. [Assessment]

163

164 **SF/Proximity**

165 *[3 Core-Tier 2 hours]*

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

167 *Topics:*

- 168 • Speed of light and computers (one foot per nanosecond vs. one GHz clocks)
- 169 • Latencies in computer systems: memory vs. disk latencies vs. across the network memory
- 170 • Caches, spatial and temporal locality, in processors and systems
- 171 • Caches, cache coherency in database, operating systems, distributed systems, and computer architecture
- 172 • Introduction into the processor memory hierarchy: registers and multi-level caches, and the formula for
- 173 average memory access time

174

175 *Learning Outcomes:*

176 [Core-Tier2]

- 177 1. Explain the importance of locality in determining performance. [Familiarity]
- 178 2. Describe why things that are close in space take less time to access. [Familiarity]
- 179 3. Calculate average memory access time and describe the tradeoffs in memory hierarchy performance in
- 180 terms of capacity, miss/hit rate, and access time. [Assessment]

181

182 **SF/Virtualization and Isolation**

183 *[2 Core-Tier 2 hours]*

184 *Topics:*

- 185 • Rationale for protection and predictable performance
- 186 • Levels of indirection, illustrated by virtual memory for managing physical memory resources
- 187 • Methods for implementing virtual memory and virtual machines

188

189 **Learning Outcomes:**

190 [Core-Tier2]

- 191 1. Explain why it is important to isolate and protect the execution of individual programs and environments
192 that share common underlying resources, including the processor, memory, storage, and network access.
193 [Familiarity]
- 194 2. Describe how the concept of indirection can create the illusion of a dedicated machine and its resources
195 even when physically shared among multiple programs and environments. [Familiarity]
- 196 3. Measure the performance of two application instances running on separate virtual machines, and determine
197 the effect of performance isolation. [Assessment]

198

199 **SF/Reliability through Redundancy**

200 [*2 Core-Tier 2 hours*]

201 **Topics:**

- 202 • Distinction between bugs and faults
- 203 • How errors increase the longer the distance between the communicating entities; the end-to-end principle
204 as it applies to systems and networks
- 205 • Redundancy through check and retry
- 206 • Redundancy through redundant encoding (error correcting codes, CRC, FEC)
- 207 • Duplication/mirroring/replicas
- 208 • Other approaches to fault tolerance and availability

209

210 **Learning Outcomes:**

211 [Core-Tier2]

- 212 1. Explain the distinction between program errors, system errors, and hardware faults (e.g., bad memory) and
213 exceptions (e.g., attempt to divide by zero). [Familiarity]
- 214 2. Articulate the distinction between detecting, handling, and recovering from faults, and the methods for their
215 implementation. [Familiarity]
- 216 3. Describe the role of error correcting codes in providing error checking and correction techniques in
217 memories, storage, and networks. [Familiarity]
- 218 4. Apply simple algorithms for exploiting redundant information for the purposes of data correction. [Usage]
- 219 5. Compare different error detection and correction methods for their data overhead, implementation
220 complexity, and relative execution time for encoding, detecting, and correcting errors. [Assessment]

221

222 **SF/Quantitative Evaluation**

223 [*Elective*]

224 **Topics:**

- 225 • Analytical tools to guide quantitative evaluation
- 226 • Order of magnitude analysis (Big O notation)
- 227 • Analysis of slow and fast paths of a system
- 228 • Events on their effect on performance (e.g., instruction stalls, cache misses, page faults)
- 229 • Understanding layered systems, workloads, and platforms, their implications for performance, and the
230 challenges they represent for evaluation
- 231 • Microbenchmarking pitfalls

232

233

234 ***Learning Outcomes:***

235 [Elective]

- 236 1. Explain the circumstances in which a given figure of system performance metric is useful. [Familiarity]
- 237 2. Explain the inadequacies of benchmarks as a measure of system performance. [Familiarity]
- 238 3. Use limit studies or simple calculations to produce order-of-magnitude estimates for a given performance
- 239 metric in a given context. [Usage]
- 240 4. Conduct a performance experiment on a layered system to determine the effect of a system parameter on
- 241 figure of system performance. [Assessment]

