

Searchable Symmetric Encryption: Optimal Locality in Linear Space via Two-Dimensional Balanced Allocations

Gilad Asharov

Moni Naor

Gil Segev

Ido Shahaf

IBM Research

Weizmann

Hebrew University

Hebrew University



STOC 2016

Cloud Storage

- We are outsourcing more and more of our data to clouds
- We trust these clouds less and less
 - Confidentiality of the data from the service provider itself
 - Protect the data from service provider security breaches



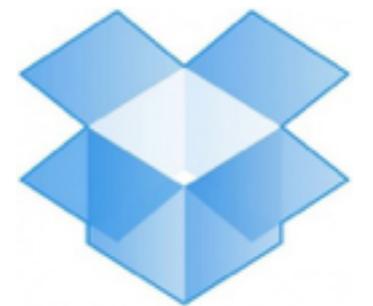
iCloud



Google Drive



OneDrive



Dropbox

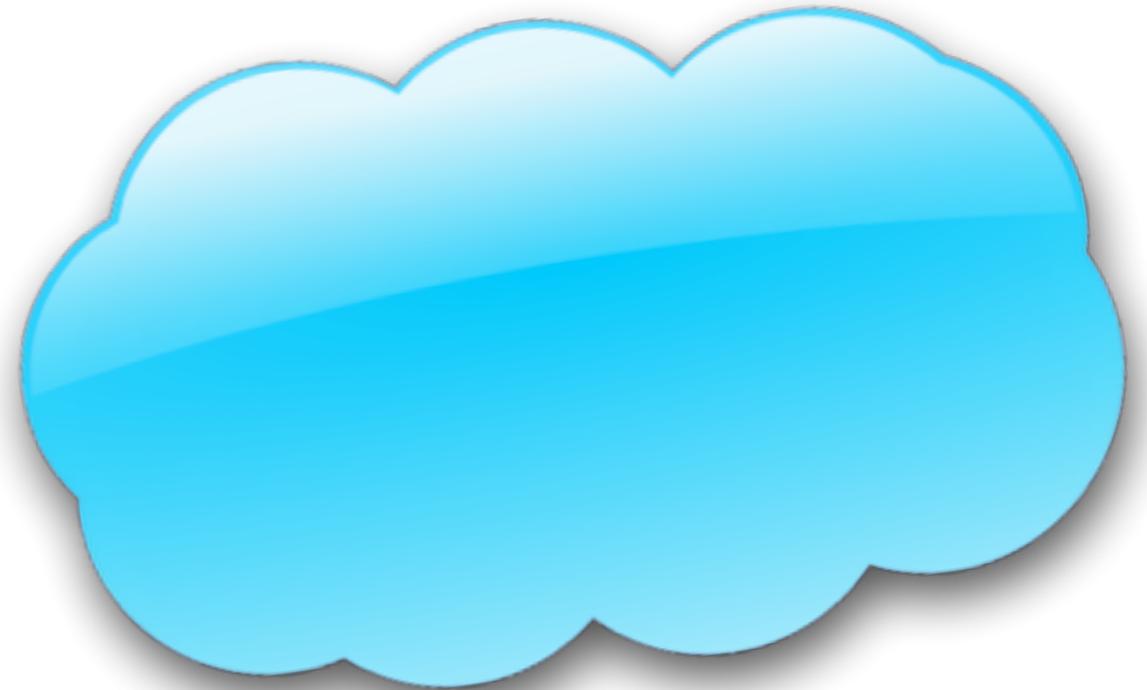
Solution: Encrypt your Data!

- But...
 - **Keyword search** is now the primary way we access our data
 - By encrypting the data - this simple operation becomes **extremely expensive**
- **How to search on encrypted data??**

Possible Solutions

- **Generic tools:** Expensive, great security
 - Functional encryption
 - Fully Homomorphic Encryption
 - Oblivious RAM*
- **More tailored solutions:** practical, security(?)
 - Property-preserving encryption
(encryption schemes that supports public tests)
 - Deterministic encryption [Bellare-Boldyreva-O'Neill06]
 - Order-preserving encryption [Agrawal-Kiernan-Srikant-Xu04]
 - Orthogonality preserving encryption [Pandey-Rouselakis04]
 - Searchable Symmetric Encryption [Song-Wagner-Perrig01]

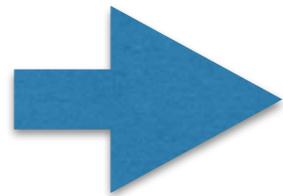
Searchable Symmetric Encryption (SSE)



Searchable Symmetric Encryption (SSE)

- **Data:** the database **DB** consists of:
 - **Keywords:** $W = \{w_1, \dots, w_n\}$ (possible keywords)
 - **Documents:** D_1, \dots, D_m (list of documents)
 - $DB(w_i) = \{id_1, \dots, id_{ni}\}$
(for every keyword w_i , list of documents / identifiers in which w_i appears)
- **Syntax of SSE:**
 - $K \leftarrow KeyGen(1^k)$ (generation of a private key)
 - $E_{DB} \leftarrow E_{DB}Setup(K, DB)$ (encrypting the database)
 - $(DB(w_i), \lambda) \leftarrow Search((K, w_i), E_{DB})$ (interactive protocol)

EDBSetup



Keyword	Records
Searchable	5,14
Symmetric	5,14,22,45,67
Encryption	1,2,3,4,5,6,7,8,9,10
Schemes	22,14

inverted index

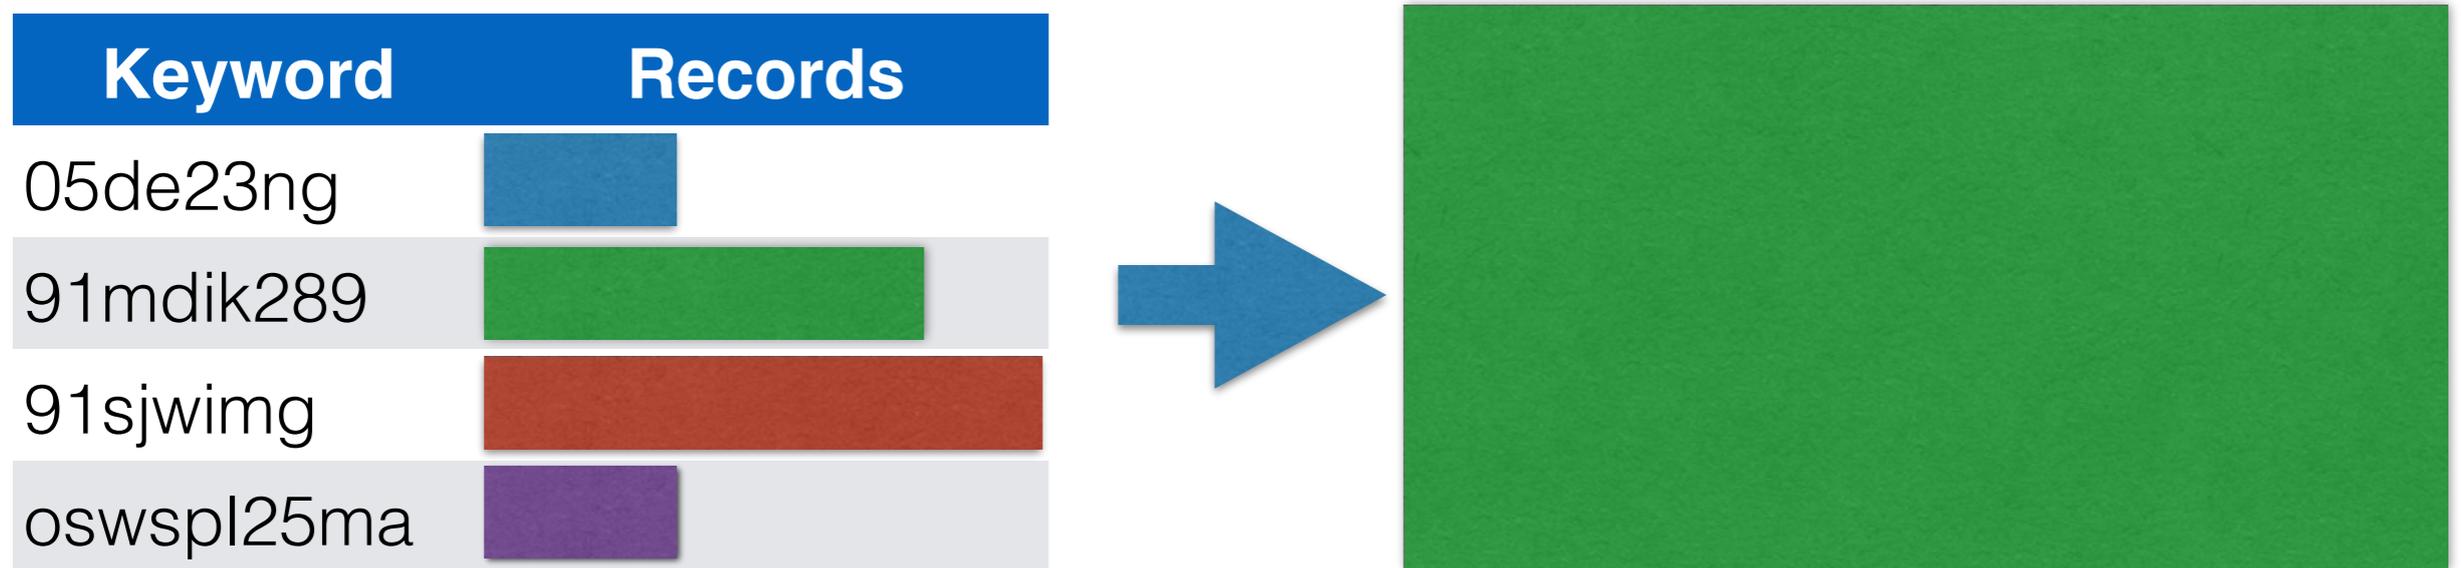
Replace each keyword w
with some $\text{PRF}_k(w)$

Keyword	Records
05de23ng	5,14
91mdik289	5,14,22,45,67
91sjwimg	1,2,3,4,5,6,7,8,9,10
oswspl25ma	22,14

encrypted index

Keyword	Records
05de23ng	
91mdik289	
91sjwimg	
oswspl25ma	

The Challenge...

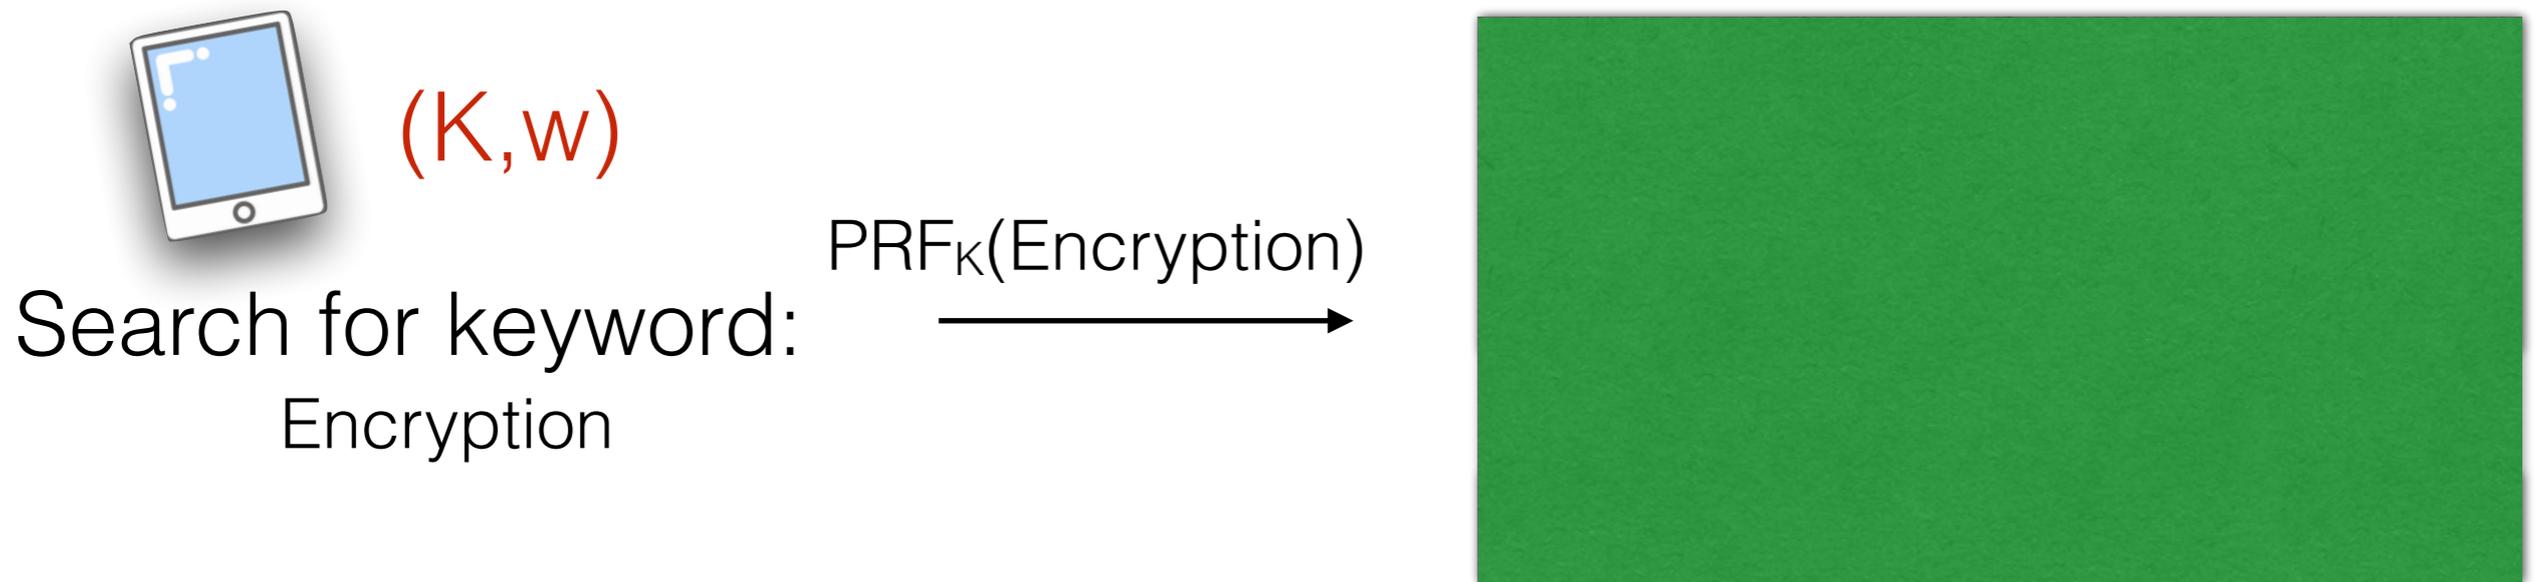


No leakage on the structure of the lists!

How to map the lists into memory?

Functionality - Search

(Allow some Leakage...)



Security Requirement:

The server should not learn anything about the structure of lists that were not queried

Security

- **Good news:** Semantic security for data; no deterministic or order preserving encryption
- But.. for reasonable performance -> **leakage** for server
- Leakage in the form of access patterns to retrieved data and queries
 - Data is encrypted but server can see intersections b/w query results
(e.g. identify popular document)
- Additional specific leakage:
 - E.g. we leak $|DB(w1)|$
 - E.g. the server learns if two documents have the same keyword
- Leads to statistical inference based on side information on data
(effect depends on application)

Mapping Lists into Memory

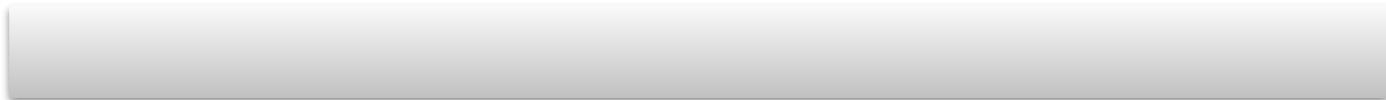
Maybe shuffle the lists?

