

Reflections on System Trustworthiness

PETER G. NEUMANN*

*Computer Science Laboratory
SRI International
EL-243, 333 Ravenswood Ave
Menlo Park, CA 94025-3493
USA*

Abstract

We examine here a range of concerns relating to computer systems and networks, with particular attention to difficulties in system development, and the resulting vulnerabilities, threats, and risks. We examine some approaches that might achieve dramatic improvements in the ability to develop, operate, and use trustworthy systems. The problems and their solutions typically require a combination of technology and social policy.

1. A Total-System Perspective	270
2. Anticipating Disasters	271
3. Trustworthiness	273
4. Risks in Trusting Untrustworthiness	275
5. Principles for Developing Trustworthy Systems	277
5.1. Saltzer–Schroeder Security Principles	277
5.2. Further Principles	280
6. System Composition: Problems and Potentials	285
6.1. Other Manifestations of Composition	290
6.2. Approaches for Predictable Composition	291
7. A Crisis in Information System Security	294
8. Optimistic Optimization	295
9. An Example: Risks in Electronic Voting Systems	297
10. The Need for Risk Awareness	300
11. Risks of Misinformation	302
12. Boon or Bane?	304
Acknowledgements	306
References	306

*Principal scientist.

1. A Total-System Perspective

Each of the following items presents pressing challenges relating to the constructive use of information technology. The totality of all the interrelated challenges requires concerted efforts that transcend the individual problems.

• *Trustworthiness.* Trustworthiness implies simply that something is worthy of being trusted to satisfy its expected requirements. Users often trust systems that are not worthy of being trusted, with respect to attributes such as system and network security, system reliability and survivability, human safety, interoperability, predictable system behavior, and other important attributes that are for the most part not receiving enough concerted attention. Computer-communication infrastructures are typically riddled with flaws. In the absence of more serious attacks, governments and system developers seem to have been lulled into a false sense of security. At present, neither proprietary nor source-available system developers are sufficiently militant in satisfying critical needs. In mass-market software, the patch mentality seems to have won out over well-designed and well-implemented systems.

• *Total system life-cycle issues.* Developing and operating trustworthy systems is inherently difficult today. Typically, a system is not likely to be trustworthy unless the relevant attributes were explicitly recognized from the beginning of system development, reflected in sound system architectures and software development, explicitly addressed in system procurements, and their fulfillment mandated throughout system operation.

• *System development practice.* Costly failures have occurred in developing large systems, such as the modernization efforts for the US Internal Revenue Service, US and UK air traffic control systems, the FBI Virtual Case File, and German TollCollect, to name just a few. Procurement and development of large-scale hardware/software systems remains a high-risk activity, with cost overruns, delays, and even abandonment of entire projects.

• *The Internet.* Increasingly, many enterprises are heavily dependent on the Internet, despite its existing limitations. Internet governance, control, and coordination create many contentious international problems. The Internet infrastructure itself is susceptible to denial-of-service attacks and compromise, while the lack of security and dependability of most attached systems also creates problems (e.g., penetrations such as open relays being used to host zombies and “bots”). Worms, viruses, and other malware are often impediments, as are ubiquitous problems of spam e-mail and phishing attacks that may result in identity theft.

• *Critical national infrastructures.* Despite some past recognition of the pervasiveness of serious vulnerabilities, critical national infrastructures such as electrical

1 power, energy, telecommunications, transportation, finance, and government conti- 1
2 nuity are typically still vulnerable to attacks and accidental collapses. For example, 2
3 massive power outages keep recurring, despite supposed improvements. Telecom- 3
4 munications outages can have severe consequences, as can transportation shutdowns 4
5 and fuel shortages. Furthermore, these infrastructures have interdependencies that 5
6 result in widespread system failures. 6

7
8 • *Accountability.* Oversight of computer activities is often as weak as oversight of 7
9 corporate practices. On the other hand, audit mechanisms must also respect privacy 8
10 needs. As one example that fails on both counts, today’s un-auditable all-electronic 9
11 voting systems are lacking in accountability in a would-be effort to protect voter 10
12 privacy. In fact, without the addition of some sort of voter-verified audit trail, they 11
13 provide no meaningful assurances that votes are correctly recorded and processed. 12
14 (For example, see the October 2004 special issue of the *Communications of the ACM*, 13
15 devoted to the integrity of election systems.) 14

16
17 • *Privacy.* Privacy is something that many people do not value until after they 16
18 have lost it. Personal privacy is relevant pervasively in our lives, especially in finan- 17
19 cial matters and health care. Some advocates of homeland security have postulated 18
20 the need to sacrifice privacy in order to attain security, although the necessity of this 19
21 tradeoff is highly debatable. Sacrificing privacy does not necessarily result in greater 20
22 security. (Benjamin Franklin’s quote is apt in this regard: “Those who would sac- 21
23 rifice liberty for security deserve neither.”) Furthermore, serious inroads to privacy 22
24 protection have occurred that may be difficult to reverse. Surveillance is becoming 23
25 more widespread, but often without adequately respecting privacy. Legitimate needs 24
26 for anonymity or at least pseudoanonymity (for example, to protect victims and le- 25
gitimate whistle-blowers) must not be suppressed or dismissed as dangerous. 26

27
28 • *Education.* In many countries, university curricula in software engineering and 27
29 trustworthiness inadequately reflect the needs of critical systems. Instruction is often 28
30 aimed at programming in the small, while more or less ignoring systems in the large. 29
31 This situation has potentially serious long-term implications worldwide. 30

32
33 As noted above, it is the totality of these problems that is of primary concern. 32
34 Simplistic local approaches are not effective. Greater foresight and pervasive system- 33
35 oriented thinking are urgently needed, along with greater private-public cooperation. 34

36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136 137 138 139 140 141 142 143 144 145 146 147 148 149 150 151 152 153 154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170 171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204 205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221 222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238 239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255 256 257 258 259 260 261 262 263 264 265 266 267 268 269 270 271 272 273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289 290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306 307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323 324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425 426 427 428 429 430 431 432 433 434 435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493 494 495 496 497 498 499 500

As Henry Petroski noted over twenty years ago [42], we generally learn less from 39
40 successes than from failures. The ACM Risks Forum [32] and *Computer-Related* 40

1 *Risks* [34] include a startling number of failures and risks, and provide a goldmine 1
2 of opportunity for anyone who wants to learn from past mistakes. Intriguingly, or 2
3 perhaps ironically, most of the content of [34] is still as relevant today as it was in 3
4 1995. The same types of failures continue to recur, and the range of causes remains 4
5 much the same. Indeed, the scope and extent of the risks has increased steadily. For 5
6 example, the ACM Risks Forum continues to report computer system development 6
7 fiascos and operational failures of aircraft, air-traffic control, defense systems, train 7
8 crashes, electrical power, telecommunications, medical health systems, and finan- 8
9 cial problems. These difficulties include problems in reliability, system survivability, 9
10 security, privacy, and human well-being. Some of these problems have been es- 10
11 calating dramatically, such as spam, phishing attacks, identity thefts, and financial 11
12 losses. 12

13 In recent years, some unusual natural disasters have occurred, such as the 9.0- 13
14 magnitude Indonesian earthquake that triggered a tsunami killing more than 200,000 14
15 people in 11 countries around the Indian Ocean, the exceptionally heavy 2006 hur- 15
16 ricane season in the Caribbean area including the devastating effects of Katrina, 16
17 and a major mudslide in La Conchita, California. Although failures of information 17
18 technology obviously had no role in *triggering* these disasters, IT systems could 18
19 play significant roles in *anticipating, detecting, monitoring, and responding to* such 19
20 events, minimizing losses of life, injuries, and consequential damages. What have 20
21 we learned from such events, especially with respect to the need for proactive con- 21
22 tingency plans? 22

23 For example, a tsunami detection and early-warning system such as had already 23
24 been deployed in the Pacific Ocean could also have been used in the Indian Ocean. 24
25 Such a system could have given timely warnings to millions of people, and could 25
26 have saved many lives *if* local authorities had citizen alerts and evacuation plans in 26
27 place. Early warnings and preparedness for hurricanes and typhoons are improving as 27
28 computer prediction of possible storm paths is becoming more accurate and as many 28
29 authorities prepare disaster response plans and train for their deployment. However, 29
30 preparedness tends to improve only after disasters have occurred (and then often 30
31 only temporarily). In the case of the mudslide in the hills above La Conchita, which 31
32 followed an awesome sequence of rainstorms, sensors in the hills were designed to 32
33 trigger advance warnings, which evidently were not taken seriously enough. (A simi- 33
34 lar slide had occurred in an adjacent area nine years earlier, and insurance companies 34
35 had already declined to provide future coverage.) 35

36 Several problems arise in connection with developing detection and warning sys- 36
37 tems. 37

38 • Institutions (especially governments, corporations, and defense departments) 38
39 tend to fashion response plans for past situations rather than for potentially dev- 39
40 astating future situations. Unless a similar disaster has recently occurred in a sim- 40

1 ilar venue under similar conditions, few people worry about low-probability high- 1
2 impact events. A comparable tendency holds for trustworthy computing. A 2
3 computer networking event not unlike a tsunami occurred in 1988—namely, the Internet 3
4 Worm [47,54] that affected about 10% of the 60,000 Internet hosts active at the time. 4
5 As a result, an emergency response team (now US-CERT) was formed to help coordi- 5
6 nate responses and warn of vulnerabilities. Prior to the year 2000, a large upgrade 6
7 effort to avoid the Y2K crisis was generally successful; the situation could have 7
8 been much more serious without the intensive remediation efforts. Today, many new 8
9 threats such as malware and terrorist attacks could easily disable critical infrastruc- 9
10 tures and the Internet. However, because the cybersecurity equivalent of a tsunami 10
11 seems extremely unlikely to many people unfamiliar with the nature of the vulnera- 11
12 bilities, there is little interest in mounting efforts to increase system trustworthiness 12
13 and engage in other preventive measures. The consequences of major meltdowns 13
14 could be very dramatic, especially if accompanied by terrorist attacks. 14

15 • Institutions tend to optimize short-term costs and ignore long-term conse- 15
16 quences. Also, farsighted analyses of what might happen are always subject to poor 16
17 assumptions, faulty reasoning, and mandates to reach self-serving conclusions. This 17
18 is discussed further in Section 8. 18
19

20 • People generally do not like to make unnecessary preparations, and often resent 20
21 taking sensible precautions. Repeated false warnings tend to inure them, with a 21
22 resulting loss of responsiveness. Even justifiable warnings that are heeded (such as the 22
23 Y2K remediation or boarding up for an oncoming hurricane) are often denigrated if 23
24 the resulting effects are only relatively minor. 24
25

26 It is clear that much greater attention needs to be devoted to predicting, detecting, 26
27 and ameliorating both natural catastrophes and unnatural computer-related misuse, 27
28 attacks, disasters, and outages. Efforts are needed to dramatically improve the trust- 28
29 worthiness of those systems on which many lives depend, and to make those systems 29
30 more tolerant to human misbehavior as well as malfunctions and natural causes. 30
31

32 3. Trustworthiness 33

34 Estimates of system trustworthiness ultimately depend on having some sort of 34
35 logical basis for confidence that a system will predictably satisfy its critical 35
36 requirements. Measures of trustworthiness are particularly important for information 36
37 security, system integrity and reliability, human safety, fault tolerance, and overall 37
38 enterprise survivability in the face of wide ranges of adversities (including malfunc- 38
39 tions, deliberate attacks, and natural causes). 39
40

1 Many lives increasingly depend on critical national infrastructures—all of which 1
 2 in turn depend in varying degrees on the predictable behavior of computer- 2
 3 communication resources. Indeed, these infrastructures often depend on the Internet 3
 4 for information and control and may be vulnerable to attacks from any attached 4
 5 computer systems. 5

6 Unless critical information system resources are sufficiently trustworthy, other 6
 7 systems are at risk from failures and subversions. Unfortunately, for many of the key 7
 8 application domains, the existing information infrastructures are lacking in trustwor- 8
 9 thiness. For example, power grids, air-traffic control, high-integrity electronic voting 9
 10 systems, the emerging US Department of Defense Secure Global Information Grid, 10
 11 national infrastructures, and many collaborative and competitive Internet-based ap- 11
 12 plications all need systems that are more trustworthy than we have today or are likely 12
 13 to have in the foreseeable future. 13

14 Numerous steps are needed to develop trustworthy systems. Consider an analogy 14
 15 with the planet’s natural environment—expectations for which are somewhat simi- 15
 16 lar to expectations for trustworthy information systems. For example, pure air and 16
 17 uncontaminated water are vital, as are the social systems that ensure them. 17

18 Although poorly chosen analogies can be misleading, the analogy with the natu- 18
 19 ral environment is appropriate. Each of the following items is applicable to both 19
 20 trustworthy information systems and natural environments. 20

