

This book is provided *FREE* with
test registration by the
Graduate Record Examinations Board.

**This practice book
contains**

- one actual full-length GRE
Computer Science Test
- test-taking strategies

Become familiar with

- test structure and content
- test instructions and
answering procedures

Compare your practice
test results with the
performance of those
who took the test at a
GRE administration.

GRE[®]

Graduate Record Examinations[®]

COMPUTER SCIENCE TEST

PRACTICE BOOK



*Educational
Testing Service*

Visit GRE Online at www.gre.org

Note to Test Takers: Keep this practice book until you receive your score report.
The book contains important information about content specifications and scoring.

Table of Contents

Purpose of the GRE Subject Tests	3
Development of the Subject Tests	3
Content of the Computer Science Test	4
Preparing for a Subject Test	5
Test-Taking Strategies	6
What Your Scores Mean	6
Practice GRE Computer Science Test	9
Scoring Your Subject Test	49
Evaluating Your Performance	52
Answer Sheet	53

Purpose of the GRE Subject Tests

The GRE Subject Tests are designed to help graduate school admission committees and fellowship sponsors assess the qualifications of applicants in specific fields of study. The tests also provide you with an assessment of your own qualifications.

Scores on the tests are intended to indicate knowledge of the subject matter emphasized in many undergraduate programs as preparation for graduate study. Because past achievement is usually a good indicator of future performance, the scores are helpful in predicting success in graduate study. Because the tests are standardized, the test scores permit comparison of students from different institutions with different undergraduate programs. For some Subject Tests, subscores are provided in addition to the total score; these subscores indicate the strengths and weaknesses of your preparation, and they may help you plan future studies.

The GRE Board recommends that scores on the Subject Tests be considered in conjunction with other relevant information about applicants. Because numerous factors influence success in graduate school, reliance on a single measure to predict success is not advisable. Other indicators of competence typically include undergraduate transcripts showing courses taken and grades earned, letters of recommendation, the GRE Writing Assessment score, and GRE General Test scores. For information about the appropriate use of GRE scores, write to GRE Program, Educational Testing Service, Mail Stop 57-L, Princeton, NJ 08541, or visit our Web site at www.gre.org/codelst.html.

Development of the Subject Tests

Each new edition of a Subject Test is developed by a committee of examiners composed of professors in the subject who are on undergraduate and graduate faculties in different types of institutions and in different regions of the United States and Canada. In selecting members for each committee, the GRE Program seeks the advice of the appropriate professional associations in the subject.

The content and scope of each test are specified and reviewed periodically by the committee of examiners. Test questions are written by the committee and by other faculty who are also subject-matter specialists and by subject-matter specialists at ETS. All questions proposed for the test are reviewed by the committee and revised as necessary. The accepted questions are assembled into a test in accordance with the content specifications developed by the committee to ensure adequate coverage of the various aspects of the field and, at the same time, to prevent overemphasis on any single topic. The entire test is then reviewed and approved by the committee.

Subject-matter and measurement specialists on the ETS staff assist the committee, providing information and advice about methods of test construction and helping to prepare the questions and assemble the test. In addition, each test question is reviewed to eliminate language, symbols, or content considered potentially offensive, inappropriate for major subgroups of the test-taking population, or likely to perpetuate any negative attitude that may be conveyed to these subgroups. The test as a whole is also reviewed to ensure that the test questions, where applicable, include an appropriate balance of people in different groups and different roles.

Because of the diversity of undergraduate curricula, it is not possible for a single test to cover all the material you may have studied. The examiners, therefore, select questions that test the basic knowledge and skills most important for successful graduate study in the particular field. The committee keeps the test up-to-date by regularly developing new editions and revising existing editions. In this way, the test content changes steadily but gradually, much like most curricula. In addition, curriculum surveys are conducted periodically to ensure that the content of a test reflects what is currently being taught in the undergraduate curriculum.

After a new edition of a Subject Test is first administered, examinees' responses to each test question are analyzed in a variety of ways to determine whether each question functioned as expected. These analyses may reveal that a question is ambiguous, requires knowledge beyond the scope of the test, or is inappropriate for the total group or a particular subgroup of examinees taking the test. Answers to such questions are not used in computing scores.

Following this analysis, the new test edition is equated to an existing test edition. In the equating process, statistical methods are used to assess the difficulty of the new test. Then scores are adjusted so that examinees who took a difficult edition of the test are not penalized, and examinees who took an easier edition of the test do not have an advantage. Variations in the number of questions in the different editions of the test are also taken into account in this process.

Scores on the Subject Tests are reported as three-digit scaled scores with the third digit always zero. The maximum possible range for all Subject Test total scores is from 200 to 990. The actual range of scores for a particular Subject Test, however, may be smaller. The maximum possible range of Subject Test subscores is 20 to 99; however, the actual range of subscores for any test or test edition may be smaller than 20 to 99. Subject Test score interpretive information is provided in *Interpreting Your GRE Scores*, which you will receive with your GRE score report, and on the GRE Web site at www.gre.org/codelst.html.

Content of the Computer Science Subject Test

The test consists of about 70 multiple-choice questions, some of which are grouped in sets and based on such materials as diagrams, graphs, and program fragments.

The approximate distribution of questions in each edition of the test according to content categories is indicated by the following outline. The percentages given are approximate; actual percentages will vary slightly from one edition of the test to another.

- I. SOFTWARE SYSTEMS AND METHODOLOGY — 35%
 - A. Data organization
 1. Data types
 2. Data structures and implementation techniques
 3. File organization (e.g., sequential, indexed, multilevel)
 - B. Program control
 1. Iteration and recursion
 2. Functions, procedures, and exception handlers
 3. Communication and synchronization
 - C. Programming languages and notation
 1. Constructs for data organization and program control
 2. Scope, binding, and parameter passing
 3. Expression evaluation

- D. Systems
 1. Compilers and interpreters
 2. Operating systems, including resource management and protection/security
 3. Networking and distributed systems
 4. System development tools
 5. System performance

II. COMPUTER ORGANIZATION AND ARCHITECTURE — 20%

- A. Logic design
 1. Implementation of combinational and sequential circuits
 2. Functional properties of digital integrated circuits
- B. Processors and control units
 1. Instruction sets
 2. Register and ALU organization
 3. Number representation
 4. Control sequencing
 5. Data paths
- C. Memories and their hierarchies
 1. Speed, capacity, cost, allocation
 2. Cache, main, secondary storage
 3. Virtual memory, paging, segmentation
- D. Communication
 1. Bus, switch, and network structures and protocols
 2. I/O
 3. Synchronization
- E. High-performance architectures
 1. Pipelining super-scalar and out-of-order execution processors
 2. Parallel computing
 3. Distributed computing

III. THEORY — 25%

- A. Automata and language theory
 1. Models of computation (finite automata, pushdown automata, Turing machines)
 2. Formal languages (regular languages, context-free languages)
 3. Decidability
- B. Design and analysis of algorithms and computational complexity
 1. Exact or asymptotic analysis of the best, worst, or average case for the time and space complexity of specific algorithms

- 2. Algorithmic design techniques (divide and conquer, dynamic programming, greedy)
- 3. Upper and lower bounds on the complexity of specific problems
- 4. NP-completeness
- C. Correctness of programs
 1. Formal specifications and assertions
 2. Verification techniques

IV. MATHEMATICAL BACKGROUND — 15%

- A. Discrete structures
 1. Mathematical logic
 2. Elementary combinatorics, including graph theory and counting arguments
 3. Elementary discrete mathematics, including number theory, discrete probability, recurrence relations
- B. Numerical mathematics
 1. Computer arithmetic, including number representations, roundoff errors, overflow and underflow
 2. Classical numerical algorithms
 3. Linear algebra

V. OTHER TOPICS — 5%

Topics including modeling and simulation, information retrieval, artificial intelligence, computer graphics, data communications, databases, VLSI, logic programming.

Preparing for a Subject Test

GRE Subject Test questions are designed to measure skills and knowledge gained over a long period of time. Although you might increase your scores to some extent through preparation a few weeks or months before you take the test, last-minute cramming is unlikely to be of further help. The following information may be helpful.

- A general review of your college courses is probably the best preparation for the test. However, the test covers a broad range of subject matter, and no one is expected to be familiar with the content of every question.

- Use this practice book to become familiar with the types of questions in the GRE Computer Science Test, paying special attention to the directions. If you thoroughly understand the directions before you take the test, you will have more time during the test to focus on the questions themselves.

Test-Taking Strategies

The questions in the practice test in this book illustrate the types of multiple-choice questions in the test. When you take the test, you will mark your answers on a separate machine-scorable answer sheet. Total testing time is two hours and fifty minutes; there are no separately timed sections. Following are some general test-taking strategies you may want to consider.