Keyword	Records
05de23ng	
91mdik289	
91sjwimg	
oswspl25ma	

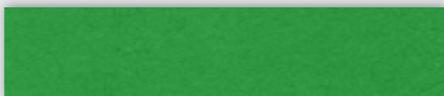
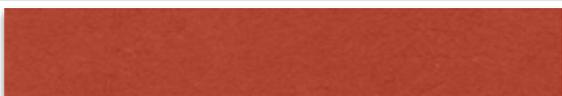
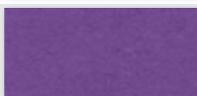


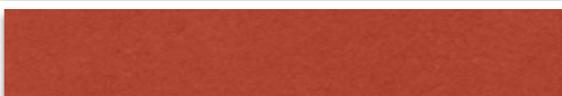
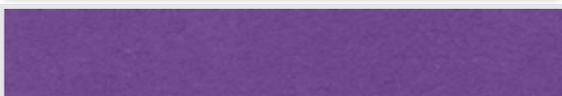
Hiding the Structure of the Lists

Maybe shuffle the lists?



Previous Constructions: Maximal Padding [CK10]

Keyword	Records
05de23ng	
91mdik289	
91sjwimg	
oswspl25ma	

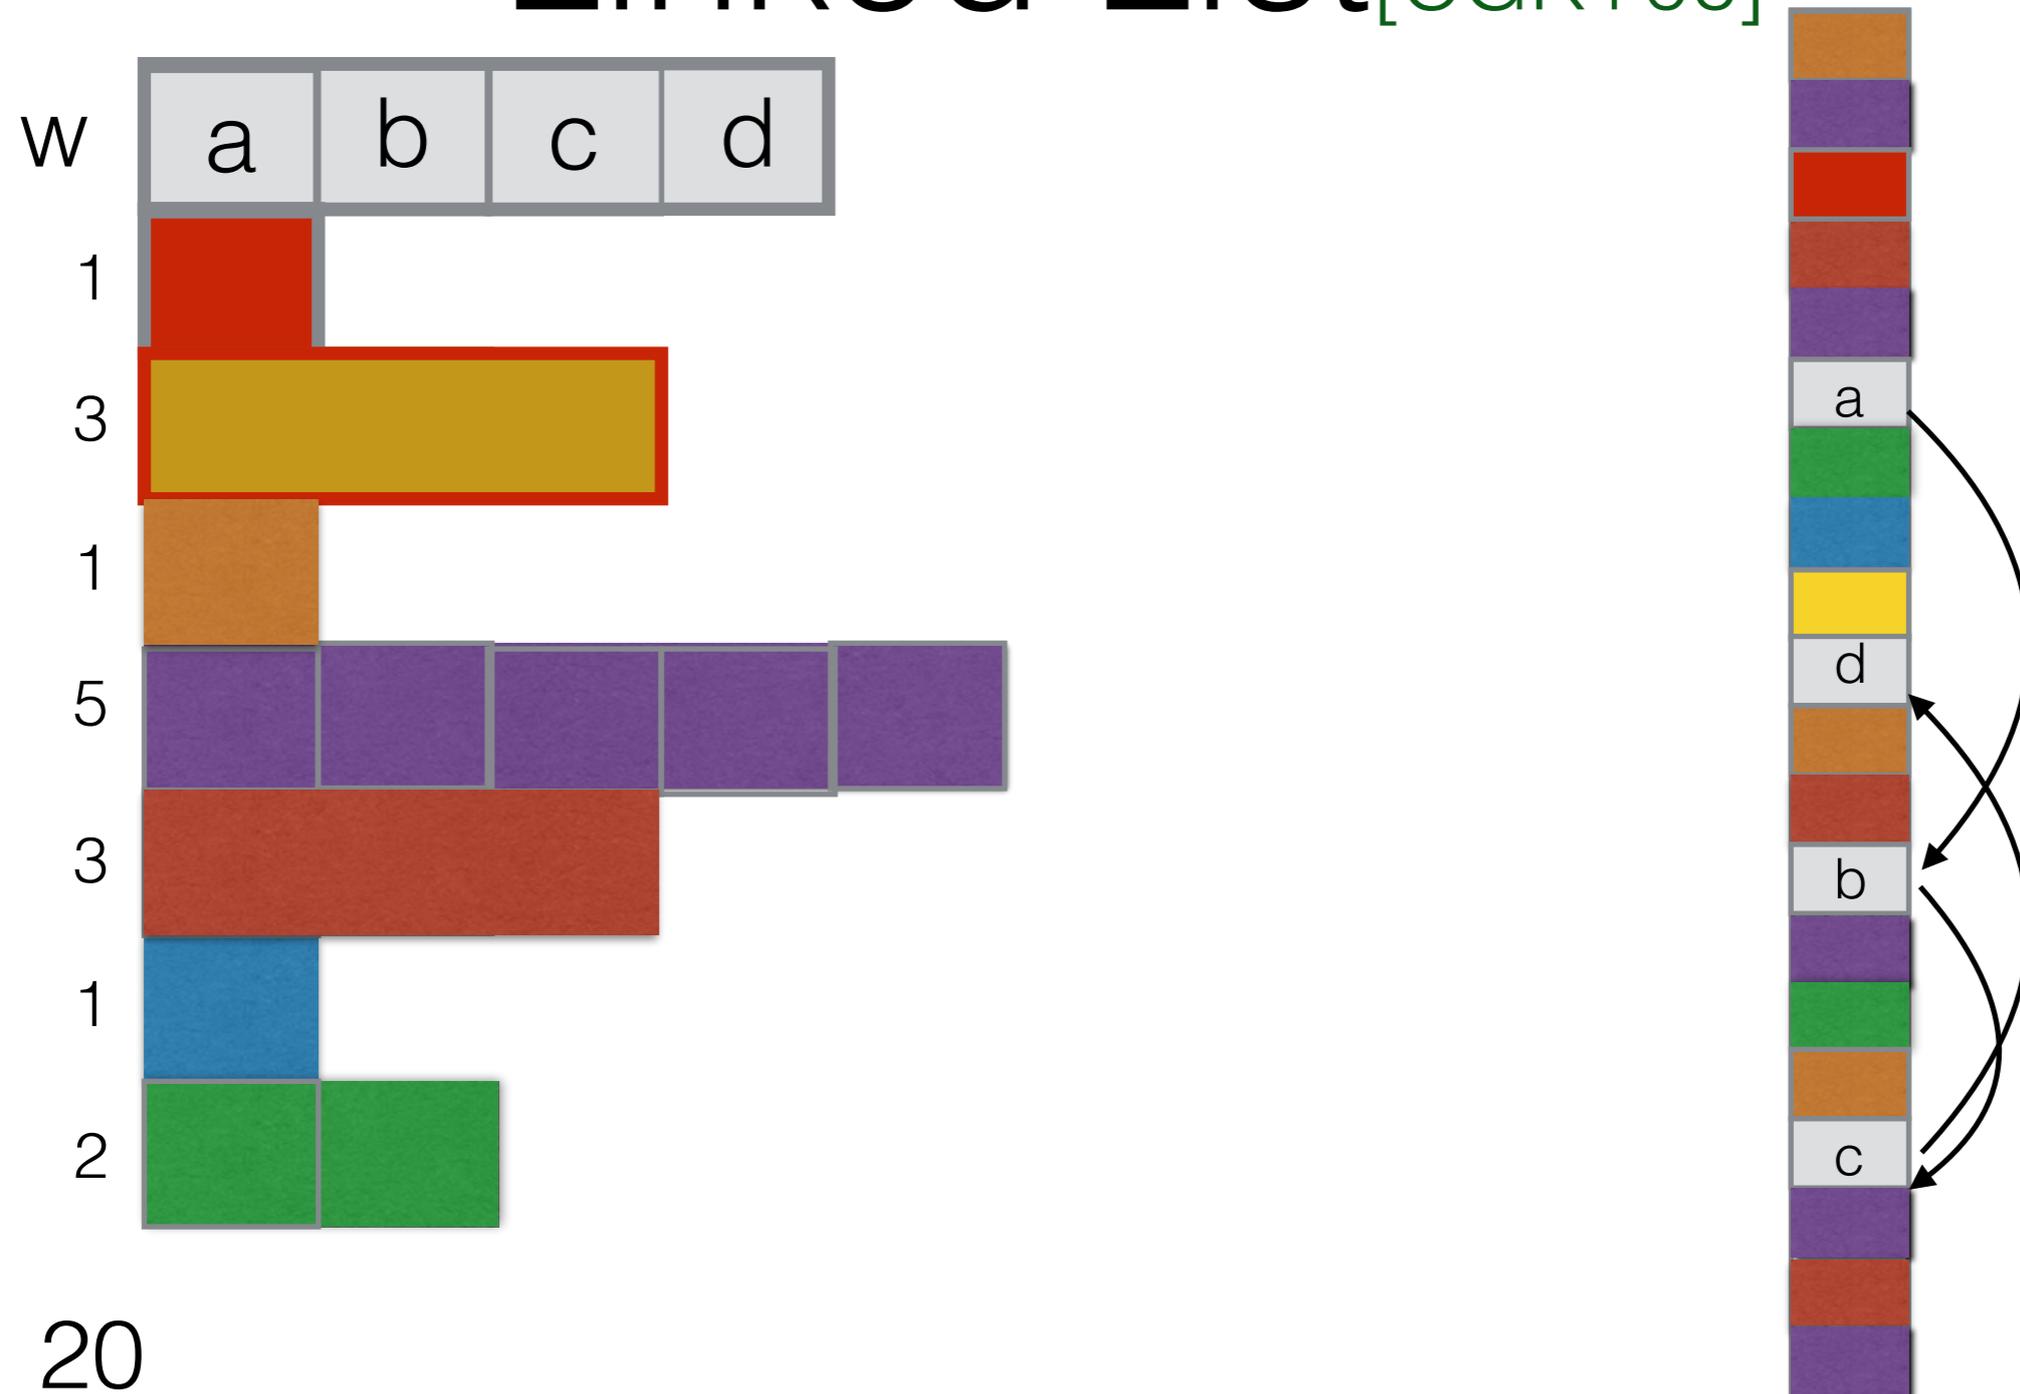
Keyword	Records
05de23ng	
91mdik289	
91sjwimg	
oswspl25ma	

- 1) Pad each list to maximal size (N?)
- 2) Store lists in random order
- 3) Pad with extra lists to hide the number of lists

Size of encrypted DB: $O(N^2)$

Previous Constructions

Linked List_[CGK+06]



Efficiency Measures

- A variant was implemented in [CJJ+13]
 - Poor performance due to... locality!



- **Space**: The overall size of the encrypted database (Want: $O(N)$)
- **Locality**: number of non-continuous memory locations the server accesses with each query (Want: $O(1)$)
- **Read efficiency**: The ratio between the number of bits the server reads with each query, and the actual size of the answer (Want: $O(1)$)

SSE and Locality [CT14]

Can we construct an SSE scheme that is optimal in **space**, **locality** and **read efficiency**?

NO!

- **Lower bound:** any scheme must be sub-optimal in either its **space** overhead, **locality** or **read efficiency**
- Impossible to construct scheme with $O(N)$ **space**, $O(1)$ **locality** and $O(1)$ **read efficiency**

Our Question:

can we construct a scheme that is nearly optimal?

Related Work

- A single keyword search
 - Related work [SWP00,Goh03,CGKO06,ChaKam10]
- Beyond single keyword search
 - Conjunctions, range queries, general boolean expression, wildcards [CJJKRS13,JKRS13,CJJJKRS14,FJKNRS15]
 - Schemes that are not based on inverted index [PKVKMCGKB14, FVKKKMB15]
- **Locality** in searchable symmetric encryption [CT14]
- Dynamic searchable symmetric encryption [.....]



Our Work

Our Results

Scheme	Space	Locality	Read Efficiency
[CGK+06,KPR12,CJJ+13]	$O(N)$	$O(n_w)$	$O(1)$
[CK10]	$O(N^2)$	$O(1)$	$O(1)$
[CT14]	$O(N \log N)$	$O(\log N)$	$O(1)$
This work I	$O(N)$	$O(1)$	$\tilde{O}(\log N)$
This work II*	$O(N)$	$O(1)$	$\tilde{O}(\log \log N)$
This work III	$O(N \log N)$	$O(1)$	$O(1)$

