

JOIN AND AGGREGATION ALGORITHMS

CSE 4/562: Database Systems | Lecture 4

DB. Sys.: T.C.B.: 15.2-15.5

Quiz!

Recap

	Filter	Merge	Derive
Column	$\pi_{A,B,\dots}(R)$ “Project”	$R \times S$ “Cartesian Product”	$\pi_{C=A+B}(R)$ “Project”
Row	$\sigma_{A<3}(R)$ “Select” or “Filter”	$R \cup S$ “Union”	$\Sigma_{A,B=COUNT()}(R)$ “Aggregate”

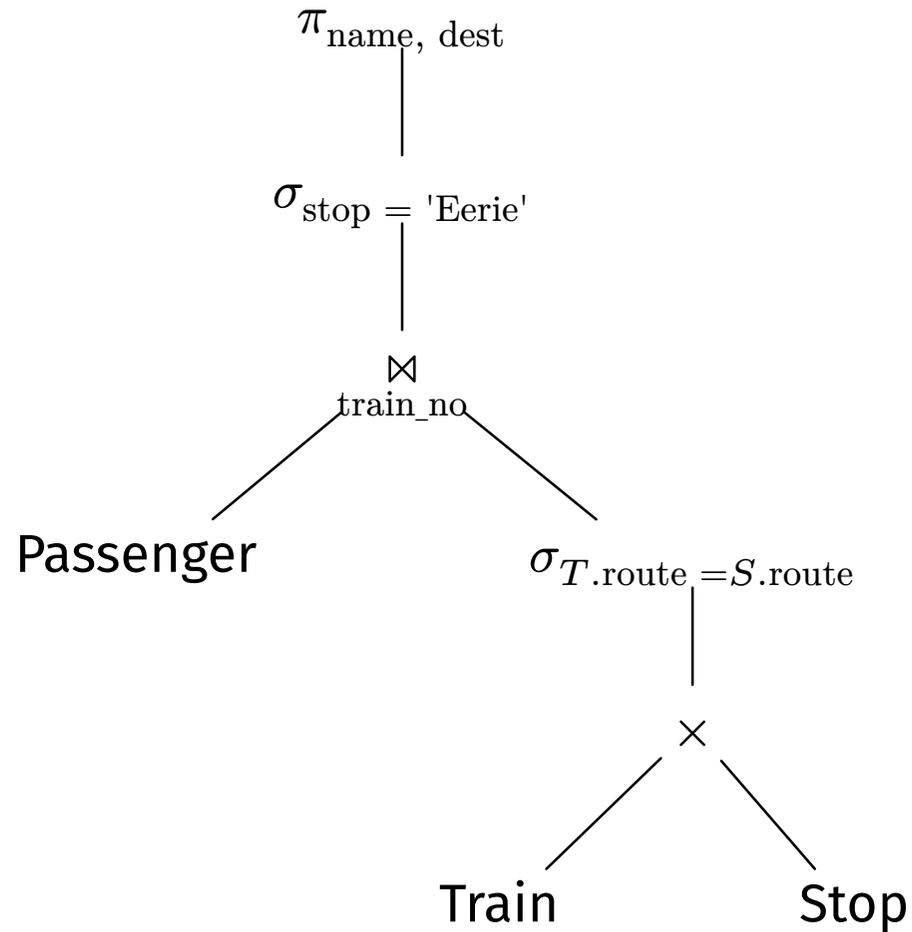
Passenger	Name	Train_no	Dest
	Athena	48	Buffalo
	Bragi	239	Albany
	Cerberus	241	Hudson

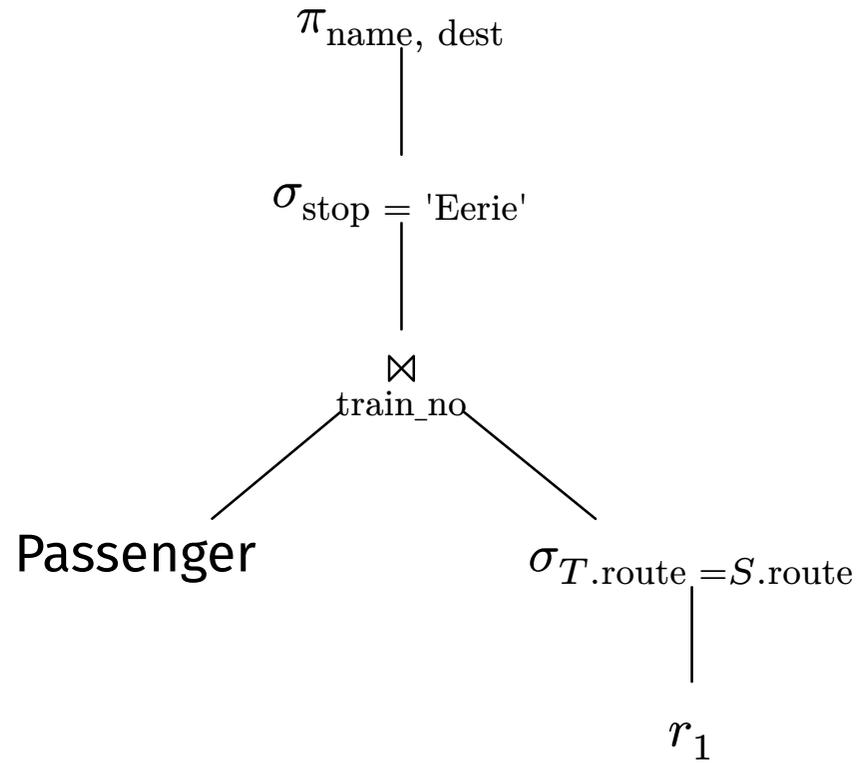
Train	Train_no	Route
	48	Lake Shore Ltd.
	239	Emp. Service
	241	Emp. Service

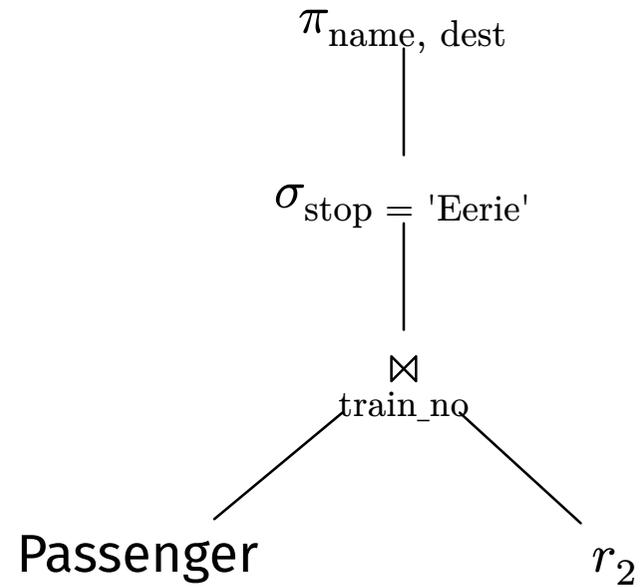
Stop	Route	Stop
	Emp. Service	NYC
	Emp. Service	Yonkers
	Emp. Service	Hudson
	Emp. Service	Albany
	Lake Shore Ltd.	Chicago
	Lake Shore Ltd.	Erie
	Lake Shore Ltd.	Buffalo

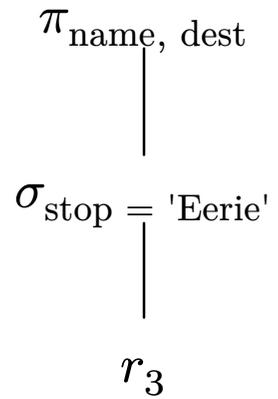
“Find me every passenger on a train that stops in Erie, and their destination”

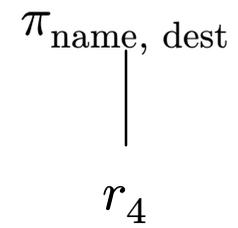
```
SELECT name, dest
FROM Passenger P, Train T, Stop S
WHERE stop = 'Erie'
      AND P.train_no = T.train_no
      AND T.route = S.route
```





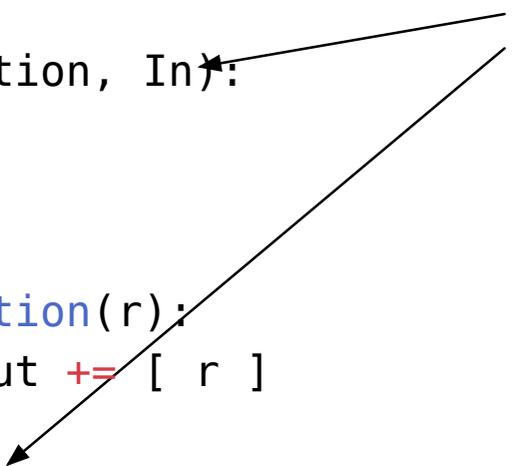






```
def filter(condition, In):  
    output = []  
  
    for r in In:  
        if condition(r):  
            output += [ r ]  
  
    return output
```

Big!



The model

- I : A limited “Internal” memory that starts off empty.
 - $|I|$ is the Memory Complexity
- E : An infinite “External” memory.

Actions

- $I[42] \leftarrow \text{Compute}(I[0], I[1], \dots)$: Adds to Runtime Complexity
- $I[12] \leftarrow \text{Read}(E[974])$: Adds to IO Complexity
 - “Read Cost”: Just reads
- $E[974] \leftarrow \text{Write}(I[12])$: Adds to IO Complexity
 - “Write Cost”: Just writes

The “IO Cost” or “IO Complexity” is Read cost + Write cost

- **Runtime Complexity:** How fast is the algorithm for a given input?
- **Memory Complexity:** How much memory do we need to run the algorithm on a given input?
- **IO Complexity:** How much data gets moved between layers of the hierarchy?

We can measure complexity...

- exactly ($3N$ pages of IO)
- asymptotically ($O(N)$ memory)

Operator-at-a-Time

Input: Complete relations

Output: Complete relation

Tradeoffs:

- Easy to implement
- Very low spatial locality (relations that don't fit in memory need to be on disk)

The Pull Model (Volcano, Iterators)

Input: Iterators over relations

Output: Iterator over relation

Tradeoffs:

- Much better spatial locality

```
class Iterator:  
    def next(self) -> Option[Record]:  
        pass
```

Operator iterators compose and the final result is read off the final iterator:

```
def print_result(query_iterator):  
    while (row := query_iterator.next()) is not None:  
        print(row)
```

Operator	Operator-at-a-time		Pull Evaluation	
	Memory	IO	Memory	IO
Table	n/a	n/a	$O(1)$	$O(n)$
Project (π)	$O(1)$	$O(n)$	$O(1)$	$O(1)$
Filter (σ)	$O(1)$	$O(n)$	$O(1)$	$O(1)$
Union (\cup)	$O(1)$	$O(n)$???	???
Product (\times)	$O(1)$ or $O(n)$	$O(n^2)$ or $O(n)$???	???
Join (\bowtie)	$O(1)$ or $O(n)$	$O(n^2)$ or $O(n)$???	???
Aggregate (Σ)	???	???	???	???

Product and Join Algorithms

```
class UnionIterator:  
    def __new__(self, in1, in2):  
        self.in1 = in1  
        self.in2 = in2  
  
    def next(self) -> Option[Record]  
        ???
```

```
class UnionIterator:
    def __new__(self, in1, in2):
        self.in1 = in1
        self.in2 = in2

    def next(self) -> Option[Record]
        if (row := self.in1.next()) is not None:
            return row
        else:
            return self.in2.next()
```

```
class UnionIterator:
    def __new__(self, in1, in2):
        self.in1 = in1
        self.in2 = in2

    def next(self) -> Option[Record]
        if (row := self.in1.next()) is not None:
            return row
        else:
            return self.in2.next()
```

What is the...

