ASSOCIATION ANALYSIS

♦ ♦ ♦

FREQUENT ITEMSETS MINING

Alexandre Termier, LIG

M2 SIF – DMV course – 2017/2018

Market basket analysis

- Analyse supermarket's transaction data
- Transaction = « market basket » of a customer
- Find which items are often bought together
  - Ex: *{bread, chocolate, butter}*
  - Ex: *{hamburger bread, tomato}* $\rightarrow$ *{steak}*
- Applications
  - Product placement
  - Cross selling (suggestion of other products)
  - Promotions

Funny example

- Most famous itemset : *{beer, diapers}*

- Found in a chain of American supermarkets

- Further study :
  - Mostly bought on Friday evenings
  - Who ? …

Example

Transactions	Items (products bought)
1	bread, butter, chocolate, vine, pencil
2	bread,butter, chocolate, pencil
3	chocolate
4	butter,chocolate
5	bread, butter, chocolate, vine
6	bread,butter, chocolate

- *{bread, butter, chocolate}* sold together in 4/6 = 66% of transactions
- *{butter, chocolate}* → *{bread}* is true in 4/5 = 80% of cases
- *{chocolate}* → *{bread, butter}* is true in 4/6 = 66% of cases

Definitions

- $A = \{a_1, \ldots, a_n\}$: **items**, A : item base
- Any $I \subseteq A$: **itemset**
  - k-itemset: itemset with k items
- $T = (t_1, \ldots, t_m), \ \forall i \ t_i \subseteq A$: **transaction database**
- $tid(t_j) = j$: transaction index
- **support**($X \in I$) = number of transactions containing itemset X
- **tidlist**($X \in I$) = list of tids of transactions containing itemset X

An itemset X is **frequent** if *support(X)* ≥ *minsup*

- Confidence of association rule $X \rightarrow Y$: $c = \dfrac{\text{support}(X \cup Y)}{\text{support}(X)}$
  $(X \cap Y = \emptyset)$

An association rule with confidence *c* **holds** if *c* ≥ *minconf*

Example, rewritten

Transactions	Items (products bought)
1	bread, butter, chocolate, vine, pencil
2	bread,butter, chocolate, pencil
3	chocolate
4	butter,chocolate
5	bread, butter, chocolate, vine
6	bread,butter, chocolate

- *{bread, butter, chocolate}*
  support = *(absolute)*
  = *(relative)*
- *{chocolate}* $\rightarrow$ *{bread, butter}*
  confidence =
- *{butter, chocolate}* $\rightarrow$ *{bread}*
  confidence =

Computing association rules

- Two steps:
  1. Compute frequent itemsets
     - Discover itemsets with support ≥ *minsup*
     - **Very expensive computationally !**
  2. Compute which association rules hold
     - Partition each itemset and discover rules with confidence ≥ *minconf*
     - Much faster than discovering itemsets

How to compute frequent itemsets ?

- Brute force approach
  - **Generate and Test** method
  - Generate all possible itemsets randomly
  - Compute their support
  - But highly combinatorial problem :
    - How many possible itemsets for 1000 items ?
  - *Infeasible in practice*

The Apriori algorithm

[Agrawal *et al.*, 93]

- Levelwise search
  - Discover frequent 1-itemsets, 2-itemsets,…

Apriori property *:*

If an itemset is not frequent, then all its supersets are not frequent

  - Ex: If *{vine, pencil}* is not frequent, then of course *{vine, pencil, chocolate}* will not be frequent
  - *Downward closure property*
  - *Anti-monotonicity property*

	A	B	C	D	E
1	X	X	X	X	X
2	X	X	X		X
3			X		
4		X	X		
5	X	X	X	X	
6	X	X	X		

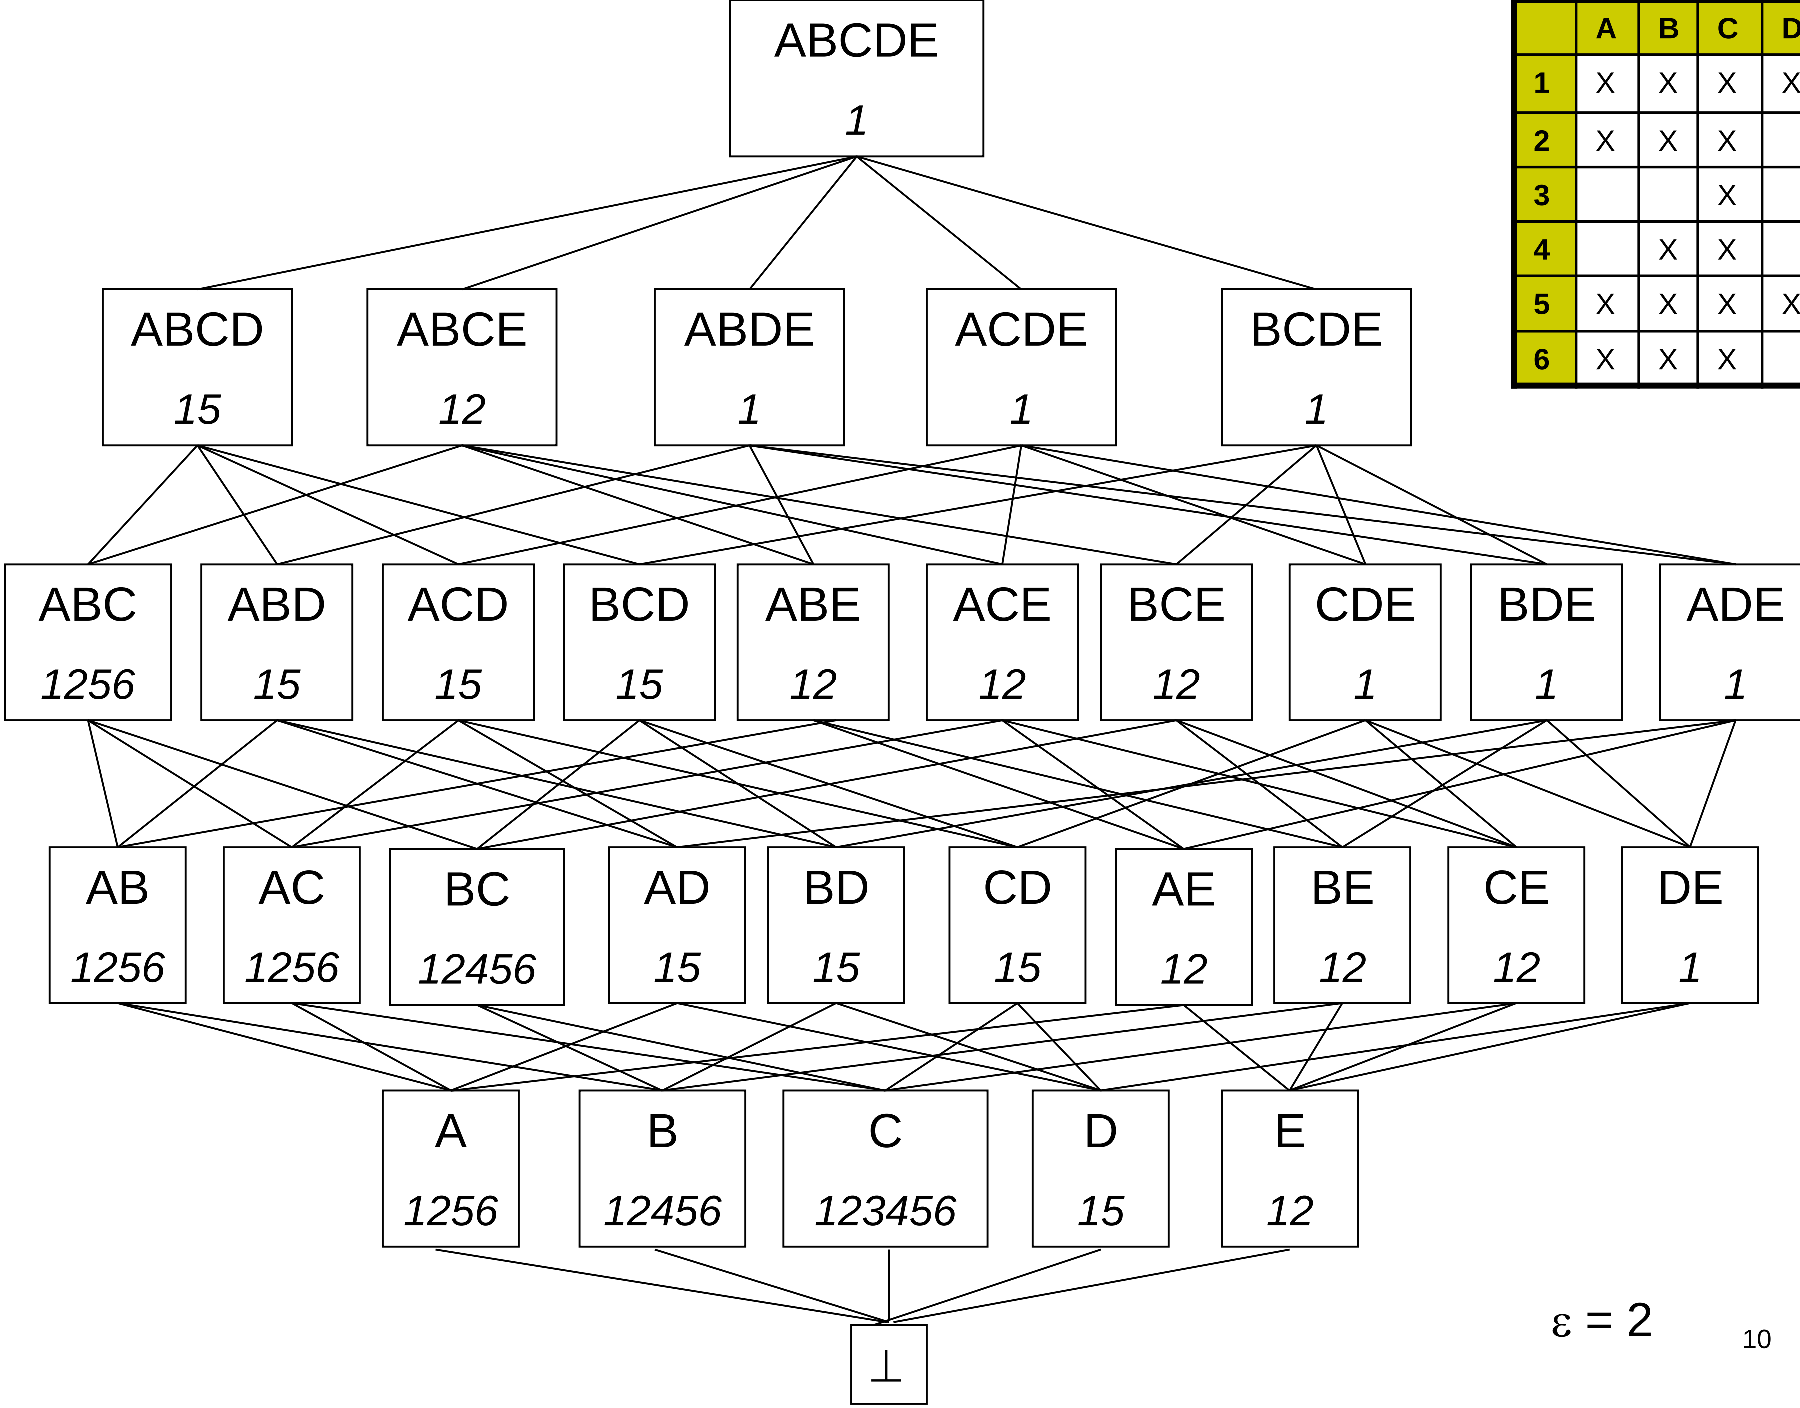

$\varepsilon = 2$

Apriori algorithm

```
Input: T, minsup

F₁ = {Frequent 1-itemsets} ;
for (k=2 ; F_{k-1} ≠ ∅ ; k++) do begin
   C_k = apriori-gen(F_{k-1}) ; // Candidates generation
   foreach transaction t ∈ T do begin
       C_t = subset(C_k, t) ; // Support counting
       foreach candidate c ∈ C_t do
           c.count++ ;
   end
   F_k = { c ∈ C_k | c.count ≥ minsup } ;
end
return ∪_k F_k ;
```

Candidate generation

- `apriori-gen`: generates candidates k-itemsets from frequent $(k\text{-}1)$-itemsets
- c (size k) = merge of $p, q \in F_{k-1}$ (both have size k-1)

- How many combinations of such p,q to build c ?