1 **Social Issues and Professional Practice (SP)**

2 While technical issues are central to the computing curriculum, they do not constitute a complete
3 educational program in the field. Students must also be exposed to the larger societal context of
4 computing to develop an understanding of the relevant social, ethical, legal and professional
5 issues. This need to incorporate the study of these non-technical issues into the ACM curriculum
6 was formally recognized in 1991, as can be seen from the following excerpt [Tucker91]:

7 *Undergraduates also need to understand the basic cultural, social, legal, and ethical*
8 *issues inherent in the discipline of computing. They should understand where the*
9 *discipline has been, where it is, and where it is heading. They should also understand*
10 *their individual roles in this process, as well as appreciate the philosophical questions,*
11 *technical problems, and aesthetic values that play an important part in the development*
12 *of the discipline.*

13 *Students also need to develop the ability to ask serious questions about the social*
14 *impact of computing and to evaluate proposed answers to those questions. Future*
15 *practitioners must be able to anticipate the impact of introducing a given product into a*
16 *given environment. Will that product enhance or degrade the quality of life? What will*
17 *the impact be upon individuals, groups, and institutions?*

18 *Finally, students need to be aware of the basic legal rights of software and hardware*
19 *vendors and users, and they also need to appreciate the ethical values that are the basis*
20 *for those rights. Future practitioners must understand the responsibility that they will*
21 *bear, and the possible consequences of failure. They must understand their own*
22 *limitations as well as the limitations of their tools. All practitioners must make a long-*
23 *term commitment to remaining current in their chosen specialties and in the discipline*
24 *of computing as a whole.*

25 As technological advances continue to significantly impact the way we live and work, the critical
26 importance of these social and professional issues continues to increase; new computer-based
27 products and venues pose ever more challenging problems each year. It is our students who
28 must enter the workforce and academia with intentional regard for the identification and
29 resolution of these problems.

30 Computer science educators may opt to deliver this core and elective material in stand-alone
31 courses, integrated into traditional technical and theoretical courses, or as special units in
32 capstone and professional practice courses. The material in this familiarity area is best covered
33 through a combination of one required course along with short modules in other courses. On the
34 one hand, some units listed as core tier-1—in particular, Social Context, Analytical Tools,
35 Professional Ethics, and Intellectual Property—do not readily lend themselves to being covered
36 in other traditional courses. Without a standalone course, it is difficult to cover these topics
37 appropriately. On the other hand, if ethical and social considerations are covered only in the
38 standalone course and not “in context,” it will reinforce the false notion that technical processes
39 are void of these other relevant issues. Because of this broad relevance, it is important that
40 several traditional courses include modules that analyze the ethical, social and professional
41 considerations in the context of the technical subject matter of the course. Courses in areas such
42 as software engineering, databases, computer networks, computer security, and introduction to
43 computing provide obvious context for analysis of ethical issues. However, an ethics-related
44 module could be developed for almost any course in the curriculum. It would be explicitly
45 against the spirit of the recommendations to have only a standalone course. Running through all
46 of the issues in this area is the need to speak to the computer practitioner’s responsibility to
47 proactively address these issues by both moral and technical actions. The ethical issues discussed
48 in any class should be directly related to and arise naturally from the subject matter of that class.
49 Examples include a discussion in the database course of data aggregation or data mining, or a
50 discussion in the software engineering course of the potential conflicts between obligations to the
51 customer and obligations to the user and others affected by their work. Programming
52 assignments built around applications such as controlling the movement of a laser during eye
53 surgery can help to address the professional, ethical and social impacts of computing. Computing
54 faculty who are unfamiliar with the content and/or pedagogy of applied ethics are urged to take
55 advantage of the considerable resources from ACM, IEEE-CS, SIGCAS (special interest group
56 on computers and society), and other organizations.

57 It should be noted that the application of ethical analysis underlies every subsection of this Social
58 and Professional knowledge area in computing. The ACM Code of Ethics and Professional
59 Conduct - www.acm.org/about/code-of-ethics - provide guidelines that serve as the basis for the

60 conduct of our professional work. The General Moral Imperatives provide an understanding of
 61 our commitment to personal responsibility, professional conduct, and our leadership roles.