- Read the test directions carefully, and work as rapidly as you can without being careless. For each question, choose the best answer from the available options.
- All questions are of equal value; do not waste time pondering individual questions you find extremely difficult or unfamiliar.
- You may want to work through the test quite rapidly, first answering only the questions about which you feel confident, then going back and answering questions that require more thought, and concluding with the most difficult questions if there is time.
- If you decide to change an answer, make sure you completely erase it and fill in the oval corresponding to your desired answer.
- Questions for which you mark no answer or more than one answer are not counted in scoring.
- As a correction for haphazard guessing, one-fourth of the number of questions you answer incorrectly is subtracted from the number of questions you answer correctly. It is improbable that mere guessing will improve your score significantly; it may even lower your score. If, however, you are not certain of the correct answer but have some knowledge of the question

and are able to eliminate one or more of the answer choices, your chance of getting the right answer is improved, and it may be to your advantage to answer the question.

- Record all answers on your answer sheet. Answers recorded in your test book will not be counted.
- Do not wait until the last five minutes of a testing session to record answers on your answer sheet.

What Your Scores Mean

Your raw score—that is, the number of questions you answered correctly minus one-fourth of the number you answered incorrectly—is converted to the scaled score that is reported. This conversion ensures that a scaled score reported for any edition of a Subject Test is comparable to the same scaled score earned on any other edition of the same test. Thus, equal scaled scores on a particular Subject Test indicate essentially equal levels of performance regardless of the test edition taken. Test scores should be compared only with other scores on the same Subject Test. (For example, a 680 on the Computer Science Test is not equivalent to a 680 on the Mathematics Test.)

Before taking the test, you may find it useful to know approximately what raw scores would be required to obtain a certain scaled score. Several factors influence the conversion of your raw score to your scaled score, such as the difficulty of the test edition and the number of test questions included in the computation of your raw score. Based on recent editions of the Computer Science Test, the table on the next page gives the range of raw scores associated with selected scaled scores for three different test editions. (Note that when the number of scored questions for a given test is greater than the range of possible scaled scores, it is likely that two or more raw scores will convert to the same scaled score.) The three test editions in the table that follows were selected to reflect varying degrees of difficulty. Examinees should note that future test editions may be somewhat more or less difficult than the test editions illustrated in the table.

**Range of Raw Scores* Needed to
Earn Selected Scaled Scores on
Three Computer Science Test Editions
That Differ in Difficulty**

Scaled Score	Raw Scores		
	Form A	Form B	Form C
800	55	48-49	45
700	39	31	28-29
600	23	19	16
500	7	6-7	4
Number of Questions Used to Compute Raw Score			
	70	70	69

*Raw Score = Number of correct answers minus one-fourth the number of incorrect answers, rounded to the nearest integer.

For a particular test edition, there are many ways to earn the same raw score. For example, on the edition listed above as “Form A,” a raw score of 39 would earn a scaled score of 700. Below are a few of the possible ways in which a scaled score of 700 could be earned on that edition.

**Examples of Ways to Earn a Scaled Score
of 700 on the Edition Labeled as “Form A”**

Raw Score	Questions Answered Correctly	Questions Answered Incorrectly	Questions Not Answered	Number of Questions Used to Compute Raw Score
39	39	0	31	70
39	42	11	17	70
39	45	23	2	70

Practice Test

To become familiar with how the administration will be conducted at the test center, first remove the answer sheet (pages 53 and 54). Then go to the back cover of the test book (page 48) and follow the instructions for completing the identification areas of the answer sheet. When you are ready to begin the test, note the time and begin marking your answers on the answer sheet.

THE GRADUATE RECORD
EXAMINATIONS®

GRE®



COMPUTER SCIENCE TEST

*Do not break the seal
until you are told to do so.*

*The contents of this test are confidential.
Disclosure or reproduction of any portion
of it is prohibited.*

THIS TEST BOOK MUST NOT BE TAKEN FROM THE ROOM.

GRE, GRADUATE RECORD EXAMINATIONS, ETS, EDUCATIONAL TESTING SERVICE and the ETS logo are registered trademarks of Educational Testing Service.

Copyright © 1995, 1996 by Educational Testing Service. All rights reserved.

Princeton, N.J. 08541

COMPUTER SCIENCE TEST

Time—170 minutes

70 Questions

Notation and Conventions:

In this test a reading knowledge of modern programming languages is assumed. The following notational conventions are used.

1. All numbers are assumed to be written in decimal notation unless otherwise indicated.
2. $\lfloor x \rfloor$ denotes the greatest integer that is less than or equal to x .
3. $\lceil x \rceil$ denotes the least integer that is greater than or equal to x .
4. $g(n) = O(f(n))$ denotes “ $g(n)$ has order at most $f(n)$ ” and means that there exist positive constants C and N such that $|g(n)| \leq Cf(n)$ for all $n > N$.
 $g(n) = \Omega(f(n))$ denotes “ $g(n)$ has order at least $f(n)$ ” and for this test means that there exist positive constants C and N such that $g(n) \geq Cf(n)$ for all $n > N$.
 $g(n) = \Theta(f(n))$ denotes “ $g(n)$ has the same order as $f(n)$ ” and means that there exist positive constants C_1 , C_2 , and N such that $C_1f(n) \leq g(n) \leq C_2f(n)$ for all $n > N$.
5. \exists denotes “there exists.”
 \forall denotes “for all.”
 \Rightarrow denotes “implies.”
 \neg denotes “not”; “ \bar{A} ” is also used as meaning “ $\neg A$.”
 \vee denotes “inclusive or”; $+$ also denotes “inclusive or.”
 \oplus denotes “exclusive or.”
 \wedge denotes “and”; also, juxtaposition of statements denotes “and,” e.g., PQ denotes “ P and Q .”
 \emptyset denotes the empty set.
6. If A and B denote sets, then:
 $A \cup B$ is the set of all elements that are in A or in B or in both;
 $A \cap B$ is the set of all elements that are in both A and B ; AB also denotes $A \cap B$;
 \bar{A} is the set of all elements not in A that are in some specified universal set.
7. In a string expression, if S and T denote strings or sets of strings, then:
An empty string is denoted by ϵ or by Λ ;
 ST denotes the concatenation of S and T ;
 $S + T$ denotes $S \cup T$ or $\{S, T\}$, depending on context;
 S^n denotes $\underbrace{SS \dots S}_n$;
 S^* denotes $\epsilon + S + S^2 + S^3 + \dots$;
 S^+ denotes $S + S^2 + S^3 + \dots$.

GO ON TO THE NEXT PAGE.

8. In a grammar:

$\alpha \rightarrow \beta$ represents a production in the grammar.


$\alpha \Rightarrow \beta$ means β can be derived from α by the application of exactly one production.


$\alpha \Rightarrow^* \beta$ means β can be derived from α by the application of zero or more productions.


Unless otherwise specified,


- (i) symbols appearing on the left-hand side of productions are nonterminal symbols; the remaining symbols are terminal symbols;
- (ii) the leftmost symbol of the first production is the start symbol;
- (iii) the start symbol is permitted to appear on the right-hand side of productions.


9. In a logic diagram:


 represents an AND element.


 represents an inclusive OR element.

 represents an exclusive OR element.

 represents a NOT element.

 represents a NAND element.

 represents a NOR element.

10.  represents a *D*-type flip-flop, which stores the value of its *D* input when clocked.

11. Binary tree traversal is defined recursively as follows:

preorder - visit the root, traverse the left subtree, traverse the right subtree

inorder - traverse the left subtree, visit the root, traverse the right subtree

postorder - traverse the left subtree, traverse the right subtree, visit the root

12. In a finite automaton diagram, states are represented by circles, with final (or accepting) states indicated by two concentric circles. The start state is indicated by the word "Start." An arc from state *s* to state *t* labeled *a* indicates a transition from *s* to *t* on input *a*. A label *a/b* indicates that this transition produces an output *b*. A label a_1, a_2, \dots, a_k indicates that the transition is made on any of the inputs a_1, a_2, \dots, a_k .

13. For a program segment *S* and predicates *P* and *Q*, a triple $\{P\}S\{Q\}$ is partially correct if, whenever *P* is true before initiation of *S*, *Q* is true upon termination of *S*. $\{P\}S\{Q\}$ is totally correct if it is partially correct and terminates for all inputs.

Given that $\{P\}S\{Q\}$ is partially correct, a precondition is any assertion that implies *P*, and a postcondition is any assertion that is implied by *Q*.

GO ON TO THE NEXT PAGE.