	A	B	C	D	E
1	X	X	X	X	X
2	X	X	X		X
3			X		
4		X	X		
5	X	X	X	X	
6	X	X	X		

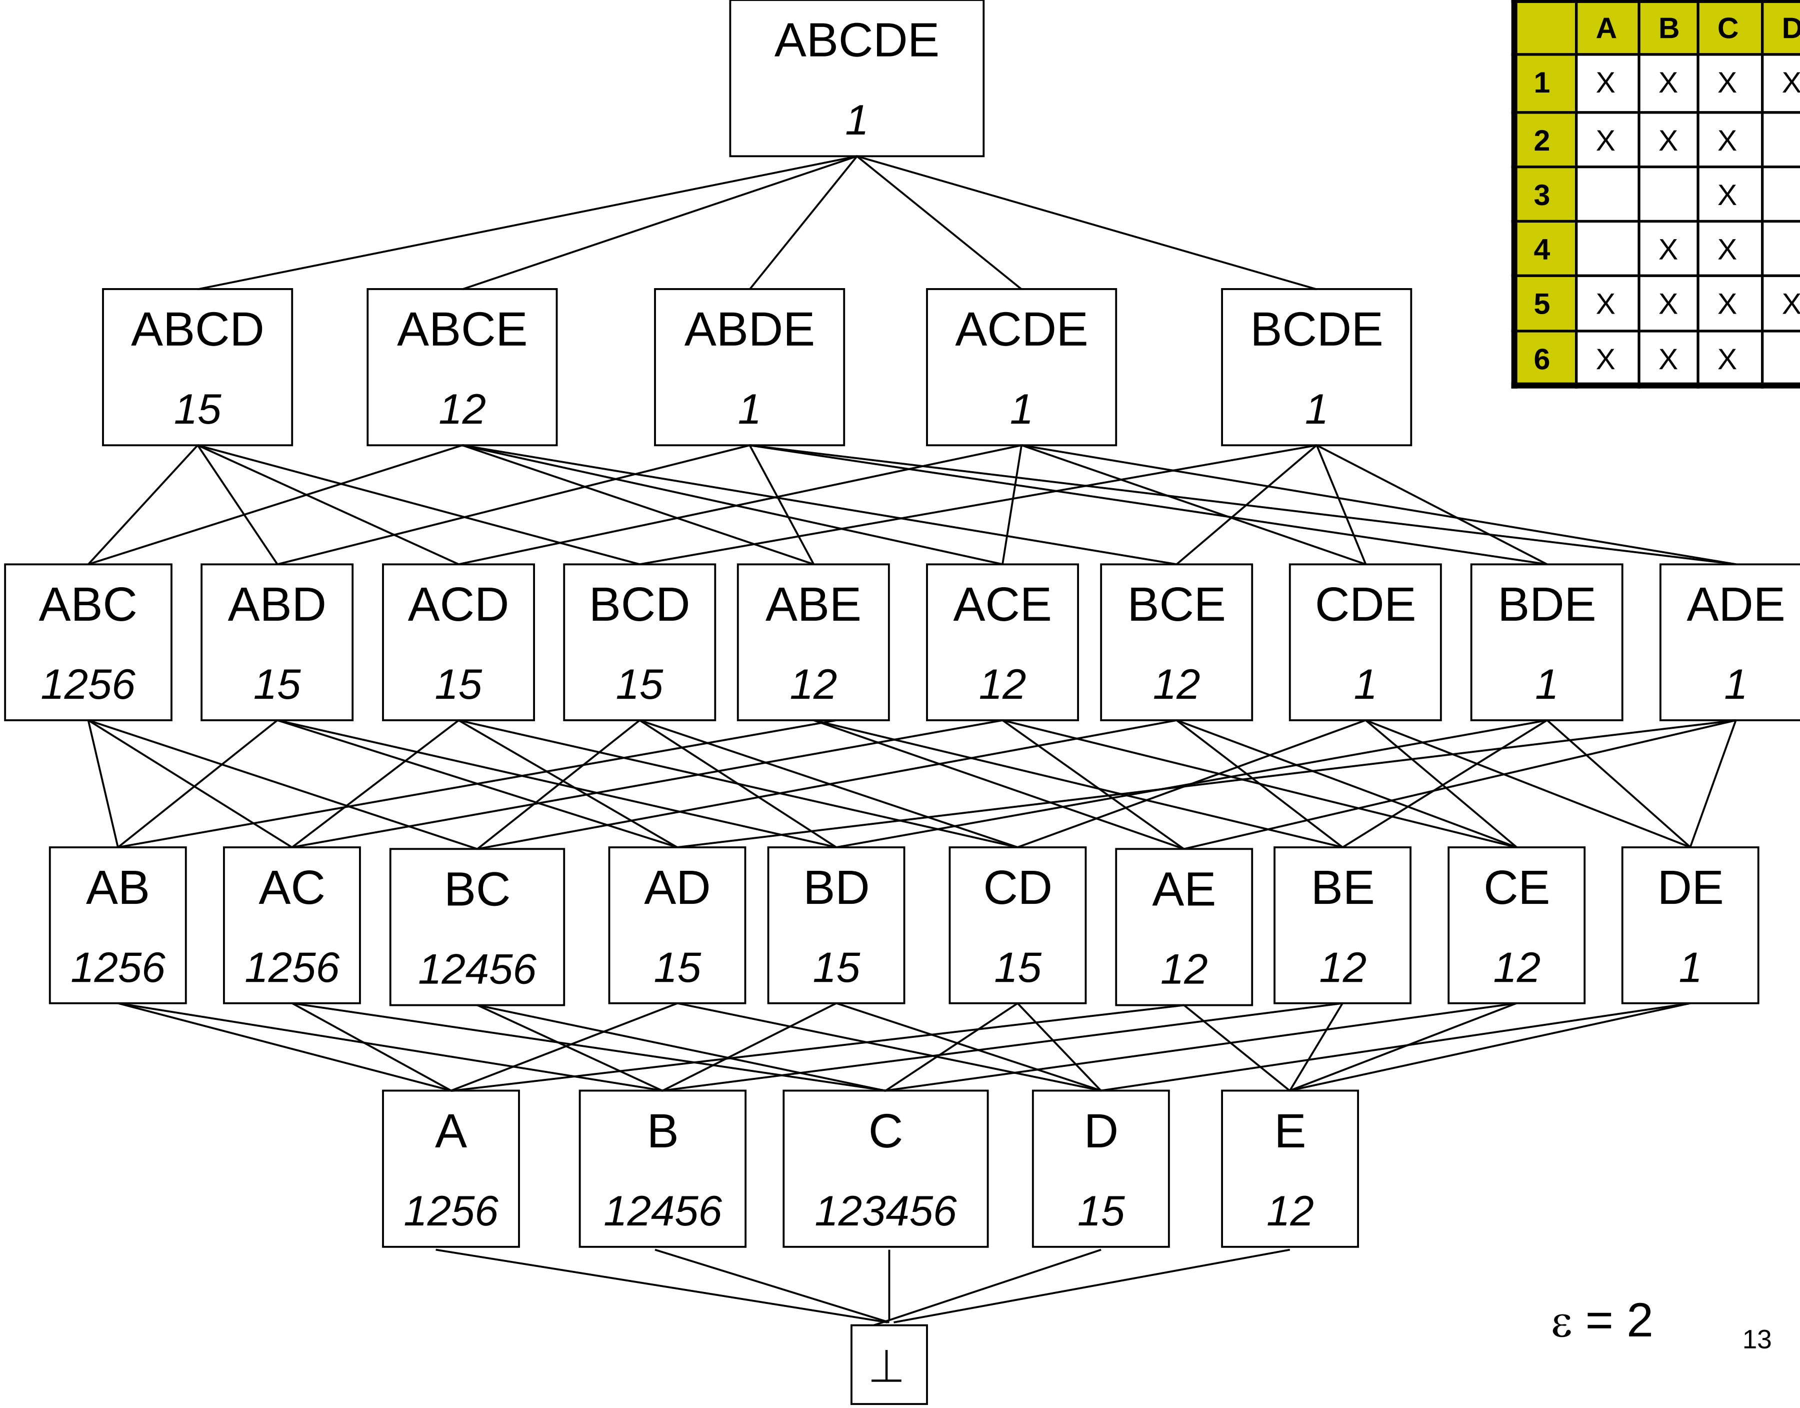

$\varepsilon = 2$

apriori-gen

```
Input: F_{k-1}

// Join step
insert into C_k
select p.item_1,p.item_2,…,p.item_{k-1},q.item_{k-1}
from p,q ∈ F_{k-1}
where p.item_1 = q.item_1,…,p.item_{k-2}=q.item_{k-2},
                                  p.item_{k-1}<q.item_{k-1}
// Prune step
foreach itemset c ∈ C_k do
   foreach (k-1)-subset s of c do
       if (s ∉ F_{k-1}) then
          delete c from C_k ;

return C_k
```

Here use of anti-monotony property !

Subset

- For each transaction *t*, find all the itemsets of C_k that are in *t*
- Brute force :
  - ```
    foreach t∈T
       foreach c ∈ Ck
             compute if c⊆t
    ```
- Too much computation !
- The Apriori solution:
  - Partition candidates into different buckets of limited size
  - Store buckets in leaves of a hash tree
  - Find candidates subset of a transaction by traversing hash tree
    ```

Hash tree construction

Max bucket size = 3

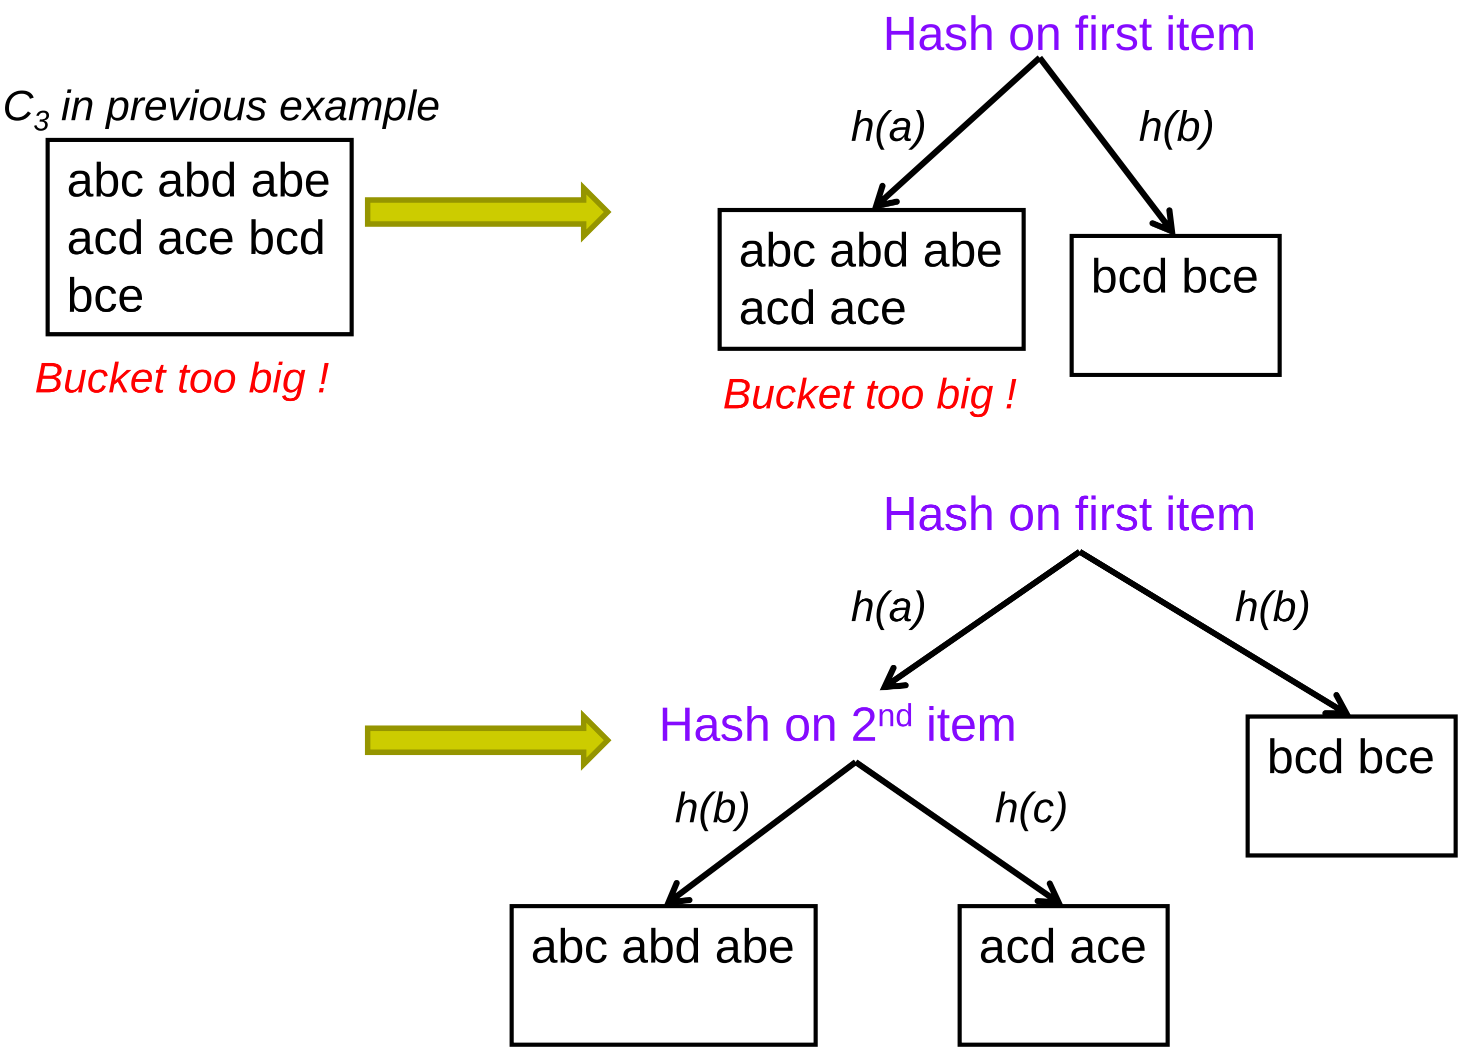

Hash tree utilisation for subset

Transaction **2'** : a c d e

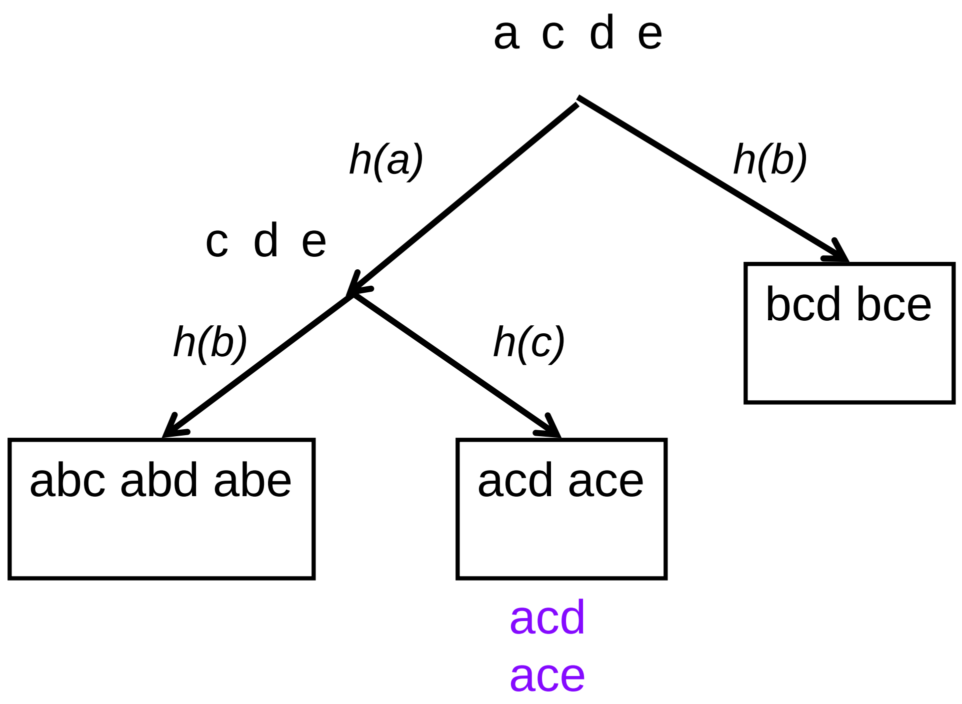

Alexandre Termier

Complexity

- apriori_gen, step k
  - Dominated by prune step ≈ $O(k.|C_k|)$
- support counting
  - ≈ $O(m.|C_k|)$ with $m = |T|$ (database size)
- → one iteration is ≈ $O(m.|C_k|)$
- Total complexity :
  - ≈ $O(m.\Sigma_k|C_k|)$
  - Worst case : candidates are all possible itemsets
    - ≈ $O(m.2^n)$ with n = number of items
- ⇒ Linear in database size
- ⇒ Exponential in number of items
- Influence of transaction width (database density) on number of traversal of hash tree

Association rules computation

- Once we have the frequent itemsets, we want the association rules.
- Reminder: we are only interested in rules that have a high confidence value

  Confidence of $X \rightarrow Y$: $c = \dfrac{\text{support}(X \cup Y)}{\text{support}(X)}$

- Let *F* be an itemset, with $|F| = k$.
  How many possible rules ?

- What is a naive solution to compute them ?

- Is it efficient ?

Monotony of confidence ?

Transactions	Items (products bought)
1	bread, butter, chocolate, vine, pencil
2	bread,butter, chocolate, pencil
3	chocolate
4	butter,chocolate
5	bread, butter, chocolate, vine
6	bread,butter, chocolate
7	bu
8	bu
9	bu

- *{chocolate} → {bread, butter}*
  confidence = 4/6 = 66%

CONFIDENCE IS NOT MONOTONE / ANTI-MONOTONE

More on monotony of confidence

- For rules coming from the same itemset, confidence is anti-monotone
  - e.g., L = {A,B,C,D}:

$$c(ABC \rightarrow D) \geq c(AB \rightarrow CD) \geq c(A \rightarrow BCD)$$

  - Confidence is anti-monotone w.r.t. number of items on the RHS of the rule

- $\rightarrow$ some pruning is possible

Association rule generation algorithm