- Memory complexity?
- IO Complexity?

```
def product(In1, In2):  
    for r in In1:  
        yield r  
    for s in In2:  
        yield s
```

```
class ProductIterator:  
    def __new__(self, in1, in2):  
        self.in1 = in1  
        self.in2 = in2  
  
    def next(self) -> Option[Record]  
        ???
```

```
class ProductIterator:
    def __new__(self, in1, in2):
        self.in1 = in1
        self.in2 = in2

        self.in2_temp = [row for row in in2]
        self.i = 0
        self.in1_row = in1.next()

    def next(self) -> Option[Record]
        if self.i >= len(self.in2_temp):
            self.i = 0
            self.in1_row = self.in1.next()
        row = self.in1_row + self.in2_temp[i]
        i += 1
        return row
```

```
class ProductIterator:
    def __new__(self, in1, in2):
        self.in1 = in1
        self.in2 = in2

        self.in2_temp = write_to_file(in2)
        self.in2 = TableIterator(self.in2_temp)
        self.in1_row = self.in1.next()

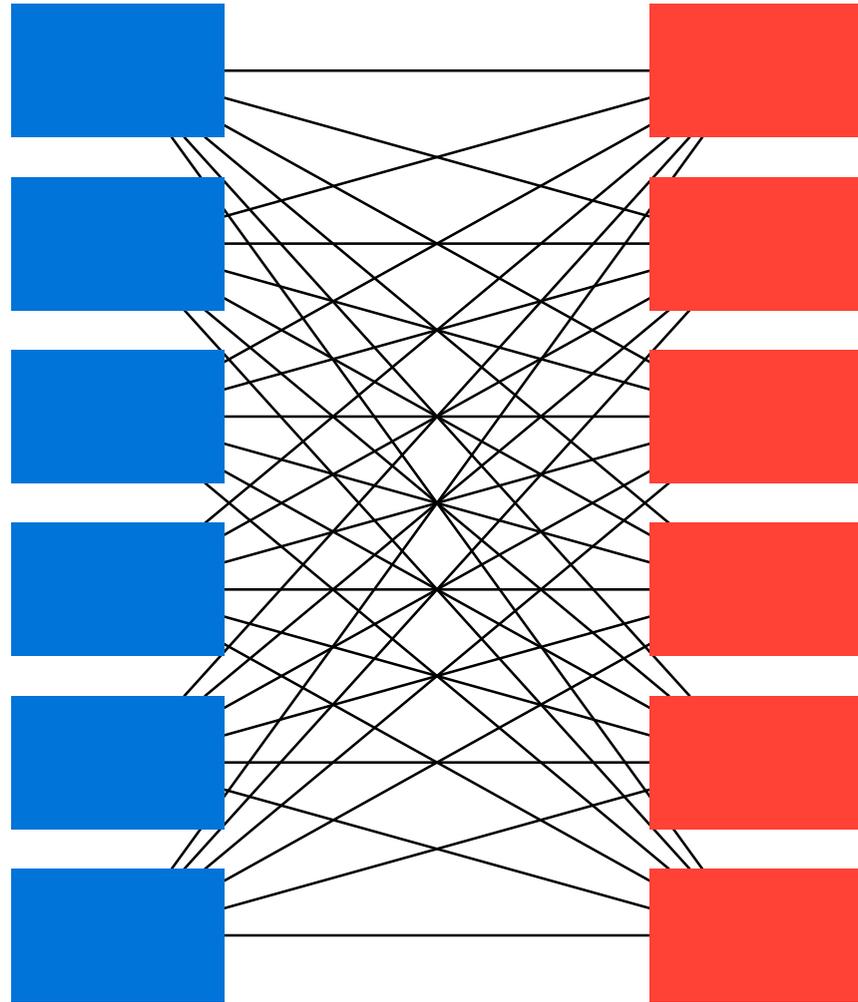
    def next(self) -> Option[Record]
        if (in2_row := self.in2.next()) is None:
            self.in2 = TableIterator(self.in2_temp)
            self.in1_row = self.in1.next()
        row = self.in1_row + self.in2.next()
        return row
```

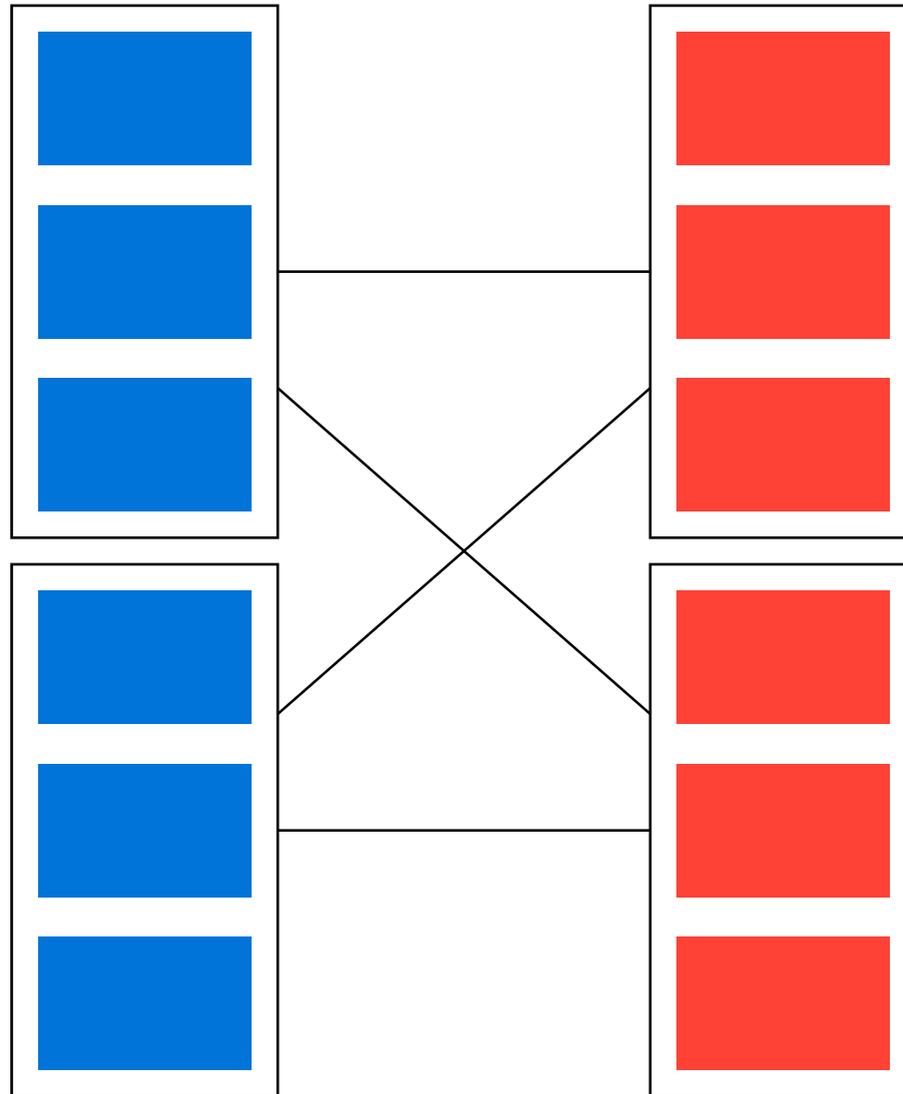
```
def product(In1, In2):  
    in2_temp = [s for s in In2]  
    for r in In1:  
        for s in in2_temp:  
            yield r + s
```

or

```
def product(In1, In2):  
    in2_temp = write_to_file(in2)  
    for r in In1:  
        for s in TableIterator(in2_temp):  
            yield r + s
```

What are the IO & Memory Complexities



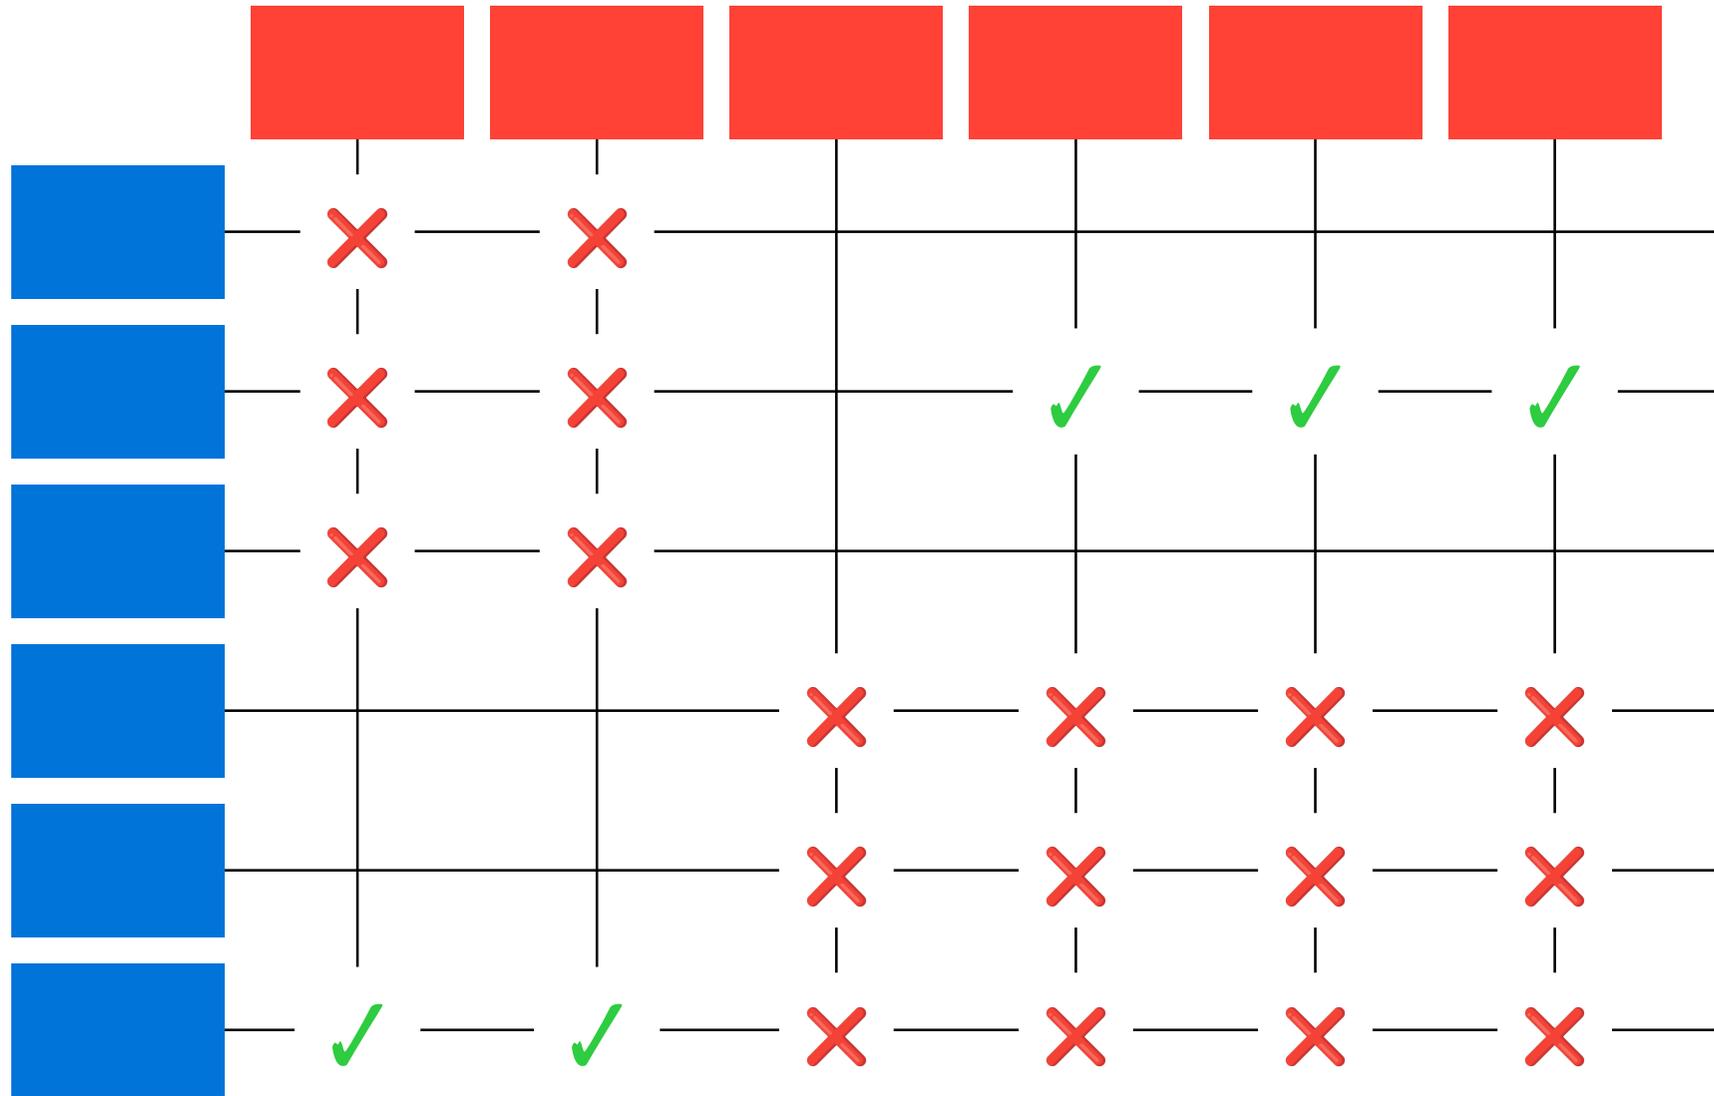


```
def product(In1, In2):  
    in2_temp = write_to_file(In2)  
    for r_batch in In1.by_batch:  
        for s in TableIterator(in2_temp):  
            for r in r_batch:  
                yield r + s
```

```
def product(In1, In2):  
    in2_temp = write_to_file(In2)  
    for r_batch in In1.by_batch:  
        for s in TableIterator(in2_temp):  
            for r in r_batch:  
                yield r + s
```

