
Final In-Class Exam

(2:35 PM – 3:50 PM : 75 Minutes)

- This exam will account for either 15% or 30% of your overall grade depending on your relative performance in the midterm and the final. The higher of the two scores (midterm and final) will be worth 30% of your grade, and the lower one 15%.
- There are three (3) questions, worth 75 points in total. Please answer all of them in the spaces provided.
- There are 16 pages including three (3) blank pages and two (2) pages of appendices. Please use the blank pages if you need additional space for your answers.
- The exam is *open slides*.

GOOD LUCK!

| Question | Pages | Score | Maximum |
|--------------------------------------|-------|-------|---------|
| 1. Parallel Prefix Sum | 2–5 | | 25 |
| 2. ϵ -Approximate Frequency | 7–9 | | 30 |
| 3. Matrix Rotation | 11–13 | | 20 |
| Total | | | 75 |

NAME: _____

QUESTION 1. [25 Points] Parallel Prefix Sum. Given a sequence of n elements $\langle x_1, x_2, \dots, x_n \rangle$ drawn from a set S with a binary associative operator \oplus (e.g., addition, multiplication, maximum, matrix product, union, etc.), the *prefix sum* problem asks one to compute a sequence of n partial sums $\langle s_1, s_2, \dots, s_n \rangle$ such that $s_i = x_1 \oplus x_2 \oplus \dots \oplus x_i$ for $1 \leq i \leq n$. In lecture 26 we studied a parallel prefix sum algorithm with $\Theta(n)$ work and $\Theta(\log^2 n)$ span¹.

In this problem we will analyze another parallel prefix sum algorithm given in Figure 1.

```

ALT-PREFIX-SUM(  $\langle x_1, x_2, \dots, x_n \rangle, \oplus$  )
(Input is a sequence of  $n$  elements  $\langle x_1, x_2, \dots, x_n \rangle$  and a binary associative operator  $\oplus$ . Output is a sequence
 $\langle s_1, s_2, \dots, s_n \rangle$  with  $s_i = x_1 \oplus x_2 \oplus \dots \oplus x_i$ , for  $1 \leq i \leq n$ . We assume  $n = 2^k$  for some integer  $k \geq 0$ .)

1. if  $n = 1$  then
2.      $s_1 \leftarrow x_1$                                 {the prefix sum of a single element is the element itself}
3. else
4.     spawn  $\langle s_1, s_2, \dots, s_{\frac{n}{2}} \rangle \leftarrow$  ALT-PREFIX-SUM (  $\langle x_1, x_2, \dots, x_{\frac{n}{2}} \rangle, \oplus$  )
                                                {sets  $s_i = x_1 \oplus x_2 \oplus \dots \oplus x_i$  for  $1 \leq i \leq \frac{n}{2}$ }
5.      $\langle s_{\frac{n}{2}+1}, s_{\frac{n}{2}+2}, \dots, s_n \rangle \leftarrow$  ALT-PREFIX-SUM (  $\langle x_{\frac{n}{2}+1}, x_{\frac{n}{2}+2}, \dots, x_n \rangle, \oplus$  )
                                                {sets  $s_{\frac{n}{2}+i} = x_{\frac{n}{2}+1} \oplus x_{\frac{n}{2}+2} \oplus \dots \oplus x_{\frac{n}{2}+i}$  for  $1 \leq i \leq \frac{n}{2}$ }
6.     sync
7.     parallel for  $i \leftarrow 1$  to  $\frac{n}{2}$  do
8.          $s_{\frac{n}{2}+i} \leftarrow s_{\frac{n}{2}} \oplus s_{\frac{n}{2}+i}$ 
                                                {extends  $s_{\frac{n}{2}+i} = x_{\frac{n}{2}+1} \oplus x_{\frac{n}{2}+2} \oplus \dots \oplus x_{\frac{n}{2}+i}$ 
to  $s_{\frac{n}{2}+i} = s_{\frac{n}{2}} \oplus x_{\frac{n}{2}+1} \oplus x_{\frac{n}{2}+2} \oplus \dots \oplus x_{\frac{n}{2}+i} = x_1 \oplus x_2 \oplus \dots \oplus x_{\frac{n}{2}+i}$ }
9. return  $\langle s_1, s_2, \dots, s_n \rangle$ 

```

Figure 1: An alternate parallel prefix sum algorithm.