```
Input: T, minsup, minconf, F_all = union of F_1..F_n

H_1 = ∅
foreach f_k ∈ F_all, k≥2 do begin
   A = (k-1)-itemsets a_k-1  such that a_k-1 ⊂ f_k ;
   foreach a_k-1 ∈ A do begin
        conf = support(f_k)/support(a_k-1) ;
        if conf ≥ minconf do begin
             output rule a_k-1 → (f_k − a_k-1) ;
             add (f_k − a_k-1) to H_1 ;
         end
   end
   ap-genrules(f_k, H_1) ;
end
```

ap-genrules

```
Input: fk, Hm : set of m-item consequents

if (k>m+1) then begin
   Hm+1 = apriori-gen(Hm) ; // Generate all possible m+1
                             itemsets
   foreach hm+1 ∈ Hm+1 do begin
      conf = support(fk)/support(fk-hm+1) ;
      if conf ≥ minconf then
         output rule fk - hm+1 → hm+1 ;
       else
         delete hm+1 from Hm+1 ;   } Pruning by anti-monotony
   end
   ap-genrules(fk, Hm+1) ;
end
```

First improvements of Apriori

- End of 90's :
  - Main memory: 64-256 MB
  - Databases: can go over 1 GB
  - Apriori : several passes over database…
  - $\Rightarrow$ need algorithms that can handle database in memory
- Partition [Savasere et al. 1995]
  - Cut the database in pieces fitting into memory, compute results for each piece and join them
- Sampling [Toivonen 1996]
  - Compute frequent itemsets on a sample of the database
- DIC [Brin et al. 1997]
  - Improves number of passes on database

Maximal frequent itemsets

Set of maximal frequent itemsets :

$MFI = \{ I \in FI \mid \forall I' \supset I \;\; I' \notin FI \}$

with FI set of frequent itemsets

- Several orders of magnitudes less MFI than FI
- Can be searched both **bottom-up** and **top-down**
- Pincer-Search [Lin & Kedem 1998]
  Max-Miner [Bayardo et al. 1998]
- BUT loss of information

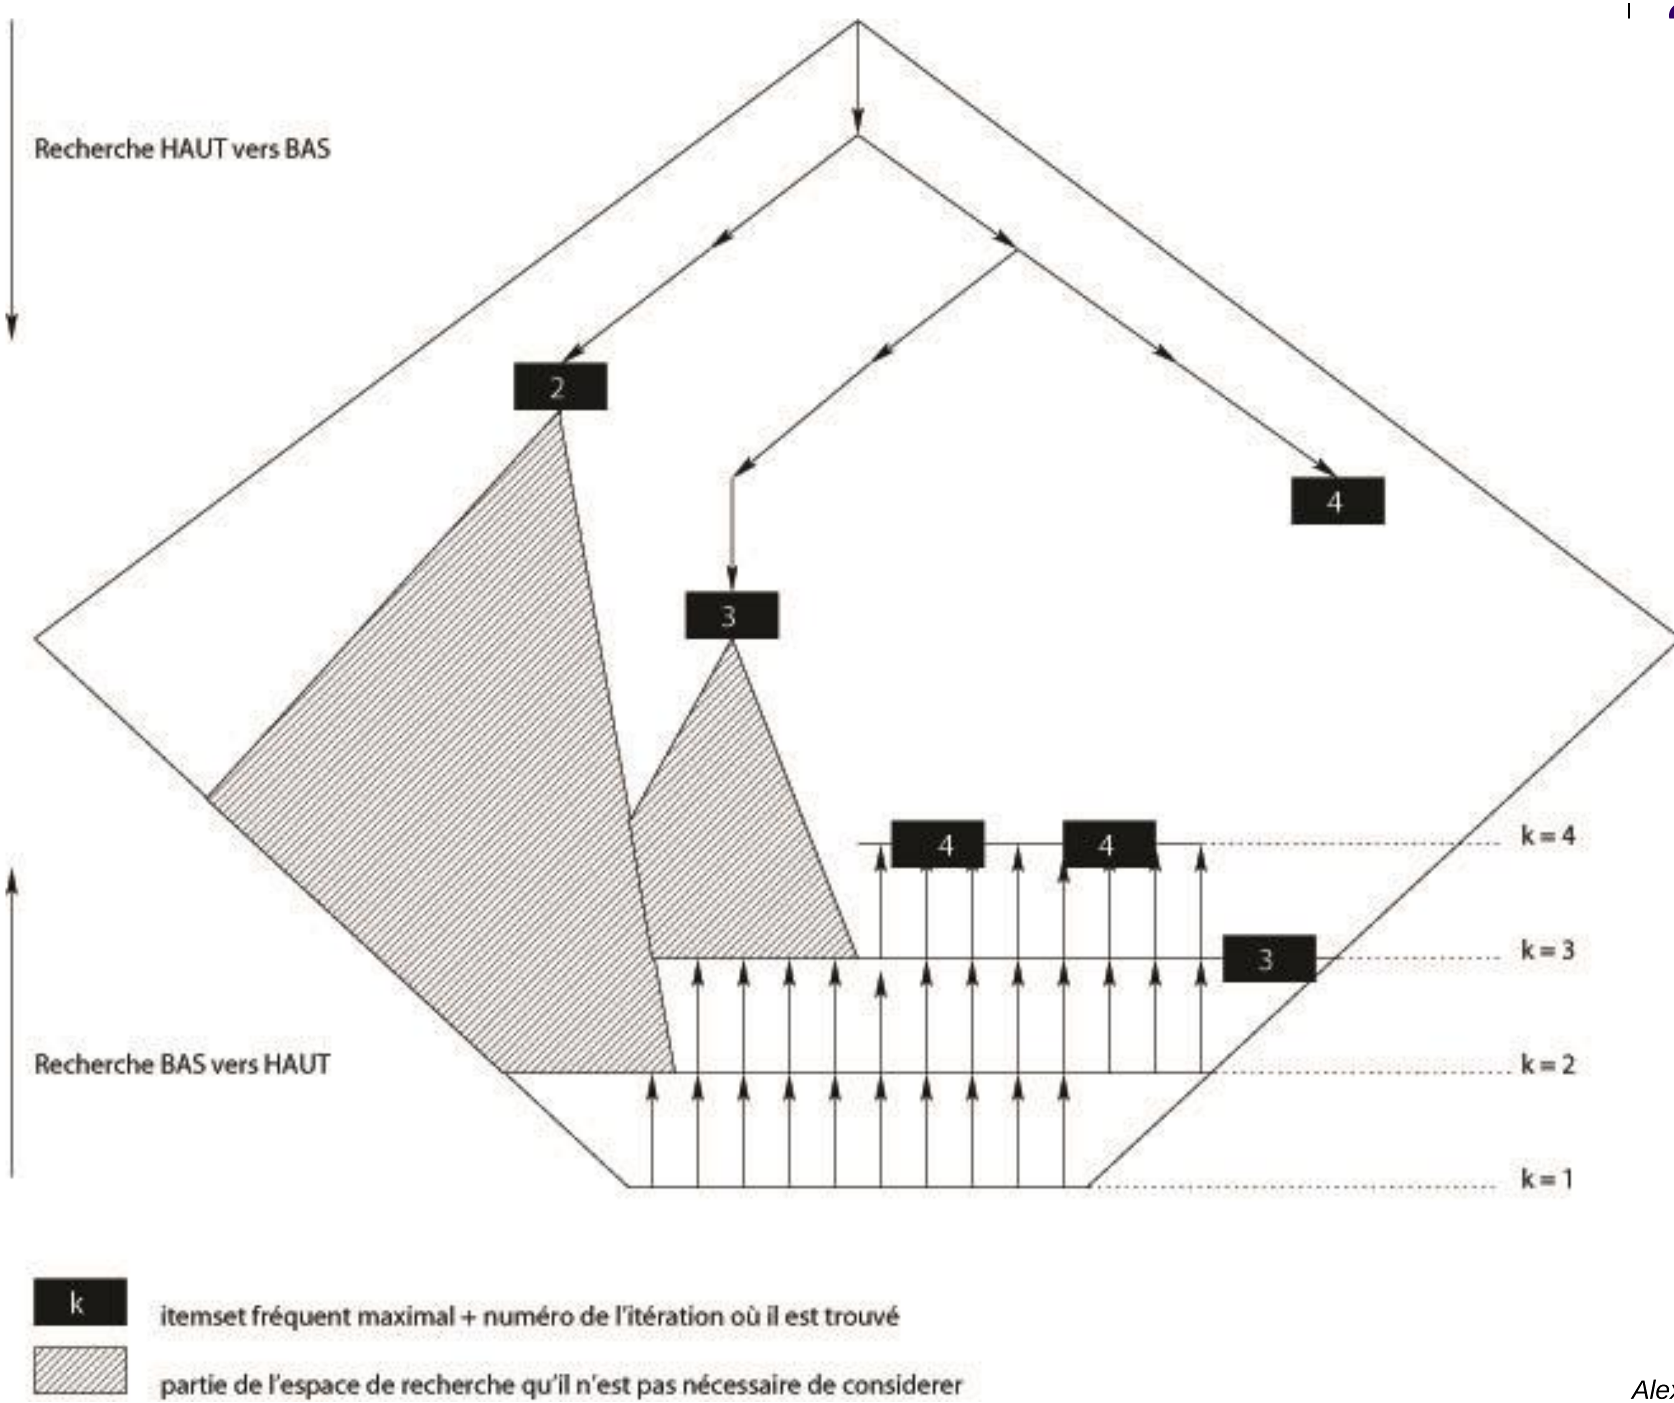

Recherche HAUT vers BAS
2
4
3
4
4
3
k = 4
k = 3
k = 2
k = 1
Recherche BAS vers HAUT
k
itemset fréquent maximal + numéro de l'itération où il est trouvé
partie de l'espace de recherche qu'il n'est pas nécessaire de considerer

The Eclat algorithm

[Zaki *et al.*, 97]

- Apriori : DB is in **horizontal** format
- Eclat introduces the **vertical** format
