Simulation Algorithms

Simulations are immensely useful, both for profit (design validation) and fun (games such as Angry Birds). Simulations can be optimized with the aid of data structures such as heap-backed priority queues, and by sorting algorithms.

Most simulations have a discrete clock, which means that the simulation timeline consists of small, equally-sized time quanta, and the universe changes at the beginning of each quanta. For example, a physical simulation of the balls on a billiard table would update the ball positions each 10 milliseconds. As long the time quanta are small, users will not realize the discrete nature of the simulation.

Discrete simulations must assume that world changes can be modeled as happening instantaneously, because the changes are computed at the beginning of each quanta. Under this assumption, we can speed up the simulation by figuring out the times when changes happen, instead of simulating every time quantum.

Events

The circuit simulation model in Problem Set 2 conveniently specified that gate output values change instantaneously (in real life, a gate's output voltage transitions continuously between the minimum voltage and the maximum voltage, over a period of time). This allowed our simulator to work by computing the effects caused by changes in gate output values, instead of simulating every time quantum. From now on, we will call these changes **events**.

For every event (gate output change), we looked at the gates whose inputs are connected to the gate's output and, if the input changes caused a change in the gates' output values, we created new events to account for these changes. In order to get the correct results during simulation, the simulator needed to consider the events in chronological order, so it used a priority queue to manage events. Events were added to the queue as they were created, and the simulation code popped the top event (with the minimum step) off the queue at each step.

Line-Sweep Algorithms

The sweep-line algorithm in Problem Set 3 involves simulating a vertical line that sweeps left-to-right through the wire layout, so we can treat it as a simulation problem. This time around, all the events can be computed directly from the input data, before starting the simulation. So we populate a Python list (array) with all the events.

```
1  class CrossVerifier(object):
2    def _events_from_layer(self, layer):
3      """Populates the sweep line events from the wire layer."""
4      left_edge = min([wire.x1 for wire in layer.wires.values()])
5      for wire in layer.wires.values():
6        if wire.is_horizontal():
7          self.events.append([left_edge, 0, wire.object_id, 'add', wire])
8        else:
9          self.events.append([wire.x1, 1, wire.object_id, 'query', wire])
```

The events are sorted according to their "time" (the x coordinate of the sweep line) in the $\_\_$init$\_\_$ method of CrossVerifier, and then $\_$compute$\_$crossings iterates over the sorted list and processes the events in "chronological" order (lines 9-10).

```
1   class CrossVerifier(object):
2     def _compute_crossings(self, count_only):
3       """Implements count_crossings and wire_crossings."""
4       if count_only:
5         result = 0
6       else:
7         result = self.result_set
8
9       for event in self.events:
10        event_x, event_type, wire = event[0], event[3], event[4]
11
12        if event_type == 'add':
13          self.index.add(KeyWirePair(wire.y1, wire))
14        elif event_type == 'query':
15          self.trace_sweep_line(event_x)
16          cross_wires = []
17          for kwp in self.index.list(KeyWirePairL(wire.y1),
18                                     KeyWirePairH(wire.y2)):
19            if wire.intersects(kwp.wire):
20              cross_wires.append(kwp.wire)
21          if count_only:
22            result += len(cross_wires)
23          else:
24            for cross_wire in cross_wires:
25              result.add_crossing(wire, cross_wire)
26
27      return result
```

The "hacked-together" code that ships with Problem Set 3 has two types of events: *add* and *query*. First off, all the horizontal wires are *added* to the range index, keyed according to their y coordinates. Then the sweep line moves across the wire layout from left to right and, every time it "hits" a vertical wire, a *query* event occurs. During a *query*, the range index is asked to provide all the horizontal wires whose y coordinates are between the y coordinates of the vertical wire's endpoints. This list of wires may contain false positives – wires that are completely to the left or completely to the right of the vertical wire, but it will definitely contain all the wires that intersect the vertical wire. Convince yourself that none of the other horizontal wires could possibly intersect the vertical wire.