What are the IO & Memory Complexities

Joins



$$h(A) = 1$$

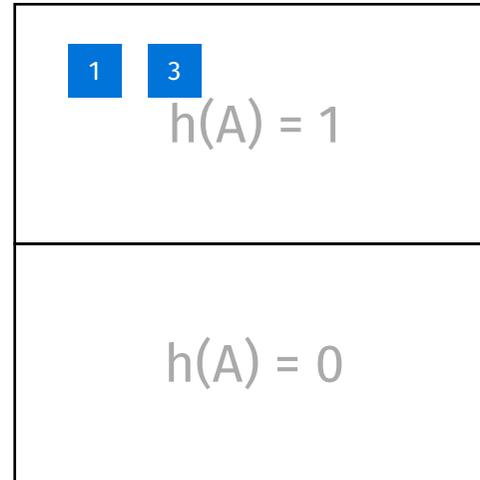
$$h(A) = 0$$

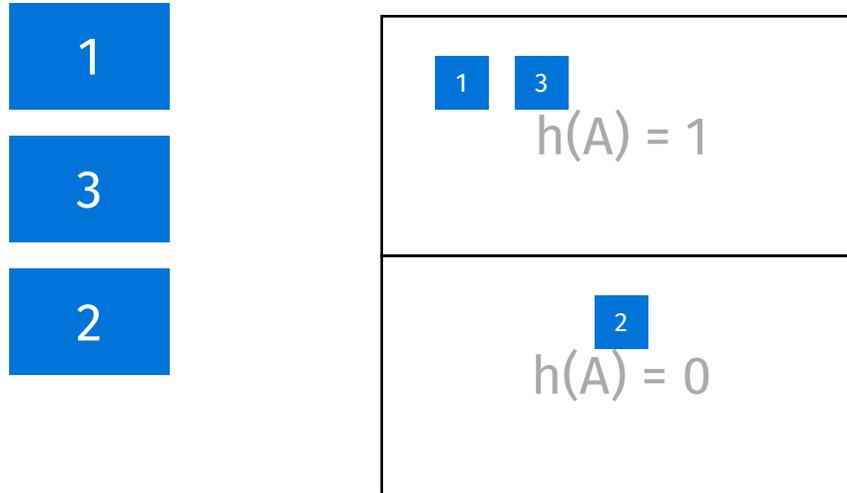
1

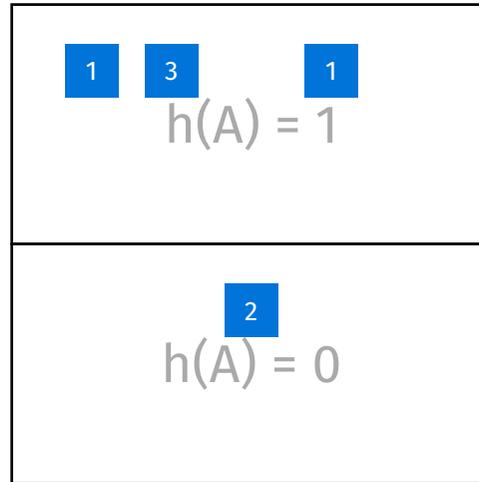
1

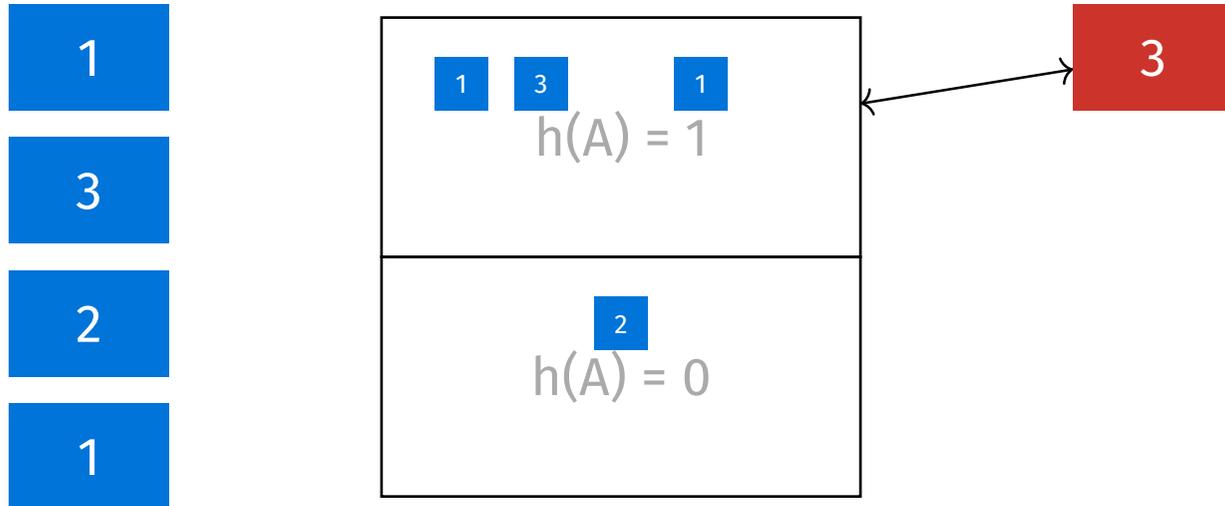
$$h(A) = 1$$

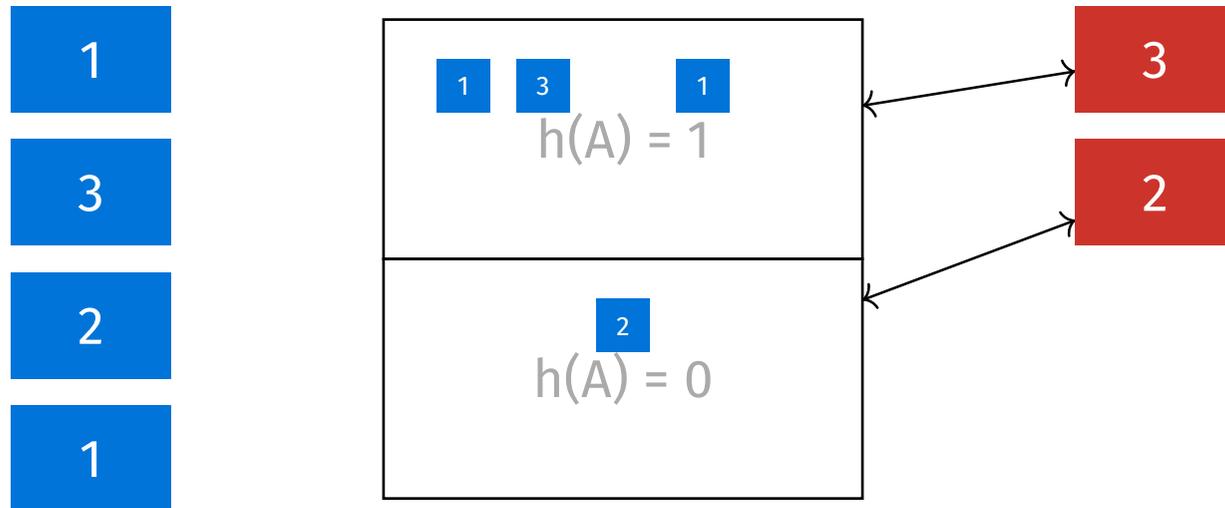
$$h(A) = 0$$

