Storage & Indexing in Modern Databases

ECS 165A – Winter 2025

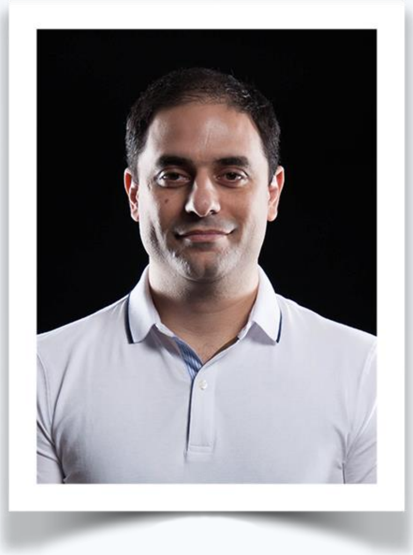

Mohammad Sadoghi
Exploratory Systems Lab
Department of Computer Science

UCDAVIS
UNIVERSITY OF CALIFORNIA

Apache ResilientDB Incubating

ExpoLab
Creativity Unfolded

How to quickly search for the desired information?

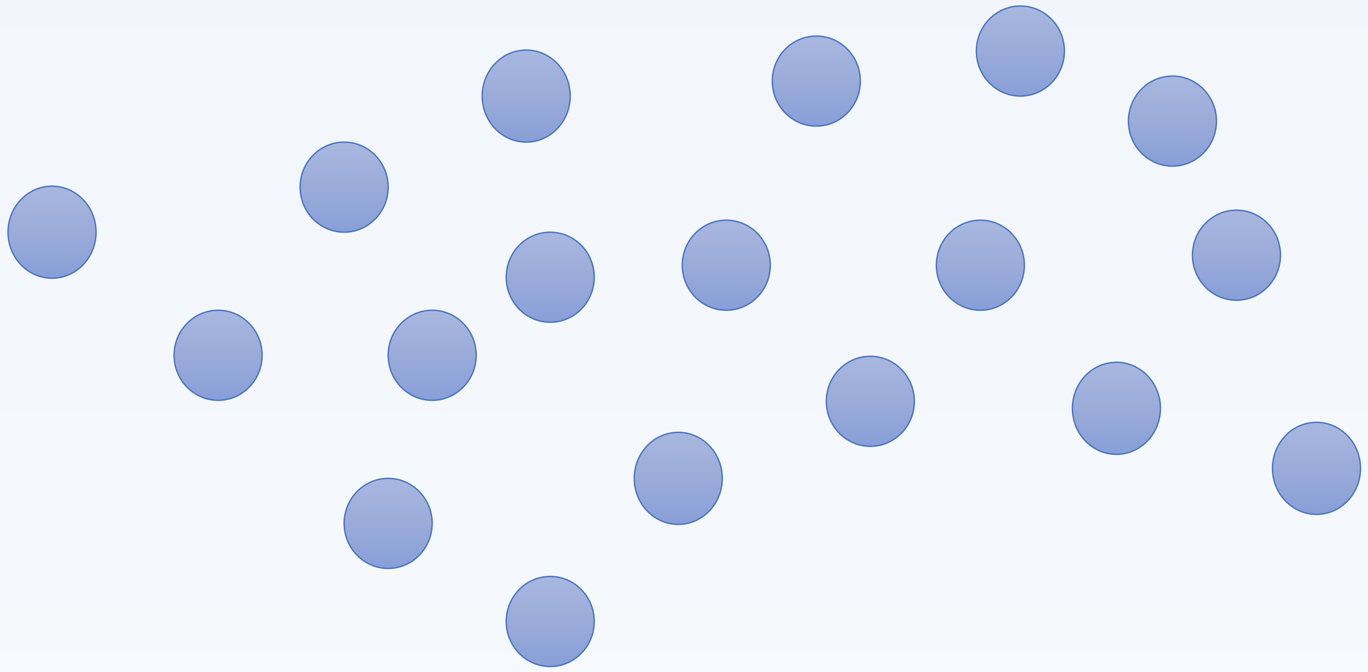

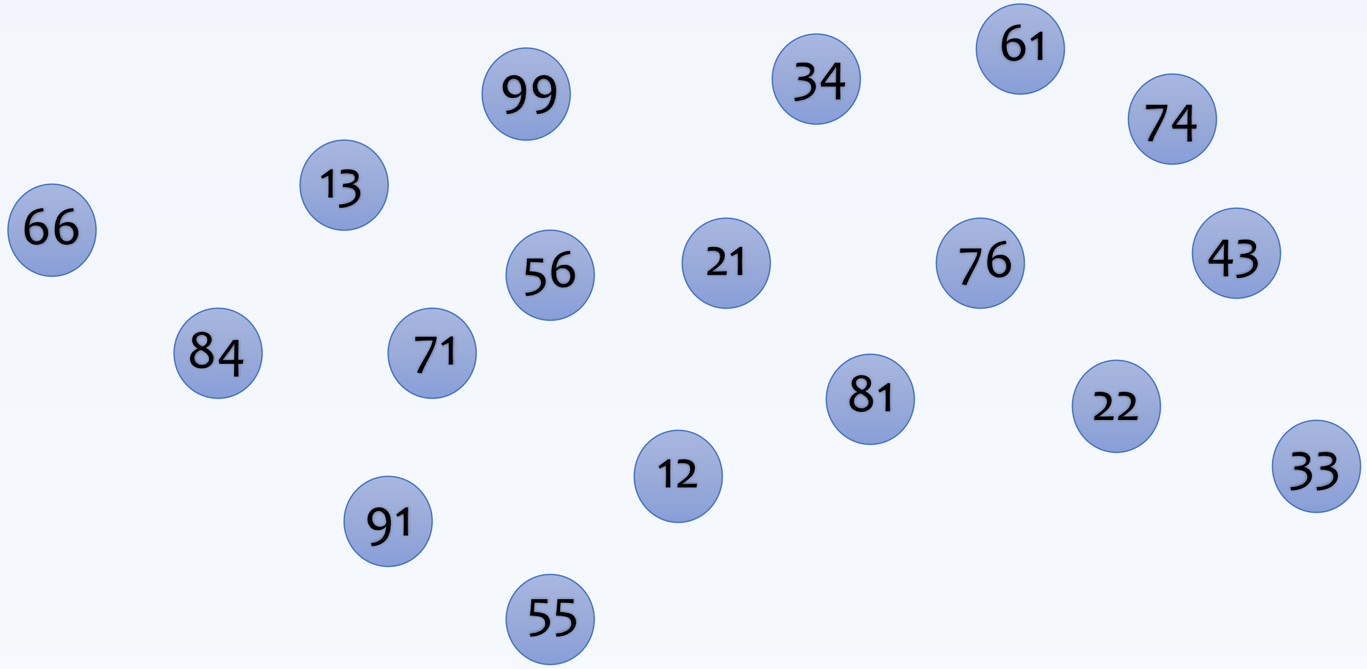
61
99
34
74
13
66
56
21
76
43
84
71
81
22
33
12
91
55

