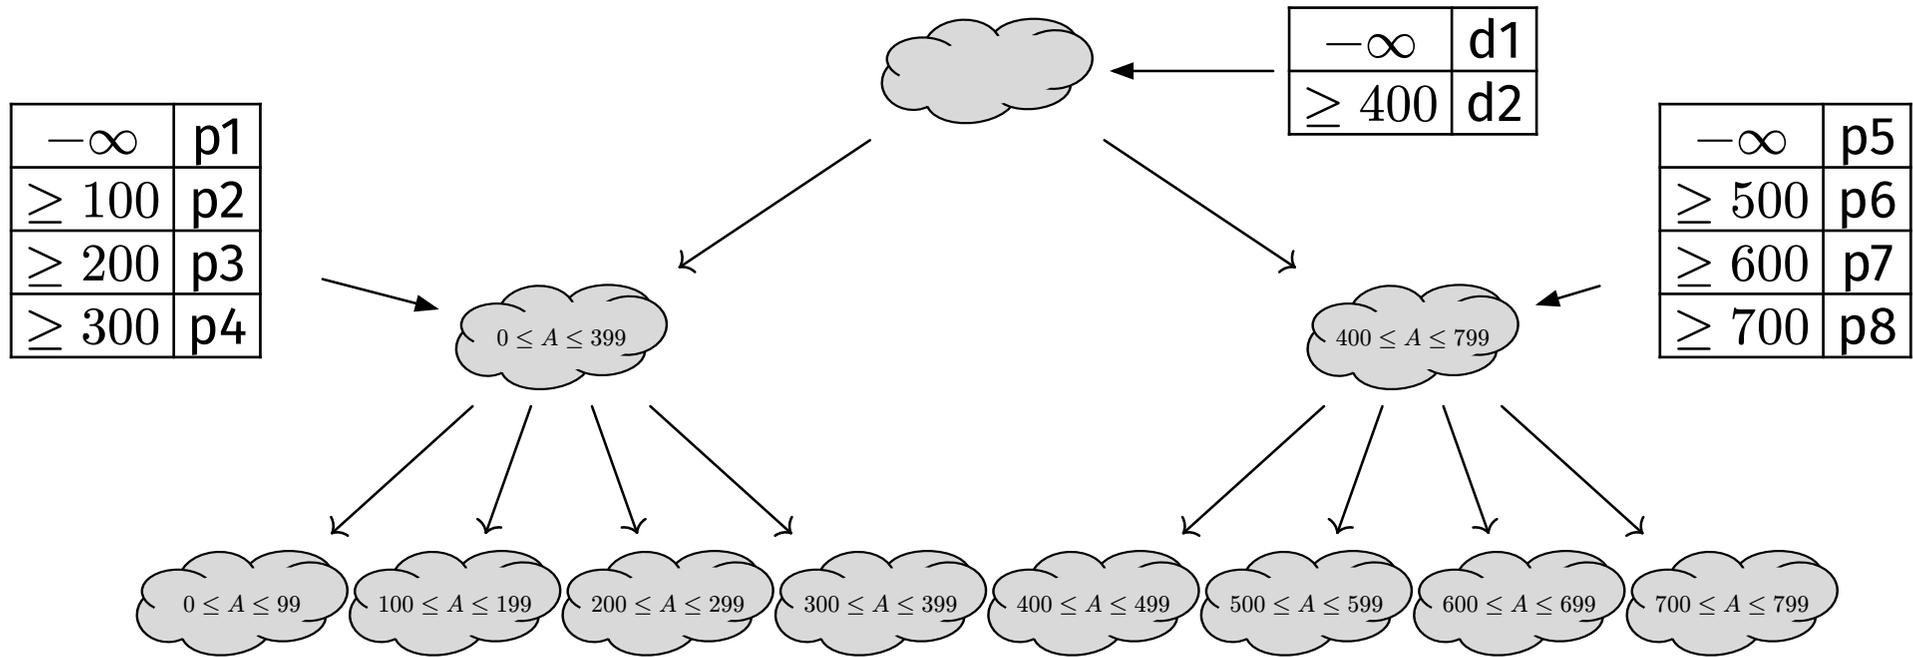


# **HASH-BASED INDEXES, COMPLEX KEYS**

CSE 4/562: Database Systems | Lecture 9

---

**DB. Sys.: T.C.B.:** Ch. 8.3-8.4, 14.1-14.2, 14.4



```
class TreeIndexScan:  
    def init():  
        # find leaf page for first record in result range  
  
    def next():  
        while current_key < last_key:  
            yield current_record  
            # step to next record
```

TreeIndexScan( $R$ ,  $\theta$ )

- IO:  $O(\log(|R|) + |\sigma_{\theta}(R)|)$
- Memory:  $O(1)$
- $\theta$ : Any 1 attribute, equality or range<sup>1</sup>