- 21 • Their critical importance is generally underappreciated until something goes 21
 22 fundamentally wrong—after which undoing the damage can be very difficult if 22
 23 not impossible. 23
- 24 • Problems can result from natural circumstances, equipment failures, human er- 24
 25 rors, malicious activity, or a combination of these and other factors. 25
- 26 • Dangerous contaminants may emerge and propagate, often unobserved. Some 26
 27 of these may remain undetected for relatively long periods of time, whereas 27
 28 others can have immediately obvious consequences. 28
- 29 • Once something has gone recognizably wrong, palliative countermeasures are 29
 30 typically fruitless—too little, too late. 30
- 31 • Your own well-being may be dramatically impeded, but there is not much you 31
 32 as an individual can do about aspects that are pervasive—perhaps international 32
 33 or even global in scope. 33
- 34 • Detection, remediation, and prevention require cooperative social efforts, such 34
 35 as public health and sanitation activities, as well as technological means includ- 35
 36 ing increased trustworthiness. 36
- 37 • Up-front preventive measures can result in significant savings and increased 37
 38 human well-being, ameliorating major problems later on. 38
- 39 39
- 40 40

- 1 ● As discussed further in Section 8, long-term thinking is relatively rare. There 1
2 is frequently little governmental or institutional emphasis on proactive prevention 2
3 of bad consequences. Many of the arguments against far-sighted planning 3
4 and remediation are skewed, being based on faulty, narrowly scoped, or short- 4
5 sighted reasoning—especially relating to short-term profits rather than long- 5
6 term savings and other benefits. 6
- 7 ● Commercial considerations tend to trump human well-being, with business 7
8 models sometimes considering protection of public welfare to be detrimental 8
9 to corporate and enterprise bottom lines. 9

10 In some contexts, pure water is becoming more expensive than oil. Fresh air is 10
11 already a crucial commodity. Short- and long-term effects of inadequately trustworth- 11
12 y information systems can similarly be costly. Proactive measures are as urgently 12
13 needed for system trustworthiness as they are for breathable air, clean water, and 13
14 environmental protection. It is difficult to remediate computer-based systems that were 14
15 not designed and implemented with trustworthiness in mind. It is also difficult to 15
16 remediate serious environmental damage. 16
17

18 Anticipating and responding to compelling long-term needs does not require extra- 18
19 ordinary foresight, whether for air, water, reversing global warming, or trustworthy 19
20 systems upon which to build infrastructures. Our long-term well-being—perhaps 20
21 even our survival—depends on our willingness to consider the future and commit- 21
22 ment to taking appropriate actions. 22
23

24 **4. Risks in Trusting Untrustworthiness** 24

25 The Internet provides ample opportunity for proving the age-old truism, “There’s 25
26 a sucker born every minute.” Carnival-style swindles and other confidence games 26
27 once limited to in-person encounters are now proliferating electronically, world- 27
28 wide, at low cost and effort. Blatantly obvious pre-Internet examples are the so-called 28
29 Nigerian-style postal scams that requested use of one’s bank account to help move 29
30 money; hoping for a proffered generous commission, the suckers are then separated 30
31 from their assets. These scams have been updated to today’s e-mail phishing and 31
32 e-mail scam attacks that efficiently harvest personal information from vastly more 32
33 people, and are considerably more sophisticated—for example, replicating a legiti- 33
34 mate website in every respect except for perhaps just one hard-to-detect bogus URL. 34
35 Indeed, it is becoming increasingly difficult to distinguish the real from the bogus, 35
36 and people continue to be victimized. 36
37

38 Many other kinds of scams, stings, and misrepresentations also exist. Deceptive 38
39 unsolicited e-mail (spam) offering bogus goods and services opens up new avenues 39
40 40

1 for fraud and identity theft. Online activities are emerging with glaring opportunities
 2 for swindles, manipulations, and assorted malfeasance, such as online auctions (with
 3 irregularities such as nondelivery and secondary criminality), an alarming increase
 4 in highly sophisticated phishing attacks, Internet gambling, and fraudulent websites
 5 (e.g., with deceptive URLs creating the appearance of legitimacy). Any of these and
 6 other situations could result in inordinate risks, such as financial ruin, blackmail,
 7 compromised democracy, or even loss of life. But it is perhaps not surprising that
 8 people fall for such schemes, particularly when the technology superficially appears
 9 genuine.

10 Today’s unauditible paperless all-electronic voting systems present significant
 11 risks (see Section 9). The risks are even greater for voting over the Internet. With
 12 independent accountability seriously lacking today, e-voting can be likened to using
 13 an off-shore gambling site not subject to any regulation and managed by unknown
 14 and unaccountable agents.

15 We tend to trust third-party relationships with banks, telephone companies, air-
 16 lines, and other service providers whose employees have in some way earned our
 17 trust, collectively or individually. But what about untrustworthy third parties? Some
 18 computer-based applications rely critically on the putative integrity and noncom-
 19 promissibility of automated trusted third parties, with little if any easily demonstrated
 20 human accountability. Examples include digital-certificate authorities, cryptographic
 21 servers, surveillance facilities, sensitive databases for law enforcement, and credit-
 22 information bureaus. With appealing short-term cost incentives for pervasive use of
 23 outsourcing, the need for demonstrably trustworthy third-party institutions becomes
 24 even more important. However, security, privacy, and accountability are often ig-
 25 nored in efforts to save money.

26 Is placing trust in offshore enterprises inherently riskier than using domestic ser-
 27 vices? Not necessarily. Corruption and inattention to detail are worldwide problems.
 28 The deciding factor here is the extent to which comprehensive oversight can be main-
 29 tained.

30 Is domestic legislation enough? Of course not. Any legislation should not be
 31 overly simplistic; for example, it should avoid seeking solely technological fixes or
 32 purely legislative solutions to deeper problems. Besides, serious complexities arise
 33 from the fact that such problems are international in scope and demand international
 34 cooperation.

35 Is there a role for liability (for flagrant misbehavior or injurious neglect) and dif-
 36 ferential insurance rates—for example, based on how well a purveyor is living up to
 37 what is expected of it? Such measures have significant potential, although they will
 38 be strongly resisted in many quarters.

39 So, how can we provide some meaningful assurance that critical entities such as
 40 direct or third parties are sufficiently trustworthy? Ideally, institutions providing,

1 controlling, managing, and monitoring potentially riskful operations should be de- 1
2 coupled from other operations, avoid conflicts of interest, and be subject to rigorous 2
3 independent oversight. Enronitis and collusion must be avoided, even if it means 3
4 that the costs are greater. Furthermore, the people involved need altruism, sufficient 4
5 foresight to anticipate the risks, and a commitment to effectively combat those risks. 5
6 At the very least, their backgrounds should be free of criminal convictions and other 6
7 activities that would cast serious suspicions on their trustworthiness. In addition, leg- 7
8 islators need to be able to see beyond the simplistic and palliative, to approaches that 8
9 address the real problems. Above all, the entire populace must become more aware 9
10 of the risks and the concerns outlined above, especially to the inherent combination 10
11 of technology and policy issues. 11

12 This conclusion should not be a surprise. Overall, there are many risks that must 12
13 be addressed. The old Latin expression “Caveat emptor” (Let the buyer beware) is 13
14 even more timely today. 14

15 5. Principles for Developing Trustworthy Systems 15

16 *Everything should be made as simple as possible—but no simpler.* 16

17 Albert Einstein 17

18 Developing trustworthy systems with complex requirements is inherently a 18
19 complex challenge. In general, simple solutions are hopelessly inadequate in such 19
20 cases. On the other hand, enormously complex systems—even if they purport to 20
21 be trustworthy—are likely to be unmanageable, from the perspective of developers, 21
22 system administrators, application operators, and end-users. 22

23 Ideally, there should be some middle ground. In particular, the recommended 23
24 approach, considered in Section 6, is to develop trustworthy systems as conceptually 24
25 sound predictable compositions of simpler components, perhaps even with provably 25
26 sound combinations of provably sound components. 26

27 In anticipation of that approach, a relevant set of principles can be helpful in 27
28 increasing trustworthiness—if the principles are used intelligently as guidelines for 28
29 system development and operation. 29

30 5.1 Saltzer–Schroeder Security Principles 30

31 The ten basic security principles formulated by Jerry Saltzer and Mike Schroeder 31
32 [51] in 1975 are all still relevant today, in a wide range of circumstances. They are 32
33 actually of broader interest than just with respect to security. For example, each one 33
34 is also relevant to reliability, survivability, and human safety. In essence, these prin- 34
35 ciples are summarized as follows (overly tersely), for present purposes: 35
36 36
37 37
38 38
39 39
40 40

- 1 ● *Economy of mechanism*: Seek design simplicity (wherever and to whatever extent it is effective). 1
- 2
- 3 ● *Fail-safe defaults*: Deny accesses unless explicitly authorized (rather than permitting accesses unless explicitly denied). 3
- 4
- 5 ● *Complete mediation*: Check every access, without exception. 5
- 6
- 7 ● *Open design*: Do not assume that design secrecy will enhance security. 7
- 8
- 9 ● *Separation of privileges*: Use separate privileges or even multiparty authorization (e.g., two keys held by different entities) to reduce misplaced trust. 9
- 10
- 11 ● *Least privilege*: Allocate minimal (separate) privileges according to need-to-know, need-to-modify, need-to-delete, need-to-use, and so on. The existence of 11
- 12 powerful mechanisms such as *superuser* is inherently dangerous. 12
- 13
- 14 ● *Least common mechanism*: Minimize the amount of mechanism common to 14
- 15 more than one user and depended on by all users. 15
- 16
- 17 ● *Psychological acceptability*: Strive for ease of use and operation—for example, 17
- 18 with easily understandable and forgiving interfaces. 18
- 19
- 20 ● *Work factors*: Make cost-to-protect commensurate with threats and expected 20
- 21 risks. 21
- 22
- 23 ● *Recording of compromises*: Provide nonbypassable tamper-resistant and tamper- 23
- 24 evident audit trails of evidence. 24

25 These are of course basic guidelines, not hard-and-fast rules—especially in light 25

26 of some potential mutual contradictions. Two fundamental caveats must be recog- 26

27 nized. First, each principle by itself may be useful in some cases and not in others. 27

28 Second, when taken in combinations, groups of principles are not necessarily all re- 28

29 inforcing; indeed, they may in some cases conflict with one another. Consequently, 29

30 development must consider appropriate use of each principle in the context of the 30

31 overall effort. Examples of a principle having both positive and negative aspects are 31

32 scattered through the following discussion. 32

33 The Saltzer–Schroeder principles grew directly out of the MIT/Honeywell/ 33

34 BellLabs Multics experience (e.g., [40]) begun in 1965 and discussed further later in 34

35 this section. Each of these principles has taken on almost mythic proportions among 35

36 the security elite, and to some extent buzzword cult status among many fringe parties. 36

37 Therefore, we do not explain each principle in detail—although considerable 37

38 depth of discussion is needed for successful application of each principle. Careful 38

39 reading of the Saltzer–Schroeder paper [51] is recommended if it is not already a 39

40 part of your library. Matt Bishop’s security books [7,8] are also useful in this regard, 40

41 placing the principles in a more general context.

42 Various caveats are considered in Section 12.

TABLE I
 APPLICABILITY OF SALTZER–SCHROEDER PRINCIPLES

Principle	Trustworthiness	Assurance
Economy of mechanism	is a vital aid to sound design. Exceptions must be handled completely.	can simplify local analysis.
Fail-safe defaults	simplifies design, use, operation, maintenance.	can simplify analysis.
Complete mediation	is vital, but beware of compromise from below.	can simplify analysis locally.
Open design	Secret designs do not preclude compromise. Open design can inspire stronger system security.	Open designs do not preclude compromise. Open design encourages independent analysis.
Separation of privileges	avoids many types of common flaws.	focuses analysis more precisely.
Least privilege	limits effects of flaws; simplifies operation.	focuses analysis more precisely.
Least common mechanism	avoids certain common flaws.	modularizes analysis.
Psychological acceptability	is relevant to usability and operations.	Ease of use is helpful, must anticipate crises.
Work factors	are misleading if systems can be compromised from outside/within/below.	give a false sense of security if unexpected compromises are ignored.
Compromise recording	is an after-the-fact diagnostic and deterrent.	is only an indirect contributor.

Table I summarizes how each of the Saltzer–Schroeder principles can contribute to the goals of trustworthiness and assurance, particularly with respect to security, reliability, and other survivability-relevant requirements. Intriguingly, most of these principles can also be helpful in enhancing composability.

In particular, complete mediation, separation of privileges, and allocation of least privilege are beneficial to composability and trustworthiness. Open design can contribute significantly to composability, when subjected to internal review and external criticism. (See Section 6.) However, conflicts persist about the importance of open design with respect to trustworthiness, with some people still clinging tenaciously to the notion that security by obscurity is sensible—despite risks of many flaws being so obvious as to be easily detected externally, even without reverse engineering. Indeed, the recent emergence of good decompilers for C and Java, along with the likelihood of similar reverse engineering tools for other languages, suggests that such attacks are becoming steadily more practical. Overall, the pretense of keeping a design se-

cret and the supposed unavailability of source code are realistically not significant deterrents, especially with ever-increasing skills among black-box system analysts. However, there are cases in which reliance on security by obscurity is unavoidable—as in the hiding of private and secret cryptographic keys, although the cryptographic algorithms and implementations can be public.

