Distributed Systems Part II
Solution to Exercise Sheet 7

Basic

1 Iterative vs. Recursive Lookup

a) In the recursive lookup there is no difference between a request originating at the node and a request being forwarded, only once the lookup is finished we need to care about forwarding the result to the previous node or returning it to the caller. This allows the same lookup logic to be reused for both. Furthermore, if the response is returned through the same path as the request was sent through, then the intermediate nodes can cache the result, potentially speeding up future lookups and distributing the load of a popular item.

b) Recursive lookups can easily be misused to mount Denial-of-Service attacks, since a single request message from an originating node is forwarded over multiple hops, each hop multiplies the impact this message had on the network. Potentially the attacker's bandwidth is multiplied by the number of hops the request is routed through. Furthermore, if the result is returned over the same path as the request, then the attacker is hidden behind a number of hops and the victim only sees traffic originating from the last hop.

2 Building a set of Hash functions

The salted hashing function derivation allows random access to any of the derived hashing functions without having to recompute the intermediate functions, as is the case in the iterative hashing function derivation scheme.

3   Multiple Skiplists

a) Figure 1 shows the structure of a multi-skiplist with 8 nodes and 3 levels. Notice that each of the lists would wrap around at the ends.

b) Unlike in the single skiplist each node now has a constant degree of $2 \cdot (l+1)$, i.e., on each level it has a right and a left neighbor, including the simple list at $l = 0$.

**c)** The number of hops is still $O(log(n))$, just like the simple skiplist.

Figure 1: A multi-skiplist with 3 levels and 8 nodes