62

63 **SP. Social Issues and Professional Practice [11 Core-Tier1 hours, 5 Core-Tier2**
 64 **hours]**

| | Core-Tier1 hours | Core-Tier2 hours | Includes Electives |
|--|------------------|------------------|--------------------|
| SP/Social Context | 1 | 2 | N |
| SP/Analytical Tools | 2 | | N |
| SP/Professional Ethics | 2 | 2 | N |
| SP/Intellectual Property | 2 | | Y |
| SP/Privacy and Civil Liberties | 2 | | Y |
| SP/Professional Communication | 1 | | Y |
| SP/Sustainability | 1 | 1 | Y |
| SP/History | | | Y |
| SP/Economies of Computing | | | Y |
| SP/Security Policies, Laws and Computer Crimes | | | Y |

65

66 **SP/Social Context**

67 *[1 Core-Tier1 hour, 2 Core-Tier2 hours]*

68 Computers and the Internet, perhaps more than any other technology, have transformed society
 69 over the past 50 years, with dramatic increases in human productivity; an explosion of options
 70 for news, entertainment, and communication; and fundamental breakthroughs in almost every
 71 branch of science and engineering.

72 *Topics:*

73 [Core-Tier1]

- 74 • Social implications of computing in a networked world (cross-reference HCI/Foundations/social models;
 75 IAS/Fundamental Concepts/social issues)
- 76 • Impact of social media on individualism, collectivism and culture.

77 [Core-Tier2]

- 79 • Growth and control of the Internet (cross-reference NC/Introduction/organization of the Internet)
- 80 • Often referred to as the digital divide, differences in access to digital technology resources and its resulting
 81 ramifications for gender, class, ethnicity, geography, and/or underdeveloped countries.

- 82 • Accessibility issues, including legal requirements
- 83 • Context-aware computing (cross-reference HC/Design for non-mouse interfaces/ ubiquitous and context-
- 84 aware)
- 85

86 **Learning Outcomes:**

87 [Core-Tier1]

- 88 1. Describe positive and negative ways in which computer technology (networks, mobile computing, cloud
- 89 computing) alters modes of social interaction at the personal level. [Familiarity]
- 90 2. Identify developers' assumptions and values embedded in hardware and software design, especially as they
- 91 pertain to usability for diverse populations including under-represented populations and the disabled.
- 92 [Familiarity]
- 93 3. Interpret the social context of a given design and its implementation. [Familiarity]
- 94 4. Evaluate the efficacy of a given design and implementation using empirical data. [Assessment]
- 95 5. Investigate the implications of social media on individualism versus collectivism and culture. [Usage]
- 96

97 [Core-Tier2]

- 98 6. Discuss how Internet access serves as a liberating force for people living under oppressive forms of
- 99 government; explain how limits on Internet access are used as tools of political and social repression.
- 100 [Familiarity]
- 101 7. Analyze the pros and cons of reliance on computing in the implementation of democracy (e.g. delivery of
- 102 social services, electronic voting). [Assessment]
- 103 8. Describe the impact of the under-representation of diverse populations in the computing profession (e.g.,
- 104 industry culture, product diversity). [Familiarity]
- 105 9. Investigate the implications of context awareness in ubiquitous computing systems. [Usage]
- 106

107 **SP/Analytical Tools**

108 *[2 Core-Tier1 hours]*

109 Ethical theories and principles are the foundations of ethical analysis because they are the
 110 viewpoints from which guidance can be obtained along the pathway to a decision. Each theory
 111 emphasizes different points such as predicting the outcome and following one's duties to others
 112 in order to reach an ethically guided decision. However, in order for an ethical theory to be
 113 useful, the theory must be directed towards a common set of goals. Ethical principles are the
 114 common goals that each theory tries to achieve in order to be successful. These goals include
 115 beneficence, least harm, respect for autonomy and justice.