Comparison Model I

We have learned in lecture that using the comparison model means we can't do searching in better than $\Omega(\log n)$ time, and we can't sort faster than $\Omega(N \log N)$. The advantage of the comparison model is that it can encompass arbitrarily complicated objects, as long as we can define an ordering relationship on them.

The events in Problem Set 3 are Python lists (arrays), because Python conveniently implements an ordering relationship on list, based on lexicographical ordering[1]. We'll call the lists used to represent events *vectors*, to distinguish them from the array that holds all the events.

The first element in an event's vector (main comparison criterion) is the event's x coordinate, so that events are processed left-to-right, modeling a vertical line that sweeps across the entire plane left-to-right.

The next two elements in an event's vector are used to break ties among events that happen at the same sweep line position. They are

1. a small integer representing the event type's priority; this ensures that *add* events at an x coordinate are processed before *query* events at that coordinate

2. a wire ID, which is a unique number; this ensures that the lexicographical comparison algorithm will never go past this element, so it will not attempt to compare Wire instance

Last, the inefficient sweep-line algorithm supplied in Problem Set 3 wants to process all the *add* events before any *query* event, so it sets the x coordinate for all the *add* events to be the minimum among the x coordinates of all wire endpoints. Together with setting the priority of *add* events set to 0, this ensures that *add* events come ahead of *query* events during comparisons.

```python
1  class CrossVerifier(object):
2    def _compute_crossings(self, count_only):
3      """Implements count_crossings and wire_crossings."""
4      if count_only:
5        result = 0
6      else:
7        result = self.result_set
8
9      for event in self.events:
10       event_x, event_type, wire = event[0], event[3], event[4]
11
12       if event_type == 'add':
13         self.index.add(KeyWirePair(wire.y1, wire))
14       elif event_type == 'query':
15         self.trace_sweep_line(event_x)
16         cross_wires = []
17         for kwp in self.index.list(KeyWirePairL(wire.y1),
18                                    KeyWirePairH(wire.y2)):
19           if wire.intersects(kwp.wire):
20             cross_wires.append(kwp.wire)
21         if count_only:
22           result += len(cross_wires)
23         else:
24           for cross_wire in cross_wires:
25             result.add_crossing(wire, cross_wire)
26
27     return result
```

---

[1]the same ordering used for words in a dictionary, e.g. Alyssa $<$ Andrew

Range Queries

The wire crossing algorithm supplied in Problem Set 3 adds all horizontal wires to a range index, which is later queried to find all the wires that can potentially intersect a vertical wire.

Range Index API

The range index concept is outlined in Problem 1 of Problem Set 3.

```
1  class RangeIndex(object):
2    """Array-based range index implementation."""
3
4    def __init__(self):
5      """Initially empty range index."""
6      self.data = []
7
8    def add(self, key):
9      """Inserts a key in the range index."""
10     if key is None:
11         raise ValueError('Cannot insert nil in the index')
12     self.data.append(key)
13
14   def remove(self, key):
15     """Removes a key from the range index."""
16     self.data.remove(key)
17
18   def list(self, first_key, last_key):
19     """List of values for the keys that fall within [first_key, last_key]."""
20     return [key for key in self.data if first_key <= key <= last_key]
21
22   def count(self, first_key, last_key):
23     """Number of keys that fall within [first_key, last_key]."""
24     result = 0
25     for key in self.data:
26       if first_key <= key <= last_key:
27         result += 1
28     return result
```

A range index supports ADD(x) and REMOVE(x) updates to build up a set of keys, and offers range queries over the keys. COUNT(l, h) returns the number of keys in the range index that belong to the $[l, h]$ range ($l \leq key \leq h$) and LIST(l, h) returns a list of the above-mentioned keys.

The code supplied in circuit2.py is optimized for size, not for speed, to convey the concept of a range index. Convince yourself that ADD runs in $O(1)$ time, but REMOVE, COUNT runs in $O(N)$ time, and LIST run in $O(N + K)$ time, where K is the number of keys in the list returned by LIST.