$\sigma_{\theta}(R)$

- IO:  $O(|R|)$
- Memory:  $O(1)$
- $\theta$ : Any condition at all

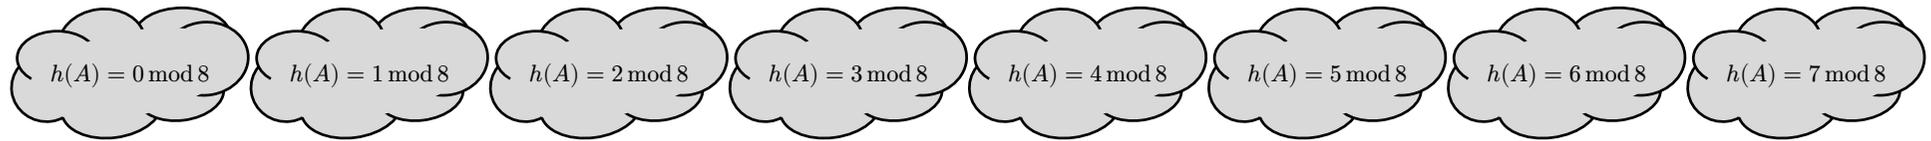
---

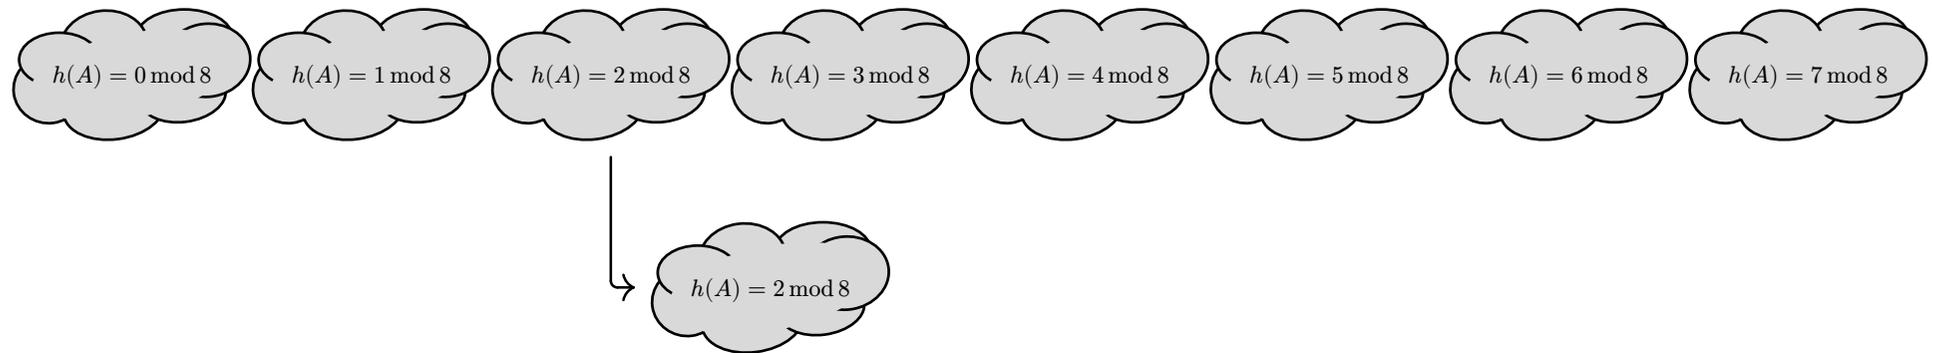
<sup>1</sup>We generalize this slightly later

A hash function  $h(k)$  is a function that is:

- **deterministic:** The same  $k$  always produces the same output.
- **(pseudo-)random:** Different  $k$ s produce uncorrelated outputs, even if the inputs are correlated.

$h(k) \bmod N$  gives you a random number in  $[0, N)$





With...

- $|R|$  records
- $P$  records per page
- $B$  buckets

---

In expectation, we will have  $\frac{|R|}{B}$  records per bucket.

In expectation each bucket will be  $\frac{|R|}{B \cdot P}$  pages.

- Each lookup/insert is an expected  $\left\lceil \frac{|R|}{B \cdot P} \right\rceil$  IOs.

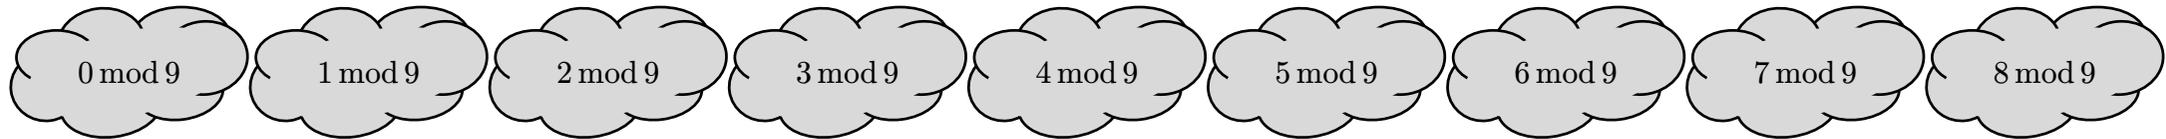
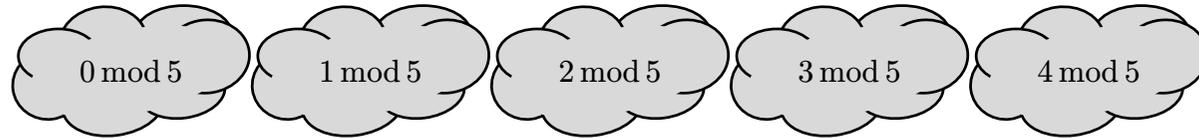
Ideally,  $\lceil \frac{|R|}{B \cdot P} \rceil = 1$  (i.e.,  $\frac{|R|}{B \cdot P} < 1$ )