116 **Topics:**

117 [Core-Tier1]

- 118 • Ethical argumentation
- 119 • Ethical theories and decision-making
- 120 • Moral assumptions and values
- 121
- 122
- 123

124 **Learning Outcomes:**

125 [Core-Tier1]

- 126 1. Evaluate stakeholder positions in a given situation. [Assessment]
- 127 2. Analyze basic logical fallacies in an argument. [Assessment]
- 128 3. Analyze an argument to identify premises and conclusion. [Assessment]
- 129 4. Illustrate the use of example and analogy in ethical argument. [Usage]
- 130 5. Evaluate ethical/social tradeoffs in technical decisions. [Assessment]

131

132 **SP/Professional Ethics**

133 *[2 Core-Tier1 hours, 2 Core-Tier2 hours]*

134 Computer ethics is a branch of practical philosophy which deals with how computing
135 professionals should make decisions regarding professional and social conduct. There are three
136 primary influences: 1) The individual's own personal code, 2) Any informal code of ethical
137 behavior existing in the work place, and 3) Exposure to formal codes of ethics.

138 **Topics:**

139 [Core-Tier1]

- 140 • Community values and the laws by which we live
- 141 • The nature of professionalism including care, attention and discipline, fiduciary responsibility, and
142 mentoring
- 143 • Keeping up-to-date as a professional in terms of familiarity, tools, skills, legal and professional framework
144 as well as the ability to self-assess and computer fluency
- 145 • Professional certification, codes of ethics, conduct, and practice, such as the ACM/IEEE-CS, SE, AITP,
146 IFIP and international societies (cross-reference IAS/Fundamental Concepts/ethical issues)
- 147 • Accountability, responsibility and liability (e.g. software correctness, reliability and safety, as well as
148 ethical confidentiality of cybersecurity professionals)

149

150 [Core-Tier2]

- 151 • The role of the professional in public policy
- 152 • Maintaining awareness of consequences
- 153 • Ethical dissent and whistle-blowing
- 154 • Dealing with harassment and discrimination
- 155 • Forms of professional credentialing
- 156 • Acceptable use policies for computing in the workplace
- 157 • Ergonomics and healthy computing environments
- 158 • Time to market and cost considerations versus quality professional standards

159

160 **Learning Outcomes:**

161 [Core-Tier1]

- 162 1. Identify ethical issues that arise in software development and determine how to address them technically
163 and ethically. [Familiarity]
- 164 2. Recognize the ethical responsibility of ensuring software correctness, reliability and safety. [Familiarity]
- 165 3. Describe the mechanisms that typically exist for a professional to keep up-to-date. [Familiarity]

- 166 4. Describe the strengths and weaknesses of relevant professional codes as expressions of professionalism and
 167 guides to decision-making. [Familiarity]
 168 5. Analyze a global computing issue, observing the role of professionals and government officials in
 169 managing this problem. [Assessment]
 170 6. Evaluate the professional codes of ethics from the ACM, the IEEE Computer Society, and other
 171 organizations. [Assessment]
 172

173 [Core-Tier2]

- 174 7. Describe ways in which professionals may contribute to public policy. [Familiarity]
 175 8. Describe the consequences of inappropriate professional behavior. [Familiarity]
 176 9. Identify progressive stages in a whistle-blowing incident. [Familiarity]
 177 10. Investigate forms of harassment and discrimination and avenues of assistance [Usage]
 178 11. Examine various forms of professional credentialing [Usage]
 179 12. Identify the social implications of ergonomic devices and the workplace environment to people's health.
 180 [Familiarity]
 181 13. Develop a computer usage/acceptable use policy with enforcement measures. [Assessment]
 182 14. Describe issues associated with industries' push to focus on time to market versus enforcing quality
 183 professional standards [Familiarity]
 184
 185

186 SP/Intellectual Property

187 *[2 Core-Tier1 hours]*

188 Intellectual property is the foundation of the software industry. The term refers to a range of
 189 intangible rights of ownership in an asset such as a software program. Each intellectual property
 190 "right" is itself an asset. The law provides different methods for protecting these rights of
 191 ownership based on their type. There are essentially four types of intellectual property rights
 192 relevant to software: patents, copyrights, trade secrets and trademarks. Each affords a different
 193 type of legal protection.