Fundamental to trustworthiness is the extent to which systems and networks can avoid being compromised by malicious or accidental human behavior and by events such as hardware malfunctions and so-called acts of God. In [35], we consider *compromise from outside*, *compromise from within*, and *compromise from below*, with fairly intuitive meanings. These notions appear throughout this report.

In theory, there are various cases where certain weak links can be avoided (such as zero-knowledge protocols that can establish a shared key without any part of the protocol requiring secrecy, Byzantine algorithms, and k -out-of- n cryptography), although in practice they may be undermined by compromises from below (involving trusted and supposedly trustworthy insiders subverting the underlying operating systems) or from outside (involving penetrations of the operating systems and masquerading as legitimate users).

From its beginning, the Multics development was strongly motivated by a set of principles—some of which were originally stated by Ted Glaser and Neumann in the first section of the first edition of the Multics Programmers' Manual in 1965. (See <http://multicians.org>.) It was also driven by extremely disciplined development. For example, no coding effort was begun until a written specification had been approved by the Multics advisory board; also, with just a few exceptions such as low-level device drivers, all the code was written in a subset of PL/I just sufficient for the needs of Multics, for which the first compiler (early PL, or EPL) had been developed by Doug McIlroy and Bob Morris.

In addition to the Saltzer–Schroeder principles, further insights on principles and discipline relating to Multics can be found in a paper by Fernando Corbató, Jerry Saltzer, and Charlie Clingen [12] and in Corbató's Turing lecture [11].

5.2 Further Principles

An earlier view of principled system development was given by Neumann in 1969 [33], relating to what is often dismissed as merely “motherhood”—but which in reality is both profound and difficult to observe in practice. The principles under consideration in that paper included automatedness, availability, convenience, debuggability, documentedness, efficiency, evolvability, flexibility, forgivingness, generality, maintainability, modularity, monitorability, portability, reliability, simplicity, and uniformity. Some of those attributes indirectly affect security and trustworthiness, whereas others affect the acceptability, utility, and long-term future of systems.

1 Considerable discussion in [33] was also devoted to (1) the risks of local optimization and the need for a more global awareness of less obvious downstream costs of development (e.g., writing code for bad—or nonexistent—specifications, and having to debug really bad code), operation, and maintenance (see Section 8); and (2) the benefits of higher-level implementation languages (which prior to Multics were rarely used for the development of operating systems [11,12]).

7 In the context of developing predictably trustworthy systems, an expanded set of principles is listed below. Although most of them might seem more or less obvious to advanced developers, there are interpretations, hidden issues, and potential pitfalls for their successful implementation. As a result, a seemingly paradoxical situation arises: understanding and experience are required in order to make optimal use of the principles. Thus, the learning experience is essentially iterative.

- 13 • *Sound architecture.* Recognizing that it is better to avoid design errors early than to attempt to fix them later, composable architectures inherently capable of evolvable, maintainable, robust implementations are required. Furthermore, good interface design is as fundamental to good architectures as is their internal designs. Both the architectural structure and the architectural interfaces (particularly the visible interfaces, but also some of the internal interfaces that must be interoperable) can benefit from careful specification.
- 21 • *Abstraction.* The primitives at any given logical or physical layer should be relevant to the functions and properties of the objects at that layer, and should mask lower-layer detail where possible. Ideally, the specification of a given abstraction should be in terms of objects meaningful at that layer, rather than requiring lower-layer (e.g., machine-dependent) concepts. Abstractions at one layer can be related to the abstractions at other layers in a variety of ways, thus simplifying the abstractions at each layer rather than collapsing different abstractions into a more complex single layer. Particularly useful examples of abstraction include trustworthiness kernels, virtual machine monitors, and similar layered defenses.
- 31 • *Modularity.* Modularity relates to the characteristic of system structures in which different entities (modules) can be relatively loosely coupled and combined to satisfy overall system requirements, whereby a module could be modified or replaced as long as the new version satisfies the given interface specification. In general, modularity is most effective when the modules reflect specific abstractions and provide encapsulation within each module (see the next item).
- 37 • *Encapsulation.* Details that are relevant to a particular abstraction should be local to that abstraction and subsequently isolated within the implementation of that abstraction and the lower layers on which the implementation depends. One example of encapsulation involves information hiding—for example, keeping

1 internal state information inaccessible to the visible interfaces [41]. Another
2 example involves masking the idiosyncrasies of physical devices from higher-
3 layer system interfaces, and from the user interfaces as well.

- 4 ● *Layered and distributed protection.* Protection (and generally defensive de-
5 sign for security, reliability, and so on) should be distributed to where it is
6 most needed, and should reflect the semantics of the objects being protected.
7 With respect to the reality of implementations that rely on—and perhaps pass
8 through—entities of different trustworthiness, layers of protection are vastly
9 preferable to flat concepts such as single sign-on (i.e., where only a single
10 authentication is required). With respect to psychological acceptability, single
11 sign-on has enormous appeal; however, it can leave enormous security vulner-
12 abilities as a result of compromise from outside, from within, or from below, in
13 both distributed and layered environments. Overall, psychological acceptability
14 can conflict with other principles, such as complete mediation, separation of
15 privileges, and least common privilege.
- 16 ● *Constrained dependency for integrity.* Dependencies on less trustworthy entities
17 should be avoided unless potential negative effects can be somehow confined or
18 constrained. However, it is possible in some cases to surmount the relative un-
19 trustworthiness of mechanisms on which certain functionality depends—as in
20 various types of trustworthiness-enhancing mechanisms (see [36]). In essence,
21 do not trust anything on which you must depend—unless you are adequately
22 satisfied with demonstrations of its trustworthiness or the ability to surmount
23 its relative untrustworthiness. This intuitive extension of Biba’s notion of mul-
24 tiple integrity [6] is considered further in Section 6.
- 25 ● *Architectural minimization of what must be trustworthy.* Appropriate trustwor-
26 thiness should be situated where it is most needed, suitable to overall system
27 requirements, rather than required uniformly across widely distributed compo-
28 nents (with potentially many weak links) or totally centralized (with creation
29 of a single weak link and forgetting other vulnerabilities). Trustworthiness is
30 expensive to implement and to ensure. Thus, significant benefits can result from
31 minimizing what has to be trustworthy. This principle can contribute notably to
32 sound architectures. In combination with economy of mechanism, this provides
33 avoidance of both bloatware and adverse dependence on less trustworthy com-
34 ponents. For example, in some cases a simple end-to-end check can determine
35 the presence of intermediate compromises and avoid the necessity of trusting
36 everything else for integrity (apart from denial-of-service attacks).
- 37 ● *Object orientation.* The OO paradigm bundles together abstraction, encapsu-
38 lation, modularity of state information, inheritance (subclasses inheriting the
39 attributes of their parent classes—e.g., for functionality and for protection), and
40

1 subtype polymorphism (subtype safety despite the possibility of application to 1
2 objects of different types). This paradigm facilitates programming generality 2
3 and software reusability, and if properly used can enhance software develop- 3
4 ment. This is a contentious topic, in that most of the OO methodologies and 4
5 languages are sloppy with respect to inheritance. 5

- 6 ● *Separation of policy and mechanism.* Statements of policy should avoid inclu- 6
7 sion of implementation-specific details. Furthermore, mechanisms should be 7
8 policy neutral where that can be advantageous in achieving functional general- 8
9 ity. However, this principle must never be used in the absence of understanding 9
10 about the range of policies that needs to be implemented. There is a tempta- 10
11 tion to avoid anticipating meaningful policies, deferring them until later in the 11
12 development—and then discovering that the desired policies cannot be realized 12
13 with the given mechanisms. This is a characteristic chicken-and-egg problem 13
14 with abstraction. 14
- 15 ● *Separation of duties.* In relation to separation of privileges, separate classes 15
16 of duties of users and computational entities should be identified, so that dis- 16
17 tinct system roles can be assigned accordingly. Distinct duties should be treated 17
18 distinctly, as in activities of system administrators, system programmers, and 18
19 unprivileged users. 19
- 20 ● *Separation of roles.* Concerning separation of privileges, the roles recognized by 20
21 protection mechanisms should correspond in some readily understandable way 21
22 to the various duties. For example, a single all-powerful superuser role intrin- 22
23 sically violates separation of duties, separation of roles, separation of privilege, 23
24 and separation of domains. The separation of would-be superuser functions into 24
25 separate roles (as in Trusted Xenix) is a good example of desirable separation. 25
26 Once again (as with single sign-on), there is a potential conflict between princi- 26
27 ples: the monolithic superuser mechanism provides economy of mechanism, but 27
28 violates other principles. In practice, all-powerful mechanisms are sometimes 28
29 unavoidable, and sometimes even desirable despite the negative consequences 29
30 (particularly if confined to a secure subenvironment). However, they should be 30
31 avoided wherever possible. 31
- 32 ● *Separation of domains.* Concerning separation of privileges, domains should 32
33 be able to enforce separate roles. For example, a single all-powerful superuser 33
34 mechanism is inherently unwise, and is in conflict with the notion of separation 34
35 of privileges. However, separation of privileges is difficult to implement if there 35
36 is inadequate separation of domains. Separation of domains can help enforce 36
37 separation of privilege, but can also provide functional separation (as in the 37
38 Multics ring structure, a kernelized operating system with a carefully designed 38
39 kernel, a capability-based architecture, or a virtual machine monitor). The prin- 39
40 40

1 ciple of least common mechanism is also somewhat related. It is desirable to 1
 2 avoid sharing of trusted multipurpose mechanisms, including executables and 2
 3 data, thereby minimizing the use of all-powerful mechanisms such as *superuser* 3
 4 and shared buffers (such as the historically seminal FORTRAN `common`). As 4
 5 one example of the flaunting of principles, exhaustion of shared resources pro- 5
 6 vides a huge source of covert storage channels, whereas the natural use of a 6
 7 common calendar clock provides a source of covert timing channels. 7

- 8 ● *Sound authentication.* Authentication is a pervasive problem. Nonbypassable 8
 9 authentication should be applicable to users, processes, procedures, and in gen- 9
 10 eral to any active entity or object. Authentication relates to evidence that the 10
 11 identity of an entity is genuine, that procedure arguments are legitimate, that 11
 12 types are properly matched when strong typing is to be invoked, and other sim- 12
 13 ilar aspects. 13
 14
- 15 ● *Sound authorization and access control.* Authorizations must be correctly and 15
 16 appropriately allocated, and nonsubvertible. Crude all-or-nothing authorizations 16
 17 are often riskful (particularly with respect to insider misuse and program- 17
 18 ming flaws). In applications for which user-group-world authorizations are 18
 19 inadequate, access-control lists and role-based authorizations may be prefer- 19
 20 able. Finer-grained access controls may be desirable in some cases, such as 20
 21 capability-based addressing and field-based database protection. However, 21
 22 knowing who has access to what at any given time should be relatively easy to 22
 23 determine. 23
- 24 ● *Administrative controllability.* The facilities by which systems and networks 24
 25 are administered must be well designed, understandable, well documented, and 25
 26 sufficiently easy to use without inordinate risks. 26
- 27 ● *Comprehensive accountability.* Well-designed and carefully implemented facil- 27
 28 ities are essential for comprehensive monitoring, auditing, interpretation, and 28
 29 automated response (as appropriate). Serious security and privacy issues must 29
 30 be addressed relating to the overall accountability processes and audit data. 30
 31

32 Similar to the summary in [Table I](#) the additional principles also tend to contribute 32
 33 to the goals of achieving composability, trustworthiness, and assurance. 33

34 At this point in the analysis, it should be no surprise that these and other principles 34
 35 can contribute in varying ways to security, reliability, survivability, and other -ilities. 35
 36 Furthermore, many of the principles and other “ilities” are linked. We cite just a few 36
 37 of the interdependencies that must be considered. 37

38 For example, authorization is of limited use without authentication, *whenever* 38
 39 *identity is important*. Similarly, authentication may be of questionable use with- 39
 40 out authorization. In some cases, authorization requires fine-grained access controls. 40

1 Least privilege requires some sort of separation of roles, duties, and domains. Sepa- 1
2 ration of duties is difficult to achieve if there is no separation of roles. Separation of 2
3 roles, duties, and domains each must rely on a supporting architecture. 3

4 The comprehensive accountability principle is particularly intricate, as it depends 4
5 critically on many other principles being invoked. For example, accountability is in- 5
6 herently incomplete without authentication and authorization—without which trace- 6
7 back to the users or originating entities is doubtful. In many cases, monitoring may 7
8 be in conflict with privacy requirements and other social considerations [16], unless 8
9 extremely stringent controls are enforceable. Separation of duties and least privilege 9
10 are particularly important here. All accountability procedures are subject to security 10
11 attacks, and are typically prone to covert channels as well. Furthermore, the proced- 11
12 ures themselves must be carefully monitored. Who monitors the monitors? (*Quis* 12
13 *auditiet ipsos audites?*) 13
14 14
15 15