Searching for 76

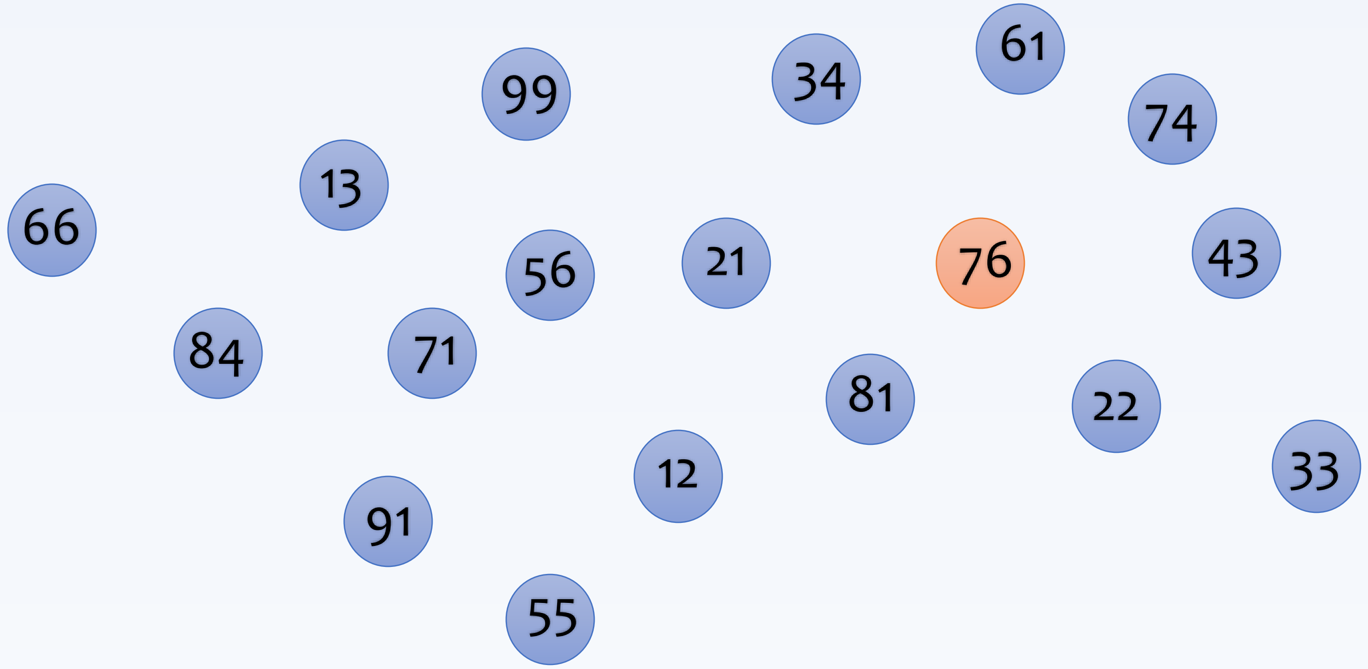

Searching for 76

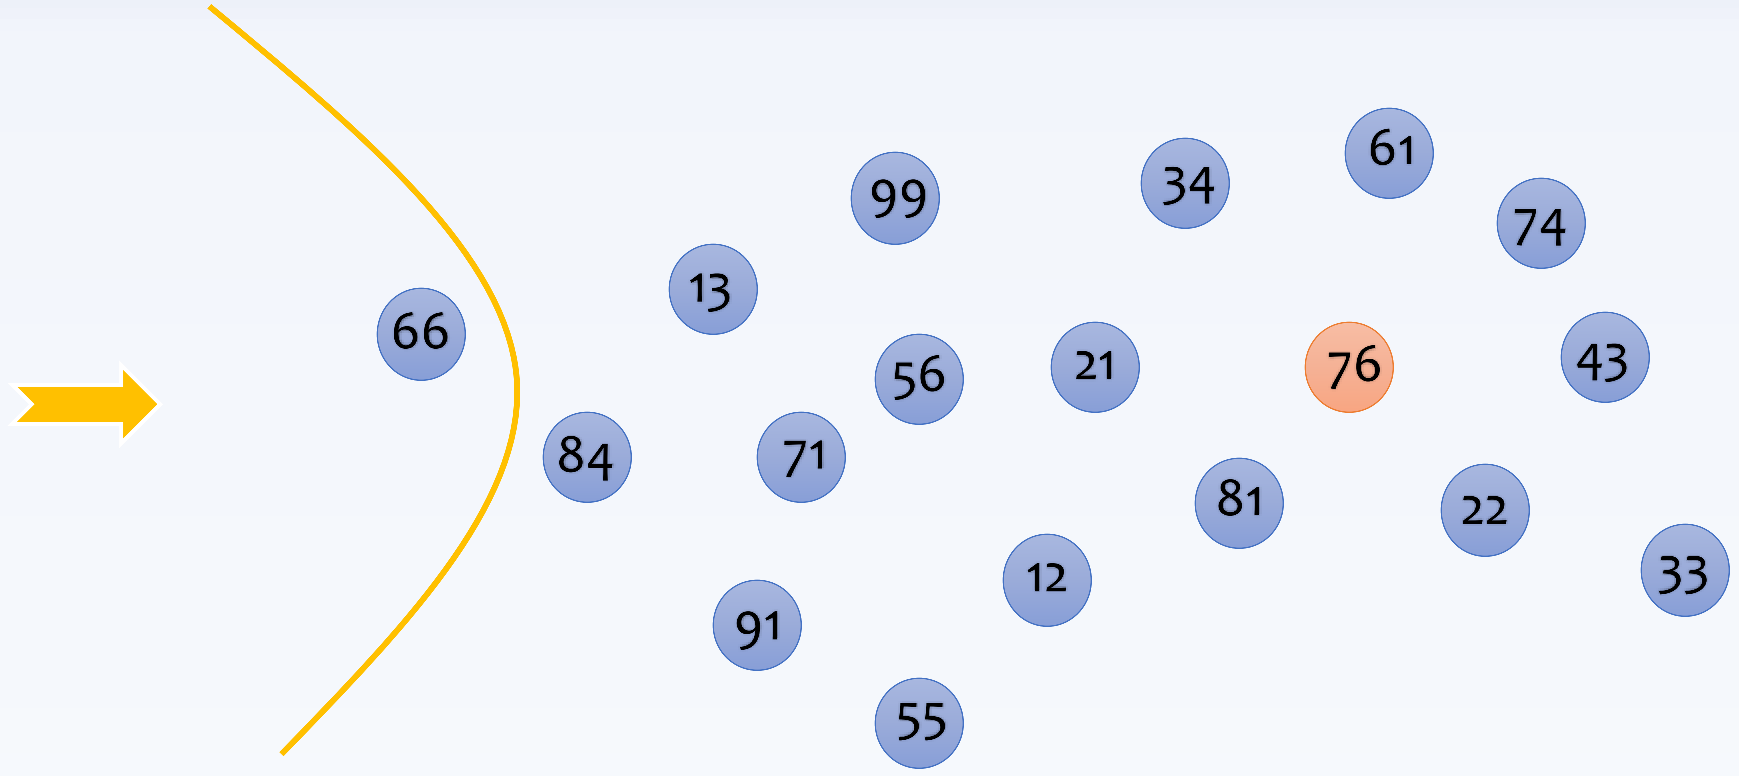

Searching for 76

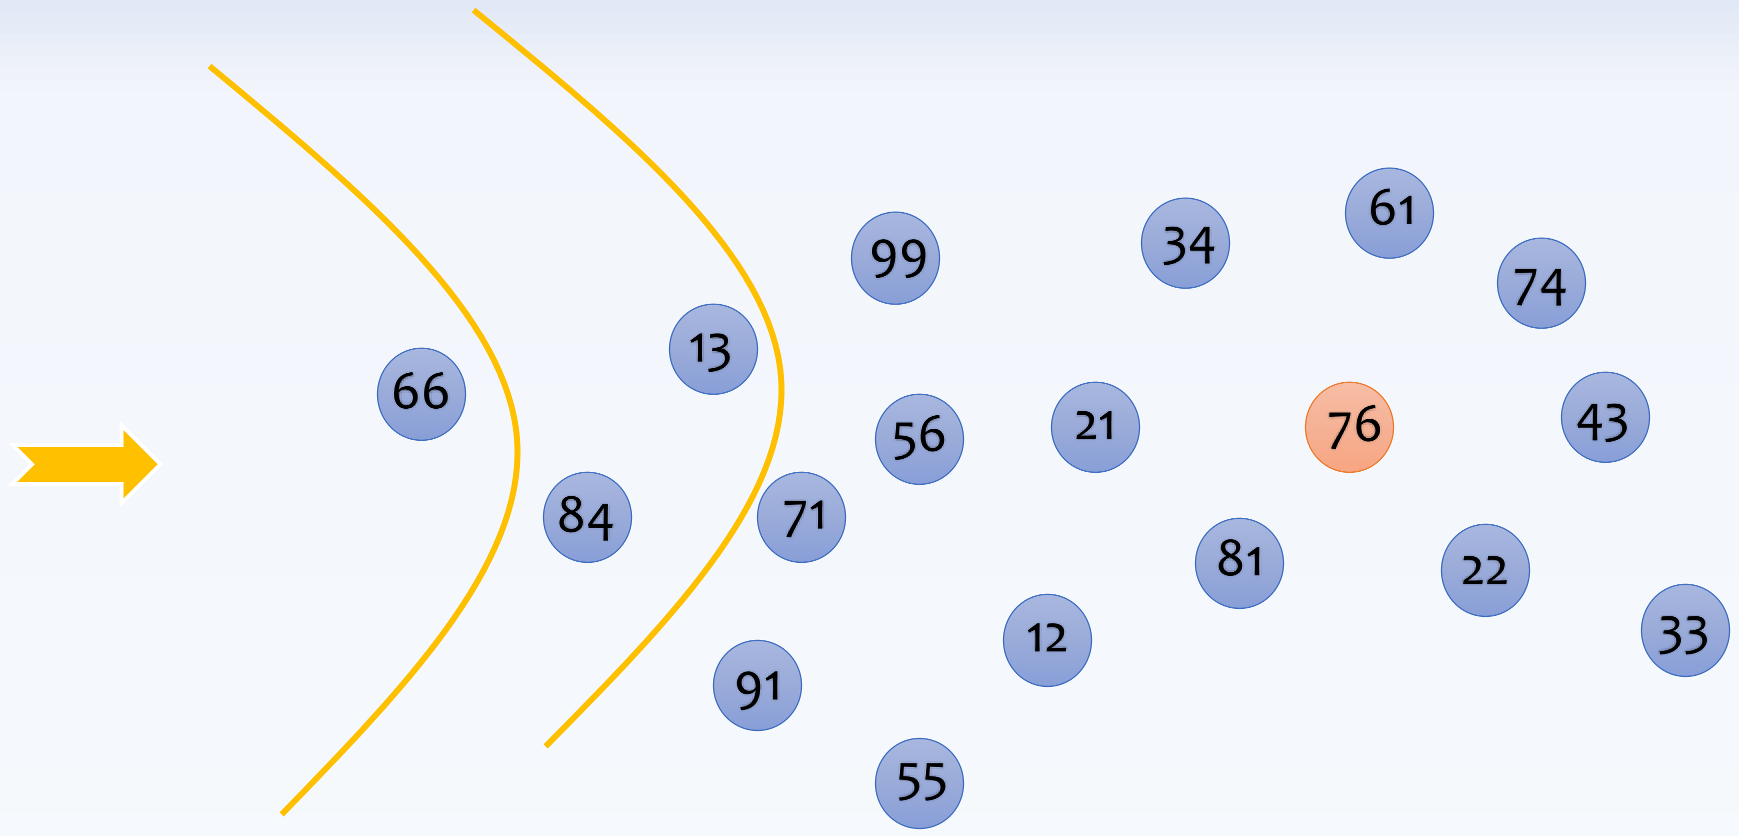

Searching for 76

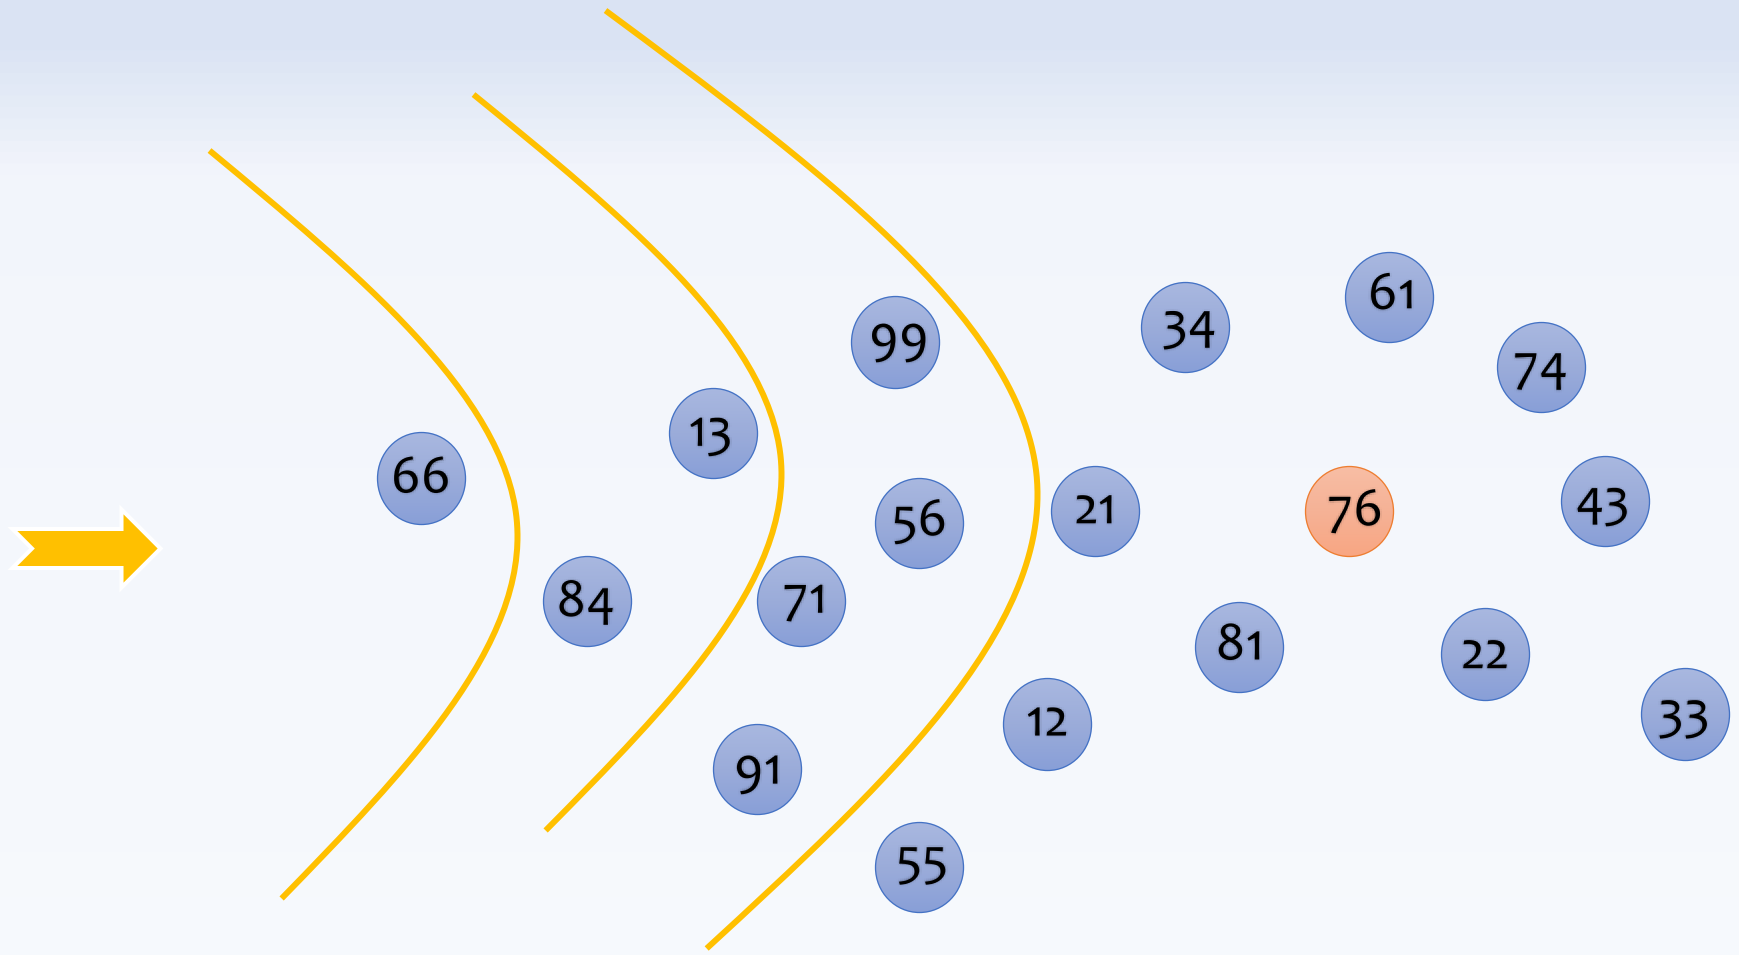

Searching for 76

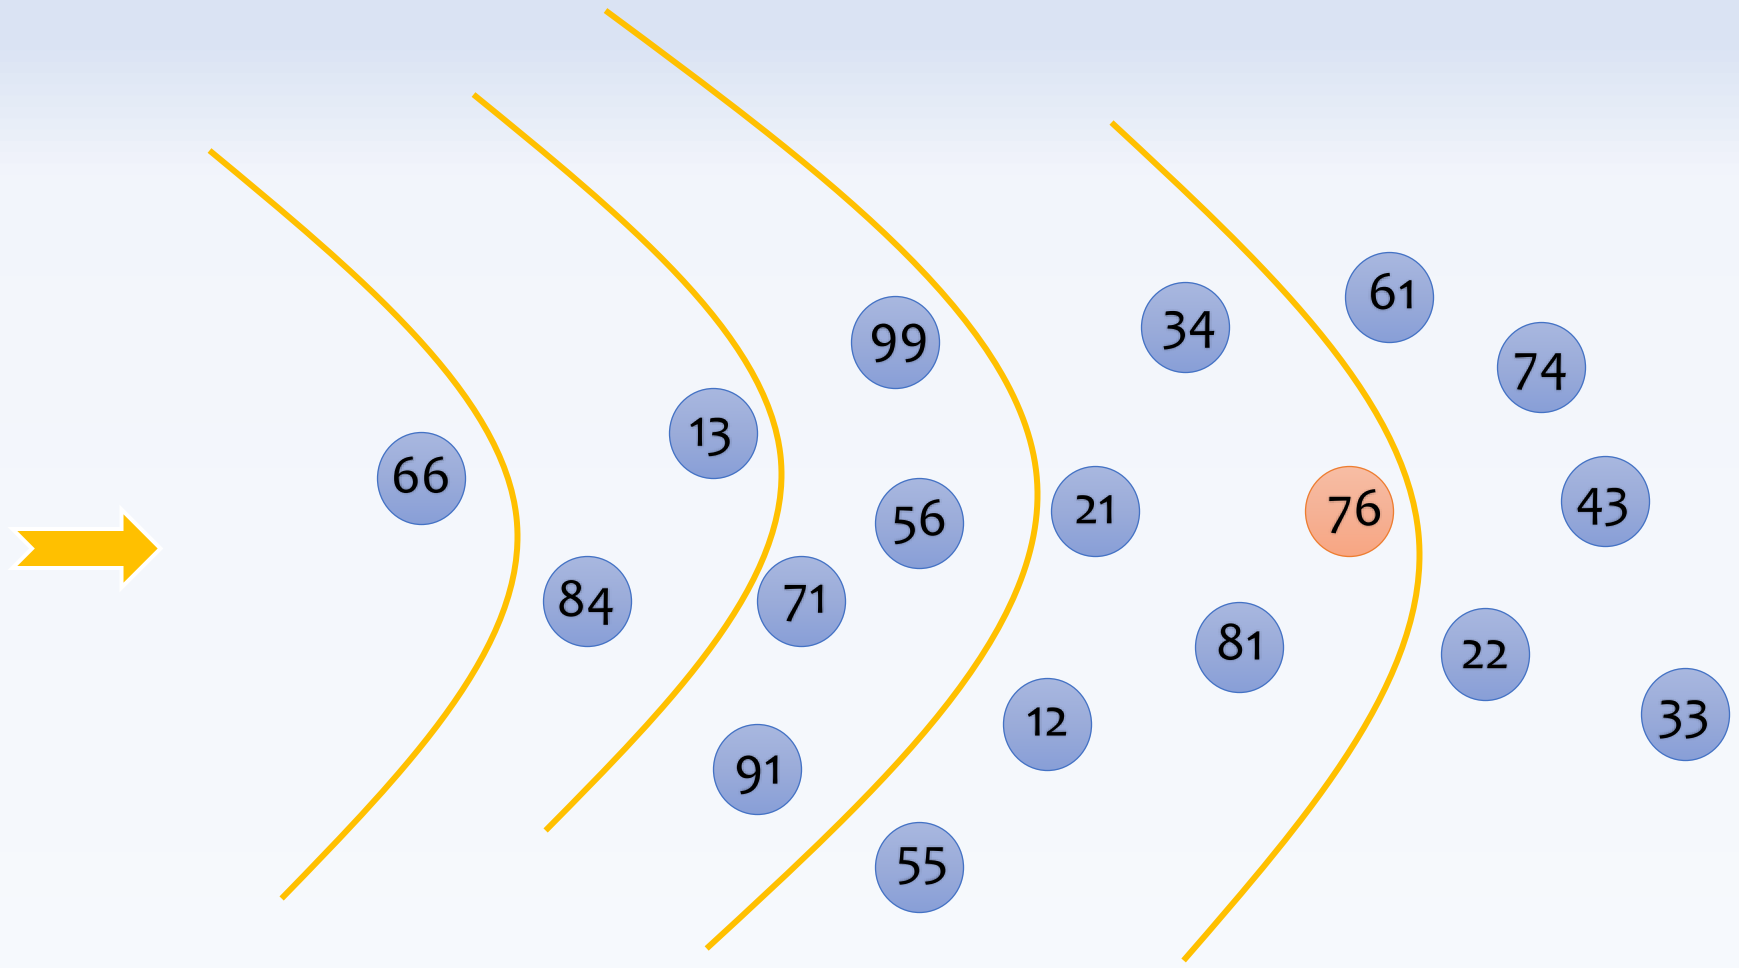

Searching for 76

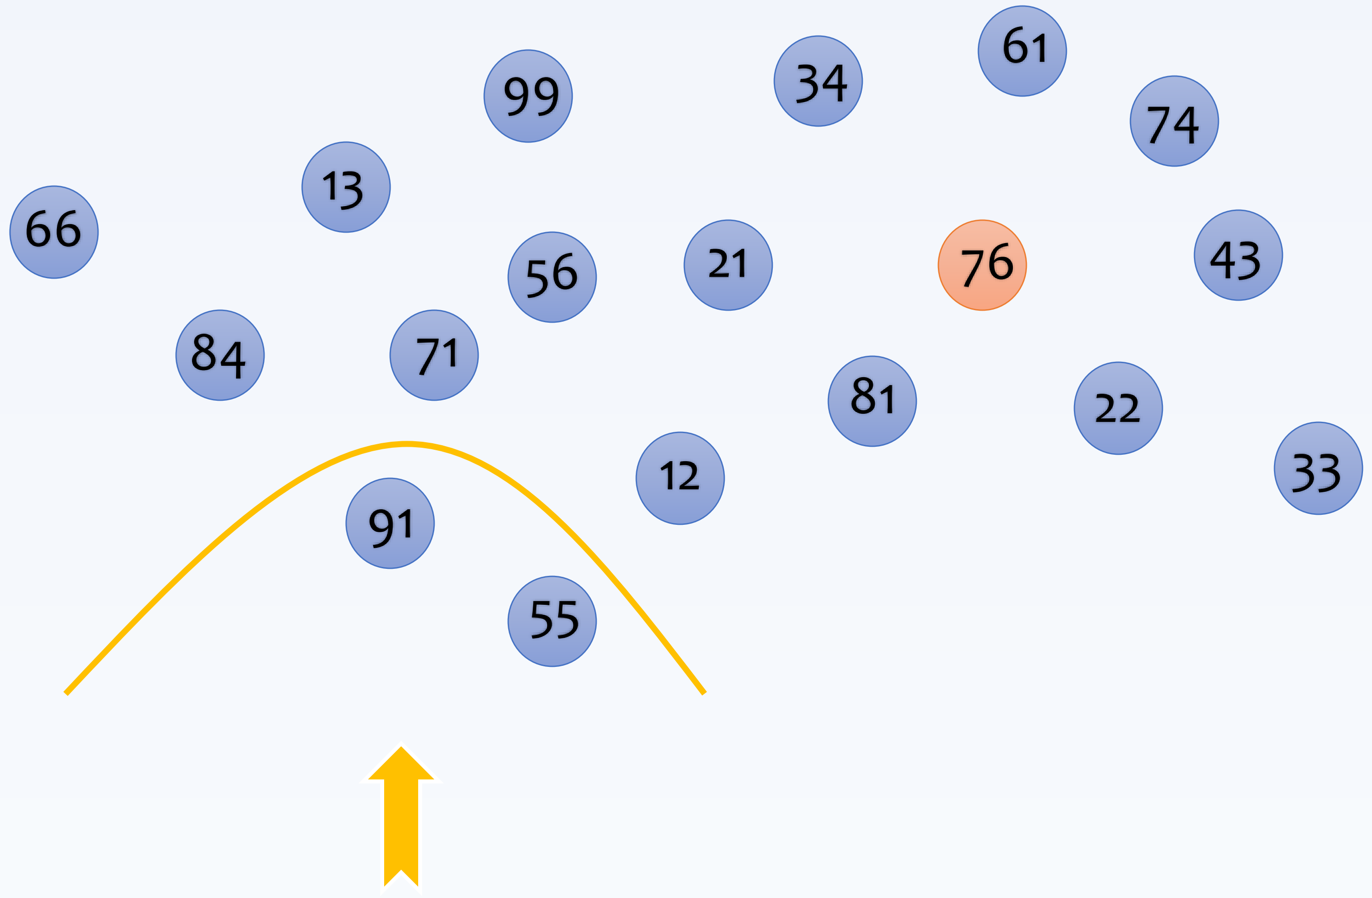

Searching for 76

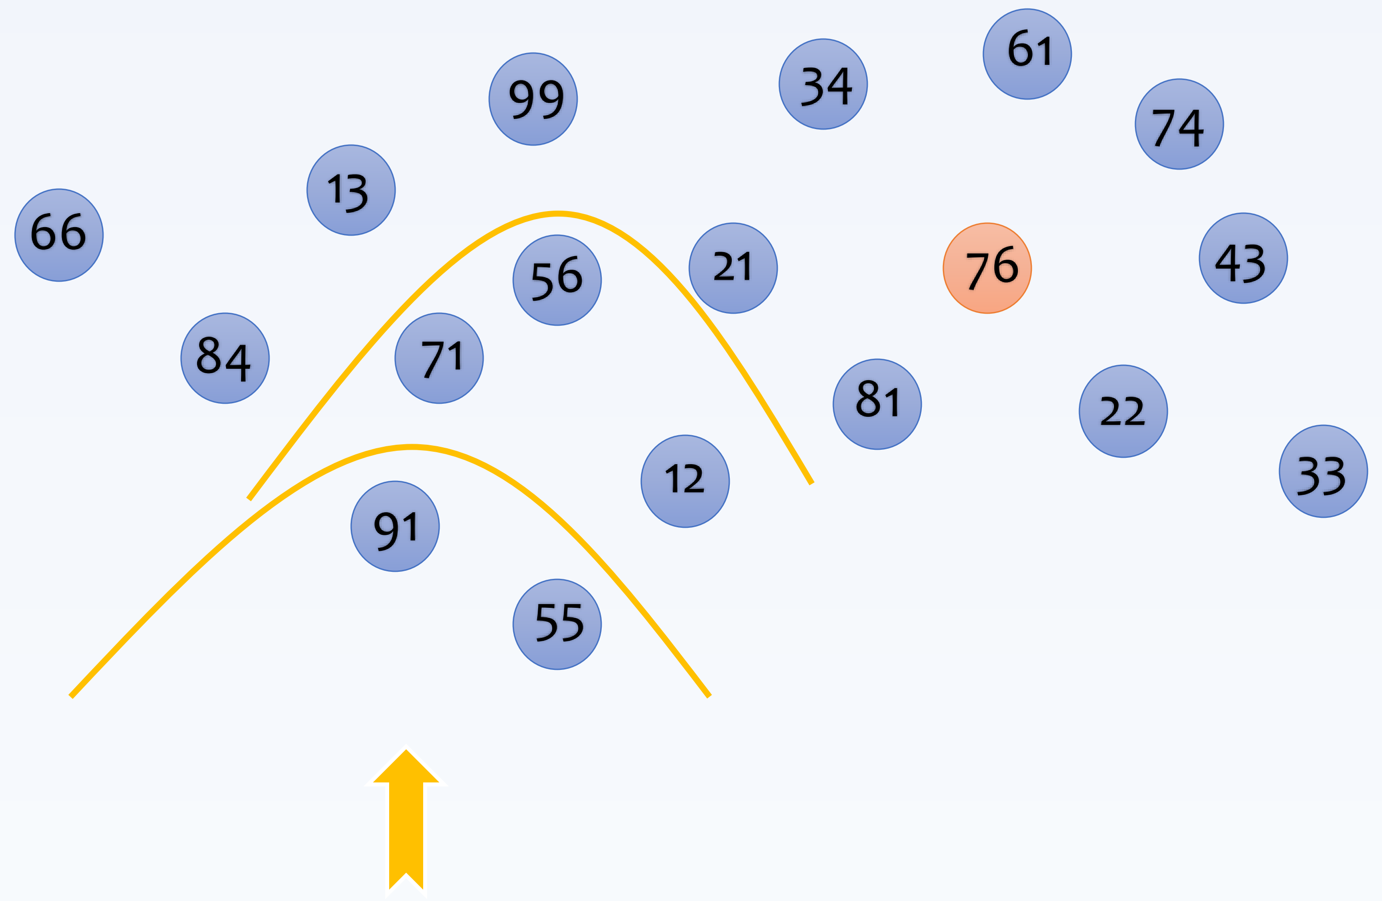

Searching for 76

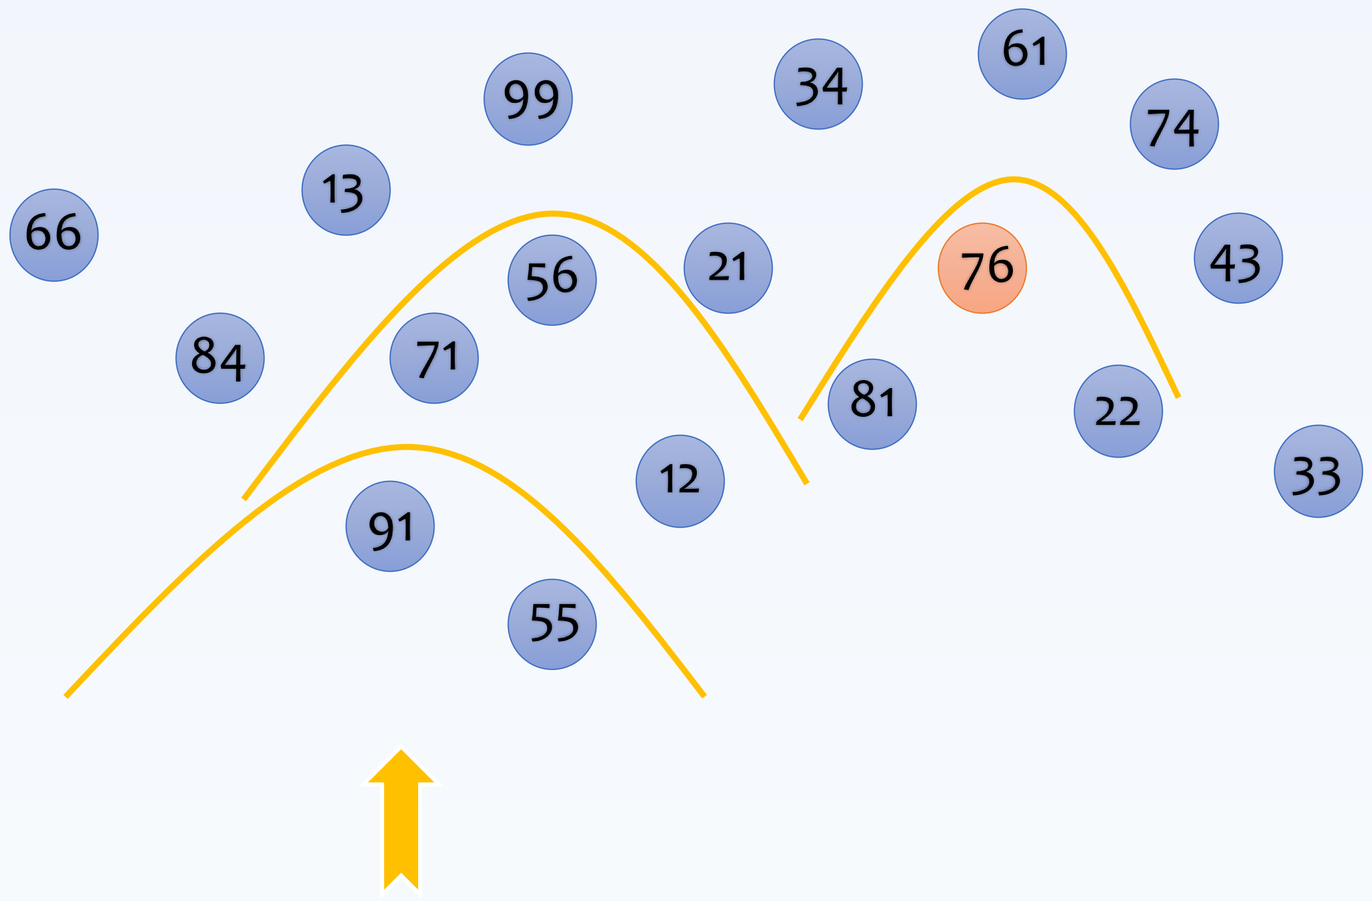

Searching for 76

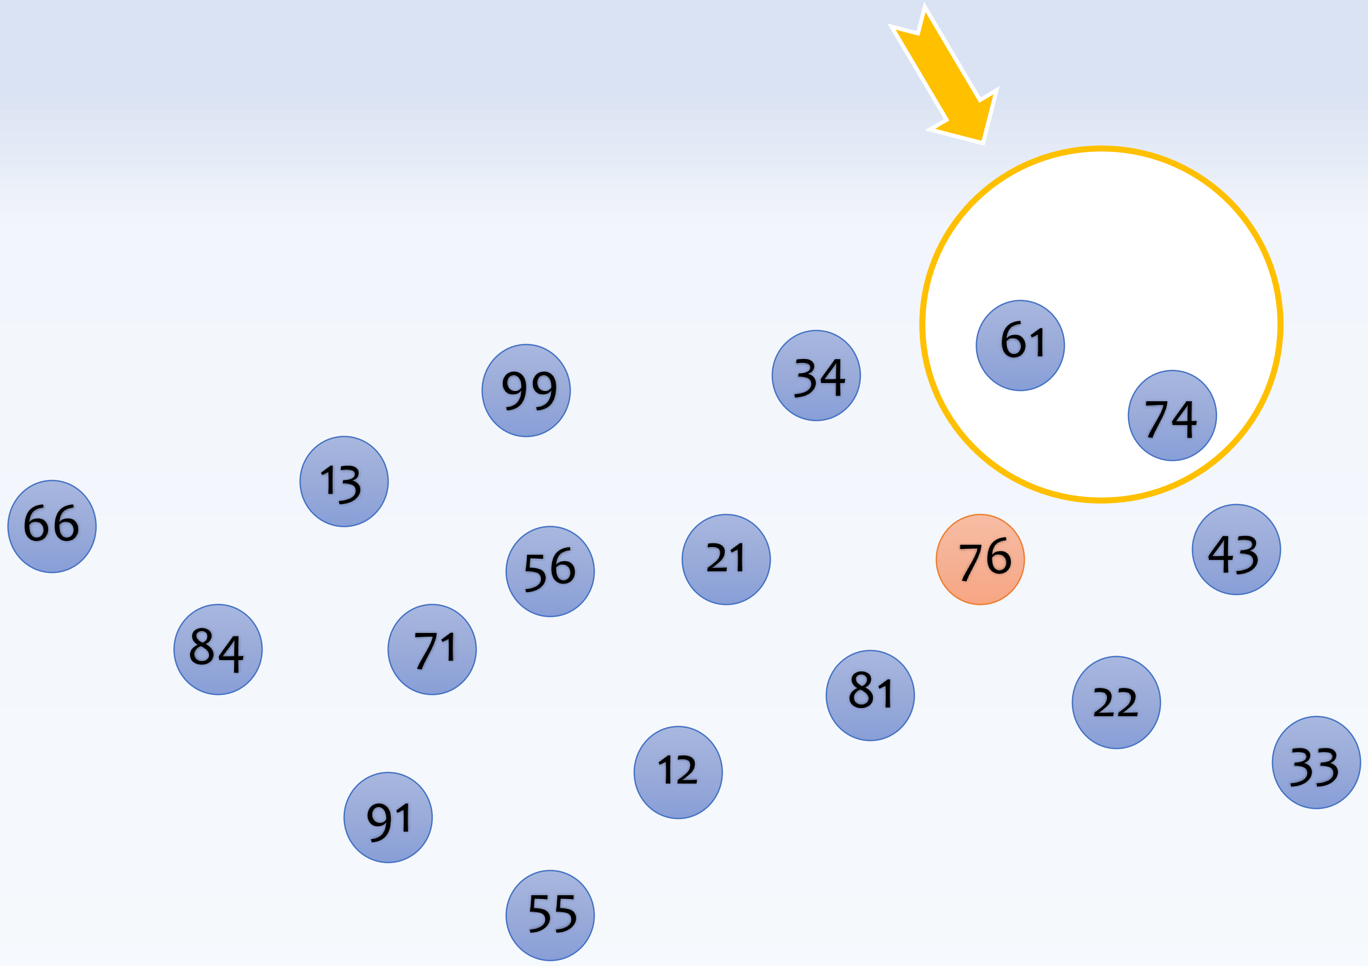

Searching for 76

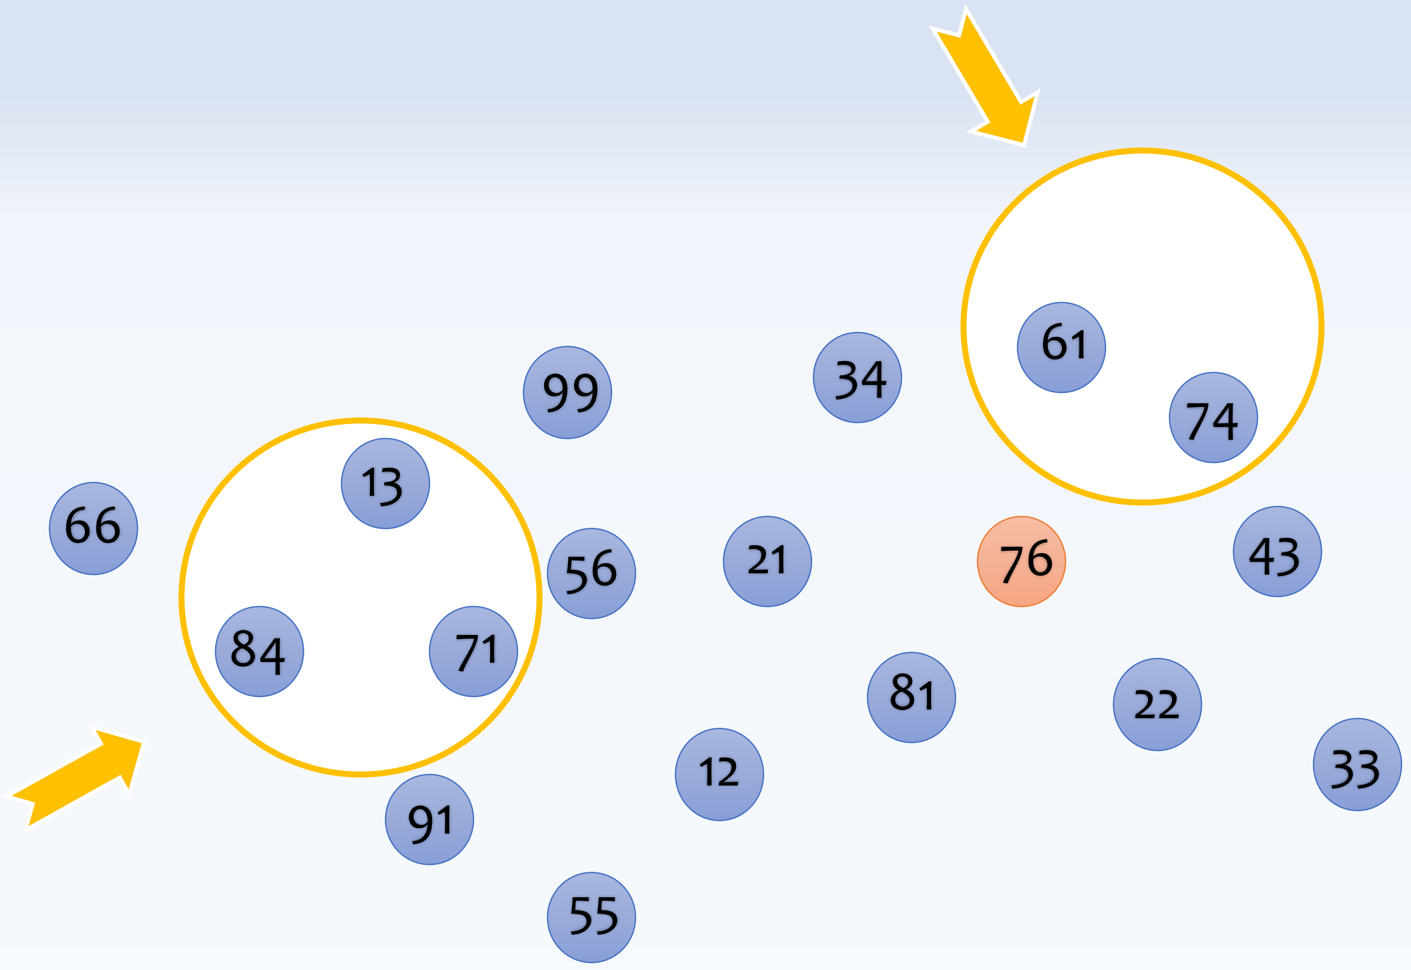

Searching for 76

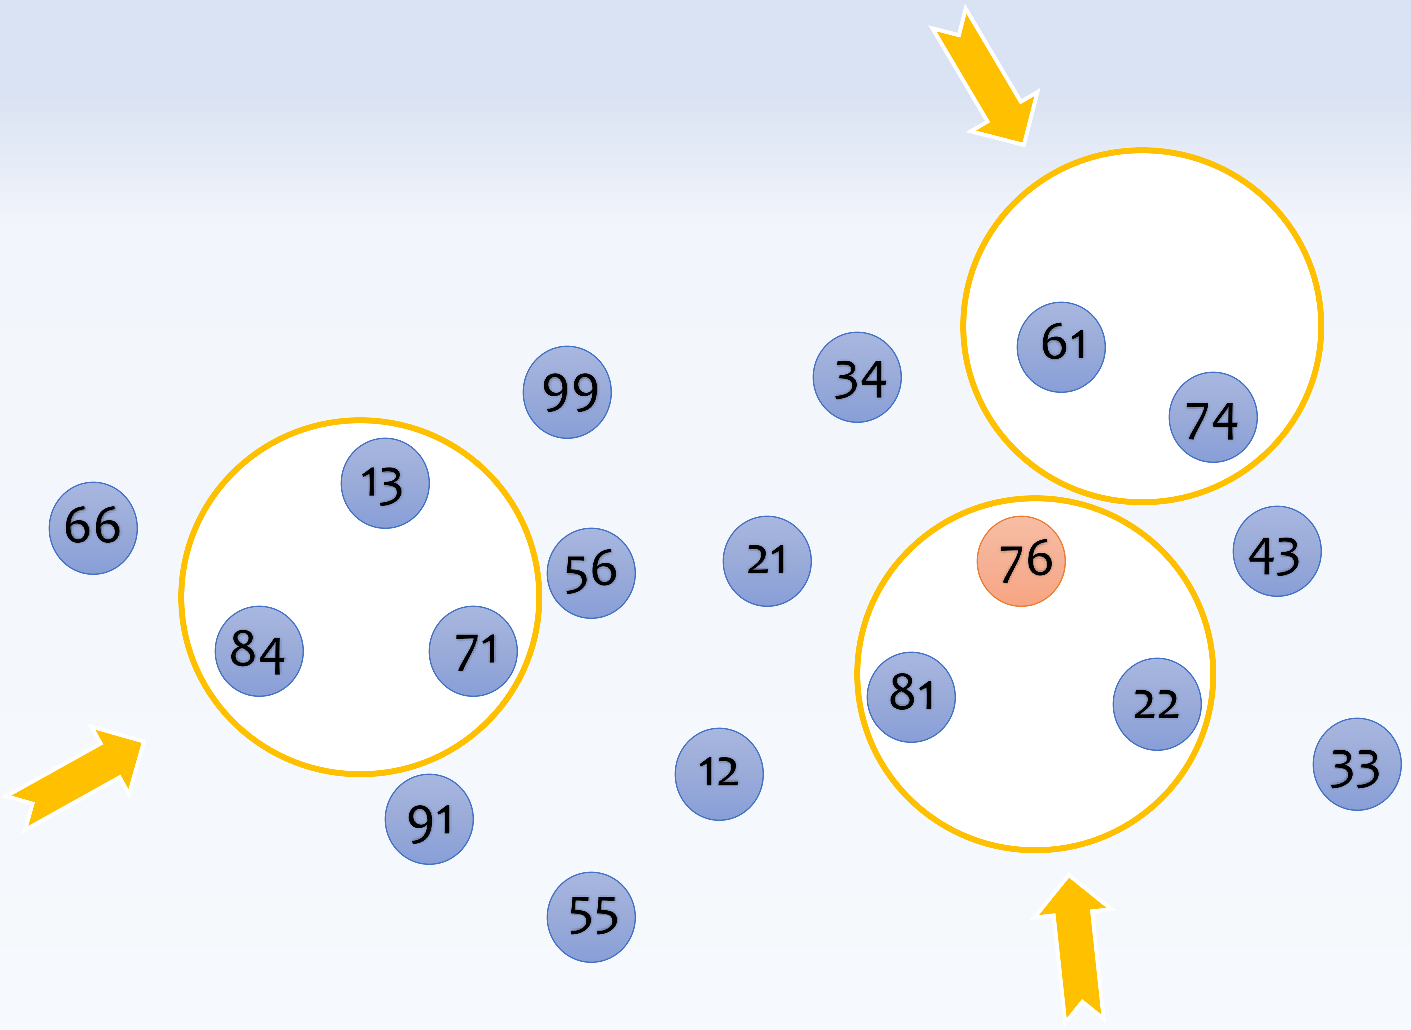

Searching for 44?
(what-if the value does not exist)
(could we have an early termination?)

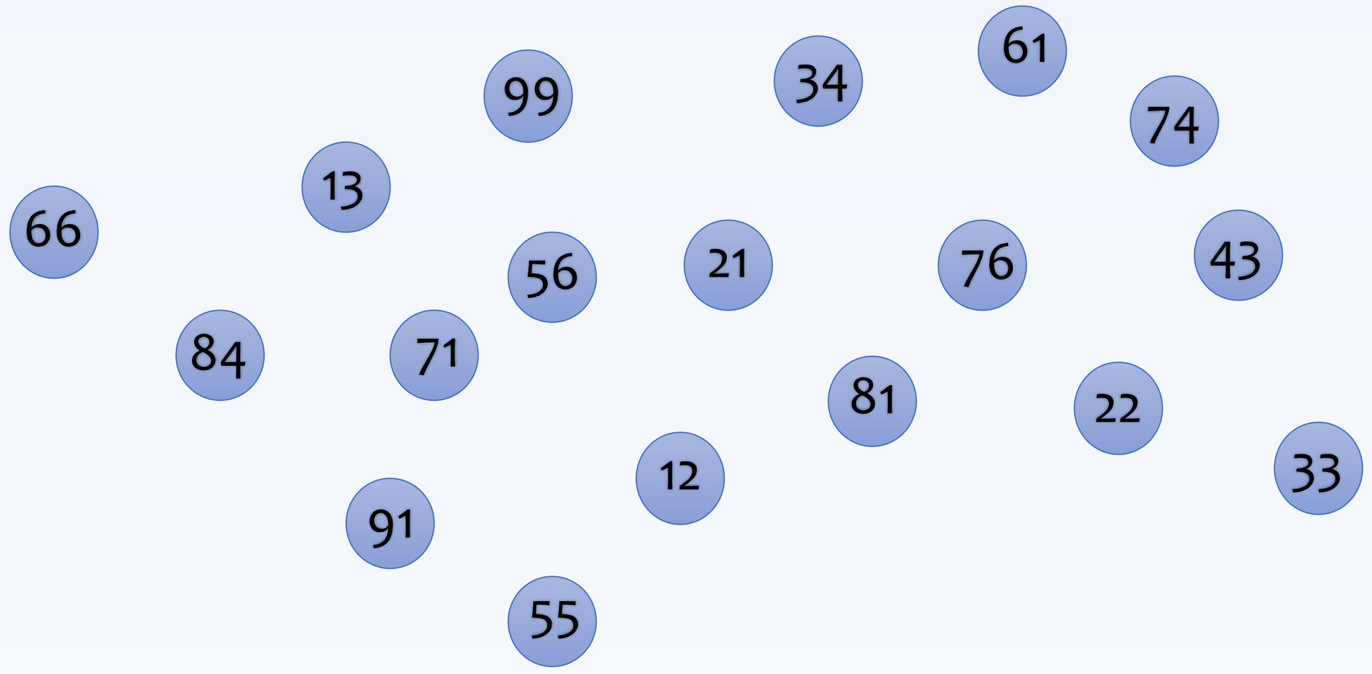

Could we impose an order to improve the search?

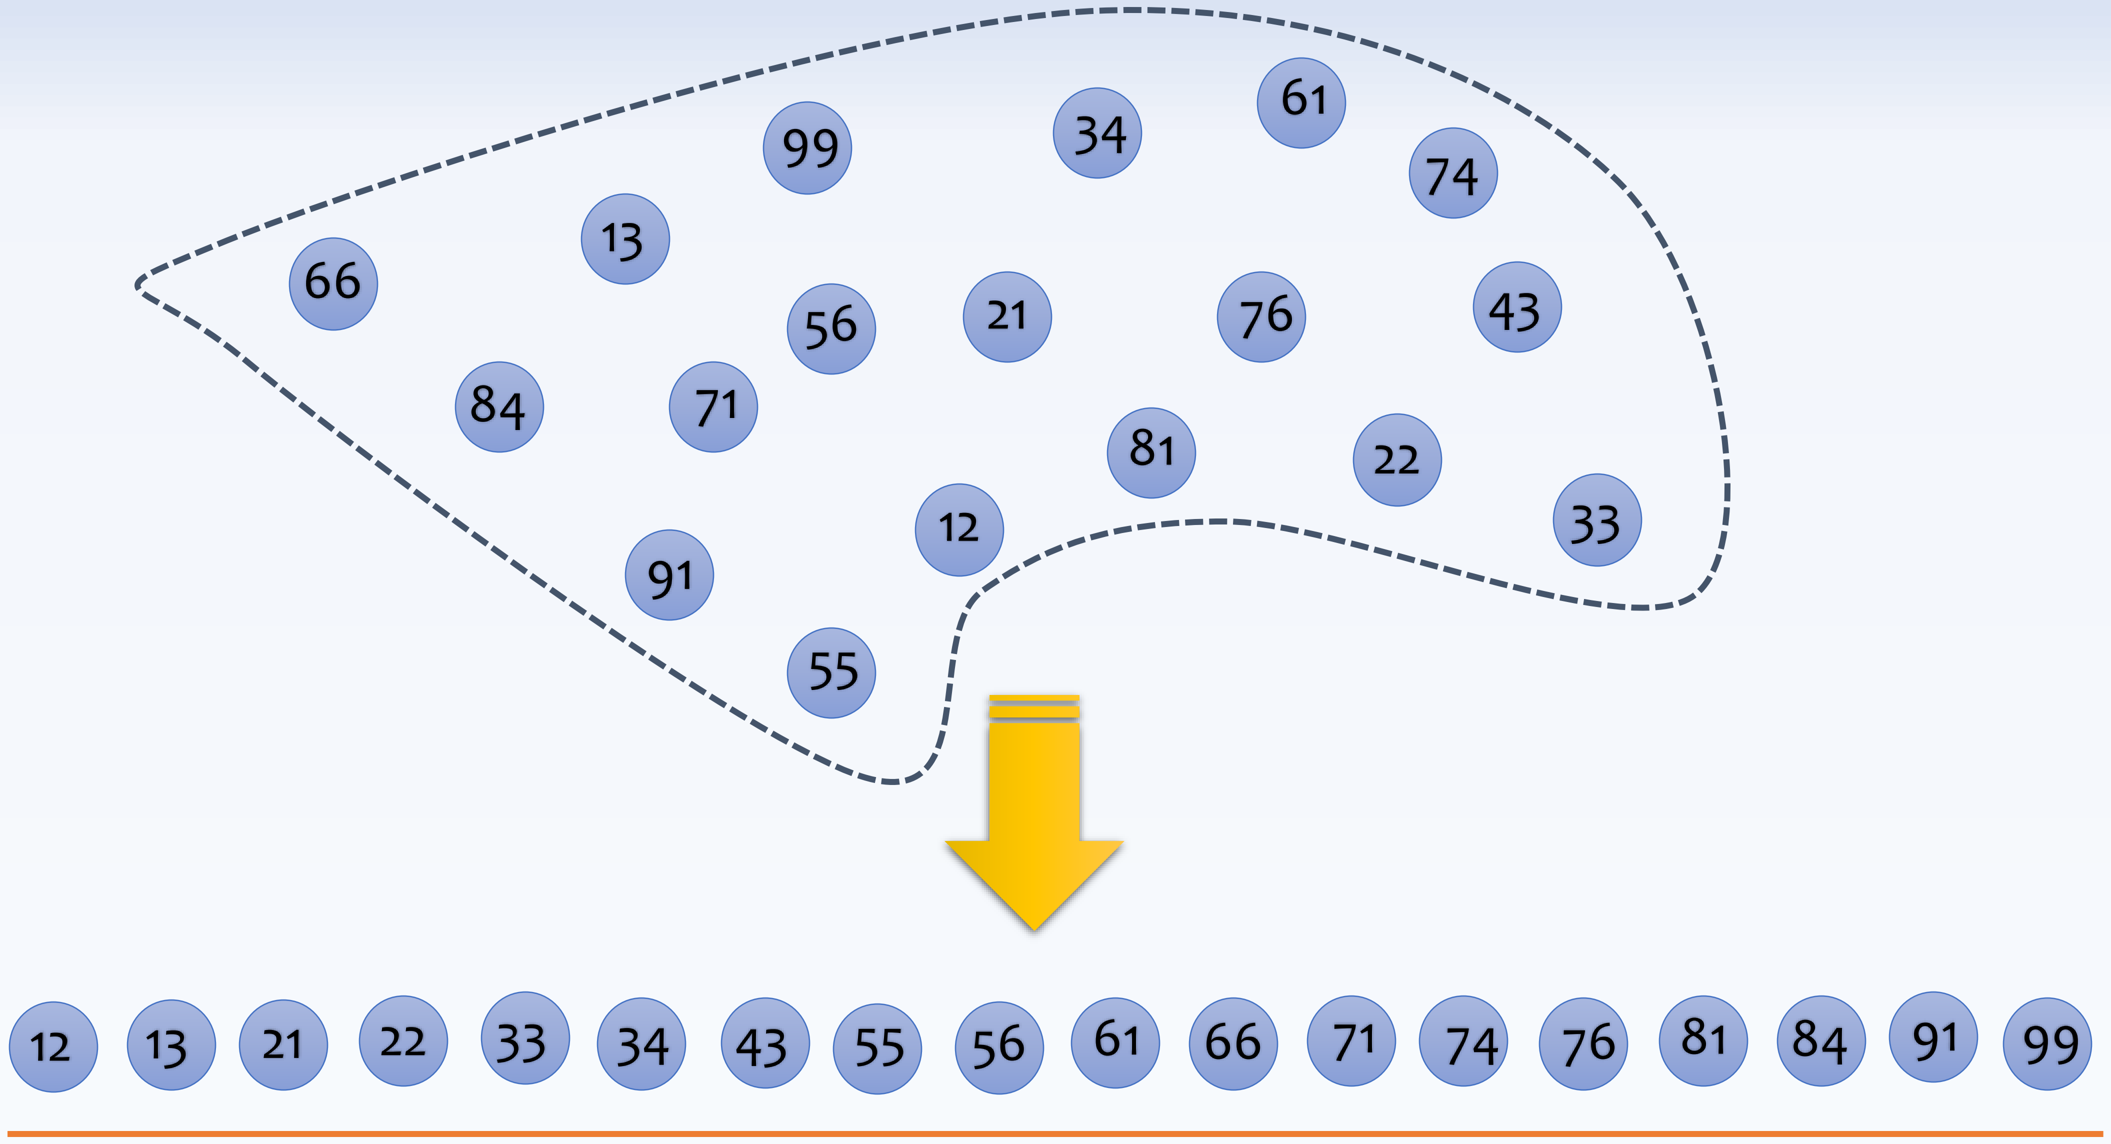
99
34
61
74
13
66
56
21
76
43
84
71
81
22
33
12
91
55
12
13
21
22
33
34
43
55
56
61
66
71
74
76
81
84
91
99

Searching for 76

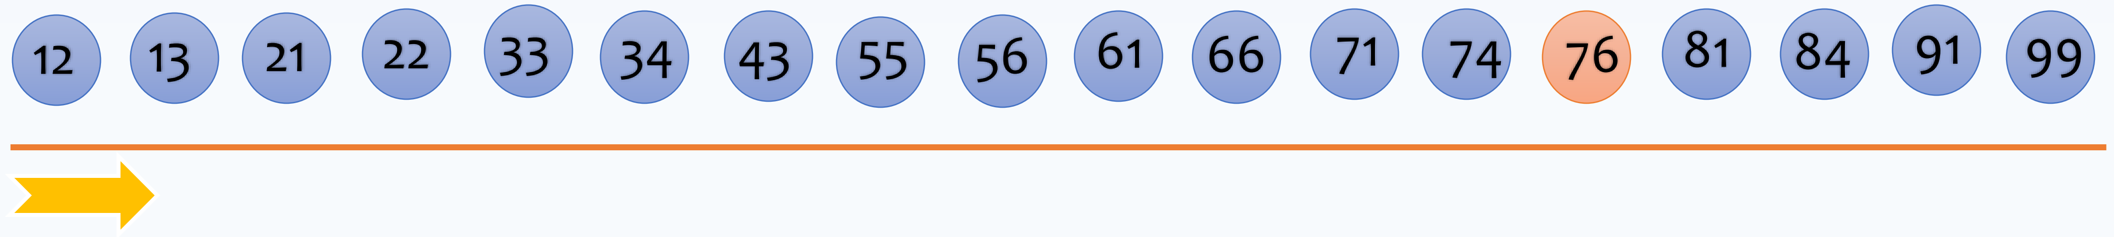

Searching for 76

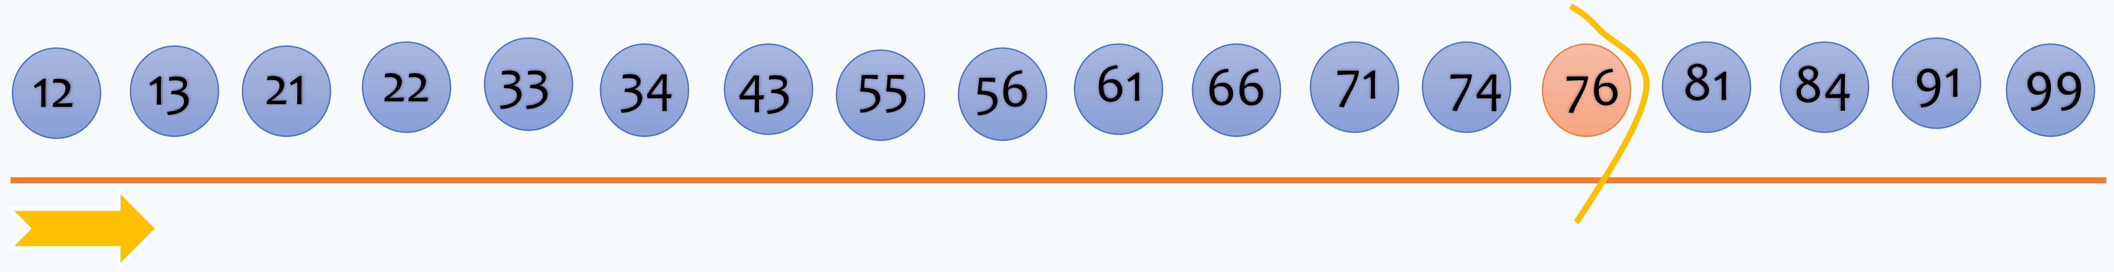

Searching for 44?
(could we have an early termination?)