  - Itemset x $\rightarrow$ tid-list(x)

	A	B	C	D	E
1	X	X	X	X	X
2	X	X	X		X
3			X		
4		X	X		
5	X	X	X	X	
6	X	X	X		

Horizontal format

A	B	C	D	E
1	1	1	1	1
2	2	2	5	2
5	4	3		
6	5	4		
	6	5		
		6		

Vertical format

Vertical format

- Support counting can be done with tid-list intersections
  - $\forall I,J$ itemsets : $tidlist(I \cup J) = tidlist(I) \cap tidlist(J)$
  - No need for costly subset tests, hash tree generation…
- Problem
  - If database is big, tidlists of the many candidates created will be big also, and will not hold in memory
- Solution
  - Partition the lattice into equivalence classes
  - In Eclat : equivalence relation = **sharing the same prefix**

Initial equivalence classes

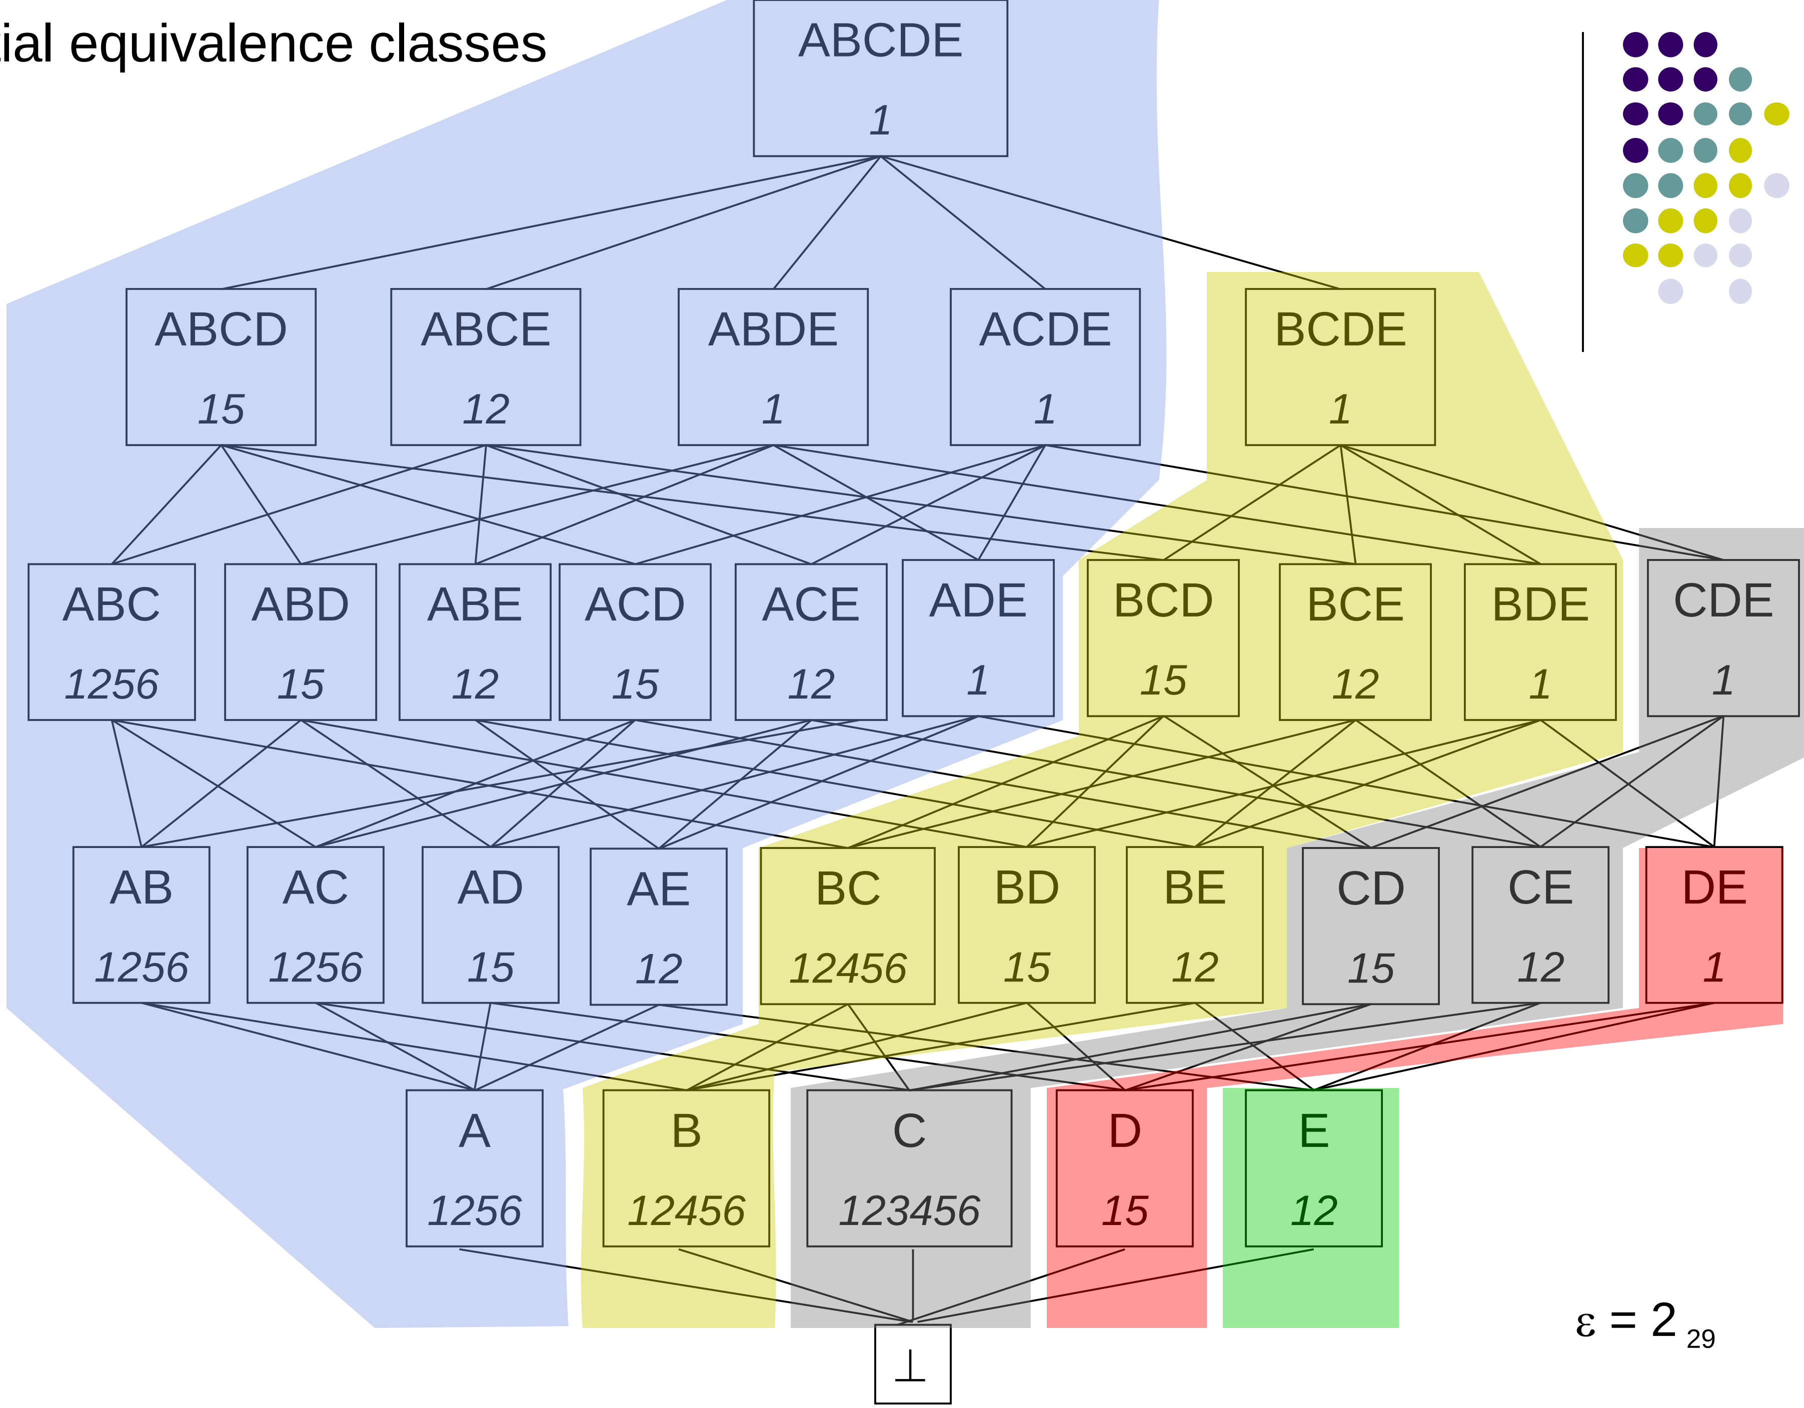

$\varepsilon = 2$

Alexandre Termier

Equivalence classes inside [A] class

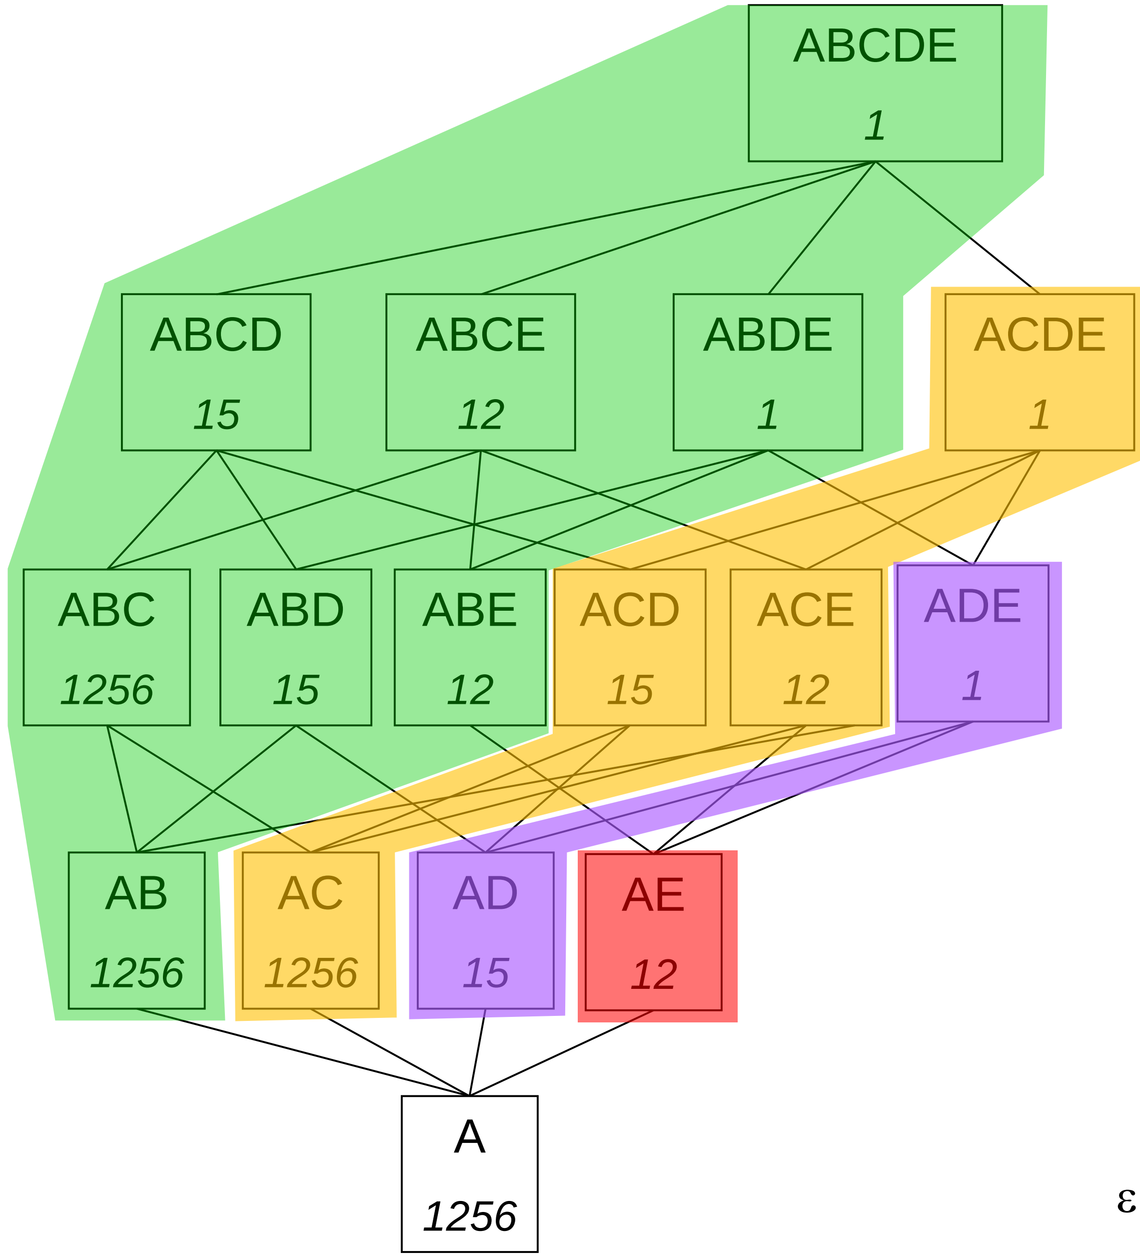

ε = 2

Alexandre Termier

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

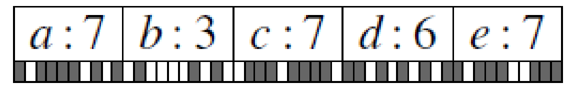

- Form a transaction list for each item. Here: bit vector representation.
  - grey: item is contained in transaction
  - white: item is not contained in transaction
- Transaction database is needed only once (for the single item transaction lists).

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

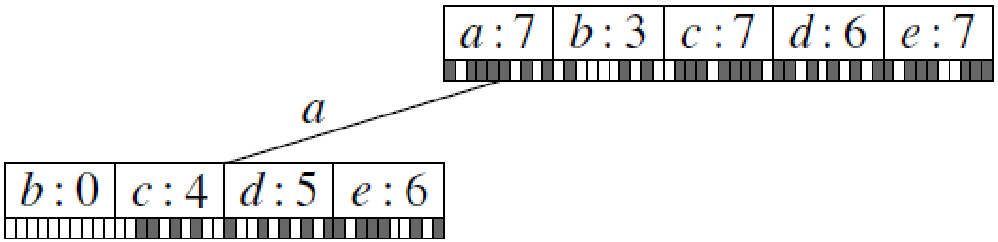

- Intersect the transaction list for item a with the transaction lists of all other items (*conditional database* for item a).
- Count the number of bits that are set (number of containing transactions). This yields the support of all item sets with the prefix a.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

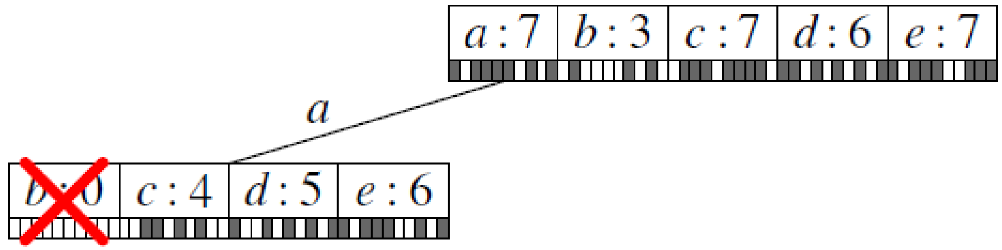

- The item set $\{a, b\}$ is infrequent and can be pruned.
- All other item sets with the prefix a are frequent and are therefore kept and processed recursively.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

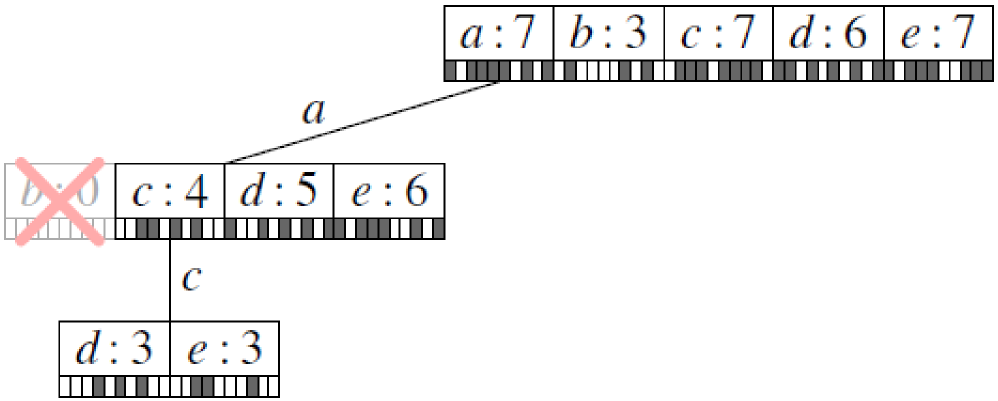

- Intersect the transaction list for the item set $\{a, c\}$ with the transaction lists of the item sets $\{a, x\}$, $x \in \{d, e\}$.