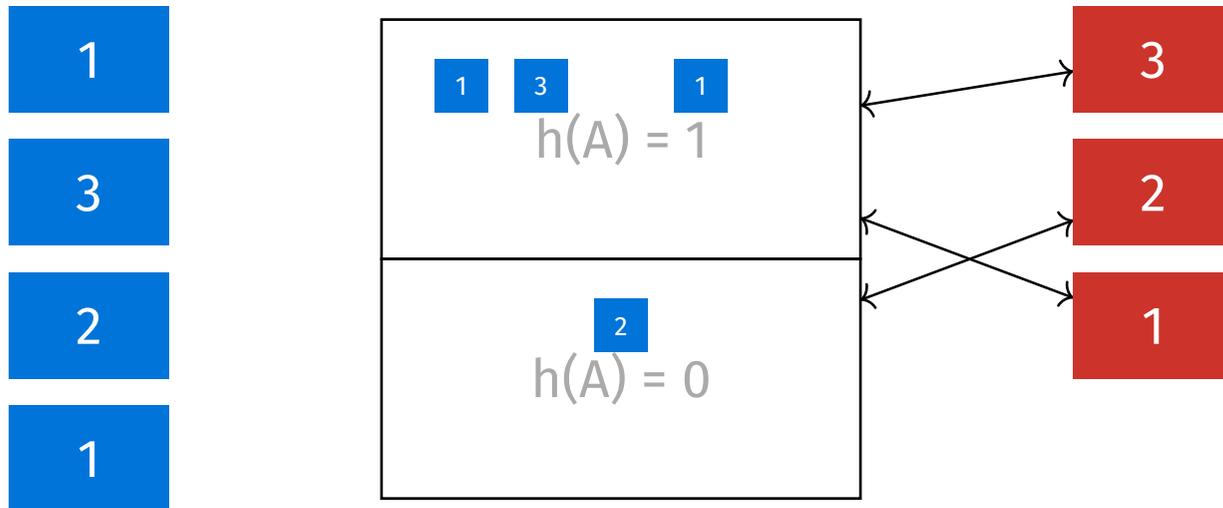


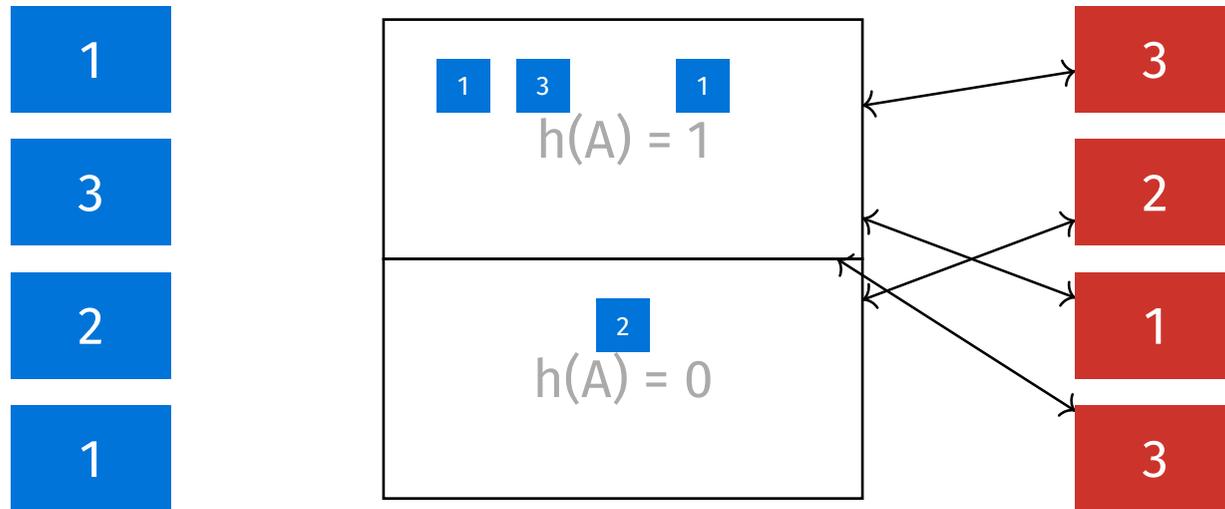


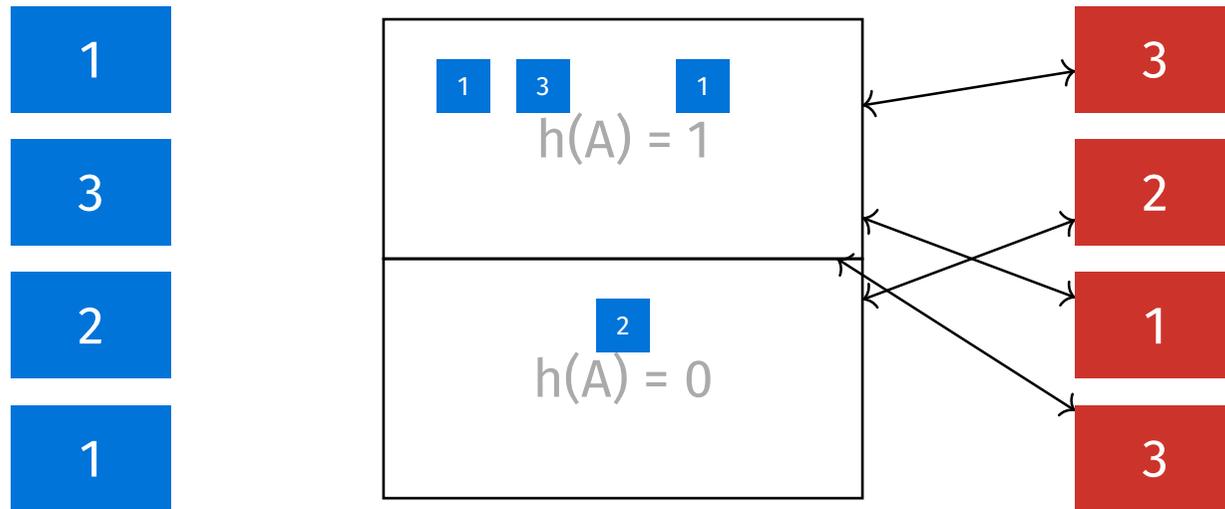




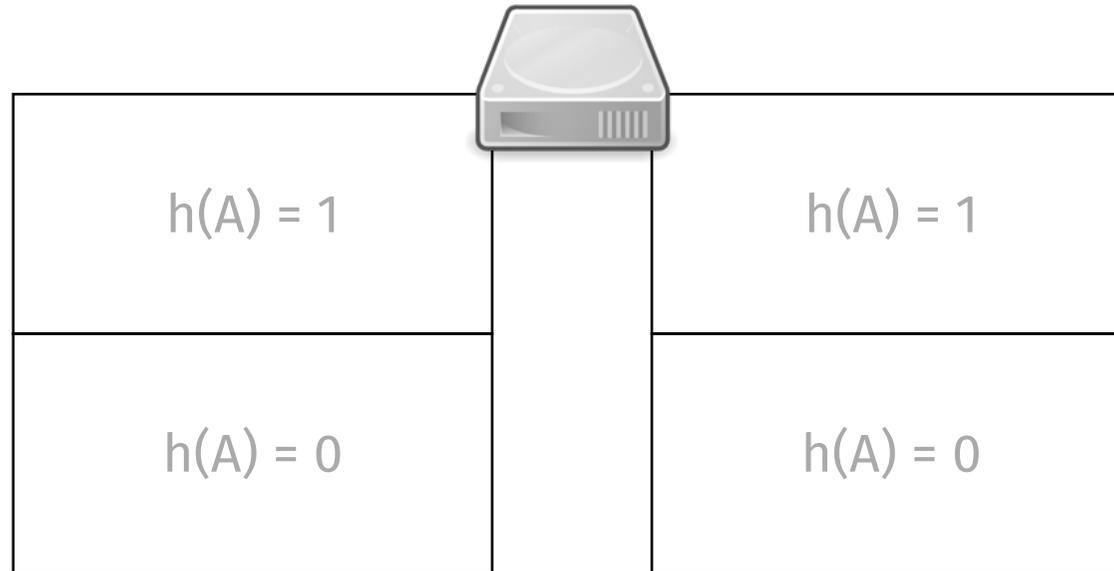




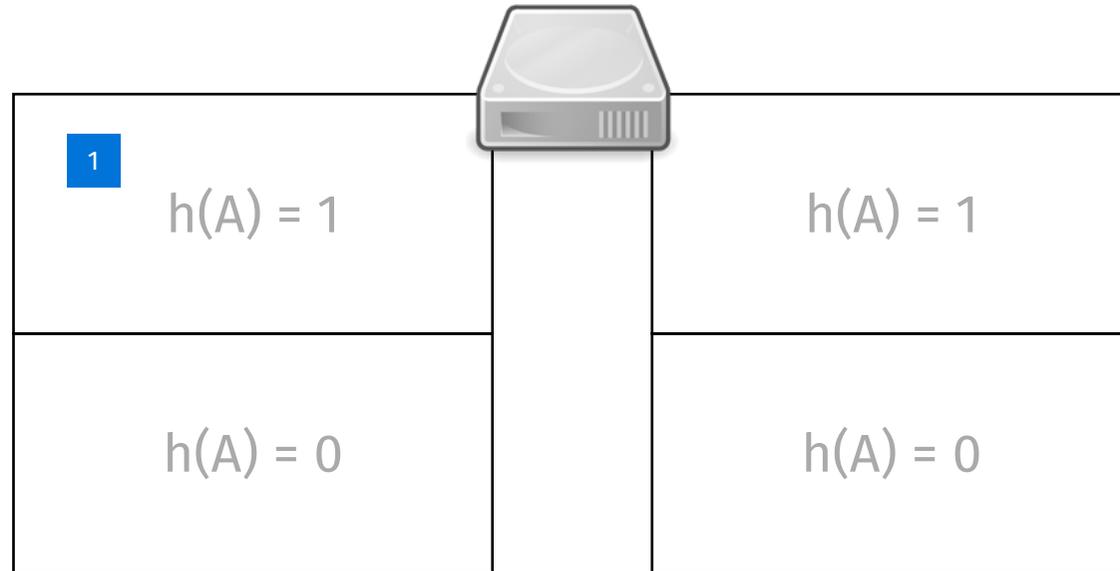


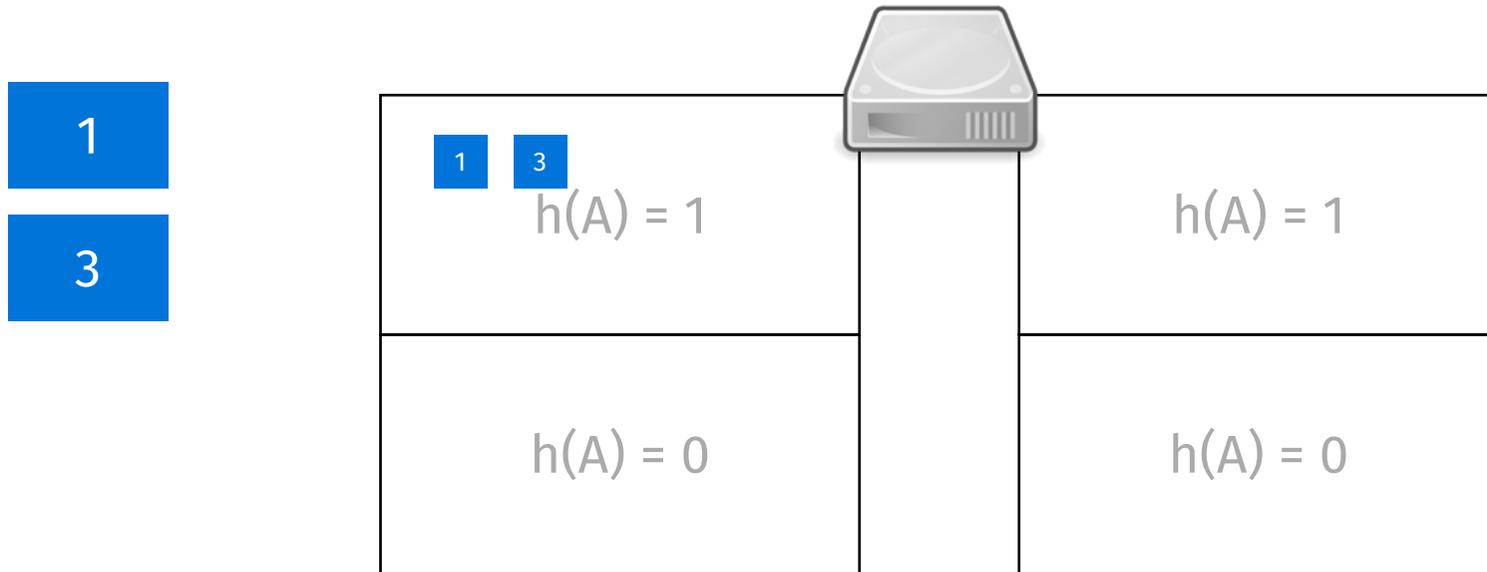


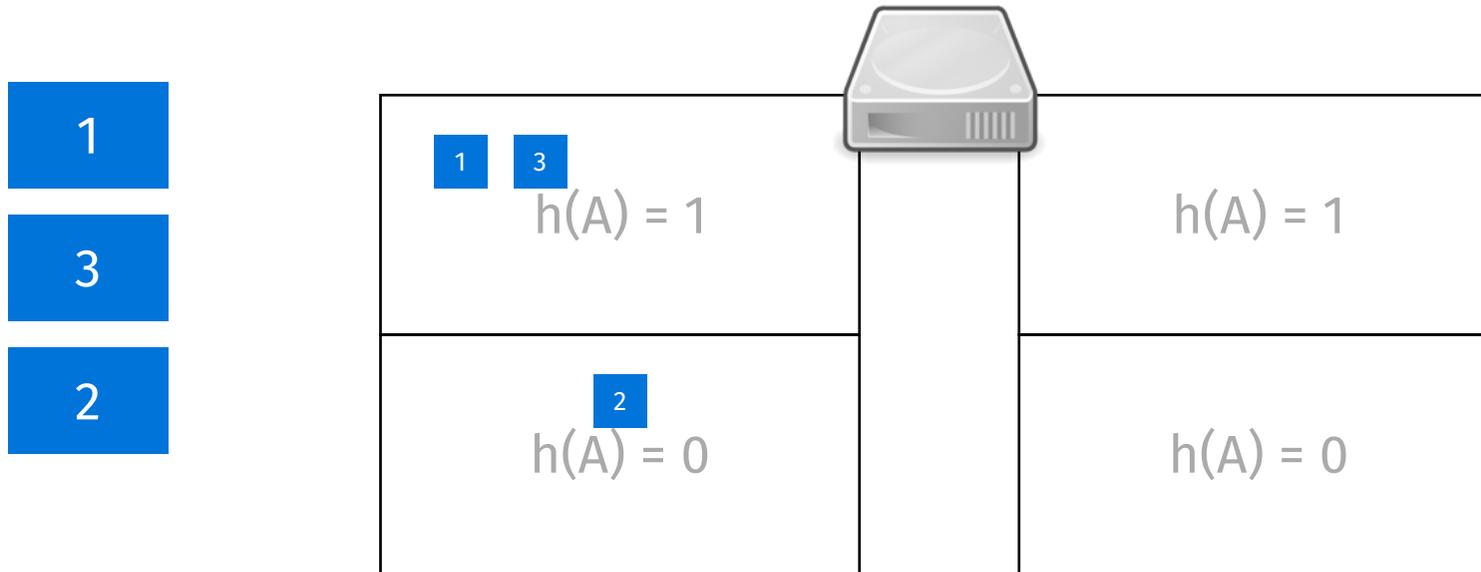
What are the IO & Memory Complexities?