$\tilde{O}(f(N)) = O(f(n) \log f(n))$

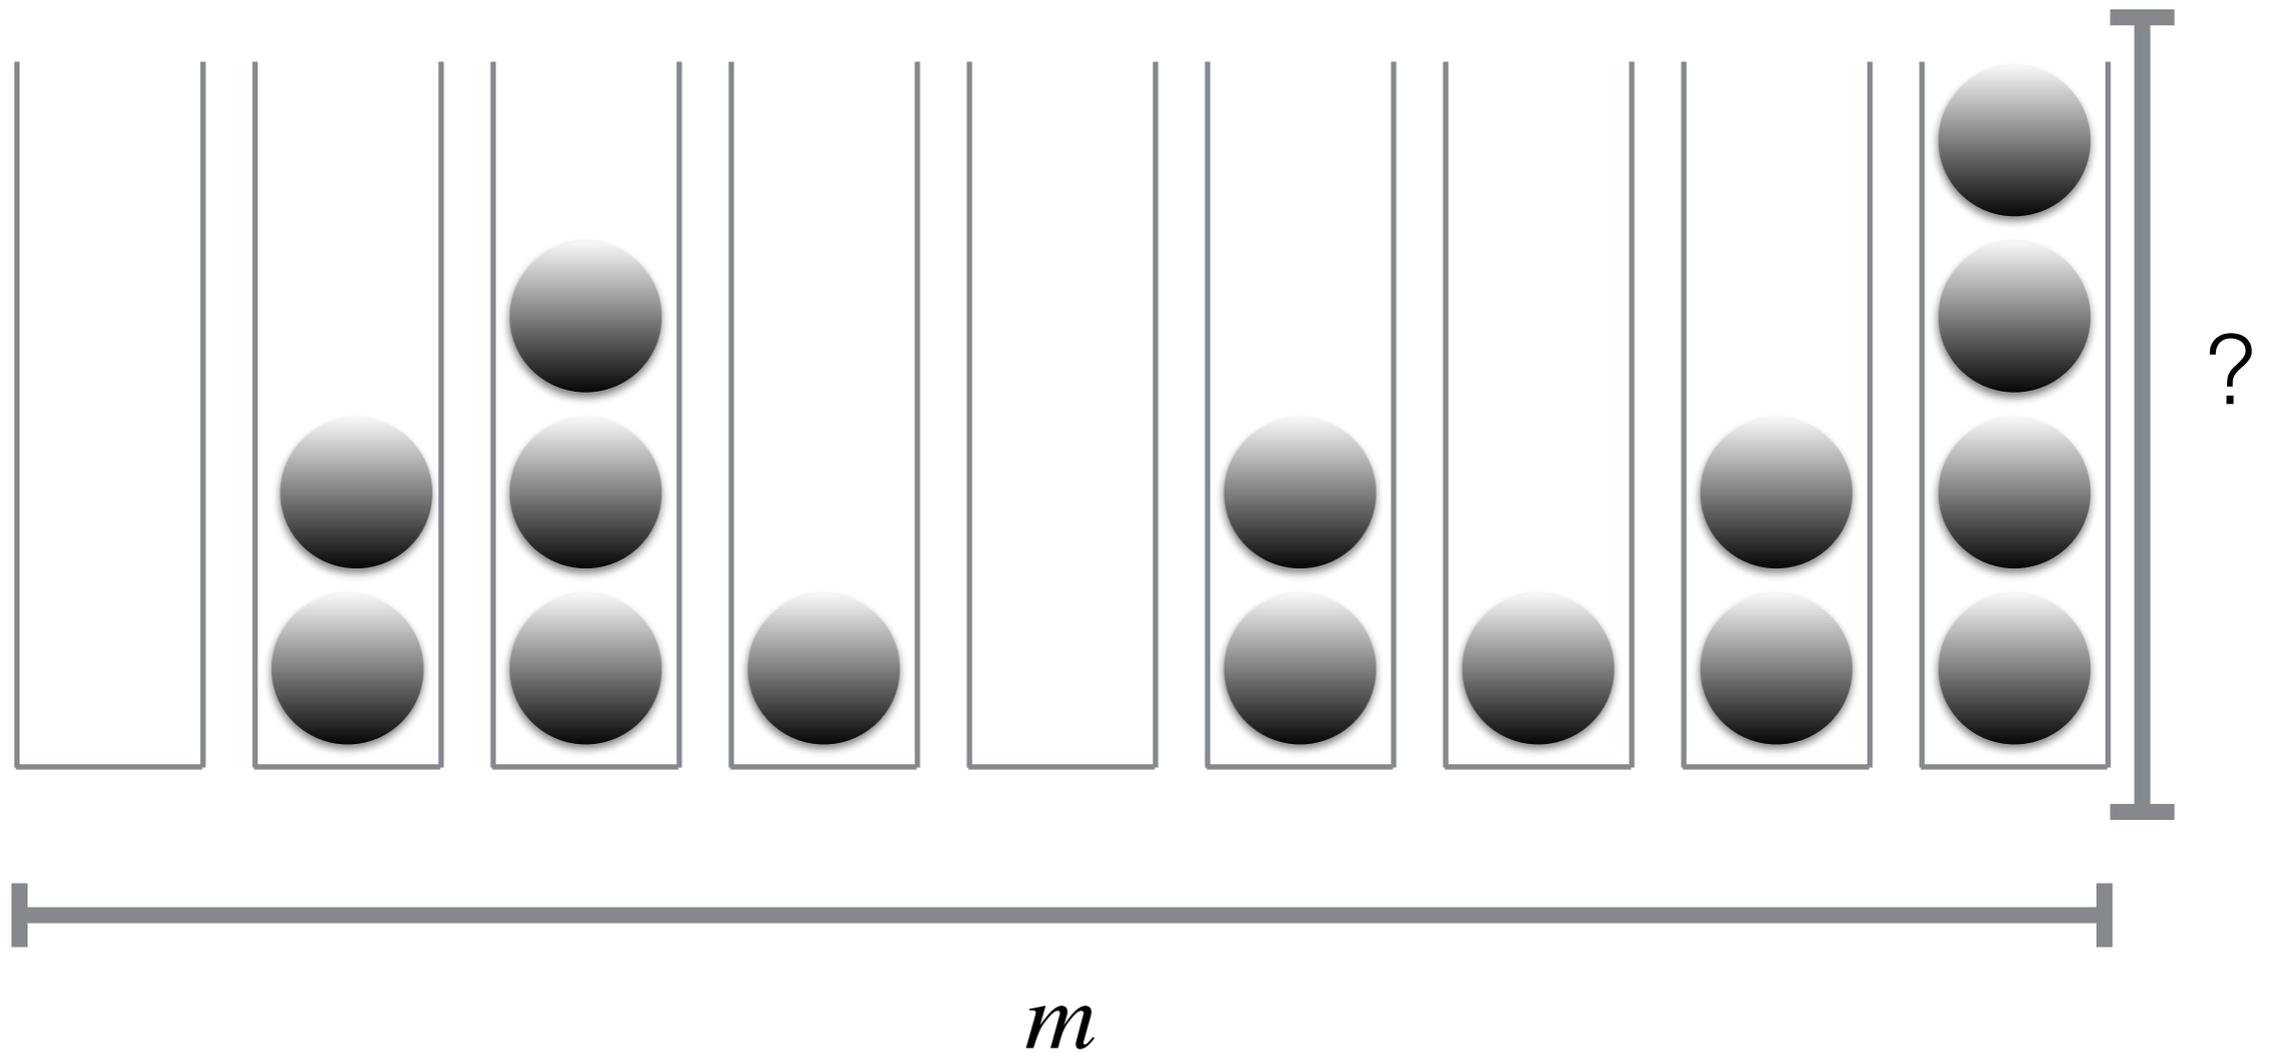
*assumes no keyword appears in more than $N^{1-1/\log \log N}$ documents

Our Approach

- We put forward a **two-dimensional** generalization of the classic **balanced allocation problem** (“balls and bins”), considering **lists of various lengths** instead of “balls” (=lists of fixed length)
 - (1) We construct efficient $2D$ balanced allocation schemes
 - (2) Then, we use cryptographic techniques to transform any such scheme into an SSE scheme

Balls and Bins

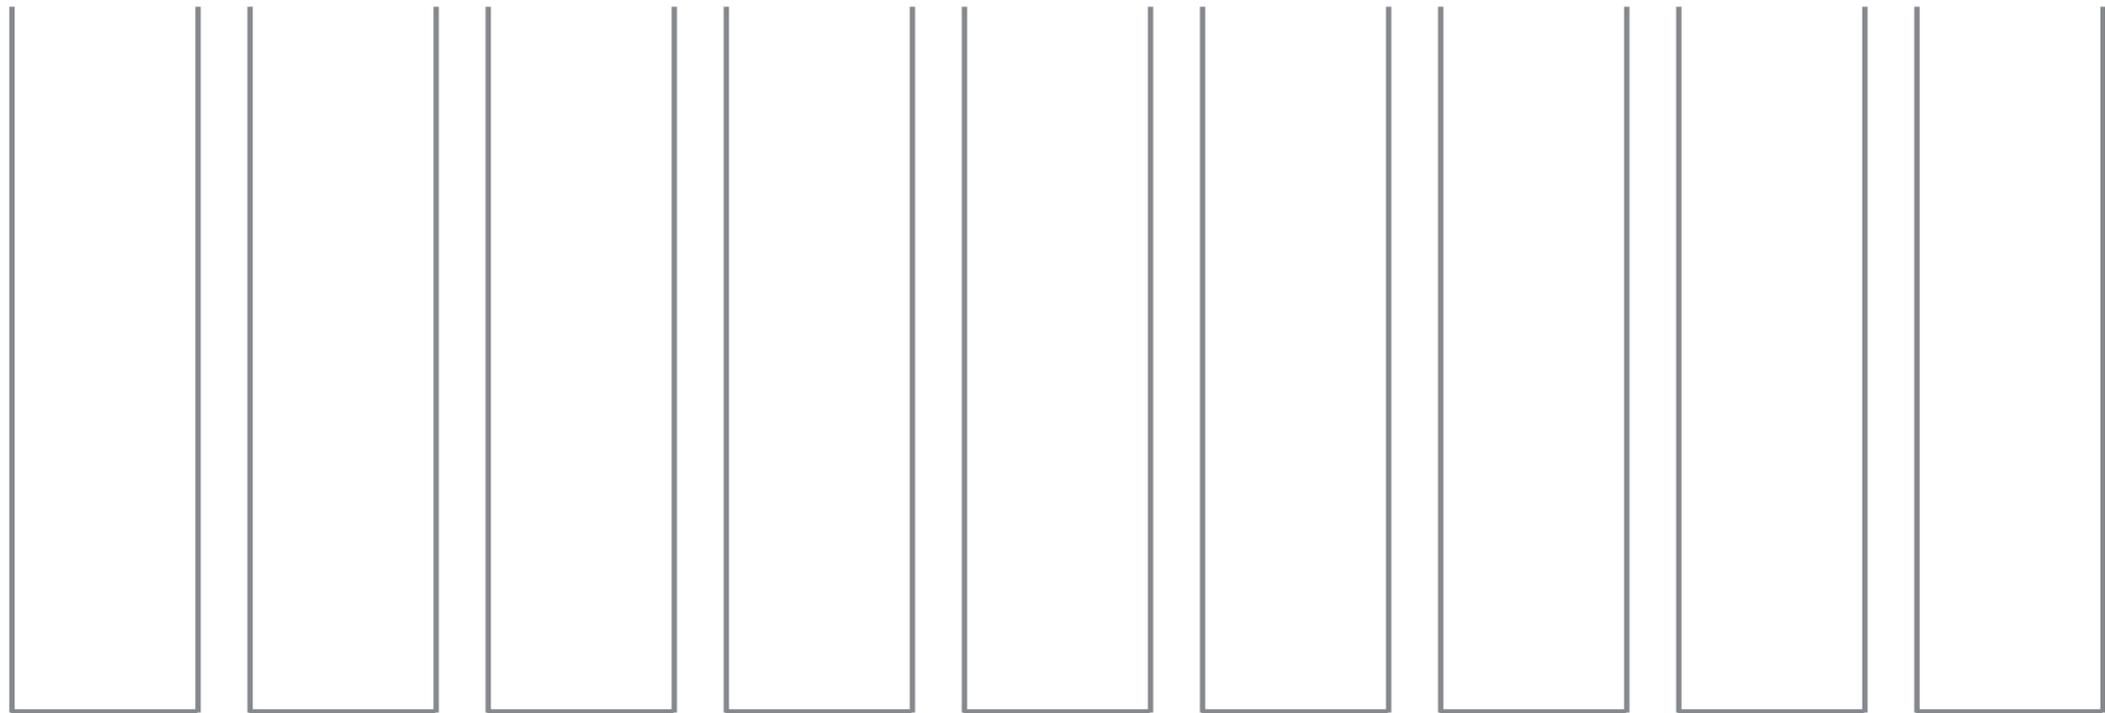
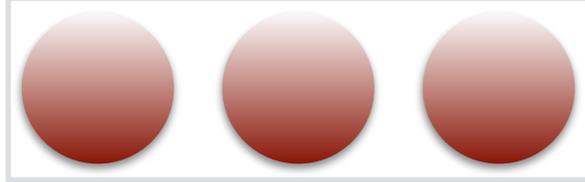
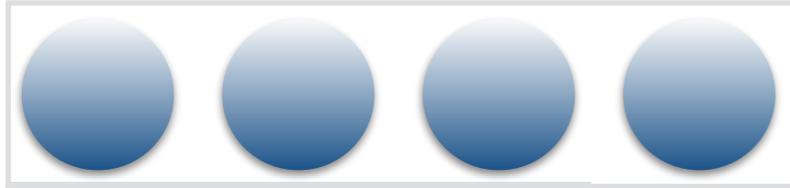
● × n



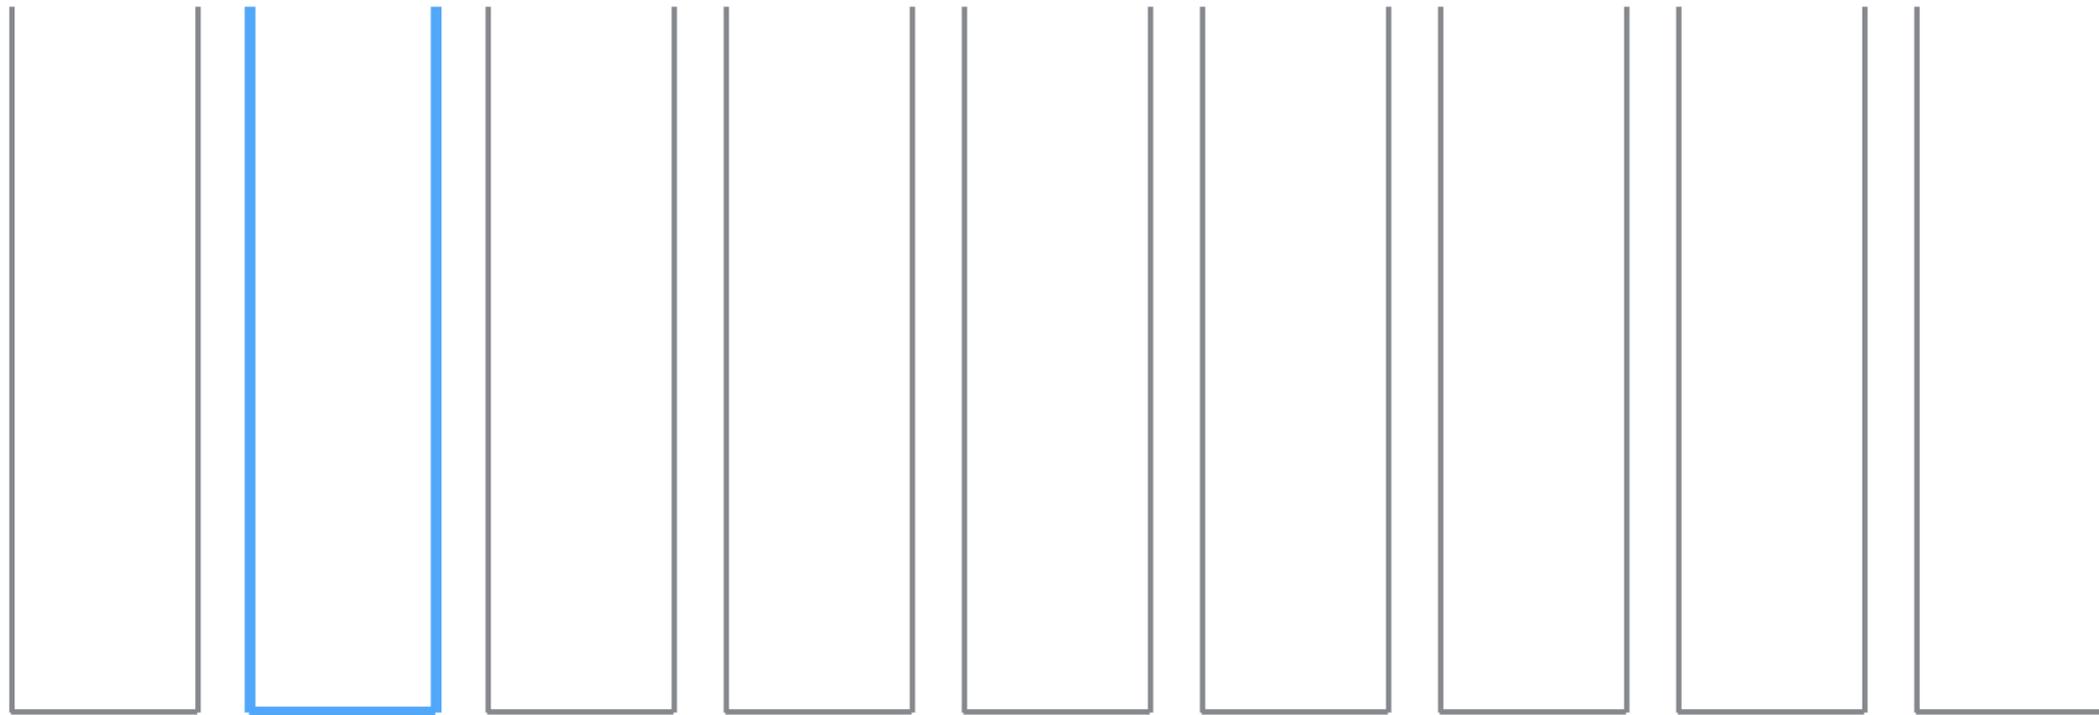
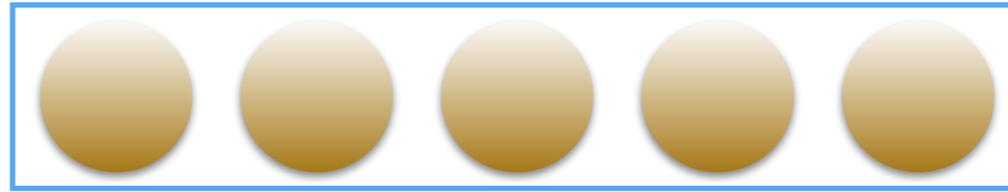
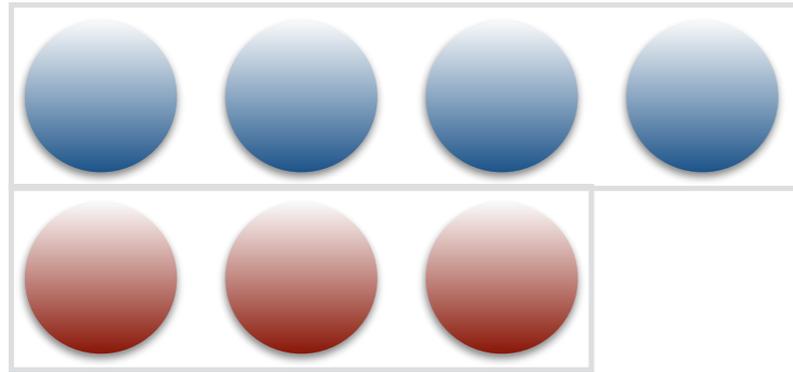
Balls and Bins (Random Allocation)