- Result: Transaction lists for the item sets $\{a, c, d\}$ and $\{a, c, e\}$.
- Count the number of bits that are set (number of containing transactions). This yields the support of all item sets with the prefix ac.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

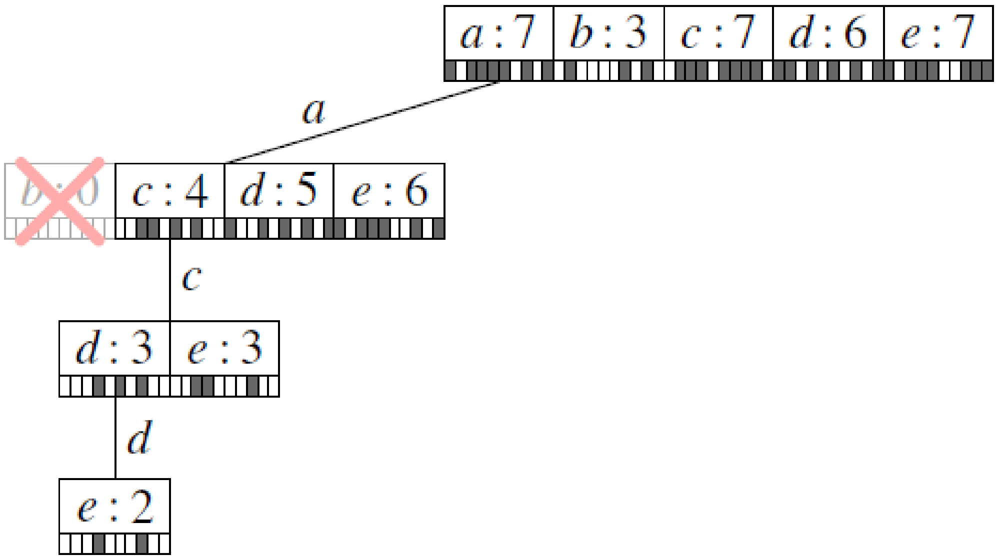

- Intersect the transaction lists for the item sets $\{a, c, d\}$ and $\{a, c, e\}$.
- Result: Transaction list for the item set $\{a, c, d, e\}$.
- With Apriori this item set could be pruned before counting, because it was known that $\{c, d, e\}$ is infrequent.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

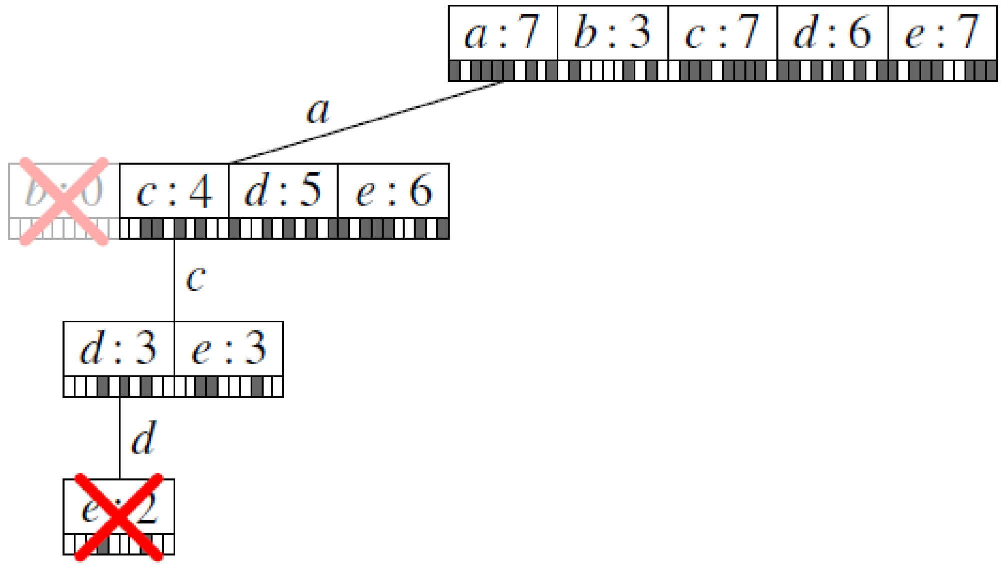

- The item set $\{a, c, d, e\}$ is not frequent (support 2/20%) and therefore pruned.
- Since there is no transaction list left (and thus no intersection possible), the recursion is terminated and the search backtracks.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

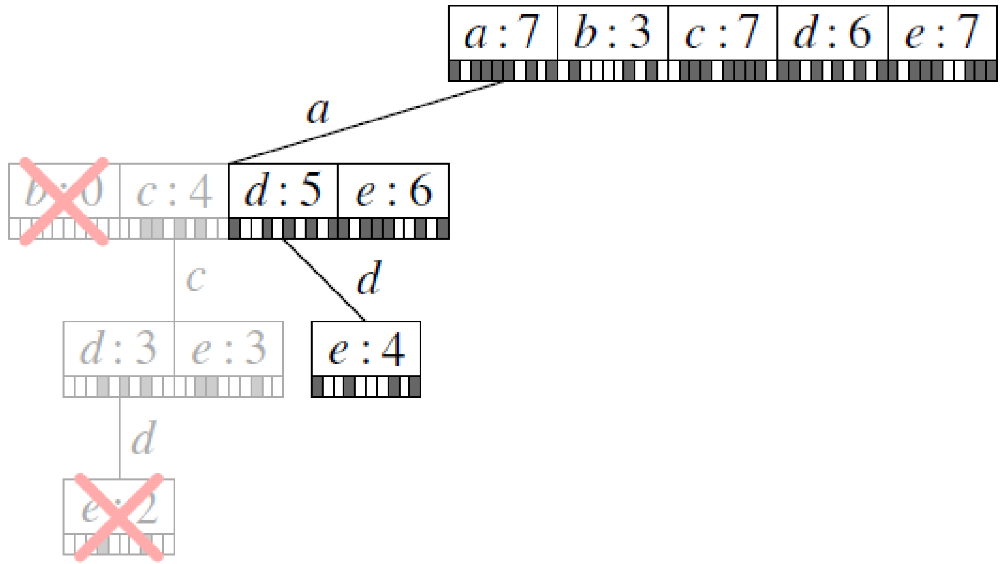

- The search backtracks to the second level of the search tree and intersect the transaction list for the item sets $\{a, d\}$ and $\{a, e\}$.
- Result: Transaction list for the item set $\{a, d, e\}$.
- Since there is only one transaction list left (and thus no intersection possible), the recursion is terminated and the search backtracks again.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

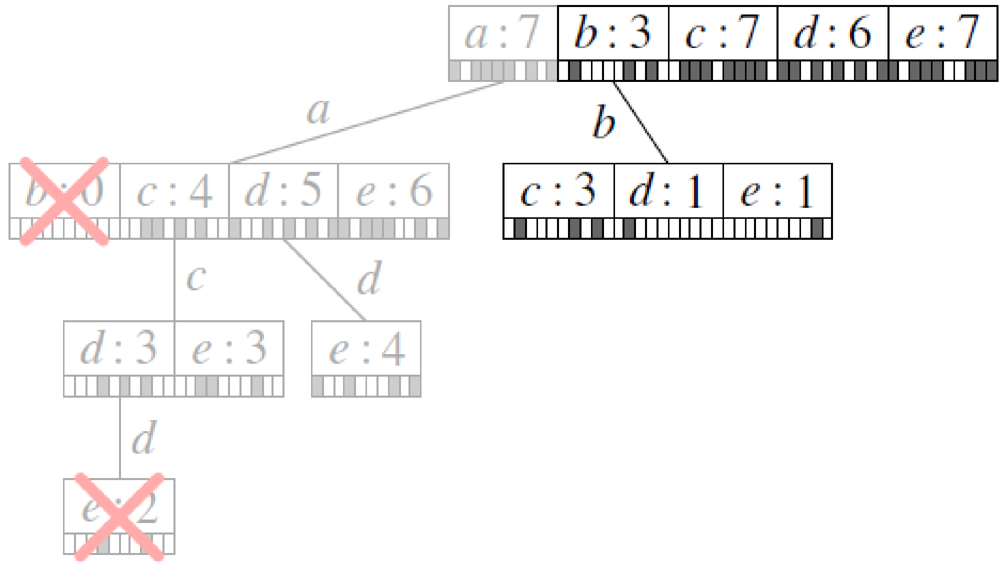

- The search backtracks to the first level of the search tree and intersect the transaction list for b with the transaction lists for c, d, and e.
- Result: Transaction lists for the item sets $\{b, c\}$, $\{b, d\}$, and $\{b, e\}$.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

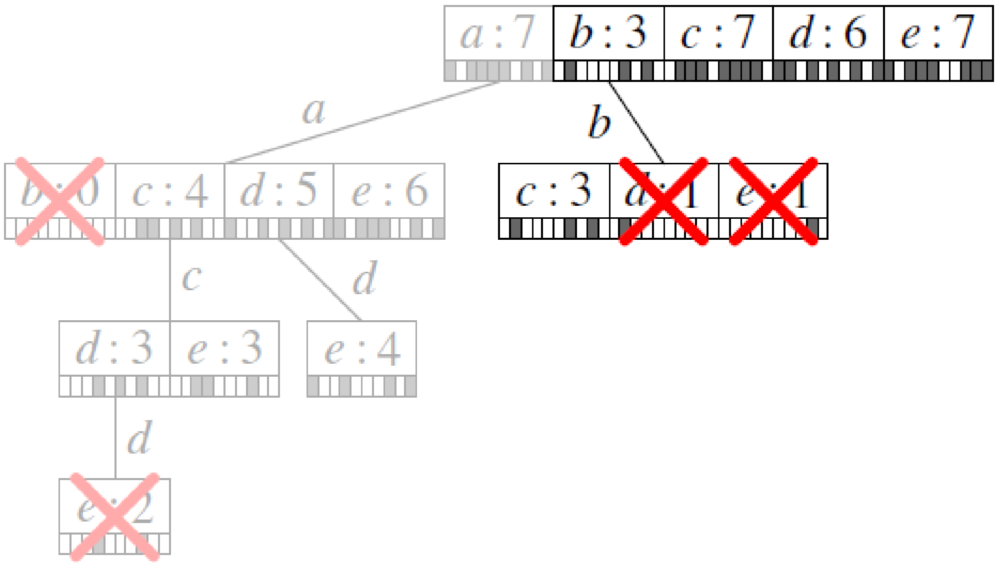

- Only one item set has sufficient support $\rightarrow$ prune all subtrees.
- Since there is only one transaction list left (and thus no intersection possible), the recursion is terminated and the search backtracks again.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

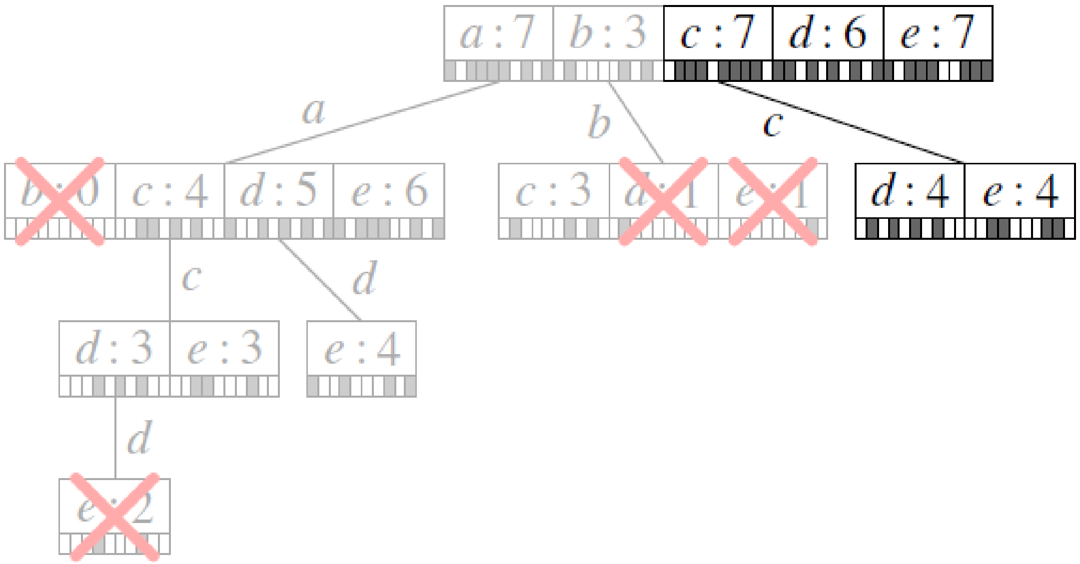