¹assuming the span of a *parallel for* loop with n iterations to be $\mathcal{O}(\log n + k)$, where k is the maximum span of a single iteration

1(a) [**7 Points**] Write down a recurrence relation describing the work done (i.e., T_1) by ALT-PREFIX-SUM, and solve it.

1(b) [**7 Points**] Write down a recurrence relation describing the span (i.e., T_∞) of ALT-PREFIX-SUM, and solve it.

1(c) [**6 Points**] Find the parallel running time (i.e., T_p) and parallelism of ALT-PREFIX-SUM.

1(d) [**5 Points**] Is ALT-PREFIX-SUM work-optimal? Why or why not?

Use this page if you need additional space for your answers.

QUESTION 2. [30 Points] ϵ -Approximate Frequency. Let $A[1 : n]$ be an array of length n containing both positive and negative numbers. Let m be the number of positive numbers in A , and let $p = \frac{m}{n}$. We are interested in estimating the value of m fast. Clearly, one can find the exact value of m in $\Theta(n)$ time simply by scanning A once and counting the number of positive numbers.

For any $\epsilon \in (0, p]$, we say that \hat{m} is an ϵ -approximation² of m provided $m - \epsilon n < \hat{m} < m + \epsilon n$.

```

APPROX-FREQ(  $A[1 : n]$ ,  $\epsilon$  )
(Inputs are an array  $A[1 : n]$  of  $n$  numbers, and a floating point parameter  $\epsilon \in (0, 1]$ . This routine chooses a sample of size  $\lceil \frac{6}{\epsilon^2} \ln n \rceil$  from  $A$  uniformly at random (with replacement), and uses that sample to estimate the number of entries of  $A$  that are positive.)

1.  $s \leftarrow \lceil \frac{6}{\epsilon^2} \ln n \rceil$                                      {size of the sample}
2.  $c \leftarrow 0$                                              {a counter that keeps track of the frequency of  $v$  in the chosen sample}
3. for  $i \leftarrow 1$  to  $m$  do                               {sample  $s$  items (with replacement) from  $A$ }
4.    $j \leftarrow \text{RANDOM}(1, n)$                              {choose an integer uniformly at random from  $[1, n]$ }
5.   if  $A[j] > 0$  then  $c \leftarrow c + 1$                    {choose  $A[j]$  as the next sample from  $A$ }
6. return  $\frac{c}{s} \times n$                                        {return the estimate}

```

Figure 2: Estimate the number of entries of $A[1 : n]$ that are positive.

This problem asks you to show that the function APPROX-FREQ given in Figure 2 which runs in $\Theta(\frac{1}{\epsilon^2} \ln n)$ worst-case time returns an ϵ -approximation of m w.h.p. in n . While analyzing the algorithm we will drop the ceiling in line 1 for simplicity, i.e., we will assume that $s = \frac{6}{\epsilon^2} \ln n$.

2(a) [5 Points] Let μ be the expected value of c right after the loop in lines 3–5 completes execution. Show that $\mu = (\frac{6}{\epsilon^2})(\frac{m}{n}) \ln n$.

²for simplicity, we have used ' $<$ ' instead of ' \leq ' in the definition of ϵ -approximation

2(b) [**12 Points**] Let \hat{c} be the exact value of c right after the loop in lines 3–5 completes execution. Prove that for $0 < \epsilon < p$ and $\delta = \frac{\epsilon}{p}$,

$$\Pr[\hat{c} \leq (1 - \delta)\mu] \leq \frac{1}{n^3} \quad \text{and} \quad \Pr[\hat{c} \geq (1 + \delta)\mu] \leq \frac{1}{n^2}.$$

2(c) [**5 Points**] let \hat{m} be the estimate of m returned by APPROX-FREQ. Argue that for $0 < \epsilon < p$, the results from part 2(b) imply the following:

$$\Pr [\hat{m} \leq m - \epsilon n] \leq \frac{1}{n^3} \text{ and } \Pr [\hat{m} \geq m + \epsilon n] \leq \frac{1}{n^2}.$$

2(d) [**8 Points**] Use your results from part 2(c) to argue that for $0 < \epsilon < p$, APPROX-FREQ returns an ϵ -approximation of m w.h.p. in n .