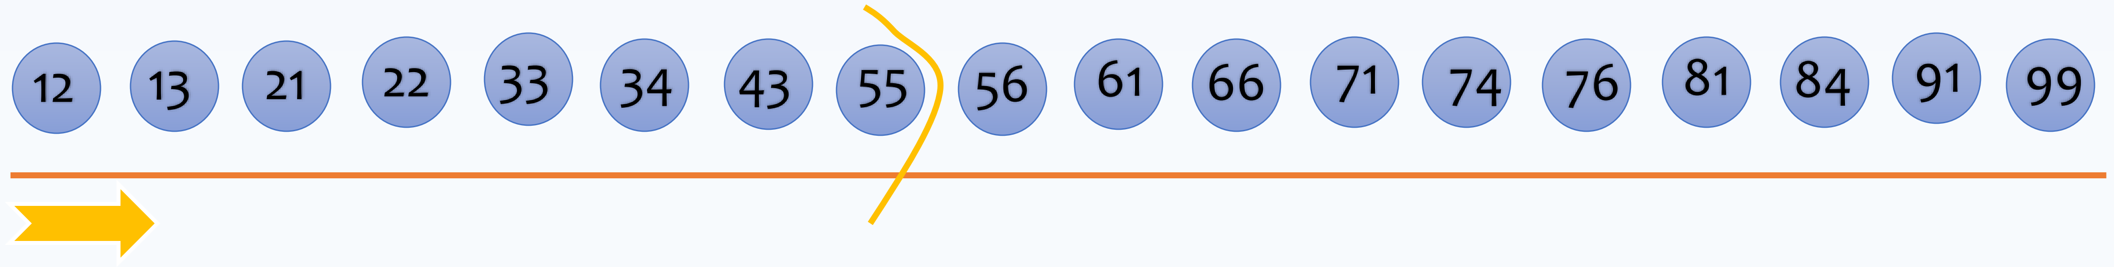

Could we impose a structure to further improve the search?

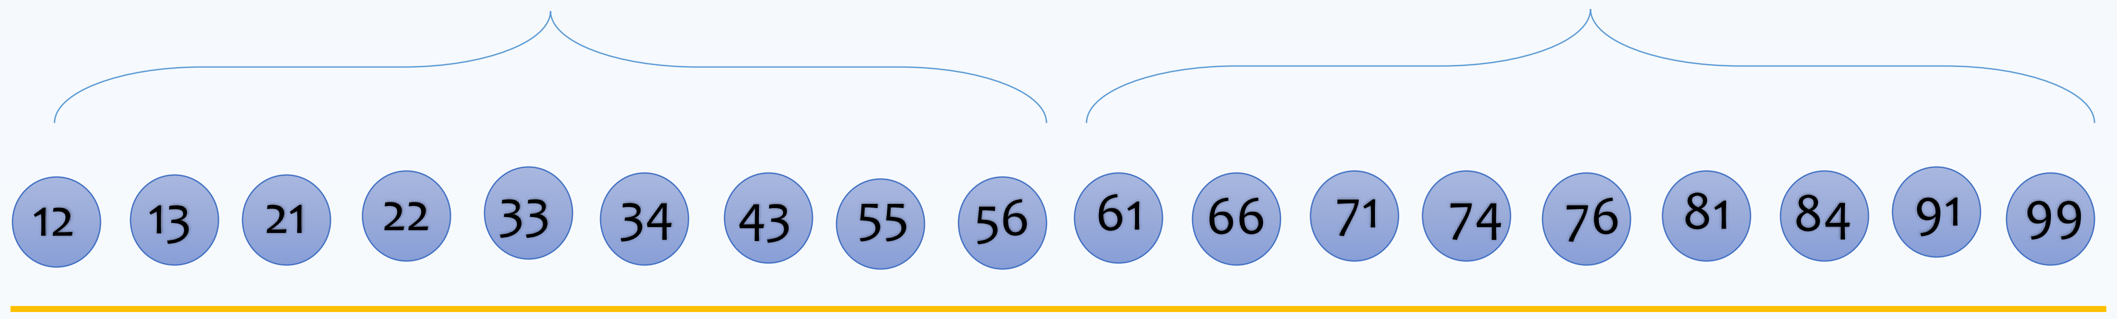
12
13
21
22
33
34
43
55
56
61
66
71
74
76
81
84
91
99

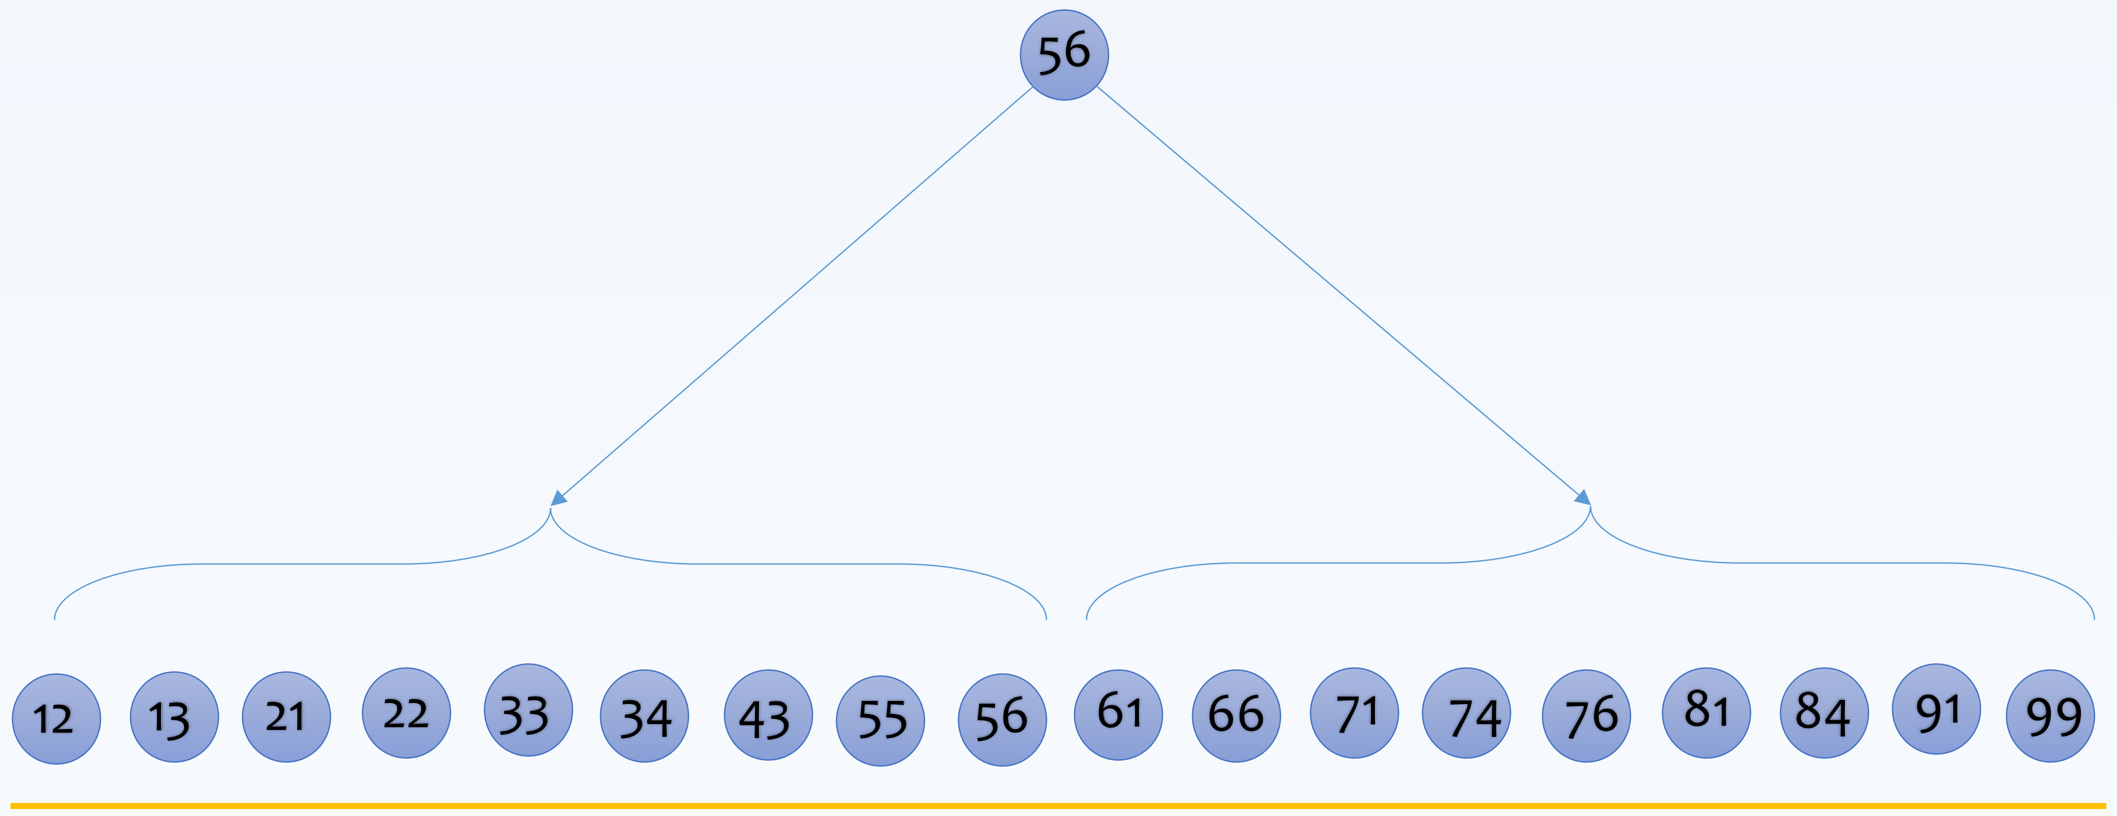
56
12
13
21
22
33
34
43
55
56
61
66
71
74
76
81
84
91
99

Searching for 76

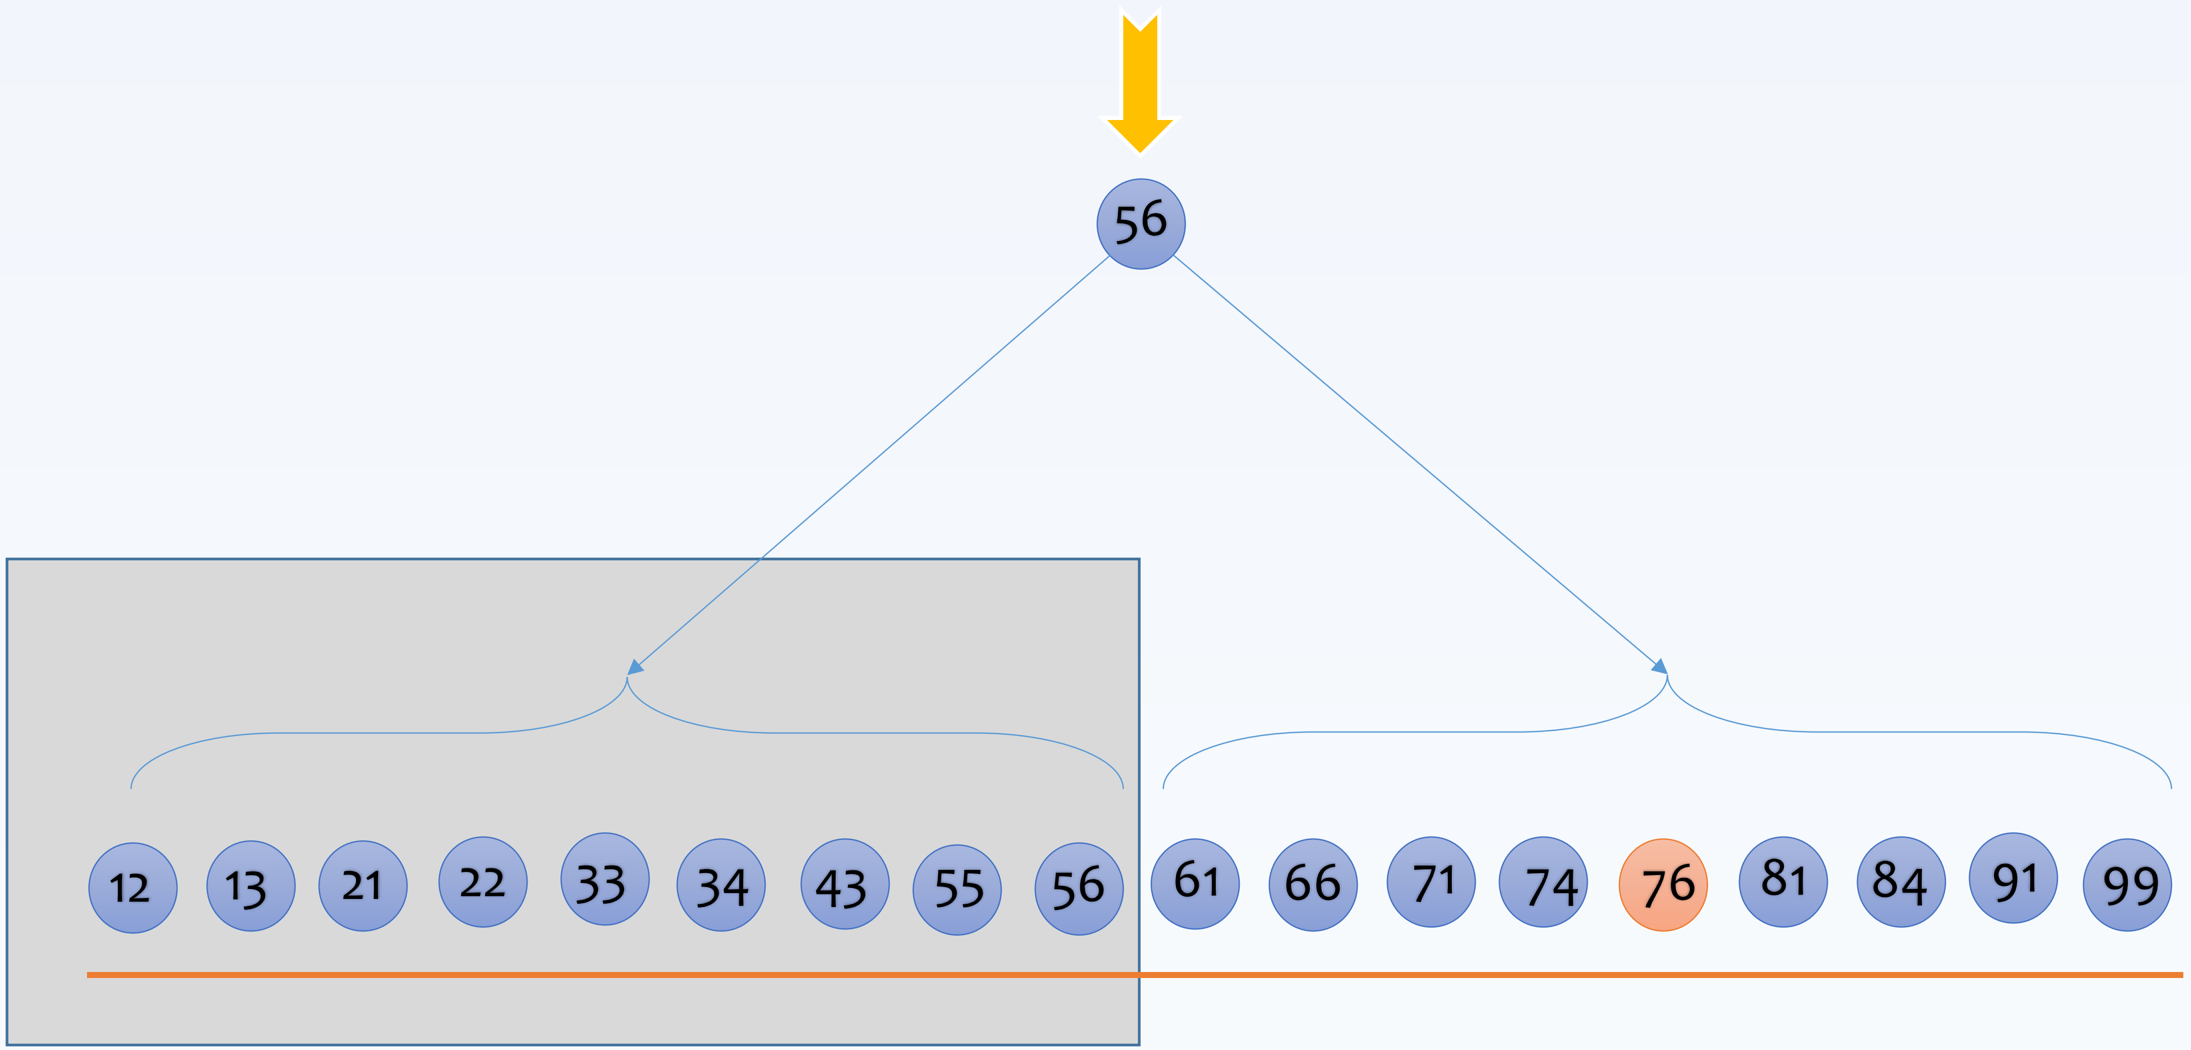

Searching for 76

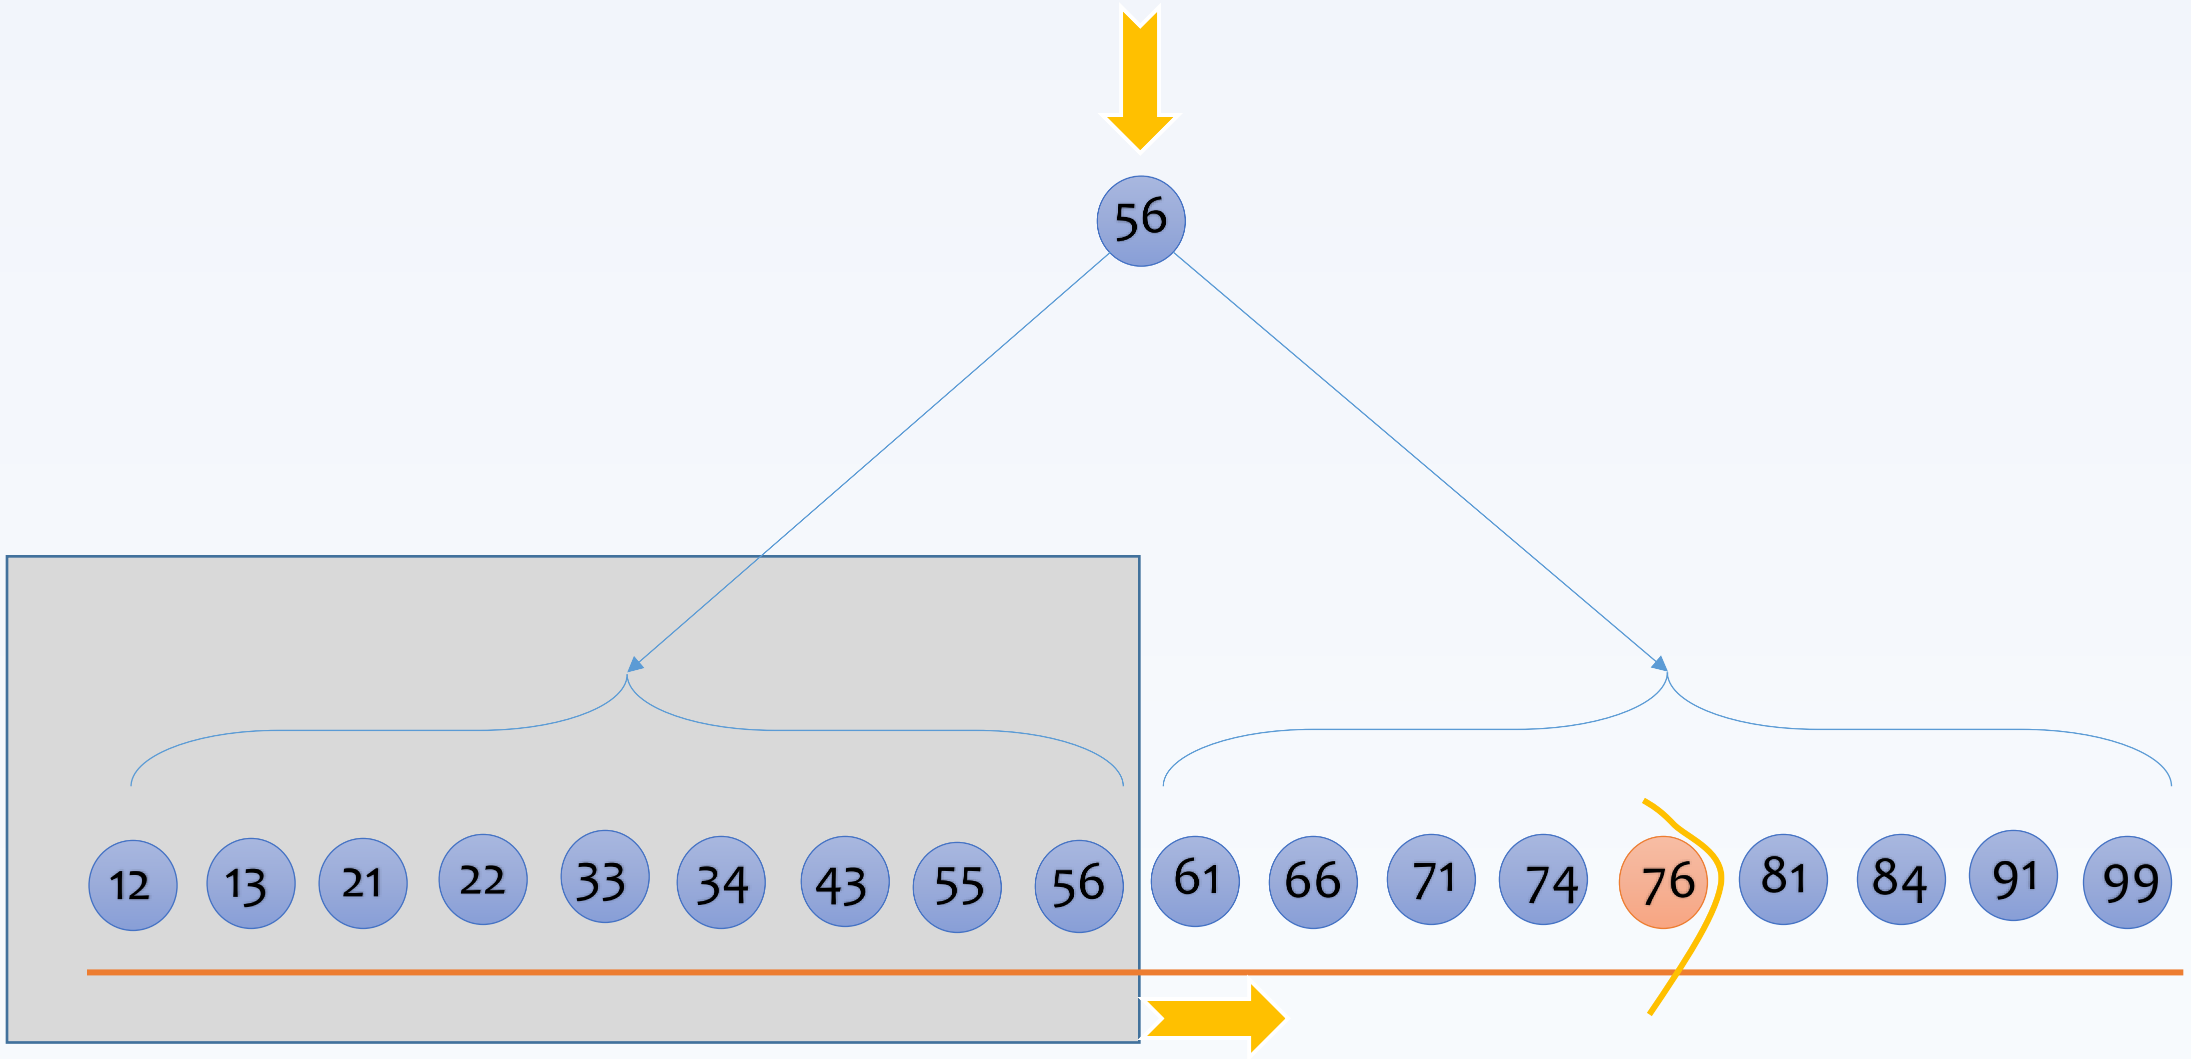

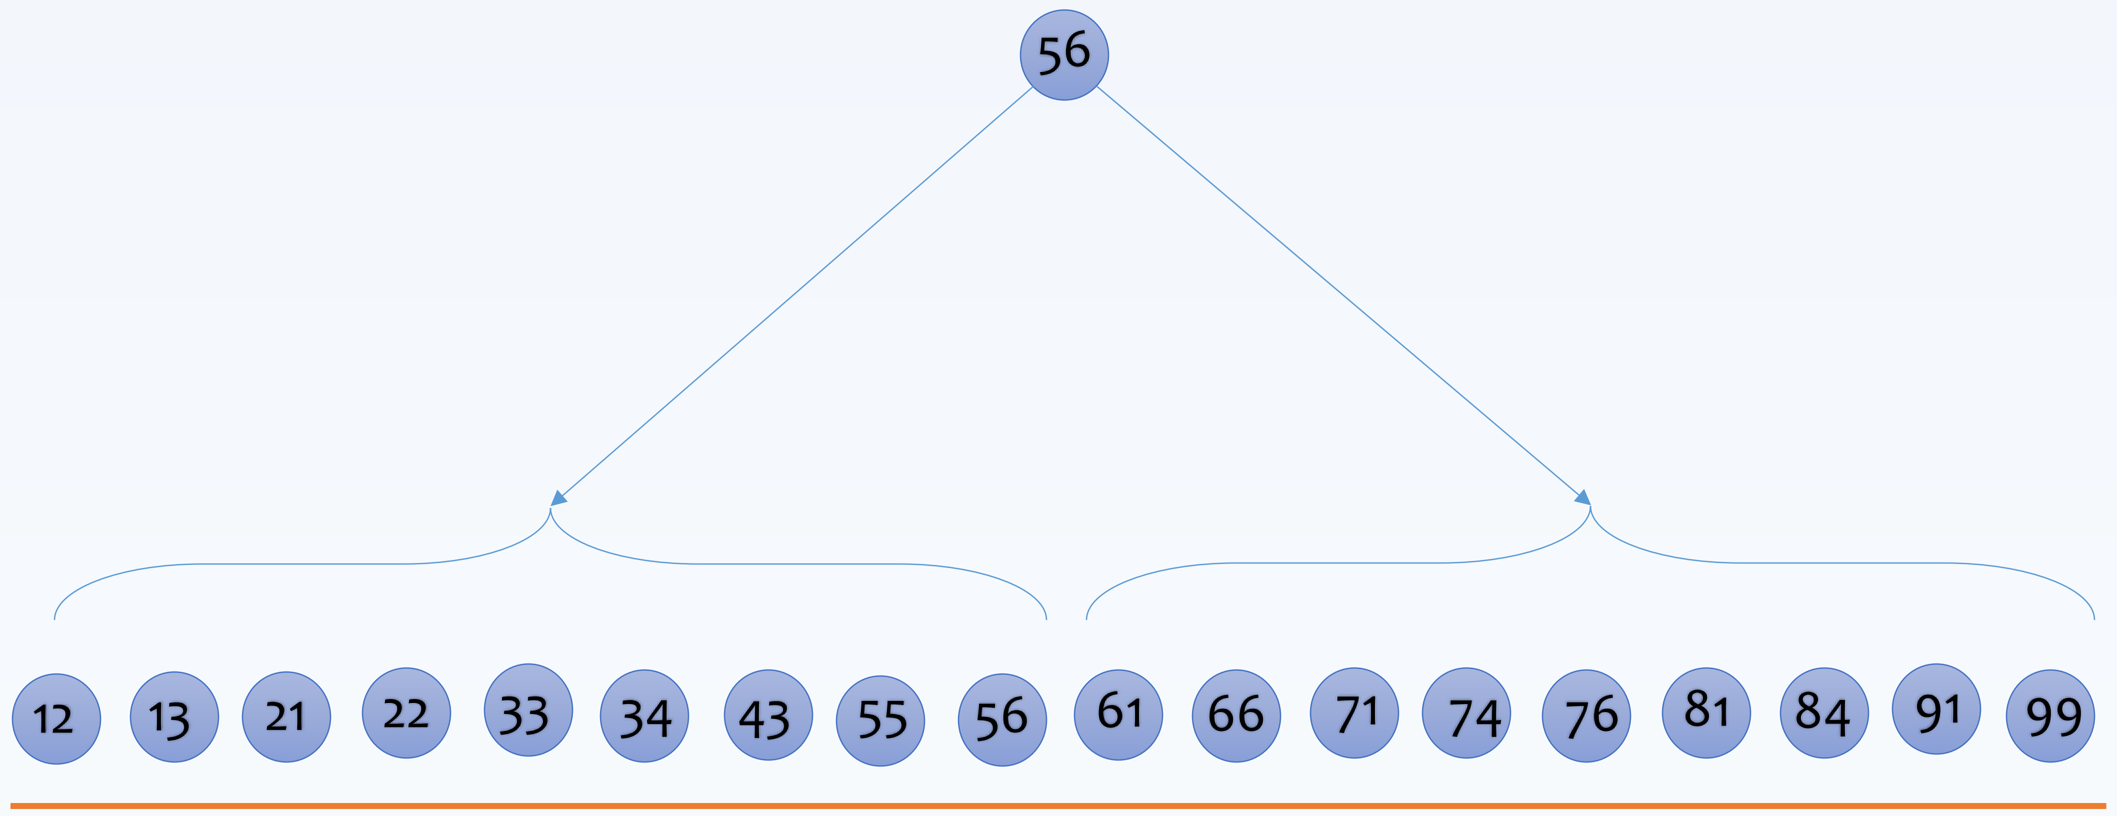
56
12
13
21
22
33
34
43
55
56
61
66
71
74
76
81
84
91
99

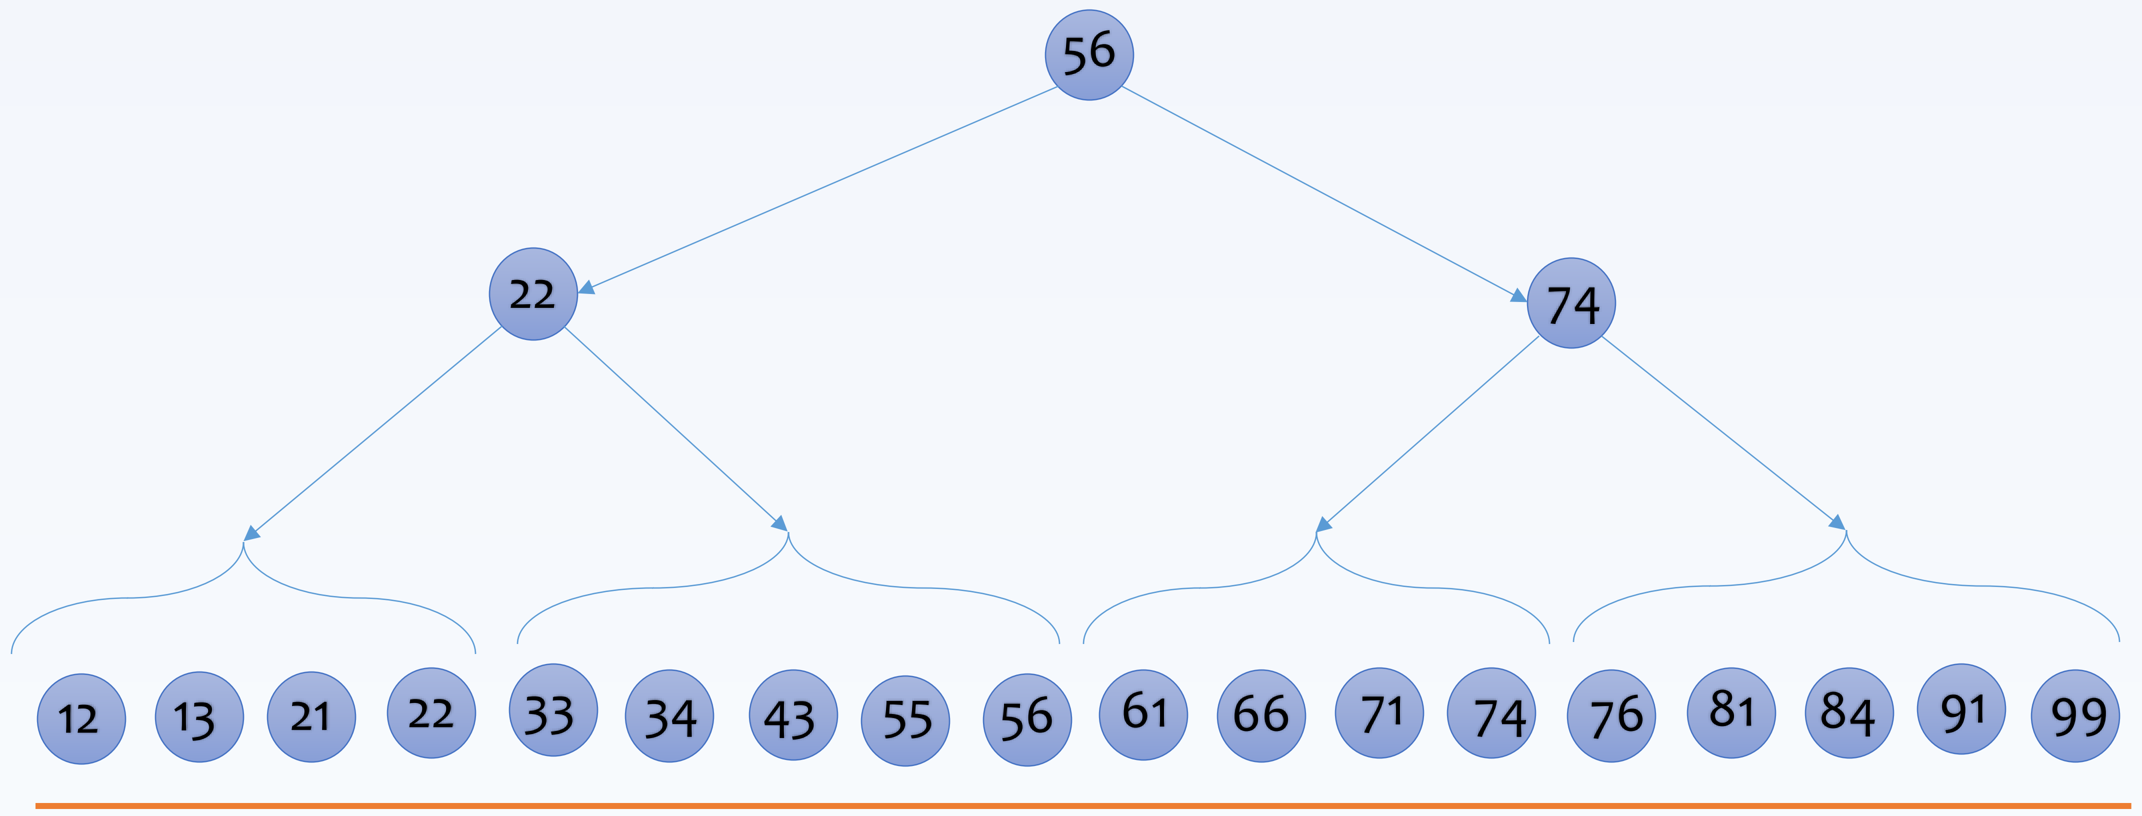
56
22
74
12
13
21
22
33
34
43
55
56
61
66
71
74
76
81
84
91
99