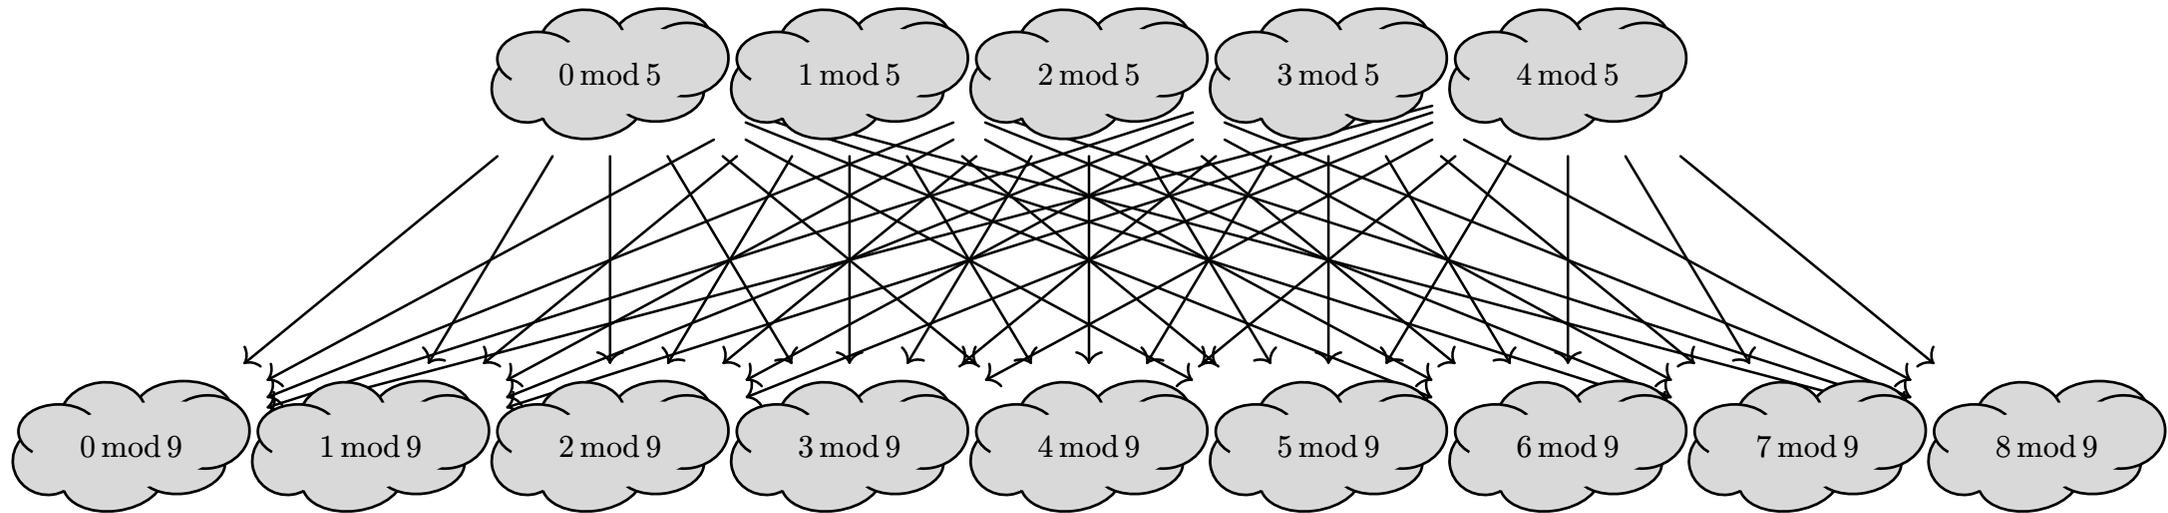
**Problem:** But  $|R|$  changes over time!

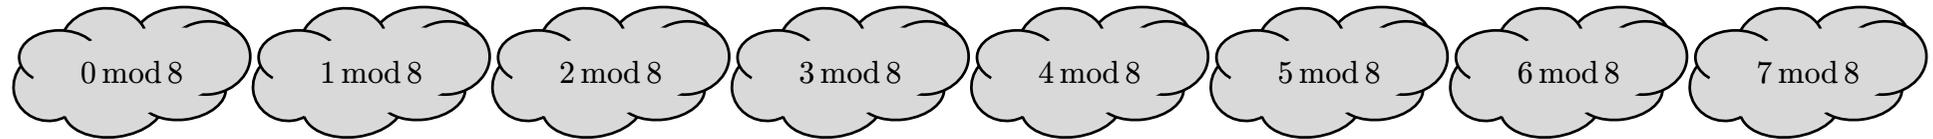
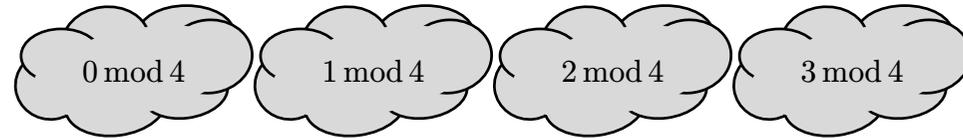
Ideally,  $\lceil \frac{|R|}{B \cdot P} \rceil = 1$  (i.e.,  $\frac{|R|}{B \cdot P} < 1$ )