Use this page if you need additional space for your answers.

QUESTION 3. [20 Points] Matrix Rotation. The *rotation* of an $n \times n$ matrix X is another $n \times n$ matrix X^R obtained by writing the i -th row of X as the $n - i + 1$ -th column of X^R for $1 \leq i \leq n$. An example is given below.

$$X = \begin{bmatrix} a_1 & a_2 & a_3 & a_4 \\ b_1 & b_2 & b_3 & b_4 \\ c_1 & c_2 & c_3 & c_4 \\ d_1 & d_2 & d_3 & d_4 \end{bmatrix} \Rightarrow X^R = \begin{bmatrix} d_1 & c_1 & b_1 & a_1 \\ d_2 & c_2 & b_2 & a_2 \\ d_3 & c_3 & b_3 & a_3 \\ d_4 & c_4 & b_4 & a_4 \end{bmatrix}$$

In this problem we will analyze the cache complexity of a couple of algorithms for rotating square matrices. We will assume that all matrices are stored in row-major order.

3(a) [5 Points] Analyze the cache complexity of ITER-MATRIX-ROTATE given in Figure 3.

```

ITER-MATRIX-ROTATE( X, Y, n )
(Input is an  $n \times n$  square matrix  $X[ 1 : n, 1 : n ]$ . This function generates the rotation of  $X$  in  $Y$ .)
1. for  $i \leftarrow 1$  to  $n$  do
2.   for  $j \leftarrow 1$  to  $n$  do
3.      $Y[ i, j ] \leftarrow X[ n - j + 1, i ]$ 

```

Figure 3: Iterative matrix rotation.

- 3(b) [**10 Points**] Complete the recursive divide-and-conquer algorithm (REC-MATRIX-ROTATE) for rotating a square matrix given in Figure 4. Analyze its cache complexity assuming a *tall* cache (i.e., $M = \Omega(B^2)$, where M is the cache size and B is the cache block size).

```

REC-MATRIX-ROTATE( X, Y, n )
(Input is an  $n \times n$  square matrix  $X[1 : n, 1 : n]$ . This function recursively generates the rotation of  $X$ 
in  $Y$ . We assume  $n = 2^k$  for some integer  $k \geq 0$ . If  $n > 1$ , let  $X_{11}$ ,  $X_{12}$ ,  $X_{21}$  and  $X_{22}$  denote the top-left,
top-right, bottom-left and bottom-right quadrants of  $X$ , respectively. Similarly for  $Y$ .)
1. if  $n = 1$  then  $Y \leftarrow X$                                 {base case: the rotation of a  $1 \times 1$  matrix is the matrix itself}
2. else                                                         {divide  $X$  and  $Y$  into quadrants, and generate the rotation of  $X$  recursively.}
3.   REC-MATRIX-ROTATE(      ,      ,      )                    {fill out}
4.   REC-MATRIX-ROTATE(      ,      ,      )                    {fill out}
5.   REC-MATRIX-ROTATE(      ,      ,      )                    {fill out}
6.   REC-MATRIX-ROTATE(      ,      ,      )                    {fill out}

```

Figure 4: Recursive matrix rotation.

3(c) [**5 Points**] Is the cache complexity result of part 4(b) optimal? Why or why not?

Use this page if you need additional space for your answers.