Searching for 76

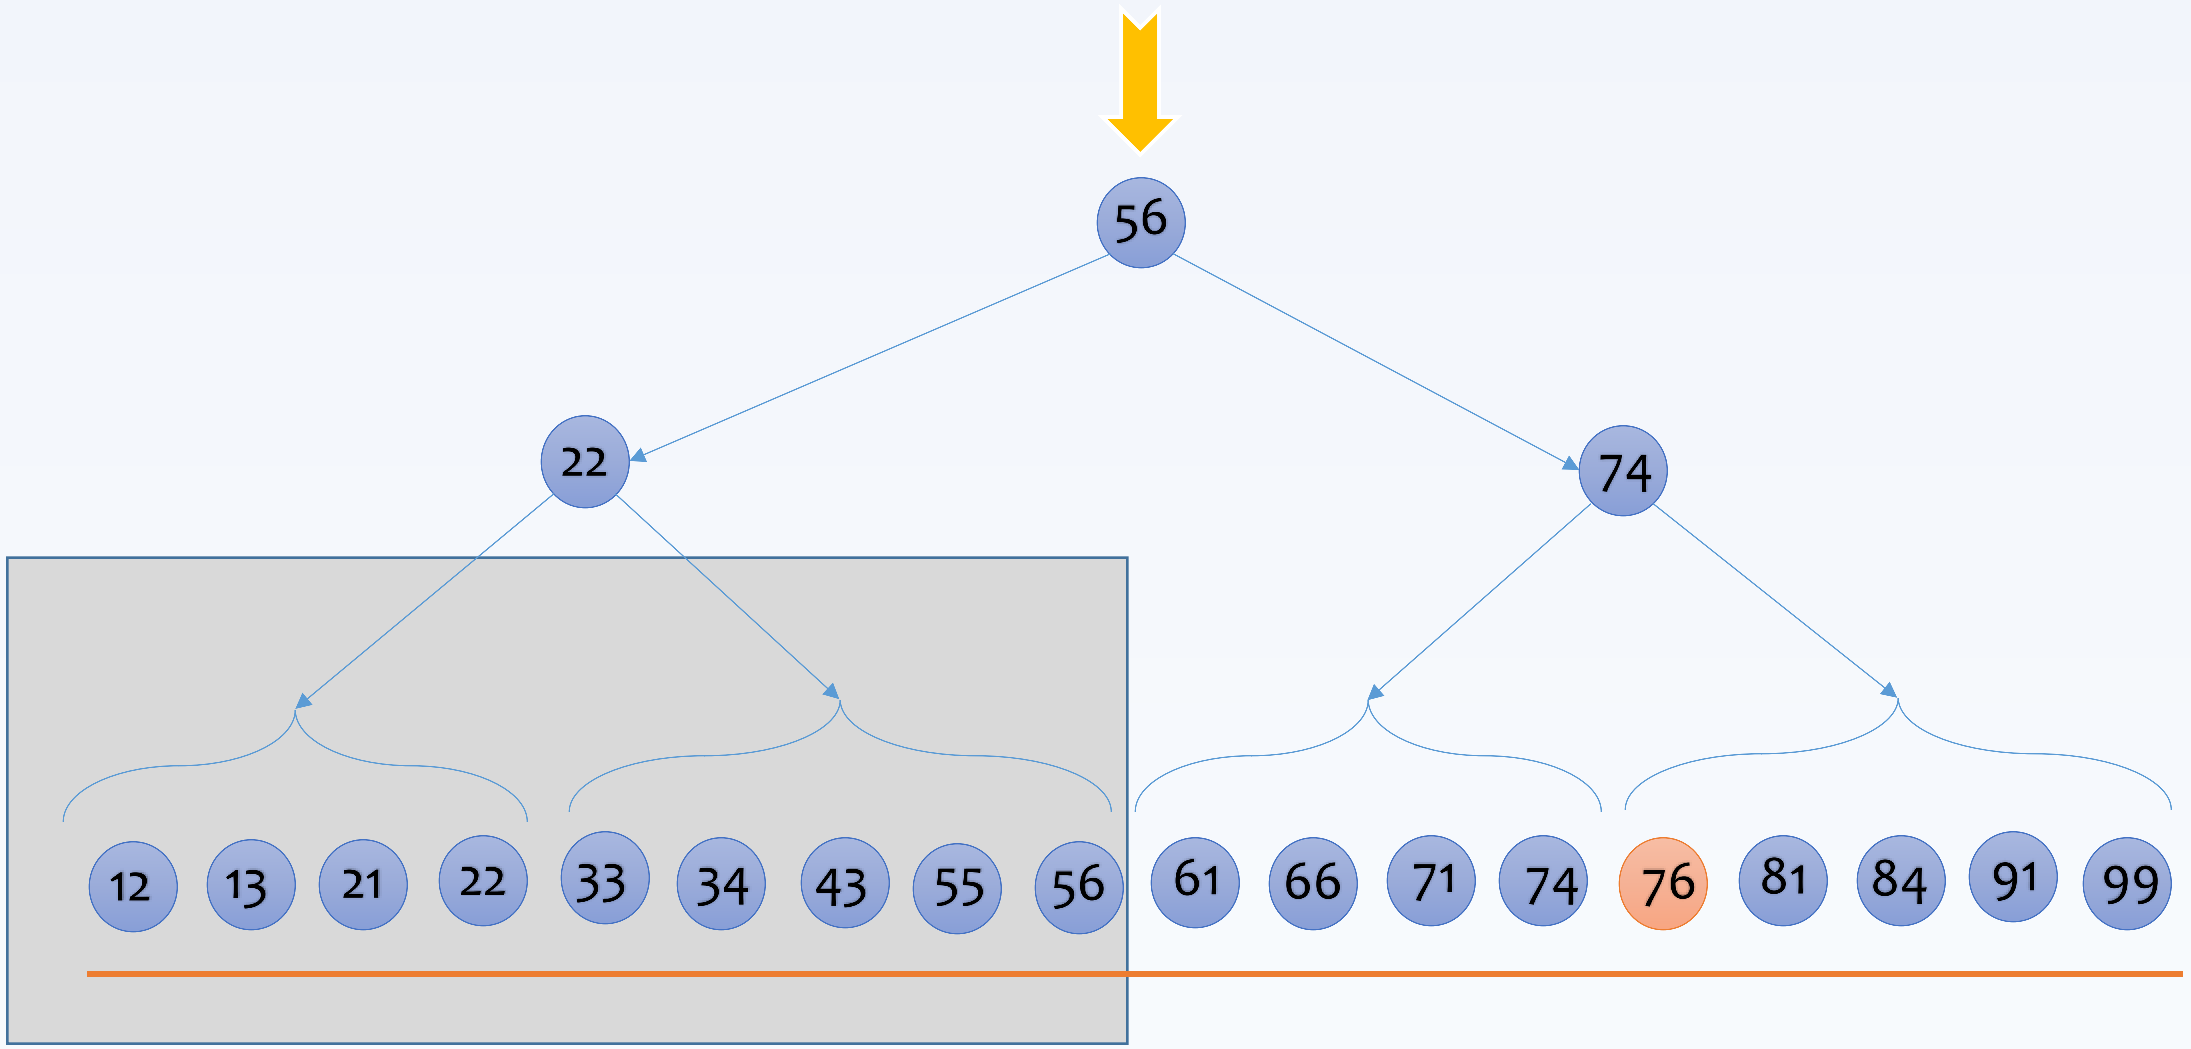

Searching for 76

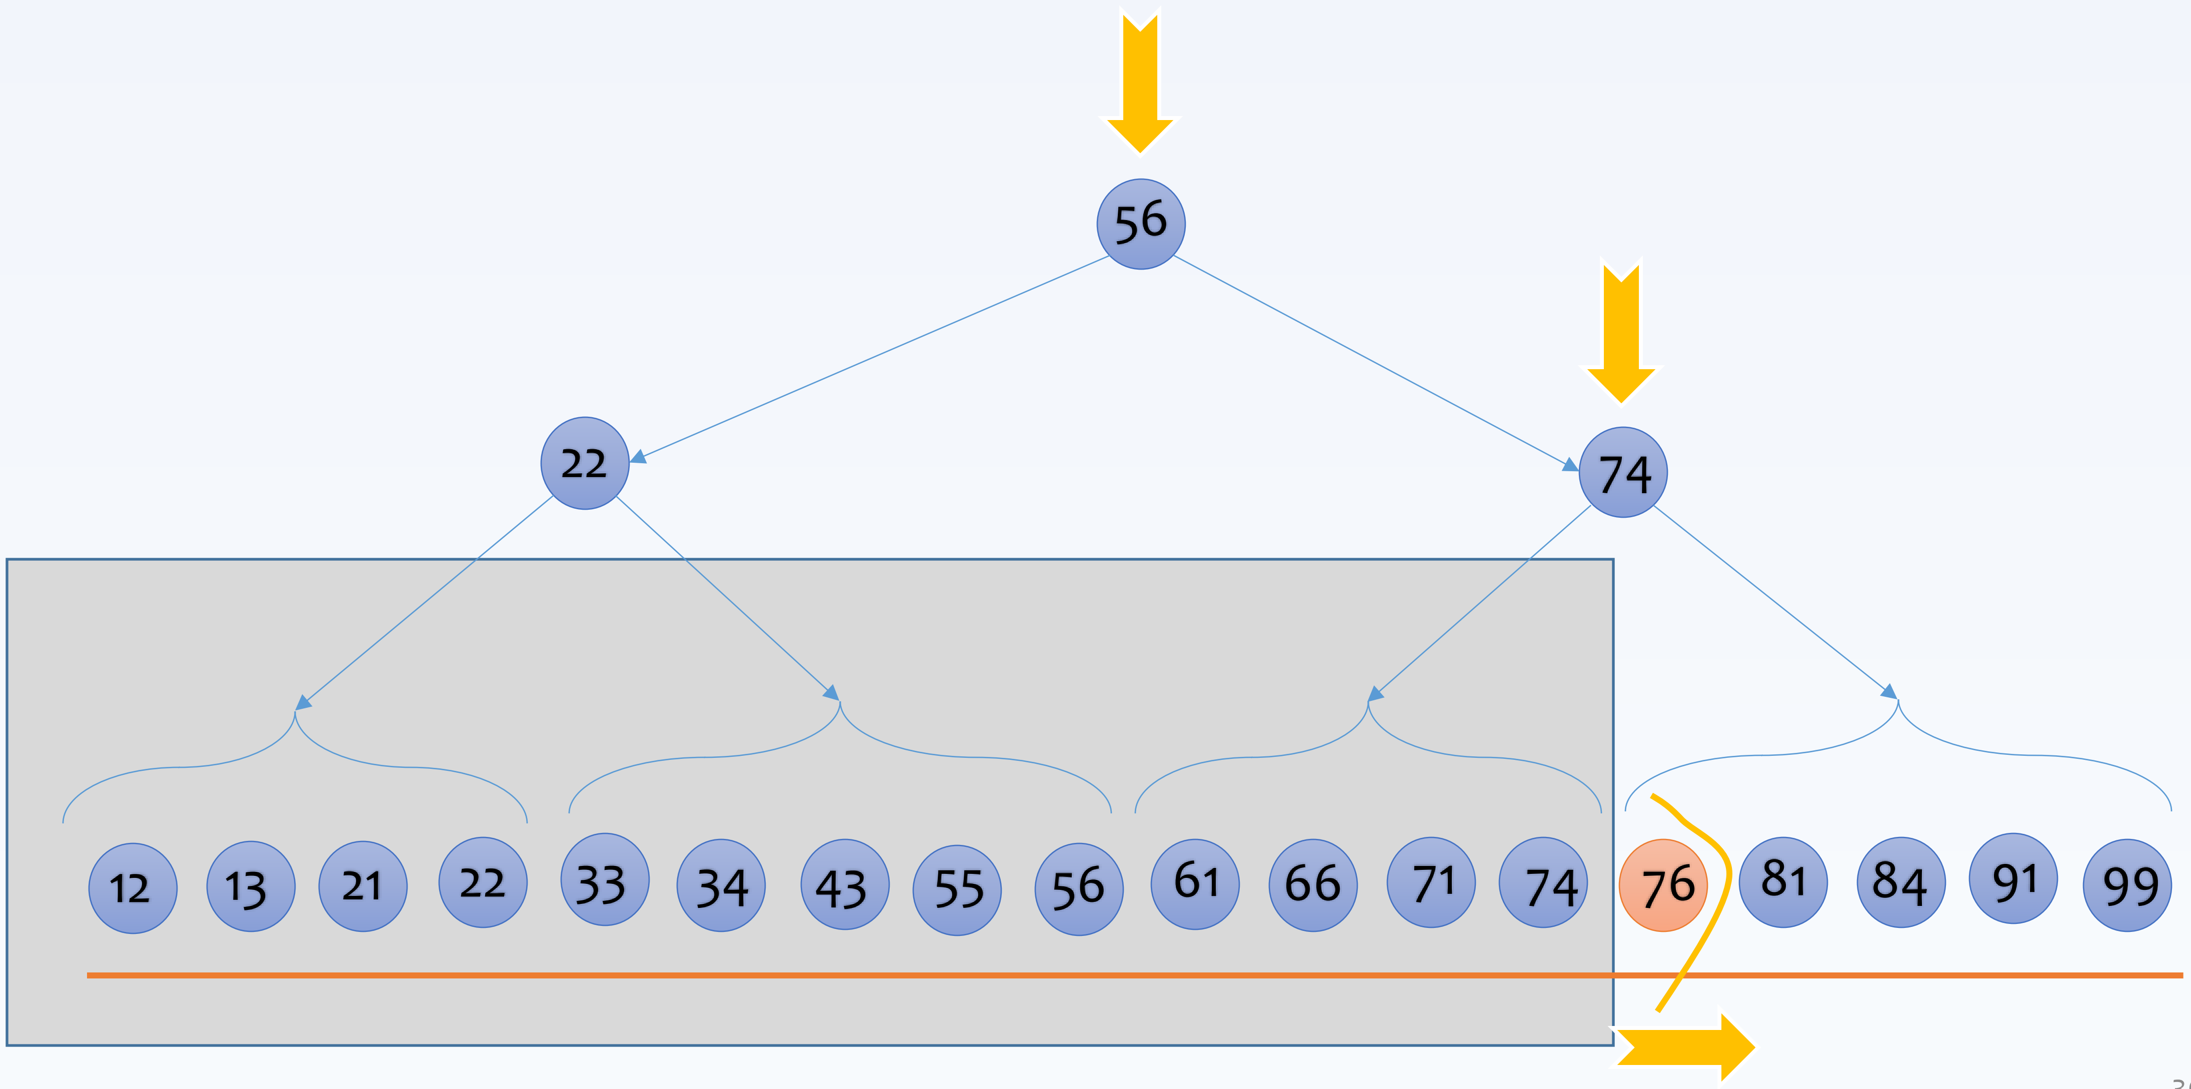

Searching for 44?
(could we have an early termination?)

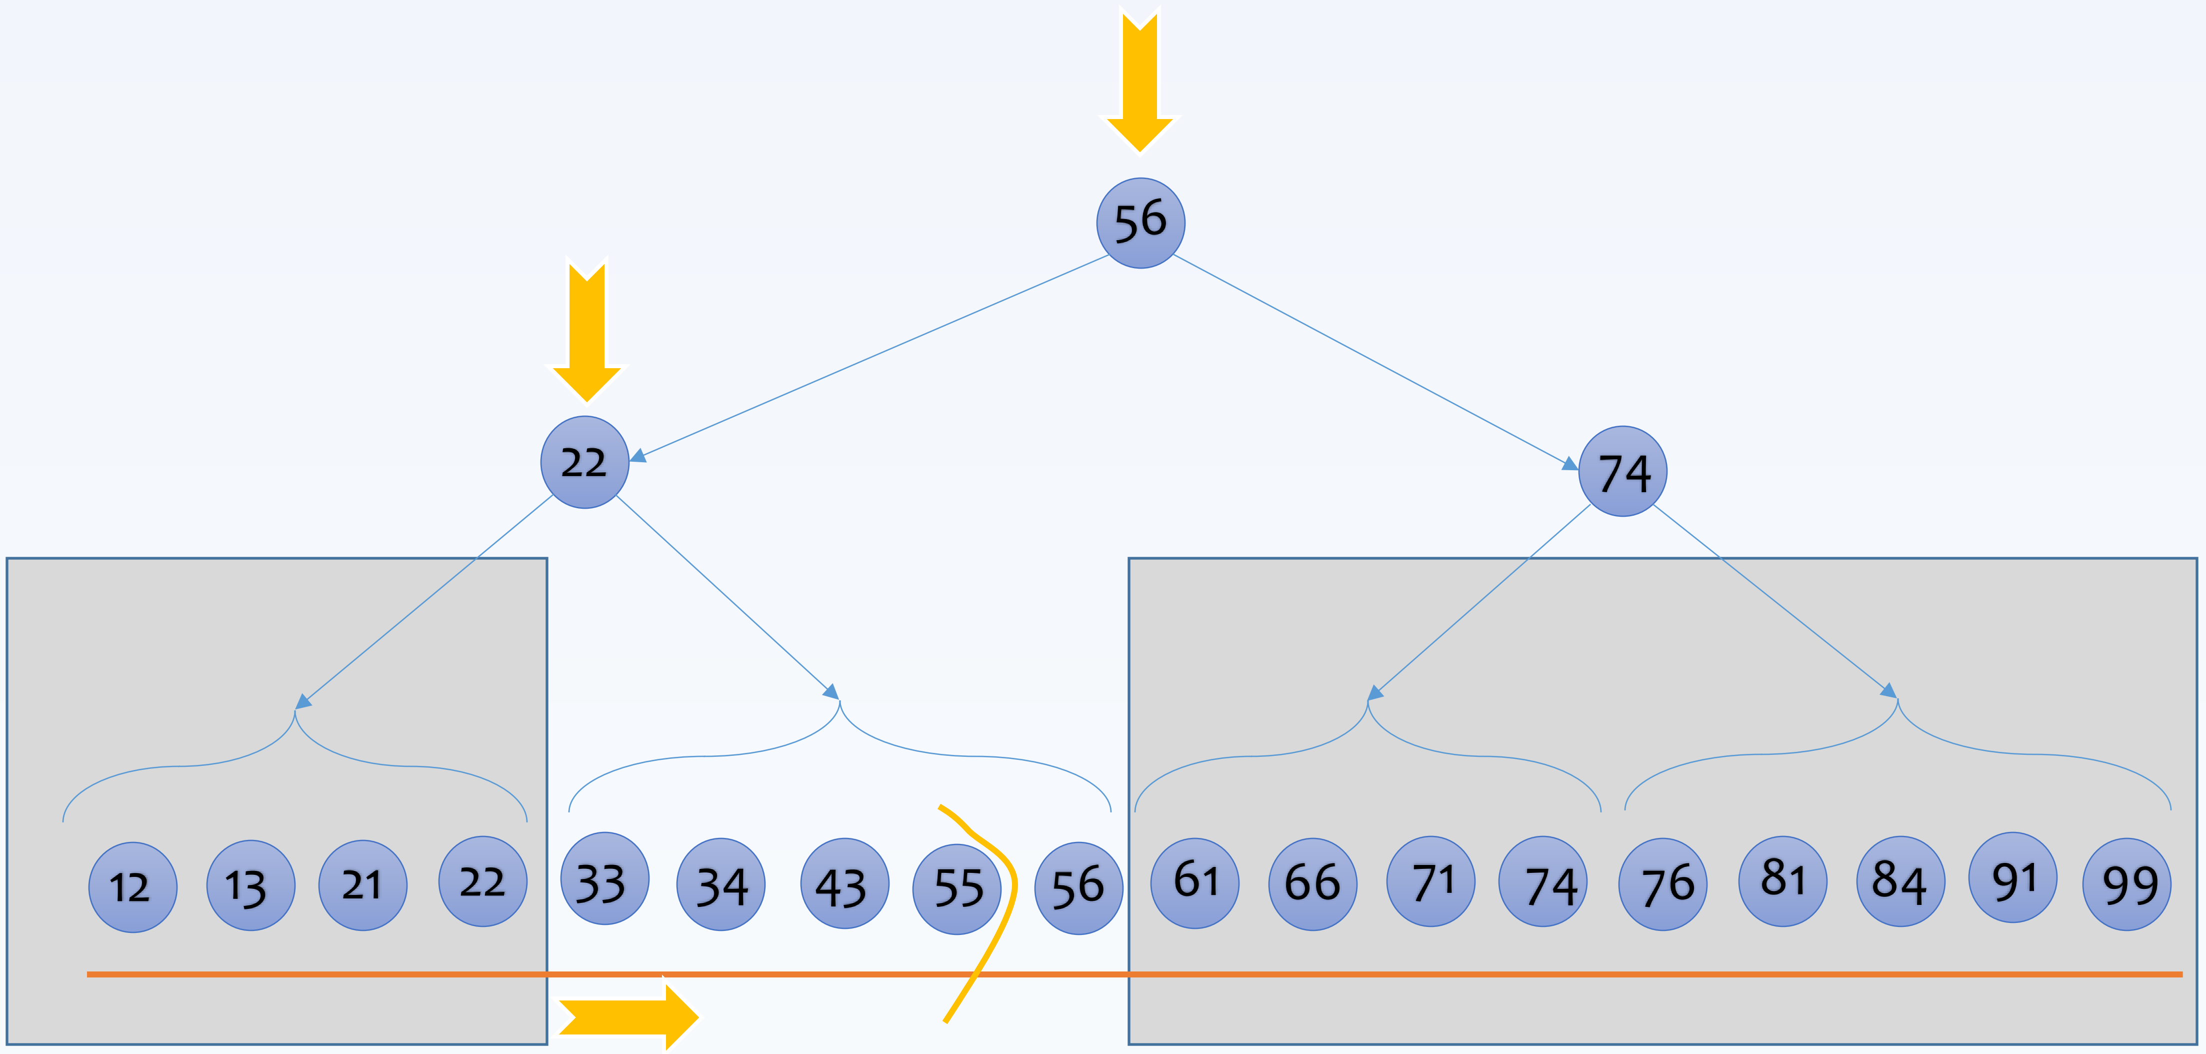

Searching for 76-91

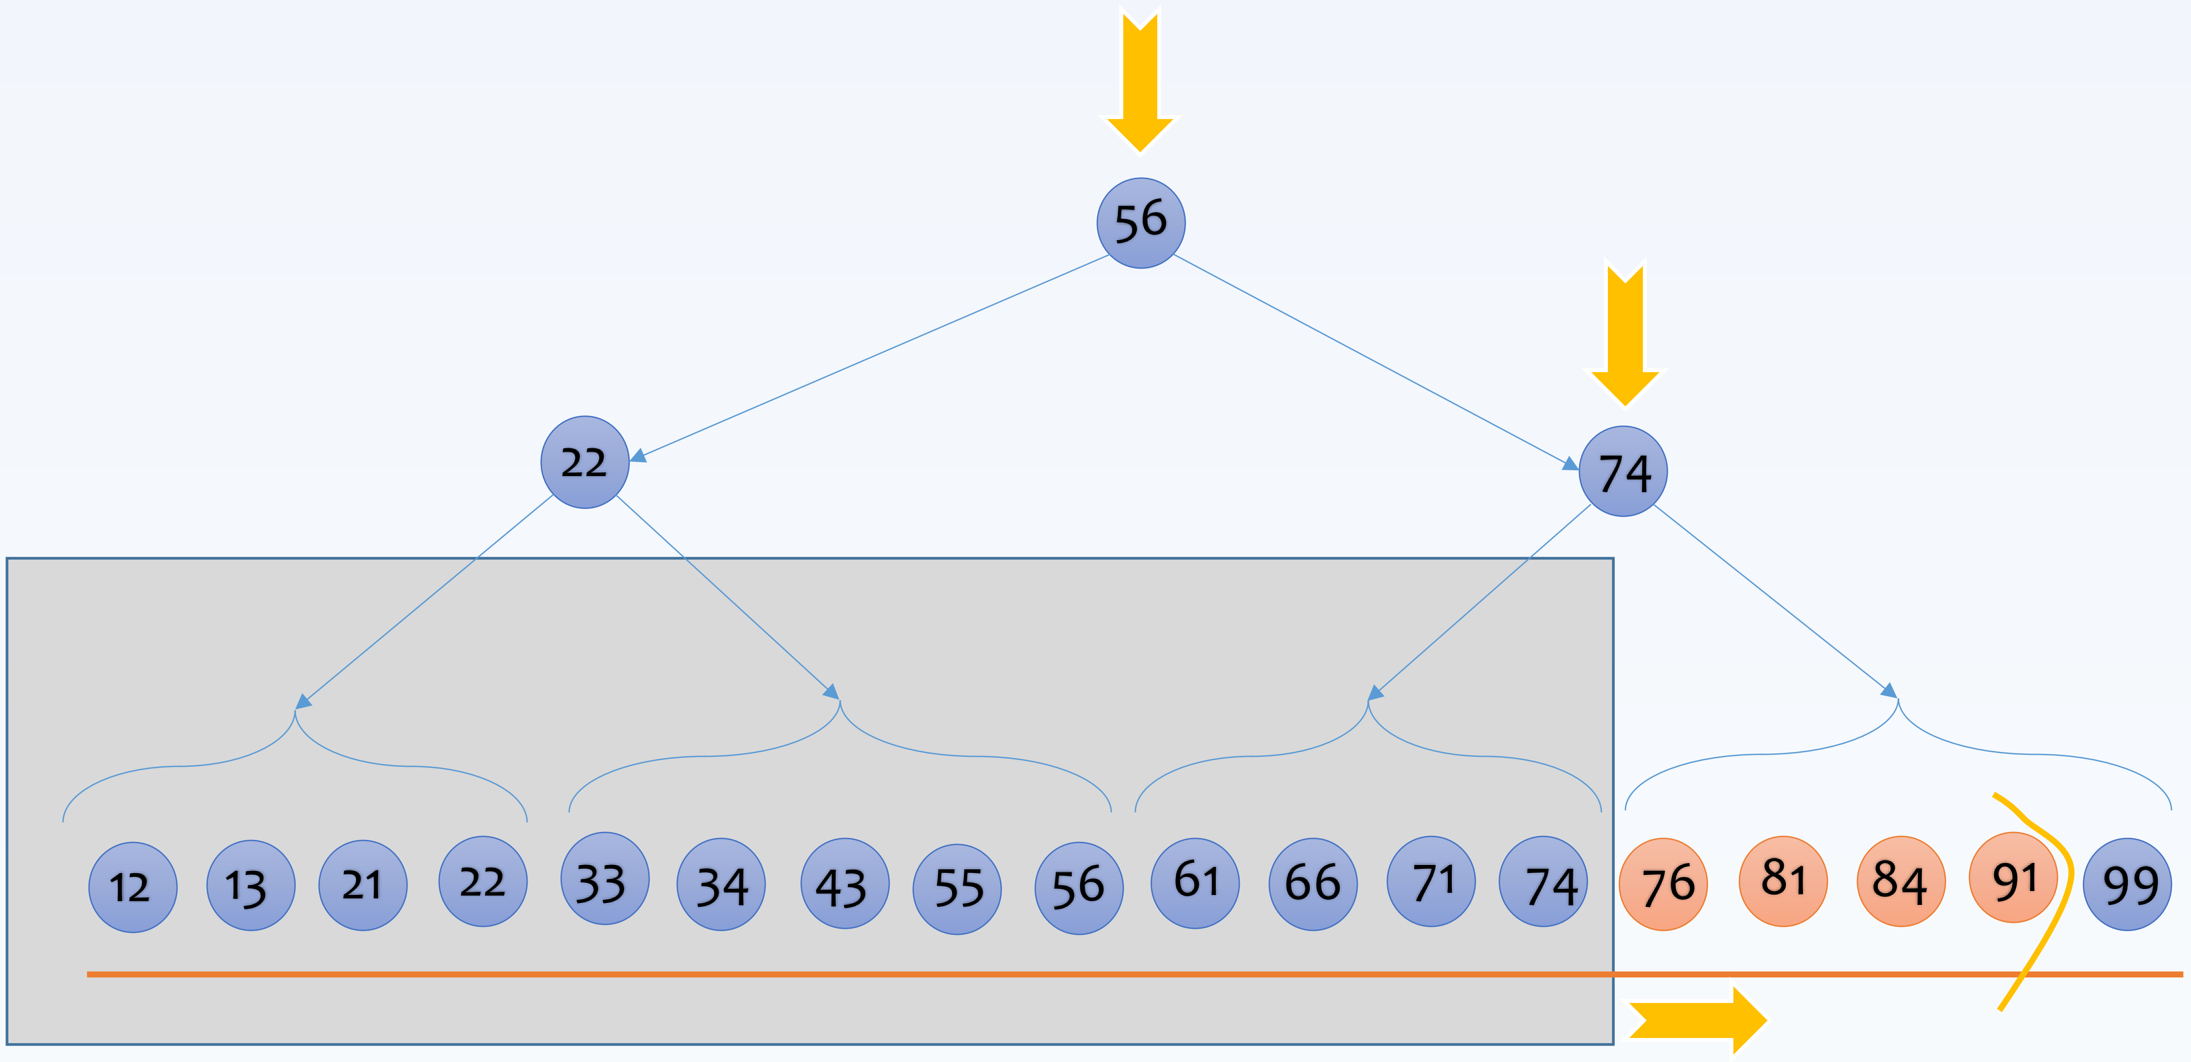

Could we spread the data cleverly to improve the search?

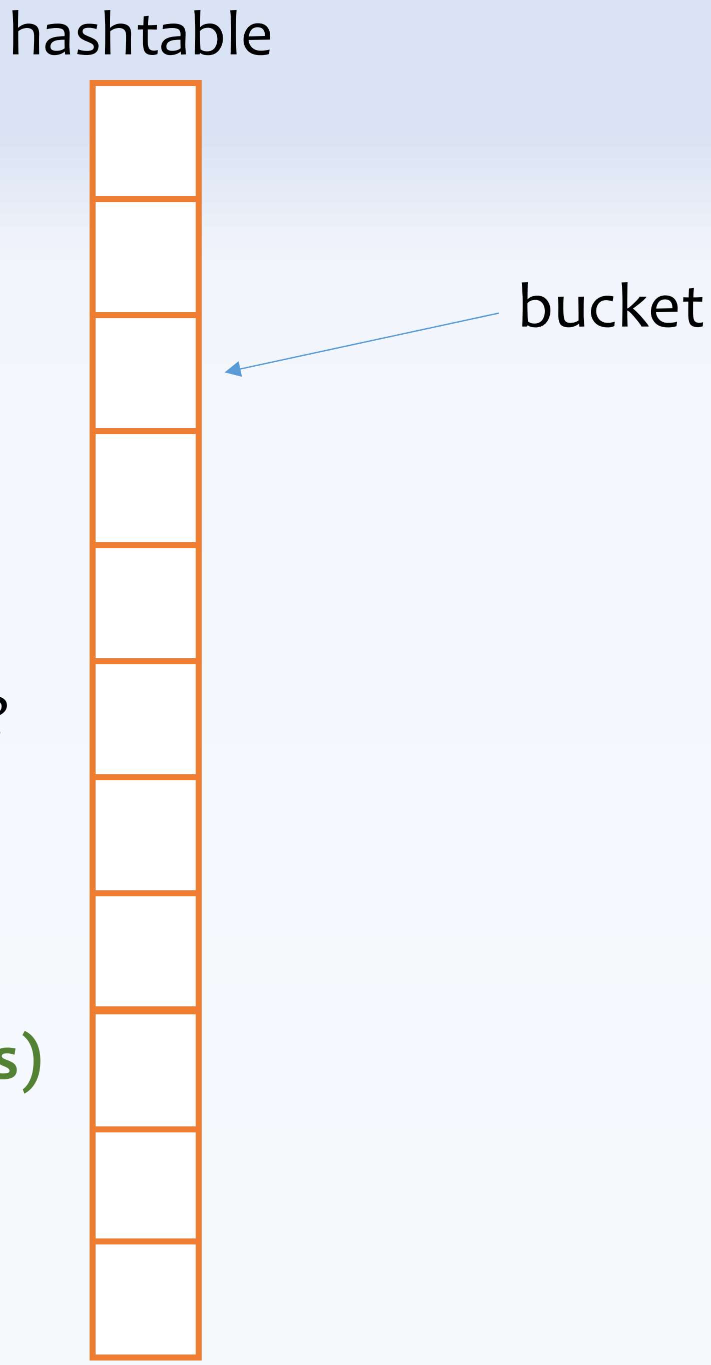

Hashing (●) = ?

