


## CMU 15-418/618: Parallel Computer Architecture and Programming Practice Exercise 5

### Problem 1: Transactions on Trees

Consider the binary search tree illustrated below.



The operations `insert` (insert value into tree, assuming no duplicates) and `sum` (return the sum of all elements in the tree) are implemented as transactional operations on the tree as shown below.

```
struct Node {
    Node *left, *right;
    int value;
};
Node* root; // root of tree, assume non-null

void insertNode(Node* n, int value) {
    if (value < n->value) {
        if (n->left == NULL)
            n->left = createNode(value);
        else
            insertNode(n->left, value);
    } else if (value > n->value) {
        if (n->right == NULL)
            n->right = createNode(value);
        else
            insertNode(n->right, value);
    } // insert won't be called with a duplicate element, so there's no else case
}

int sumNode(Node* n) {
    if (n == null) return 0;
    int total = n->value;
    total += sumNode(n->left);
    total += sumNode(n->right);
    return total;
}

void insert(int value) { atomic { insertNode(root, value); } }
int sum() { atomic { return sumNode(root); } }
```

Consider the following four operations are executed against the tree in parallel by different threads.

```
insert(10);  
insert(25);  
insert(24);  
int x = sum();
```

A. (2 pts) Consider different orderings of how these four operations could be evaluated. Please draw all possible trees that may result from execution of these four transactions. (Note: it's fine to draw only subtrees rooted at node 20 since that's the only part of the tree that's effected.)

B. (2 pts) Please list all possible values that may be returned by `sum()`.


C. (2 pts) Do your answers to parts A or B change depending on whether the implementation of transactions is optimistic or pessimistic? Why or why not?

- D. (2 pts) Consider an implementation of **lazy, optimistic** transactional memory that manages transactions at the granularity of tree nodes (the read and writes sets are lists of nodes). Assume that the transaction `insert(10)` commits when `insert(24)` and `insert(25)` are currently at node 20, and `sum()` is at node 40. Which of the four transactions (if any) are aborted? **Please describe why.**
- E. (2 pts) Assume that the transaction `insert(25)` commits when `insert(10)` is at node 15, `insert(24)` has already modified the tree but not yet committed, and `sum()` is at node 3. Which transactions (if any) are aborted? **Again, please describe why.**
- F. (2 pts) Now consider a transactional implementation that is **pessimistic** with respect to writes (check for conflict on write) and **optimistic** with respect to reads. The implementation also employs a “writer wins” conflict management scheme – meaning that the transaction issuing a conflicting write will not be aborted (the other conflicting transaction will). Describe how a **livelock problem** could occur in this code.

G. (2 pts) Give one livelock avoidance technique that an implementation of a pessimistic transactional memory system might use. You only need to summarize a basic approach, but make sure your answer is clear enough to refer to how you'd schedule the *transactions*.


## Problem 2: Interconnects

Consider sending two 256-bit packets in the Omega network below. Packet A is sent from node 4 to node 0, and packet B from node 6 to node 1. The network uses worm-hole routing with flits of size 32 bits. All network links can transmit 32 bits in a clock. **The first flit of both packets leaves the respective sending node at the same time. If flits from A and B must contend for a link, flits from packet A always get priority.**



- A. (2 pts) What is the latency of transmitting packet A to its destination?
- B. (2 pts) What is the latency of transmitting packet B to its destination? Please describe your calculation—switches are numbered to help your explanation.)
- C. (2 pts) If the network used store-and-forward routing, what would be the minimum latency transmitting one packet through the network? (Assume this packet is the only traffic on the network.)

- D. (2 pts) Now consider sending packet A from node 2 to node 0 and packet B from node 5 to 3 on the unidirectional ring interconnect shown below. Assuming the conditions from part A (32-bits send over a link per clock, worm-hole routing, same packet and flit sizes, both messages sent at the same time, packet A prioritized over packet B), what is the minimum latency for message A to arrive at its destination? Message B?