Directions: Each of the questions or incomplete statements below is followed by five suggested answers or completions. Select the one that is the best of the choices offered and then mark the corresponding space on the answer sheet.

1. Any set of Boolean operators that is sufficient to represent all Boolean expressions is said to be complete. Which of the following is NOT complete?

(A) {AND, NOT} (B) {NOT, OR} (C) {AND, OR} (D) {NAND} (E) {NOR}

2. Which of the following decimal numbers has an exact representation in binary notation?

(A) 0.1 (B) 0.2 (C) 0.3 (D) 0.4 (E) 0.5

3. Bob writes down a number between 1 and 1,000. Mary must identify that number by asking "yes/no" questions of Bob. Mary knows that Bob always tells the truth. If Mary uses an optimal strategy, then she will determine the answer at the end of exactly how many questions in the worst case?


(A) 1,000 (B) 999 (C) 500 (D) 32 (E) 10

GO ON TO THE NEXT PAGE.

4. If x is a string, then x^R denotes the reversal of x . If x and y are strings, then $(xy)^R =$

- (A) xy^R (B) yx^R (C) y^Rx (D) x^Ry^R (E) y^Rx^R
-

5. A procedure that printed the binary tree








in postorder would produce as output

- (A) A B C D E F (B) A B D E C F (C) D B E A F C (D) D E B F C A (E) D E F B C A
-

GO ON TO THE NEXT PAGE.

6. Which of the following is NOT a binary search tree?



GO ON TO THE NEXT PAGE.

7. Consider the following Pascal-like program fragment.

```
var i, j : integer ;  
procedure P(k, m : integer) ;  
begin  
  k := k - m ;  
  m := k + m ;  
  k := m - k  
end ;  
i := 2 ;  
j := 3 ;  
P(i, j) ;
```

If both parameters to P are passed by reference, what are the values of i and j at the end of the program fragment?


- (A) $i = 0, j = 2$ (B) $i = 1, j = 5$ (C) $i = 2, j = 3$ (D) $i = 3, j = 2$
(E) None of the above
-

8. A starvation-free job-scheduling policy guarantees that no job waits indefinitely for service. Which of the following job-scheduling policies is starvation-free?

- (A) Round-robin (B) Priority queuing (C) Shortest job first (D) Youngest job first
(E) None of the above
-

GO ON TO THE NEXT PAGE.

9. Consider a singly linked list of the form



where F is a pointer to the first element in the list and L is a pointer to the last element in the list. The time of which of the following operations depends on the length of the list?

- (A) Delete the last element of the list.
 - (B) Delete the first element of the list.
 - (C) Add an element after the last element of the list.
 - (D) Add an element before the first element of the list.
 - (E) Interchange the first two elements of the list.
-

10.


```
p := 1; k := 0 ;
while k < n do
begin
  p := 2 * p ;
  k := k + 1
end ;
```

For the program fragment above involving integers p , k , and n , which of the following is a loop invariant; i.e., true at the beginning of each execution of the loop and at the completion of the loop?

- (A) $p = k + 1$ (B) $p = (k + 1)^2$ (C) $p = (k + 1)2^k$ (D) $p = 2^k$ (E) $p = 2^{k+1}$
-

GO ON TO THE NEXT PAGE.

11. Consider an output-producing, deterministic finite state automaton (DFA) of the kind indicated in the figure below, in which it is assumed that every state is a final state.



Assume that the input is at least four bits long. Which of the following is(are) true?

- I. The last bit of the output depends on the start state.
- II. If the input ends with "1100", then the output must end with "1".
- III. The output cannot end with "1" unless the input ends with "1100".

- (A) I only
- (B) II only
- (C) I and II only
- (D) II and III only
- (E) I, II, and III

GO ON TO THE NEXT PAGE.

12. A particular BNF definition for a “word” is given by the following rules.

```
<word> ::= <letter> | <letter><pairlet> | <letter><pairdig>
<pairlet> ::= <letter><letter> | <pairlet><letter><letter>
<pairdig> ::= <digit><digit> | <pairdig><digit><digit>
<letter> ::= a|b|c|. . .|y|z
<digit> ::= 0|1|2|. . .|9
```

Which of the following lexical entities can be derived from <word> ?

- I. word
 - II. words
 - III. c22
- (A) None
(B) I and II only
(C) I and III only
(D) II and III only
(E) I, II, and III

GO ON TO THE NEXT PAGE.

Questions 13-14 relate to the following C-like program.

```
#include <stdio.h>
main()
{
    float sum = 0.0, j = 1.0, i = 2.0;
    while (i/j > 0.001)
    {
        j = j + j;
        sum = sum + i/j;
        printf("%f \n", sum);
    }
}
```

13. How many lines of output does the program produce?

- (A) 0-9 (B) 10-19 (C) 20-29 (D) 30-39 (E) More than 39

14. Which of the following is the integer that best approximates the last number printed?

- (A) 0 (B) 1 (C) 2 (D) 3 (E) 4
-

15. An integer c is a common divisor of two integers x and y if and only if c is a divisor of x and c is a divisor of y . Which of the following sets of integers could possibly be the set of all common divisors of two integers?

- (A) $\{-6, -2, -1, 1, 2, 6\}$
(B) $\{-6, -2, -1, 0, 1, 2, 6\}$
(C) $\{-6, -3, -2, -1, 1, 2, 3, 6\}$
(D) $\{-6, -3, -2, -1, 0, 1, 2, 3, 6\}$
(E) $\{-6, -4, -3, -2, -1, 1, 2, 3, 4, 6\}$
-

GO ON TO THE NEXT PAGE.

16. Consider the following grammar.

$S ::= AB$
 $A ::= a$
 $A ::= BaB$
 $B ::= bbA$

Which of the following is FALSE?

- (A) The length of every string produced by the grammar is even.
 - (B) No string produced by the grammar has an odd number of consecutive b's.
 - (C) No string produced by the grammar has three consecutive a's.
 - (D) No string produced by the grammar has four consecutive b's.
 - (E) Every string produced by the grammar has at least as many b's as a's.
-

17. A particular parallel program computation requires 100 seconds when executed on a single processor. If 40 percent of this computation is "inherently sequential" (i.e., will not benefit from additional processors), then the theoretically best possible elapsed times for this program running with 2 and 4 processors, respectively, are

- (A) 20 and 10 seconds
 - (B) 30 and 15 seconds
 - (C) 50 and 25 seconds
 - (D) 70 and 55 seconds
 - (E) 80 and 70 seconds
-

GO ON TO THE NEXT PAGE.

18.

Virtual Page	Actual Page
0	3
1	—
2	4
3	0

The lookup page table shown above is for a job in a paged virtual storage system with a page size of 1,024 locations. Each virtual address is in the form $[p, d]$ where p and d are the page number and the displacement in that page, respectively.

A virtual address of $[0, 514]$ maps to an actual address of

- (A) 514 (B) 1024 (C) 3586 (D) 4514
(E) none of the above

19. If the expression $((2 + 3) * 4 + 5 * (6 + 7) * 8) + 9$ is evaluated with $*$ having precedence over $+$, then the value obtained is the same as the value of which of the following prefix expressions?

- (A) $+ + * + 2 3 4 * * 5 + 6 7 8 9$
(B) $+ * + + 2 3 4 * * 5 + 6 7 8 9$
(C) $* + + 2 3 4 * * 5 + + 6 7 8 9$
(D) $* + + + 2 3 4 * * 5 + 6 7 8 9$
(E) $+ * + * 2 3 4 + + 5 * 6 7 8 9$

20. Let P be a procedure that for some inputs calls itself (i.e., is recursive). If P is guaranteed to terminate, which of the following statements must be true?

- I. P has a local variable.
 - II. P has an execution path where it does not call itself.
 - III. P either refers to a global variable or has at least one parameter.
- (A) I only
(B) II only
(C) I and II only
(D) II and III only
(E) I, II, and III

GO ON TO THE NEXT PAGE.

21. Consider a computer design in which multiple processors, each with a private cache memory, share global memory using a single bus. This bus is the critical system resource.

Each processor can execute one instruction every 500 nanoseconds as long as memory references are satisfied by its local cache. When a cache miss occurs, the processor is delayed for an additional 2,000 nanoseconds. During half of this additional delay, the bus is dedicated to serving the cache miss. During the other half, the processor cannot continue, but the bus is free to service requests from other processors. On average, each instruction requires 2 memory references. On average, cache misses occur on 1 percent of references.

What proportion of the capacity of the bus would a single processor consume, ignoring delays due to competition from other processors?