16 6. System Composition: Problems and Potentials 16

17 17
18 18
19 The challenge of developing systems with realistic trustworthiness requirements 19
20 is inherently complex, despite persistent advice to keep it simple. However, consider 20
21 the goal of building trustworthy systems using predictably sound compositions of 21
22 well-designed components along with analysis of the properties that are preserved by, 22
23 transformed by, or emerging from the compositions. Conceptually, that can greatly 23
24 simplify and improve development. Indeed, composition is seemingly theoretically 24
25 relatively straightforward to achieve—especially if we follow the guidance of David 25
26 Parnas, Edsger Dijkstra, and others. Unfortunately, there is a huge gap between the- 26
27 ory and common practice: system compositions at present are typically *ad hoc*, based 27
28 on the intersection of potentially incompatible component properties, and dependent 28
29 on untrustworthy components that were not designed for interoperability and whose 29
30 behavior can undermine the compositions—often resulting in unexpected results 30
31 and risks. In practice, it is particularly difficult to determine all potentially nega- 31
32 tive effects of compositions of arbitrary components that were not designed with 32
33 composition explicitly in mind. 33

34 Composition is a concept that is meaningful with respect to many entities, includ- 34
35 ing requirements, specifications, protocols, implemented components, and analytic 35
36 results such as evaluations and formal proofs. In many cases, the composition of 36
37 different entities may have unpleasant results. 37

38 Other problems may arise because of the order in which operations are carried out, 38
39 even though the operations may be theoretically commutative or in some broader 39
40 sense equivalent (perhaps producing different but nevertheless acceptable results). 40

1 For example, consider the combination of error-correcting coding (which adds redun- 1
2 dancy), compression (which removes redundancy), and cryptography (which ideally 2
3 makes meaningful content look essentially random). Compressing after encrypting 3
4 makes little sense, because there is little apparent redundancy. Similarly, compress- 4
5 ing after adding redundancy for error correction also makes little sense, because it 5
6 vitiates the overall error correction. Thus, if such a combination were to be effective, 6
7 compression should precede encryption, which then should be followed by error- 7
8 correcting coding. 8

9 With regard to subsystem composition, the following are particular concerns. 9

- 10 • *Composability and compositionality.* A distinction is sometimes made between 10
11 two concepts pertaining to composition. *Composability* relates to the 11
12 predictability of the preservation or transformation of existing properties under 12
13 composition. *Compositionality* refers to the predictability of properties that 13
14 emerge as a result of compositions. 14
15
- 16 • *Inadequate requirements.* If stated requirements do not explicitly demand that 16
17 subsystems and other components be developed in ways that encourage compat- 17
18 ibility and interoperability, composability is likely to be difficult to achieve. 18
19 Furthermore, poorly defined requirements are likely to hinder composability. 19
- 20 • *Nonexistent or inappropriate specifications.* If system and subsystem specifi- 20
21 cations do not adequately define the relationships among interfaces, inputs, 21
22 internal state information and state transitions, outputs, and exception condi- 22
23 tions, and if those specifications are oblivious to critical relationships with 23
24 related functionality, determining to what extent composability is possible be- 24
25 comes much more difficult. Composition of underconstrained specifications is 25
26 an inherent problem, because the extent to which the components compose is ill- 26
27 defined; supposed demonstrations of composability may actually be meaning- 27
28 less. Overly constrained specifications (e.g., including unnecessarily low-level 28
29 and possibly incompatible details) are also often an impediment to compos- 29
30 ability. Shared state information across components is a particular source of 30
31 potential problems. 31
- 32 • *Properties that exist beyond what is defined by stated individual subsystem 32
33 interface specifications.* Assuming the presence of meaningful specifications, 33
34 inadequacies of the specifications and inconsistencies between specifications 34
35 and implementations are characteristic problems. In general, specifications are 35
36 always inherently incomplete with respect to defining what should *not* happen, 36
37 even if they are fairly good at defining what should happen. (Abstraction is a 37
38 very important technique for simplifying specifications, but it suppresses detail 38
39 that may include undesirable aspects of behavior and may therefore negatively 39
40 affect compositional properties.) In addition, programming languages and com- 40

1 pilers themselves provide very few if any guarantees that something beyond 1
2 what is expected cannot occur. Examples include shared-buffer interactions and 2
3 unanticipated information residues from one invocation of a subsystem to a sub- 3
4 sequent or concurrent invocation of the same subsystem; buffer overflows and 4
5 other cases of inadequate bounds checks and inadequate runtime validation; 5
6 inadequate authentication; improper initialization and finalization; improper en- 6
7 capsulation, which can result in interference and other unexpected interactions; 7
8 race conditions; covert channels; and intentionally planted Trojan horses. This 8
9 list represents just the nose of the camel. All these problems can impair com- 9
10 posability. As one example, various Windows operating systems are actually 10
11 relatively modular (which is essential for orderly development), but the mod- 11
12 ules are not sufficiently encapsulated to prevent adverse effects resulting from 12
13 composition. 13

- 14 ● *Properties that manifest themselves only as a result of combinations of sub-* 14
15 *systems.* Examples include adverse *emergent* properties (i.e., disruptive or even 15
16 constructive effects that are not evident in any of the individual subsystems but 16
17 that arise only when the subsystems are combined); adverse feedback interac- 17
18 tions between subsystems, such as infinite loops or dependence on functionality 18
19 that is less trustworthy; emergent covert channels that do not exist in any of 19
20 the subsystems in isolation; mutual incompatibilities in the interfaces—perhaps 20
21 resulting from internal state interference; global failure modes resulting from 21
22 local faults, as in the 1980 ARPANET collapse [48] and the 1990 AT&T long- 22
23 distance collapse (e.g., see [34]); so-called “man-in-the-middle” attacks (which 23
24 might alternatively be called untrustworthy interpositions), in which an inter- 24
25 poser can simulate the actions of each component; and other failure modes that 25
26 arise only in the overall system context. A fascinating noncomposability situ- 26
27 ation is noted in attempts to combine encryption with digital signatures [4]: 27
28 signatures are composable with public-key cryptography, but *not* with symmetric 28
29 cryptography, in which case security may break down. These impediments 29
30 to composability can arise essentially everywhere throughout the development 30
31 life cycle—for example, incompatibilities among different requirements and 31
32 policies, undesirable interactions in specifications and implementations, and 32
33 difficulties in reconfiguration and maintenance. 33
- 34 ● *Multivendor and multiteam incompatibilities.* In the interest of having heteroge- 34
35 neous architectures that enable mixing and matching of alternative components, 35
36 it may be desirable to use multiple system developers. However, incompati- 36
37 bilities among interface assumptions, the existence of proprietary internal and 37
38 external interfaces, and extreme performance degradations resulting from the 38
39 inability to optimize across components can all result in difficulties in compos- 39
40 ing components. 40

- 1 • *Scalability*. Composability typically creates many issues of scalability. For 1
2 example, performance may degrade badly or nonpredictably as multiple sub- 2
3 systems are conjoined. Composability can lead to a wide range of expected 3
4 performance implications—for example, linear, multiplicative, or exponential 4
5 in the number of composed subsystems. In practice, even further degradations 5
6 can result—for example, from design or implementation flaws or indirect ef- 6
7 fects of the composition, such as unrecognized dependence on substantively 7
8 slow interactions. Obviously, infinite loops and standstill deadlocks (“deadly 8
9 embraces”) are limiting cases of degradation, and often arise as a result of sub- 9
10 system compositions. 10
- 11 • *Human issues*. Above all, people are the ultimate source of many problems. 11
12 The supposed “good guys” can accidentally have profoundly negative effects 12
13 on composability, through poor system conception, inadequate requirements, 13
14 lack of comprehensive and accurate specifications, bad software-engineering 14
15 practice, misuse or bad choices of programming languages, badly managed de- 15
16 velopment, and sloppy operational practice (for example). Insider “bad guys” 16
17 can have various negative effects on the desired composability, such as installing 17
18 Trojan horses during development, operation, and reconfiguration that impair 18
19 interoperability and compromise security. Human activities can also directly 19
20 impair enterprise interoperability [18]. Outsider “bad guys” are generally less 20
21 likely to negatively affect composability externally, except as a result of pen- 21
22 etrations (through which they effectively become bad insiders), subversion of 22
23 the development process, tampering, and denials of service (often without any 23
24 internal access required). 24

25 There are many desiderata for achieving predictably assured composition, relating 25
26 to requirements, specifications, implementations, programming languages, configu- 26
27 ration information, and analyses thereof. Several relevant issues are noted below. 27
28

- 29 • *Compatibility and interoperability*. Compatibility implies merely the ability to 29
30 coexist within a common framework, whereas interoperability additionally im- 30
31 plies the ability to work together without adverse side effects. Both are essential 31
32 prerequisites for composability. 32
- 33 • *Web interoperability*. In recent years, considerable effort has been devoted to 33
34 ward establishing a common definition of a *Web portal* concept that would 34
35 facilitate universal interoperability providing access to Web services. As one 35
36 example, Michael Alan Smith [53] has proposed a hierarchical General Portal 36
37 Model that attempts to unify seventeen different definitions from the litera- 37
38 ture. From the top, the layers address process interfaces (process identification, 38
39 transformation), resource discovery (resource identification, resource location, 39
40 resource binding), and network interfaces (security, network access). In this 40

1 context, a portal implies an “infrastructure providing secure, customizable, per- 1
2 sonalizable, integrated access to dynamic content from a variety of sources, in 2
3 a variety of formats, *wherever it is needed*.” Among other approaches is that of 3
4 a service-oriented architecture (e.g., [24]). 4

- 5 ● *Consistency and completeness of the interface specifications.* Externally dis- 5
6 cernible functional behavior should be precisely what is specified, implying 6
7 bilateral consistency of behavior with respect to the functional specifications. 7
8 That is, the subsystem must do what it is supposed to do, *and nothing else* 8
9 *beyond what is specified*. However, because specifications are inherently in- 9
10 complete, many system failures (in security, reliability, performance, and so on) 10
11 can result from events that occur outside the scope of specifications and thus are 11
12 undetectable by any analyses based on those specifications. 12
- 13 ● *Independence of specification abstractions.* As noted above, abstraction can 13
14 be an enormous aid to composability of specifications, as well as to assur- 14
15 ance proofs. However, it is essential that the details not explicitly represented 15
16 by each abstraction be independent of the details of other abstractions. Oth- 16
17 erwise, composability will most likely be impaired. One elegant example of 17
18 provable composability is seen in the orthogonality theorem of Chander, Dean, 18
19 and Mitchell [9], which provides soundness and completeness proofs for a trust 19
20 management kernel with a clean separation between authorization and struc- 20
21 tured distributed naming. 21
- 22 ● *Timing and synchronization issues.* In general, Lamport-style safety properties 22
23 (i.e., nothing bad happens) compose better than liveness properties (something 23
24 good eventually happens with certainty) [25], but this boundary is blurred by 24
25 the inclusion of timing constraints, which are technically safety properties, but 25
26 generally not composable. It is also blurred by the existence of properties that 26
27 are neither safety nor liveness—such as information flow. Furthermore, time 27
28 (whether real time or relative time) is typically common to different abstrac- 28
29 tions, which is a reason that synchronization and timing constraints can present 29
30 serious impediments to facile composition. For example, see Kopetz [23] on 30
31 composability in the Time-Triggered Architecture. 31
- 32 ● *Explicit state visibility and information hiding.* If a subsystem is stateless (i.e., 32
33 it does not remember any of its own state information from one invocation to 33
34 the next), then it is less likely to have adverse interactions when that subsystem 34
35 is composed with other subsystems—although there are always issues such as 35
36 noncommutativity of operations and interference during concurrent execution. 36
37 In addition, nontrivial recovery, as in selective rollback, may be unnecessary. 37
38 However, statelessness is often not a desirable goal—although stack disciplines 38
39 effectively separate the internal state information from the subsystem itself and 39
40 40

1 simplify composability. Assuming that a subsystem is stateful (i.e., it retains 1
 2 at least some of its own state information from one incarnation to the next), 2
 3 there is a choice between the classical notion of information hiding and ex- 3
 4 plicit external visibility of state information (which tends to make explicit any 4
 5 residues that might impair compositionality). On the other hand, because in- 5
 6 formation hiding typically masks internal state information, it can hinder facile 6
 7 composability if there are any implicitly shared states. However, this should 7
 8 be detectable with sensible specifications and implementation. (For example, 8
 9 pointers, loosely bound aliases, and other indirect references tend to create 9
 10 problems.) Thus, the separation of common stateful entities can greatly facil- 10
 11 itate composition. Information hiding is also very desirable for other reasons, 11
 12 including isolation, security, system integrity, and tamper resistance. 12

13 One interesting historical approach is found in the formal specifications of 13
 14 SRI's Provably Secure Operating System (PSOS [19,38,39]), in which certain 14
 15 state information is hidden but from which the state information that is explic- 15
 16 itly visible at the module interface is derived. Because hidden state information 16
 17 could not be accessed outside of the module (information hiding), it could not 17
 18 be referenced in any other module specification. As a result, there can be no 18
 19 module state residues or other state information that can be accessible to other 19
 20 modules or subsequent invocations of the same module beyond what is explic- 20
 21 itly declared as visible. This greatly increases the composability of modules 21
 22 and the analysis of potential interactions. It also rules out certain characteristic 22
 23 design flaws simply because it is impossible to represent them in the specifi- 23
 24 cations, even accidentally! (Note that bad implementations can introduce bugs 24
 25 that are not definable in specifications.) 25
 26