- E. (2 pts) **THIS QUESTION IS UNRELATED TO THE PREVIOUS ONES.** Consider a parallel version of the 2D grid solver problem from class. The implementation uses a 2D tiled assignment of grid cells to processors. (Recall the grid solver updates all the red cells of the grid based on the value of each cell's neighbors, then all the black cells). Since the grid solver requires communication between processors, you choose to buy a computer with a crossbar interconnect. Your friend observes your purchase and suggests there is another network topology that would have provided the same performance at a lower cost. What is it? (Why is the performance the same?)

### Problem 3: Concurrency Using Fine-Grained Locks

**NOTE: THIS QUESTION IS FOR EXTRA PRACTICE ONLY—IT WILL NOT BE GRADED BY THE STAFF ALTHOUGH SOLUTIONS WILL BE PROVIDED.**

**NOTE: this problem uses the same setup as Problem 1, but does not depend on its answer. See Problem 1 for the setup, but don't worry if you weren't able to answer Problem 1.**

- A. Consider an alternative implementation of the thread safe binary search tree from the previous problem that uses fine-grained (per-node) locking to synchronize the insert operations. The code below is a correct implementation of insert. Notice that it does not use hand-over-hand locking. **Assuming that insert is the only operation allowed on the tree**, describe why it's okay for the implementation to be in the state where it holds no locks at certain points during traversal.

```
struct Node {
    Node *left, *right;
    Lock* lock;
    int value;
};

Node* root; // root of tree, assume not null

void insertNode(Node* n, int value) {

    lock(n->lock);

    // assume n is not null
    if (value < n->value) {
        if (n->left == NULL) {
            n->left = createNode(value);
            unlock(n->lock);
        } else {
            unlock(n->lock);
            insertNode(n->left, value);
        }
    } else if (value > n->value) {
        if (n->right == NULL) {
            n->right = createNode(value);
            unlock(n->lock);
        } else {
            unlock(n->lock);
            insertNode(n->right, value);
        }
    }
    // insert won't be called with a duplicate element,
    // so there's no else case
}

void insert(int value) {
    insertNode(root, value);
}
```

B. (10 pts) Consider a system that processes a queue of tree manipulation commands.

```
insert(10);
insert(25);
print sum();
insert(24);
insert(100);
print sum();
insert(1);
...
```

The system can execute the commands concurrently as desired, provided that **the results of all sum operations** are consistent with the commands being executed in SERIAL ORDER. That is, in the example above, the first print must reflect the insertion of 10 and 25, but it must not reflect the insertion of 24, 100, and 1. Give an implementation that respects these semantics and enables as much concurrency as possible.

**This is a tricky problem. First try and find a solution that allows inserts() that come before sum() in the queue to execute concurrently with the sum(). We'll give good partial credit for that. For full credit (tricky) we're looking for a solution that allows sum() to execute concurrently with inserts() before and after it in the queue. Note your implementation will likely need to modify BOTH the implementation of insert() and sum().**

Starter code is given on the next page to give you room.



```

struct Node {
    Node *left, *right;
    int value;
    Lock lock;
};

Node* root; // root of tree

void insertNode(Node* n, int value) {

    if (value < n->value) {
        if (n->left == NULL) {
            n->left = createNode(value);
        } else {
            insertNode(n->left, value);
        }
    } else if (value > n->value)
        if (n->right == NULL) {
            n->right = createNode(value);
        } else {
            insertNode(n->right, value);
        }
    } // insert won't be called with a duplicate element, so there's no else case

void insert(int value) {
    insertNode(root, value);
}

int sumNode(Node* n) {
    if (n == null)

        return 0;

    Node* left;
    Node* right;
    left = n->left;
    right = n->right

    int leftSum;
    int rightSum;

    if (left)
        leftSum = sumNode(n->left);
    else
        leftSum = 0;
    if (right)
        leftSum = sumNode(n->right);
    else
        rightSum = 0;

    return n->value + leftSum + rightSum;
}

int sum() {
    return sumNode(root);
}

```