- (A) $\frac{1}{50}$ (B) $\frac{1}{27}$ (C) $\frac{1}{25}$ (D) $\frac{2}{27}$ (E) $\frac{1}{5}$
-

22. In multiprogrammed systems it is advantageous if some programs such as editors and compilers can be shared by several users.

Which of the following must be true of multiprogrammed systems in order that a single copy of a program can be shared by several users?

- I. The program is a macro.
- II. The program is recursive.
- III. The program is reentrant.

- (A) I only
(B) II only
(C) III only
(D) II and III only
(E) I, II, and III
-

23. A particular disk unit uses a bit string to record the occupancy or vacancy of its tracks, with 0 denoting vacant and 1 denoting occupied. A 32-bit segment of this string has the hexadecimal value D4FE2003. The percentage of occupied tracks for the corresponding part of the disk, to the nearest percent, is


- (A) 12% (B) 25% (C) 38% (D) 44% (E) 62%
-

GO ON TO THE NEXT PAGE.

24. A program that checks spelling works in the following way. A hash table has been defined in which each entry is a Boolean variable initialized to *false*. A hash function has been applied to each word in the dictionary, and the appropriate entry in the hash table has been set to *true*. To check the spelling in a document, the hash function is applied to every word in the document, and the appropriate entry in the hash table is examined. Which of the following is (are) correct?
- I. *true* means the word was in the dictionary.
 - II. *false* means the word was not in the dictionary.
 - III. Hash table size should increase with document size.
- (A) I only
 - (B) II only
 - (C) I and II only
 - (D) II and III only
 - (E) I, II, and III

GO ON TO THE NEXT PAGE.

25.



The finite automaton above recognizes a set of strings of length 6. What is the total number of strings in the set?

- (A) 18 (B) 20 (C) 30 (D) 32
 (E) None of the above

26. Let S be the statement:

```
for i := 1 to N do V[i] := V[i] + 1
```

Which of the following perform(s) the same changes to V as S ?

- I. $i := 0$;
 while $i \leq N$ do
 begin $i := i + 1$; $V[i] := V[i] + 1$ end
- II. $i := 1$;
 while $i < N$ do
 begin $V[i] := V[i] + 1$; $i := i + 1$ end
- III. $i := 0$;
 while $i < N$ do
 begin $V[i + 1] := V[i + 1] + 1$; $i := i + 1$ end

- (A) I only
 (B) II only
 (C) III only
 (D) II and III only
 (E) I, II, and III


GO ON TO THE NEXT PAGE.

27. `var i, j, x : integer ;
 read(x) ;
 i := 1 ; j := 1 ;
 while i < 10 do
 begin
 j := j * i ;
 i := i + 1 ;
 if i = x then exit
 end`

For the program fragment above, which of the following statements about the variables i and j must be true after execution of the fragment?

- (A) $(j = (x - 1)!) \wedge (i \geq x)$
- (B) $(j = 9!) \wedge (i = 10)$
- (C) $(j = 10!) \wedge (i = 10)$
- (D) $((j = 10!) \wedge (i = 10)) \vee ((j = (x - 1)!) \wedge (i = x))$
- (E) $((j = 9!) \wedge (i \geq 10)) \vee ((j = (x - 1)!) \wedge (i = x))$

28.



State 0 is both the starting state and the accepting state.

Each of the following is a regular expression that denotes a subset of the language recognized by the automaton above EXCEPT

- (A) $0^*(11)^*0^*$
- (B) $0^*1(10^*1)^*1$
- (C) $0^*1(10^*1)^*10^*$
- (D) $0^*1(10^*1)0(100)^*$
- (E) $(0^*1(10^*1)^*10^* + 0^*)^*$

GO ON TO THE NEXT PAGE.

Questions 29-30 are based on the following program heading and type declaration written in a Pascal-like language with garbage collection.

```
program Main(input, output);
type
  Link = ↑Cell;
  Cell = record
    Info : integer;
    Next : Link
  end;
```

29. If the rest of the program is as follows,

```
var
  p1, p2 : Link;
procedure A;
var
  p3 : Link;
begin
(ii)   new(p2);
(iii)  new(p3);
(iv)   new(p3↑.Next)
end;
begin {Main}
(i)    new(p1);
      A;
(v)    ...
end.
```

then, at line (v), which of the allocated *Cell*s is (are) available for reclamation by the garbage collector?

- (A) None
- (B) Only the one allocated by (iii)
- (C) Only the ones allocated by (iii) and (iv)
- (D) Only the ones allocated by (ii), (iii), and (iv)
- (E) Those allocated by (i), (ii), (iii), and (iv)

GO ON TO THE NEXT PAGE.

30. In another program written in a Pascal-like language with the same heading and type declaration, assume that a choice must be made between the following two options.

Option *P*: the procedure *Push* defined by

```
procedure Push(n : integer ; h : Link) ;
var
  p : Link ;
begin
  new(p) ;
  p↑.Info := n ;
  p↑.Next := h ;
  h := p
end ;
```

to be used in the main program using a procedure call of the form

```
Push(n, Head) ;
```

Option *F*: the function *Push* defined by

```
function Push(n : integer ; h : Link) : Link ;
var
  p : Link ;
begin
  new(p) ;
  p↑.Info := n ;
  p↑.Next := h ;
  Push := p
end ;
```

to be used in the main program using an assignment of the form

```
Head := Push(n, Head) ;
```

Suppose *Head* points to the first *Cell* of a list of *Cells*, and a new *Cell* is to be added at the front of the list, leaving *Head* pointing to the newly added *Cell*.

Which of the following implementations will FAIL to accomplish this objective?

- (A) Option *P* with parameters passed by value
- (B) Option *F* with parameters passed by value
- (C) Option *P* with parameters passed by reference
- (D) Option *F* with parameters passed by reference
- (E) Option *P* with parameters passed by name

GO ON TO THE NEXT PAGE.


31. An XY flip-flop operates as indicated by the following table.

<u>Inputs</u>		<u>Current State</u>	<u>Next State</u>
X	Y	Q	Q
0	0	Q	1
0	1	Q	\bar{Q}
1	1	Q	0
1	0	Q	Q

Which of the following expresses the next state in terms of the X and Y inputs and the current state Q ?

- (A) $(\bar{X} \wedge \bar{Q}) \vee (\bar{Y} \wedge Q)$
- (B) $(\bar{X} \wedge Q) \vee (\bar{Y} \wedge \bar{Q})$
- (C) $(X \wedge \bar{Q}) \vee (Y \wedge Q)$
- (D) $(X \wedge \bar{Q}) \vee (\bar{Y} \wedge Q)$
- (E) $(X \wedge \bar{Q}) \vee (\bar{Y} \wedge \bar{Q})$

32. A black-and-white computer graphics display is divided up into an array of pixels as shown below.



Each of the pixels can take on one of eight gray levels ranging from 0 (white) to 7 (black). In order to prevent sharp discontinuities of shade, the software system that causes pictures to be displayed enforces the rule that the gray levels of two adjacent pixels cannot differ by more than two. How many of the 64 possible assignments of gray levels to two adjacent pixels satisfy this rule?

- (A) 24
- (B) 32
- (C) 34
- (D) 40
- (E) 64

GO ON TO THE NEXT PAGE.

33. A doubly linked list is declared as

```
Element = record
           Value   : integer ;
           Fwd, Bwd : ↑Element
           end ;
```

where *Fwd* and *Bwd* represent forward and backward links to adjacent elements of the list.

Which of the following segments of code deletes the element pointed to by *X* from the doubly linked list, if it is assumed that *X* points to neither the first nor the last element of the list?

(A) $X↑.Bwd↑.Fwd := X↑.Fwd ;$
 $X↑.Fwd↑.Bwd := X↑.Bwd$

(B) $X↑.Bwd↑.Fwd := X↑.Bwd ;$
 $X↑.Fwd↑.Bwd := X↑.Fwd$

(C) $X↑.Bwd↑.Bwd := X↑.Fwd ;$
 $X↑.Fwd↑.Fwd := X↑.Bwd$

(D) $X↑.Bwd↑.Bwd := X↑.Bwd ;$
 $X↑.Fwd↑.Fwd := X↑.Fwd$

(E) $X↑.Bwd := X↑.Fwd ;$
 $X↑.Fwd := X↑.Bwd$

GO ON TO THE NEXT PAGE.

34. Processes $P1$ and $P2$ have a producer-consumer relationship, communicating by the use of a set of shared buffers:

```
 $P1$  : repeat  
    obtain an empty buffer  
    fill it  
    return a full buffer  
forever
```

```
 $P2$  : repeat  
    obtain a full buffer  
    empty it  
    return an empty buffer  
forever
```

Increasing the number of buffers is likely to do which of the following?