- n balls, m bins
 - Choose for each ball one bin uniformly at random
 - **$m=n$** : with high probability - there is no bin with more than $\frac{\log n}{\log \log n} \cdot (1 + o(1))$
 - **$m=n/\log n$** : with overwhelming probability, there is no bin with load greater than $\tilde{O}(\log n)$

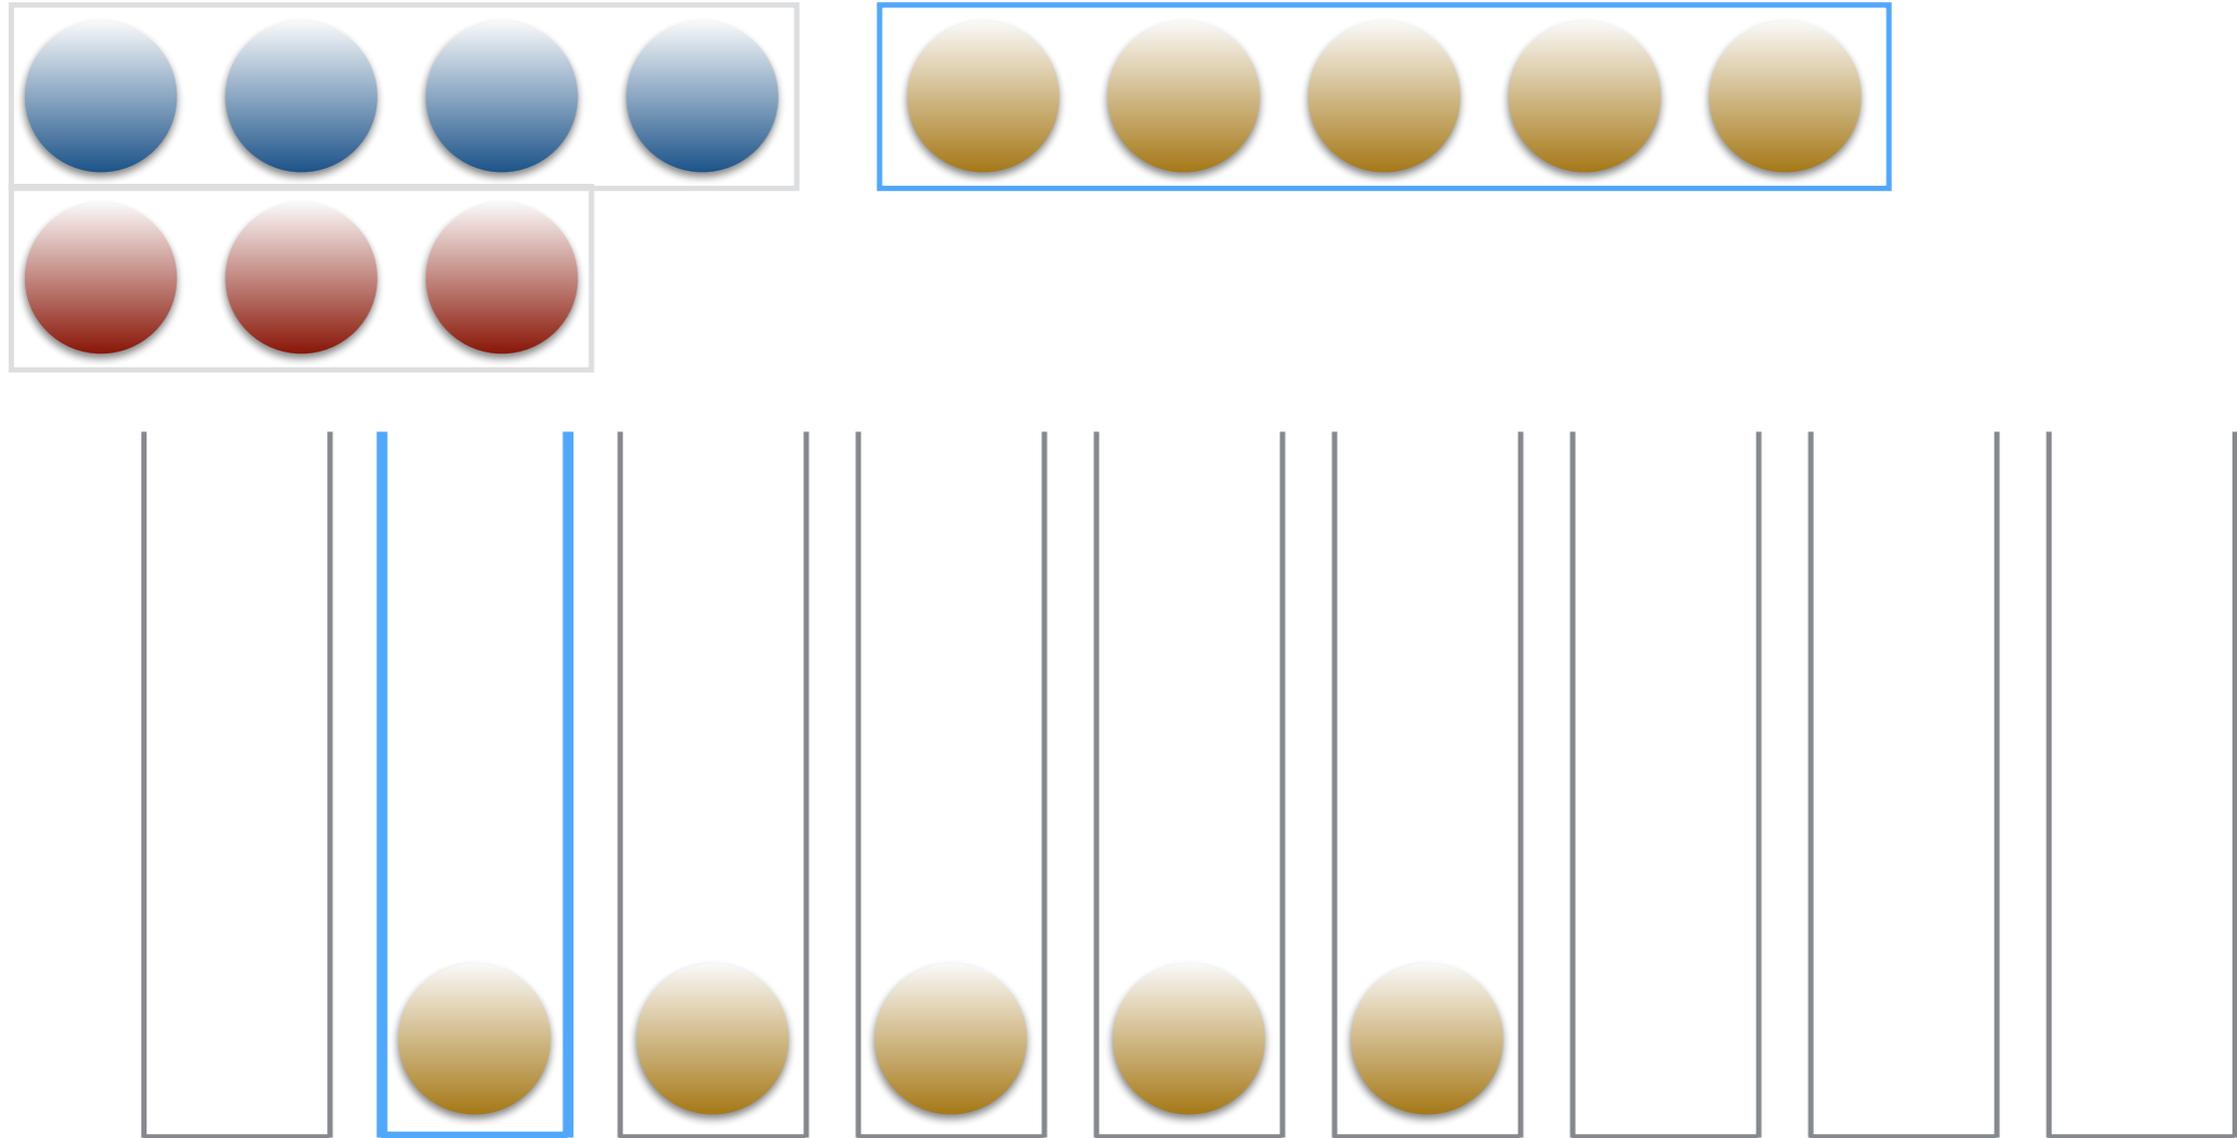
Two-Dimensional Allocation



Two-Dimensional Allocation

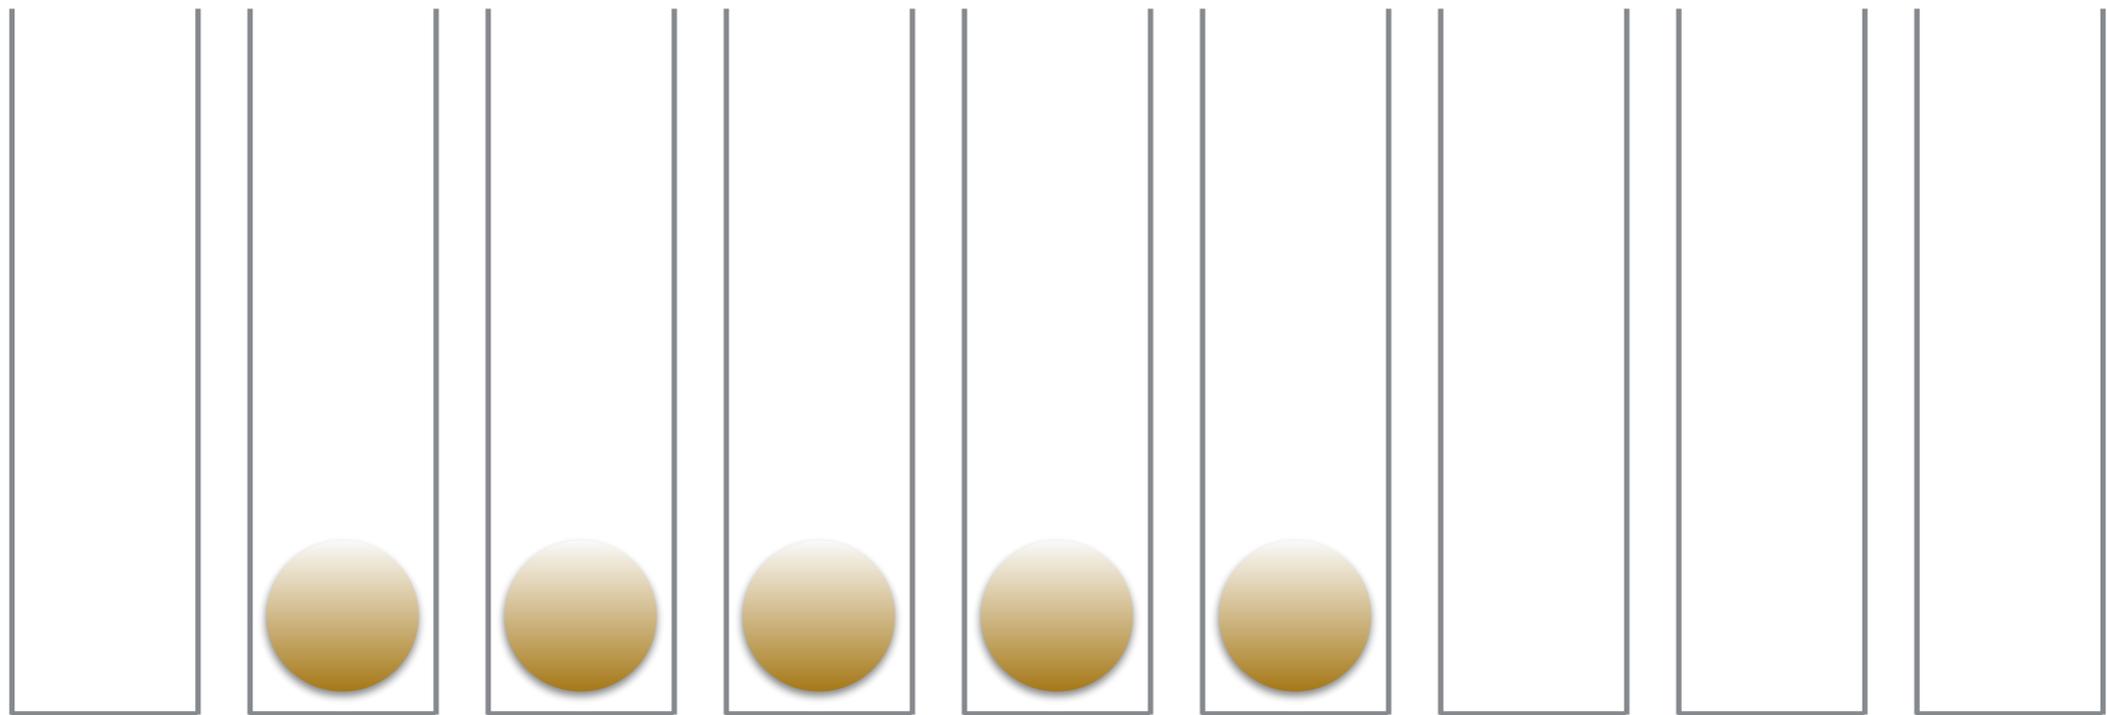
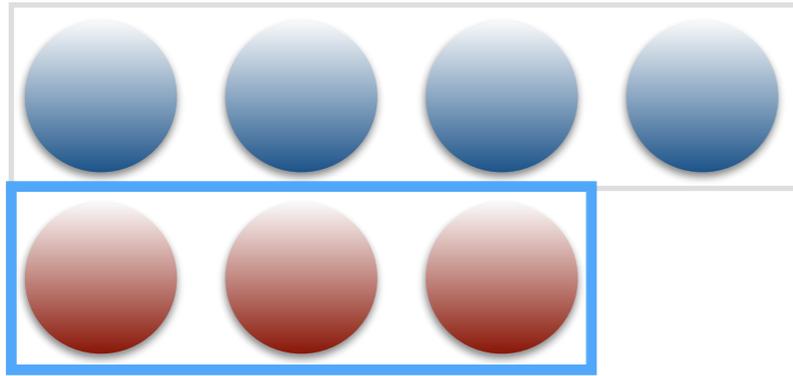