- Backtrack to the first level of the search tree and intersect the transaction list for c with the transaction lists for d and e.
- Result: Transaction lists for the item sets $\{c, d\}$ and $\{c, e\}$.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

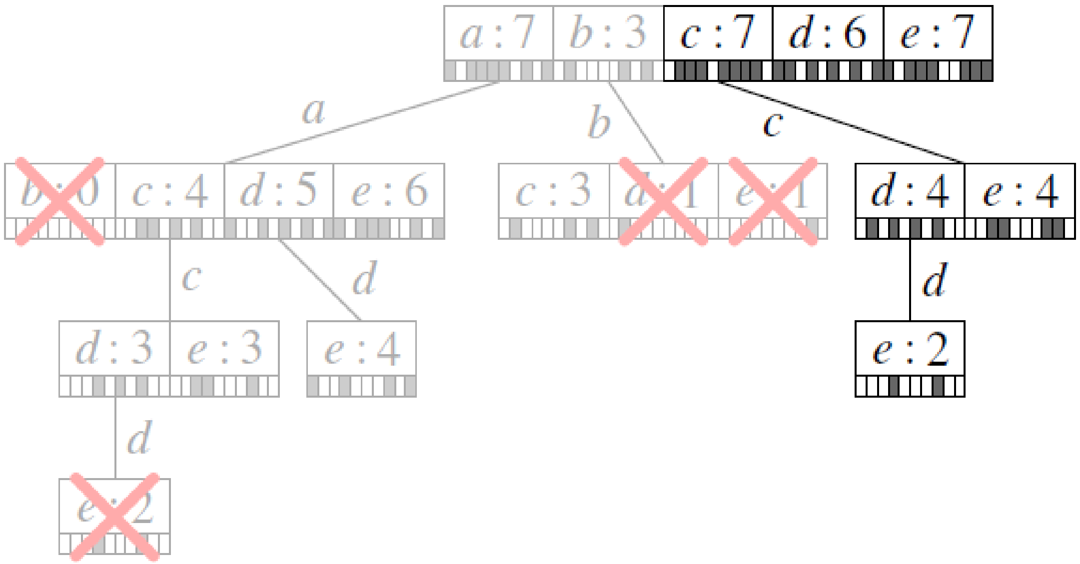

- Intersect the transaction list for the item sets $\{c, d\}$ and $\{c, e\}$.
- Result: Transaction list for the item set $\{c, d, e\}$.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

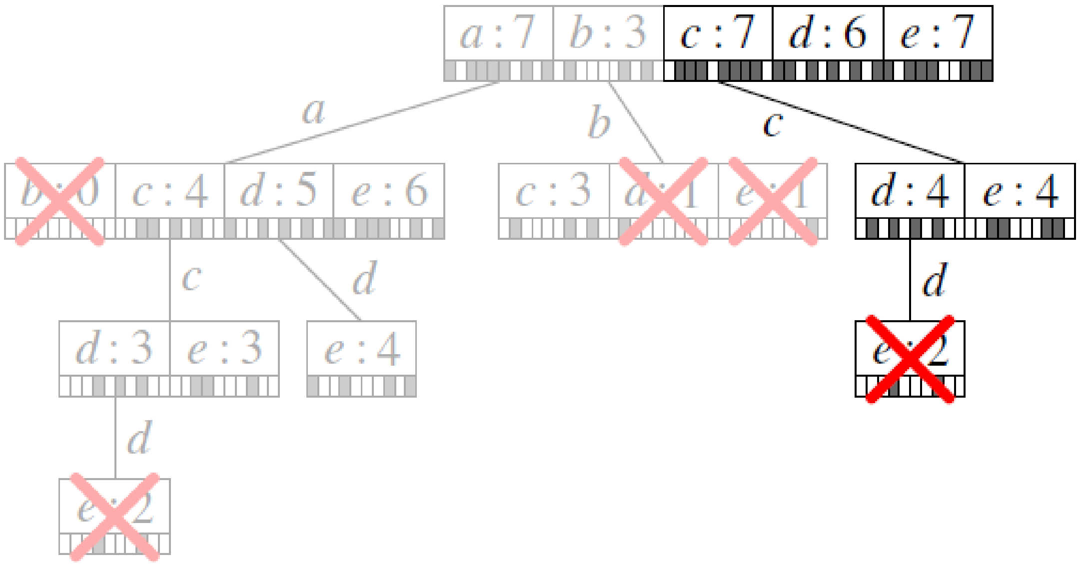

- The item set $\{c, d, e\}$ is not frequent (support 2/20%) and therefore pruned.
- Since there is no transaction list left (and thus no intersection possible), the recursion is terminated and the search backtracks.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

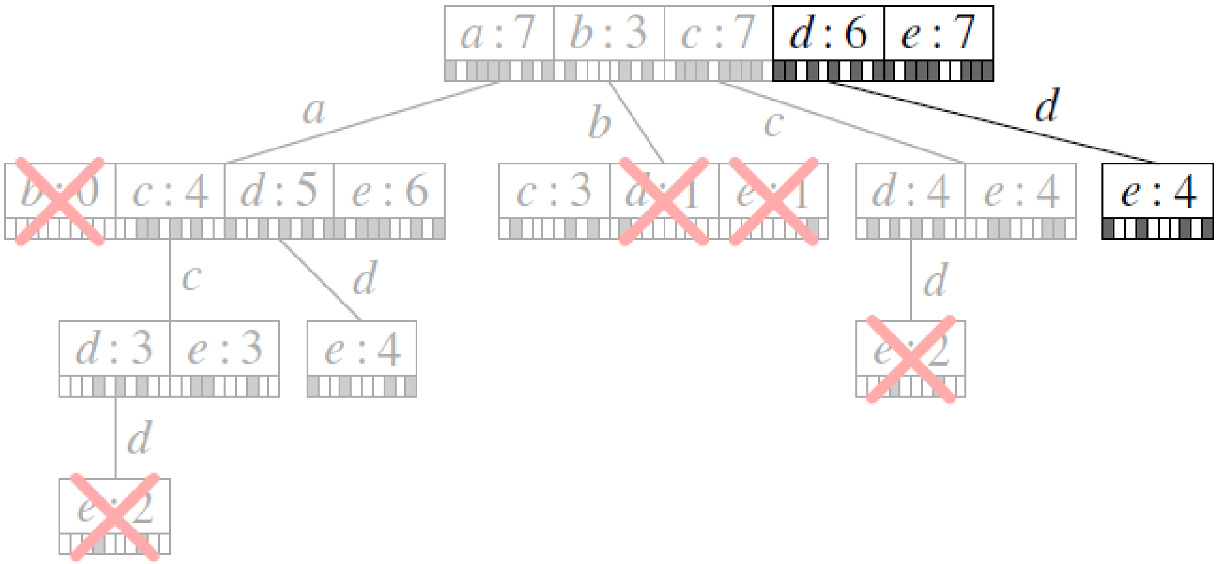

- The search backtracks to the first level of the search tree and intersect the transaction list for d with the transaction list for e.
- Result: Transaction list for the item set $\{d, e\}$.
- With this step the search is finished.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

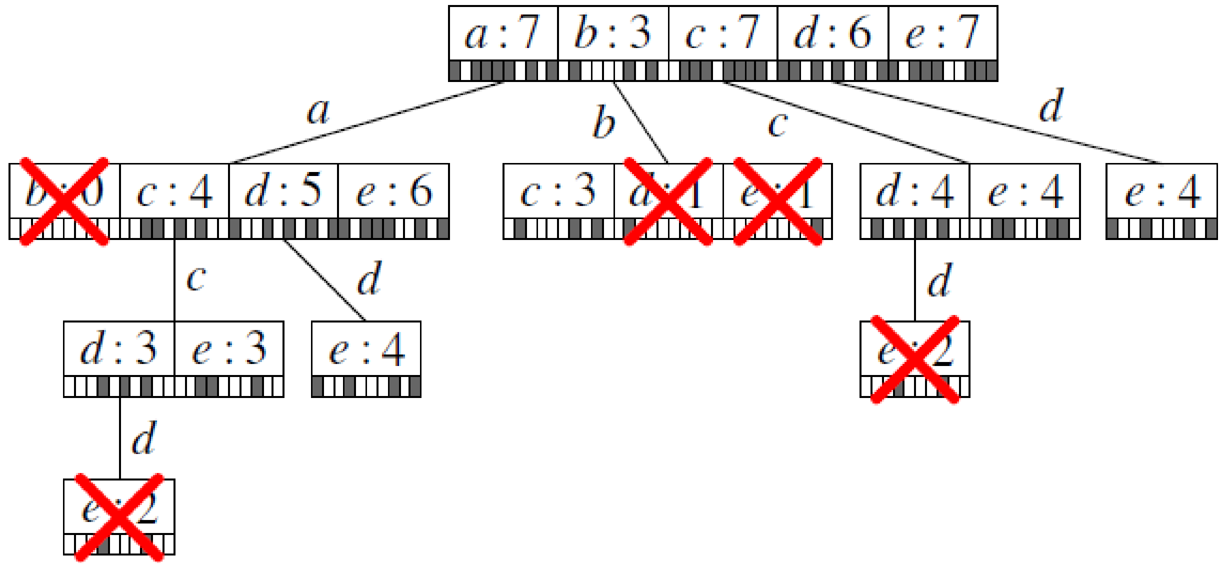

- The found frequent item sets coincide, of course, with those found by the Apriori algorithm.
- However, a fundamental difference is that Eclat usually only writes found frequent item sets to an output file, while Apriori keeps the whole search tree in main memory.

Eclat: Depth-First Search

1: $\{a, d, e\}$
2: $\{b, c, d\}$
3: $\{a, c, e\}$
4: $\{a, c, d, e\}$
5: $\{a, e\}$
6: $\{a, c, d\}$
7: $\{b, c\}$
8: $\{a, c, d, e\}$
9: $\{b, c, e\}$
10: $\{a, d, e\}$

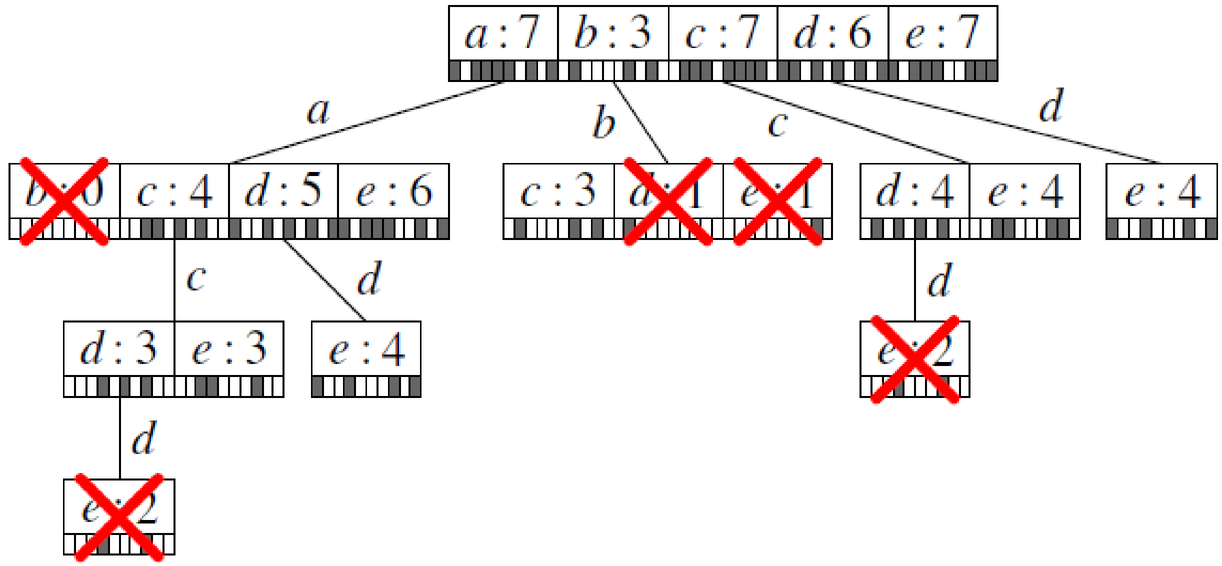

- Note that the item set $\{a, c, d, e\}$ could be pruned by Apriori without computing its support, because the item set $\{c, d, e\}$ is infrequent.

- The same can be achieved with Eclat if the depth-first traversal of the prefix tree is carried out from right to left *and* computed support values are stored.
It is debatable whether the expected gains justify the memory requirement.

Eclat algorithm