- I. Increase the rate at which requests are satisfied (throughput)
- II. Decrease the likelihood of deadlock
- III. Increase the ease of achieving a correct implementation

- (A) I only
- (B) II only
- (C) III only
- (D) II and III only
- (E) I, II, and III

GO ON TO THE NEXT PAGE.

35. Which of the following instruction-set features is NOT generally considered an obstacle to aggressive pipelining of an integer unit?
- (A) Condition codes set by every instruction
 - (B) Variable-length encoding of instructions
 - (C) Instructions requiring widely varying numbers of cycles to execute
 - (D) Several different classes (sets) of registers
 - (E) Instructions that read several arguments from memory
-

GO ON TO THE NEXT PAGE.

Questions 36-38 are based on an implementation of an abstract type “set of integers” by a concrete type *IntSet* defined as follows.

```
type
  IntSet = record
    Last: 0..Max;
    V: array[1..Max] of integer
  end;
```

The “add element” operation on an object *S* is implemented by storing the value of the element in $S.V[S.Last + 1]$ and incrementing $S.Last$, unless $Last = Max$, in which case an error flag is raised. Note that duplicate values may appear in an object’s concrete representation but are hidden in its abstract representation.

36. Let A be the representation function from *IntSet* into “set of integers” such that, for each concrete object S of type *IntSet*, $A(S)$ is the set of integers represented by S . $A(S)$ can be written as
- (A) $\{i \mid i \in 1..S.Last\}$
 - (B) $\{i \mid i \in S.Last..Max\}$
 - (C) $\{S.V[i] \mid i \in 1..S.Last\}$
 - (D) $\{S.V[i] \mid i \in 1..Max\}$
 - (E) $\{S.V[i] \mid i \in 1..S.V[S.Last]\}$
37. An advantage of choosing this implementation of “set of integers” is that adding an element to a set is a constant-time operation. Which of the following is a disadvantage of this implementation?
- (A) Adding elements to very small sets could cause error flags to be raised.
 - (B) Deleting elements from very large sets could cause error flags to be raised.
 - (C) Determining whether a set is empty will require nonconstant time.
 - (D) Constructing the union of two sets will require quadratic time in the size of the set being constructed.
 - (E) Deleting an element from a set will require exponential time in the size of the set from which the element is deleted.

GO ON TO THE NEXT PAGE.

38. The following code fragment mutates concrete *IntSet* objects.

```
procedure P(var S: IntSet, x: integer) ;
  var k: integer ;
begin
  k := 1 ;
  while k <= S.Last do
  begin
    if S.V[k] = x then
    begin
      S.V[k] := S.V[S.Last] ;
      S.Last := S.Last - 1
    end
    else k := k + 1
  end
end ;
```

Which of the following abstract operations of “set of integers” does *P* implement?

- (A) Add *x* to *S*
 - (B) Delete *x* from *S*
 - (C) Intersect {*x*} and *S*
 - (D) Union {*x*} and *S*
 - (E) Make a copy of *S*
-

39. To compute the matrix product M_1M_2 , where M_1 has p rows and q columns and where M_2 has q rows and r columns, takes time proportional to pqr , and the result is a matrix of p rows and r columns. Consider the product of three matrices $N_1N_2N_3$ that have, respectively, w rows and x columns, x rows and y columns, and y rows and z columns. Under what condition will it take less time to compute the product as $(N_1N_2)N_3$ (i.e., multiply the first two matrices first) than to compute it as $N_1(N_2N_3)$?

- (A) There is no such condition; i.e., they will always take the same time.
 - (B) $1/x + 1/z < 1/w + 1/y$
 - (C) $x > y$
 - (D) $1/w + 1/x < 1/y + 1/z$
 - (E) $w + x > y + z$
-

GO ON TO THE NEXT PAGE.

40. Which of the following is usually NOT represented in a subroutine's activation record frame for a stack-based programming language?
- (A) Values of local variables
 - (B) A heap area
 - (C) The return address
 - (D) Stack pointer for the calling activation record
 - (E) Information needed to access nonlocal variables
-

41. Let A be a finite nonempty set with cardinality n . The number of subsets $S \subseteq A$ having odd cardinality is
- (A) n
 - (B) $2^{\frac{n}{2}}$
 - (C) 2^{n-1}
 - (D) 2^n
 - (E) not determinable except in terms of whether n is even or odd
-

GO ON TO THE NEXT PAGE.

42. A certain algorithm A has been shown to have running time $O(N^{2.5})$, where N is the size of the input. Which of the following is NOT true about algorithm A ?
- (A) There exist constants C_1 and C_2 such that for all N the running time is less than $C_1N^{2.5} + C_2$ seconds.
 - (B) For all N , there may be some inputs for which the running time is less than $N^{2.4}$ seconds.
 - (C) For all N , there may be some inputs for which the running time is less than $N^{2.6}$ seconds.
 - (D) For all N , there may be some inputs for which the running time is more than $N^{2.4}$ seconds.
 - (E) For all N , there may be some inputs for which the running time is more than $N^{2.6}$ seconds.
-

43. A data structure is comprised of nodes each of which has exactly two pointers to other nodes, with no null pointers. The following C program is to be used to count the number of nodes accessible from a given node. It uses a mark field, assumed to be initially zero for all nodes. There is a statement missing from this code.

```
struct test {int info, mark; struct test *p, *q; }
int nodecount(struct test *a)
{
    if (a->mark) return 0 ;
    return nodecount(a->p) + nodecount(a->q) + 1 ;
}
```


Which change should be made to make the program work properly?

- (A) Add " $a \rightarrow \text{mark} = 1;$ " as the first statement.
 - (B) Add " $a \rightarrow \text{mark} = 1;$ " after the "if" statement.
 - (C) Add " $a \rightarrow \text{mark} = 1;$ " as the last statement.
 - (D) Add " $a \rightarrow \text{mark} = 0;$ " after the "if" statement.
 - (E) Add " $a \rightarrow \text{mark} = 0;$ " as the last statement.
-

GO ON TO THE NEXT PAGE.

44. Of the following, which best characterizes computers that use memory-mapped I/O?
- (A) The computer provides special instructions for manipulating I/O ports.
 - (B) I/O ports are placed at addresses on the bus and are accessed just like other memory locations.
 - (C) To perform an I/O operation, it is sufficient to place the data in an address register and call the channel to perform the operation.
 - (D) Ports are referenced only by memory-mapped instructions of the computer and are located at hardwired memory locations.
 - (E) I/O can be performed only when memory management hardware is turned on.
-


Questions 45-46 are based on the graph below, which shows measured values of throughput *versus* multiprogramming level for a particular computer system. (“Throughput” is the rate at which requests are satisfied; “multiprogramming level” is the number of requests competing for system resources.)



45. At lower multiprogramming levels, throughput increases as multiprogramming level increases. This phenomenon is best explained by the fact that as multiprogramming level increases
- (A) the system overhead increases
 - (B) some system resource begins to saturate (i.e., to be utilized 100%)
 - (C) I/O activity per request remains constant
 - (D) the average time spent in the system by each request increases
 - (E) the potential for concurrent activity among system resources increases
46. At intermediate multiprogramming levels, the rate of increase of throughput with multiprogramming levels decreases. This phenomenon is best explained by the fact that as multiprogramming level increases
- (A) I/O activity per request remains constant
 - (B) some system resource begins to saturate (i.e., to be utilized 100%)
 - (C) the utilization of memory improves
 - (D) the average time spent in the system by each request increases
 - (E) the potential for concurrent activity among system resources increases
-

GO ON TO THE NEXT PAGE.

47. If *Tree1* and *Tree2* are the trees indicated below,



which traversals of *Tree1* and *Tree2*, respectively, will produce the same sequence of node names?

- (A) Preorder, postorder
- (B) Postorder, inorder
- (C) Postorder, postorder
- (D) Inorder, inorder
- (E) Postorder, preorder

48. Let A be a finite set with m elements, and let B be a finite set with n elements. The number of distinct functions mapping A into B is


- (A) n^m
- (B) $n!/(n - m)!$
- (C) $n!$
- (D) $n!/(m!(n - m)!)$
- (E) 2^{n+m}

GO ON TO THE NEXT PAGE.

49.

$$S \rightarrow aS \mid b$$

The “parsing automaton” below is for the context-free grammar with the productions indicated above.



Each state includes certain “items,” which are productions with dots in their right sides. The parser using this automaton, with $X_1X_2 \dots X_n$ on the stack, reduces by production $A \rightarrow \alpha$ if and only if there is a path, labeled $X_1X_2 \dots X_n$ from the start state to a state that includes the item $A \rightarrow \alpha \cdot$ (note the dot at the right end). Which of the following stack contents causes the parser to reduce by some production?