A first cut at optimizing range indexes would use a sorted array.

```
1  class BlitRangeIndex(object):
2    """Sorted array-based range index implementation."""
3
```

```
4    def __init__(self):
5      """Initially empty range index."""
6      self.data = []
7
8    def add(self, key):
9      """Inserts a key in the range index."""
10     if key is None:
11         raise ValueError('Cannot insert None in the index')
12     self.data.insert(self._binary_search(key), key)
13
14   def remove(self, key):
15     """Removes a key from the range index."""
16     index = self._binary_search(key)
17     if index < len(self.data) and self.data[index] == key:
18       self.data.pop(index)
19
20   def list(self, low_key, high_key):
21     """List of values for the keys that fall within [low_key, high_key]."""
22     low_index = self._binary_search(low_key)
23     high_index = self._binary_search(high_key)
24     return self.data[low_index:high_index]
25
26   def count(self, low_key, high_key):
27     """Number of keys that fall within [low_key, high_key]."""
28     low_index = self._binary_search(low_key)
29     high_index = self._binary_search(high_key)
30     return high_index - low_index
31
32   def _binary_search(self, key):
33     """Binary search for the given key in the sorted array."""
34     low, high = 0, len(self.data) - 1
35     while low <= high:
36       mid = (low + high) // 2
37       mid_key = self.data[mid]
38       if key < mid_key:
39         high = mid - 1
40       elif key > mid_key:
41         low = mid + 1
42       else:
43         return mid
44     return high + 1
```

This implementation is a bit more complicated, but runs a lot faster. COUNT runs in $O(\log N)$ time, and LIST runs in $O(\log(N) + K)$ time, but in exchange ADD runs in $O(N)$ time. REMOVE still runs in $O(N)$ time, like before. In a query-heavy workload, `BlitRangeIndex` will vastly outperform `RangeIndex`. If there are a lot of updates, the performance won't be too bad, because the $O(N)$ operations in ADD and REMOVE, as well as the $O(K)$ operation in LIST, are built-in Python functions, which are most likely implemented in C, and therefore have very small constant factors. For many practical cases, this implementation will do very well. However, on really large data sets, this code will be out-performed by an implementation of an algorithm with better

asymptotic performance.

Comparison Model II

We can use the comparison model to introduce an ordering relationship between all the wires in the range index – horizontal wires should be ordered by their y coordinate.

The first instinct when implementing this ordering relationship would be to augment the `Wire` class with the right magic methods, so `Wire` instances can be compared using the standard Python operators $<$, $<=$, $>$ and $>=$. However, this approach is "dirty" – What if we later want to order vertical wires by their x coordinate? How about ordering all the wires by their names? Moreover, the approach falls apart when it comes time to do a range query – how would we create keys representing the y coordinates of a vertical wire's endpoints? "Fake" `Wire` instances come to mind, but that's also quite hackish.

Instead, we create a new kind of objects, `KeyWirePair`, which are used as the range index keys.

```
1   class KeyWirePair(object):
2     """Wraps a wire and the key representing it in the range index.
3
4     Once created, a key-wire pair is immutable."""
5
6     def __init__(self, key, wire):
7       """Creates a new key for insertion in the range index."""
8       self.key = key
9       self.wire = wire
10      self.wire_id = wire.object_id
```

A `KeyWirePair` instance that is inserted into the range index has its `key` set to the wires' y coordinate. Comparisons use the `key` field as the primary criterion, and then use `wire_id` (a wire's unique ID) to break ties. Two keys with the same `wire_id` represent the same wire, so they are equal. It doesn't make sense to insert the same wire into the range index twice, so the data structure used to implement the range index doesn't have to deal with equal keys.

```
1   class KeyWirePair(object):
2     def __lt__(self, other):
3       return (self.key < other.key or
4               (self.key == other.key and self.wire_id < other.wire_id))
5
6     def __le__(self, other):
7       return (self.key < other.key or
8               (self.key == other.key and self.wire_id <= other.wire_id))
9
10    def __gt__(self, other):
11      return (self.key > other.key or
12              (self.key == other.key and self.wire_id > other.wire_id))
13
14    def __ge__(self, other):
15      return (self.key > other.key or
16              (self.key == other.key and self.wire_id >= other.wire_id))
```

```
17
18    def __eq__(self, other):
19      return self.key == other.key and self.wire_id == other.wire_id
20
21    def __ne__(self, other):
22      return self.key == other.key and self.wire_id == other.wire_id
```

Note the special method names (e.g., `__le__` means *less than or equal to* and is used to implement the <= operator).