194 *Topics:*

195 [Core-Tier1]

- 196 • Philosophical foundations of intellectual property
 197 • Intellectual property rights (cross-reference IM/Information Storage and Retrieval/intellectual property and
 198 protection)
 199 • Intangible digital intellectual property (IDIP)
 200 • Legal foundations for intellectual property protection
 201 • Digital rights management
 202 • Copyrights, patents, trade secrets, trademarks
 203 • Plagiarism
 204

205 [Elective]

- 206 • Foundations of the open source movement
 207 • Software piracy
 208

209

210 **Learning Outcomes:**

211 [Core-Tier1]

- 212 1. Discuss the philosophical bases of intellectual property. [Familiarity]
- 213 2. Discuss the rationale for the legal protection of intellectual property. [Familiarity]
- 214 3. Describe legislation aimed at digital copyright infringements. [Familiarity]
- 215 4. Critique legislation aimed at digital copyright infringements [Assessment]
- 216 5. Identify contemporary examples of intangible digital intellectual property [Familiarity]
- 217 6. Justify uses of copyrighted materials. [Assessment]
- 218 7. Evaluate the ethical issues inherent in various plagiarism detection mechanisms. [Assessment]
- 219 8. Interpret the intent and implementation of software licensing. [Familiarity]
- 220 9. Discuss the issues involved in securing software patents. [Familiarity]
- 221 10. Characterize and contrast the concepts of copyright, patenting and trademarks. [Assessment]

222
223 [Elective]

- 224 11. Identify the goals of the open source movement. [Familiarity]
- 225 12. Identify the global nature of software piracy. [Familiarity]

226

227 **SP/ Privacy and Civil Liberties**

228 *[2 Core-Tier1 hours]*

229 Electronic information sharing highlights the need to balance privacy protections with
230 information access. The ease of digital access to many types of data makes privacy rights and
231 civil liberties more complex, differing among the variety of cultures worldwide.

232 **Topics:**

233 [Core-Tier1]

- 234 • Philosophical foundations of privacy rights (cross-reference IS/Fundamental Issues/philosophical issues)
- 235 • Legal foundations of privacy protection
- 236 • Privacy implications of widespread data collection for transactional databases, data warehouses,
237 surveillance systems, and cloud computing (cross reference IM/Database Systems/data independence;
238 IM/Data Mining/data cleaning)
- 239 • Ramifications of differential privacy
- 240 • Technology-based solutions for privacy protection (cross-reference IAS/Fundamental Concepts/data
241 protection laws)

242
243 [Elective]

- 244 • Privacy legislation in areas of practice
- 245 • Civil liberties and cultural differences
- 246 • Freedom of expression and its limitations

247
248 **Learning Outcomes:**

249 [Core-Tier1]

- 250 1. Discuss the philosophical basis for the legal protection of personal privacy. [Familiarity]
- 251 2. Evaluate solutions to privacy threats in transactional databases and data warehouses. [Assessment]
- 252 3. Recognize the fundamental role of data collection in the implementation of pervasive surveillance systems
253 (e.g., RFID, face recognition, toll collection, mobile computing). [Familiarity]

- 254 4. Recognize the ramifications of differential privacy. [Familiarity]
255 5. Investigate the impact of technological solutions to privacy problems. [Usage]

256
257 [Elective]

- 258 6. Critique the intent, potential value and implementation of various forms of privacy legislation.
259 [Assessment]

- 260 7. Identify strategies to enable appropriate freedom of expression. [Familiarity]

261

262 **SP/ Professional Communication**

263 *[1 Core-Tier1 hour]*

264 Professional communication conveys technical information to various audiences who may have
265 very different goals and needs for that information. Effective professional communication of
266 technical information is rarely an inherited gift, but rather needs to be taught in context
267 throughout the undergraduate curriculum.

268 **Topics:**

269 [Core-Tier1]

- 270 • Reading, understanding and summarizing technical material, including source code and documentation
271 • Writing effective technical documentation and materials
272 • Dynamics of oral, written, and electronic team and group communication (cross-reference
273 HCI/Collaboration and Communication/group communication; SE/Project Management/team participation)
274 • Communicating professionally with stakeholders
275 • Utilizing collaboration tools (cross-reference HCI/ Collaboration and Communication/online communities;
276 IS/Agents/collaborative agents)

277
278 [Elective]

- 279 • Dealing with cross-cultural environments (cross-reference HCI/User-Centered Design and Testing/cross-
280 cultural evaluation)
281 • Tradeoffs of competing risks in software projects, such as technology, structure/process, quality, people,
282 market and financial

283

284 **Learning Outcomes:**

285 [Core-Tier1]