(returns a value between 1 to n, where n is the number of buckets)

Inserting 81

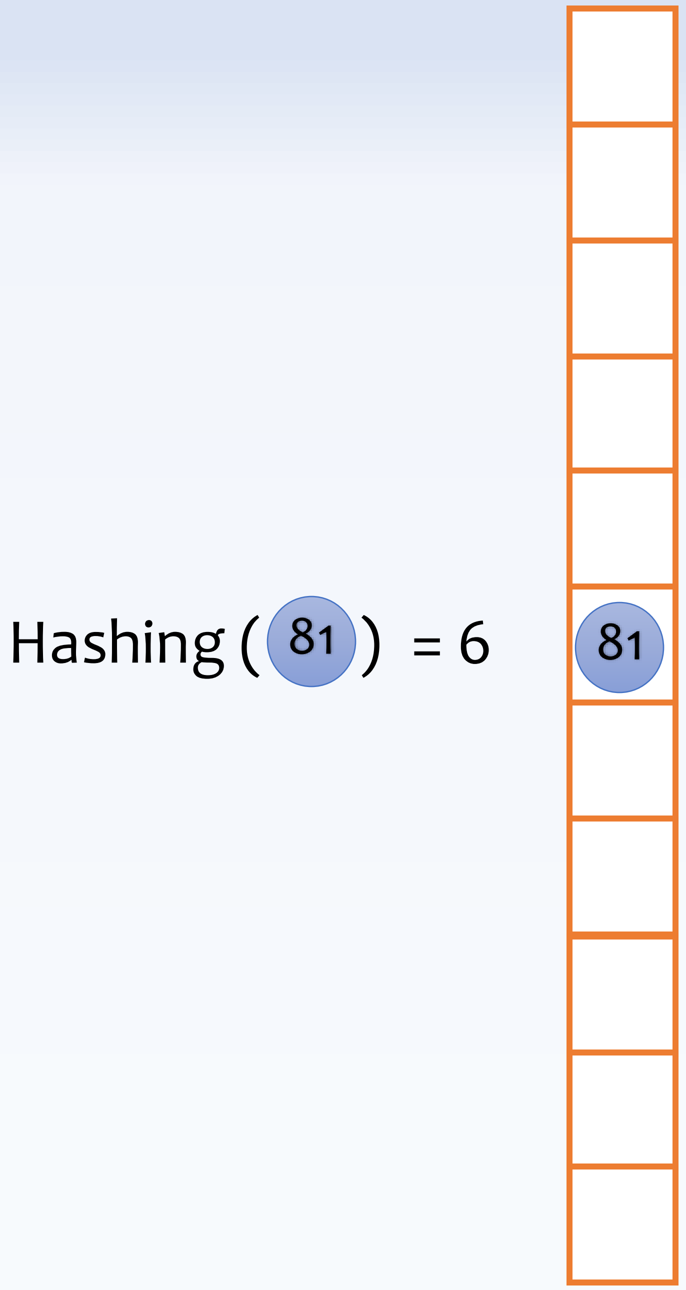

Inserting 43

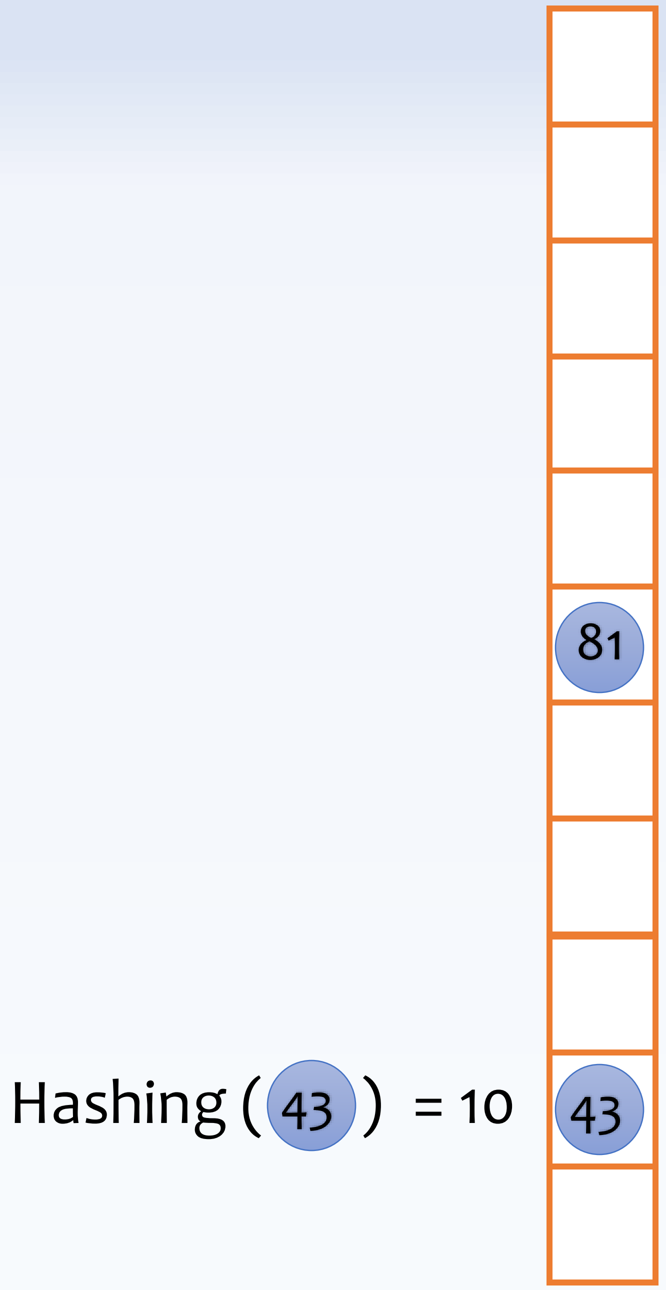

Inserting 76

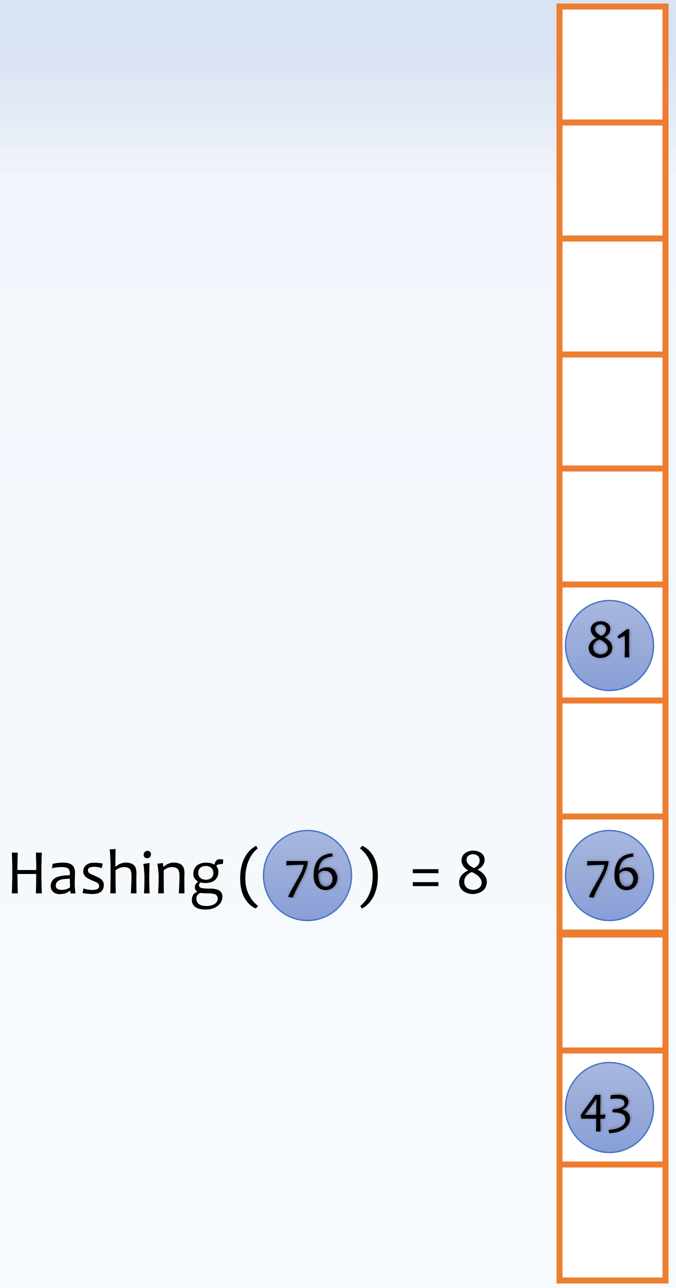

Inserting 91

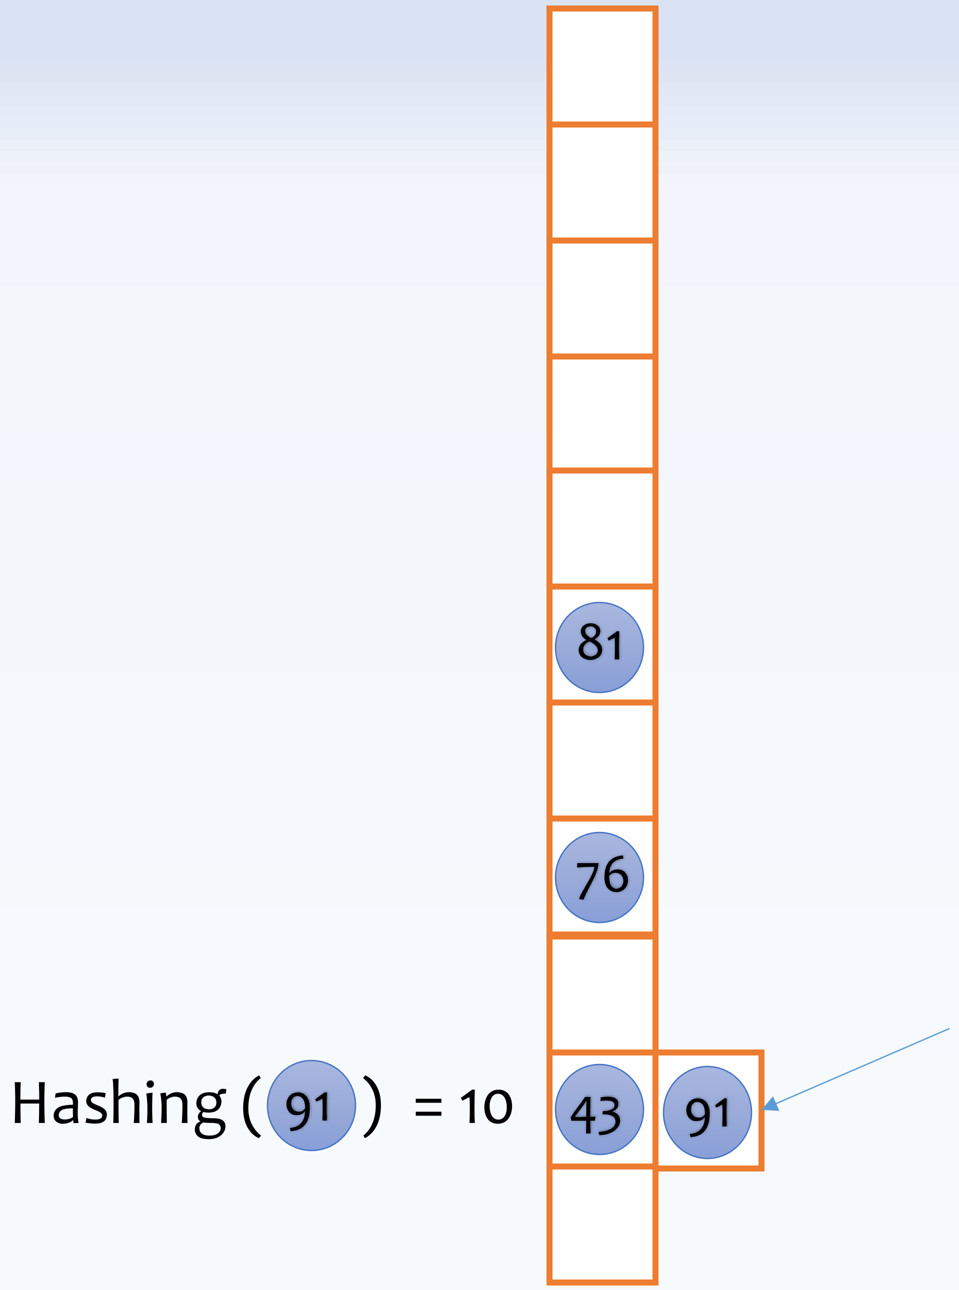

collisions
(when multiple values
hash to the same bucket)

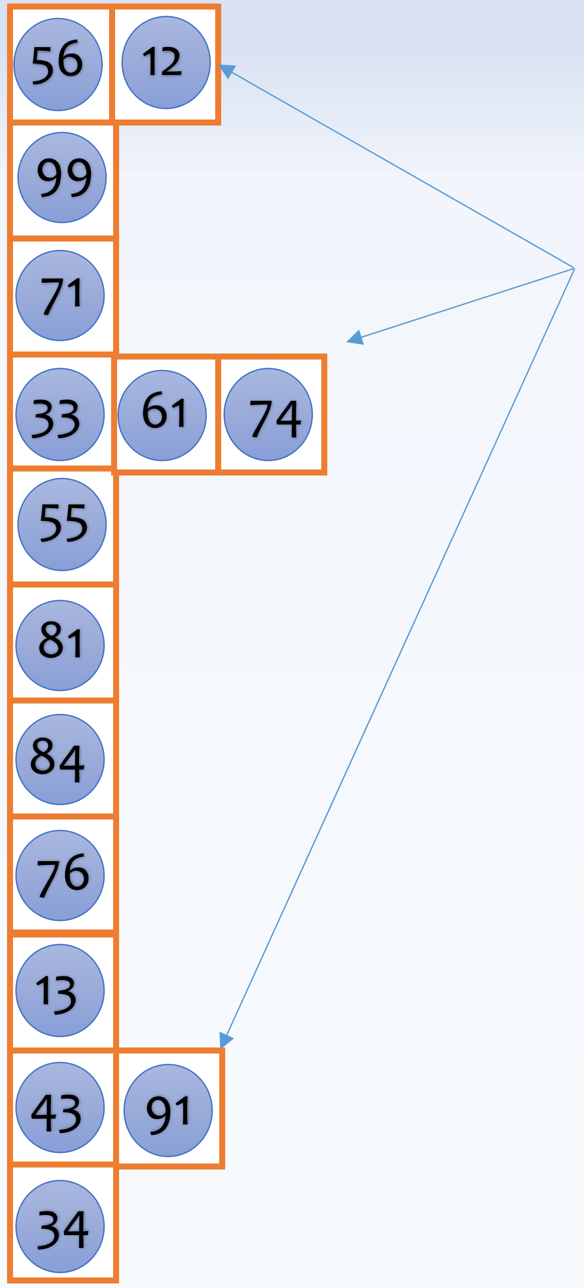

collisions
(when multiple values
hash to the same bucket)

Searching for 76

(now we can have a constant lookup cost)

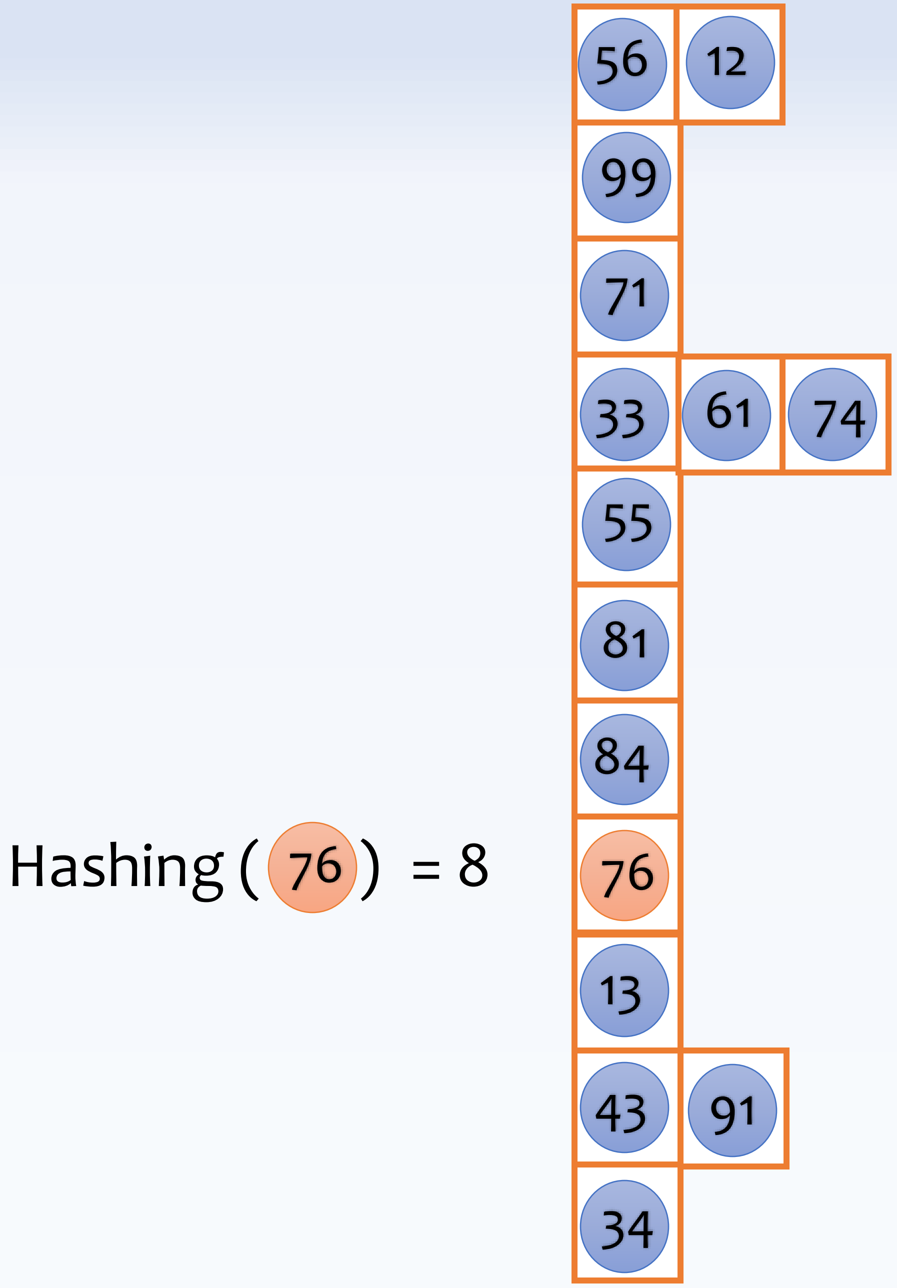

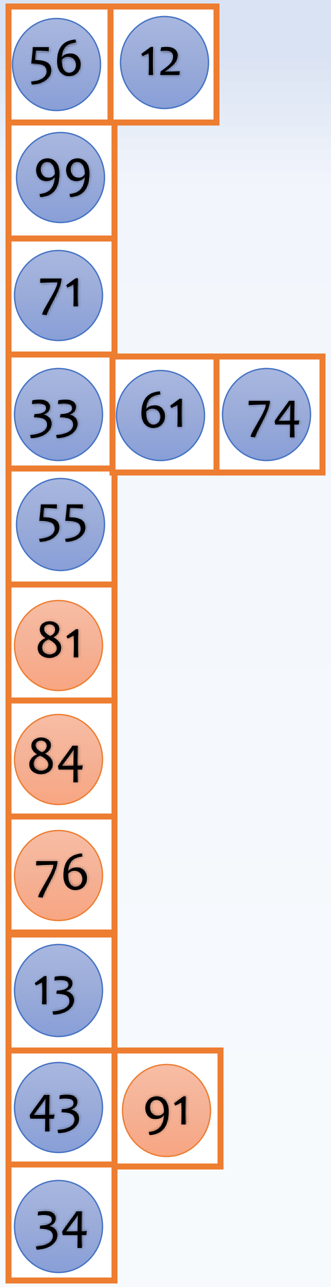

Searching for 76-91?

Could we instead search for 76, 77, 78, ..., 90, 91?

Searching for 76-91

Could we instead search for 76, 77, 78, ..., 90, 91?

Hashing (76) = 8

Hashing (77) = 1

Hashing (78) = 3

⋮

Hashing (81) = 6

⋮

Hashing (84) = 7

⋮

Hashing (90) = 8

Hashing (91) = 10

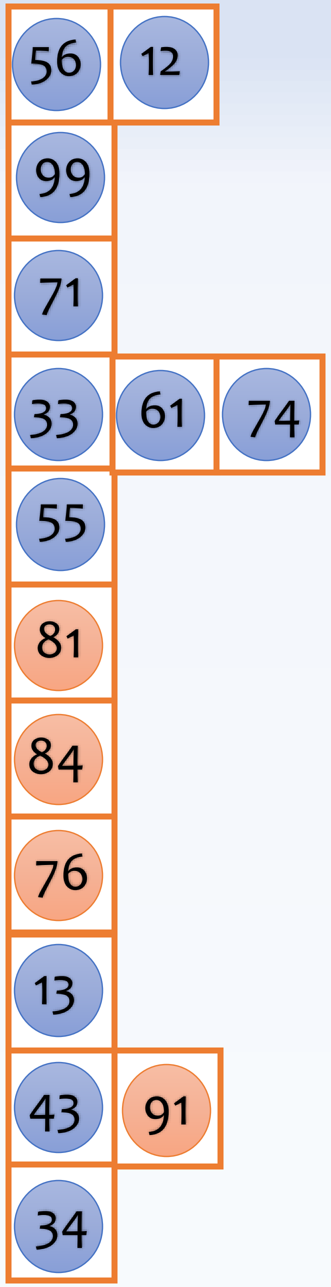

Searching for 76-91

How about 76.01, 76.02, 76.03, … ?
(simply not practical)

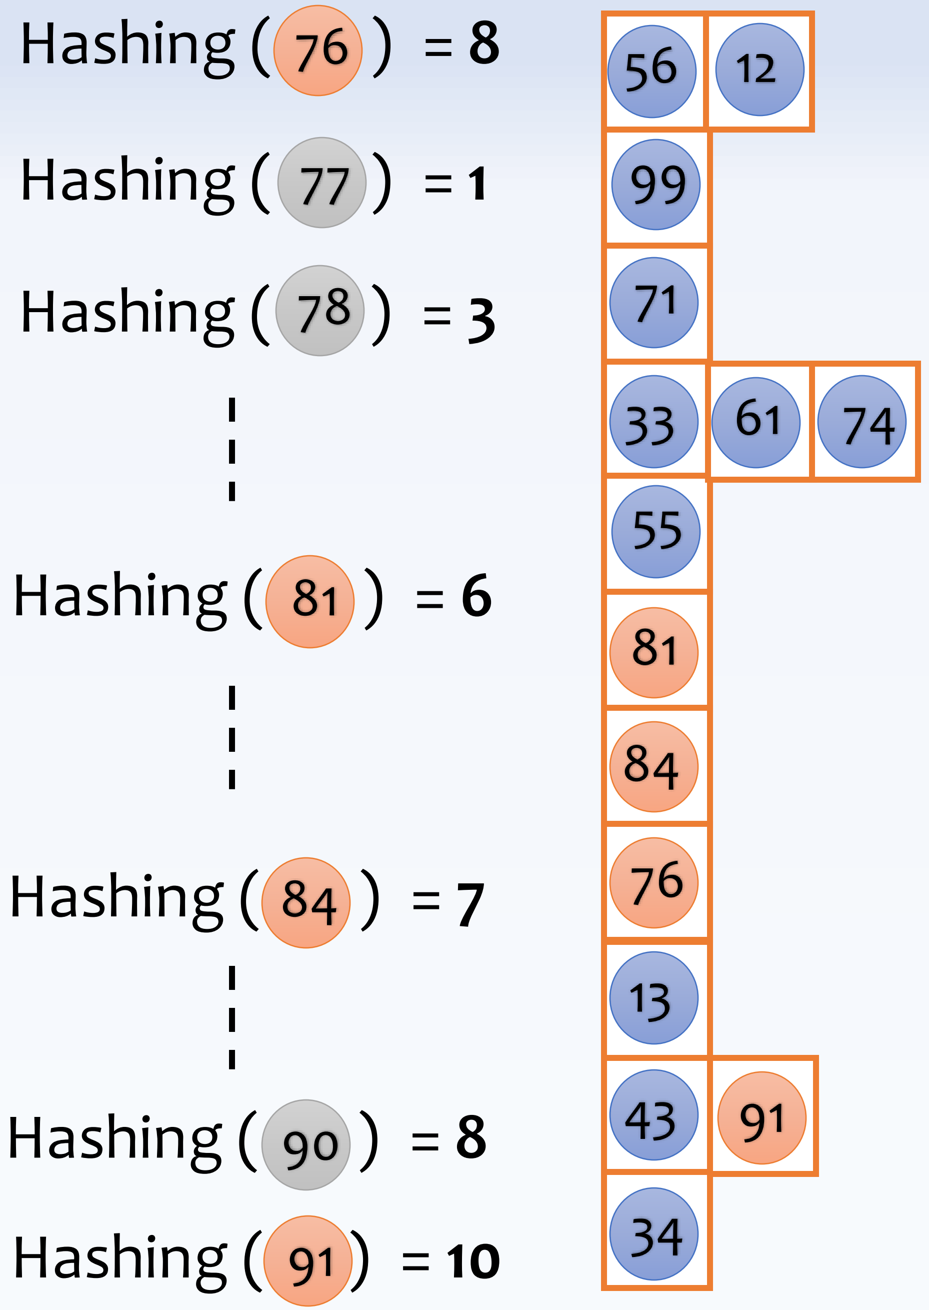

Could we imagine a new design to support searching for a range of values efficiently?