27 6.1 Other Manifestations of Composition 27

28 As noted at the beginning of this section, composition is not limited only to com- 28
 29 ponents. It has other manifestations as well. 29
 30

- 31 • *Policy composition.* Serious problems can result when different policies are 31
 32 in conflict or otherwise do not compose properly—especially if that lack of 32
 33 composability is not discovered until much later in development. Furthermore, 33
 34 attempting to compose policies often results in emergent properties that are not 34
 35 evident from the constituent policies. For example, see work by Virgil Gligor 35
 36 et al. with respect to the composability of separation-of-duty policies [21] and 36
 37 application-specific security policies [20]. Gligor notes (among other things) 37
 38 that policy composability does not necessarily imply the usefulness of the re- 38
 39 sulting policies, and that existing compositionality criteria are not always re- 39
 40 listic. Preventing denials of service is a particularly thorny policy; besides, 40

1 policies that do not address denials of service are inherently incomplete. Of 1
2 considerable interest is work by Heiko Mantel relating to the general compos- 2
3 ability of secure system policies and components [28] (e.g., flow properties that 3
4 are preserved under refinement [27]). Many past efforts are of particular interest 4
5 to the research community, such as [2,29]. 5

- 6 ● *Protocol composition.* There is also ongoing work on protocol composability— 6
7 for example, see [13]. An interesting research challenge might be to consider 7
8 a particular collection of protocols (e.g., for authentication, encryption, and in- 8
9 tegrity preservation) and prove that they are mutually composable, subject to 9
10 certain constraints; the proofs could also be extended to demonstrating that their 10
11 modular implementations would be composable. 11
- 12 ● *Proof composition.* A book on compositionality of proofs [15] is worth care- 12
13 ful reading for anyone interested in formal verification and high assurance of 13
14 systems. 14
- 15 ● *Certification composition.* Rushby [49] has characterized some of the main is- 15
16 sues relating to the modular certification of an aircraft that is derived from 16
17 separate certification of its components, based on an extension of a formal ver- 17
18 ification approach. The crucial elements involve separation of assumptions and 18
19 guarantees (based on “assume-guarantee reasoning”) into normal and abnormal 19
20 cases. 20
21

22 6.2 Approaches for Predictable Composition 23

24 The following approaches can enhance the likelihood of predictable compositions. 24
25

- 26 ● *Dependency analysis.* In many systems, unrecognized interdependencies among 26
27 different components can hinder composability. Similar comments are relevant 27
28 to contradictory or otherwise incompatible interdependencies among policies, 28
29 models, separately compiled software, and even proofs. Identifying such de- 29
30 pendencies and removing them or otherwise neutralizing them would be a 30
31 considerable aid to composability, 31
- 32 ● *Constrained and guarded dependency strategies.* The principle of constrained 32
33 dependency for integrity is introduced in Section 5. Deterministic linearization 33
34 or other suitable prioritization of intersubsystem dependencies (such as a lat- 34
35 tice ordering) can avoid many adverse dependency problems, such as often 35
36 result from misguided locking strategies and search strategies, compatibility 36
37 mismatches in system upgrades, and unanticipated distributed interactions. For 37
38 example, in Dijkstra’s THE system paper [17], the use of a linearly ordered 38
39 hierarchical locking structure guaranteed that no deadly embraces could oc- 39
40 cur *between* two different layers of abstraction (although in subsequent years a 40

1 deadly embrace was occasionally discovered *within* a particular layer). As another
2 example, Biba's multilevel integrity [6] (MLI) requires in essence that
3 no computational entity (e.g., user, program, process, or data) may depend
4 on any other entities that are deemed less trustworthy (i.e., that are poten-
5 tially less highly trusted) with respect to integrity. In the broadened sense of
6 dependence considered here, the strict lattice ordering of multilevel integrity
7 attributes implied by Biba may be relaxed if any relative untrustworthiness can
8 be masked by creative system architecture or otherwise transcended—as in the
9 trustworthiness-enhancing mechanisms enumerated in [36] as well as other archi-
10 tectural approaches such as isolation kernels and virtual machine monitors.
11 Also, see Abadi et al. [1] for a formalization of dependency.

- 12 ● *Functional consistency among layers of abstraction.* The 1977 Robinson–Levitt
13 paper [45] on hierarchical formal specifications introduced the concept of formal
14 mappings between different layers of functional specifications that repre-
15 sent abstract implementations of each layer as a function of the lower layers.
16 Formal proofs at one layer can be derived by using the mapping functions
17 together with the formal specifications at appropriate layers. The relatively
18 unsung Robinson–Levitt mapping analysis is actually quite far-reaching, and
19 can be used directly to relate properties of a composed system to individ-
20 ual properties of its subsystems. As noted above with respect to correctness
21 and completeness of interface specifications, this approach is of course limited
22 by any incompleteness in the functional specifications and mapping functions.
23 The Robinson–Levitt approach was part of the SRI Hierarchical Development
24 Methodology (HDM) [46] used in the Provably Secure Operating System [19,
25 38,39] project in the 1970s. An extremely impressive new application of this
26 approach in a modern setting has been developed by John Rushby and Rance
27 DeLong [50], which uses the interpretation mechanism of SRI's current formal
28 methods environment (PVS), and applies it to high-assurance separation ker-
29 nels (which explicitly provide both isolation and controlled sharing) as well as
30 virtual-machine architectures. An earlier informal application of explicit inter-
31 layer relationships is found in the analysis of the interlayer dependencies in the
32 Honeywell/Secure Computing Corporation (SCC) LOGical Coprocessor Ker-
33 nel (LOCK) [52]. (PSOS's type-enforcement was the precursor of several SCC
34 systems, including the Sidewinder firewall.)
- 35 ● *Operating system and programming language approaches.* Program modular-
36 ity, recursive and nested procedure-call protocols, clean stack disciplines, and
37 the absence of unintended residues can all greatly enhance composability. Virtu-
38 alized multiprocessing and rigorously enforced virtual machine separation has
39 considerable possibilities in enabling extremely efficient distributed process-
40 ing by abstracting out many of the usual pitfalls, especially when distributed

1 across networked systems. There is an important role for sound programming 1
2 languages that naturally enforce modular separation with abstraction and en- 2
3 capsulation, compilers that efficiently enforce the programming-language mod- 3
4 ularity and strong typing, systems that provide efficient interprocedure and 4
5 interprocess control flow, and optimizing compilers that do not throw out the baby 5
6 with the bathwater (e.g., by prematurely binding entities that need to remain 6
7 separated until later, creating less easily analyzed object code, seriously imped- 7
8 ing debugging, or compromising security separations provided by architectural 8
9 encapsulations and programming languages). (However, well-implemented ag- 9
10 gressive optimizers are less likely to violate security than programmers are.) As 10
11 one example, SPARK (the SPADE Ada Kernel, based on the Southampton Pro- 11
12 gram Analysis Development Environment) provides a language-based approach 12
13 to improving security and safety. Correctness-preserving transformations that 13
14 survive compilation and optimization are another approach with significant 14
15 promise. In particular, optimizing compilers must be fairly farsighted not to 15
16 compromise the integrity of source code in the context of its system execution, 16
17 although careful modularity with abstraction and encapsulation can diminish 17
18 some of those possible effects. An alternative approach to assuring the sound- 18
19 ness of the optimization is the translation validation approach considered at 19
20 NYU [55], in which a validation tool confirms that the object code produced by 20
21 the optimizer is a correct translation of the source code. 21
22

- 23 • *Principled designs, implementations, and use.* As a summary of this section, 23
24 the Saltzer–Schroeder principles and the further principles discussed above are 24
25 potentially extremely beneficial to the attainment of security. Techniques par- 25
26 ticularly relevant to composability include abstraction, hierarchical layering, 26
27 encapsulation, design diversity, composability, pervasive authentication, and ac- 27
28 cess control, as well as administrative and operational controllability, pervasive 28
29 accountability and recovery, separation of policy and mechanism, assignment 29
30 of least privilege, separation of concerns, separation of roles, separation of du- 30
31 ties, and separation of domains. The object-oriented paradigm also has some 31
32 merit, especially strong typing. (However, the would-be inheritance of imple- 32
33 mentations without strict inheritance of specification subclasses tends to impede 33
34 composability. Every subclass instance must meet the specifications of all its 34
35 superclasses, or else all verifications of uses of the superclasses are unsup- 35
36 ported.) 36
37

38
39 Several recent proceedings are worthy of consideration with regard to composable 39
40 system architecture and software engineering [43,22,26,14]. 40

7. A Crisis in Information System Security

Section 4 considers risks in trusting entities that might not actually be trustworthy. Nevertheless, flawed systems that can cause more security and reliability problems than they solve are in widespread use.

Untrustworthy mass-market software might be used so extensively for various reasons, even if the source code is proprietary and the vendor can arbitrarily download questionable software changes without user intervention. Sometimes this is a path of least resistance (with few perceived alternatives) or obliviousness. Or perhaps it has the appearance of saving money *in the short term*. In some cases it is mandated organizationally—ostensibly to simplify procurement, administration, and maintenance, or because of a desire to remain within the monolithic mainstream. Often security, reliability, and the risks of networking are considered less important, or there is a belief that the free market will provide a cure. But the simplest answer is probably “because it’s there.” However, irrespective of any reasons *why* people might be willing to use flawed software, in certain cases it might be wiser *not to use it* at all—especially where the risks are considerable.

In my fourth testimony (August 2001) in five years for committees of the US House of Representatives, I made the following statement—amplifying similar statements made in earlier years:

“Although there have been advances in the research community on information security, trustworthiness, and dependability, the overall situation in practice appears to continually be getting worse, relative to the increasing threats and risks—for a variety of reasons. The information infrastructure is still fundamentally riddled with security vulnerabilities, affecting end-user systems, routers, servers, and communications; new software is typically flawed, and many old flaws still persist; worse yet, patches for residual flaws often introduce new vulnerabilities. There is much greater dependence on the Internet, for Governmental use as well as private and corporate use. Many more systems are being attached to the Internet all over the world, with ever increasing numbers of users—some of whom have decidedly ulterior motives. Because so many systems are so easily interconnectable, the opportunities for exploiting vulnerabilities and the ubiquity of the sources of threats are also increased. Furthermore, even supposedly stand-alone systems are often vulnerable. Consequently, the risks are increasing faster than the amelioration of those risks.”

In many respects, the situation does not seem to be getting better. The continuing flurry of viruses, worms, and system crashes raises the level of disruption to users and institutions. The incessant flow of identified vulnerability reports and the further existence of flaws that are not widely known suggest serious problems. The continual needs for installing copious patches in mass-market software (and the iterative problems they sometimes cause) suggest that we are not converging. Putting

1 the blame on inadequate system administration seems fatuous. Various exploitations 1
2 of flaws (such as worms and viruses) are further examples of endemic problems in 2
3 vulnerable systems that can be exploited. Unfortunately, too many people seem to be 3
4 oblivious to the underlying security problems. 4

5 Suggestions that we need to raise the bar may be countered with the argument that 5
6 past attacks have not really been serious, and we have had few pervasive disasters 6
7 of information system security, so why should we worry? Unfortunately, Murphy's 7
8 Law suggests that if it can happen, it eventually will. Also, the general overem- 8
9 phasis on short-term costs allows long-term concerns to be ignored. (See the next 9
10 section.) 10

11 The Free Software/Open Source movements have been touted as possible al- 11
12 ternatives to the inflexibilities of closed-source proprietary code. Indeed, GNU- 12
13 Linux/BSD Unix variants are gaining considerable credibility, and are seemingly 13
14 less susceptible to malware attacks. However, by itself, availability of source code 14
15 is not a panacea, and sound software engineering is still essential. Even if an en- 15
16 tire system has been subjected to extremely rigorous open evaluation and stringent 16
17 operational controls, that may not be enough to ensure adequate behavior. 17

18 Many users and application developers have grown accustomed to flaky software, 18
19 perhaps because they do not have to meet critical requirements (as in nuclear power 19
20 control, power distribution, and flight and air-traffic control) and suffer no liability 20
21 for disasters. Perhaps it is time to follow the adage of "Just Say No" to bad software, 21
22 and to demand that software development be dramatically improved. 22

23 Many different approaches to software system development can be found in 23
24 practice, such as object-oriented programming, aspect-oriented programming, agile 24
25 software development, service-oriented architecture, design patterns, model-based 25
26 design, event-driven architecture, clean-room development, extreme programming, 26
27 formal methods, a long list of methodologies named after their progenitors, and so 27
28 on. The discipline of these and other approaches can be very helpful, but trustworthi- 28
29 ness demands much more than conventional software. Principled approaches are just 29
30 one more step forward, and need to be coupled with sound development practices. 30
31
32

33 8. Optimistic Optimization 33

34
35 Many people (corporate executives, managers, developers, and so on) tend to ig- 35
36 nore the long-term implications of decisions made for short-term gains, often based 36
37 on overly optimistic or fallacious assumptions. In principle, much greater bene- 37
38 fits can result from far-sighted vision based on realistic assumptions. For example, 38
39 serious environmental effects (including global warming, water and air pollution, 39
40 pesticide toxicity, and adverse genetic engineering) are largely ignored in pursuit 40

