11.4, 2(d)

Number of runs = $10,000,000/320 = 31250$
Number of passes (excluding Pass 0) = $\lceil log_8 31250 \rceil = 5$ (we have eight "input" buffers)
Number of blocks (32 pages each) we read per pass = $10,000,000/32 = 312,500$
Number of blocks (64 pages each) we write per pass = $10,000,000/64 = 156,250$

Total time taken by one pass	=	(average seek time + average rotational delay) * number of blocks read and write + transfer time for all pages
	=	$[(10+5) * (312,500 + 156,250) + 2 * 10,000,000]ms$
	=	$2.7 * 10^4 s$

Total time = $5 * 2.7 * 10^4 = 13.5 * 10^4 s \approx 37.5 hours$

12.4, 4

Number of pages of relation R = $10,000/10 = 1000$ pages
Number of pages of relation S = $2,000/10 = 200$ pages

Cost of joining R and S using hash join = $3 * (1000 + 200) = 3600$ pages

This cost will remain the same as long as the smaller relation (i.e. S) fits into the buffer pool.
We want $B > \lceil \sqrt{200f} \rceil$. If f = 1.2, the minimum number of B is 16.

14.4, 2(d)

Cost metric is the number of I/O.

Number of pages containing tuples that meet the condition E.title = 'CFO'
= 10,000 * 10% = 1,000 pages.

Number of pages containing tuples that meet the condition E.title = 'CFO' and E.dname='Toy'
= 10,000 * 5% = 500 pages.

Use the B+ tree to find the first page that contains tuple with condition E.title = 'CFO'.
Since the index is clustered, we will do a linear scan. The cost of the scan is 1,000 pages.
The cost of the output is 500 pages. The total cost is 1,500 pages + # I/O to traverse the
B+-Tree from the root to the leaf.