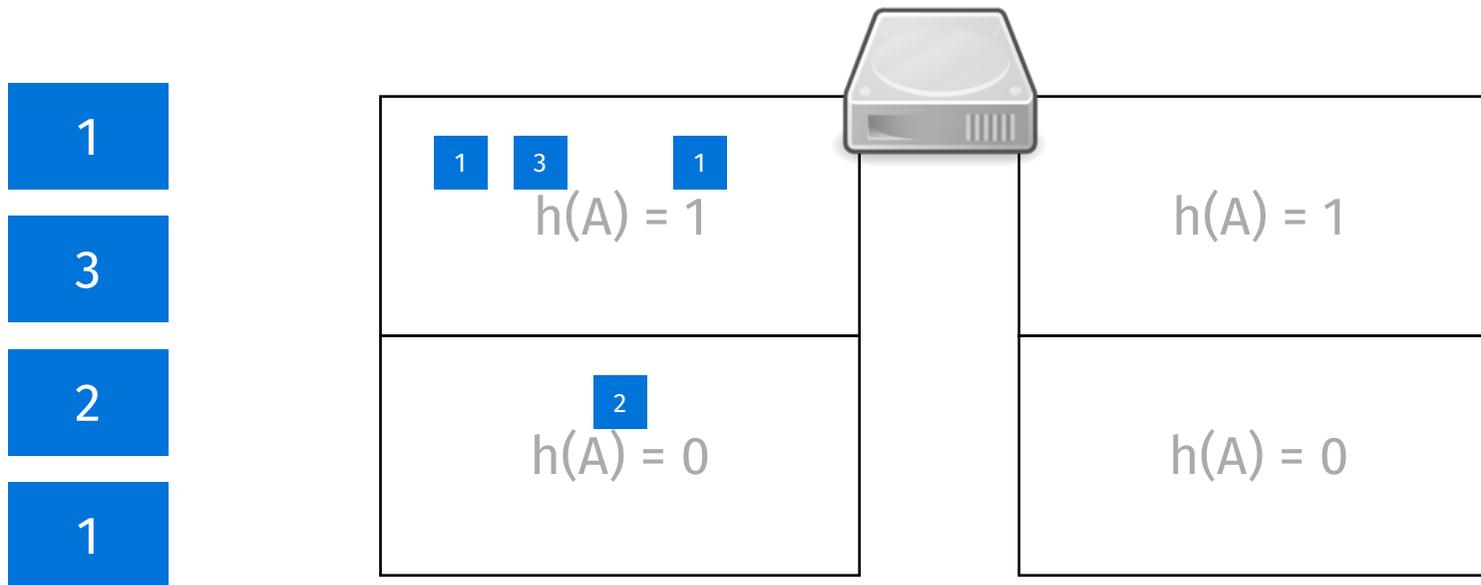


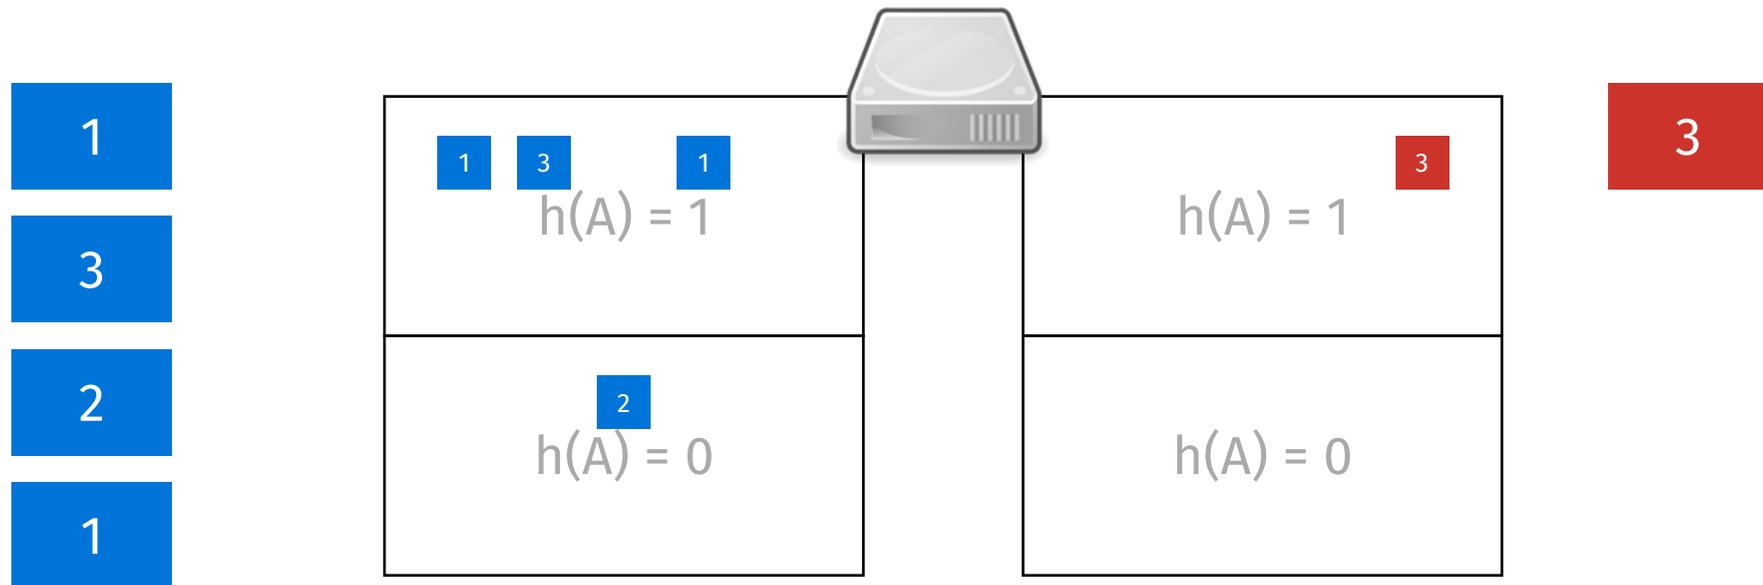
1

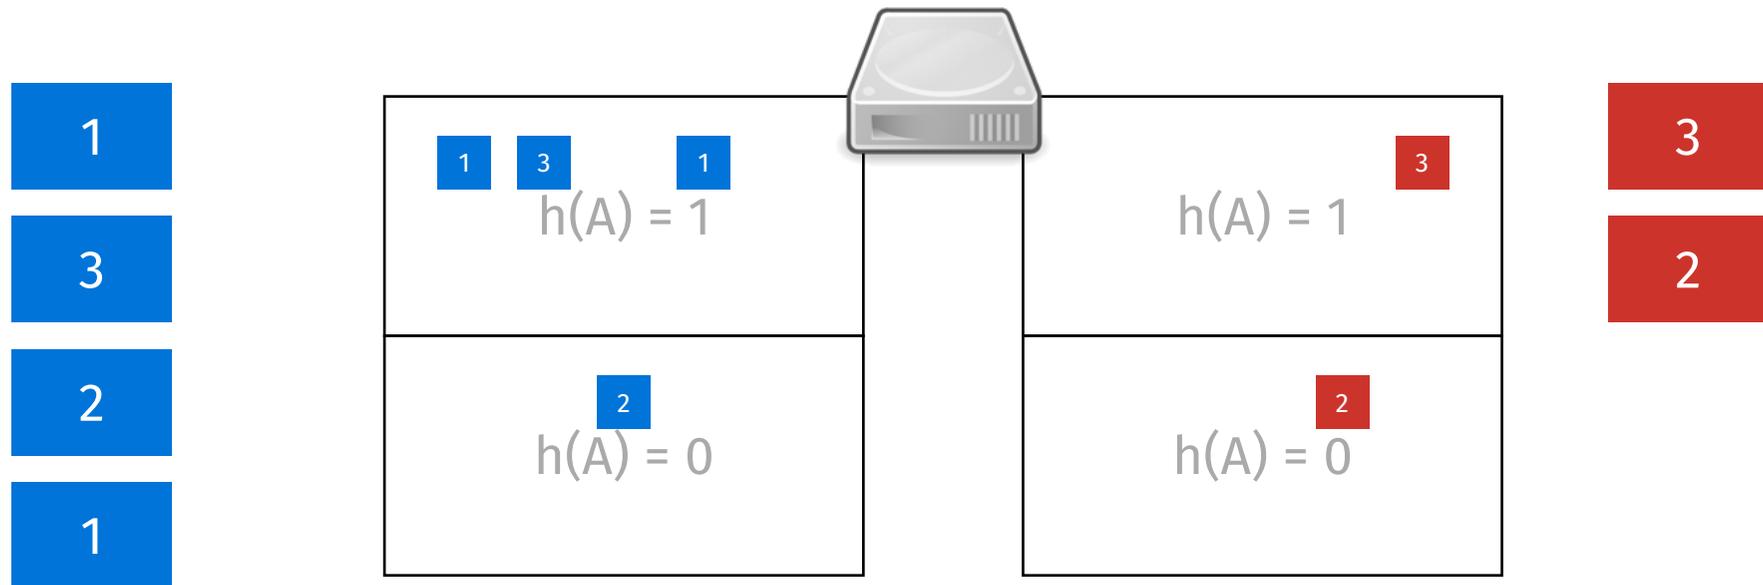


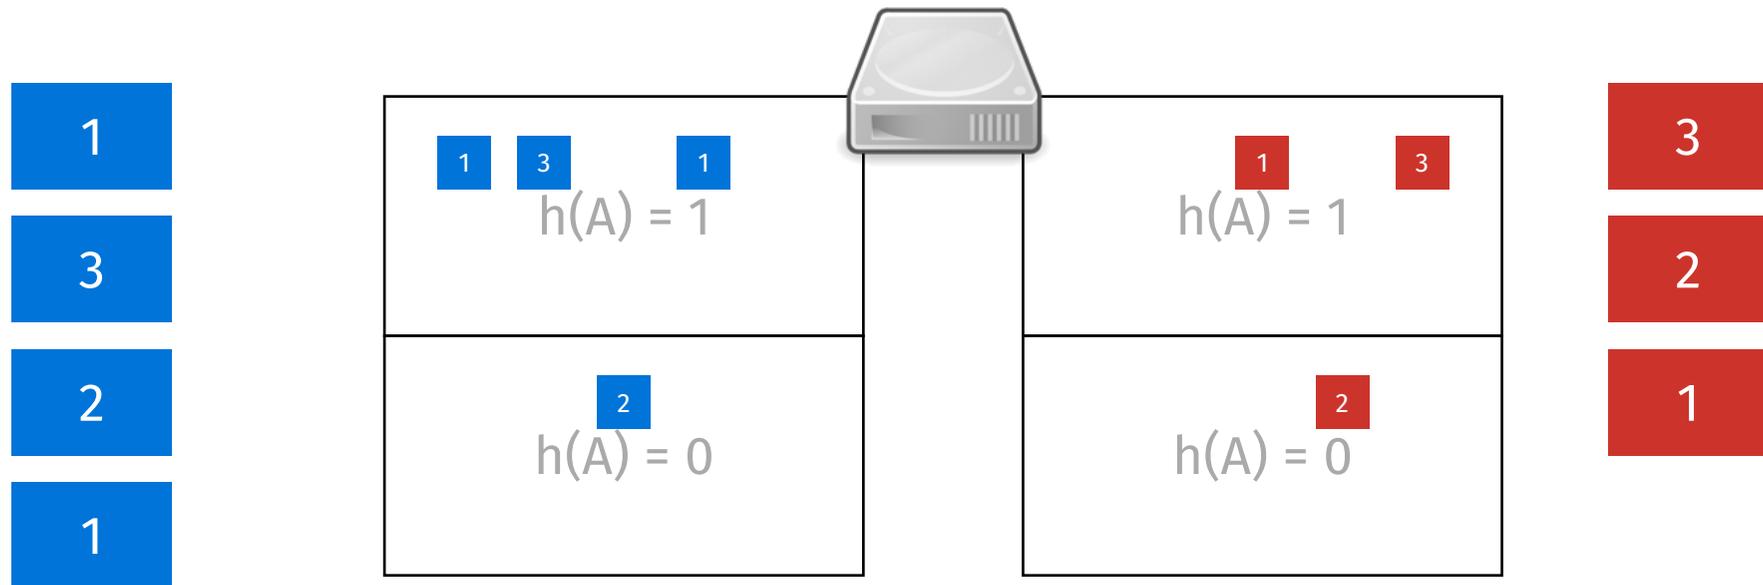


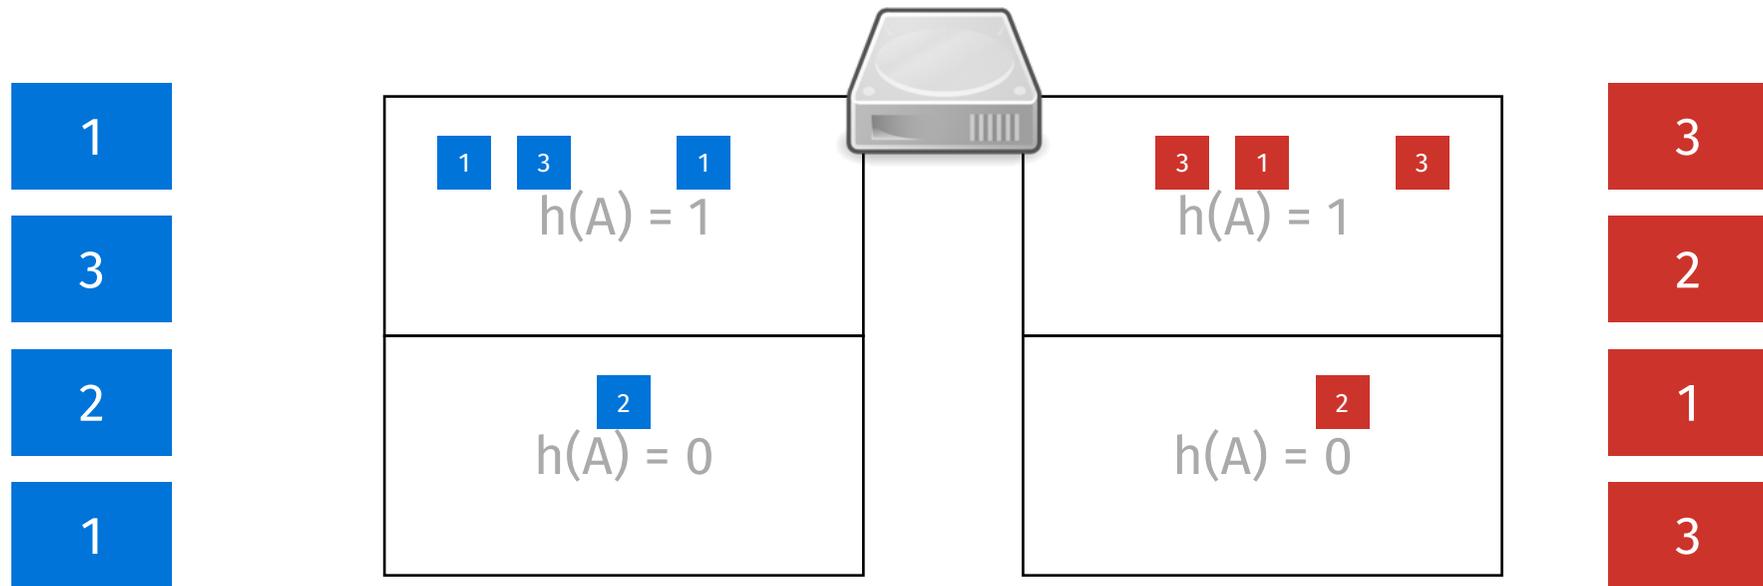