- (A) a (B) aa (C) bb (D) aaS (E) ϵ

50. Let $k \geq 2$. Let L be the set of strings in $\{0,1\}^*$ such that $x \in L$ if and only if the number of 0's in x is divisible by k and the number of 1's in x is odd. The minimum number of states in a deterministic finite automaton (DFA) that recognizes L is

- (A) $k + 2$ (B) $2k$ (C) $k \log k$ (D) k^2 (E) 2^k

GO ON TO THE NEXT PAGE.

51.



In the diagram above, the inverter and the *AND*-gates labeled 1 and 2 have delays of 9, 10, and 12 nanoseconds, respectively. Wire delays are negligible. For certain values of *a* and *c*, together with a certain transition of *b*, a glitch (spurious output) is generated for a short time, after which the output assumes its correct value. The duration of the glitch is

- (A) 7 ns (B) 9 ns (C) 11 ns (D) 13 ns (E) 31 ns
-

52. Of the following sorting algorithms, which has a running time that is LEAST dependent on the initial ordering of the input?

- (A) Insertion sort
 - (B) Quicksort
 - (C) Merge sort
 - (D) Selection sort
 - (E) Shellsort
-

GO ON TO THE NEXT PAGE.

53. A certain well-known computer family represents the exponents of its floating-point numbers as “excess-64” integers; i.e., a typical exponent $e_6e_5e_4e_3e_2e_1e_0$ represents the number $e = -64 + \sum_{i=0}^6 2^i e_i$. Two such exponents are input to a conventional 7-bit parallel adder. Which of the following should be accomplished in order to obtain a sum that is also in excess-64 notation?
- (A) The most significant adder output bit should be complemented.
 - (B) An end-around carry should be generated.
 - (C) The adder outputs should be bitwise complemented.
 - (D) A low-order carry should be inserted.
 - (E) The adder output should be left unchanged.
-

Questions 54-55 are based on the following information.

Consider a virtual memory with M resident pages and a periodic page reference string

$$p_1, p_2, \dots, p_N, p_1, p_2, \dots, p_N, p_1, p_2, \dots, p_N$$


of N distinct requests.

Assume that the string of N requests was constructed randomly, and assume that initially none of the pages are resident.

54. Assume that $N = 2M$ and FIFO is used. If the string p_1, p_2, \dots, p_N is repeated three times, then the number of page faults is
- (A) $N/2$
 - (B) N
 - (C) $N + 3$
 - (D) $2N$
 - (E) $3N$
55. If the string of N requests is repeated many times, which of the following is (are) true?
- I. The fewest page faults occur when $M \geq N$.
 - II. LRU and FIFO will always result in the same number of page faults.
 - III. LIFO will always result in the fewest possible page faults.
- (A) I only
 - (B) III only
 - (C) I and II
 - (D) I and III
 - (E) II and III
-

GO ON TO THE NEXT PAGE.

56.



The logic circuit above is used to compare two unsigned 2-bit numbers, $X_1 X_0 = X$ and $Y_1 Y_0 = Y$, where X_0 and Y_0 are the least significant bits. (A small circle on any line in a logic diagram indicates logical NOT.)

Which of the following always makes the output Z have the value 1 ?

- (A) $X > Y$ (B) $X < Y$ (C) $X = Y$ (D) $X \geq Y$ (E) $X \neq Y$
-

57. It is known that the language $L \subseteq \{a,b\}^*$ that consists of all strings that contain an equal number of a 's and b 's is context-free. Let M be the regular language a^*b^* . Which of the following is (are) true?

- I. $L \cap M$ is a context-free language.
- II. $L \cap M$ is a regular language.
- III. $L \cap M = \{a^n b^m \mid n \text{ is a positive integer less than integer } m\}$.

- (A) None
 (B) I only
 (C) III only
 (D) I and III
 (E) II and III
-

58. Of the following, which gives the best upper bound for the value of $f(N)$ where f is a solution to the recurrence

$$f(2N + 1) = f(2N) = f(N) + \log N \text{ for } N \geq 1,$$

with $f(1) = 0$?

- (A) $O(\log N)$ (B) $O(N \log N)$ (C) $O(\log N) + O(1)$ (D) $O((\log N)^2)$ (E) $O(N)$
-

GO ON TO THE NEXT PAGE.

59. Let I denote the formula: $(q \Rightarrow p) \Rightarrow (p \Rightarrow q)$
Let II denote the formula: $(p \Rightarrow q) \wedge q$

Which of the following is true?

- (A) I is not a tautology and II is not satisfiable.
 - (B) I is not a tautology and II is satisfiable.
 - (C) I is satisfiable and II is not satisfiable.
 - (D) I is a tautology and II is satisfiable.
 - (E) I is satisfiable and II is a tautology.
-

60. Consider a data type whose elements are integers and whose operations are INSERT, DELETE, and FINDCLOSEST, with FINDCLOSEST(y) defined to be some element x in the current set such that $|x - y| \leq |x_i - y|$ for all x_i in the current set. Let

$$T = \max(T_{\text{INSERT}}, T_{\text{DELETE}}, T_{\text{FINDCLOSEST}})$$

where T_{OP} denotes the worst-case time complexity for the given operation OP . Which of the following data structures would be best to use in order to minimize T ?

- (A) A sorted list
 - (B) An unordered list
 - (C) An implicit heap
 - (D) An AVL tree
 - (E) A hash table with conflicts resolved by a linked list
-

61. A block of 105 words of memory is used for dynamic storage for objects of sizes 3 and 10 words. The operations supported by the storage are:

```
 $x := alloc(n)$  {Allocate any block of  $n$  consecutive words and }  
                {return its starting address ; note that 3 and }  
                {10 are the only legal values for  $n$ . }  
                }
```

```
 $free(y)$  {Make the previously allocated block starting at }  
          {address  $y$  available for re-use. }  
          }
```

At a point where a certain request for allocation of a block of words cannot be granted because of a lack of a sufficiently long block of consecutive words to fill that request, what is the minimum possible number of words that might actually be in use?

- (A) 10 (B) 24 (C) 53 (D) 96 (E) 103
-

GO ON TO THE NEXT PAGE.

62. Languages with a structure that implements abstract data types (e.g., a C++ class) can prevent access to components of this structure by all operations except those that are part of this structure. However, definitions of such a structure often contain declarations of components of the structure (e.g., the header file for a C++ class may contain declarations of its private components).

For such a language, an object's name could be bound at run time to stack storage for its component values (direct representation) or to a stack pointer referencing heap storage for its component values (indirect representation). Which of the following statements about comparisons between direct and indirect representations is (are) true?

- I. Indirect representation noticeably increases compilation time.
- II. Direct representation decreases the time needed to access components of a variable.
- III. When the storage size of some private component of a variable changes, indirect representation minimizes the number of recompilations of source modules that must be performed.

- (A) I only
- (B) III only
- (C) I and II only
- (D) II and III only
- (E) I, II, and III

GO ON TO THE NEXT PAGE.

Questions 63-64 refer to the following information.

An array $A[1..n]$ is said to be k -ordered if

$$A[i - k] \leq A[i] \leq A[i + k]$$

for each i such that $k < i \leq n - k$. For example, the array 1 4 2 6 3 7 5 8 is 2-ordered.

63. In a 2-ordered array of $2N$ elements, what is the maximum number of positions that an element can be from its position if the array were 1-ordered?

- (A) $2N - 1$ (B) 2 (C) $N/2$ (D) 1 (E) N

64. In an array of $2N$ elements that is both 2- and 3-ordered, what is the maximum number of positions that an element can be from its position if the array were 1-ordered?

- (A) $2N - 1$ (B) 2 (C) $N/2$ (D) 1 (E) N
-

GO ON TO THE NEXT PAGE.

65. Consider a queue between the two processes indicated below. N is the length of the queue; e, f , and b are semaphores.

```
init : e := N ; f := 0 ; b := 1 ;
```

```
process 1 :      process 2 :
loop            loop
  P(e)           P(f)
  P(b)           P(b)
  enqueue       dequeue
  V(b)           V(b)
  V(f)           V(e)
end loop        end loop
```

Which of the following statements is (are) true?

- I. The purpose of semaphore f is to ensure that *dequeue* is not executed on an empty queue.
- II. The purpose of semaphore e is to ensure that deadlock does not occur.
- III. The purpose of semaphore b is to provide mutual exclusion for queue operations.

- (A) I only (B) III only (C) I and III only (D) II and III only (E) I, II, and III
-

66. Church's thesis equates the concept of "computable function" with those functions computable by, for example, Turing machines. Which of the following is true of Church's thesis?