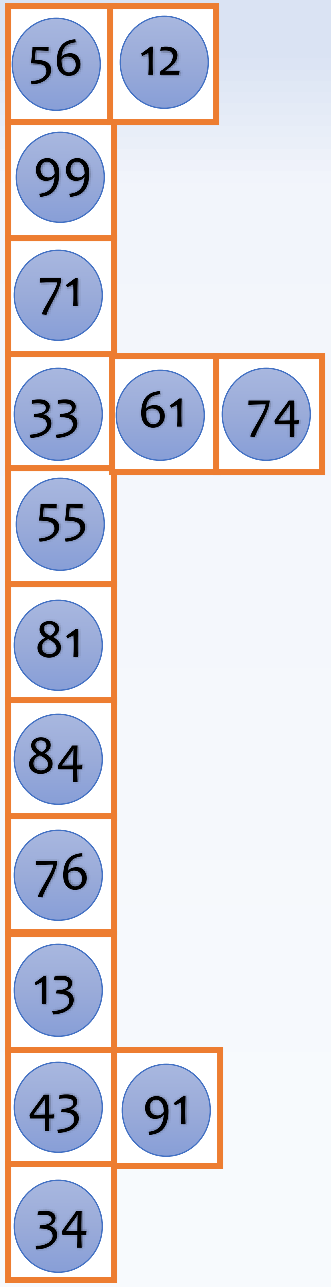

Let’s promote a subset of values as seeds

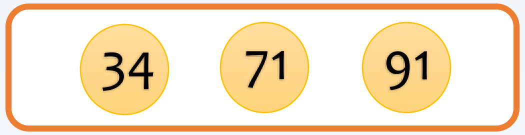

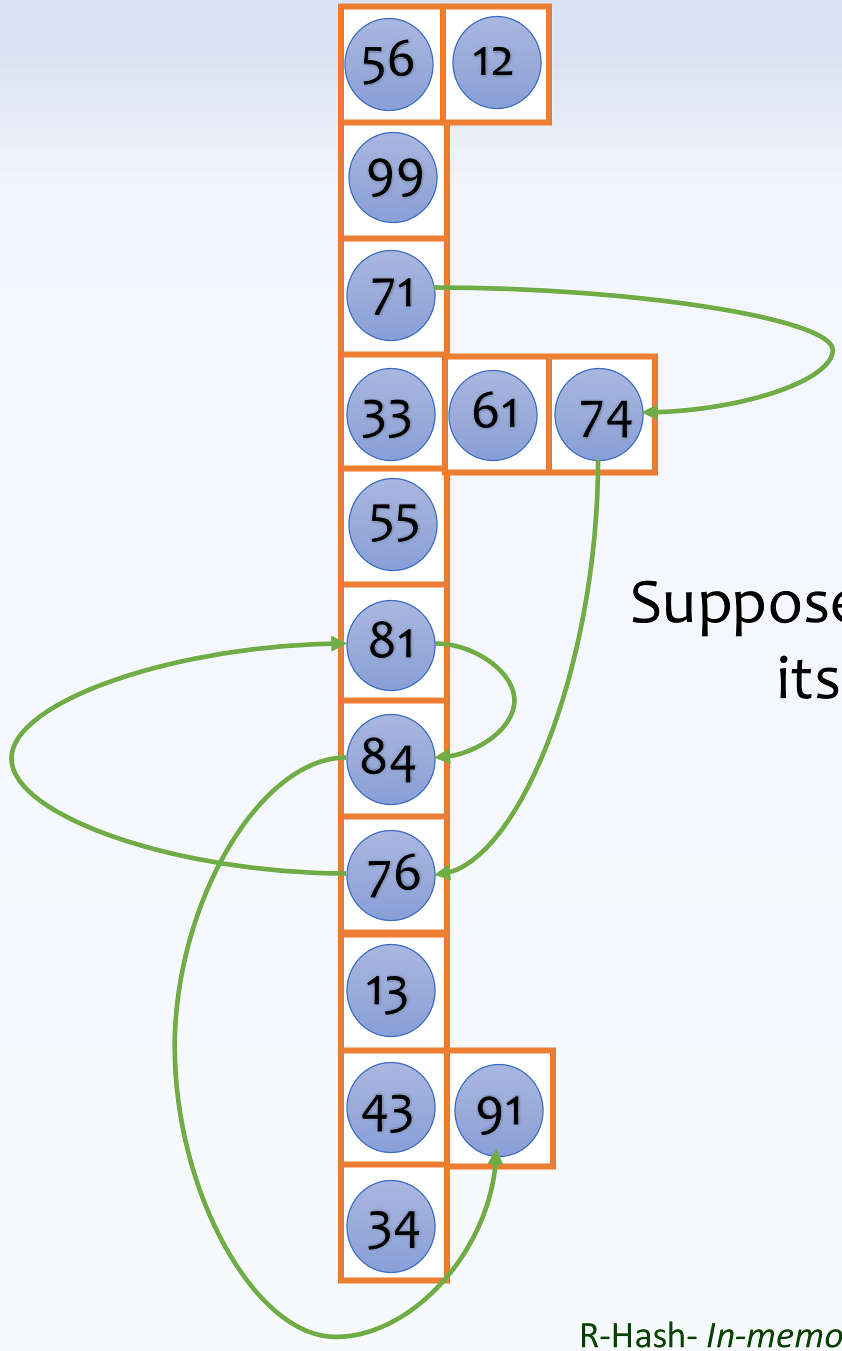

Let’s promote a subset of values as seeds

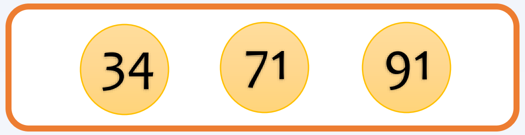

Suppose every value points to its next larger value

Searching for 76-91

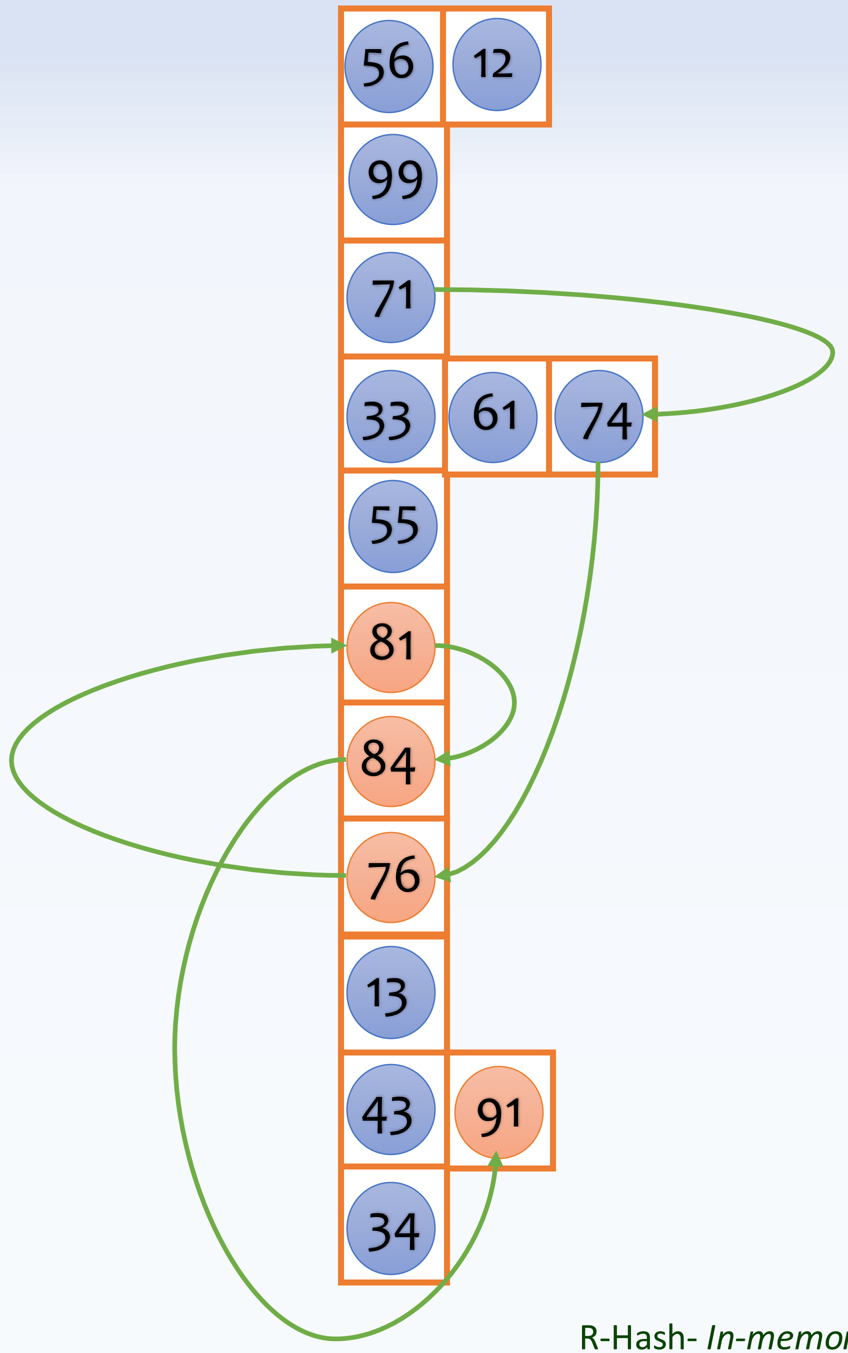

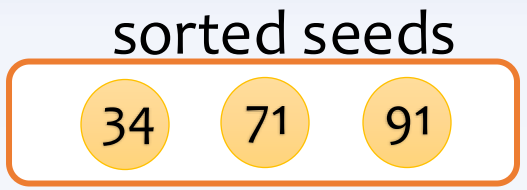

Searching for 76-91

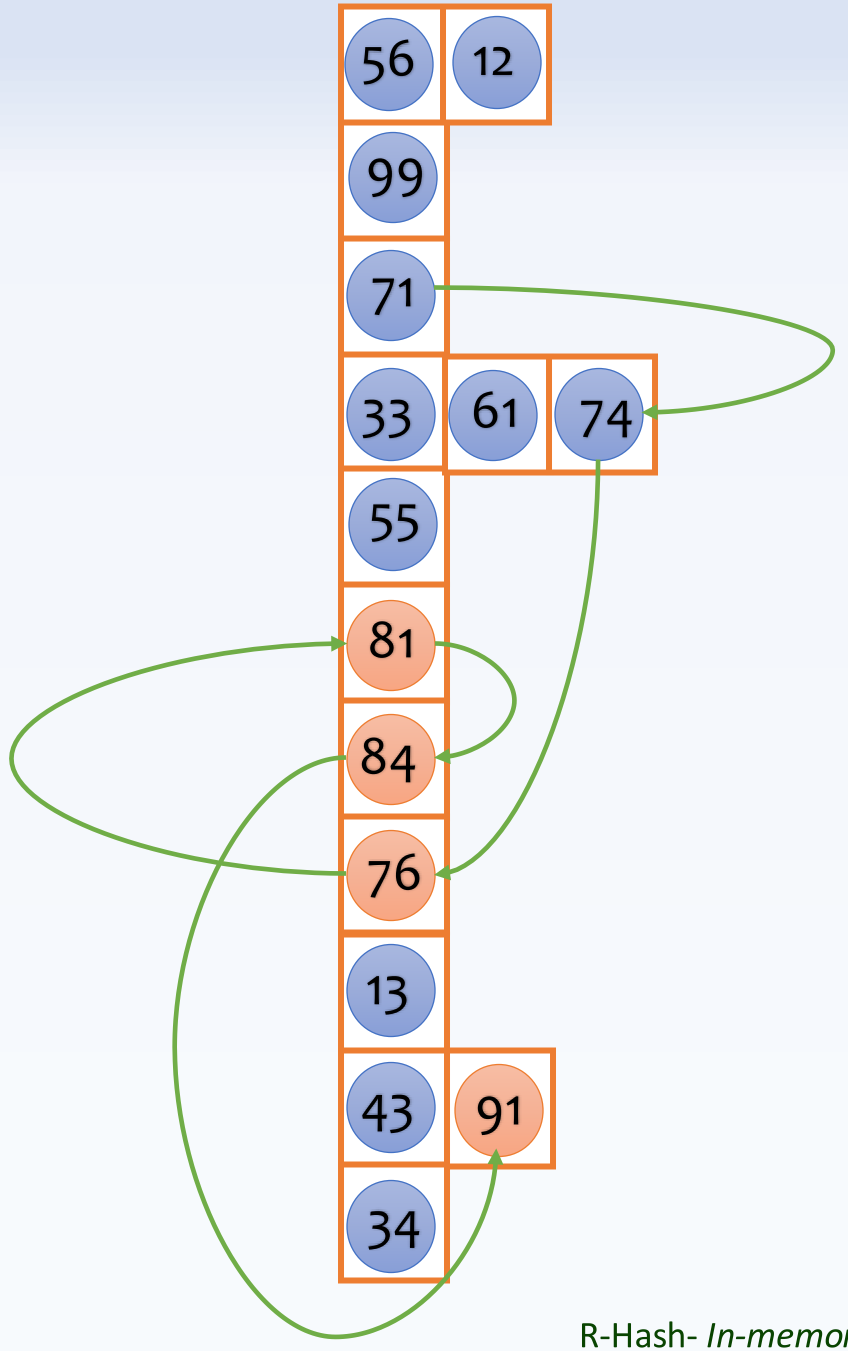

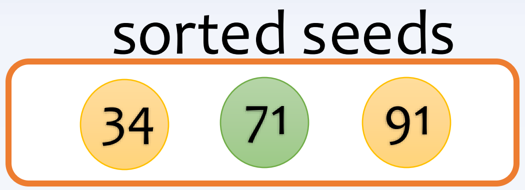

Find the largest seed smaller than 76: 71

Searching for 76-91

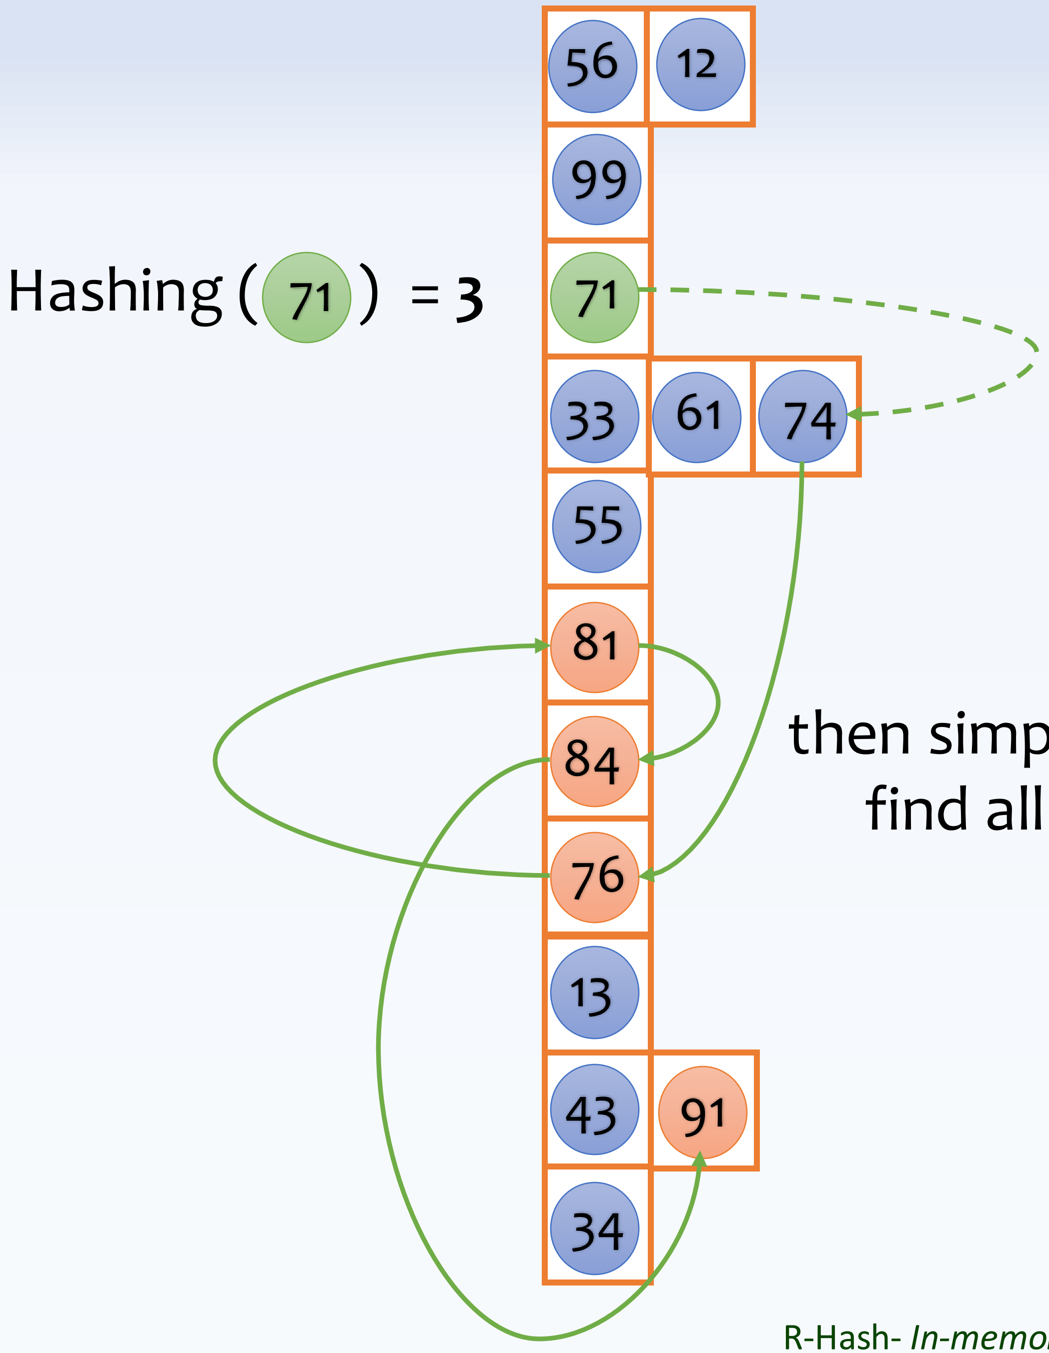

sorted seeds

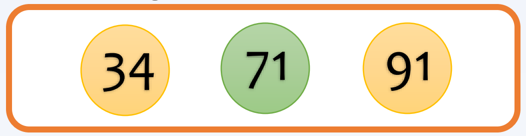

Find the largest seed smaller than 76: 71

then simply follow the pointers to find all values between 76-91

Inserting 79

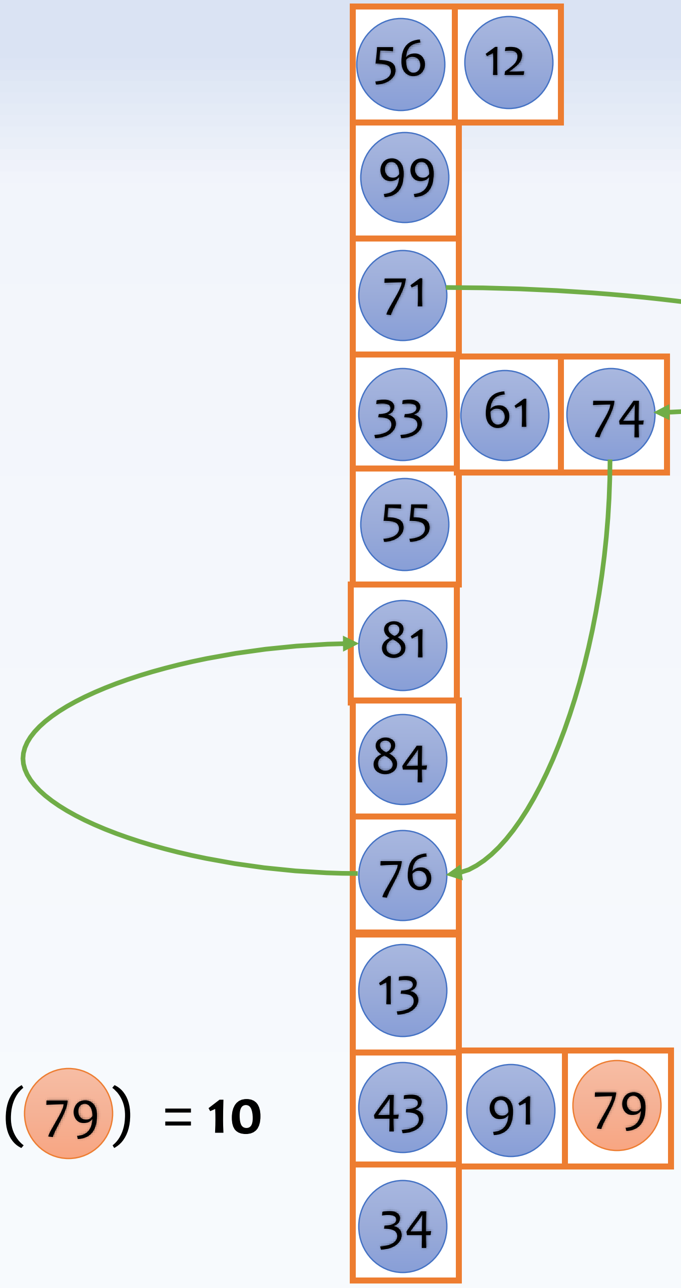

Hashing (79) = **10**

Inserting 79

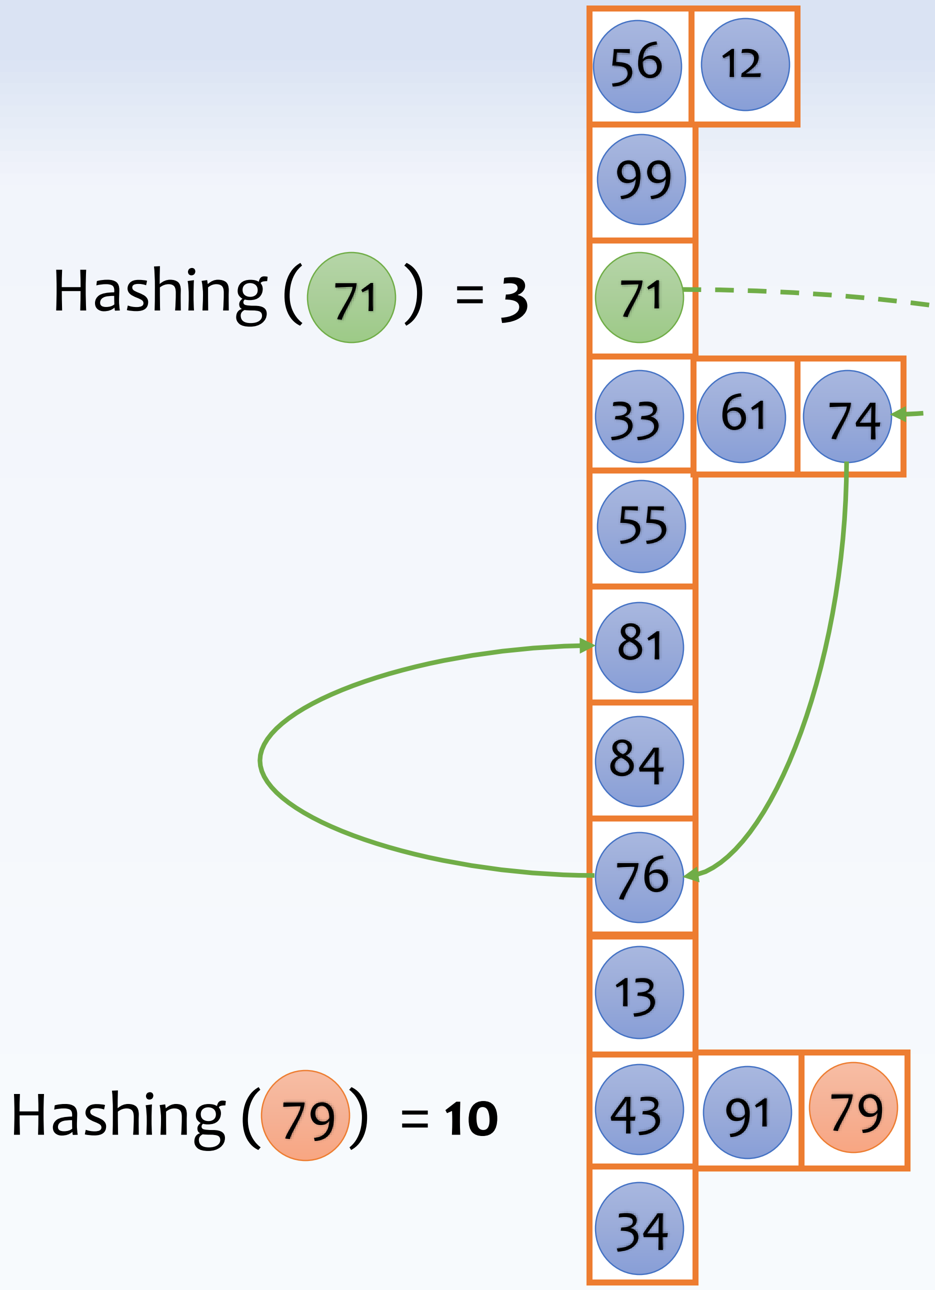

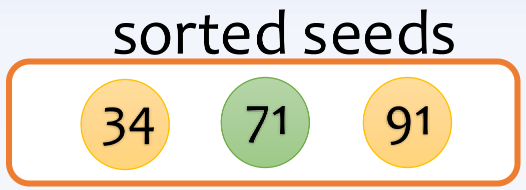

Find the largest seed smaller than 79: 71

Inserting 79

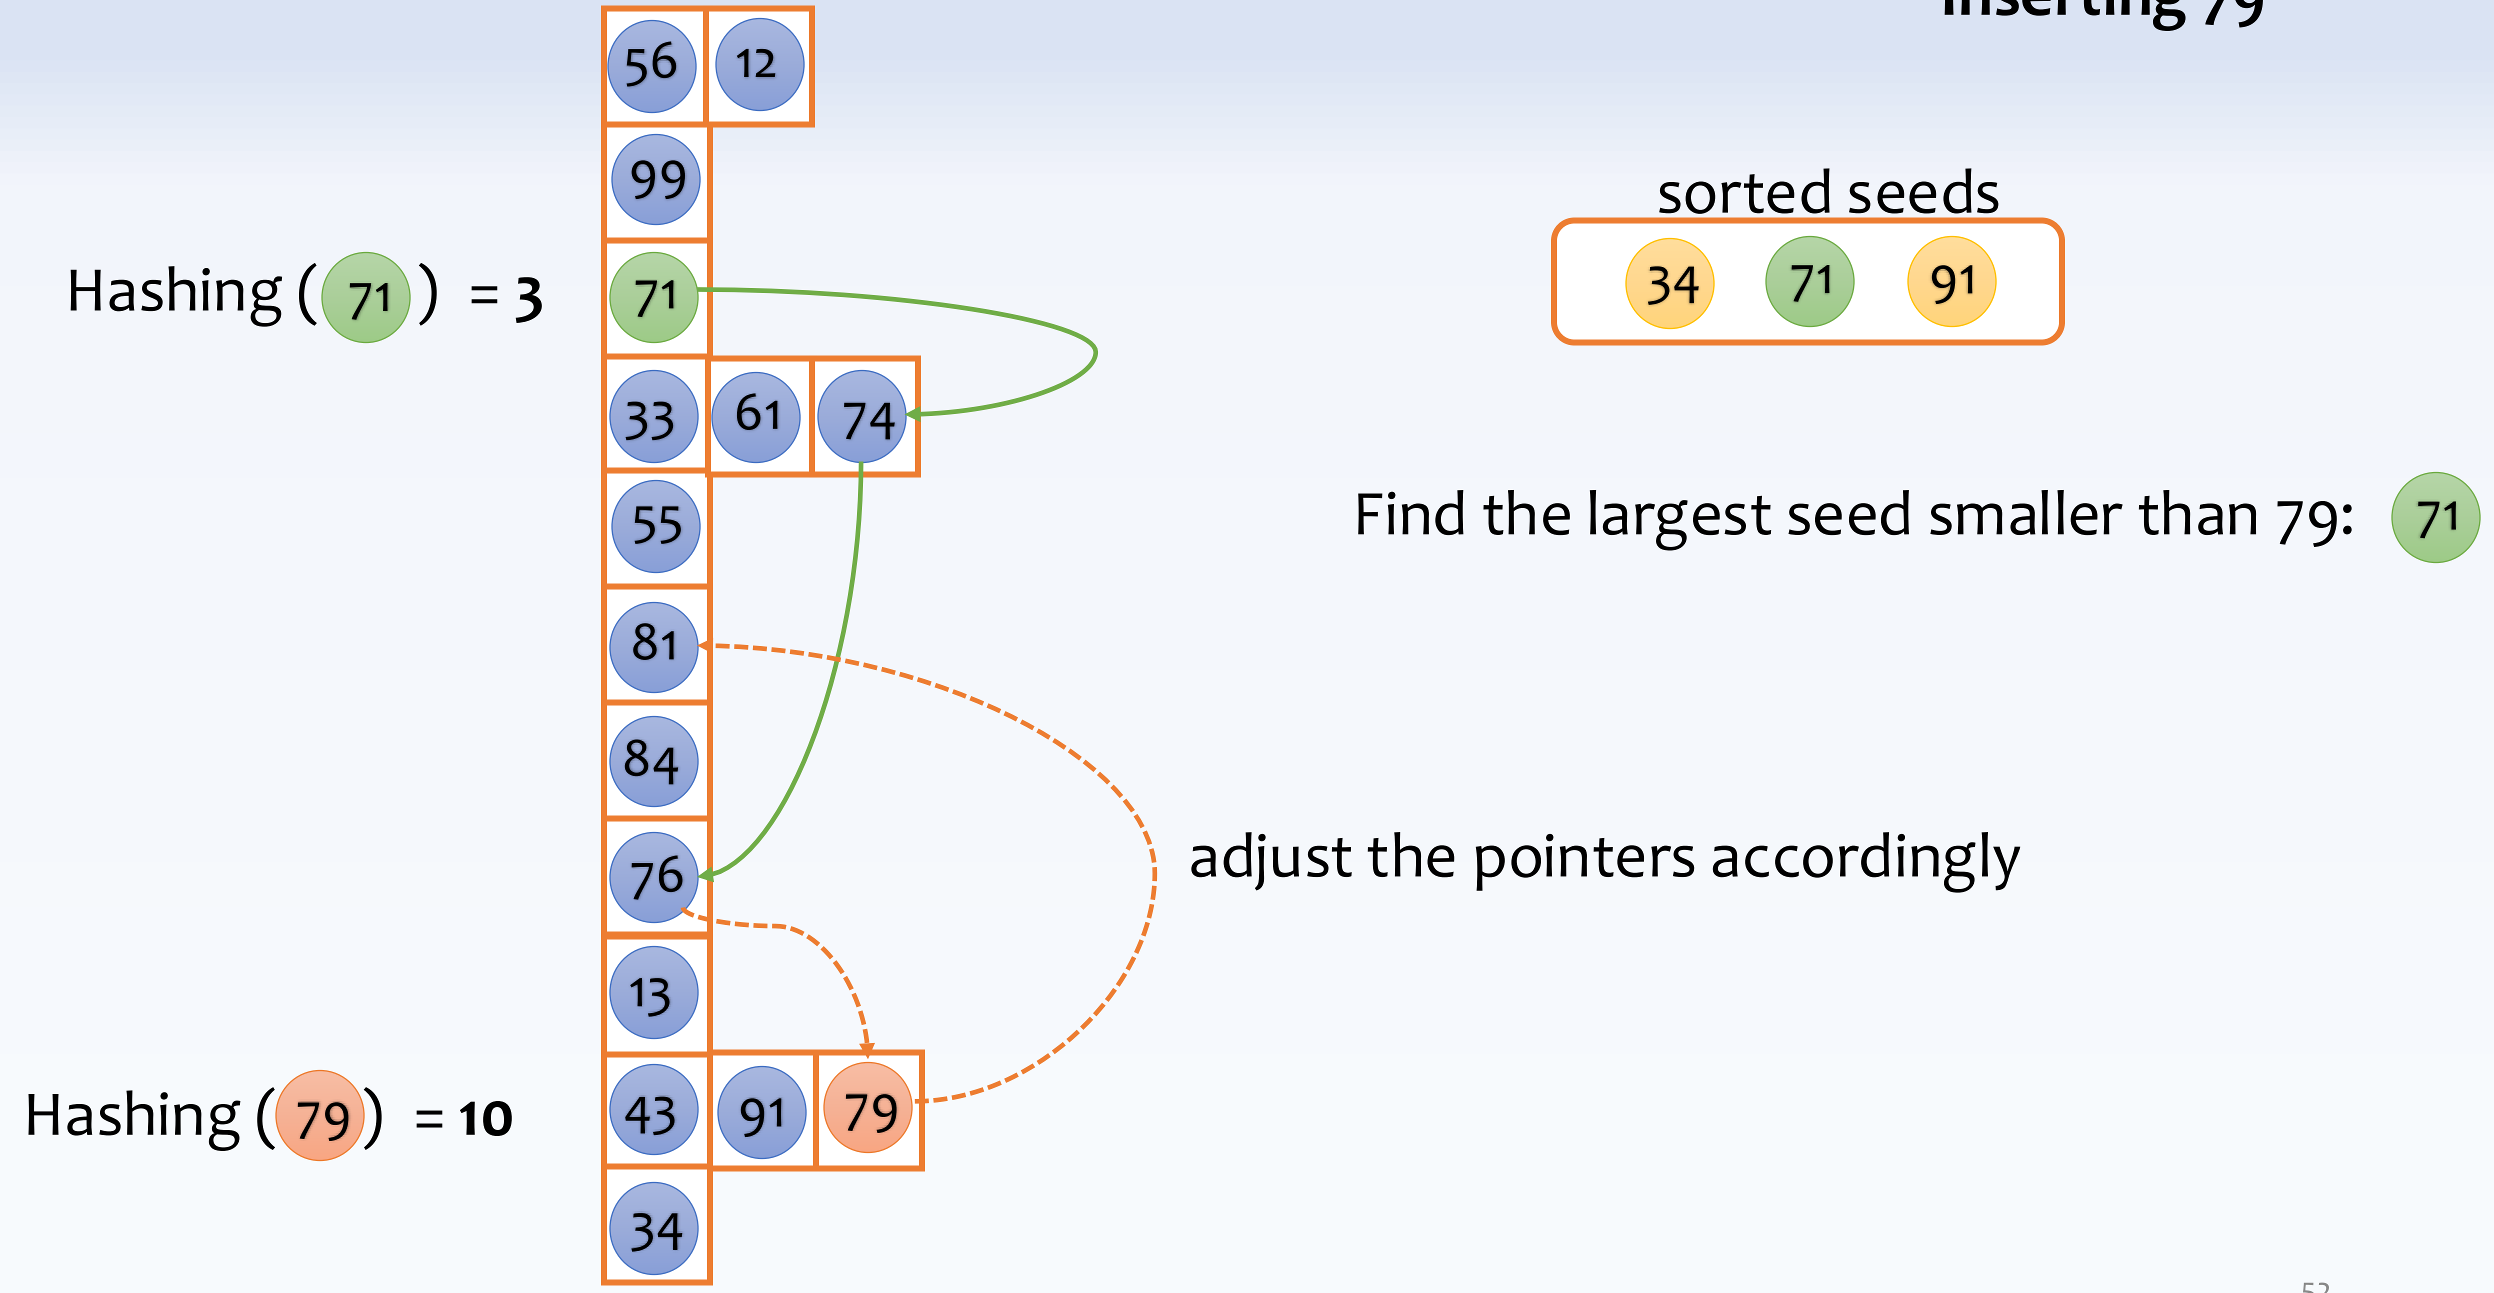

Find the largest seed smaller than 79:

adjust the pointers accordingly

Database Storage Layouts
(how likely that we need an index for range queries?)

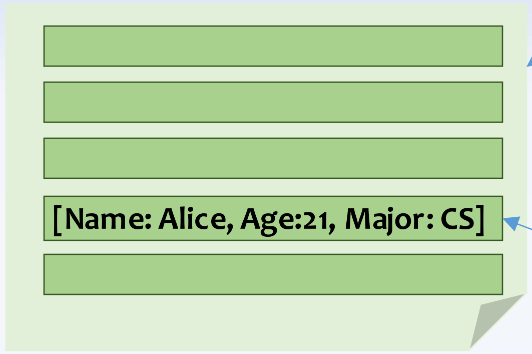

database pages
(containing a set of records)

a database record, e.g.,
[Name: Alice, Age:21, Major: CS]

Row-based Layout

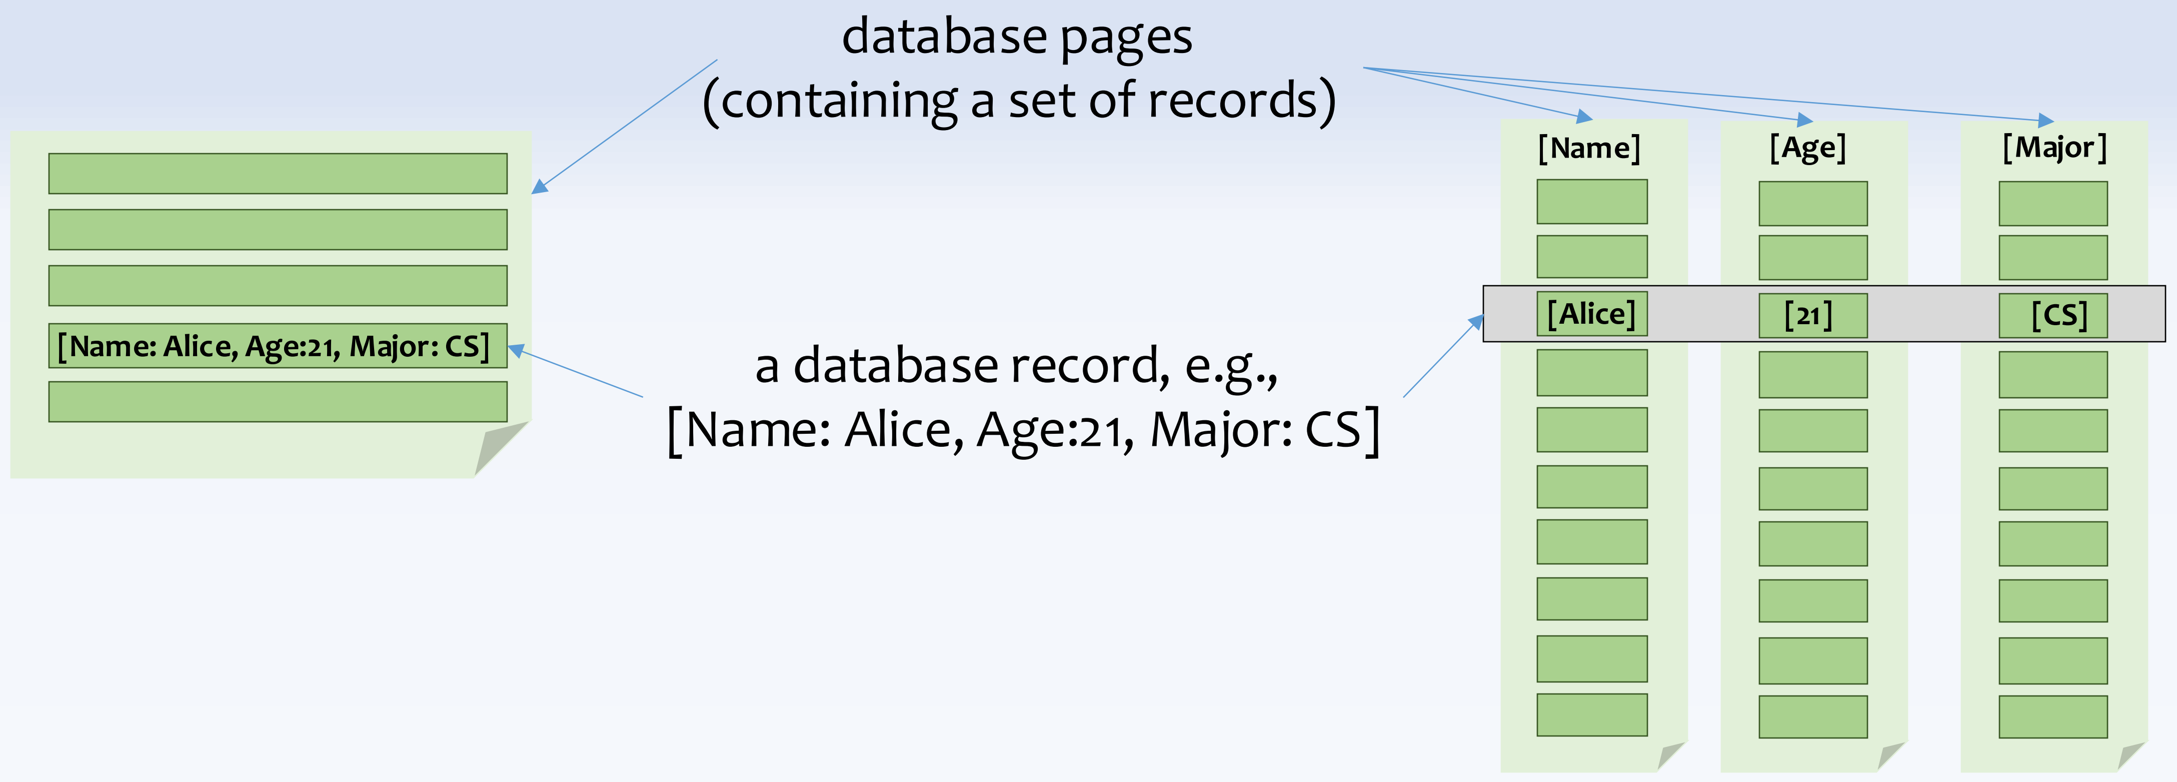
database pages
(containing a set of records)
[Name: Alice, Age:21, Major: CS]
a database record, e.g.,
[Name: Alice, Age:21, Major: CS]
[Name]
[Age]
[Major]
[Alice]
[21]
[CS]