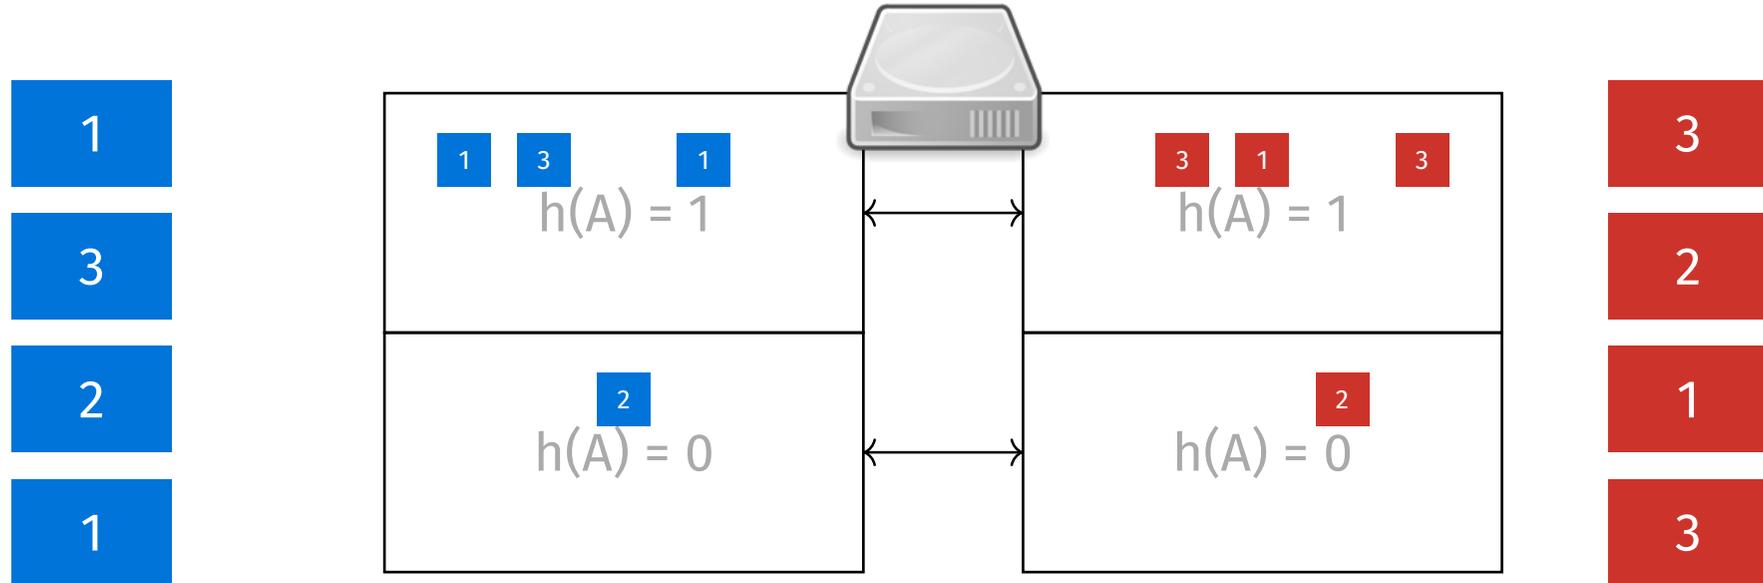










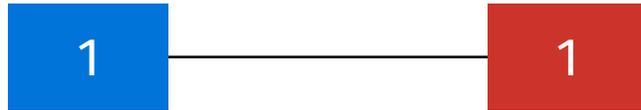


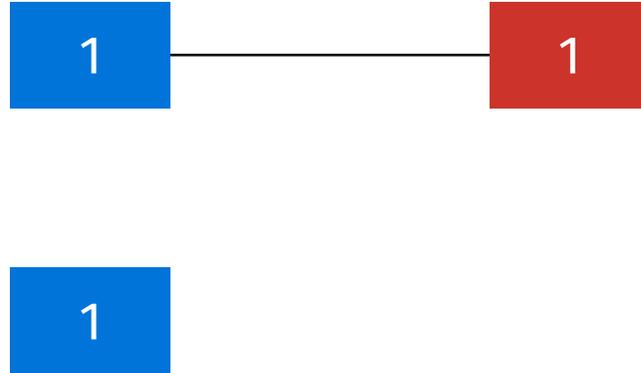
What are the IO & Memory Complexities?