1 of short-term profits. However, conservation and environmental protection appear 1
2 much more relevant when considered in the context of long-term costs and benefits. 2
3 Furthermore, governments are besieged by intense short-sighted lobbying by special 3
4 interests. Insider financial manipulations have serious long-term economic effects. 4
5 Research funding has been increasingly focusing on short-term returns, to the detri- 5
6 ment of the future. 6

7 Computer system development is a particularly frustrating example. Most system 7
8 developers are unable or unwilling to confront life-cycle issues up front and in 8
9 the large, although it is clear that up-front investments can yield enormous bene- 9
10 fits later in the life cycle. In particular, defining requirements carefully and wisely 10
11 at the beginning of a development effort can greatly enhance the entire subsequent 11
12 life cycle and reduce its costs. This process should ideally anticipate all essential 12
13 requirements explicitly, including (for example) security, reliability, scalability, and 13
14 relevant application-specific needs such as evolvability, maintainability, usability, 14
15 interoperability, and enterprise survivability. Many such requirements are typically 15
16 extremely difficult to add once system development is well underway. Furthermore, 16
17 certain types of requirements tend to change; thus, system architectures and inter- 17
18 faces should be relatively flaw-free and inherently adaptable without introducing 18
19 further flaws. Insisting on principled software engineering (such as modular abstrac- 19
20 tion, encapsulation, and type safety), sensible use of sound programming languages, 20
21 and use of appropriate support tools can significantly reduce the frequency of soft- 21
22 ware bugs. All these up-front investments can also reduce the subsequent costs of 22
23 debugging, integration, system administration, and long-term evolution—if sensibly 23
24 invoked. 24

25 Consideration of the value of up-front efforts is a decades-old concept. However, 25
26 it is often widely ignored or done badly, for a variety of reasons—such as short-term 26
27 profitability, rush to market, lack of commitment to quality, lack of liability concerns, 27
28 ability to shift late life-cycle costs to customers, inadequate education, experience 28
29 and training, and unwillingness to pursue other than seemingly easy answers. 29

30 Overly optimistic development plans that ignore these issues tend to win out over 30
31 more realistic plans, but can lead to difficulties later on—for developers, system 31
32 users, and even innocent bystanders. The past is littered with systems that did not 32
33 work properly and people who did not perform according to the assumptions embed- 33
34 ded in the development and operational life cycles. (An example is seen in the mad 34
35 rush to low-integrity paperless electronic voting systems with essentially no opera- 35
36 tional accountability, discussed in Section 9.) The lessons of past failures are widely 36
37 ignored. Instead, we have a *caveat emptor* culture, with developers and vendors dis- 37
38 claiming all warranties and liability. 38

39 Many would-be solutions result in part from short-sighted approaches. Firewalls, 39
40 virus checkers, and spam filters all have some benefits, but also some problems. 40

1 Firewalls would be more effective if they were not required to pass all sorts of executable content, such as ActiveX and JavaScript—but many users want those features
2 enabled. (To date, viruses and worms have been rather benign, considering the full
3 potential of really malicious code.) However, active content and malware would be
4 much less harmful in a well-architected environment that could constrain executable
5 content in some sort of “sandbox” that has rigidly limited effects.
6

7 Spammers seem to adapt very rapidly to whatever defenses they encounter. For
8 example, they can test their current offerings against existing anti-spam products and
9 adapt accordingly. Furthermore, domestic legislation may simply drive spammers
10 offshore, without reducing the pain.

11 Better incentives are needed for far-sighted optimization, in larger contexts and
12 over longer periods of time, with realistic assumptions and appropriate architectural
13 flexibility to adapt to changing requirements. Achieving this will require many
14 changes in research and development agendas, software and system development
15 cultures, educational programs, laws, economy, commitment, and perhaps most
16 important—in obtaining well-documented success stories to show the way for others.
17 Particularly in critical applications, if it is not worth doing sensibly, perhaps it is not
18 worth doing at all. But as David Parnas has said, let’s not just preach motherhood;
19 let’s teach people how to be good mothers.
20

21 **9. An Example: Risks in Electronic Voting Systems**

22
23

24 The challenge of ensuring election system integrity provides a paradigmatic example
25 of the considerations of the previous sections. The election process is an
26 end-to-end phenomenon whose trustworthiness typically depends on the integrity
27 of every step in the process. Unfortunately, each of those steps represents various
28 potential weak links that can be compromised in many ways, accidentally and intentionally,
29 technologically or otherwise; each step must be safeguarded from the outset
30 and auditable throughout the entire process.

31 Irregularities reported in the 2000 and 2004 US national elections span the entire
32 process, concerning voter registration, disenfranchisement and harassment of
33 legitimate voters, huge delays in certain precincts, unbalanced distribution of voting
34 equipment, absence of provisional ballots (required by the Help America Vote
35 Act), mishandling of absentee ballots, and problems in casting and counting ballots
36 for e-voting as well as other modes of casting and counting votes. Some machines
37 could not be booted. Some machines lost votes because of programming problems,
38 or recorded more votes than voters. Some touch-screen machines altered the intended
39 vote from one candidate to another. The integrity of the voting technologies themselves
40 is limited by weak evaluation standards, secret evaluations that are paid for

1 by the vendors, all-electronic systems that lack voter-verified audit trails and mean- 1
2 ingful recountability, unaudited post-certification software changes, even runtime 2
3 system or data alterations, and human error and misuse. (Gambling machines are 3
4 held to much higher standards.) Other risks arise from partisan vendors and elec- 4
5 tion officials. Furthermore, statistically significant divergences between exit polls 5
6 and unaudited results created questions in certain states. All these concerns add to 6
7 uncertainties about the integrity of the overall election processes. 7

8 With modern technology, the voting process could be more robust. Whether or not 8
9 the potential weak links are mostly technological, the process can certainly be made 9
10 significantly more trustworthy. Indeed, it seems to be better in many other countries 10
11 than in the US; for example, Ireland, India, and the Netherlands seem to be taking 11
12 integrity challenges seriously. As technologists, we should be helping to ensure that 12
13 is the case—for example, by participating in the standards process or perhaps by 13
14 aiding the cause of available source code and publicly accessible evaluations. How- 14
15 ever, the end-to-end nature of the problem includes many people whose accidental 15
16 or intentional behavior can alter the integrity of the overall process, and thus creates 16
17 many nontechnological risks. 17

18 With respect to computers used in elections, the principles outlined here would 18
19 enable considerable improvements in trustworthiness if they were observed in prac- 19
20 tice. For example, architecturally minimizing the parts of the total system that must 20
21 be trusted would by itself be a huge improvement, thereby reducing the extent of 21
22 the weak links. The same is true of the principle of separating policy and mecha- 22
23 nism. 23

24 The importance of understanding the idiosyncrasies of mechanisms and human 24
25 interfaces, and indeed understanding the entire process, is illustrated by the 2000 25
26 Presidential election—with respect to hanging chad, dimpled chad, uncleaned chad 26
27 slots, butterfly-ballot layouts, and the human procedures underlying voter registration 27
28 and balloting. Clearly, the entire election process has vulnerabilities, including the 28
29 technology and the surrounding administration. Looking into the future, a new ed- 29
30 ucational problem will arise if preferential balloting becomes more widely adopted, 30
31 whereby preferences for competing candidates are prioritized and the votes for the 31
32 lowest-vote candidate are iteratively reallocated according to the specified priorities. 32
33 This concept has many merits, although it certainly further complicates ballot layouts 33
34 and voter awareness! 34

35 Alternative approaches have been proposed to existing voting systems (which have 35
36 typically been lever machines, optically scanned paper, and paperless unaudit- 36
37 able direct-recording computer systems). In approximate order of increasing concep- 37
38 tual complexity, these include (with examples of each) paper-based systems (Ben 38
39 Adida [3], David Chaum [10], Ron Rivest [44]), cryptographic solutions (Andy 39
40 40

1 Neff [31], Josh Benaloh [5]), and voter-verified paper audit trails (VVPATs) (as an
2 add-on for existing all-electronic systems, as proposed by Rebecca Mercuri [30]).

3 The VVPAT approach attempts to overcome the lack of integrity in existing direct-
4 recording systems, but creates further complexity in the process. It is primarily a
5 short-term fix to the current situation, in which proprietary software and proprietary
6 evaluations against inherently incomplete voluntary standards provide relatively little
7 system integrity. Cryptographic approaches require considerable care in design,
8 analysis, implementation, and assurance, but also have the potential to avoid paper
9 records—if the end-to-end systems could be made sufficiently trustworthy. On the
10

11
12 TABLE II
13 APPLICABILITY OF PRINCIPLES TO ELECTIONS

Principle	Computerization	Human Procedures
Economy of mechanism (+ sound architecture)	Simplistic mechanisms are dangerous. Complex systems need extensive analysis and predictable composability.	Operational simplicity is essential for poll workers. Perspicuous risk assessment is desirable throughout.
Fail-safe defaults	can simplify operation and improve trustworthiness.	can mitigate against insider misuse, fraud, and errors.
Complete mediation	can be useful in principled system architectures.	Weakness in depth requires end-to-end oversight.
Open design (+ openness generally)	Proprietary closed-source software and evaluations are inherently suspect.	Diverse oversight is essential throughout the entire process, especially over weak links.
Separation of privileges	can reduce insider misuse, human error, system failures.	can avoid centralized vested control throughout.
Least privilege (+ reduced needs for trust) (+ constrained dependency)	eschews root-privilege misuse, bootload subversion, trusting untrustworthiness. Avoid software built on subvertible underpinnings.	is important throughout the entire process, obviates allocation of excessive trust. Do not trust potentially untrustworthy people.
Least common mechanism	Beware of common flaws and common fault modes.	Separate roles may simplify assurance.
Psychological acceptability	Voter- and official-friendly systems can be helpful.	Ease of use and operation can help if it is not simplistic.
Work factors (+ objective risk analyses)	must encompass all systems, not just limited to strength of cryptography/authentication.	must encompass the entire process end-to-end, including developers and operators.
Compromise recording (+ pervasive monitoring)	Tamper-resistant audit trails are critical whenever results are suspect, and may help disincentivize fraud.	Manual procedures need oversight against compromise from outside/within/below, not just when suspicions arise.

1 other hand, the proposed paper-based systems have considerable conceptual simplic- 1
 2 ity and avoid many of the integrity problems of computer-based systems. However, 2
 3 these approaches address primarily only the vote recording and counting parts of the 3
 4 election process. End-to-end integrity must also include voter registration, voter and 4
 5 vote authentication, and postprocessing. 5

6 **Table II** tersely summarizes the potential relevance of principles (left column) 6
 7 for overall system architectures and development, for both computer-related systems 7
 8 (middle column) and operational procedures (right column) throughout the election 8
 9 process. The table represents a broadening of the Saltzer–Schroeder principles to ad- 9
 10 dress some additional aspects (suggested by what follows the *plus* sign in parentheses 10
 11 in the *principle* column). It thus generalizes the original principles somewhat to in- 11
 12 clude related concepts discussed herein that reach farther than what was originally 12
 13 covered by Saltzer and Schroeder. It also reflects on the fact that these principles are 13
 14 relevant to trustworthiness overall—including (for example) many types of human 14
 15 errors and system failures that are not just limited to security issues. However, it 15
 16 does not remind the reader that this set of principles is only part of what is needed. 16
 17 Ultimately, expertise, experience, and good judgment are essential. 17
 18
 19

20 **10. The Need for Risk Awareness** 20

21
 22 Around the world, our lives are increasingly dependent on technology. What 22
 23 should be the responsibilities of technologists regarding technological and nontech- 23
 24 nological issues? 24

25 • Solving real-world problems often requires technological expertise as well as 25
 26 sufficient understanding of a range of economic, social, political, national, and in- 26
 27 ternational implications. Although it may be natural to want to decouple technology 27
 28 from the other issues, such problems typically cannot be solved by technology alone. 28
 29 They need to be considered in the broader context. 29

30 • Although experts in one area may not be qualified to evaluate detailed would-be 30
 31 solutions in other areas, their own experience may be sufficient to judge the concep- 31
 32 tual merits of such solutions. For example, demonstrable practical impossibility or 32
 33 fundamental limitations of the concept, or the existence of serious conflicts of inter- 33
 34 est of the participants, or an obvious lack of personal and system-wide integrity are 34
 35 causes for concern. 35

36 • Ideally, we need more open, holistic, and interdisciplinary examinations of the 36
 37 underlying problems and their proposed solutions. (For example, see [37].) 37

38 Many concerns arise in important computer-related application areas, such as avi- 38
 39 ation, health care, defense, homeland security, law enforcement and intelligence— 39
 40 with similar conclusions. In each area, a relevant challenge is that of developing 40