Row-based Layout

Column-based Layout

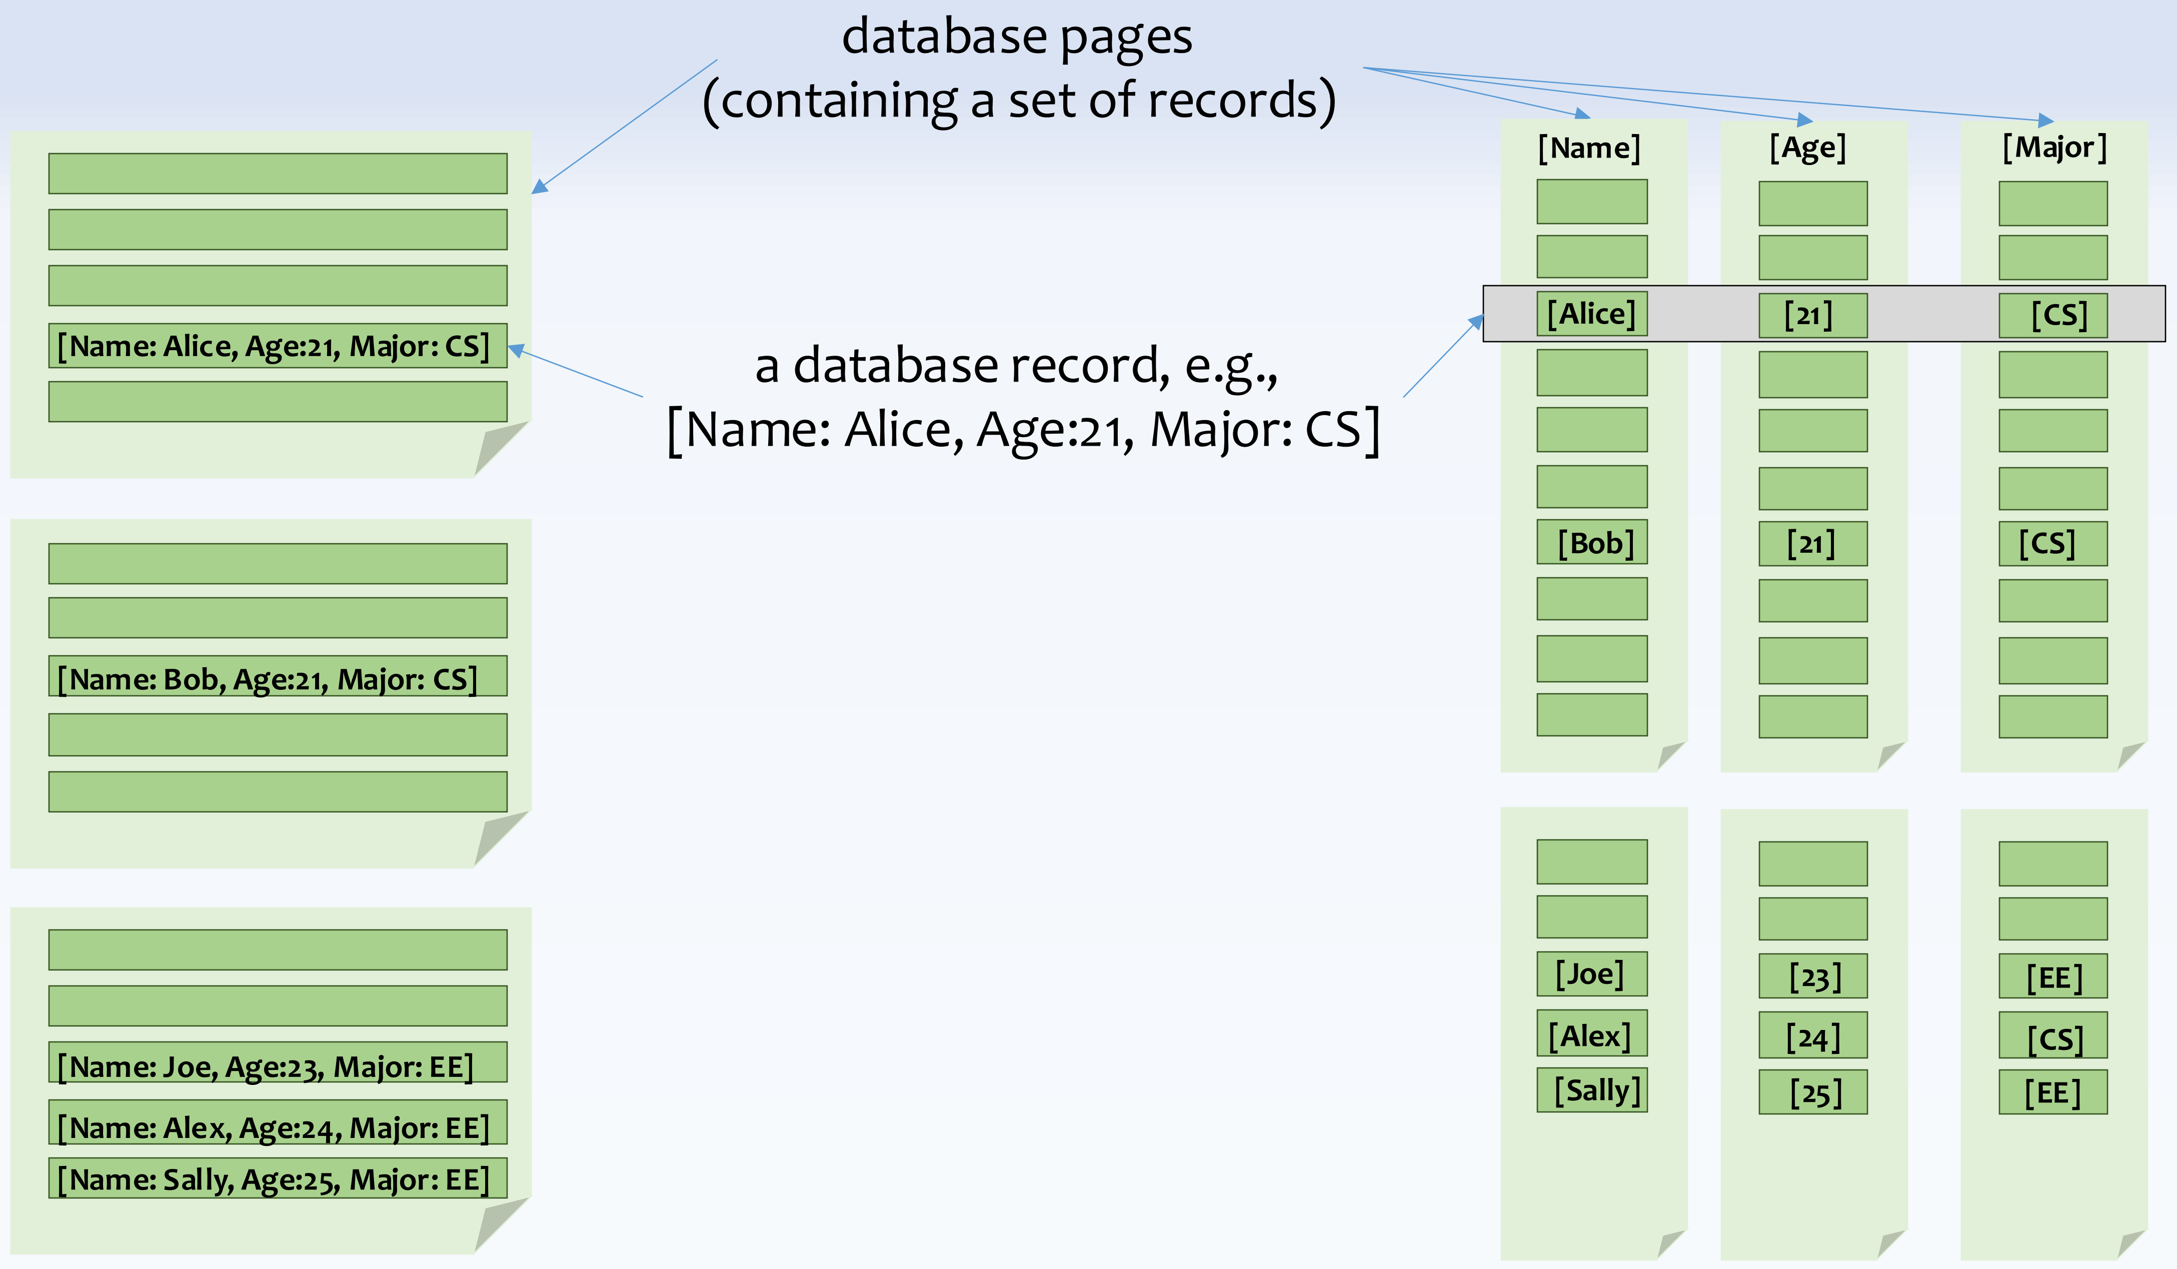
database pages
(containing a set of records)
a database record, e.g.,
[Name: Alice, Age:21, Major: CS]
[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
[Name]
[Age]
[Major]
[Alice]
[21]
[CS]
[Bob]
[21]
[CS]
[Joe]
[Alex]
[Sally]
[23]
[24]
[25]
[EE]
[CS]
[EE]

Row-based Layout

Column-based Layout

Searching for all students between the age of 21 to 24 (may return many students)

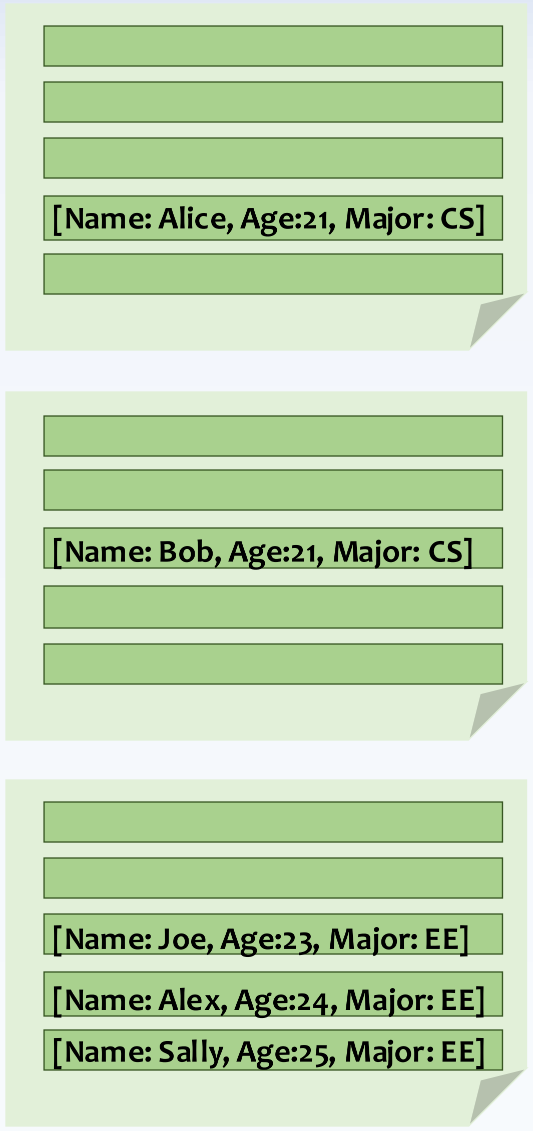
[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]

Row-based Layout

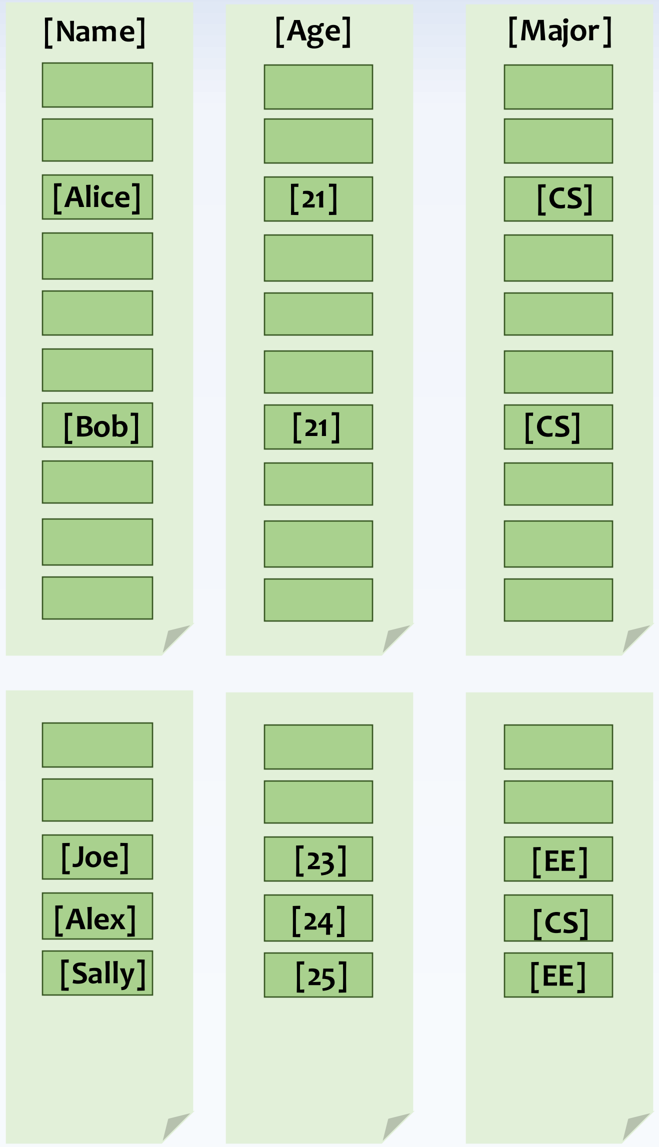

Column-based Layout

Searching for all students between the age of 21 to 24 (may return many students)

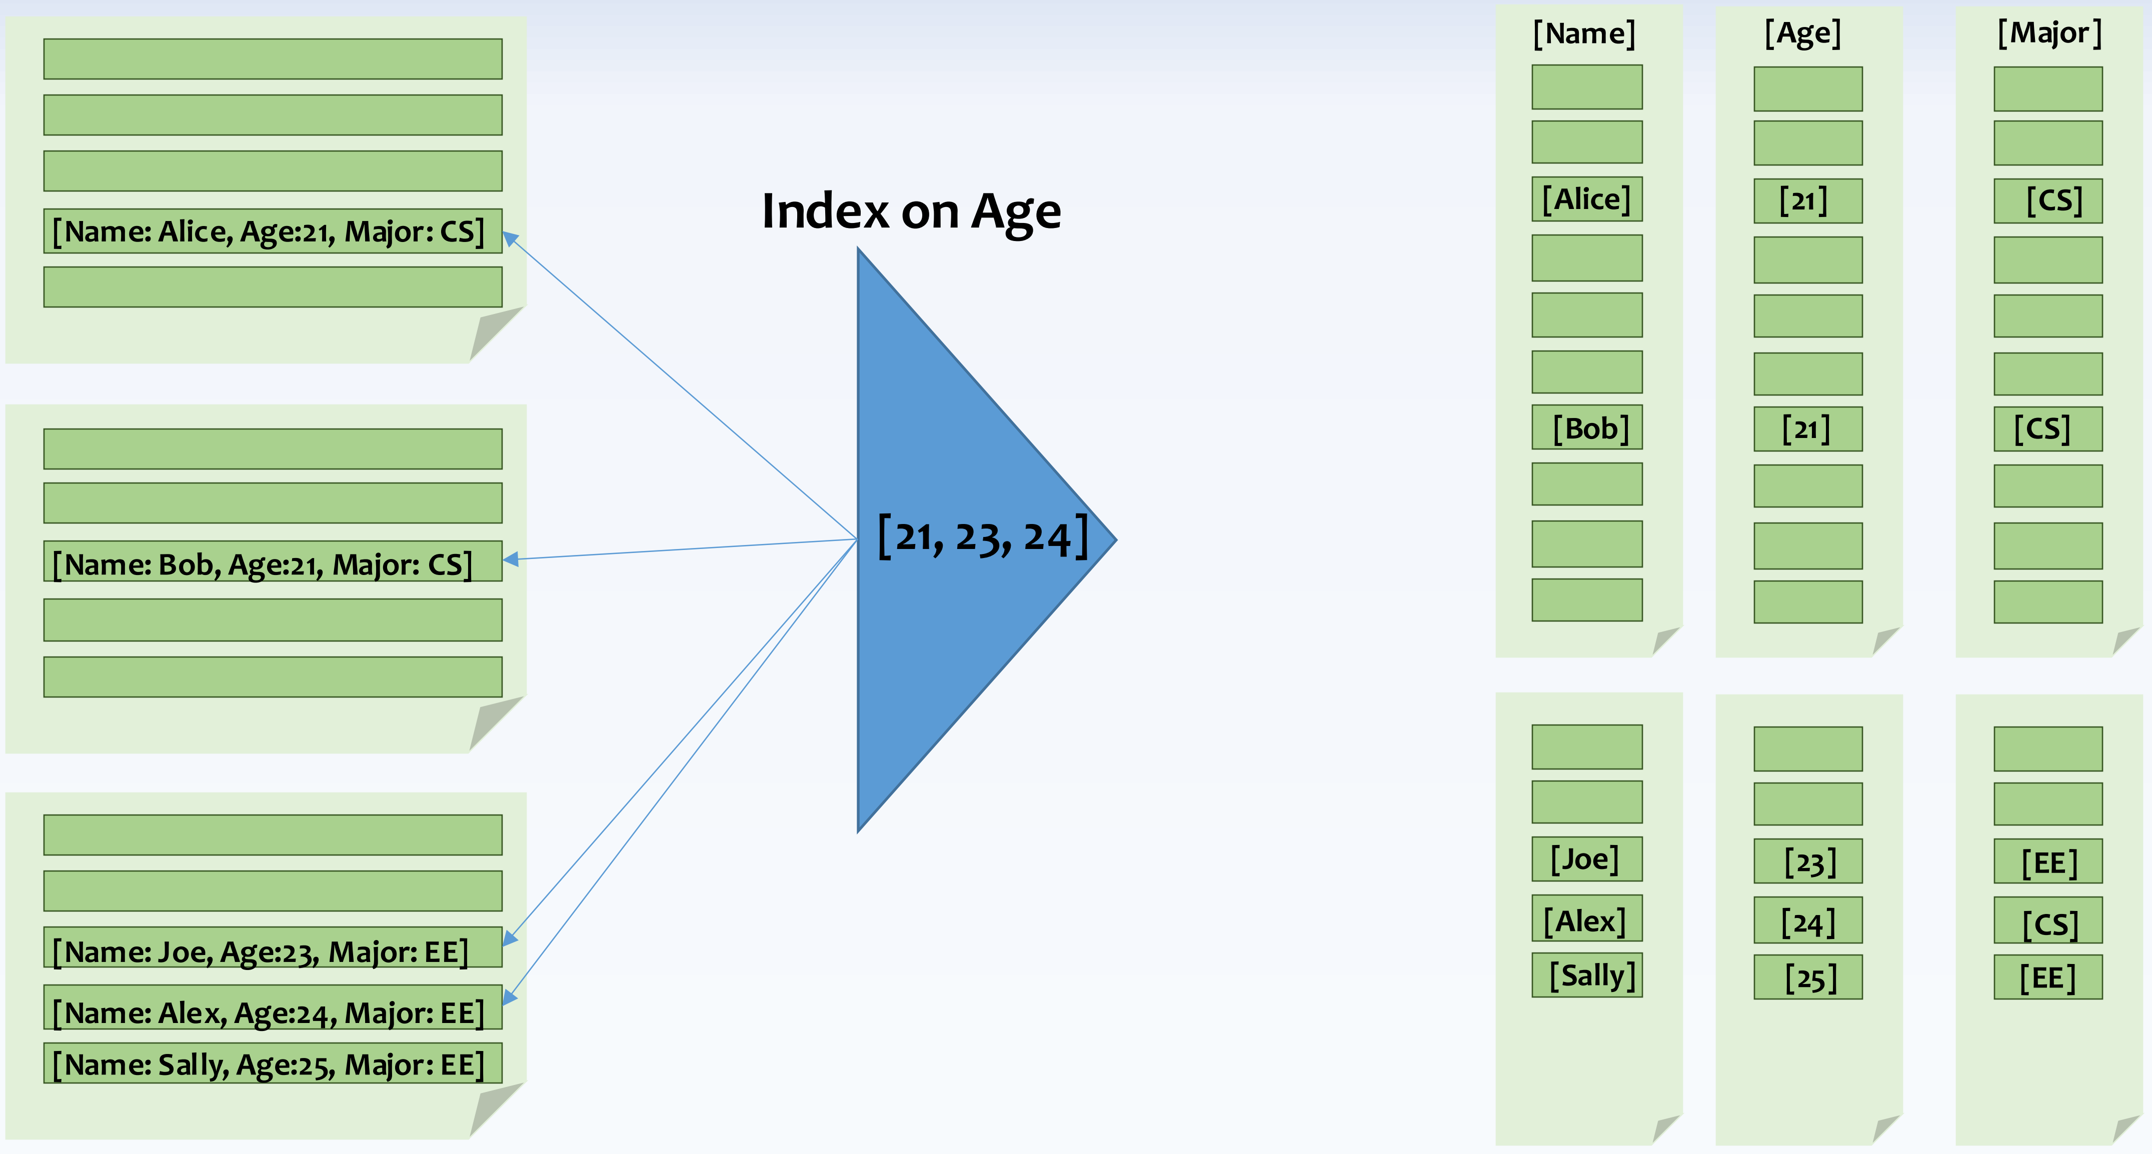

[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
Index on Age
[21, 23, 24]
[Name]
[Alice]
[Bob]
[Joe]
[Alex]
[Sally]
[Age]
[21]
[21]
[23]
[24]
[25]
[Major]
[CS]
[CS]
[EE]
[CS]
[EE]

Row-based Layout

Column-based Layout

Searching for all students between the age of 21 to 24 (may return many students)

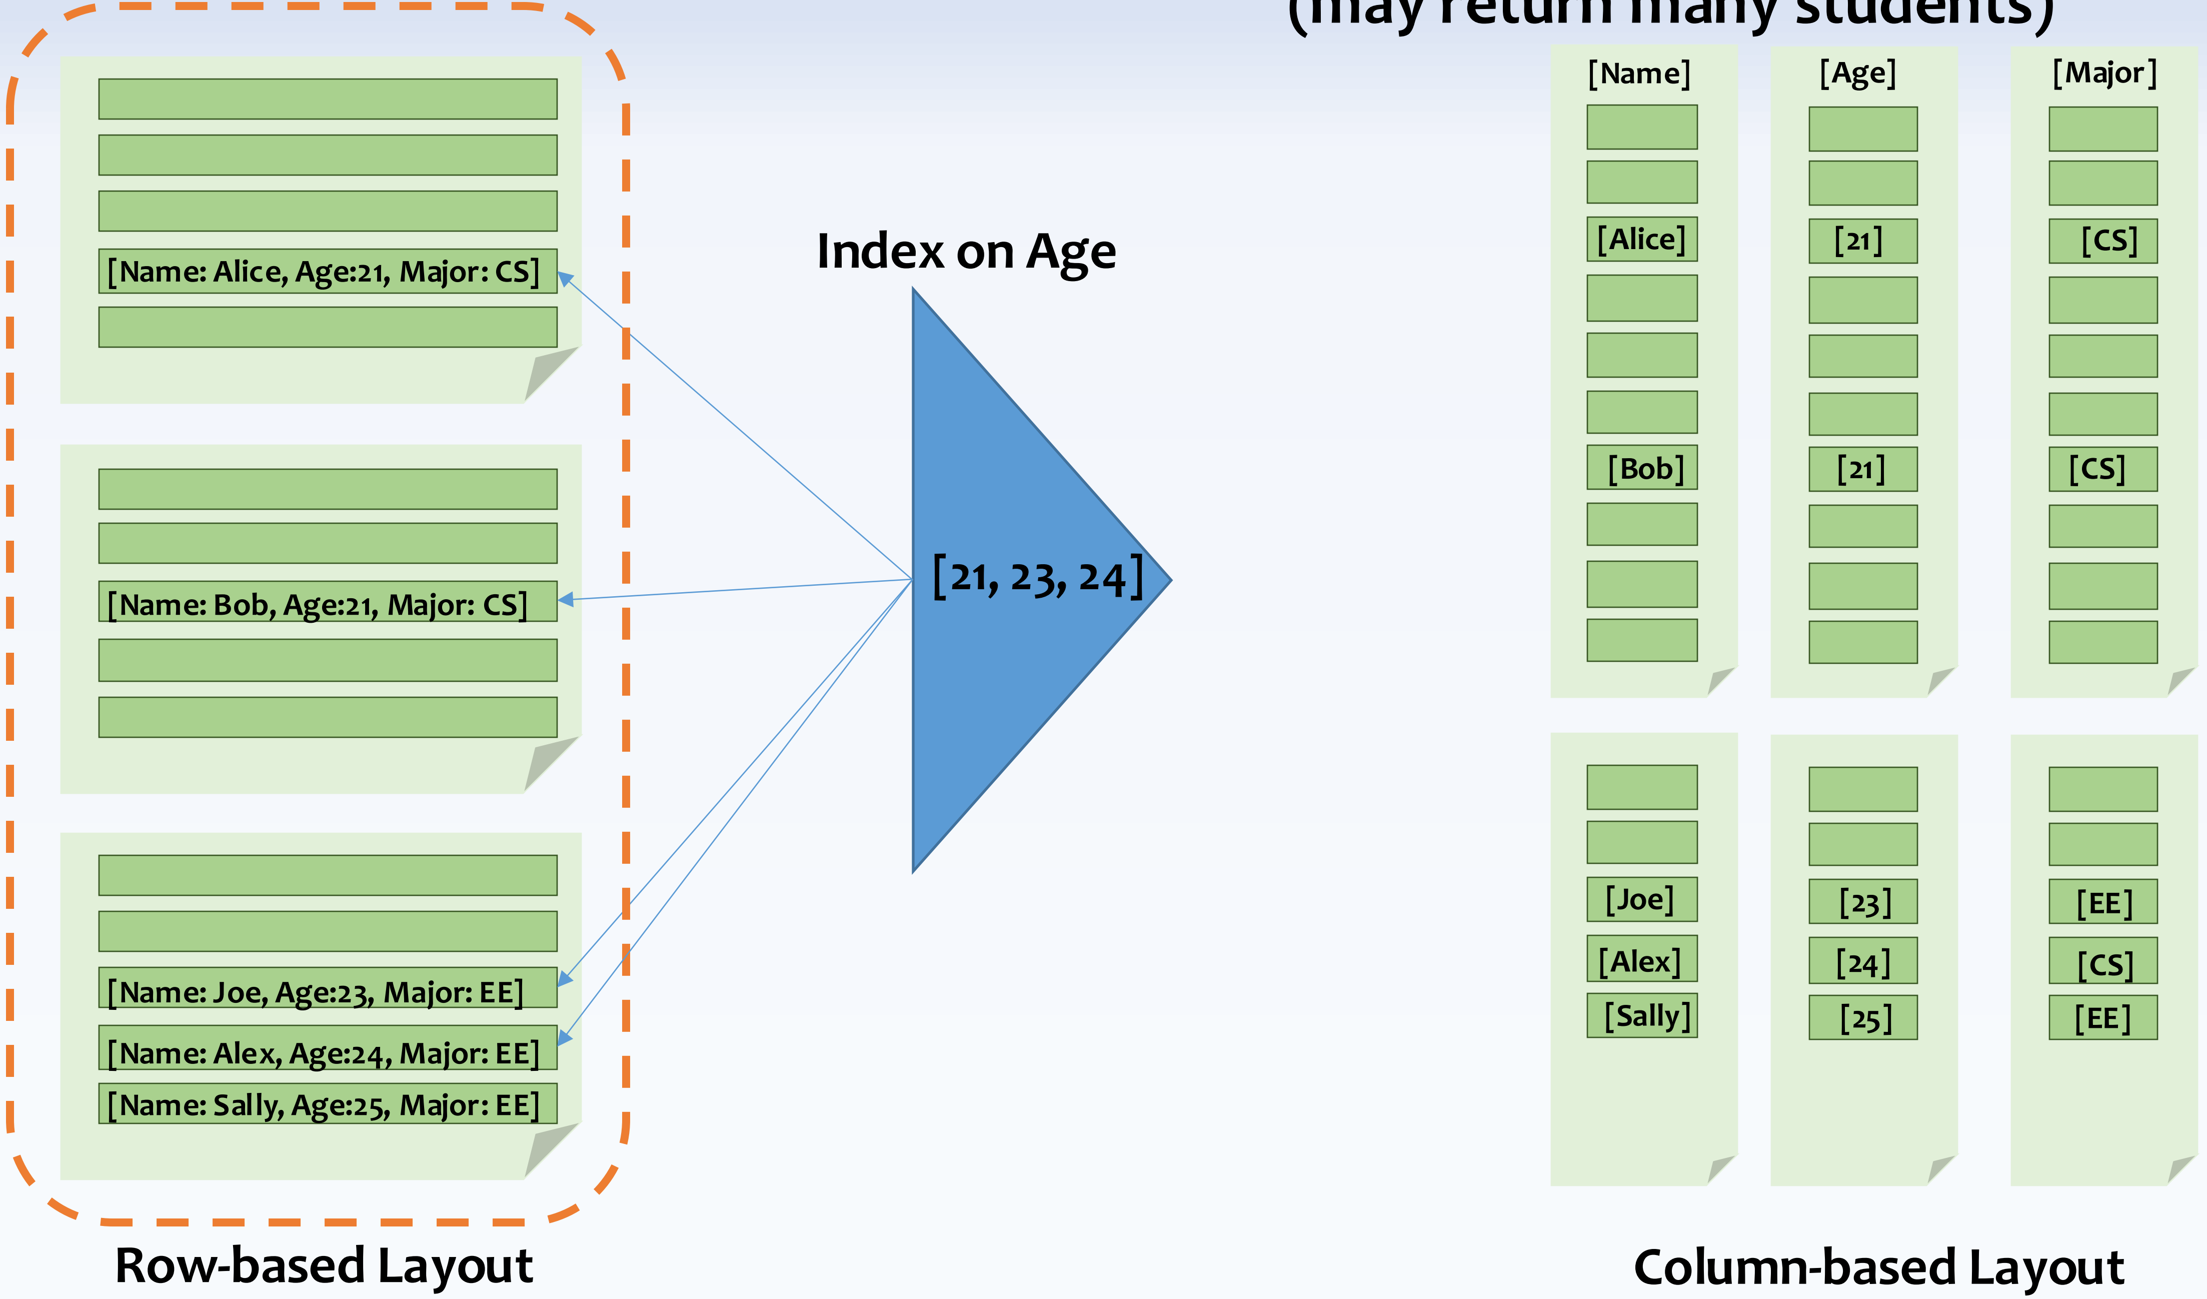

[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
Row-based Layout
Index on Age
[21, 23, 24]
[Name]
[Age]
[Major]
[Alice]
[21]
[CS]
[Bob]
[21]
[CS]
[Joe]
[23]
[EE]
[Alex]
[24]
[CS]
[Sally]
[25]
[EE]
Column-based Layout

Searching for all students between the age of 21 to 24 (may return many students)