Two-Dimensional Allocation

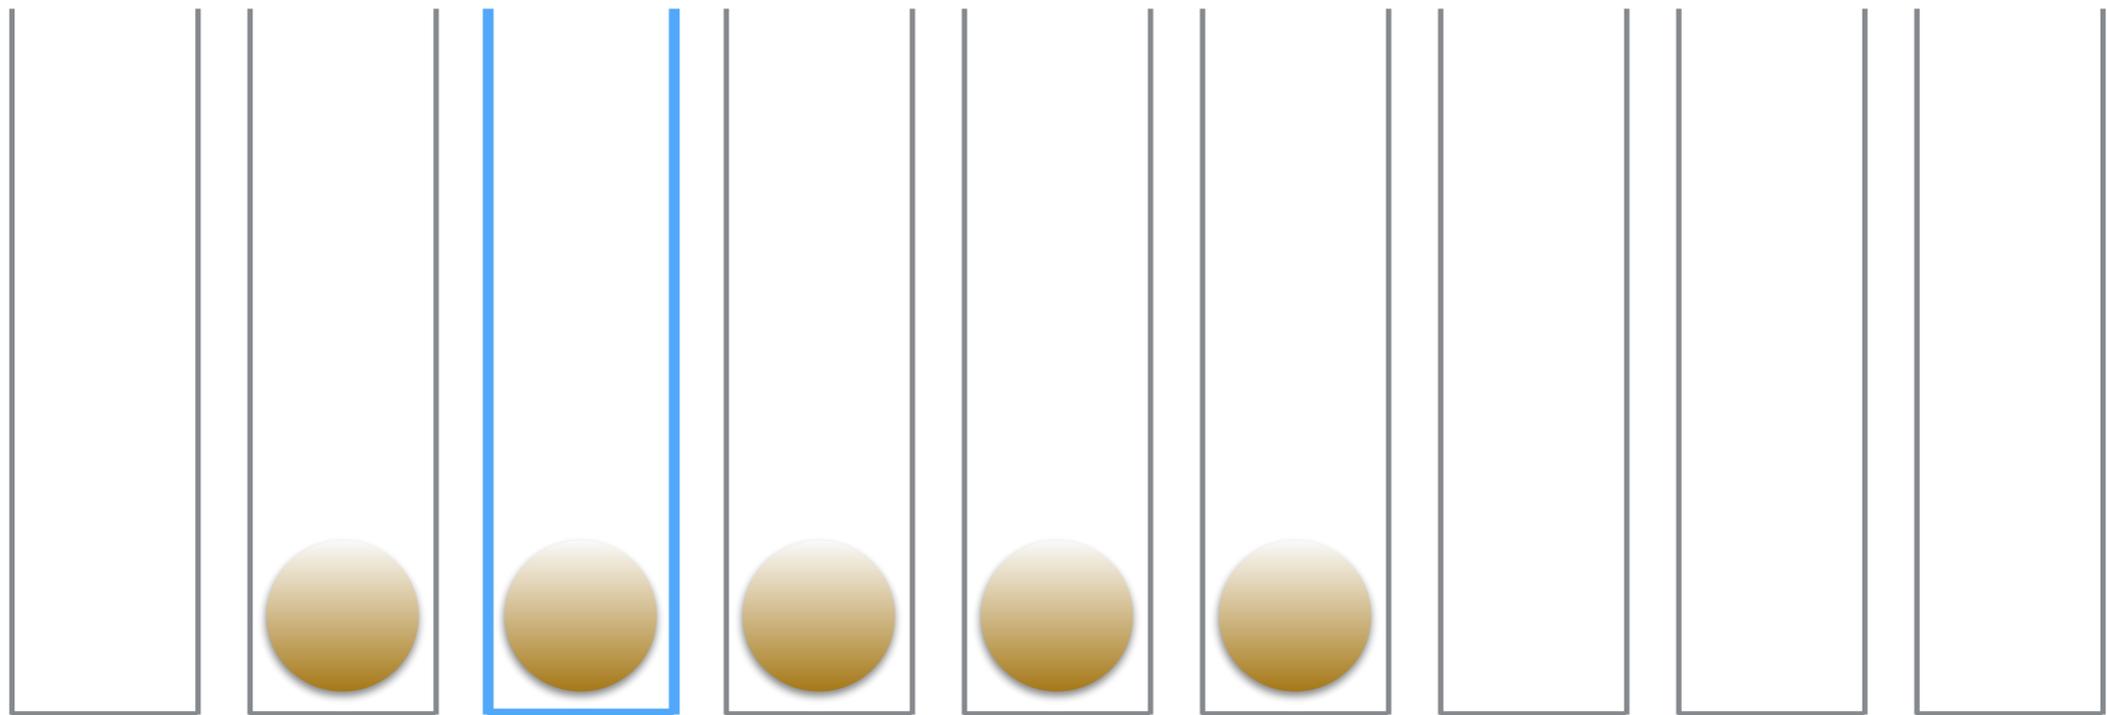
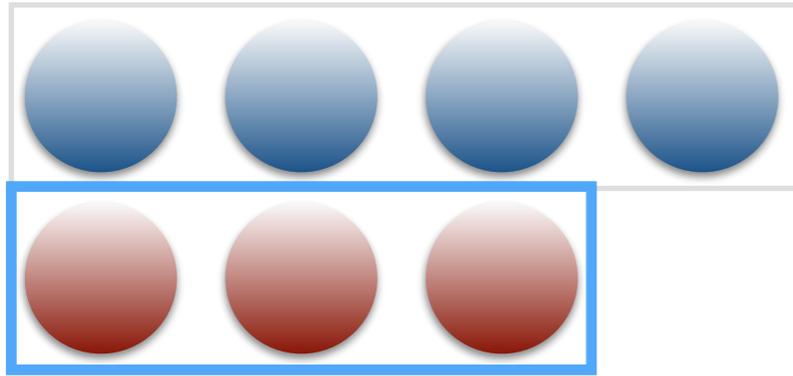


Place the whole list according to a ***single*** probabilistic choice!

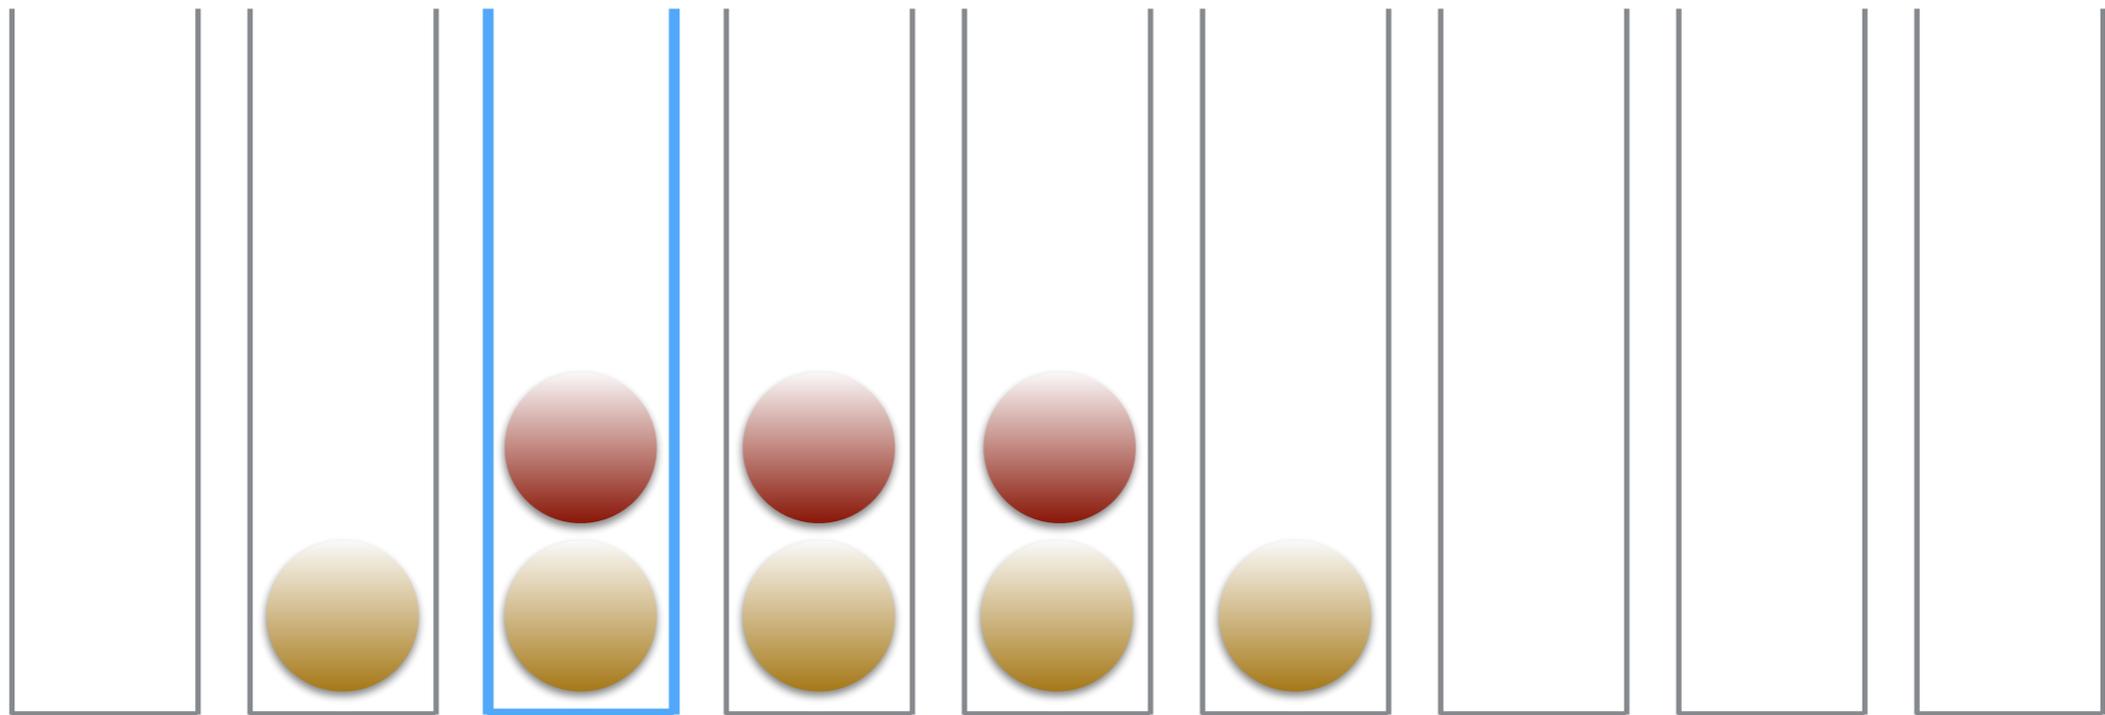
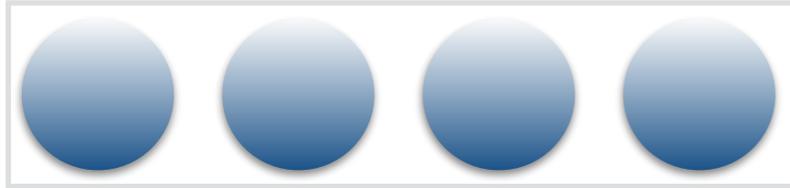
Two-Dimensional Allocation



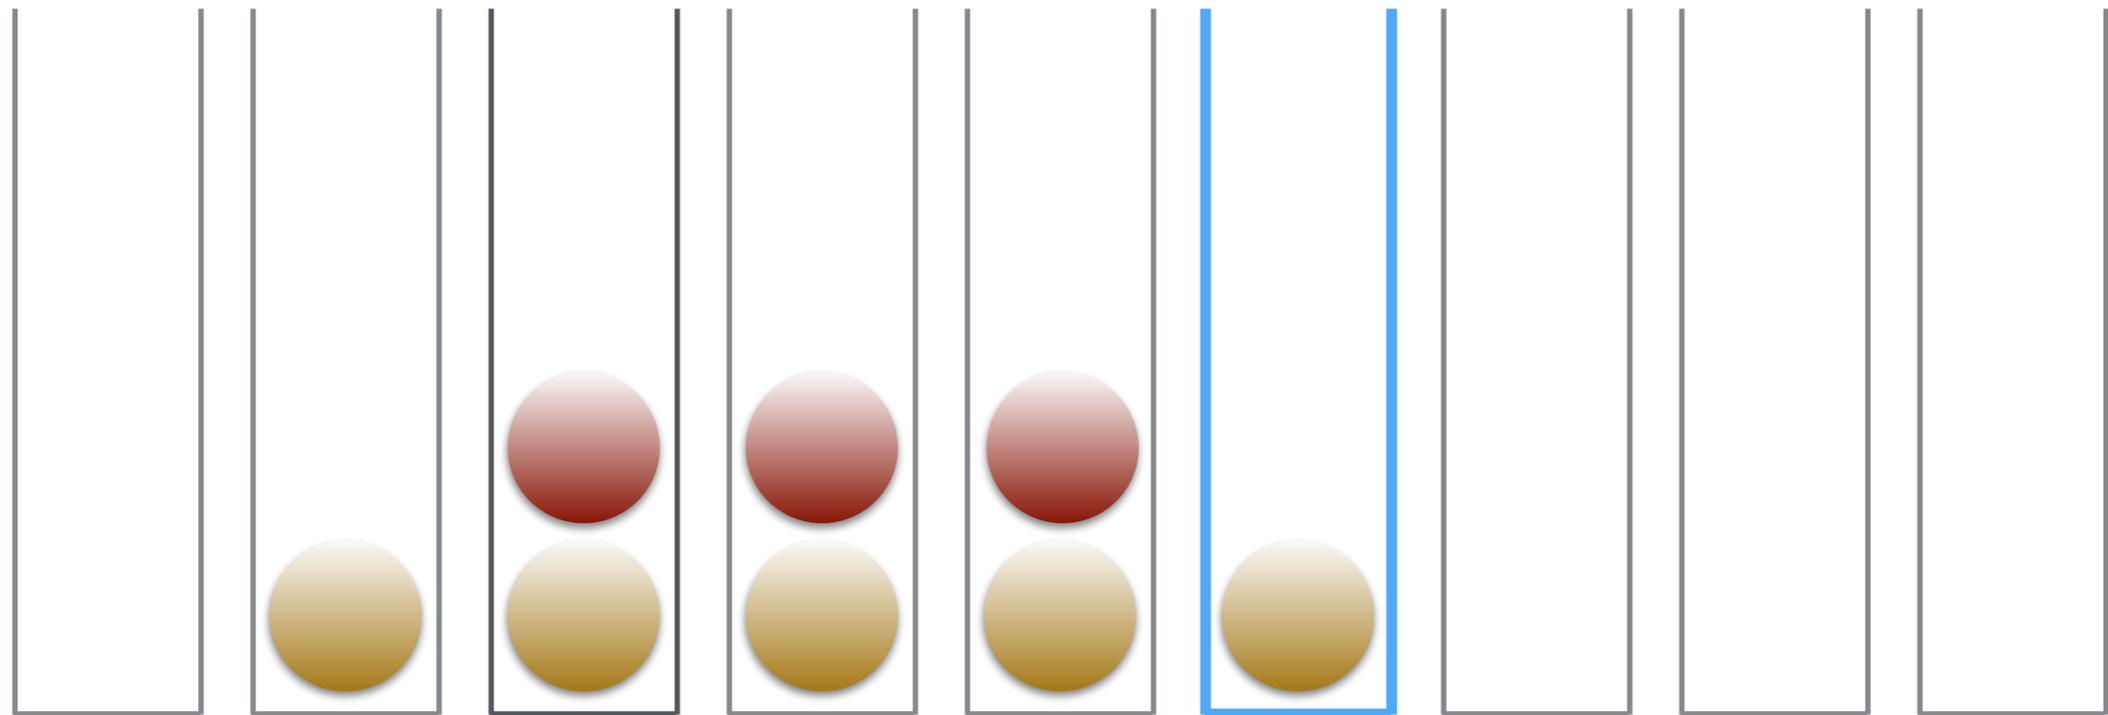
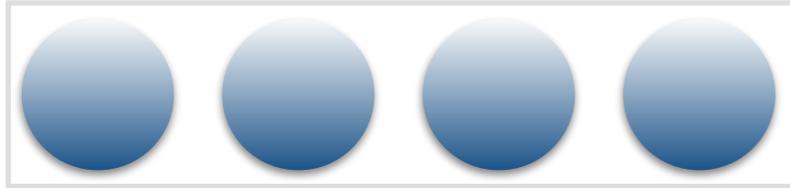
Two-Dimensional Allocation



