

Name: _____

This exam is open book, open notes, but there can be no sharing of any material. You can use the Internet, but only as a library. If you are not sure if something is allowed, check with me.

COS 480 students must answer all questions that are not marked with a \star in 40 minutes. The points for each question are indicated in parentheses next to the question number. Questions marked with a \star may also be answered, for extra credit.

COS 580 students must answer all questions, including those marked with a \star , in 60 minutes. Each question is worth $2/3$ times the points indicated in parentheses.

Some questions will use the database instance depicted below. The two title rows of each table indicate the attribute names and types. The **Trees** relation indicates, for each tree, its common name, botanical name, type, typical trunk diameter (in inches), typical height (in feet), and minimum and maximum zones. The zones refer to regions of the U.S. where the trees are likely to grow well. The **Places** relation indicates, for each place, the city, state, population, zone (as above), subzone (a finer subdivision of zones), and average minimum temperature. Your answers to questions that ask for queries should work for all instances of databases conforming to the given schema, not only the one depicted below.

These tables are repeated on the last page of the test. (You may detach that page and use it for reference. There is no need to reattach it.)

Trees

name	botname	ttype	dia	height	minz	maxz
varchar(50)	varchar(50)	varchar(25)	float	float	int	int
White Pine	Pinus strobus	coniferous	30.0	90.0	3	8
Pitch Pine	Pinus rigida	coniferous	18.0	35.0	5	7
Bigtooth Aspen	Populus grandidentata	deciduous	15.0	70.0	3	5
Quaking Aspen	Populus tremuloides	deciduous	13.0	67.5	1	8

Places

city	state	pop	zone	subzone	mintemp
varchar(20)	varchar(20)	integer	integer	char(1)	float
Orono	Maine	9112	5	b	-15.0
Bangor	Maine	31473	5	a	-15.1
Bar Harbor	Maine	4820	5	b	-14.0
Caribou	Maine	8312	4	a	-25.4
Van Buren	Maine	2631	3	a	-35.6
Tucson	Arizona	486699	8	a	39.0

For notational convenience in relational algebra, we shall abbreviate as follows:

Trees(name, botname, ttype, dia, height, minz, maxz) $T(N, B, T, D, H, M, X)$

Places(city, state, pop, zone, subzone, mintemp) $P(C, S, P, Z, Y, L)$

1. (1 pt.) Write your name in the space provided above.
2. (9 pts.) Indicate the result of evaluating each of the following SQL queries.

(a) `select avg(height) as AvgHt
 from Trees
 where botname like 'Pinus%';`

(b) `select name, ttype
 from Trees
 where minz >= all
 (select zone from Places where state = 'Maine');`

3. (10 pts.) Indicate the result of evaluating each of the following relational-algebra expressions (Recall the shorthand notation for table- and attribute-names from page 1.)

(a) $\sigma_{C='Van\ Buren'} \pi_{CSN} \sigma_{M \leq Z \wedge Z \leq X} (T \times P)$

(b) $\pi_{CC'ZYY'} \sigma_{Z=Z' \wedge Y < Y'} (P \times \rho_{P'(C',S',P',Z',Y',L')} P)$

4. (10 pts.) Write SQL queries as directed below.

(a) Write a SQL query to find trees and places such that the place's zone is the minimum zone for the tree. The output should consist of tuples of the form (n, c, s) , where n is the common name of a tree whose minimum zone equals the zone of the city c in state s .

(b) Write a SQL query to find pairs of trees that have overlapping zones and whose typical trunk diameters differ by five or more feet. The output should consist of tuples of the form (n_1, n_2) , where n_1 and n_2 are the common names of two trees whose zones overlap and n_2 's typical trunk diameter is at least 5 ft. greater than that of n_1 . (We say the zones of two trees overlap if the intervals formed by their `minz` and `maxz` attributes overlap or touch.)

5. (10 pts.) For each part below, write a relational algebra query that is equivalent to the SQL query in the corresponding part of Question 4.

(a)

(b)

6. (10 pts) ★ Given relations $R(A, B, C, D, E)$ and $S(D, E)$, the *quotient* of R and S is a relation consisting of tuples (a, b, c) such that for every tuple (d, e) in S , there is a tuple (a, b, c, d, e) in R :

$$R \div S = \{(a, b, c) \mid \forall (d, e) \in S : (a, b, c, d, e) \in R\}$$

Express $R \div S$ using only the six basic relational-algebra operators. [Hint: You need only the projection, cross product, and set difference operators.]

7. (10 pts.) * Write a SQL query that is equivalent to $R \div S$ as defined in Question 6. [Hint: It may be easier to think of the SQL query without referring to the algebra query.]

Scratch page

Material here will not be graded. You may detach and discard this page.

Trees

name	botname	ttype	dia	height	minz	maxz
varchar(50)	varchar(50)	varchar(25)	float	float	int	int
White Pine	Pinus strobus	coniferous	30.0	90.0	3	8
Pitch Pine	Pinus rigida	coniferous	18.0	35.0	5	7
Bigtooth Aspen	Populus grandidentata	deciduous	15.0	70.0	3	5
Quaking Aspen	Populus tremuloides	deciduous	13.0	67.5	1	8

Places

city	state	pop	zone	subzone	mintemp
varchar(20)	varchar(20)	integer	integer	char(1)	float
Orono	Maine	9112	5	b	-15.0
Bangor	Maine	31473	5	a	-15.1
Bar Harbor	Maine	4820	5	b	-14.0
Caribou	Maine	8312	4	a	-25.4
Van Buren	Maine	2631	3	a	-35.6
Tucson	Arizona	486699	8	a	39.0

For notational convenience in relational algebra, we shall abbreviate as follows:

Trees(name, botname, ttype, dia, height, minz, maxz) $T(N, B, T, D, H, M, X)$

Places(city, state, pop, zone, subzone, mintemp) $P(C, S, P, Z, Y, L)$