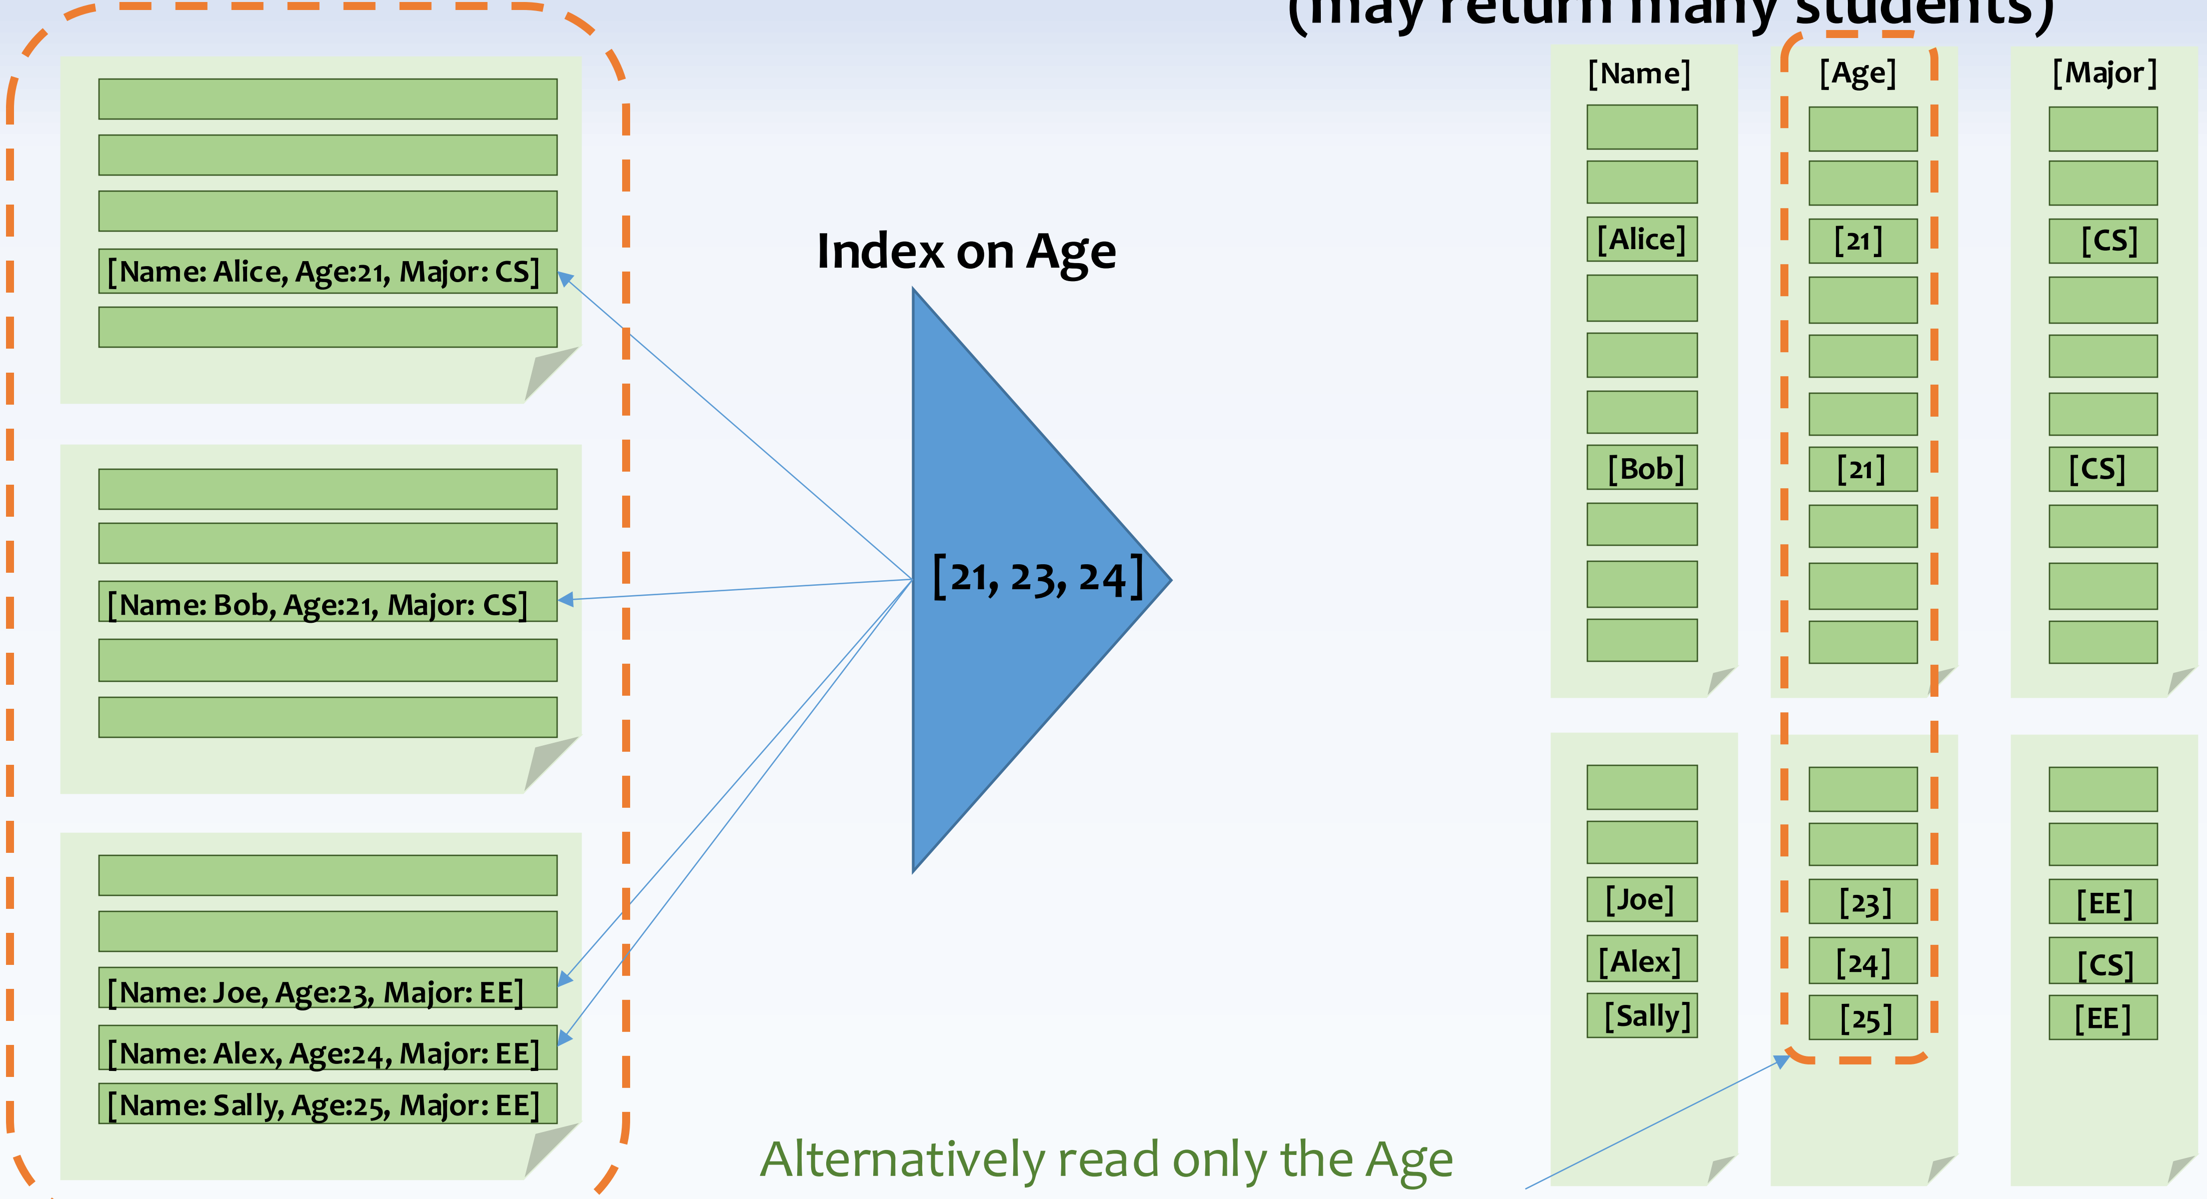

[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
Index on Age
[21, 23, 24]
[Name]
[Alice]
[Bob]
[Joe]
[Alex]
[Sally]
[Age]
[21]
[21]
[23]
[24]
[25]
[Major]
[CS]
[CS]
[EE]
[CS]
[EE]

Alternatively read only the Age column to find the relevant values

Row-based Layout

Column-based Layout

Searching for all students between the age of 21 to 24 (may return many students)

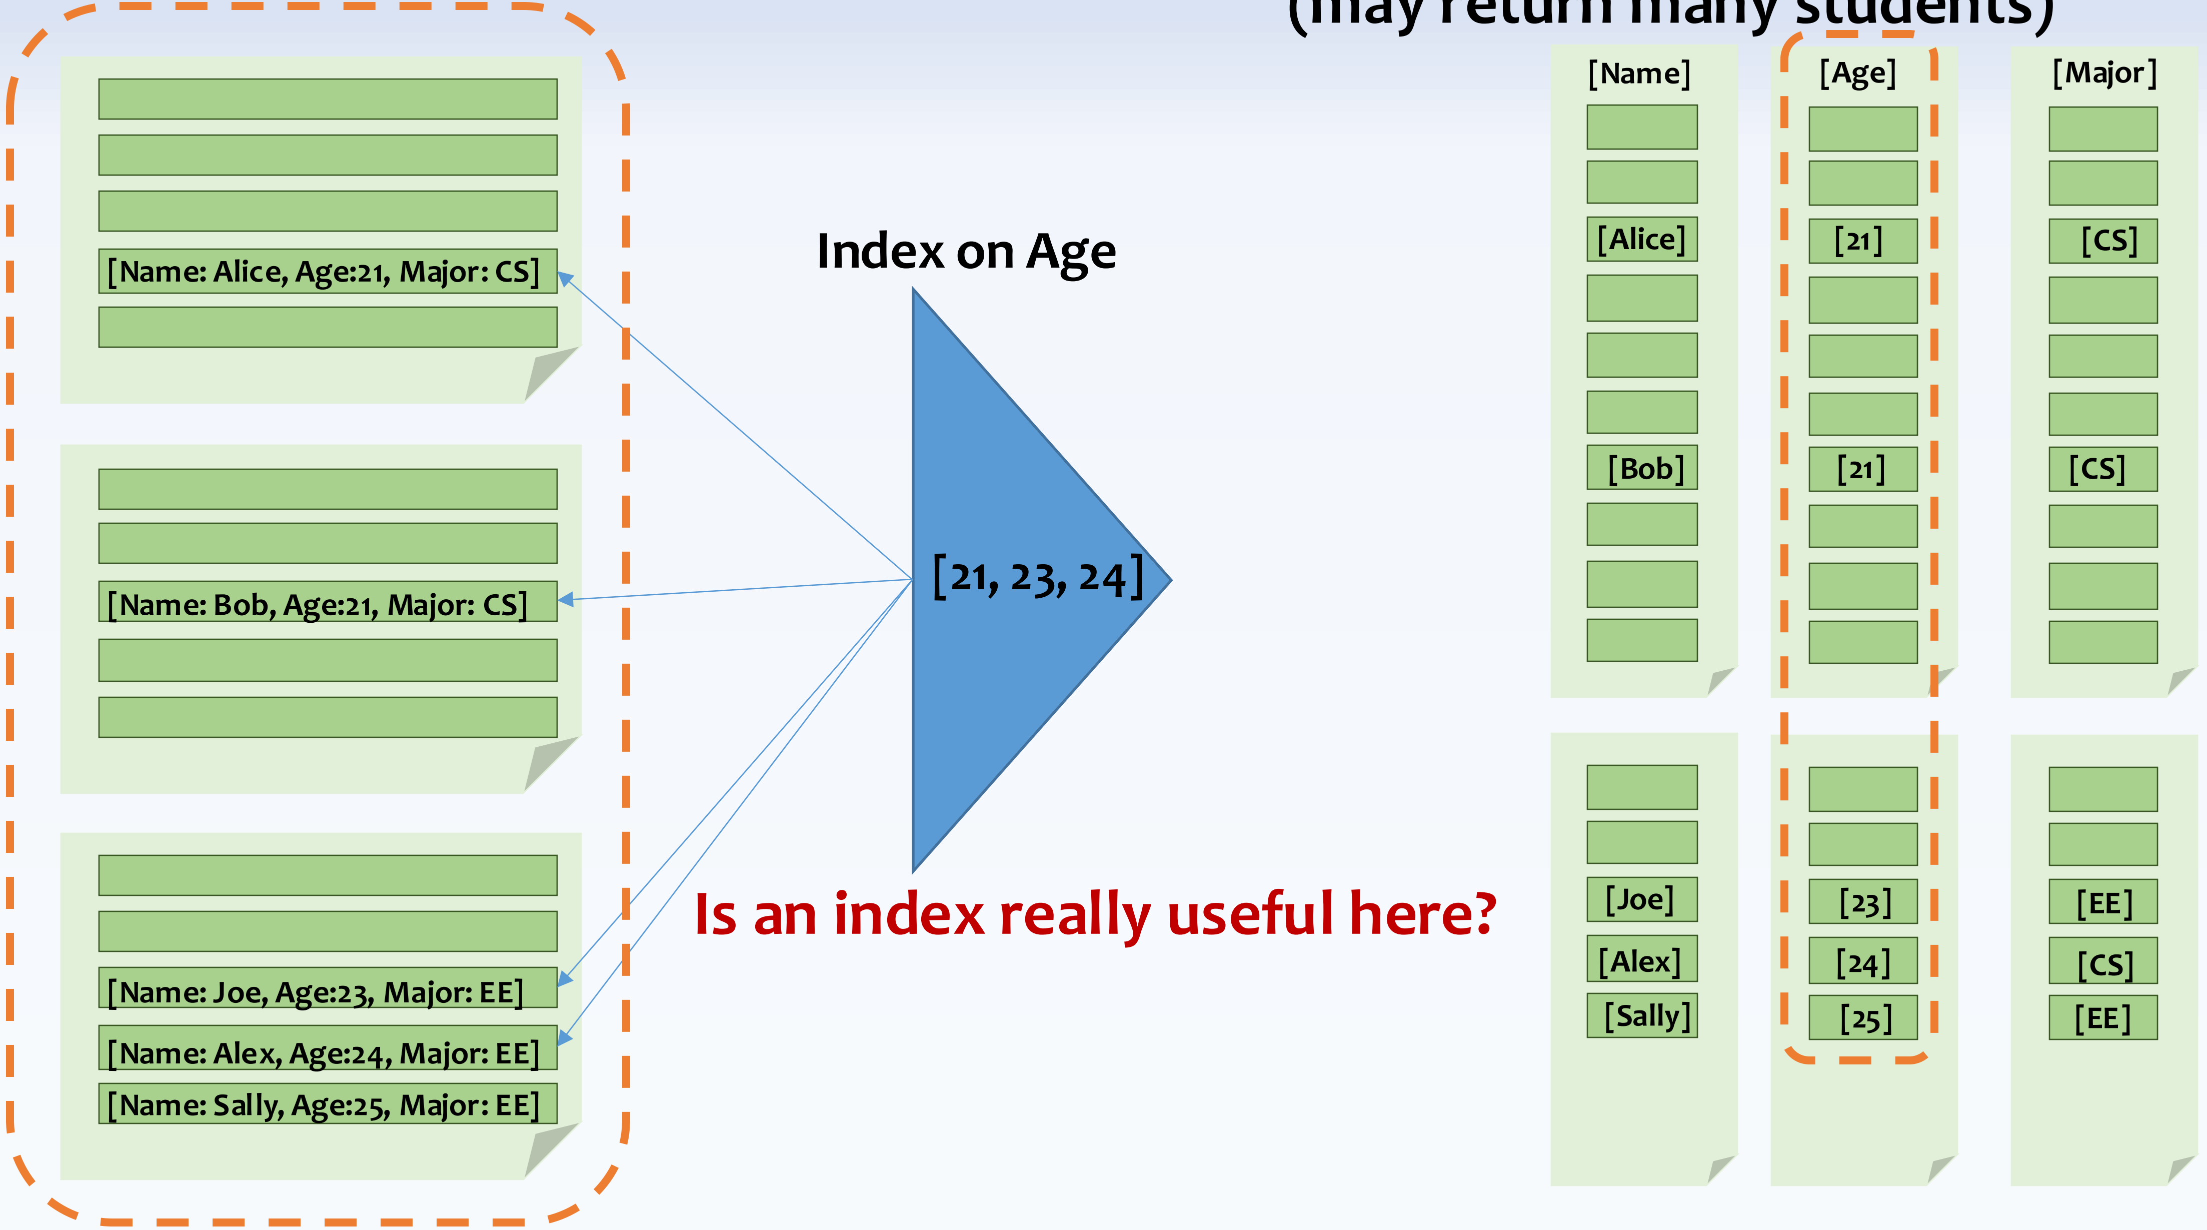
[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
Index on Age
[21, 23, 24]
[Name]
[Alice]
[Bob]
[Joe]
[Alex]
[Sally]
[Age]
[21]
[21]
[23]
[24]
[25]
[Major]
[CS]
[CS]
[EE]
[CS]
[EE]

Is an index really useful here?

Row-based Layout

Column-based Layout

Searching for all students over the age of 24 (may return only a few students)

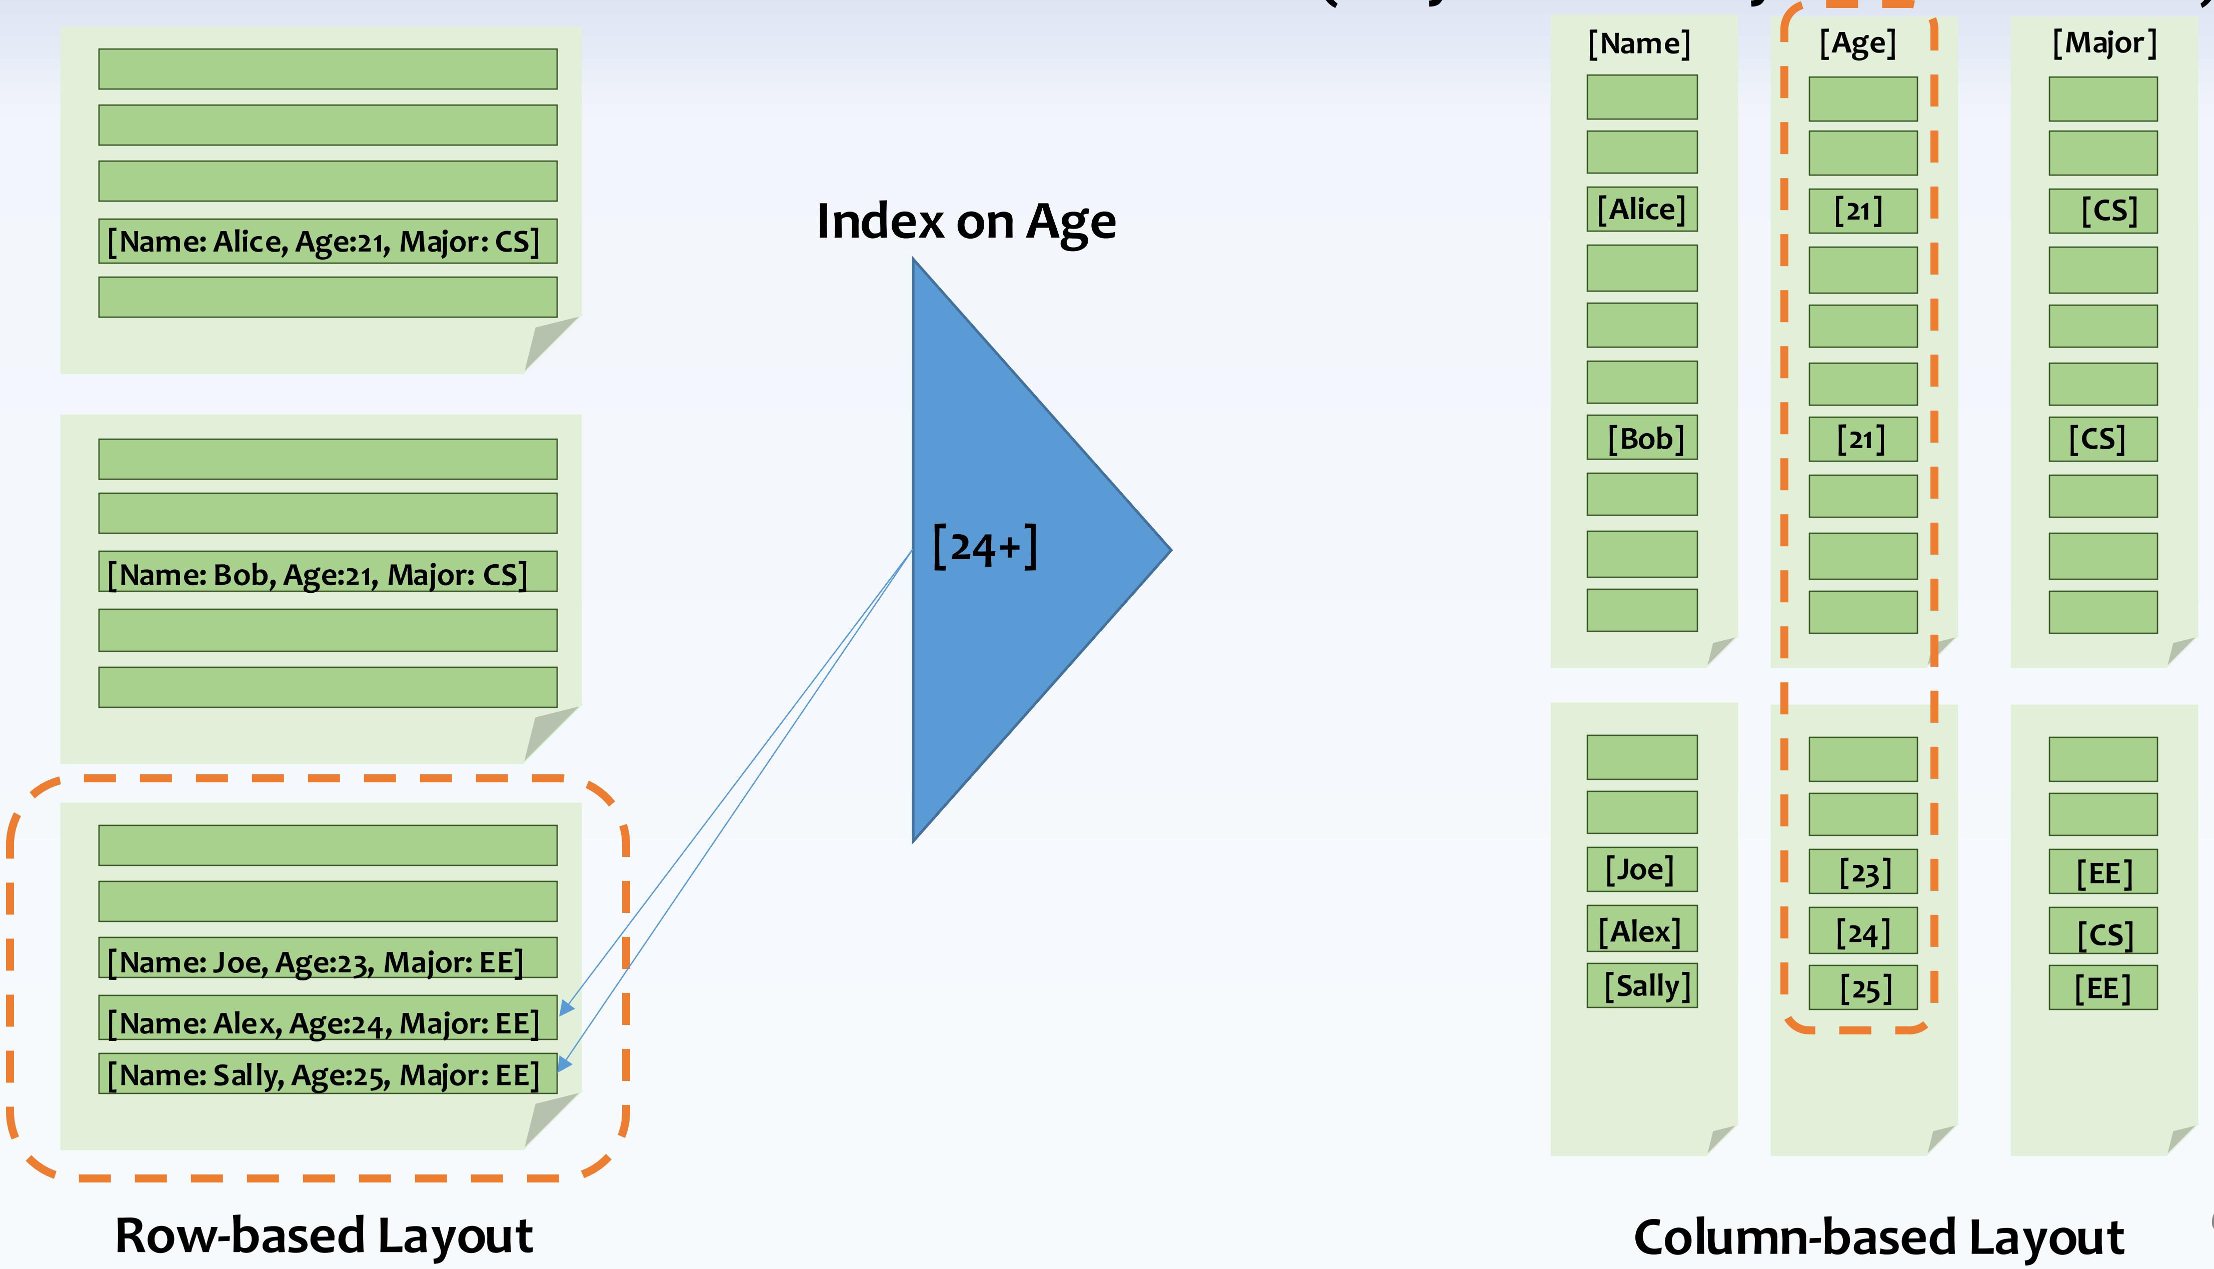

[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
Row-based Layout
Index on Age
[24+]
[Name]
[Alice]
[Bob]
[Joe]
[Alex]
[Sally]
[Age]
[21]
[21]
[23]
[24]
[25]
[Major]
[CS]
[CS]
[EE]
[CS]
[EE]
Column-based Layout

Searching for all students over the age of 24 (may return only a few students)

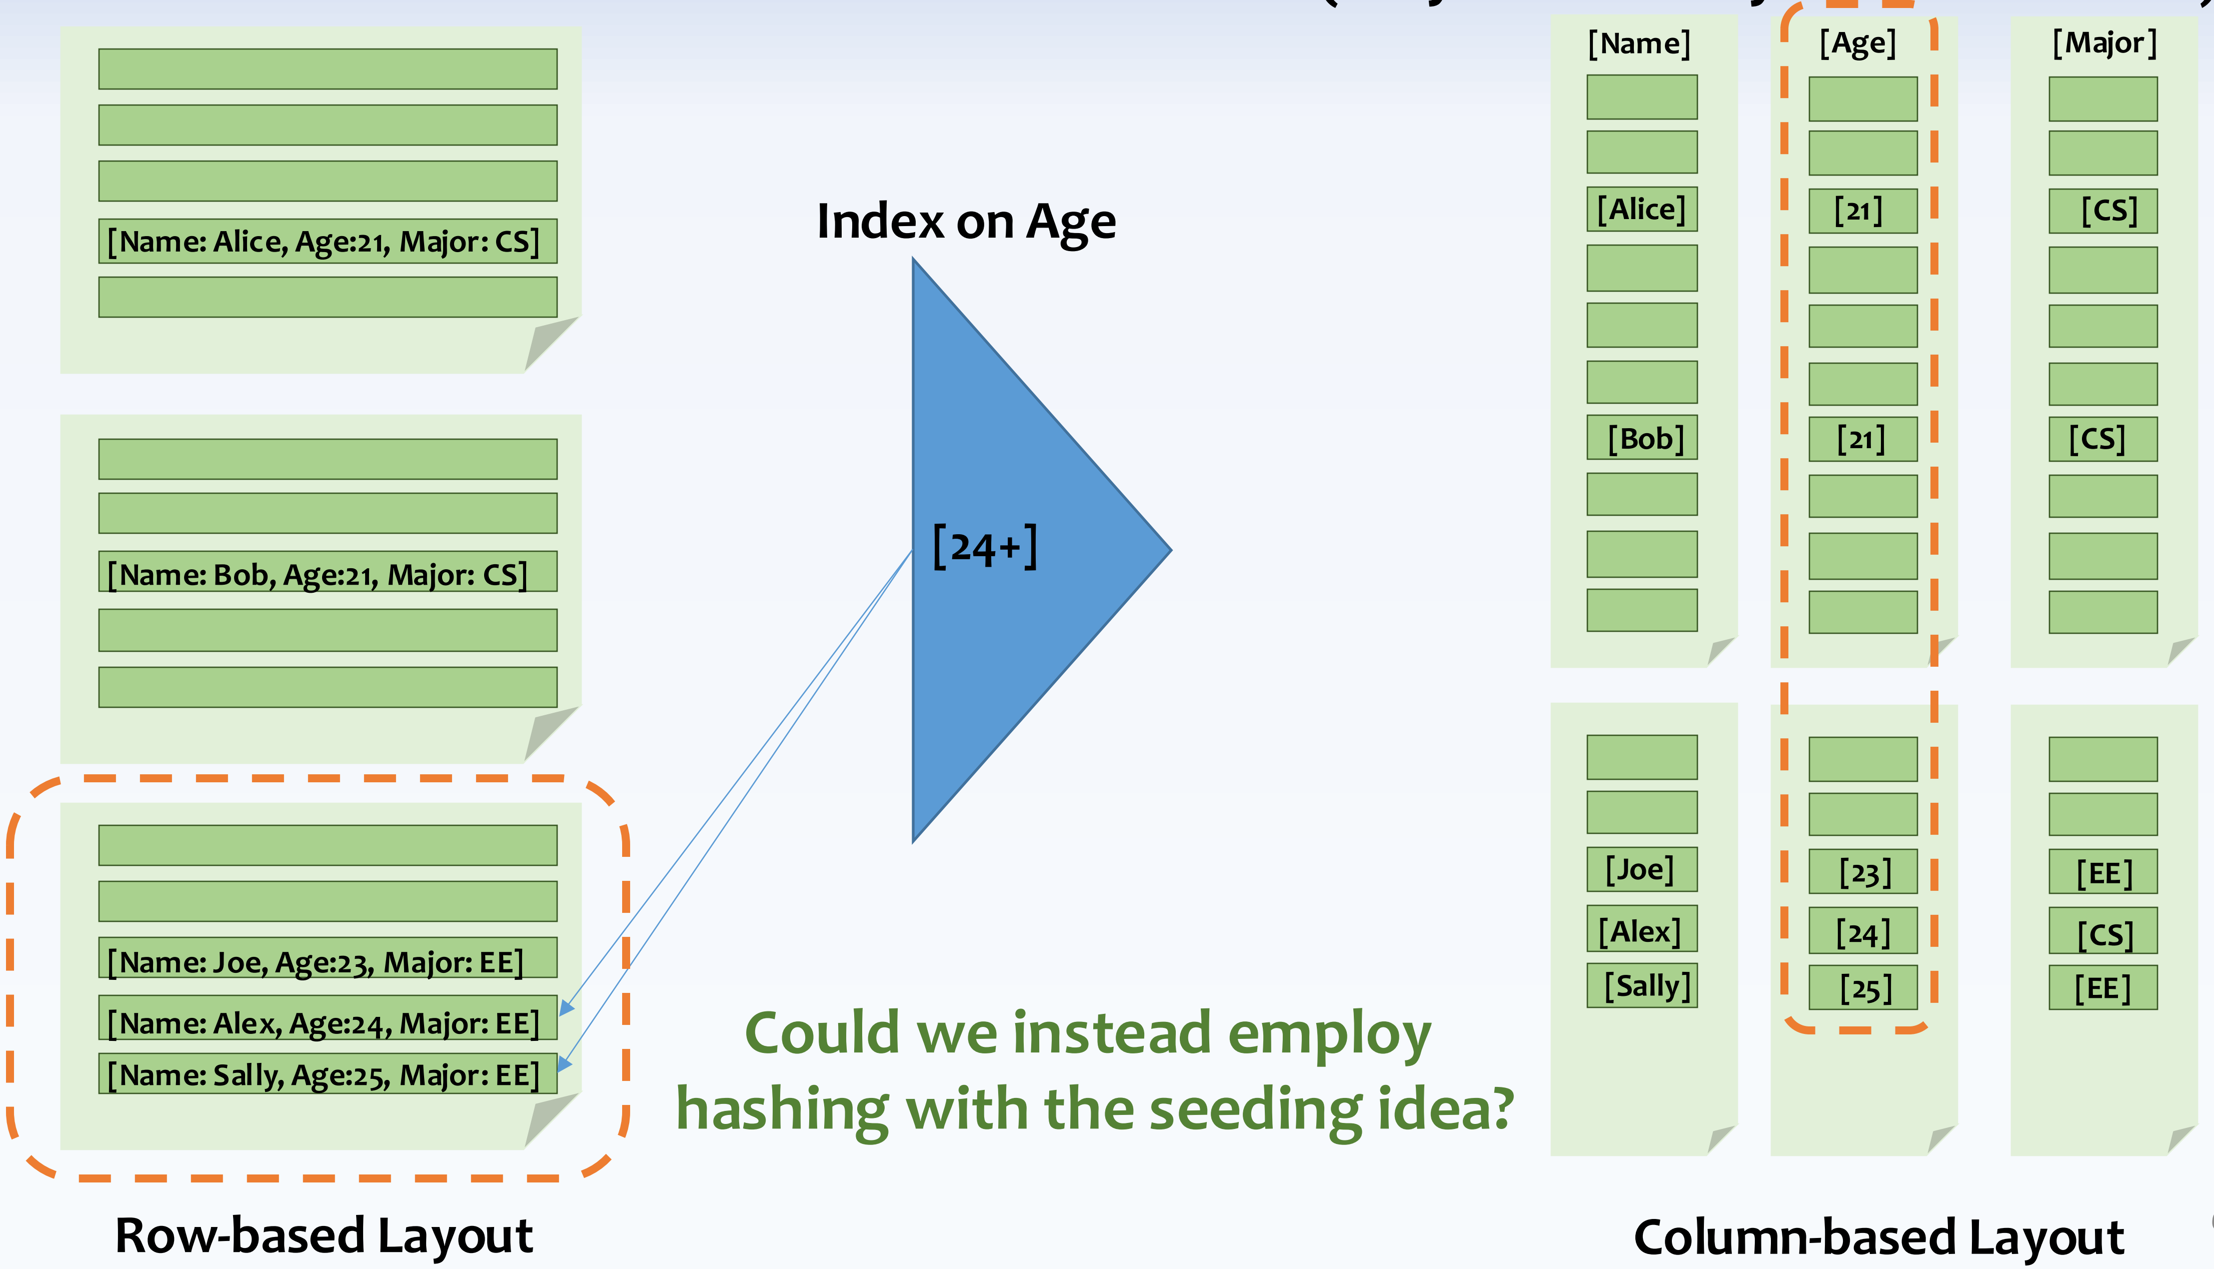
[Name: Alice, Age:21, Major: CS]
[Name: Bob, Age:21, Major: CS]
[Name: Joe, Age:23, Major: EE]
[Name: Alex, Age:24, Major: EE]
[Name: Sally, Age:25, Major: EE]
Index on Age
[24+]
[Name]
[Alice]
[Bob]
[Joe]
[Alex]
[Sally]
[Age]
[21]
[21]
[23]
[24]
[25]
[Major]
[CS]
[CS]
[EE]
[CS]
[EE]
Row-based Layout
Column-based Layout

Could we instead employ hashing with the seeding idea?

Thank You
Questions?