1 and operating end-to-end trustworthy environments capable of satisfying stringent 1
2 requirements for human safety, reliability, system integrity, information security, and 2
3 privacy, in which many technological and nontechnological issues must be addressed 3
4 throughout the computer systems and operational practices. Overall, technologists 4
5 need to provide adequate trustworthiness in our socially important information systems 5
6 by technological and other means. Research and development communities interna- 6
7 tionally have much to offer in achieving trustworthy computer-communication 7
8 systems. However, they also have the responsibility of being aware of the other im- 8
9 plications of the use of these systems. 9

10 A deeper knowledge of fundamental principles of computer technology and their 10
11 implications will be increasingly essential in the future, for a wide spectrum of indi- 11
12 viduals and groups, each with its own particular needs. Our lives are becoming ever 12
13 more dependent on understanding computer-related systems and the risks involved. 13
14 Although this may sound like a meta-motherhood statement, wise implementation of 14
15 motherhood is decidedly nontrivial—especially with regard to risks. 15

16 Computer scientists who are active in creating the groundwork for the future need 16
17 to better understand system issues in the large, especially the practical limitations of 17
18 theoretical approaches. System designers and developers need broader and deeper 18
19 knowledge—including those people responsible for the human interfaces used in 19
20 inherently riskful operational environments; interface design is often critical. Partic- 20
21 ularly in those systems that are not wisely conceived and implemented, operators and 21
22 users of the resulting systems also need an understanding of certain fundamentals. 22
23 Corporation executives need an understanding of various risks and countermeasures. 23
24 In each case, knowledge must increase dramatically over time, to reflect rapid evolu- 24
25 tion. Fortunately, the fundamentals do not change as quickly as the widget of the day, 25
26 which suggests that pervasive emphasis on education and ongoing training is needed 26
27 with respect to the concepts of this chapter. 27

28 An alternative view suggests that many technologies can be largely hidden from 28
29 view, and that people need not understand (or indeed, might prefer not to know) 29
30 the inner workings. For example, David Parnas's early papers on abstraction, en- 30
31 capsulation, and information hiding are important in this regard. Although masking 31
32 complexity is certainly possible in theory, in practice we have seen too many occa- 32
33 sions (for examples, see the ACM Risks Forum archives) in which the occurrence of 33
34 inadequately anticipated exceptions resulted in disasters. The complexities arising in 34
35 handling exceptions apply ubiquitously, to defense, medical systems, transportation 35
36 systems, personal finance, security, to our dependence on critical infrastructures that 36
37 can fail—and to anticipating the effects of such exceptions in design, implementa- 37
38 tion, and operation. 38

39 Thus, computer-related education is vital for everyone. The meaning of the Latin 39
40 word “educere” (to educate) is literally “to lead forth.” However, in general, many 40

1 people do not have an adequate perception of the risks and their potential implica- 1
 2 tions. When, for example, the information media tell us that air travel is safer than 2
 3 automobile travel (on a passenger-mile basis, perhaps), the comparison may be less 3
 4 important than the concept that both could be significantly improved. When we are 4
 5 told that electronic commerce is secure and reliable, we need to recognize the cases 5
 6 in which it is not. 6

7 With considerable foresight and wisdom, Vint Cerf has repeatedly said that “The 7
 8 Internet is for Everyone.” The Internet can provide a fertile medium for learning for 8
 9 anyone who wants to learn, but it also creates serious opportunities for the unchecked 9
 10 perpetuation of misinformation and counterproductive learning that will need to be 10
 11 unlearned. 11

12 In general, we learn what is most valuable to us from personal experience, not 12
 13 by being force-fed lowest-common-denominator details. In that spirit, it is impor- 13
 14 tant that education, training, and practical experiences provide motivations for true 14
 15 learning. For technologists, education needs to have a pervasive systems orientation 15
 16 that encompasses concepts of software and system engineering, security, and reli- 16
 17 ability, as well as stressing the importance of suitable human interfaces. For everyone 17
 18 else, there needs to be much better appreciation of the sociotechnical and economic 18
 19 implications—including the risks issues. Above all, a sense of vision of the bigger 19
 20 picture is perhaps what is most needed. 20
 21
 22
 23

24 **11. Risks of Misinformation** 24
 25

26 The problems of online misinformation are evidently worsening, because of the 26
 27 growth of the Internet and our ever increasing dependence on online systems. In- 27
 28 formation technology is a double-edged sword—perhaps even more so than many 28
 29 other technologies. In the hands of enlightened individuals, institutions, and govern- 29
 30 ments, its use can be enormously beneficial. In other hands, it can be detrimental. 30
 31 Unfortunately, the dichotomy is often in the eye of the beholder, perhaps depending 31
 32 on one’s objectives (e.g., personal financial gains, corporate profits, global economic 32
 33 well-being, politics, privacy, and environmental concerns). 33

34 Given a collection of online information, many people behave as if it is inherently 34
 35 authentic and accurate. This myth applies not only to websites, but also to many types 35
 36 of special-purpose databases, such as those found in law enforcement, motor vehicle 36
 37 departments, medicine, insurance, social security, credit information, and homeland 37
 38 security. We have seen many cases in which misinformation (e.g., false flight data, 38
 39 erroneous medical records, undeleted acquittals, or tampered files) has resulted in 39
 40 serious consequences. The same is true of imprecise information (e.g., resulting in 40

1 false arrests, or affecting everyone with a particular name such as “David Nelson” 1
2 who attempts to board an airplane). 2

3 Although an individual can occasionally observe that personal information about 3
4 one’s self is incorrect, more typically such erroneous information is hidden from the 4
5 individual in question, possibly in diversely inaccurate versions. Overall, it is usually 5
6 impossible for one to ensure that all such instances are correct—especially when 6
7 mirrored in unknown sites all over the world. Furthermore, it is difficult to determine 7
8 whether or not online information about anything else is authoritative. Worse yet, the 8
9 volume of questionable information is growing at an extraordinary rate, and attempts 9
10 to update substantive misinformation often have little effect—especially with the 10
11 persistence of incorrect cached versions. 11

12 We increasingly rely on the Internet for many purposes, including education and 12
13 enlightenment, irrespective of whether the sources are accurate. Oft-repeated overly 13
14 simplistic sound-bite mantras seem to be popular. Furthermore, some people seem 14
15 eager to waste time and energy that could be better spent elsewhere—or to drop 15
16 out. There is a tendency for entrenched positions to remain fixed. Are we losing our 16
17 ability to listen openly to other views and engage in constructive thought? 17

18 Another problem involves the inaccessibility of vital information. We seem to 18
19 have evolved into a mentality of “If it is not on the Internet, it does not exist.” Even 19
20 though there are many more data bytes available today than ever before, search en- 20
21 gines reportedly find only a small percentage of those pages, almost none of the 21
22 database-driven dynamic Web pages, and very little of what is in most public li- 22
23 braries. Copyright restrictions and proprietary claims further limit what is available. 23
24 For example, professional society digital libraries tend to be accessible only to those 24
25 members who pay to subscribe. Furthermore, overzealous filtering blocks many 25
26 authoritative sources of information. Are our education and information gathering 26
27 suffering from a lowest-common-denominator process? 27

28 The propagation of misinformation has long been a problem in conventional print 28
29 and broadcast media, but represents another problem that is exacerbated by the speed 29
30 and bandwidth of the Internet. In general, widely held beliefs in supposedly valid 30
31 information tend to take on lives of their own as urban myths; they tend to be trusted 31
32 far beyond what is reasonable, even in the presence of well-based demonstrations of 32
33 their invalidity. 33

34 In the face of such rampant misinformation, the truth can be difficult to accept, 34
35 partly because it can be so difficult to ascertain, partly because it can seem so 35
36 starkly inconsistent with popular misinformation, and partly because people want 36
37 to believe in simple answers. Thus, we are revisiting classical problems that might 37
38 now be considered as E-Epistemology, involving the nature and fundamentals of on- 38
39 line knowledge—especially with reference to its limits and validity. However, there 39
40 are some possible remedies, such as epistemic educational processes that teach us 40

1 how to evaluate information objectively. For websites, this might entail examining 1
 2 who are the sponsors, what affiliations are implied, where the information comes 2
 3 from, whether multiple seemingly reinforcing items all stem from the same incorrect 3
 4 source, whether purported website security and privacy claims are actually justified, 4
 5 and so on. 5
 6

7 8 **12. Boon or Bane?** 8 9

10 Predicting the long-term effects of computers is both difficult and easy: it is easy 10
 11 to predict the future (often mistakenly), but very difficult to be correct. Here are some 11
 12 suggested possible visions of the future. 12

13 • Computers play an increasing role in enabling and mediating communication 13
 14 between people. They have great potential for improving communication, but there is 14
 15 a real risk that they will simply overload us, keeping us from really communicating. 15
 16 We already receive far more information than we can process. A lot of it is noise. 16
 17 Will computers help us to communicate or will they interfere? 17

18 • Computers play an ever-increasing role in our efforts to educate our young. In 18
 19 some countries, educators want to have computers in every school, or even one on 19
 20 every desk. Computers can help in certain kinds of learning, but it takes time to 20
 21 learn the arcane set of conventions that govern their use. Even worse, many children 21
 22 become so immersed in the cartoon world created by computers that they accept it 22
 23 as real, losing interest in other things. Will computers really improve our education, 23
 24 or will children be consumed by them? 24

25 • Computers play an ever increasing role in our war-fighting. Most modern 25
 26 weapon systems depend on computers. Computers also play a central role in military 26
 27 planning and exercises. Perhaps computers will eventually do the fighting and 27
 28 protect human beings. We might even hope that wars would be fought with simulators, 28
 29 not weapons. On the other hand, computers in weapon systems might simply make 29
 30 us more efficient at killing each other and impoverishing ourselves. Will computers 30
 31 result in more slaughter or a safer world? 31

32 • Information processing can help to create and preserve a healthy environment. 32
 33 Computers can help to reduce the energy and resources we expend on such things 33
 34 as transportation and manufacturing, as well as improve the efficiency of buildings 34
 35 and engines. However, they also use energy, and their production and disposal create 35
 36 pollution. They seem to inspire increased consumption, creating what some ancient 36
 37 Chinese philosophers called “artificial desires.” Will computers eventually improve 37
 38 our environment or make it less healthy? 38

39 • By providing us with computational power and good information, computers 39
 40 have the potential to help us think more effectively. On the other hand, bad informa- 40

tion can mislead us, irrelevant information can distract us, and intellectual crutches can cripple our reasoning ability. We may find it easier to surf the Web than to think. Will computers ultimately enhance or reduce our ability to make good decisions?

- Throughout history, many people have tried to eliminate artificial and unneeded distinctions among people. We have begun to learn that everyone has much in common—men and women of all colors, races, and nationalities. Computers have the power to make borders irrelevant, to hide surface differences, and to help us overcome long-standing prejudices. However, they also facilitate the creation of isolated, antisocial groups that may spread hatred and false information. Will computers ultimately improve our understanding of other peoples or lead to more misunderstanding and hatred?

- Computers can help us to grow more food, build more houses, invent better medicines, and satisfy other basic human needs. They can also distract us from our real needs and make us hunger for more computers and more technology, which we then produce at the expense of more essential commodities. Will computers ultimately enrich us or leave us poorer?

- Computers can be used in potentially dangerous systems to make them safer. They can monitor motorists, nuclear plants, and aircraft. They can control medical devices and machinery. Because they do not fatigue and are usually vigilant, they can make our world safer. On the other hand, the software that controls these systems and the people involved may actually be untrustworthy. Bugs are not the exception; they are the norm. Will computers ultimately make us safer or increase our level of risk?

Much of the accumulated wisdom summarized in this chapter is not particularly new. But it is also not widely practiced. Many people are so busy advancing and applying technology that they do not look either back or forward. We should look back to recognize what we have learned about computer-related risks (e.g., [34]). We must look forward to anticipate the future effects of our efforts, including unanticipated combinations of seemingly harmless phenomena. Evidence over the past decades suggests we are not responding adequately to the challenges. Predilections for short-term optimization without regard for long-term costs abound. We must strive to make sure that we maximize the benefits and minimize the harm. Among other things, we must build stronger and more robust computer systems while remaining acutely aware of the risks associated with their use. Perhaps disciplined observance of the content of this chapter can help provide an impetus for the considerable culture change that is required for the development of trustworthy systems, networks, and enterprises in the future.

ACKNOWLEDGEMENTS

This chapter was written in part under National Science Foundation Grant Number 0524111. Sections 5 and 6 are based on a report sponsored by Douglas Maughan when he was the US Defense Advanced Research Projects Agency (DARPA) Program Manager for the Composable High-Assurance Trustworthy System (CHATS) program. Joshua Levy contributed some incisive comments on the final draft.

The author is very grateful to members of the ACM Committee on Computers and Public Policy, who over the years have been extraordinarily constructive in guiding the Inside Risks columns of the *Communications of the ACM*, which provided a springboard for some of the material here. In particular, Peter Denning, Jim Horning, David Parnas, Jerry Saltzer, and Lauren Weinstein have been especially helpful, as has Tom Lambert at *CACM*. In particular, Section 12 is based on a column that appeared in the March 2001 *CACM*, coauthored with David Parnas.

Section 3 was inspired by an article by Tim Batchelder, “An Anthropology of Air”, *Townsend Letter for Doctors and Patients*, pp. 105–106, November 2005. “Because [air] is negative space, it is difficult to see the value in preserving it.”