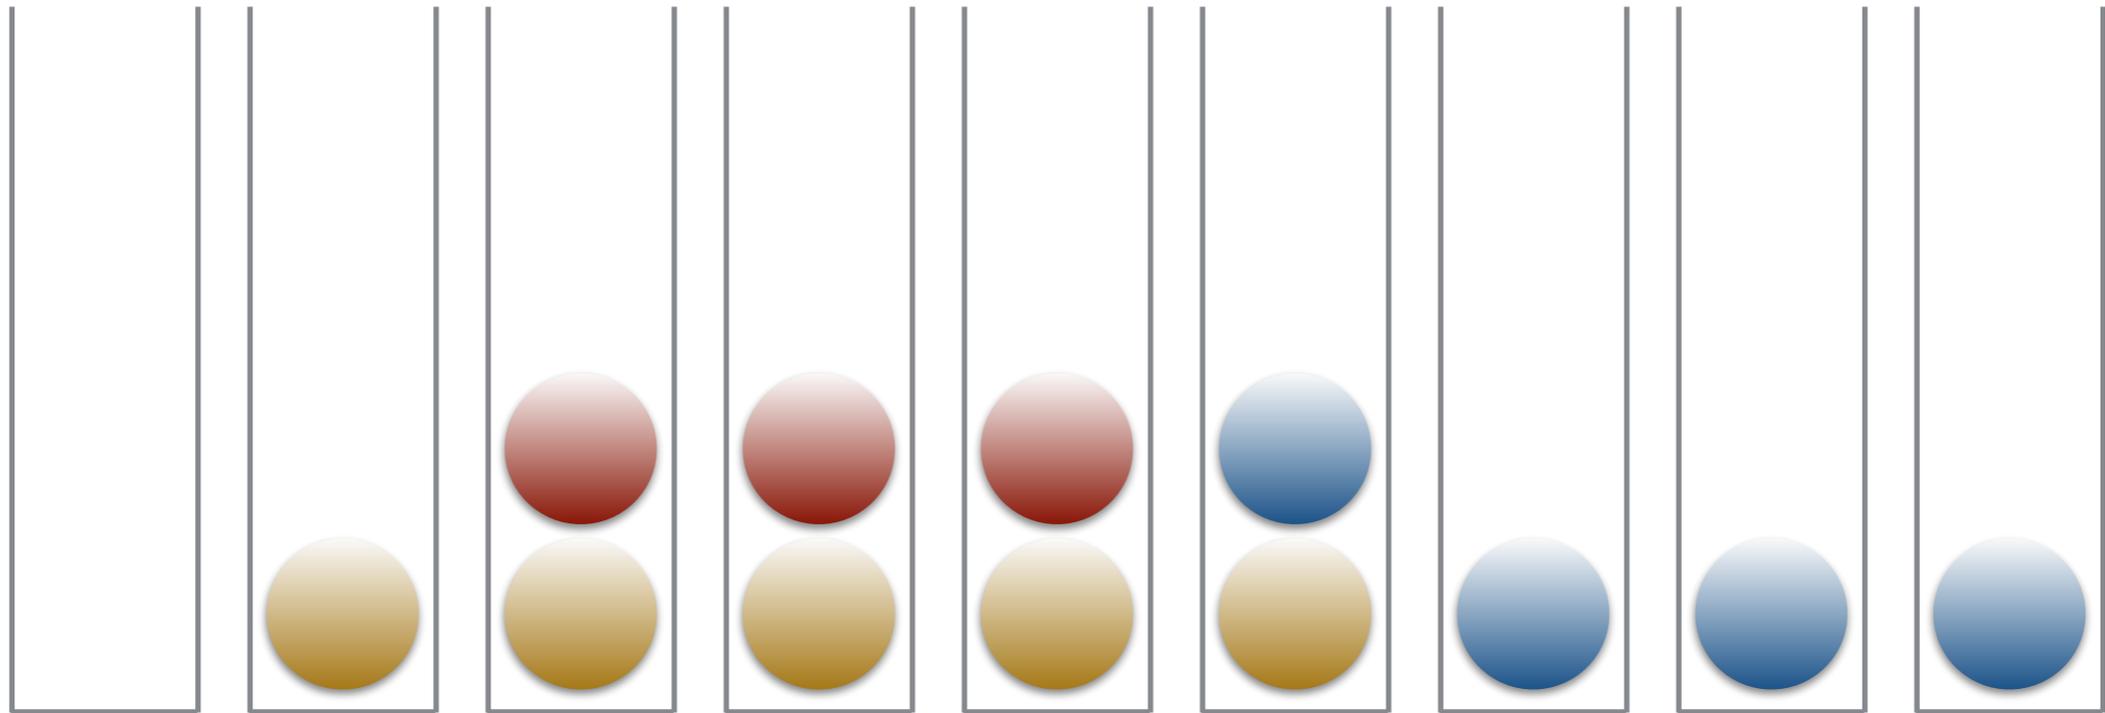
Two-Dimensional Allocation



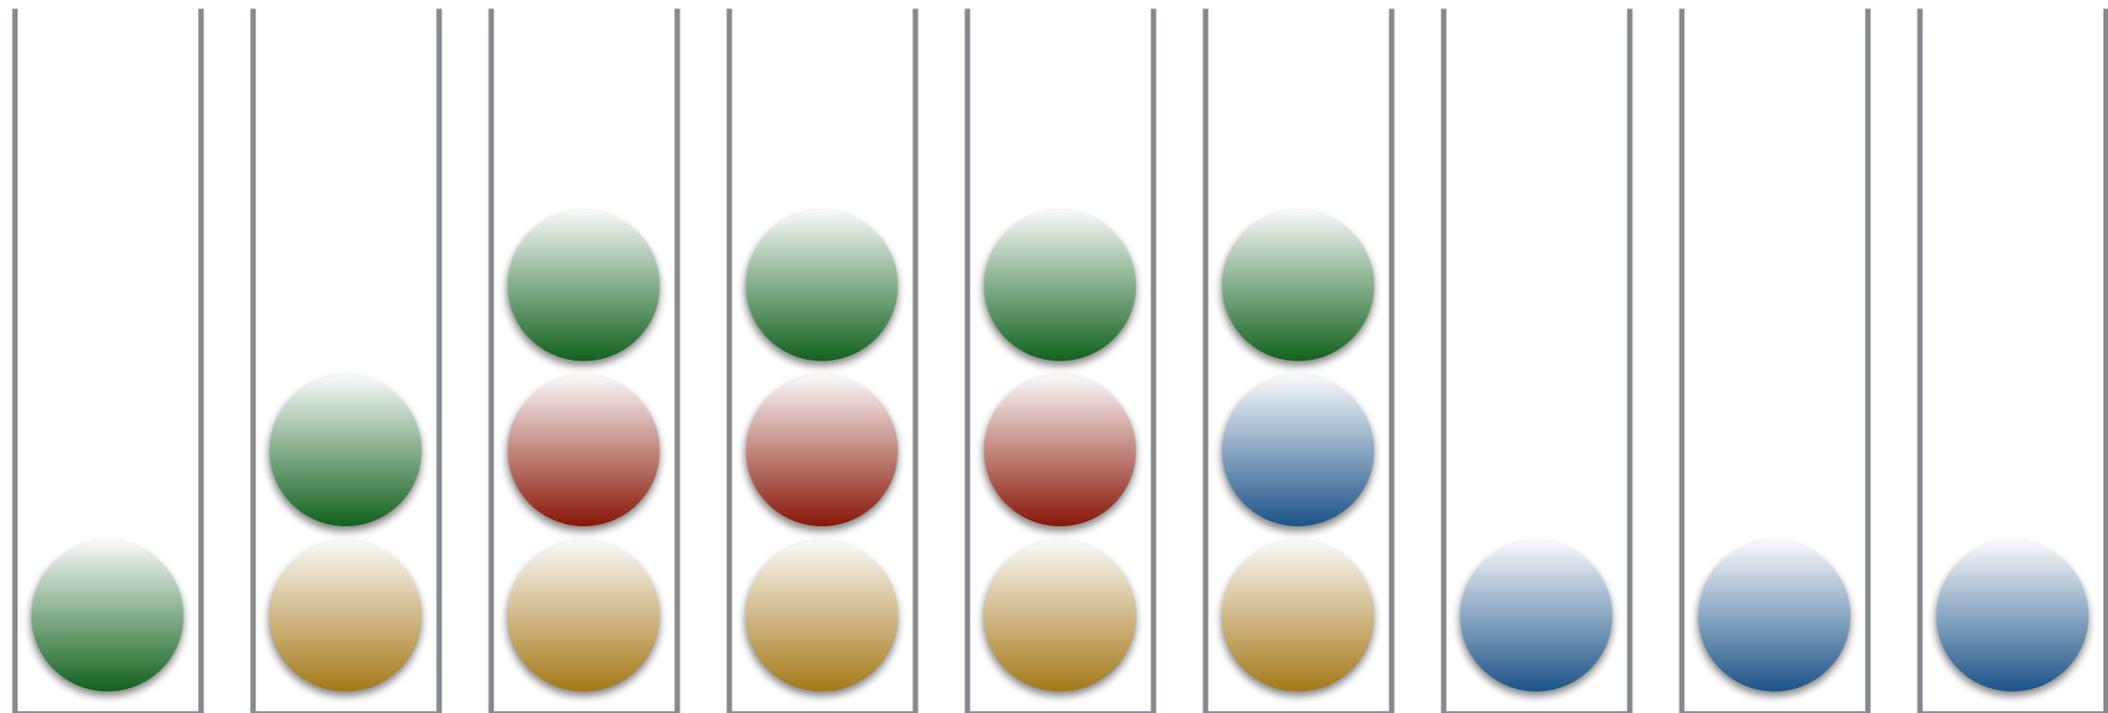
Two-Dimensional Allocation



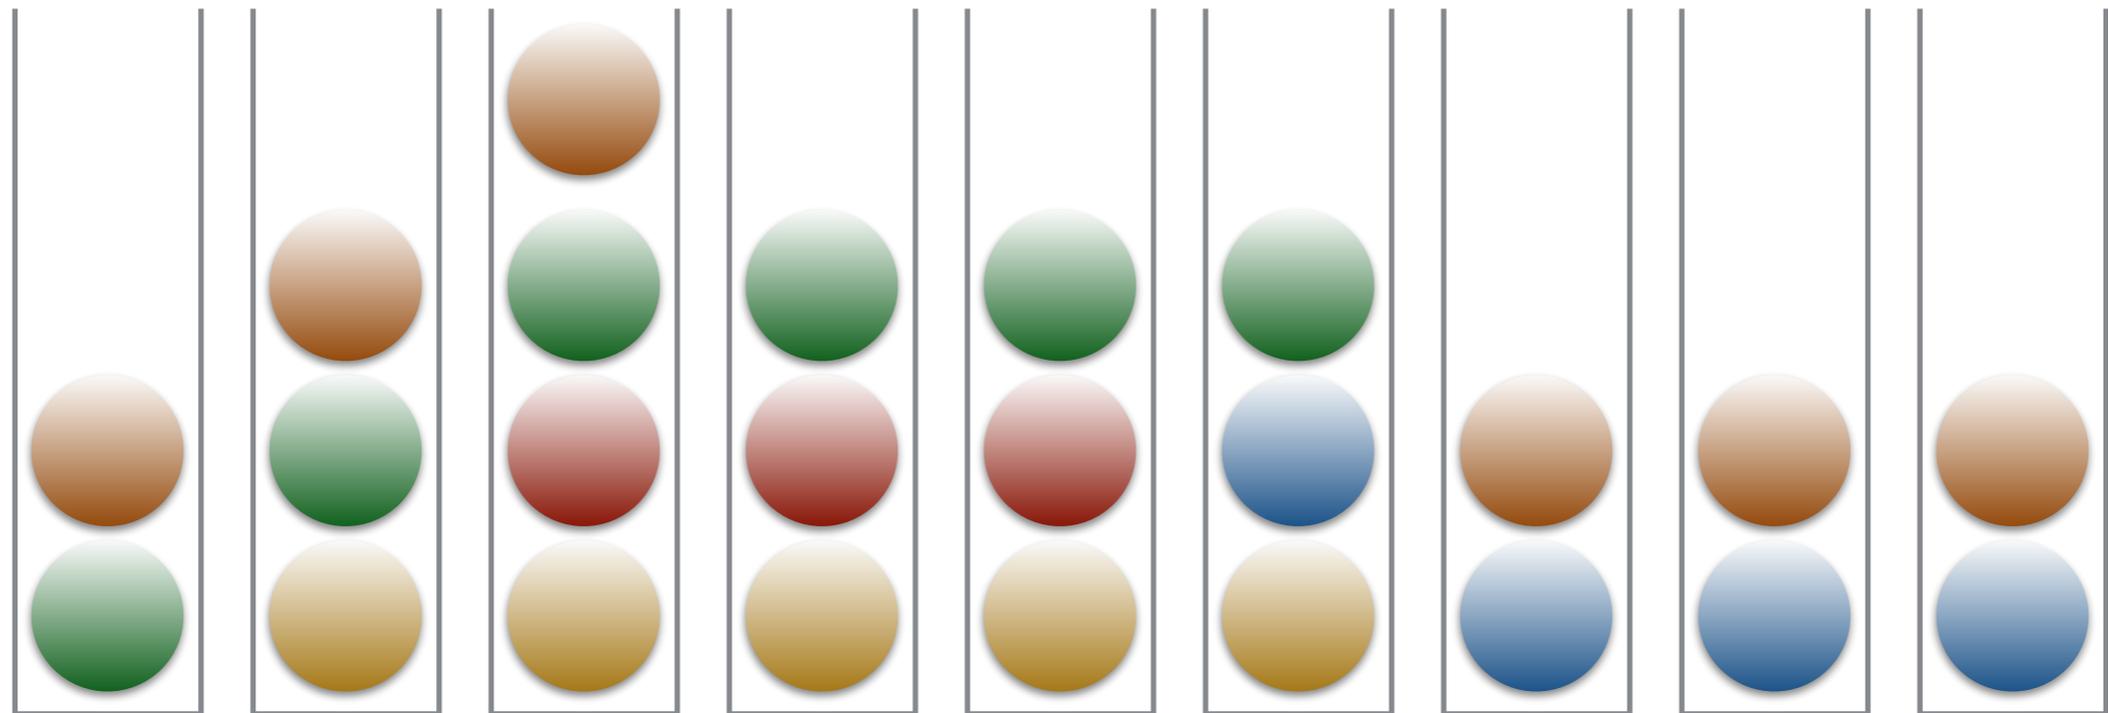
Two-Dimensional Allocation



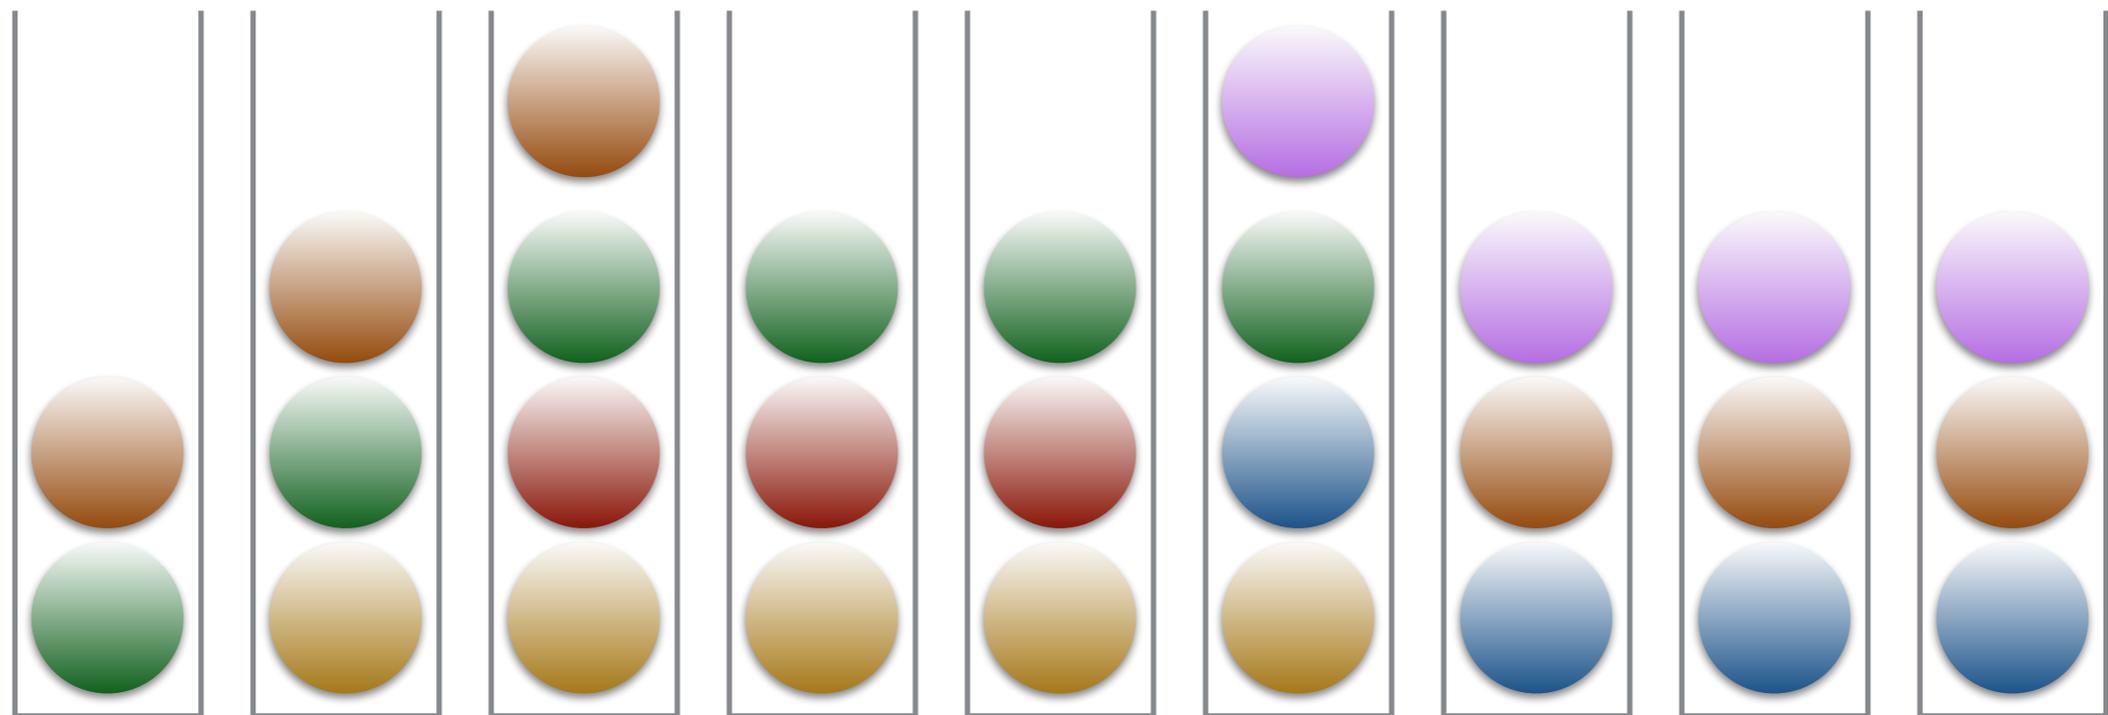
Two-Dimensional Allocation



Two-Dimensional Allocation



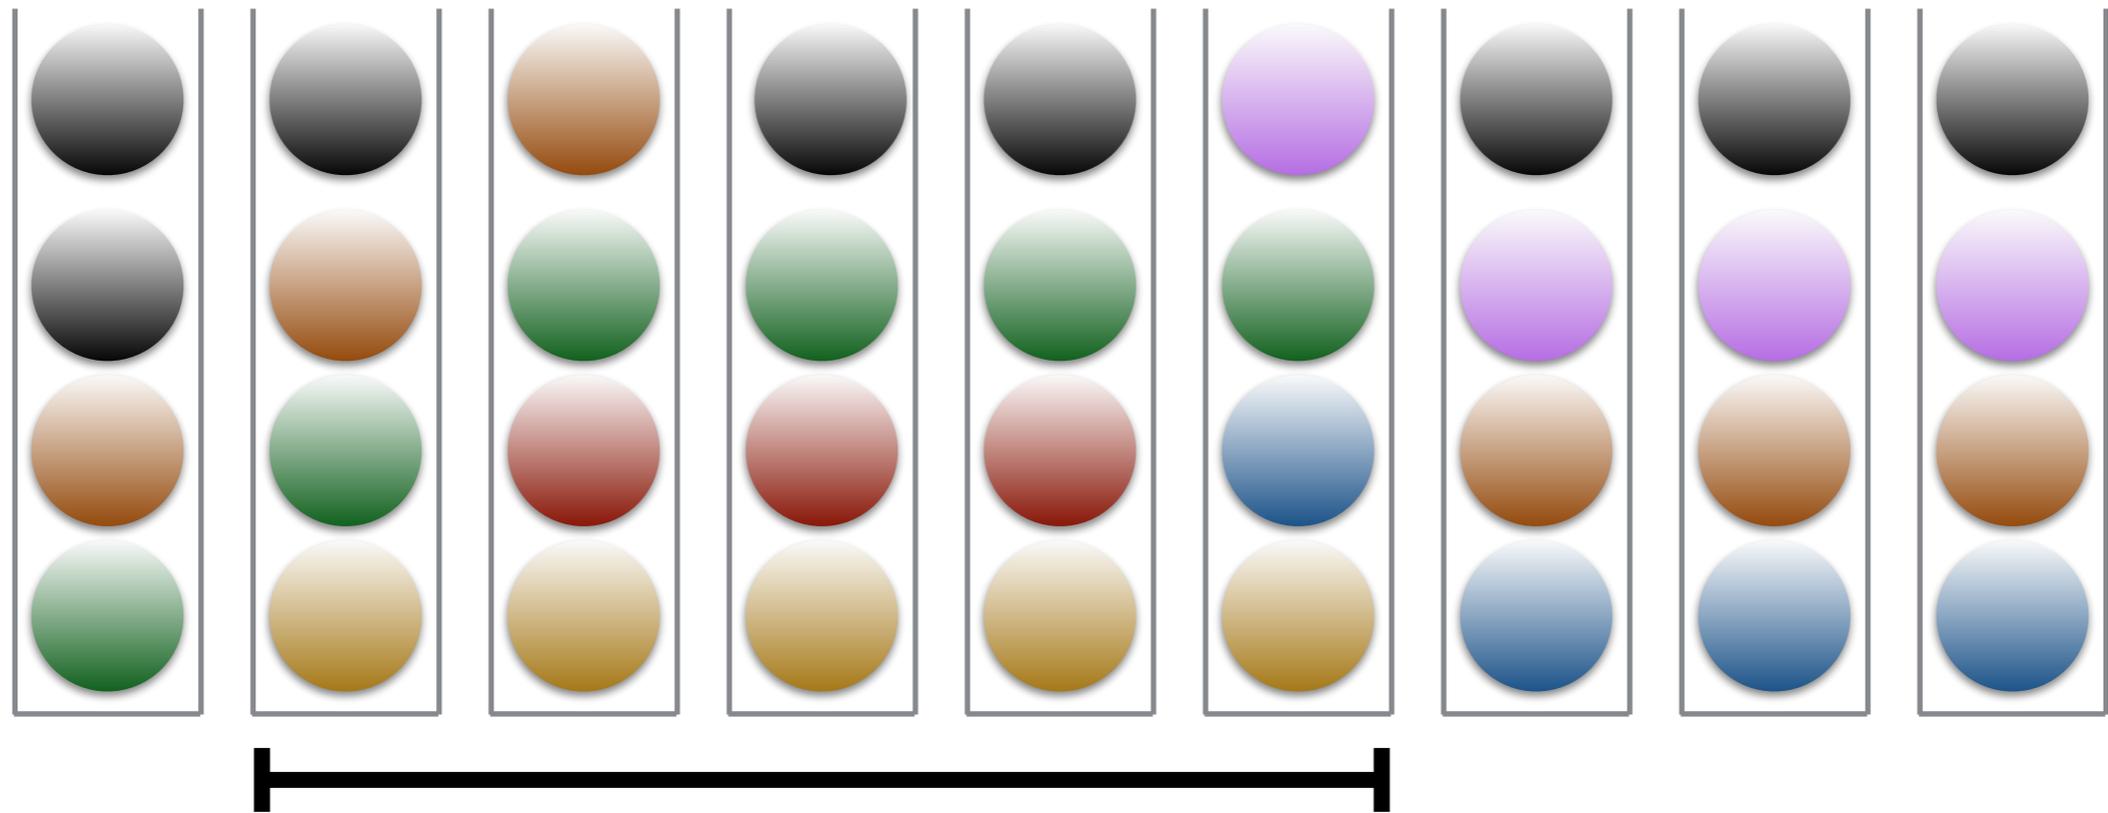
Two-Dimensional Allocation



What is the maximal load?

How Do We Search?

Search(●)