```
Input: T, minsup

compute L1 and L2 // like apriori
Transform T in vertical representation
CE2 = Decompose L2 in equivalence classes
forall E2∈CE2 do
  compute_frequent(E2)
end forall

return ∪kFk ;
```

compute_frequent(E_{k-1})

```
forall itemsets I1 and I2 in Ek-1 do
   if |tidlist(I1)∩tidlist(I2)| ≥ minsup then
      Lk ← Lk∪{I1∪I2}
   end if
end forall

CEk = Decompose Lk in equivalence classes
forall Ek∈CEk do
   compute_frequent(Ek)
end forall
```

The FP-growth approach

- FP-Growth : Frequent Pattern Growth
- No candidate generation
- Compress transaction database into FP-tree (Frequent Pattern Tree)
  - Extended prefix-tree
- Recursive processing of *conditional databases*
- Can be one order of magnitude faster than Apriori

FP-tree

- Compact structure for representing DB and frequent itemsets

1. Composed of :
   - root
   - item-prefix subtrees
   - frequent-item-header array
2. Node =
   - item-name
   - count *// number of transactions containing path reaching this node*
   - node-link *// next node having same item-name*
3. Entry in frequent-item-header array =
   - item-name
   - head of node-link *// pointer to first node having item-name*

- Both an **horizontal** (prefix-tree) and a **vertical** (node links) structure

FP-tree example (1/2)

Original transaction database	Items ordered by frequency	Transactions reordered by item frequency (infrequent item **F** pruned)	Transactions sorted lexicographically
A B C D E F	C : 6	C B A D E	C
A B C E	B : 5	C B A E	C B
C	A : 4	C	C B A
B C	D : 2	C B	C B A D
A B C D	E : 2	C B A D	C B A D E
A B C	F : 1	C B A	C B A E

minsup = 2

FP-tree example (2/2)

C
C B
C B A
C B A D
C B A D E
C B A E

Transactions sorted lexicographically

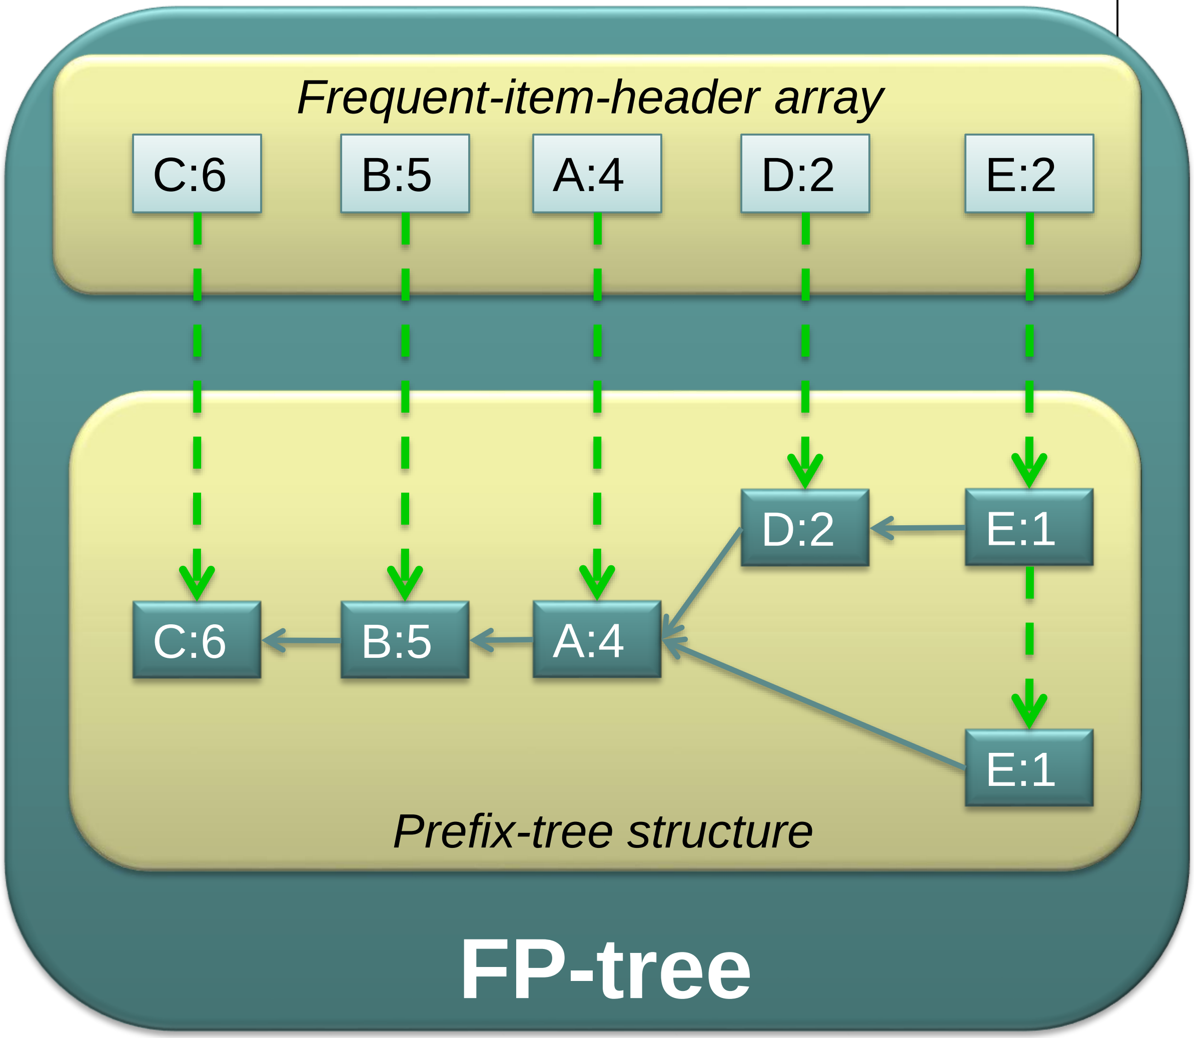

Alexandre Termier

Exercise

- Draw the FP-tree for the following DB : *(minsup = 3)*

A D F
A C D E
B D
B C D
B C
A B D
B D E
B C E G
C D F
A B D

FP-Growth: Preprocessing the Transaction Database

1	2	3	4	5
$a\ d\ f$	d: 8	$d\ a$	$d\ b$	
$a\ c\ d\ e$	b: 7	$d\ c\ a\ e$	$d\ b\ c$	
$b\ d$	c: 5	$d\ b$	$d\ b\ a$	FP-tree
$b\ c\ d$	a: 4	$d\ b\ c$	$d\ b\ a$	(see next slide)
$b\ c$	e: 3	$b\ c$	$d\ b\ e$	
$a\ b\ d$	f: 2	$d\ b\ a$	$d\ c$	
$b\ d\ e$	g: 1	$d\ b\ e$	$d\ c\ a\ e$	
$b\ c\ e\ g$		$b\ c\ e$	$d\ a$	
$c\ d\ f$		$d\ c$	$b\ c$	
$a\ b\ d$	$s_{\min} = 3$	$d\ b\ a$	$b\ c\ e$	

1. Original transaction database.
2. Frequency of individual items.
3. Items in transactions sorted descendingly w.r.t. their frequency and infrequent items removed.
4. Transactions sorted lexicographically in ascending order (comparison of items is the same as in preceding step).
5. Data structure used by the algorithm (details on next slide).

Transaction Representation: FP-Tree

- Build a **frequent pattern tree (FP-tree)** from the transactions (basically a prefix tree with links between branches for items).
- Frequent single item sets can be read directly from the FP-tree.

Simple Example Database

1	$a\ d\ f$	4	$d\ b$
	$a\ c\ d\ e$		$d\ b\ c$
	$b\ d$		$d\ b\ a$
	$b\ c\ d$		$d\ b\ a$
	$b\ c$		$d\ b\ e$
	$a\ b\ d$		$d\ c$
	$b\ d\ e$		$d\ c\ a\ e$
	$b\ c\ e\ g$		$d\ a$
	$c\ d\ f$		$b\ c$
	$a\ b\ d$		$b\ c\ e$

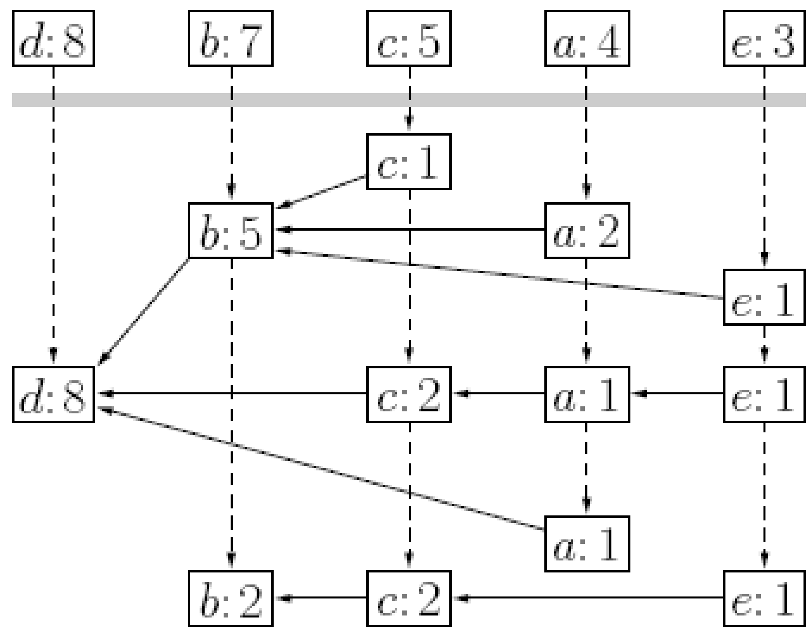

FP-tree

FP-Growth