- (A) It was first proven by Alan Turing.
 - (B) It has not yet been proven, but finding a proof is a subject of active research.
 - (C) It can never be proven.
 - (D) It is now in doubt because of the advent of parallel computers.
 - (E) It was never believed, but was assumed in order to simplify certain undecidability results.
-

67. Consider the following C code.

```
int f(int x)
{ if (x < 1) return 1 ;
  else return f(x - 1) + g(x) ;
}

int g(int x)
{ if (x < 2) return 1 ;
  else return f(x - 1) + g(x/2) ;
}
```

Of the following, which best describes the growth of $f(x)$ as a function of x ?

- (A) Logarithmic (B) Linear (C) Quadratic (D) Cubic (E) Exponential
-

GO ON TO THE NEXT PAGE.

68. Let $G = (V, E)$ be a connected, undirected graph, and let a and b be two distinct vertices in V . Let P_1 be the problem of finding a shortest simple path between a and b , and let P_2 be the problem of finding a longest simple path between a and b .

Which of the following statements about P_1 and P_2 is true?

- (A) Both P_1 and P_2 can be solved in polynomial time.
 - (B) P_1 can be solved in polynomial time but P_2 is not known to be solvable in polynomial time.
 - (C) P_1 is not known to be solvable in polynomial time but P_2 can be solved in polynomial time.
 - (D) It is not known whether either P_1 or P_2 can be solved in polynomial time.
 - (E) It is known that neither P_1 nor P_2 can be solved in polynomial time.
-

69. Let T denote a nonempty binary tree in which every node either is a leaf or has two children. Then

$n(T)$ denotes the number of non-leaf nodes of T (where $n(T) = 0$ if T is a leaf),

$h(T)$ denotes the height of T (where $h(T) = 0$ if T is a leaf),

T_L denotes the left subtree of T , and T_R denotes the right subtree of T .

If F is a function defined by

$$F(T) = \begin{cases} 0 & \text{if } T \text{ is a leaf} \\ F(T_L) + F(T_R) + \min(h(T_L), h(T_R)) & \text{otherwise,} \end{cases}$$

then $F(T) =$

- (A) $n(T) + h(T) - 1$
 - (B) $n(T) + h(T)$
 - (C) $n(T) + h(T) + 1$
 - (D) $n(T) - h(T) - 1$
 - (E) $n(T) - h(T)$
-

70. If DFA denotes “deterministic finite automata” and NDFA denotes “nondeterministic finite automata,” which of the following is FALSE?

- (A) For any language L , if L can be recognized by a DFA, then \bar{L} can be recognized by a DFA.
 - (B) For any language L , if L can be recognized by an NDFA, then \bar{L} can be recognized by an NDFA.
 - (C) For any language L , if L is context-free, then \bar{L} is context-free.
 - (D) For any language L , if L can be recognized in polynomial time, then \bar{L} can be recognized in polynomial time.
 - (E) For any language L , if L is decidable, then \bar{L} is decidable.
-

IF YOU FINISH BEFORE TIME IS CALLED, YOU MAY CHECK YOUR WORK ON THIS TEST.

I

NOTE: To ensure prompt processing of test results, it is important that you fill in the blanks exactly as directed.

SUBJECT TEST

A. Print and sign your full name in this box:

PRINT: _____
(LAST) (FIRST) (MIDDLE)

SIGN: _____

Copy this code in box 6 on your answer sheet. Then fill in the corresponding ovals exactly as shown.

6. TITLE CODE				
2	9	5	2	1
0	0	0	0	0
1	1	1	1	1
2	2	2	2	2
3	3	3	3	3
4	4	4	4	4
5	5	5	5	5
6	6	6	6	6
7	7	7	7	7
8	8	8	8	8
9	9	9	9	9

Copy the Test Name and Form Code in box 7 on your answer sheet.

TEST NAME Computer Science

FORM CODE GR9629

GRADUATE RECORD EXAMINATIONS SUBJECT TEST

B. The Subject Tests are intended to measure your achievement in a specialized field of study. Most of the questions are concerned with subject matter that is probably familiar to you, but some of the questions may refer to areas that you have not studied.

Your score will be determined by subtracting one-fourth the number of incorrect answers from the number of correct answers. Questions for which you mark no answer or more than one answer are not counted in scoring. If you have some knowledge of a question and are able to rule out one or more of the answer choices as incorrect, your chances of selecting the correct answer are improved, and answering such questions will likely improve your score. It is unlikely that pure guessing will raise your score; it may lower your score.

You are advised to use your time effectively and to work as rapidly as you can without losing accuracy. Do not spend too much time on questions that are too difficult for you. Go on to the other questions and come back to the difficult ones later if you can.

YOU MUST INDICATE ALL YOUR ANSWERS ON THE SEPARATE ANSWER SHEET. No credit will be given for anything written in this examination book, but you may write in the book as much as you wish to work out your answers. After you have decided on your response to a question, fill in the corresponding oval on the answer sheet. BE SURE THAT EACH MARK IS DARK AND COMPLETELY FILLS THE OVAL. Mark only one answer to each question. No credit will be given for multiple answers. Erase all stray marks. If you change an answer, be sure that all previous marks are erased completely. Incomplete erasures may be read as intended answers. Do not be concerned that the answer sheet provides spaces for more answers than there are questions in the test.

Example:

What city is the capital of France?

- (A) Rome
- (B) Paris
- (C) London
- (D) Cairo
- (E) Oslo

Sample Answer

(A) ● (C) (D) (E)

(A) ✗ (C) (D) (E)

(A) ✗ (C) (D) (E)

(A) ● (C) (D) (E)

(A) ● (C) (D) (E)

CORRECT ANSWER
PROPERLY MARKED

IMPROPER MARKS

DO NOT OPEN YOUR TEST BOOK UNTIL YOU ARE TOLD TO DO SO.

Scoring Your Subject Test

Computer Science Test scores typically range from 520 to 800. The range for different editions of a given test may vary because different editions are not of precisely the same difficulty. The differences in ranges among different editions of a given test, however, usually are small. This should be taken into account, especially when comparing two very high scores. In general, differences between scores at the 99th percentile should be ignored. **The score conversion table on page 51 shows the score range for this edition of the test only.**

The worksheet on page 50 lists the correct answers to the questions. Columns are provided for you to mark whether you chose the correct (C) answer or an incorrect (I) answer to each question. Draw a line across any question you omitted, because it is not counted in the scoring. At the bottom of the page, enter the total number correct and the total number incorrect. Divide the total incorrect by 4 and subtract the resulting number from the total correct. This is the adjustment made for guessing. Then round the result to the nearest whole number. This will give you your raw total score. Use the total score conversion table to find the scaled total score that corresponds to your raw total score.

Example: Suppose you chose the correct answers to 38 questions and incorrect answers to 22. Dividing 22 by 4 yields 5.5. Subtracting 5.5 from 38 equals 32.5, which is rounded to 33. The raw score of 33 corresponds to a scaled score of 670.

Worksheet for the Computer Science Test, Form GR9629 Only
Answer Key and Percentages* of Examinees
Answering Each Question Correctly

QUESTION Number	Answer	P +	TOTAL	
			C	I
1	C	64		
2	E	69		
3	E	81		
4	E	82		
5	D	88		
6	E	78		
7	D	76		
8	A	78		
9	A	70		
10	D	83		
11	D	43		
12	D	65		
13	B	69		
14	C	39		
15	C	52		
16	D	56		
17	D	76		
18	C	69		
19	A	67		
20	D	76		
21	B	21		
22	C	64		
23	D	70		
24	B	28		
25	B	42		
26	C	72		
27	E	68		
28	D	64		
29	C	52		
30	A	64		
31	A	56		
32	C	48		
33	A	83		
34	A	39		
35	D	26		

QUESTION Number	Answer	P +	TOTAL	
			C	I
36	C	54		
37	A	25		
38	B	70		
39	B	29		
40	B	46		
41	C	33		
42	E	61		
43	B	62		
44	B	49		
45	E	60		
46	B	60		
47	B	74		
48	A	33		
49	D	32		
50	B	30		
51	A	47		
52	C	41		
53	A	20		
54	E	57		
55	C	35		
56	A	56		
57	B	43		
58	D	17		
59	B	45		
60	D	28		
61	B	24		
62	D	40		
63	E	25		
64	D	29		
65	C	45		
66	C	22		
67	E	42		
68	B	39		
69	E	22		
70	C	45		

Correct (C) _____

Incorrect (I) _____

Total Score:

C - I/4 = _____

Scaled Score (SS) = _____

* The P+ column indicates the percentage of Computer Science Test examinees that answered each question correctly; it is based on a sample of December 1996 examinees selected to represent all Computer Science Test examinees tested between October 1, 1993, and September 30, 1996.