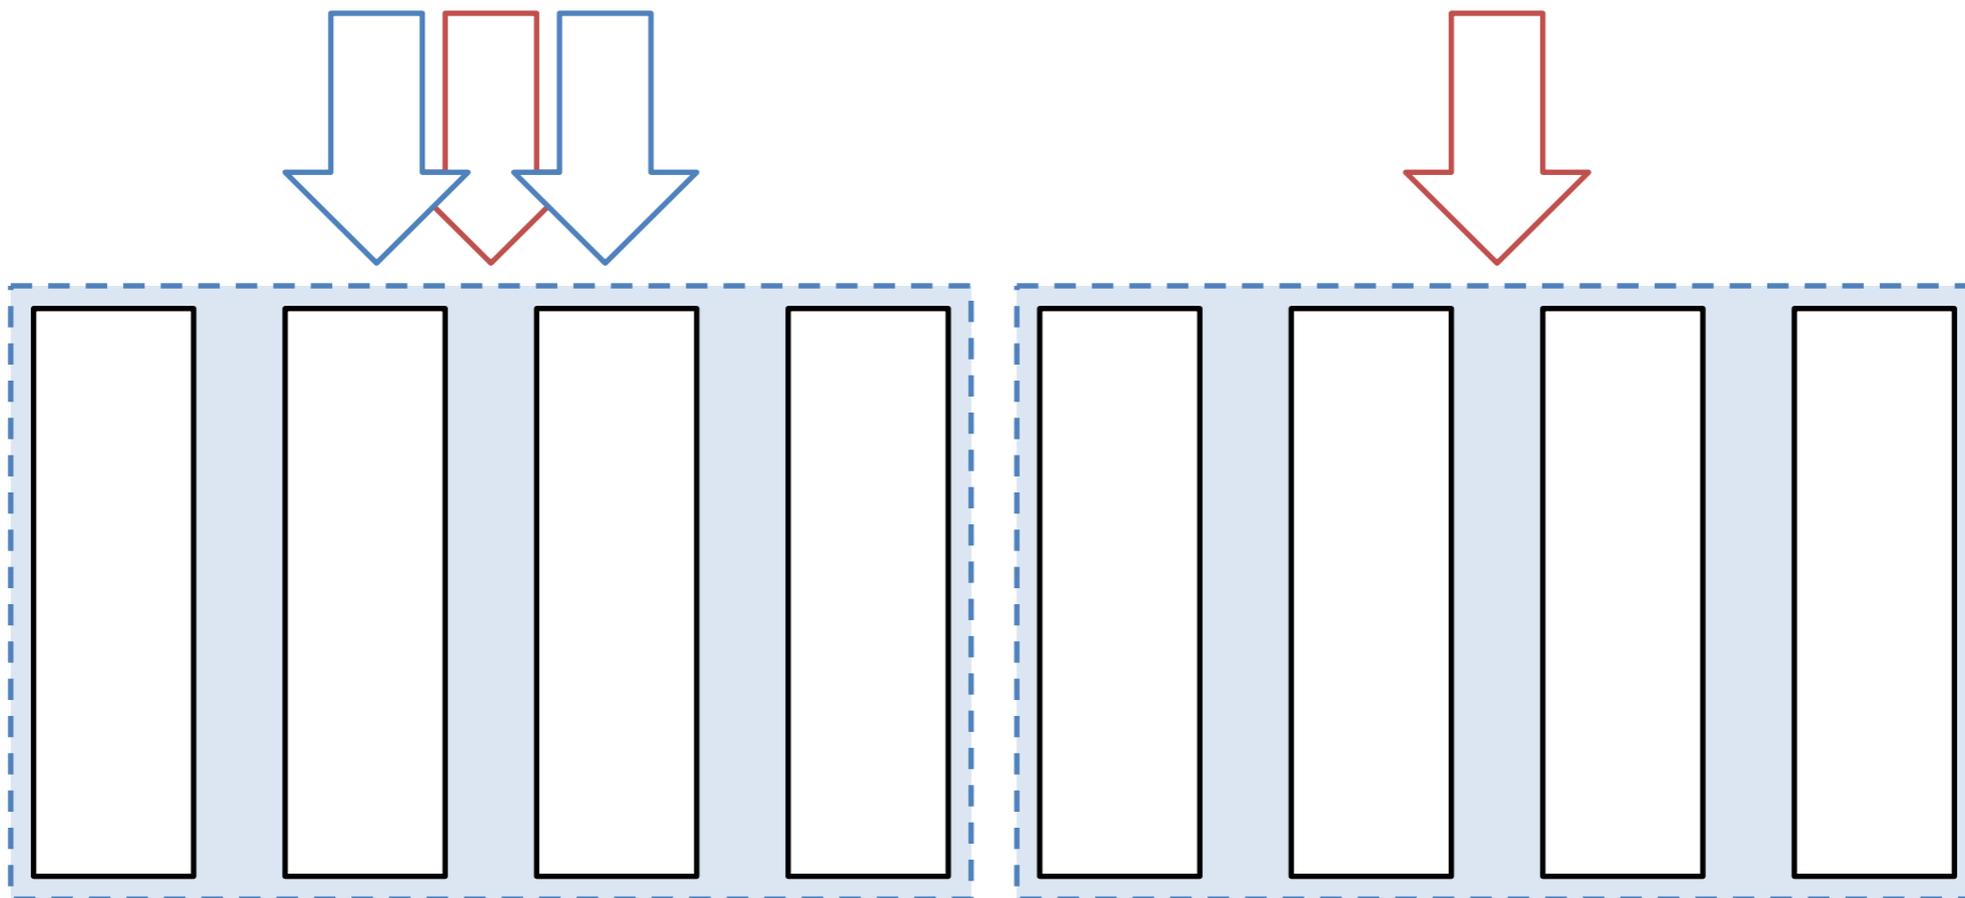
Our First Scheme: 2D Random Allocation

- **Theorem:** Set **#Bins = $N/O(\log N \log \log N)$** . Then, with an overwhelming probability, the maximal load is **$3 \log N \log \log N$**
- **Main Challenge** (compared to 1D case):
Heavy dependencies between the elements of the same list
- **This yields an SSE scheme with:**
 - Space: **#Bins x BinSize = $O(N)$**
 - Locality: **$O(1)$**
 - Read efficiency: **$\tilde{O}(\log n)$**