Lindsay Marshall at Newcastle has been extremely helpful in providing a nicely searchable archive facility for the ACM Risks Forum. (The official archives of all issues of the ACM Risks Forum since its inception in August 1985 can be found at <http://www.risks.org>, which indirections to <http://catless.ncl.ac.uk/Risks/>, and at <ftp://www.sri.com/risks/>.)

Will Tracz has consistently encouraged the author to contribute salient risks-related highlights to the *ACM SIGSOFT Software Engineering Notes*, ever since Will took over the editorship from Neumann’s founding stint (1976–1993).

The citations given herein represent just the tip of an enormous literature iceberg. Additional references relevant to this chapter can easily be gleaned by searching the Web or by browsing my website.

REFERENCES

- [1] Abadi M., Banerjee A., Heintze N., Riecke J.G., “A core calculus of dependency”, in: *POPL '99, Proc. of the 26th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, Texas, January 20–22, 1999*, pp. 147–160.
- [2] Abadi M., Lamport L., “Composing specifications”, in: de Bakker J.W., de Roever W.-P., Rozenberg G. (Eds.), *Stepwise Refinement of Distributed Systems: Models, Formalisms, Correctness, REX Workshop, Mook, The Netherlands, May–June 1989*, in: *Lecture Notes in Comput. Sci.*, vol. 230, Springer-Verlag, Berlin, 1989, pp. 1–41.
- [3] Adida B., Neff C.A., “Ballot casting assurance”, in: *Workshop on Electronic Voting Technology Workshop, Vancouver, BC, Canada, August 2006*, USENIX.

- 1 [4] **An J.H., Dodis Y., Rabin T.**, “On the security of joint signature and encryption”, in: *Advances in Cryptology, EUROCRYPT 2002, Amsterdam, The Netherlands, in: Lecture Notes* 1
2 *in Comput. Sci.*, Springer-Verlag, Berlin, May 2002, pp. 83–107. 2
3
4 [5] **Benaloh J.**, “Simple verifiable elections”, in: *Workshop on Electronic Voting Technology* 3
4 *Workshop, Vancouver, BC, Canada, August 2006*, USENIX. 4
5
6 [6] **Biba K.J.**, “Integrity considerations for secure computer systems”, Technical Report MTR 5
6 3153, The Mitre Corporation, Bedford, Massachusetts, June 1975. Also available from
7 USAF Electronic Systems Division, Bedford, Massachusetts, as ESD-TR-76-372, April
8 1977. 8
9
10 [7] **Bishop M.**, *Computer Security: Art and Science*, Addison–Wesley, Reading, MA, 2002. 9
11 [8] **M. Bishop**, *Introduction to Computer Security*, Addison–Wesley, Reading, MA, 2004. 10
12 [9] **Chander A., Dean D., Mitchell J.C.**, “Reconstructing trust management”, *J. Computer* 11
12 *Security* **12** (1) (January 2004) 131–164. 12
13 [10] **Chaum D.**, “Secret-ballot receipts and transparent integrity: Improving voter con- 13
14 fidence & electronic voting at polling places”, Technical report, March 2002, 14
15 <http://www.chaum.org>. 15
16 [11] **F.J. Corbató**, “On building systems that will fail (1990 Turing Award Lecture, with a 16
17 following interview by Karen Frenkel)”, *Commun. ACM* **34** (9) (September 1991) 72–90. 16
18 [12] **Corbató F.J., Saltzer J., Clingen C.T.**, “Multics: The first seven years”, in: *Proc. of the* 17
18 *Spring Joint Computer Conference, vol. 40, Montvale, New Jersey*, AFIPS Press, 1972. 18
19 [13] **Datta A., Küsters R., Mitchell J.C., Ramanathan A., Shmatikov V.**, “Unifying 19
20 equivalence-based definitions of protocol security”, in: *Proc. of the ACM SIGPLAN and* 20
21 *IFIP WG 1.7 Fourth Workshop on Issues in the Theory of Security, Oakland, California*, 21
22 IEEE Computer Society, April 2004. 22
23 [14] **de Boer F.S.**, et al. (Eds.), *Formal Methods for Components and Objects, 4th Inter-* 23
24 *national Symposium, Lecture Notes in Comput. Sci.*, vol. 4111, Springer-Verlag, Berlin, 24
25 November 2005. 25
26 [15] **de Roeper W.-P., de Boer F., Hanneman U., Hooman J., Lakhnech Y., Poel M., Zwiers J.**, 26
27 *Concurrency Verification: Introduction to Compositional and Noncompositional Methods*, 27
28 *Cambridge Tracts Theoret. Comput. Sci.*, vol. 54, Cambridge University Press, New York, 28
29 NY, 2001. 29
30 [16] **Denning D.E., Neumann P.G., Parker D.B.**, “Social aspects of computer security”, in: 30
31 *Proc. of the 10th National Computer Security Conference*, September 1987. 31
32 [17] **Dijkstra E.W.**, “The structure of the THE multiprogramming system”, *Commun.* 32
33 *ACM* **11** (5) (May 1968). 33
34 [18] **Faughn A.W.**, “Interoperability: Is it achievable?”, Technical report, Harvard University 34
35 PIRP report, 2001. 35
36 [19] **Feiertag R.J., Neumann P.G.**, “The foundations of a Provably Secure Operating System 36
37 (PSOS)”, in: *Proc. of the National Computer Conference*, AFIPS Press, 1979, pp. 329–334, 37
38 <http://www.csl.sri.com/neumann/psos.pdf>. 38
39 [20] **Gligor V.D., Gavrilá S.I.**, “Application-oriented security policies and their composition”, 39
40 in: *Proc. of the 1998 Workshop on Security Paradigms, Cambridge, England*, 1998. 40
41 [21] **Gligor V.D., Gavrilá S.I., Ferraiolo D.**, “On the formal definition of separation-of-duty 41
42 policies and their composition”, in: *Proc. of the 1998 Symposium on Security and Privacy*, 42
43 *Oakland, California*, IEEE Computer Society, May 1998. 43
44

- [22] Gorton I., et al. (Eds.), *Component-Based Software Engineering, 9th International Symposium, Lecture Notes in Comput. Sci.*, vol. 4063, Springer-Verlag, Berlin, June/July 2006.
- [23] Kopetz H., “Composability in the time-triggered architecture”, in: *Proc. of the SAE World Congress, Detroit, Michigan*, SAE Press, 2000, pp. 1–8.
- [24] Krafzig D., Banke K., Slama D., *Enterprise SOA Service-Oriented Architecture Best Practices*, Prentice Hall, Upper Saddle River, NJ, 2004.
- [25] Lamport L., “A simple approach to specifying concurrent program systems”, *Commun. ACM* **32** (1) (January 1989) 32–45.
- [26] Löwe W., et al. (Eds.), *Software Composition, 5th International Workshop, Lecture Notes in Comput. Sci.*, vol. 4089, Springer-Verlag, Berlin, March 2006.
- [27] Mantel H., “Preserving information flow properties under refinement”, in: *Proc. of the 2001 Symposium on Security and Privacy, Oakland, California*, IEEE Computer Society, May 2001, pp. 78–91.
- [28] Mantel H., “On the composition of secure systems”, in: *Proc. of the 2002 Symposium on Security and Privacy, Oakland, California*, IEEE Computer Society, May 2002, pp. 88–101.
- [29] McCullough D., “A hookup theorem for multilevel security”, *IEEE Trans. Software Engineering* **16** (6) (June 1990).
- [30] Mercuri R., “Electronic vote tabulation checks and balances”, PhD thesis, Department of Computer Science, University of Pennsylvania, 2001. <http://www.notablesoftware.com/evote.html>.
- [31] Neff C.A., “A verifiable secret shuffle and its application to e-voting”, in: *Proc. of the ACM Conference on Computer and Communications Security, Philadelphia, Pennsylvania*, November 2001, pp. 116–125.
- [32] Neumann P.G., “Illustrative risks to the public in the use of computer systems and related technology, index to RISKS cases”, Technical report, Computer Science Laboratory, SRI International, Menlo Park, California. Updated regularly at <http://www.csl.sri.com/neumann/illustrative.html> also in .ps and .pdf form for printing in a denser format.
- [33] Neumann P.G., “The role of motherhood in the pop art of system programming”, in: *Proc. of the ACM Second Symposium on Operating Systems Principles, Princeton, New Jersey*, ACM, October 1969, pp. 13–18, <http://www.multicians.org/pgn-motherhood.html>.
- [34] Neumann P.G., *Computer-Related Risks*, ACM Press, New York, and Addison–Wesley, Reading, MA, 1995.
- [35] Neumann P.G., “Practical architectures for survivable systems and networks”, Technical report, Final Report, Phase Two, Project 1688, SRI International, Menlo Park, California, June 2000. <http://www.csl.sri.com/neumann/survivability.html>.
- [36] Neumann P.G., “Principled assuredly trustworthy composable architectures”, Technical report, Computer Science Laboratory, SRI International, Menlo Park, California, December 2004. <http://www.csl.sri.com/neumann/chats4.html>, .pdf and .ps.
- [37] Neumann P.G., “Holistic systems”, *ACM Software Engineering Notes* **31** (6) (November 2006).
- [38] Neumann P.G., Boyer R.S., Feiertag R.J., Levitt K.N., Robinson L., “A provably secure operating system: The system, its applications, and proofs”, Technical report, Computer

- 1 Science Laboratory, SRI International, Menlo Park, California, May 1980, 2nd edition, Re- 1
2 port CSL-116. 2
- 3 [39] Neumann P.G., Feiertag R.J., “PSOS revisited”, in: *Proc. of the 19th Annual Com-* 3
4 *puter Security Applications Conference (ACSAC 2003), Classic Papers section, Las Vegas,* 4
5 *Nevada*, IEEE Computer Society, December 2003, pp. 208–216, <http://www.acsac.org/> and 5
6 <http://www.csl.sri.com/neumann/psos03.pdf>. 6
- 7 [40] Organick E.I., *The Multics System: An Examination of Its Structure*, MIT Press, Cam- 7
8 bridge, MA, 1972. 8
- 9 [41] Parnas D.L., “On the criteria to be used in decomposing systems into modules”, *Commun.* 9
10 *ACM* **15** (12) (December 1972). 10
- 11 [42] Petroski H., *To Engineer Is Human: The Role of Failure in Successful Design*, St. Martin’s 11
12 Press, New York, 1985. 12
- 13 [43] Reussner R.H., et al. (Eds.), *Architecting Systems with Trustworthy Components, Inter-* 13
14 *national Seminar, Dagstuhl, Germany, Lecture Notes in Comput. Sci.*, vol. 3938, Springer- 14
15 Verlag, Berlin, December 2004. 15
- 16 [44] Rivest R.L., “The threeballot voting system”, Technical report, MIT, Cambridge, MA, 16
17 October 2006. 17
- 18 [45] Robinson L., Levitt K.N., “Proof techniques for hierarchically structured programs”, 18
19 *Commun. ACM* **20** (4) (April 1977) 271–283. 19
- 20 [46] Robinson L., Levitt K.N., Silverberg B.A., *The HDM Handbook*, Computer Science Lab- 20
21 oratory, SRI International, Menlo Park, California, June 1979 (three volumes). 21
- 22 [47] Rochlis J.A., Eichin M.W., “With microscope and tweezers: The Worm from MIT’s per- 22
23 spective”, *Commun. ACM* **32** (6) (June 1989) 689–698. 23
- 24 [48] Rosen E., “Vulnerabilities of network control protocols”, *ACM SIGSOFT Software Engi-* 24
25 *neering Notes* **6** (1) (January 1981) 6–8. 25
- 26 [49] Rushby J., “Modular certification”, Technical report, Computer Science Laboratory, SRI 26
27 International, Menlo Park, California, June 2002. 27
- 28 [50] Rushby J.M., DeLong R., “Compositional assurance for security: A MILS integration 28
29 protection profile”, Technical report, Computer Science Laboratory, SRI International, 29
30 Menlo Park, California, December 2006. 30
- 31 [51] Saltzer J.H., Schroeder M.D., “The protection of information in computer systems”, *Proc.* 31
32 *IEEE* **63** (9) (September 1975) 1278–1308. 32
- 33 [52] Saydjari O.S., Beckman J.M., Leaman J.R., “LOCKing computers securely”, in: *10th* 33
34 *National Computer Security Conference, Baltimore, Maryland*, 21–24 September 1987, pp. 34
35 129–141; 35
36 Reprinted in Turn R. (Ed.), *Advances in Computer System Security*, vol. 3, Artech House, 36
37 Dedham, Massachusetts, 1988. 37
- 38 [53] Smith M.A., “Portals: Toward an application framework for interoperability”, *Commun.* 38
39 *ACM* **47** (10) (October 2004) 93–97. 39
- 40 [54] Spafford E.H., “The Internet Worm: Crisis and aftermath”, *Commun. ACM* **32** (6) (June 40
1989) 678–687. 40
- [55] Zuck L., Pnueli A., Fang Y., Goldberg B., “VOC: A translation validator for optimizing compilers”, in: *Electronic Notes in Theoretical Computer Science*, 2002. 40