`KeyWirePair` instances used for range queries have the `wire` field set to `None`. To facilitate range queries, we subclass `KeyWirePair` with `KeyWirePairL` and `KeyWirePairH`, which are meant to be used as the low and high keys in range queries.

```
1  class KeyWirePairL(KeyWirePair):
2    def __init__(self, key):
3      self.key = key
4      self.wire = None
5      self.wire_id = -1000000000
6
7  class KeyWirePairH(KeyWirePair):
8    def __init__(self, key):
9      self.key = key
10     self.wire = None
11     self.wire_id = 1000000000
```

Asides from setting `wire` to `None`, the special classes set the `wire_id` field to special values that are guaranteed to be smaller / greater than all the IDs of real wires. This means that a `KeyWirePairL` with a certain key (y coordinate) will be considered lower than all the keys for wires with the same y coordinate and, conversely, a `KeyWirePairH` will be higher than the keys of wires with the same y coordinate. This implementation trick eliminates some boundary cases for range queries, because the keys in the index are guaranteed to not be equal to the keys used in queries.

The need of being able to bound wire IDs in `KeyWirePairL` and `KeyWirePairH` pushed us away from Python's built-in `id()` function, which provides unique IDs for all Python objects. Instead, the verifier implements its own ID scheme for `Wire` instances.

```
1  class Wire(object):
2    def __init__(self, name, x1, y1, x2, y2):
3      # Normalize the coordinates.
4      if x1 > x2:
5        x1, x2 = x2, x1
6      if y1 > y2:
7        y1, y2 = y2, y1
8
9      self.name = name
10     self.x1, self.y1 = x1, y1
11     self.x2, self.y2 = x2, y2
12     self.object_id = Wire.next_object_id()
13
14    # Next number handed out by Wire.next_object_id()
```

```
15    _next_id = 0
16
17    @staticmethod
18    def next_object_id():
19      """Returns a unique numerical ID to be used as a Wire's object_id."""
20      id = Wire._next_id
21      Wire._next_id += 1
22      return id
```

Wire IDs are consecutive non-negative integers, so it's reasonable to assume they will be bounded by 1,000,000,000 (a billion), because storing a billion wires would require a lot of RAM.

LIST **Range Queries in BSTs**

Given a binary search tree (BST), we want to implement LIST($tree, l, h$) which produces a list of all the keys between l and h ($l \leq key \leq h$), in $O(\log(N) + K)$ time, where N is the number of keys in the tree, and K is the number of keys output by LIST.

We will convince ourselves that the pseudo-code below will do the trick.

LIST($tree, l, h$)

1   $lca = $ LCA($tree, l, h$)
2   $result = [\,]$
3   NODE-LIST($lca, l, h, result$)
4   **return** $result$

First off we identify the smallest subtree that contains all the keys between l and h. Intuitively, LCA computes the root of that subtree. Formally, LCA computes the lowest-common ancestor of nodes with keys l and h. If l and h do not exist in the tree, LCA returns the lowest-common ancestor of the two nodes that would be created by inserting l and h.

LCA($tree, l, h$)

1   $node = tree.root$
2   **until** $node == $ NIL or ($l \leq node.key$ and $h \geq node.key$)
3   **if** $l < node.key$
4           $node = node.left$
5   **else**
6           $node = node.right$
7   **return** $node$

Once the subtree's root is found by LCA, NODE-LIST performs an in-order traversal of the subtree, but prunes the subtrees that are guaranteed to contain keys outside the $[l, h]$ range.

NODE-LIST($node, l, h, result$)

```
1   if node == NIL
2       return
3   if node.key ≥ l
4       NODE-LIST(node.left, l, h, result)
5   if l ≤ node.key and node.key ≤ h
6       ADD-KEY(result, node.key)
7   if node.key ≤ h
8       NODE-LIST(node.right, l, h, result)
```

6.006 Introduction to Algorithms
Fall 2011