**Score Conversions and Percents Below*
For GRE Computer Science Test, Form GR9629 Only**

TOTAL SCORE					
Raw Score	Scaled Score	%	Raw Score	Scaled Score	%
70	910	99	36-37	690	60
68-69	900	99	35	680	56
67	890	99	33-34	670	53
65-66	880	99	32	660	49
64	870	99	30-31	650	46
62-63	860	98	29	640	42
61	850	97	27-28	630	39
59-60	840	96	25-26	620	35
58	830	95	24	610	32
56-57	820	94	22-23	600	29
55	810	93	21	590	25
53-54	800	91	19-20	580	23
52	790	89	18	570	19
50-51	780	86	16-17	560	17
48-49	770	85	15	550	15
47	760	82	13-14	540	12
45-46	750	79	12	530	11
44	740	77	10-11	520	9
42-43	730	74	9	510	7
41	720	71	7-8	500	6
39-40	710	67	6	490	5
38	700	64	4-5	480	4
			2-3	470	3
			1	460	2
			0	450	1

*Percentage scoring below the scaled score is based on the performance of 13,161 examinees who took the Computer Science Test between October 1, 1993, and September 30, 1996. Due to changes in the test-taking population, the percentile rank data have also changed. To obtain current percentile rank information, visit the GRE Web site at www.gre.org/codelst.html, or contact the GRE Program.

Evaluating Your Performance

Now that you have scored your test, you may wish to compare your performance with the performance of others who took this test. Both the worksheet on page 50 and the table on page 51 use performance data from GRE Computer Science Test examinees.

The data in the worksheet on page 50 are based on the performance of a sample of the examinees who took this test in December 1996. This sample was selected to represent the total population of GRE Computer Science Test examinees tested between October 1993 and September 1996. The numbers in the column labeled “P+” on the worksheet indicate the percentages of examinees in this sample who answered each question correctly. You may use these numbers as a guide for evaluating your performance on each test question.

The table on page 51 contains, for each scaled score, the percentage of examinees tested between October 1993 and September 1996 who received lower scores. Interpretive data based on the scores earned by examinees tested in this three-year period were used by admissions officers in the 1997-98 testing year. These percentages appear in the score conversion table in a

column to the right of the scaled scores. For example, in the percentage column opposite the scaled score of 730 is the number 74. This means that 74 percent of the GRE Computer Science Test examinees tested between October 1993 and September 1996 scored lower than 730. To compare yourself with this population, look at the percentage next to the scaled score you earned on the practice test. *Note:* due to changes in the test-taking population, the percentile rank data have also changed. To obtain current percentile rank information, visit the GRE Web site at www.gre.org/codelst.html, or contact the GRE Program.

It is important to realize that the conditions under which you tested yourself were not exactly the same as those you will encounter at a test center. It is impossible to predict how different test-taking conditions will affect test performance, and this is only one factor that may account for differences between your practice test scores and your actual test scores. By comparing your performance on this practice test with the performance of other GRE Computer Science Test examinees, however, you will be able to determine your strengths and weaknesses and can then plan a program of study to prepare yourself for taking the GRE Computer Science Test under standard conditions.

SUBJECT TEST

COMPLETE THE
CERTIFICATION STATEMENT,
THEN TURN ANSWER SHEET
OVER TO SIDE 1.

CERTIFICATION STATEMENT

Please write the following statement below, DO NOT PRINT.

"I certify that I am the person whose name appears on this answer sheet. I also agree not to disclose the contents of the test I am taking today to anyone." Sign and date where indicated.

SIGNATURE: _____ DATE: _____ / _____ / _____
Month Day Year

BE SURE EACH MARK IS DARK AND COMPLETELY FILLS THE INTENDED SPACE AS ILLUSTRATED HERE: ●
YOU MAY FIND MORE RESPONSE SPACES THAN YOU NEED. IF SO, PLEASE LEAVE THEM BLANK.

115	A	B	C	D	E	147	A	B	C	D	E	179	A	B	C	D	E	211	A	B	C	D	E
116	A	B	C	D	E	148	A	B	C	D	E	180	A	B	C	D	E	212	A	B	C	D	E
117	A	B	C	D	E	149	A	B	C	D	E	181	A	B	C	D	E	213	A	B	C	D	E
118	A	B	C	D	E	150	A	B	C	D	E	182	A	B	C	D	E	214	A	B	C	D	E
119	A	B	C	D	E	151	A	B	C	D	E	183	A	B	C	D	E	215	A	B	C	D	E
120	A	B	C	D	E	152	A	B	C	D	E	184	A	B	C	D	E	216	A	B	C	D	E
121	A	B	C	D	E	153	A	B	C	D	E	185	A	B	C	D	E	217	A	B	C	D	E
122	A	B	C	D	E	154	A	B	C	D	E	186	A	B	C	D	E	218	A	B	C	D	E
123	A	B	C	D	E	155	A	B	C	D	E	187	A	B	C	D	E	219	A	B	C	D	E
124	A	B	C	D	E	156	A	B	C	D	E	188	A	B	C	D	E	220	A	B	C	D	E
125	A	B	C	D	E	157	A	B	C	D	E	189	A	B	C	D	E	221	A	B	C	D	E
126	A	B	C	D	E	158	A	B	C	D	E	190	A	B	C	D	E	222	A	B	C	D	E
127	A	B	C	D	E	159	A	B	C	D	E	191	A	B	C	D	E	223	A	B	C	D	E
128	A	B	C	D	E	160	A	B	C	D	E	192	A	B	C	D	E	224	A	B	C	D	E
129	A	B	C	D	E	161	A	B	C	D	E	193	A	B	C	D	E	225	A	B	C	D	E
130	A	B	C	D	E	162	A	B	C	D	E	194	A	B	C	D	E	226	A	B	C	D	E
131	A	B	C	D	E	163	A	B	C	D	E	195	A	B	C	D	E	227	A	B	C	D	E
132	A	B	C	D	E	164	A	B	C	D	E	196	A	B	C	D	E	228	A	B	C	D	E
133	A	B	C	D	E	165	A	B	C	D	E	197	A	B	C	D	E	229	A	B	C	D	E
134	A	B	C	D	E	166	A	B	C	D	E	198	A	B	C	D	E	230	A	B	C	D	E
135	A	B	C	D	E	167	A	B	C	D	E	199	A	B	C	D	E	231	A	B	C	D	E
136	A	B	C	D	E	168	A	B	C	D	E	200	A	B	C	D	E	232	A	B	C	D	E
137	A	B	C	D	E	169	A	B	C	D	E	201	A	B	C	D	E	233	A	B	C	D	E
138	A	B	C	D	E	170	A	B	C	D	E	202	A	B	C	D	E	234	A	B	C	D	E
139	A	B	C	D	E	171	A	B	C	D	E	203	A	B	C	D	E	235	A	B	C	D	E
140	A	B	C	D	E	172	A	B	C	D	E	204	A	B	C	D	E	236	A	B	C	D	E
141	A	B	C	D	E	173	A	B	C	D	E	205	A	B	C	D	E	237	A	B	C	D	E
142	A	B	C	D	E	174	A	B	C	D	E	206	A	B	C	D	E	238	A	B	C	D	E
143	A	B	C	D	E	175	A	B	C	D	E	207	A	B	C	D	E	239	A	B	C	D	E
144	A	B	C	D	E	176	A	B	C	D	E	208	A	B	C	D	E	240	A	B	C	D	E
145	A	B	C	D	E	177	A	B	C	D	E	209	A	B	C	D	E	241	A	B	C	D	E
146	A	B	C	D	E	178	A	B	C	D	E	210	A	B	C	D	E	242	A	B	C	D	E

IF YOU DO NOT WANT THIS ANSWER SHEET TO BE SCORED

If you want to cancel your scores from this administration, complete A and B below. You will not receive scores for this test; however, you will receive confirmation of this cancellation. No record of this test or the cancellation will be sent to the recipients you indicated, and there will be no scores for this test on your GRE file. Once a score is canceled, it cannot be reinstated.

To cancel your scores from this test administration, you must:

A. fill in both ovals here . . . ○ — ○ B. sign your full name here: _____

TR	TW	TFS	TCS	1R	1W	1FS	1CS	2R	2W	2FS	2CS
FOR ETS USE ONLY				3R	3W	3FS	3CS	4R	4W	4FS	4CS
				5R	5W	5FS	5CS	6R	6W	6FS	6CS