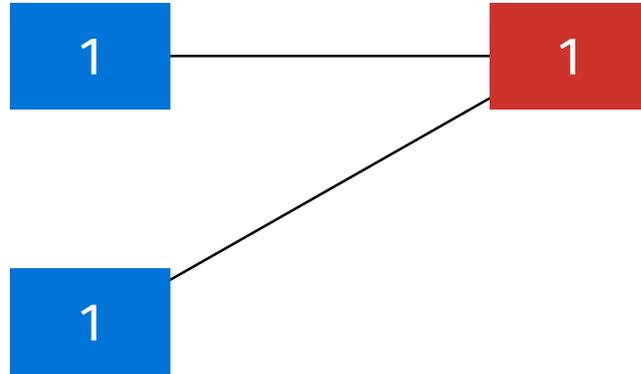
1

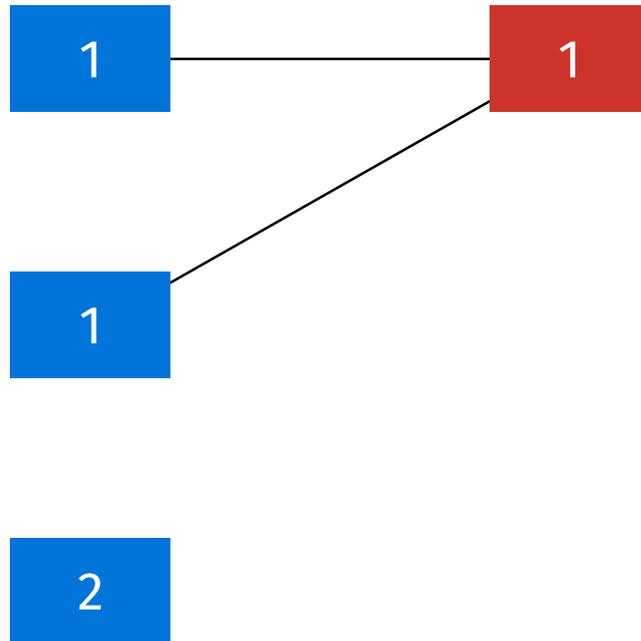
1

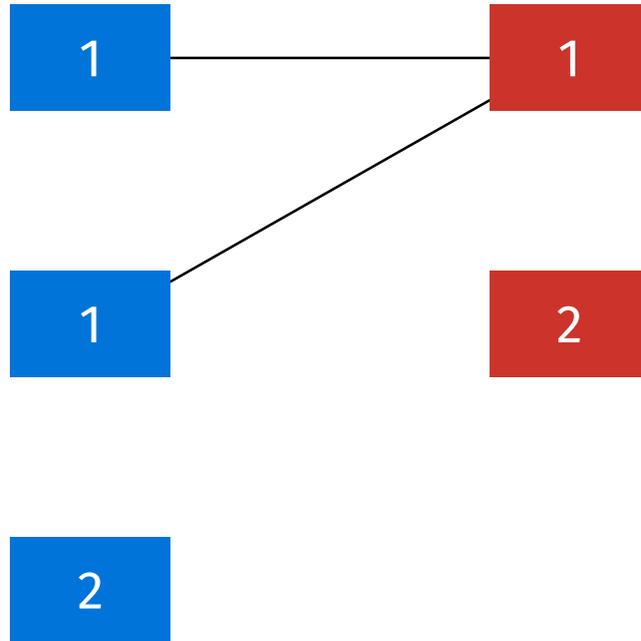
1

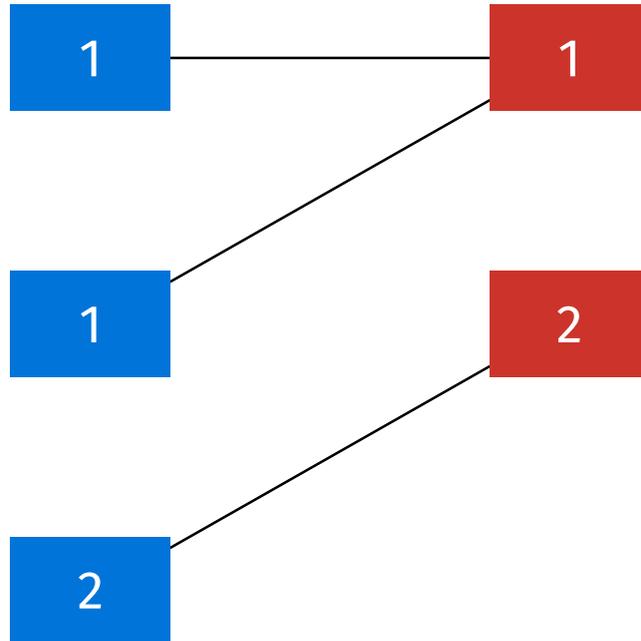


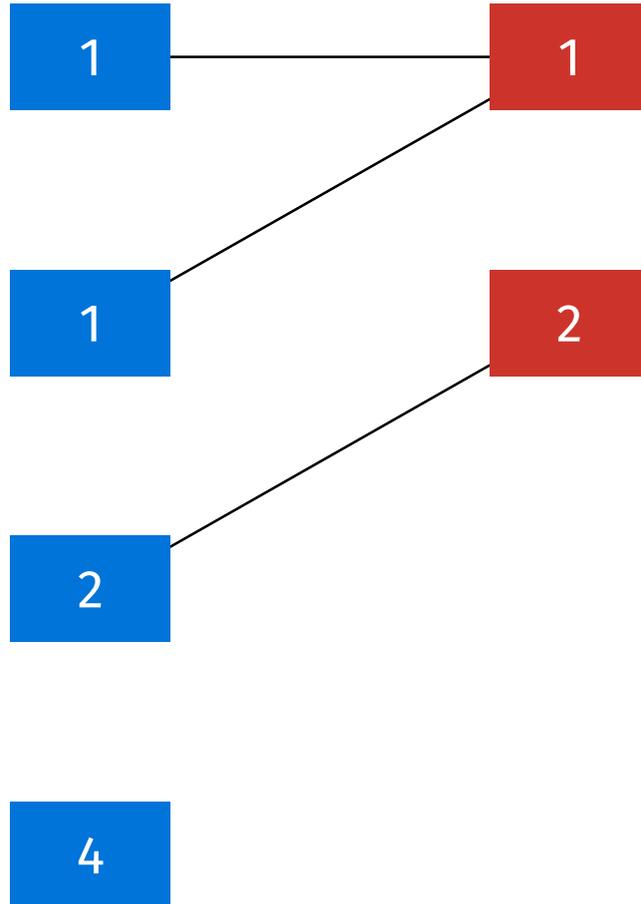


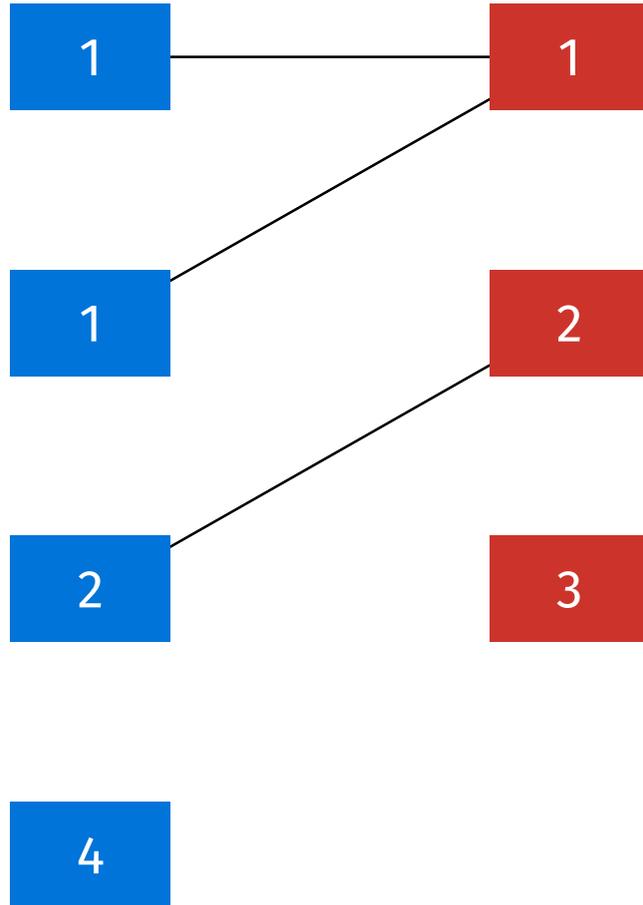


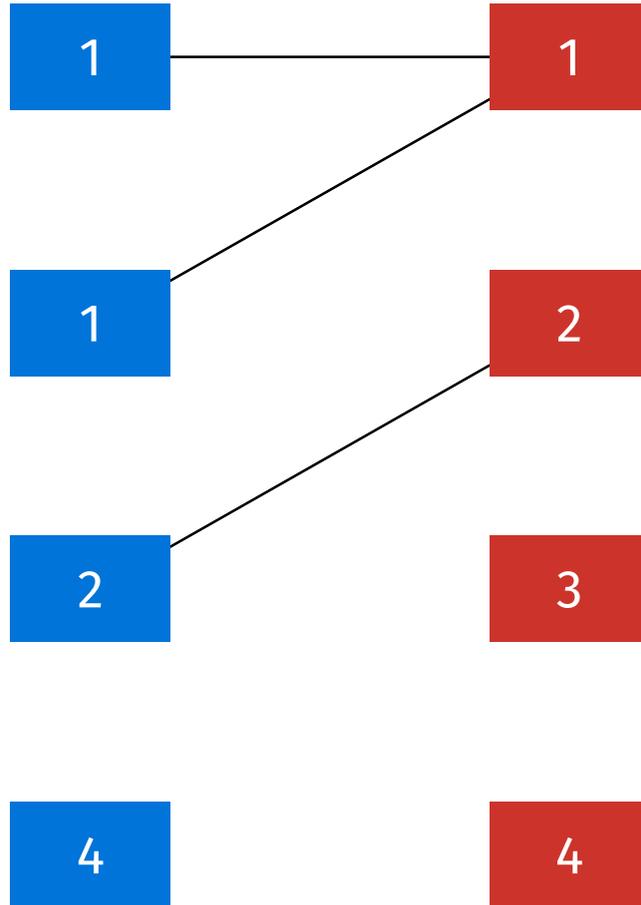


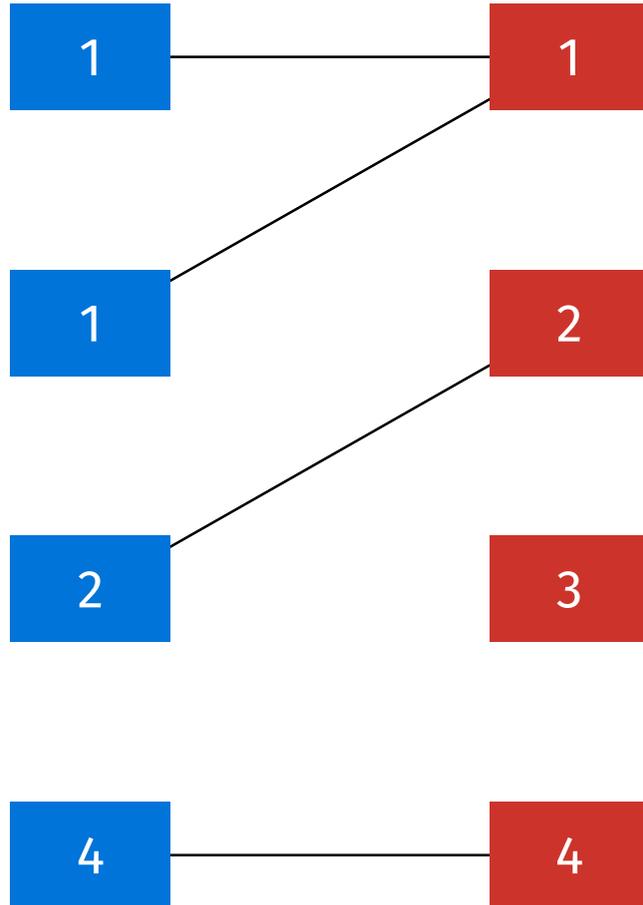












```
def sort_merge_join(key, In1, In2):  
    In1 = In1.sorted();    In2 = In2.sorted()  
    row_1 = In1.next();    row_2 = In2.next()  
  
    while row_1 is not None and row_2 is not None:  
        if row_1[key] == row_2[key]:  
            yield row_1 + row_2  
        elif row_1[key] < row_2[key]:  
            row_1 = In1.next()  
        else:  
            row_2 = In2.next()
```

What are the Memory and IO complexities?