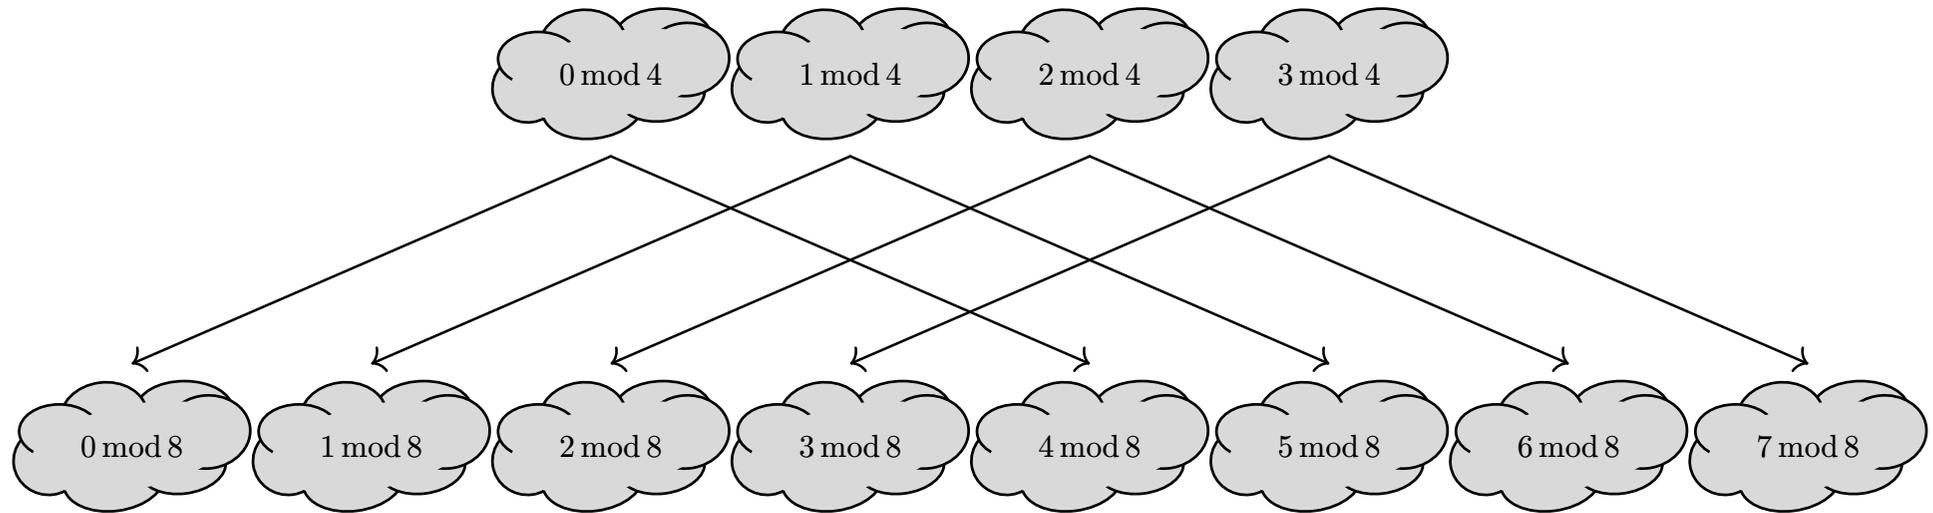
**Problem:** But  $|R|$  changes over time!