APPENDIX I: SOME ELEMENTARY PROBABILITY RESULTS

Given an event A , $\Pr[A]$ denotes the probability of occurrence of A . By \bar{A} we denote the opposite or complement of event A . Then $\Pr[\bar{A}]$ denotes the probability of event A not occurring. Clearly,

$$0 \leq \Pr[A], \Pr[\bar{A}] \leq 1 \quad \text{and} \quad \Pr[\bar{A}] = 1 - \Pr[A].$$

Given two events A and B ,

- $A \cap B$ is the event of both A and B occurring, and
- $A \cup B$ is the event of at least one of A and B occurring.

Then the corresponding complements are as follows:

$$\overline{A \cap B} = \bar{A} \cup \bar{B} \quad \text{and} \quad \overline{A \cup B} = \bar{A} \cap \bar{B}.$$

If A and B are mutually exclusive (i.e., both cannot occur simultaneously³), then $\Pr[A \cap B] = 0$. You might find the following relationship useful:

$$\Pr[A \cup B] = \Pr[A] + \Pr[B] - \Pr[A \cap B].$$

Observe that if A and B are mutually exclusive, the relationship given above reduces to:

$$\Pr[A \cup B] = \Pr[A] + \Pr[B].$$

³e.g., if A is the event $(x < 5)$ and B is the event $(x > 5)$ then both A and B cannot be true (i.e., cannot occur) at the same time

APPENDIX II: USEFUL TAIL BOUNDS

Markov's Inequality. Let X be a random variable that assumes only nonnegative values. Then for all $\delta > 0$, $Pr[X \geq \delta] \leq \frac{E[X]}{\delta}$.

Chebyshev's Inequality. Let X be a random variable with a finite mean $E[X]$ and a finite variance $Var[X]$. Then for any $\delta > 0$, $Pr[|X - E[X]| \geq \delta] \leq \frac{Var[X]}{\delta^2}$.

Chernoff Bounds. Let X_1, \dots, X_n be independent Poisson trials, that is, each X_i is a 0-1 random variable with $Pr[X_i = 1] = p_i$ for some p_i . Let $X = \sum_{i=1}^n X_i$ and $\mu = E[X]$. Following bounds hold:

Lower Tail:

- for $0 < \delta < 1$, $Pr[X \leq (1 - \delta)\mu] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^\mu$
- for $0 < \delta < 1$, $Pr[X \leq (1 - \delta)\mu] \leq e^{-\frac{\mu\delta^2}{2}}$
- for $0 < \gamma < \mu$, $Pr[X \leq \mu - \gamma] \leq e^{-\frac{\gamma^2}{2\mu}}$

Upper Tail:

- for any $\delta > 0$, $Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right)^\mu$
- for $0 < \delta < 1$, $Pr[X \geq (1 + \delta)\mu] \leq e^{-\frac{\mu\delta^2}{3}}$
- for $0 < \gamma < \mu$, $Pr[X \geq \mu + \gamma] \leq e^{-\frac{\gamma^2}{3\mu}}$

APPENDIX III: THE MASTER THEOREM

Let $a \geq 1$ and $b > 1$ be constants, let $f(n)$ be a function, and let $T(n)$ be defined on the nonnegative integers by the recurrence

$$T(n) = \begin{cases} \Theta(1), & \text{if } n \leq 1, \\ aT\left(\frac{n}{b}\right) + f(n), & \text{otherwise,} \end{cases}$$

where, $\frac{n}{b}$ is interpreted to mean either $\lfloor \frac{n}{b} \rfloor$ or $\lceil \frac{n}{b} \rceil$. Then $T(n)$ has the following bounds:

Case 1: If $f(n) = \mathcal{O}(n^{\log_b a - \epsilon})$ for some constant $\epsilon > 0$, then $T(n) = \Theta(n^{\log_b a})$.

Case 2: If $f(n) = \Theta(n^{\log_b a} \log^k n)$ for some constant $k \geq 0$, then $T(n) = \Theta(n^{\log_b a} \log^{k+1} n)$.

Case 3: If $f(n) = \Omega(n^{\log_b a + \epsilon})$ for some constant $\epsilon > 0$, and $af\left(\frac{n}{b}\right) \leq cf(n)$ for some constant $c < 1$ and all sufficiently large n , then $T(n) = \Theta(f(n))$.