For $R \bowtie S$

- **Product** (aka Nested Loop): $|R| + |R| \cdot |S|$ extra IO, $O(1)$ mem, General
- **Block Nested Loop**: General, $|R| + \frac{|R| \cdot |S|}{|B|}$ extra IO, $O(|B|)$ mem, General
- **1-Pass Hash Join**: No extra IO, $O(|R|)$ mem Equijoin Only
- **2-Pass (Grace) Hash Join**: $2 \cdot (|R| + |S|)$ extra IO, $O(1)$ mem, Equijoin Only
- **Sort/Merge Join**: No extra IO, $O(1)$ mem, Equijoin Only, Data must be sorted

Aggregation

- `Init()`: The “default” value of the aggregate
- `Accumulate(agg, value)`: Incorporate a new value
- `Finalize(agg) -> value`: Post-process the aggregate

Aggregate `Init()` `Accumulate(a, v)` `Finalize(a)`

Sum

Aggregate	Init()	Accumulate(a, v)	Finalize(a)
Sum	0	$a + v$	identity

Aggregate	Init()	Accumulate(a, v)	Finalize(a)
Sum	0	$a + v$	identity
Count			

Aggregate	<code>Init()</code>	<code>Accumulate(a, v)</code>	<code>Finalize(a)</code>
Sum	0	$a + v$	identity
Count	0	$a + 1$	identity

Aggregate	Init()	Accumulate(a, v)	Finalize(a)
Sum	0	$a + v$	identity
Count	0	$a + 1$	identity
Max			

Aggregate	Init()	Accumulate(a, v)	Finalize(a)
Sum	0	$a + v$	identity
Count	0	$a + 1$	identity
Max	$-\infty$	$\max(a, v)$	identity

Aggregate	Init()	Accumulate(a, v)	Finalize(a)
Sum	0	$a + v$	identity
Count	0	$a + 1$	identity
Max	$-\infty$	$\max(a, v)$	identity
Average			

Aggregate	<code>Init()</code>	<code>Accumulate(a, v)</code>	<code>Finalize(a)</code>
Sum	0	$a + v$	identity
Count	0	$a + 1$	identity
Max	$-\infty$	$\max(a, v)$	identity
Average	{s: 0, c: v}	{s: a.s+v, c: a.c+1}	a.s/a.c

```
def aggregate(group_by, In):  
    groups = defaultdict(Init)  
  
    for row in In:  
        groups[row[group_by]] = Accumulate(groups[row[group_by]],  
                                           row)  
  
    for (key, value) in groups:  
        yield (key, Finalize(value))
```

What are the Memory and IO complexities?

**What if there's not
enough memory?**

$$h(A) = 1$$

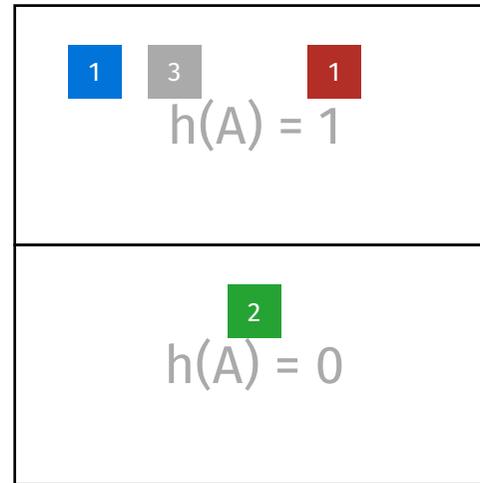
$$h(A) = 0$$

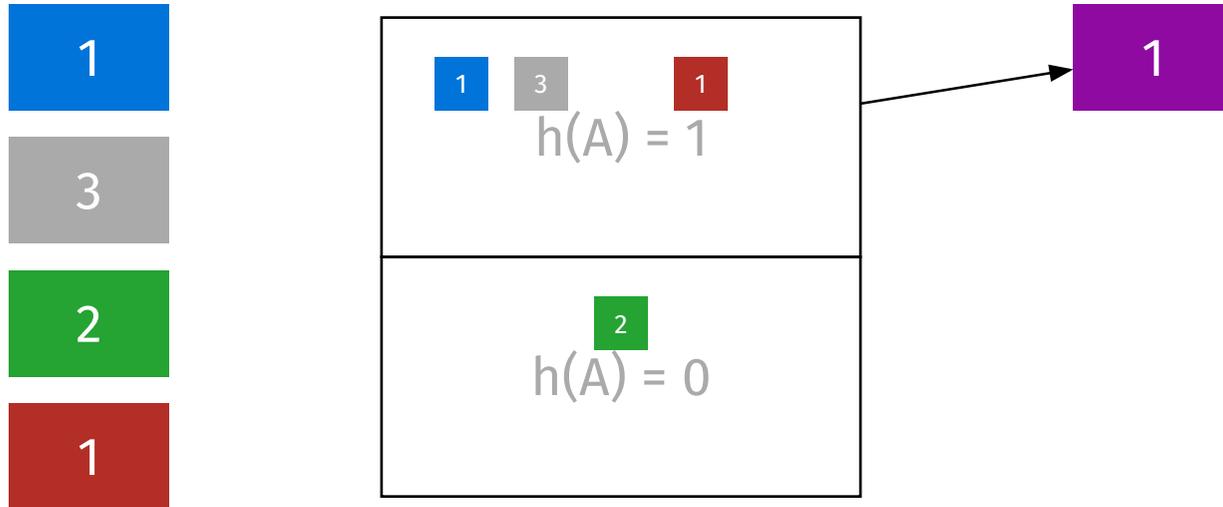
1

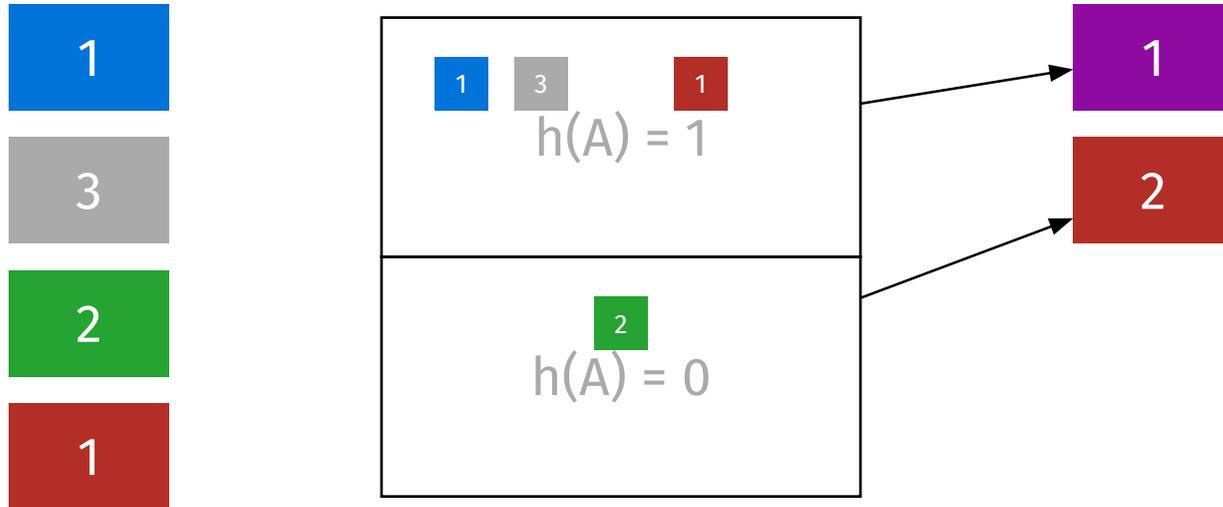
1

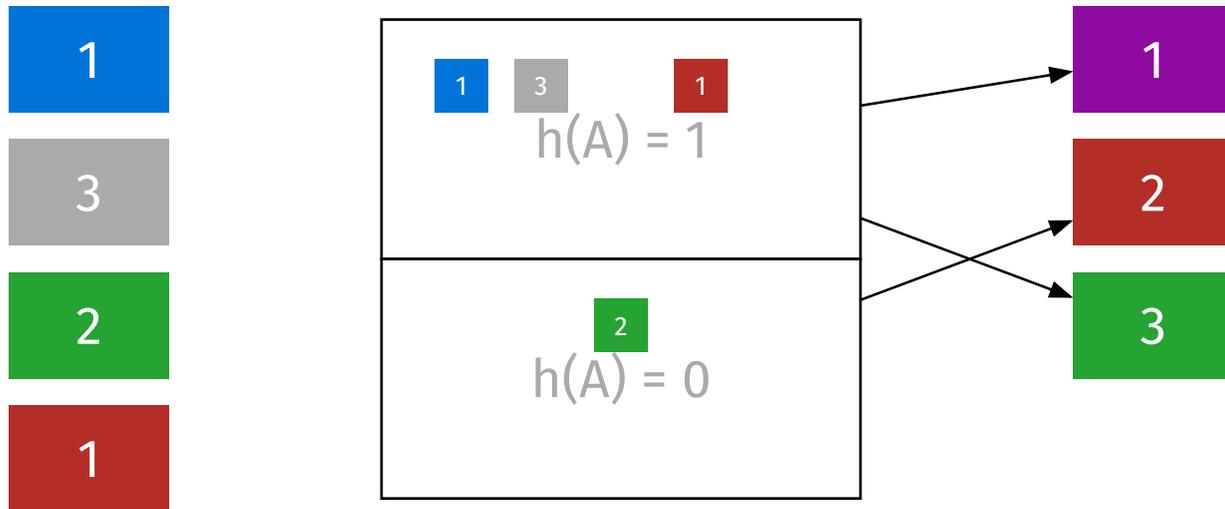
$$h(A) = 1$$

$$h(A) = 0$$









```
def partitioned_aggregate(group_by, In):  
    partitions = [ TempFile() for i in range(PARTITIONS) ]  
  
    for row in In:  
        partitions[ row[group_by].hash() % PARTITIONS ].write(row)  
  
    for partition_file in partitions:  
        groups = defaultdict(Init)  
        for row in partition_file:  
            groups[row[group_by]] = Accumulate(groups[row[group_by]], row)  
        for (key, value) in groups:  
            yield (key, Finalize(value))
```

What are the Memory and IO complexities?

```
def partitioned_aggregate(group_by, In):  
    In = In.sorted()  
    group = None;      agg = None  
    while row := In.next() is not None:  
        if row[group_by] != group:  
            if group is not None:  
                yield group, agg  
            group = row[group_by];  agg = Init()  
        agg = Accumulate(agg, row)  
  
    if group is not None:  
        yield group, Finalize(agg)
```

What are the Memory and IO complexities?

Deliverables

- AI Quiz and Checkpoint 0 Past Due!
 - Reach out to me if you haven't finished it yet!
- Checkpoint 1 due Monday!