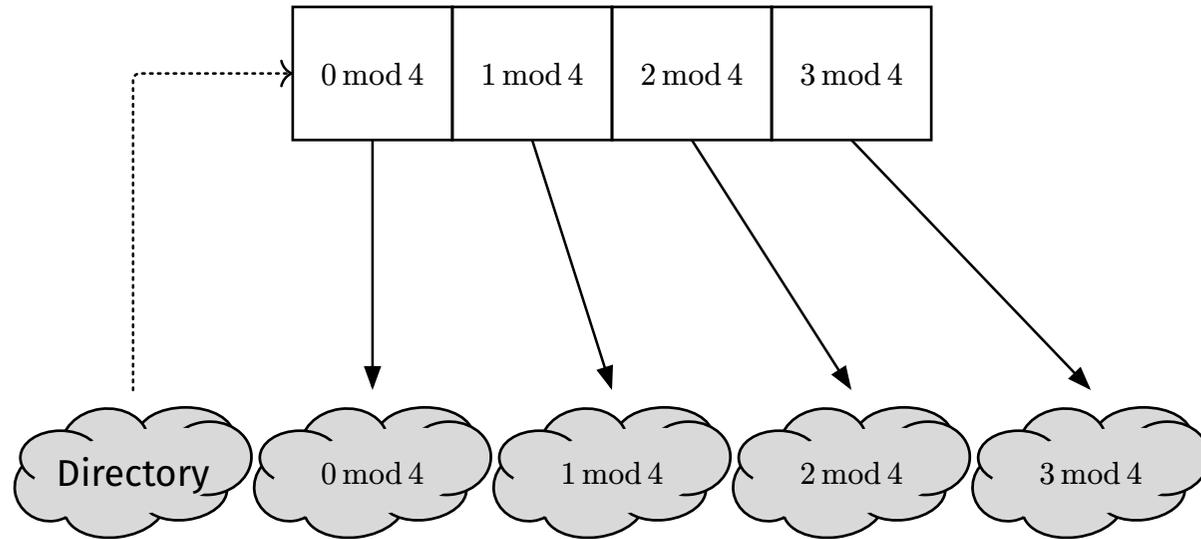
- We can't control  $P$
- Our only option is to change  $B$