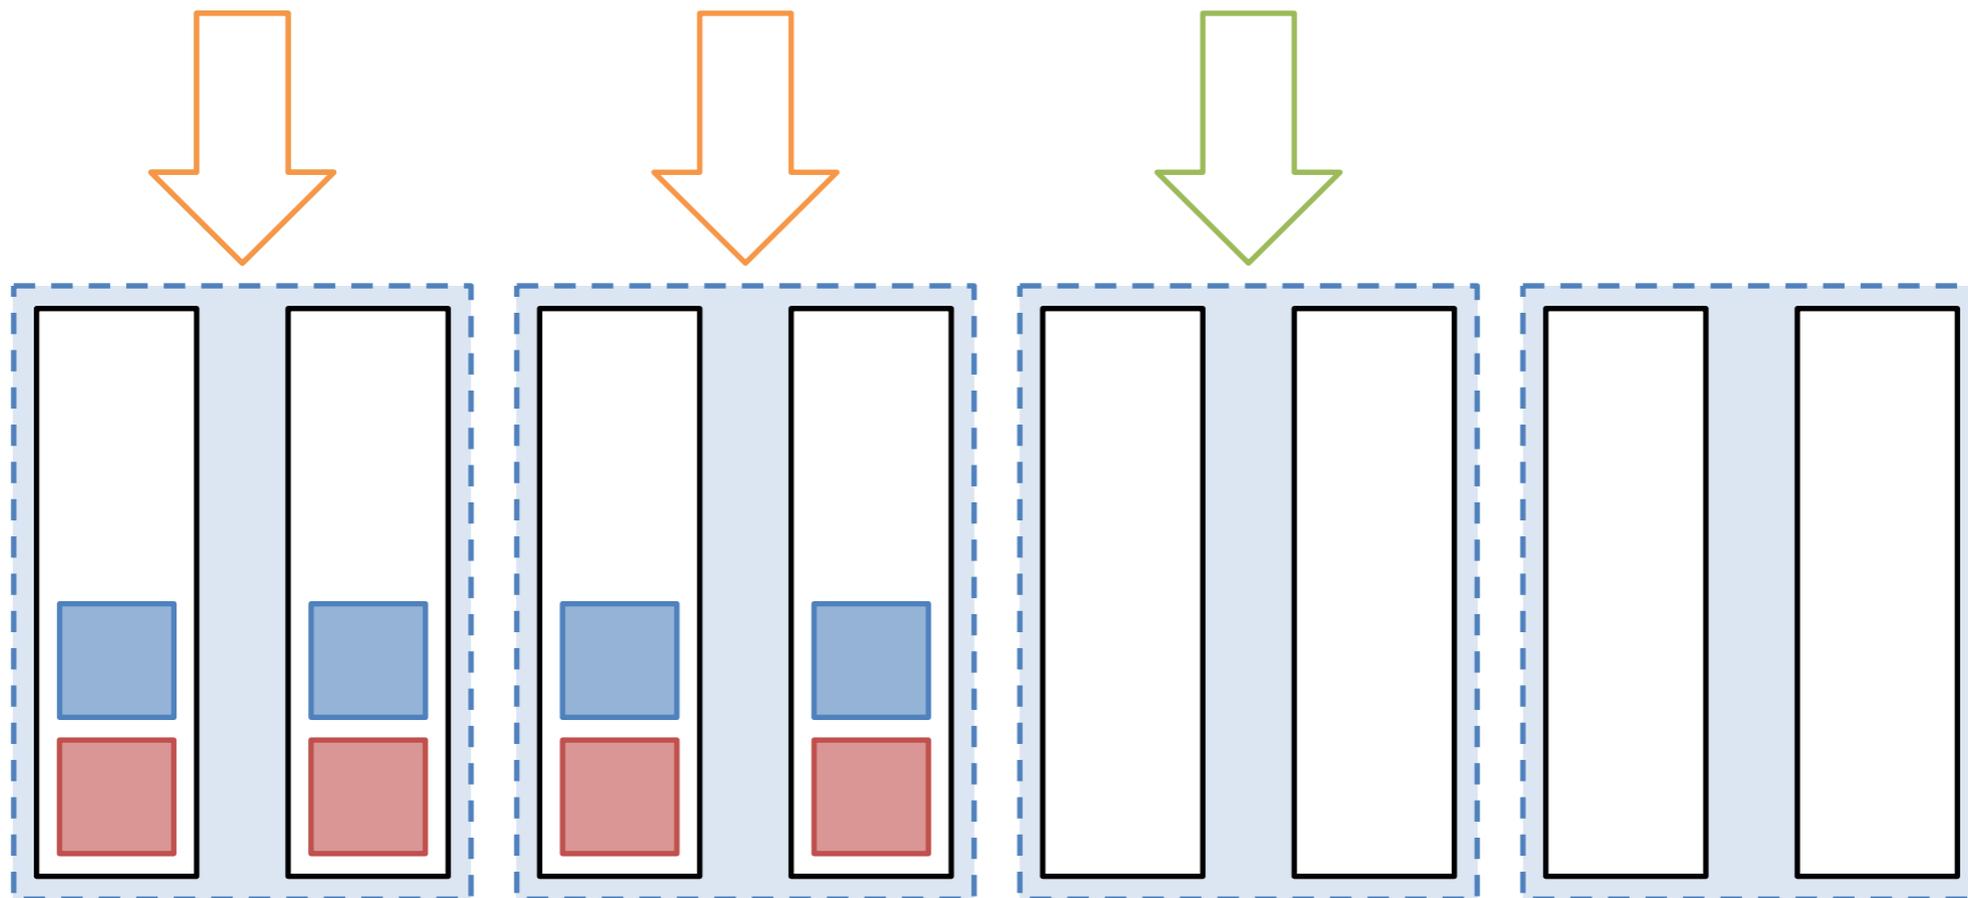
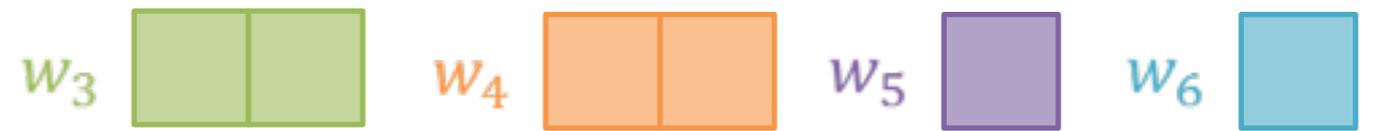
The Power of Two Choices

- In the classic “balls and bins” [ABKU99]:
 - If we choose **one** random bin for each ball, then the maximal load is $O(\log N / \log \log N)$
 - If we choose **two** random bins for each ball, and place the ball in the least loaded one, then the maximal load is $O(\log \log N)$
 - Exponential improvement!
- Can we adapt the two-choice paradigm to the 2D case?

2D Two-Choice Allocation

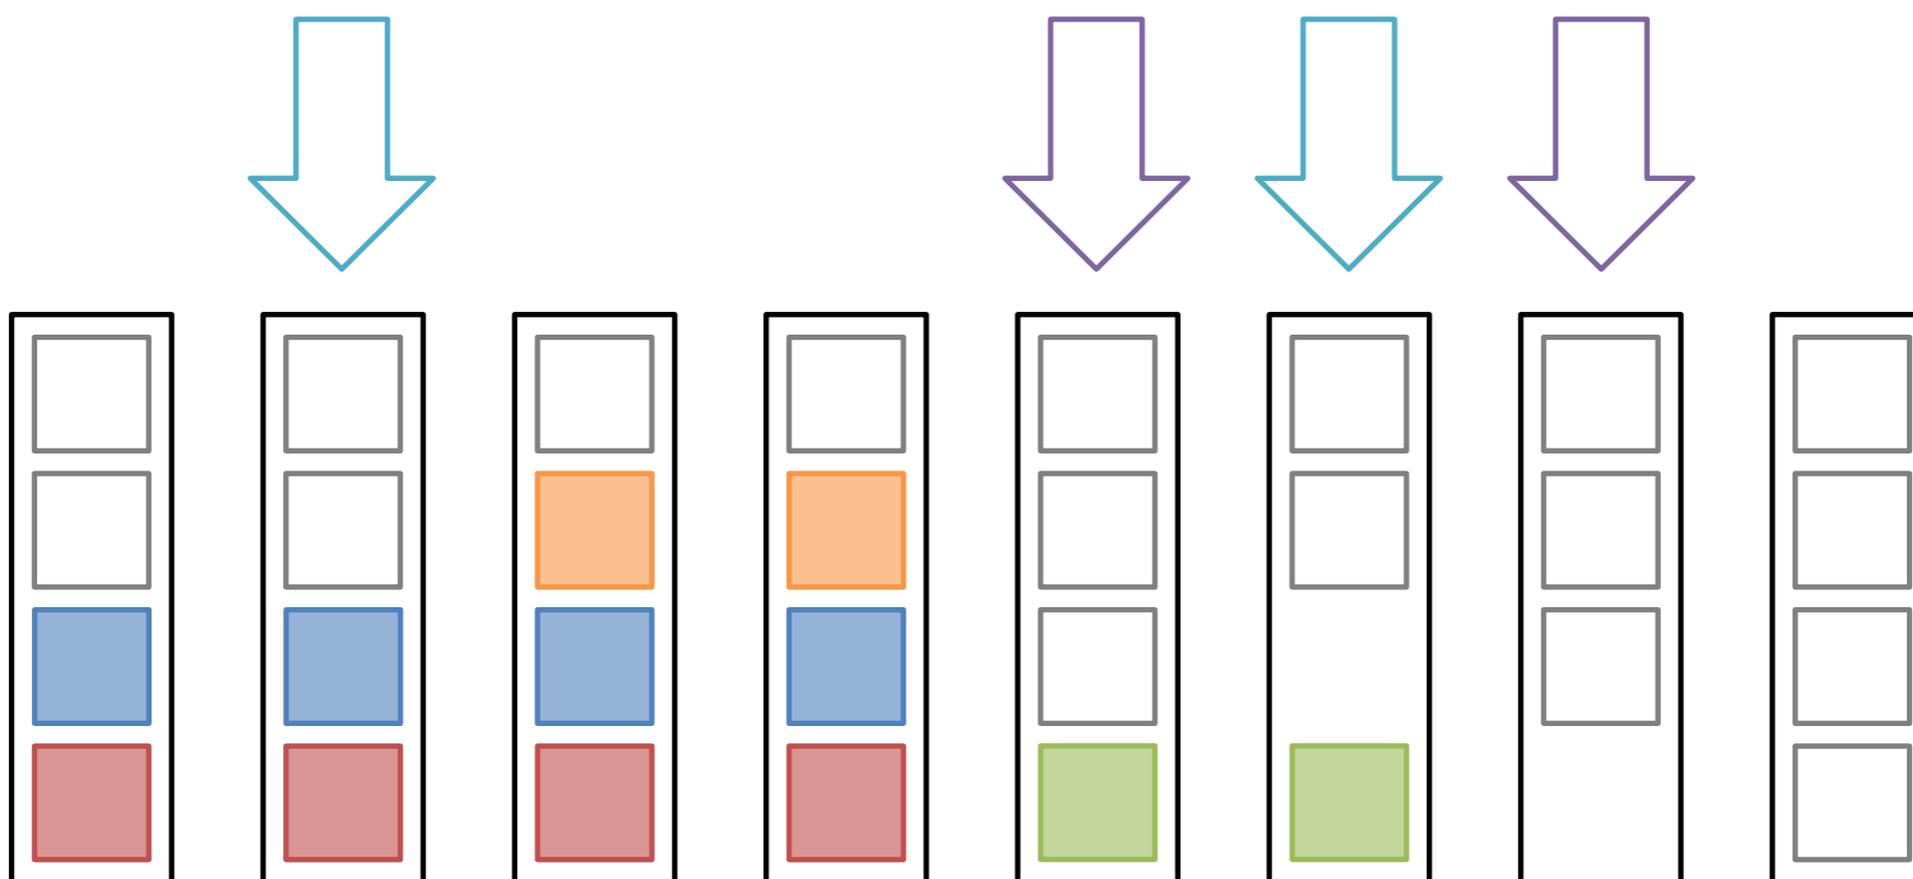


2D Two-Choice Allocation

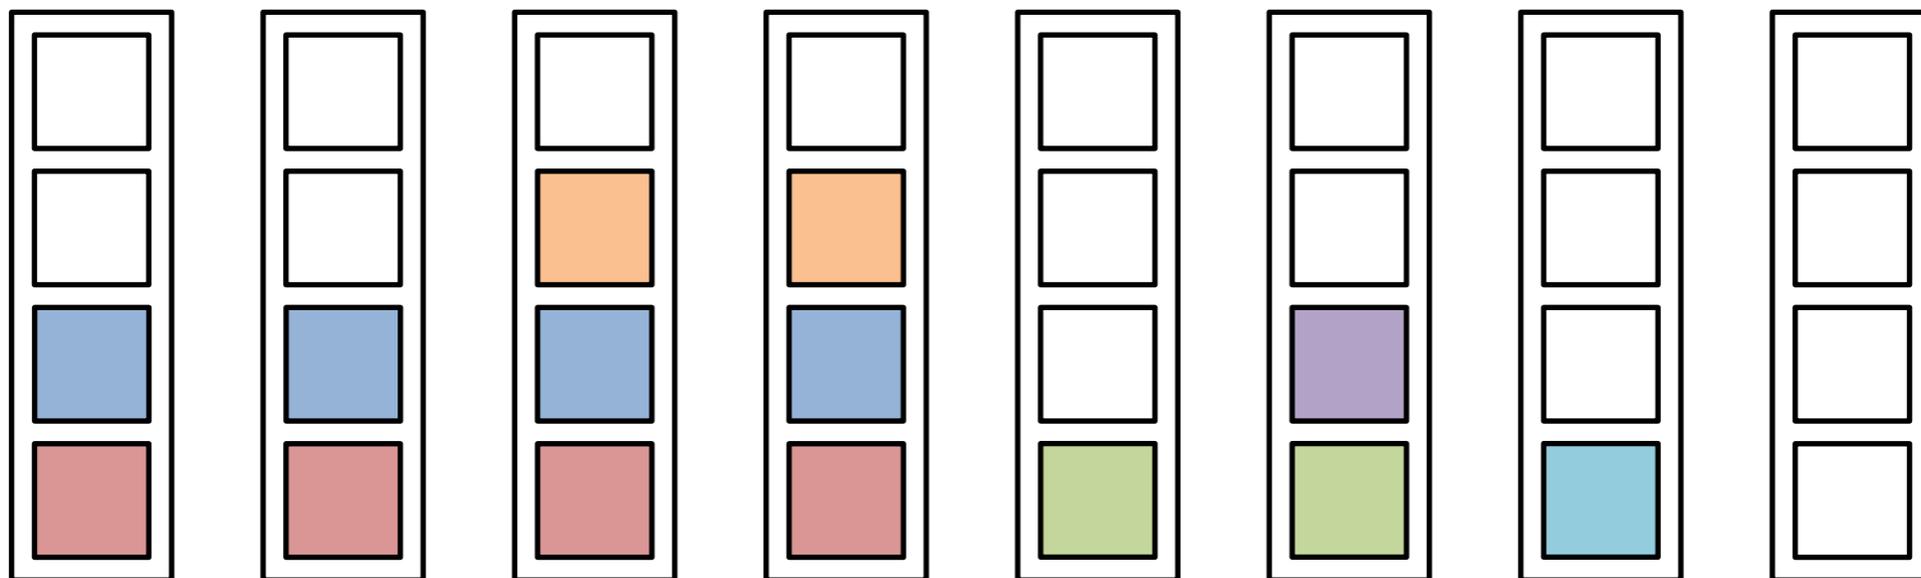


2D Two-Choice Allocation

w_5  w_6 



2D Two-Choice Allocation



2D Two-Choice Allocation

Theorem: Assume all lists are of length at most $N^{1-1/\log\log N}$,
and set **#Bins = $N/(\log\log N (\log\log\log N)^2)$** .
Then, with an overwhelming probability, the maximal load is
 $O(\log\log N (\log\log\log N)^2)$

- **Main Challenge:** (compared to 1D case):
 - Many challenges...
- This yields an SSE scheme with:
 - Space: **#Bins x BinSize = $O(N)$**
 - Read efficiency: **$2\text{BinSize} = \tilde{O}(\log\log N)$**
 - Locality: **$\tilde{O}(1)$**

Summary

- **Our approach:** SSE via two-dimensional balanced allocations

Scheme	Space	Locality	Read Efficiency
This work I	$O(N)$	$O(1)$	$\tilde{O}(\log N)$
This work II*	$O(N)$	$O(1)$	$\tilde{O}(\log \log N)$
This work III	$O(N \log N)$	$O(1)$	$O(1)$

Nice combination between DS and Cryptography

Thank You!