CS 7260
Scribe Notes: 12 November 2012
Scriber: Surabhi Potnis

The data structure used for GPS clock is a balanced binary search tree. The keys of this tree are the GPS virtual finish times. Each node of the tree contains the following fields :-

1. Key ie. virtual finish time of the last packet
2. Length
3. Drop
4. Area

Length, Drop and Area are determined from the ladder structure.

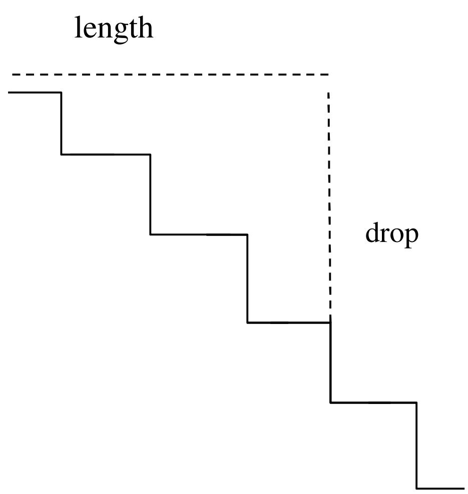

If one lists the keys (virtual finish times) of the nodes traversed in an in-order fashion they will be in non-decreasing order.

Area of the root = entire area

Area of the internal nodes depends on the area of their parent.

parent.area=left.area + right.area + (left.drop * right.length)

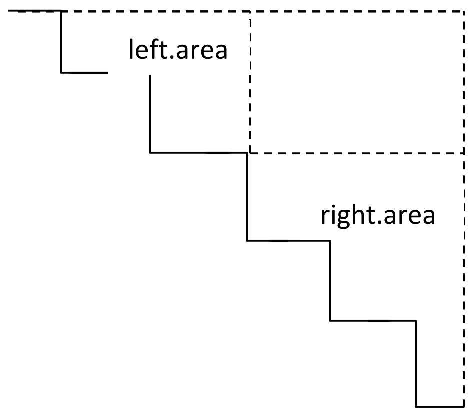

This equation will be considered in the balance function of the tree. Balancing is done by AV tree rotations.

Invariants for balancing the tree :-

Invariant 1 If one lists the keys (virtual finish times) of the nodes traversed in an in-order fashion they will be in non-decreasing order.

Invariant 2 : parent.area=left.area + right.area + (left.drop * right.length)

Movement of nodes on addition of packets : When new packets are added the virtual finish time of the node increases. To keep the tree balanced the position of this node must be changed. The old node is deleted and a new node with the new key is added. Deletion of the node affects the invariant conditions for the parent and ancestors but not siblings. Both insert and delete cause path of destruction, hence we have 2 paths of destruction of each node movement.

<u>Stack Reference pdf or Re-use distance</u>

This is another augmented data structure for GPS. It is based on the LRU (least recently used) cache replacement algorithm.

Consider the following sequence of memory accesses.

a b c a b c

This initial reuse distance is set to 0 or infinity depending on the convention we want to follow.

The reuse distance between the two ‘a’s is 3. Similarly the reuse distance between the two ‘b’s and two ‘c’s is also 3. We store these reuse distances in a histogram. The augmented data structure is used to store this histogram.

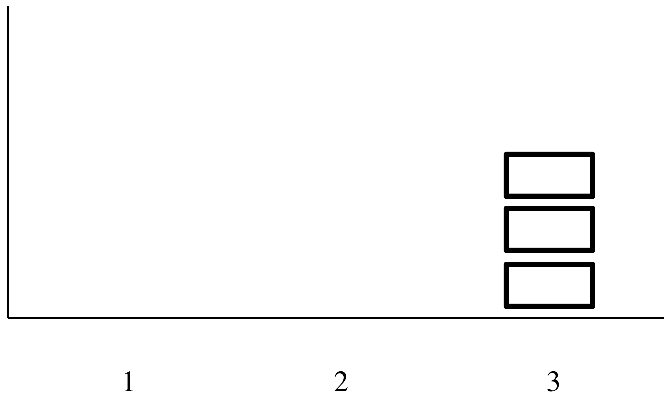

Strong Locality theorem : If the reuse distance histogram is non-increasing then LRU is the best non-lookahead cache replacement policy no matter how large the cache is.