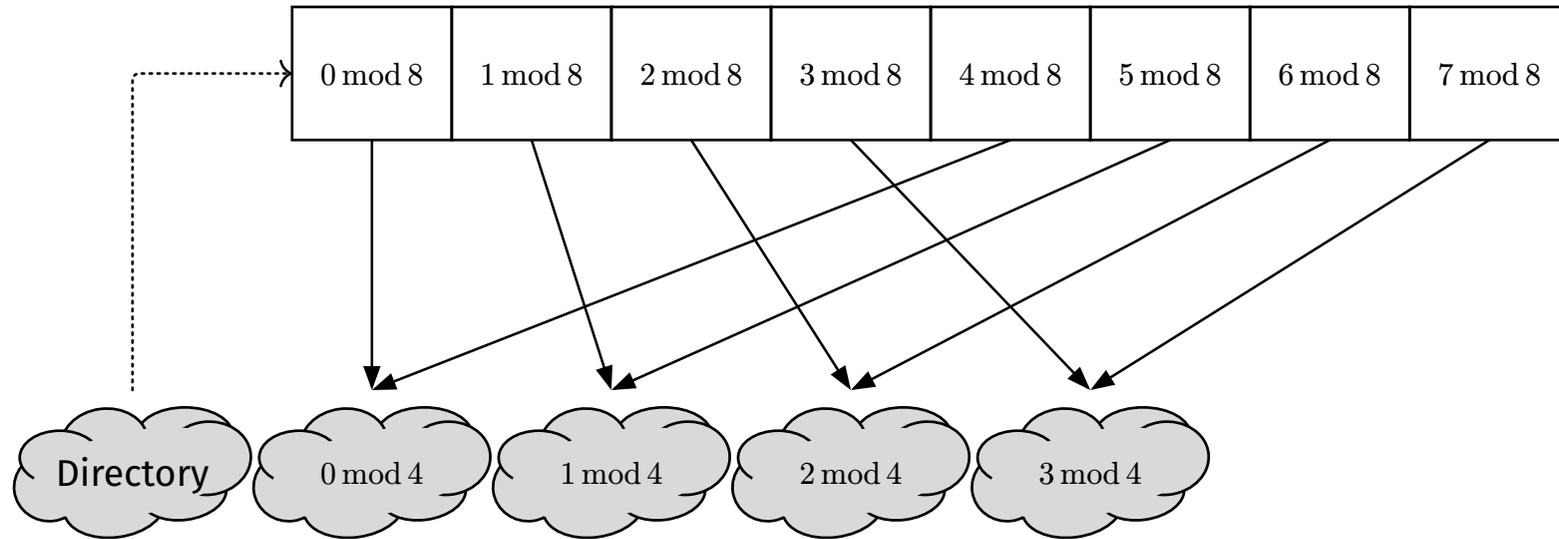


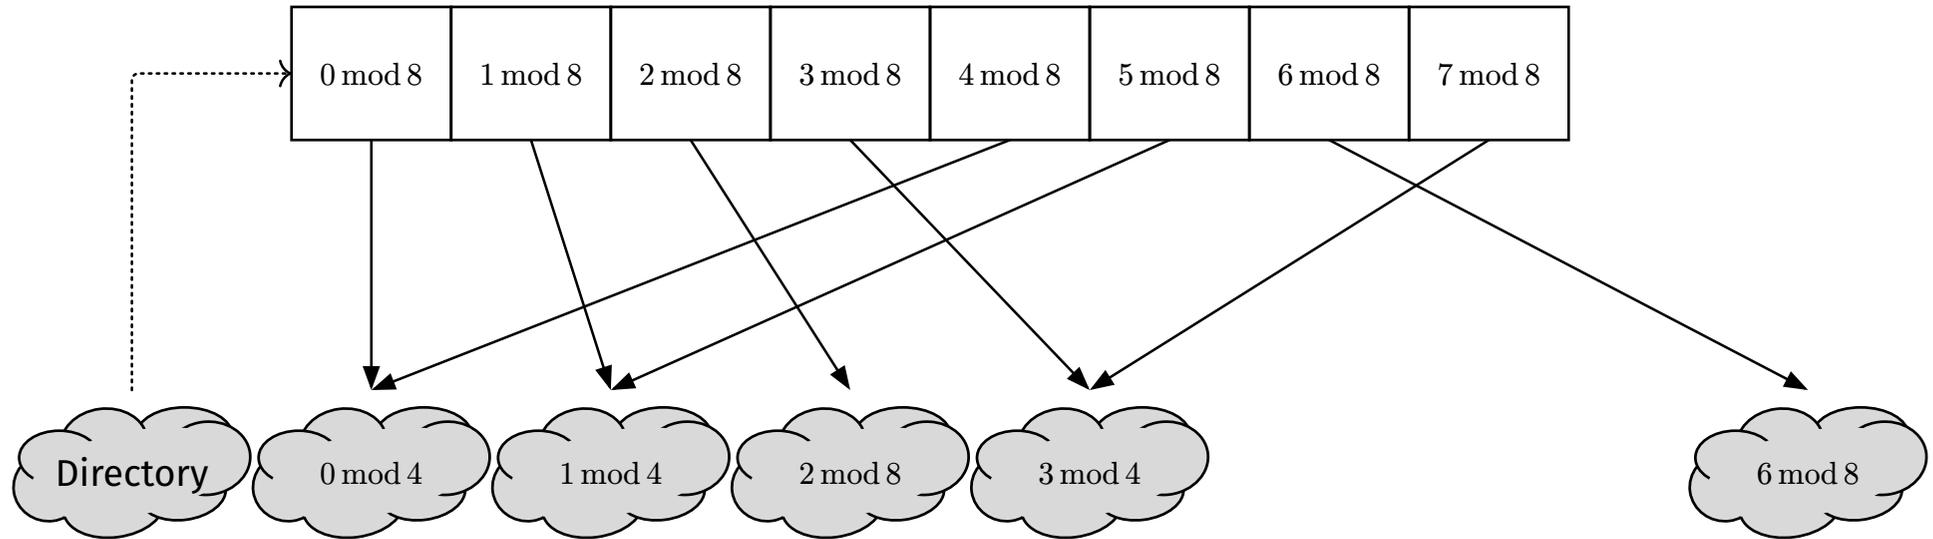


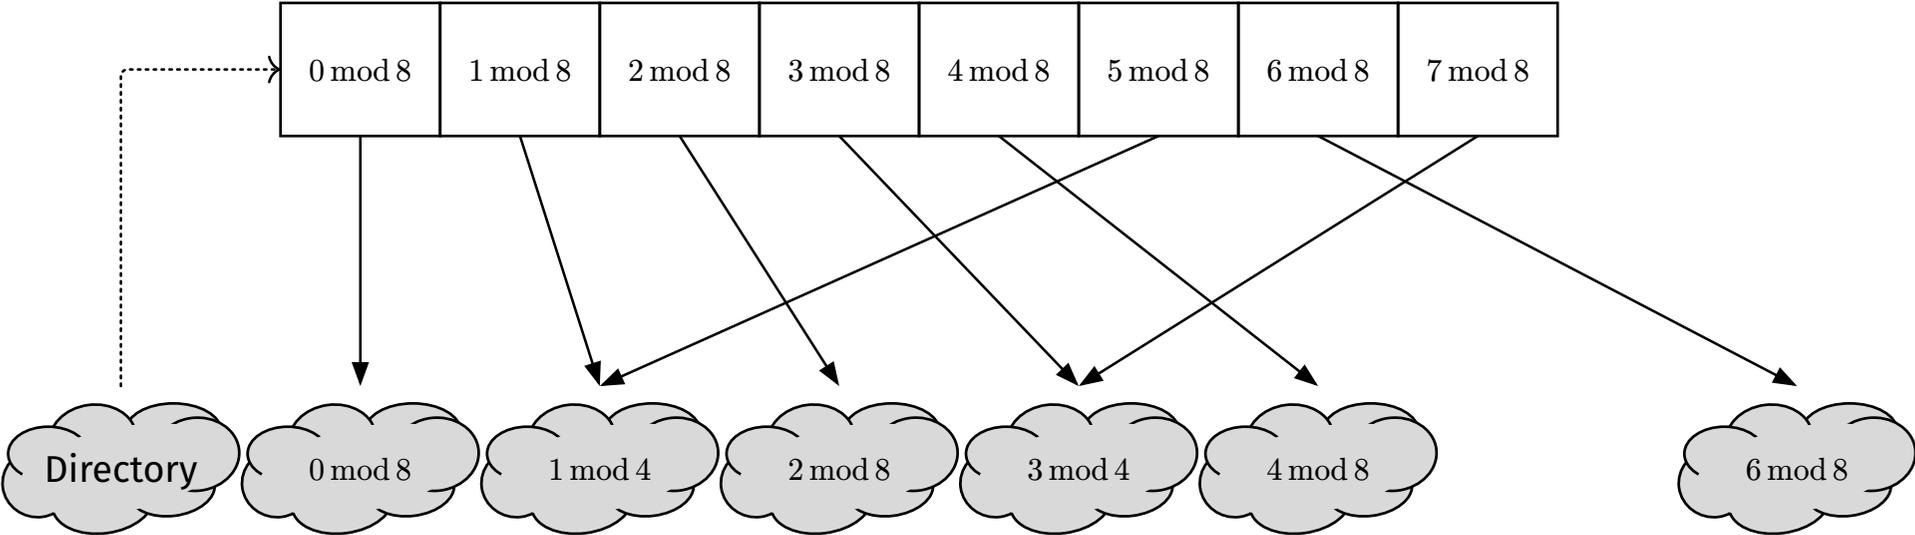


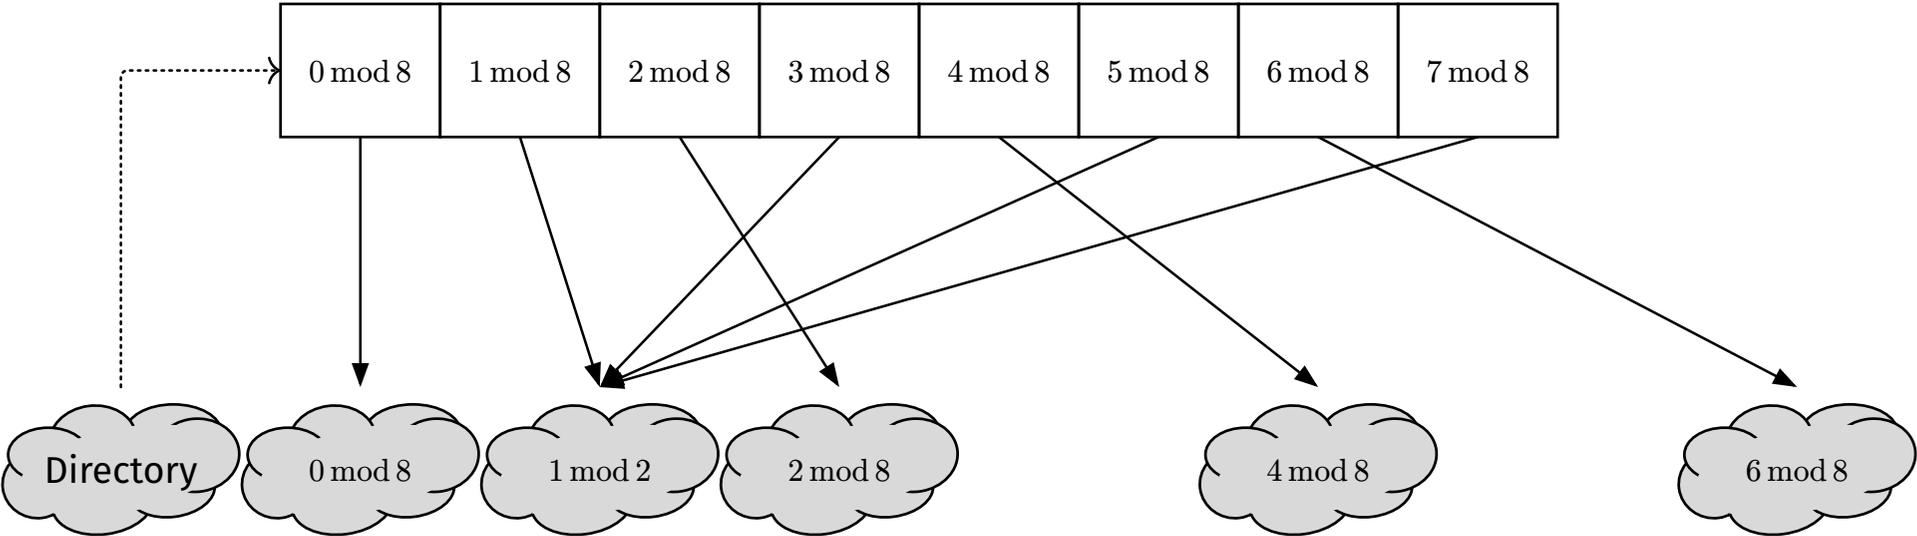












## **Split Page $i$ from $\text{mod } N$ to $\text{mod } 2N$**

1. Double the size of the directory if it is not already at least  $2N$  entries
2. Rehash everything that was  $i \bmod N$  to  $i \bmod 2N$  and  $i + N \bmod 2N$
3. Update directory pointers  $i$  and  $i + N$

## **Merge pages $i$ and $i + N$ from $\text{mod } 2N$ to $\text{mod } N$**

1. Merge everything that was  $i \bmod 2N$  and  $i + N \bmod 2N$  into  $i \bmod N$ .
2. Update directory pointer  $i + N$  (and  $i + N + 2N, \dots$  if necessary)
3. Halve the size of the directory if possible

The total IO cost of doubling is still approximately  $O(N)$ ,  
but the cost is more spread out.