- 286 1. Write clear, concise, and accurate technical documents following well-defined standards for format and for
287 including appropriate tables, figures, and references. [Usage]
288 2. Evaluate written technical documentation to detect problems of various kinds. [Assessment]
289 3. Develop and deliver a good quality formal presentation. [Assessment]
290 4. Plan interactions (e.g. virtual, face-to-face, shared documents) with others in which they are able to get
291 their point across, and are also able to listen carefully and appreciate the points of others, even when they
292 disagree, and are able to convey to others that they have heard. [Usage]
293 5. Describe the strengths and weaknesses of various forms of communication (e.g. virtual, face-to-face, shared
294 documents) [Familiarity]
295 6. Examine appropriate measures used to communicate with stakeholders involved in a project. [Usage]
296 7. Compare and contrast various collaboration tools. [Assessment]

297

298

- 299 [Elective]
- 300 8. Discuss ways to influence performance and results in cross-cultural teams. [Familiarity]
- 301 9. Examine the tradeoffs and common sources of risk in software projects regarding technology,
- 302 structure/process, quality, people, market and financial. [Usage]
- 303 10. Evaluate personal strengths and weaknesses to work remotely as part of a multinational team. [Assessment]
- 304

305 **SP/ Sustainability**

306 *[1 Core-Tier1 hour, 1 Core-Tier2 hour]*

307 Sustainability is characterized by the United Nations as “development that meets the needs of the

308 present without compromising the ability of future generations to meet their own needs.”

309 Sustainability was first introduced in the CS2008 curricular guidelines. Topics in this emerging

310 area can be naturally integrated into other familiarity areas and units, such as human-computer

311 interaction and software evolution.

312 **Topics:**

313 [Core-Tier1]

- 314 • Being a sustainable practitioner by taking into consideration cultural and environmental impacts of
- 315 implementation decisions (e.g. organizational policies, economic viability, and resource consumption).
- 316 • Explore global social and environmental impacts of computer use and disposal (e-waste)
- 317

318 [Core-Tier2]

- 319 • Environmental impacts of design choices in specific areas such as algorithms, operating systems, networks,
- 320 databases, programming languages, or human-computer interaction (cross-reference SE/Software
- 321 Evaluation/software evolution)
- 322

323 [Elective]

- 324 • Guidelines for sustainable design standards
- 325 • Systemic effects of complex computer-mediated phenomena (e.g. telecommuting or web shopping)
- 326 • Pervasive computing. Information processing that has been integrated into everyday objects and activities,
- 327 such as smart energy systems, social networking and feedback systems to promote sustainable behavior,
- 328 transportation, environmental monitoring, citizen science and activism.
- 329 • Conduct research on applications of computing to environmental issues, such as energy, pollution, resource
- 330 usage, recycling and reuse, food management, farming and others.
- 331 • How the sustainability of software systems are interdependent with social systems, including the
- 332 knowledge and skills of its users, organizational processes and policies, and its societal context (e.g. market
- 333 forces, government policies).
- 334

335 **Learning Outcomes:**

336 [Core-Tier1]

- 337 1. Identify ways to be a sustainable practitioner [Familiarity]
- 338 2. Illustrate global social and environmental impacts of computer use and disposal (e-waste) [Usage]
- 339
- 340

- 341 [Core-Tier2]
342 3. Describe the environmental impacts of design choices within the field of computing that relate to algorithm
343 design, operating system design, networking design, database design, etc. [Familiarity]
344 4. Investigate the social and environmental impacts of new system designs through projects. [Usage]
345
346 [Elective]
347 5. Identify guidelines for sustainable IT design or deployment [Familiarity]
348 6. List the sustainable effects of telecommuting or web shopping [Familiarity]
349 7. Investigate pervasive computing in areas such as smart energy systems, social networking, transportation,
350 agriculture, supply-chain systems, environmental monitoring and citizen activism. [Usage]
351 8. Develop applications of computing and assess through research areas pertaining to environmental issues
352 (e.g. energy, pollution, resource usage, recycling and reuse, food management, farming) [Assessment]
353

354 **SP/ History**

355 *[Elective]*

356 This history of computing is taught to provide a sense of how the rapid change in computing
357 impacts society on a global scale. It is often taught in context with foundational concepts, such as
358 system fundamentals and software developmental fundamentals.