**Supporting multiple keys**

To answer  $A = 1 \wedge B = 2 \wedge C = 3$ :

Sort data with  $(A, B, C)$  as a key, **lexically** sorted.

- All records with  $A = 1$  are contiguous
- All records with  $A = 1 \wedge B = 2$  are contiguous
- All records with  $A = 1 \wedge B = 2 \wedge C = 3$  are contiguous
- All records with  $A = 1 \wedge B = 2 \wedge C \geq 3$  are contiguous

A lexically sorted tree index supports:

- Point lookups of the key
- Range lookups of all records matching an exact **prefix** of the key
- As above, but where the last attribute in the prefix is a range query

**How about multiple sort  
orders?**

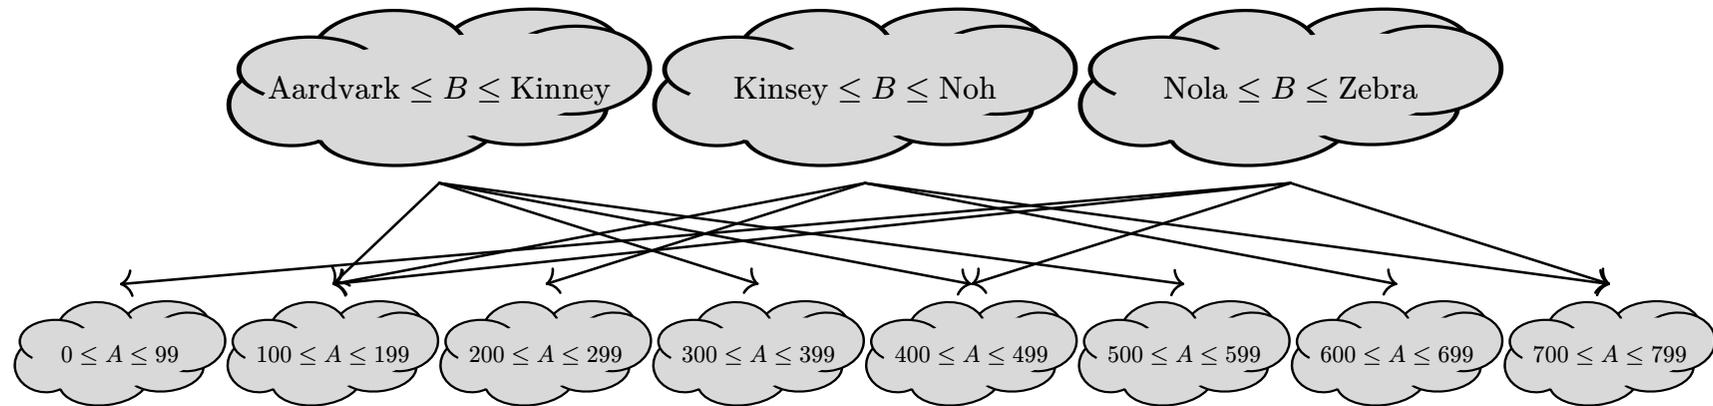
Can we build a single B+ Tree that supports both of these?

```
SELECT * FROM R WHERE R.A > 3
```

and

```
SELECT * FROM R WHERE R.B = 2
```

**Idea:** Leaves point to pages (and optionally record IDs) holding matching records.



## Clustering

- **Clustered Index:** Data layout is fully correlated with index attribute
- **Unclustered Index:** Data layout is not fully correlated with index attribute
  - Often called a “Secondary” index

## Density

- **Dense Index:** Index stores an entry for each record.
  - (key, page\_id, slot\_id)
- **Sparse Index:** Index stores references to ranges of records.
  - (lower\_bound, upper\_bound, page\_id)

How about ...  $\times_{\theta} S'$ ?

$$R \bowtie_{R.B > S.B} S$$

$$R \bowtie_{S.\text{low} < R.B < S.\text{high}} S$$

**Goal:** Compute  $R \bowtie_{R.B=S.B} S$

1. Build a hash table over  $S$
2. For each row  $r \in R$  use the hash table to find rows  $s \in S$  where  $r.B = s.B$

**Goal:** Compute  $R \bowtie_{R.B=S.B} S$

1. Build a ~~hash table~~ *index* over  $S$
2. For each row  $r \in R$  use the ~~hash table~~ *index* to find rows  $s \in S$  where  $r.B = s.B$

**Goal:** Compute  $R \bowtie_{R.B > S.B} S$

1. Build an *index* over  $S$
2. For each row  $r \in R$  use the *index* to find rows  $s \in S$  where  $r.B > s.B$

**Goal:** Compute  $R \bowtie_{R.B > S.B} S$

1. ~~Build an *index* over  $S$~~  (use an index that already exists)
2. For each row  $r \in R$  use the *index* to find rows  $s \in S$  where  $r.B > s.B$

**Goal:** Compute  $R \bowtie_{\theta} S$

```
for r in R:  
    for s in IndexScan(S,  $\theta$ [r]):  
        yield r + s
```

- Checkpoint 2 posted; Autolab up
- Checkpoint 1 solutions posted
- Quiz 1, 2 results posted