359 *Topics:*

- 360 • Prehistory—the world before 1946
- 361 • History of computer hardware, software, networking (cross-reference AR/Digital logic and digital systems/
362 history of computer architecture)
- 363 • Pioneers of computing
- 364 • History of Internet
- 365

366 *Learning Outcomes:*

- 367 1. Identify significant continuing trends in the history of the computing field. [Familiarity]
- 368 2. Identify the contributions of several pioneers in the computing field. [Familiarity]
- 369 3. Discuss the historical context for several programming language paradigms. [Familiarity]
- 370 4. Compare daily life before and after the advent of personal computers and the Internet. [Assessment]
- 371

372 **SP/ Economies of Computing**

373 *[Elective]*

374 Economics of computing encompasses the metrics and best practices for personnel and financial
375 management surrounding computer information systems. Cost benefit analysis is covered in the
376 Information Assurance and Security Knowledge Area under Risk Management.

377 *Topics:*

- 378 • Monopolies and their economic implications
- 379 • Effect of skilled labor supply and demand on the quality of computing products
- 380 • Pricing strategies in the computing domain

- 381 • The phenomenon of outsourcing and off-shoring software development; impacts on employment and on
- 382 economics
- 383 • Consequences of globalization for the computer science profession
- 384 • Differences in access to computing resources and the possible effects thereof
- 385 • Costing out jobs with considerations on manufacturing, hardware, software, and engineering implications
- 386 • Cost estimates versus actual costs in relation to total costs
- 387 • Entrepreneurship: prospects and pitfalls
- 388 • Use of engineering economics in dealing with finances
- 389

390 ***Learning Outcomes:***

- 391 1. Summarize the rationale for antimonopoly efforts. [Familiarity]
- 392 2. Identify several ways in which the information technology industry is affected by shortages in the labor
- 393 supply. [Familiarity]
- 394 3. Identify the evolution of pricing strategies for computing goods and services. [Familiarity]
- 395 4. Discuss the benefits, the drawbacks and the implications of off-shoring and outsourcing. [Familiarity]
- 396 5. Investigate and defend ways to address limitations on access to computing. [Usage]
- 397

398 **SP/ Security Policies, Laws and Computer Crimes**

399 ***[Elective]***

400 While security policies, laws and computer crimes are important, it is essential they are viewed
 401 with the foundation of other Social and Professional knowledge units, such as Intellectual
 402 Property, Privacy and Civil Liberties, Social Context, and Professional Ethics. Computers and
 403 the Internet, perhaps more than any other technology, have transformed society over the past 50
 404 years. At the same time, they have contributed to unprecedented threats to privacy; whole new
 405 categories of crime and anti-social behavior; major disruptions to organizations; and the large-
 406 scale concentration of risk into information systems.

407 ***Topics:***

- 408 • Examples of computer crimes and legal redress for computer criminals (cross-reference IAS/Digital
- 409 Forensics/rules of evidence)
- 410 • Social engineering, identity theft and recovery (cross-reference HCI/Human Factors and Security/trust,
- 411 privacy and deception)
- 412 • Issues surrounding the misuse of access and breaches in security
- 413 • Motivations and ramifications of cyber terrorism and criminal hacking, “cracking”
- 414 • Effects of malware, such as viruses, worms and Trojan horses
- 415 • Crime prevention strategies
- 416 • Security policies (cross-reference IAS/Security Policy and Governance/security policies)
- 417

418 ***Learning Outcomes:***

- 419 1. List classic examples of computer crimes and social engineering incidents with societal impact.
- 420 [Familiarity]
- 421 2. Identify laws that apply to computer crimes [Familiarity]
- 422 3. Describe the motivation and ramifications of cyber terrorism and criminal hacking [Familiarity]
- 423 4. Examine the ethical and legal issues surrounding the misuse of access and various breaches in security
- 424 [Usage]

- 425
- 426
- 427
- 428
- 429
5. Discuss the professional's role in security and the trade-offs involved. [Familiarity]
 6. Investigate measures that can be taken by both individuals and organizations including governments to prevent or mitigate the undesirable effects of computer crimes and identity theft [Usage]
 7. Write a company-wide security policy, which includes procedures for managing passwords and employee monitoring. [Usage]