```
FP-growth(FP, prefix)

foreach frequent item x in increasing order of frequency do
   prefix = prefix ∪ x
   Dx = ∅
   count_x = 0
   foreach node-link nl_x of x do
        D_x = D_x ∪ {transaction of path reaching x, with
                     count for each item = nl_x.count, without x}
        count_x += nl_x.count
   end
   if count_x ≥ minsup then
        output (prefix ∪ x)
        FP_x = FP-tree constructed from D_x
        FP-growth(FP_x, prefix)
   end if
end
```

FP-Growth example

C : 6
B : 5
A : 4
D : 2
E : 2

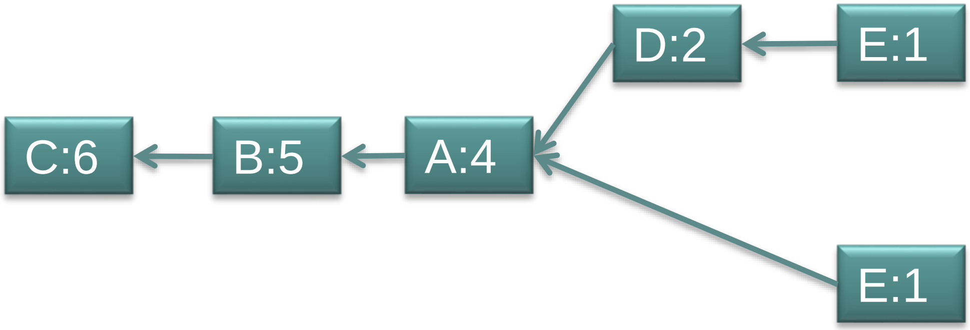

Original FP-tree

Start with item E (lowest support)

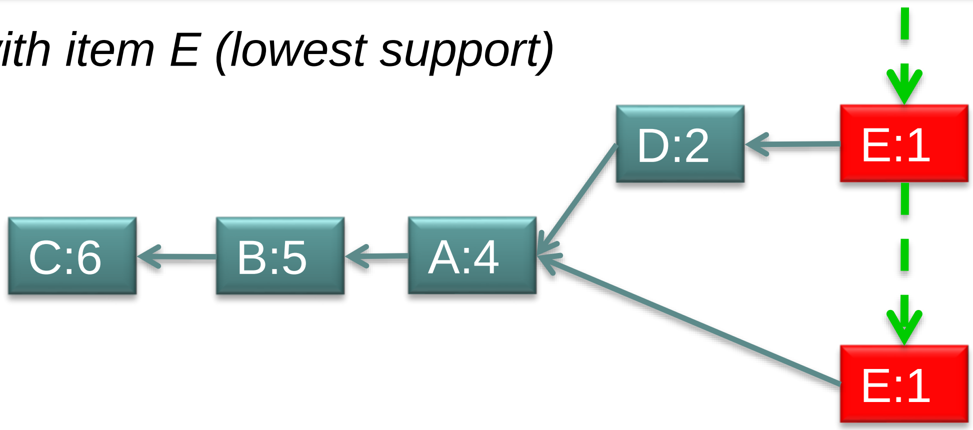

$count_E = 1+1 = 2$

$\Rightarrow$ E is frequent

$\Rightarrow$ Output E

Conditional FP-tree for E

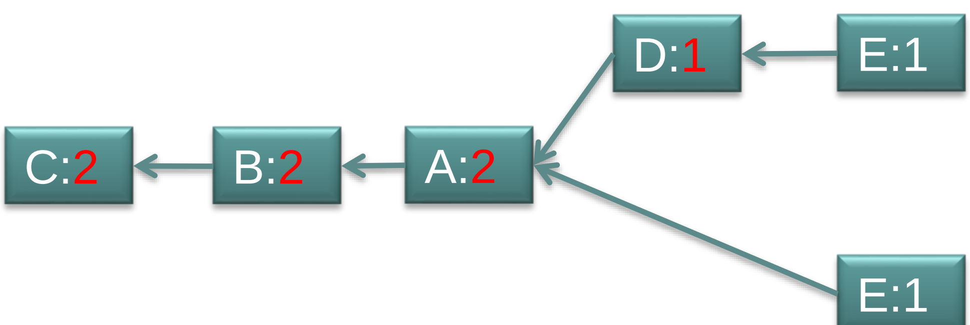

- update counts
  $\rightarrow$ only transactions containing E
- drop E

Alexandre Termier

FP-Growth example (cont.)

Conditional FP-tree for E

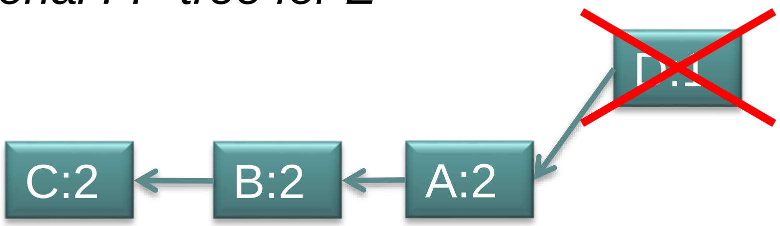

D not frequent here
→ do not consider DE

Loop on ~~AE~~, BE, CE

The rest is left as exercise…

For AE :

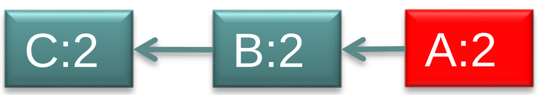

$count_{AE}$= 2

⇒ AE is frequent

⇒ Output AE

Conditional FP-tree for AE:

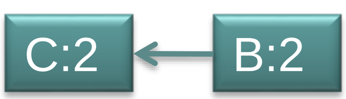

For BAE :

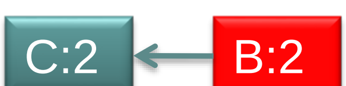

$count_{BAE}$= 2

⇒ BAE is frequent

⇒ Output BAE

Conditional FP-tree for BAE:

For CBAE :

$count_{CBAE}$= 2

⇒ CBAE is frequent

⇒ Output CBAE

Experiments: Execution Times

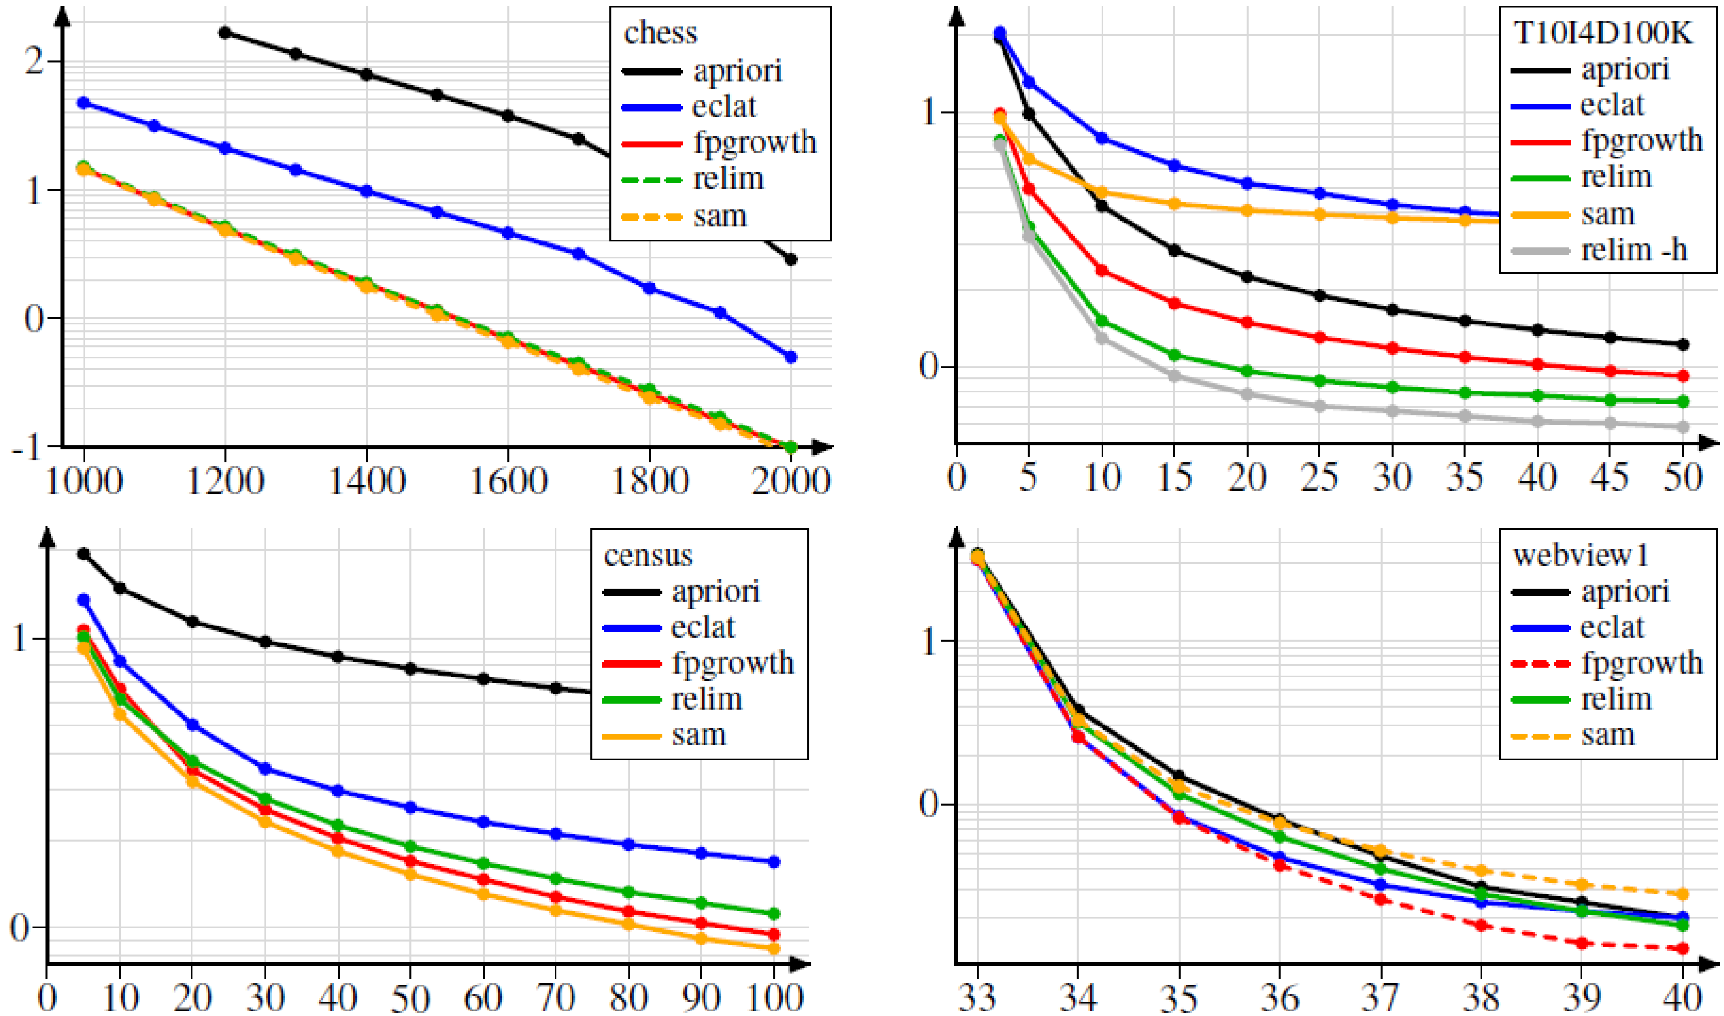

Decimal logarithm of execution time in seconds over absolute minimum support.

Problems of frequent itemsets

- Large computation time
- For low support values, huge number of frequent itemsets
- Lots of redundant information

Simple example:

Tid	Transaction
1	A B C D
2	A B C
3	A B C E

minsup = 2

FIS	Support	Tidlist
A B C	3	{1,2,3}
A B	3	{1,2,3}
A C	3	{1,2,3}
B C	3	{1,2,3}
A	3	{1,2,3}
B	3	{1,2,3}
C	3	{1,2,3}

REDUNDANT

Closed frequent itemsets

[Pasquier *et al.*, 99]

- We have seen that there is **loss of information** with maximal frequent itemsets

- Lets consider equivalence classes for frequent itemsets sharing the same tidlist
- The closed frequent itemsets are the maximums of these equivalence classes

Set of closed frequent itemsets :

$$CFI = \{ I \in FI \mid \forall I' \in FI \text{ tq } tidlist(I')=tidlist(I) \;\; I' \subseteq I \}$$

with FI set of frequent itemsets

$\Rightarrow$ Sets are ordered by inclusion: $MFI \subseteq CFI \subseteq FI$

	A	B	C	D	E
1	X	X	X	X	X
2	X	X	X		X
3			X		
4		X	X		
5	X	X	X	X	
6	X	X	X		

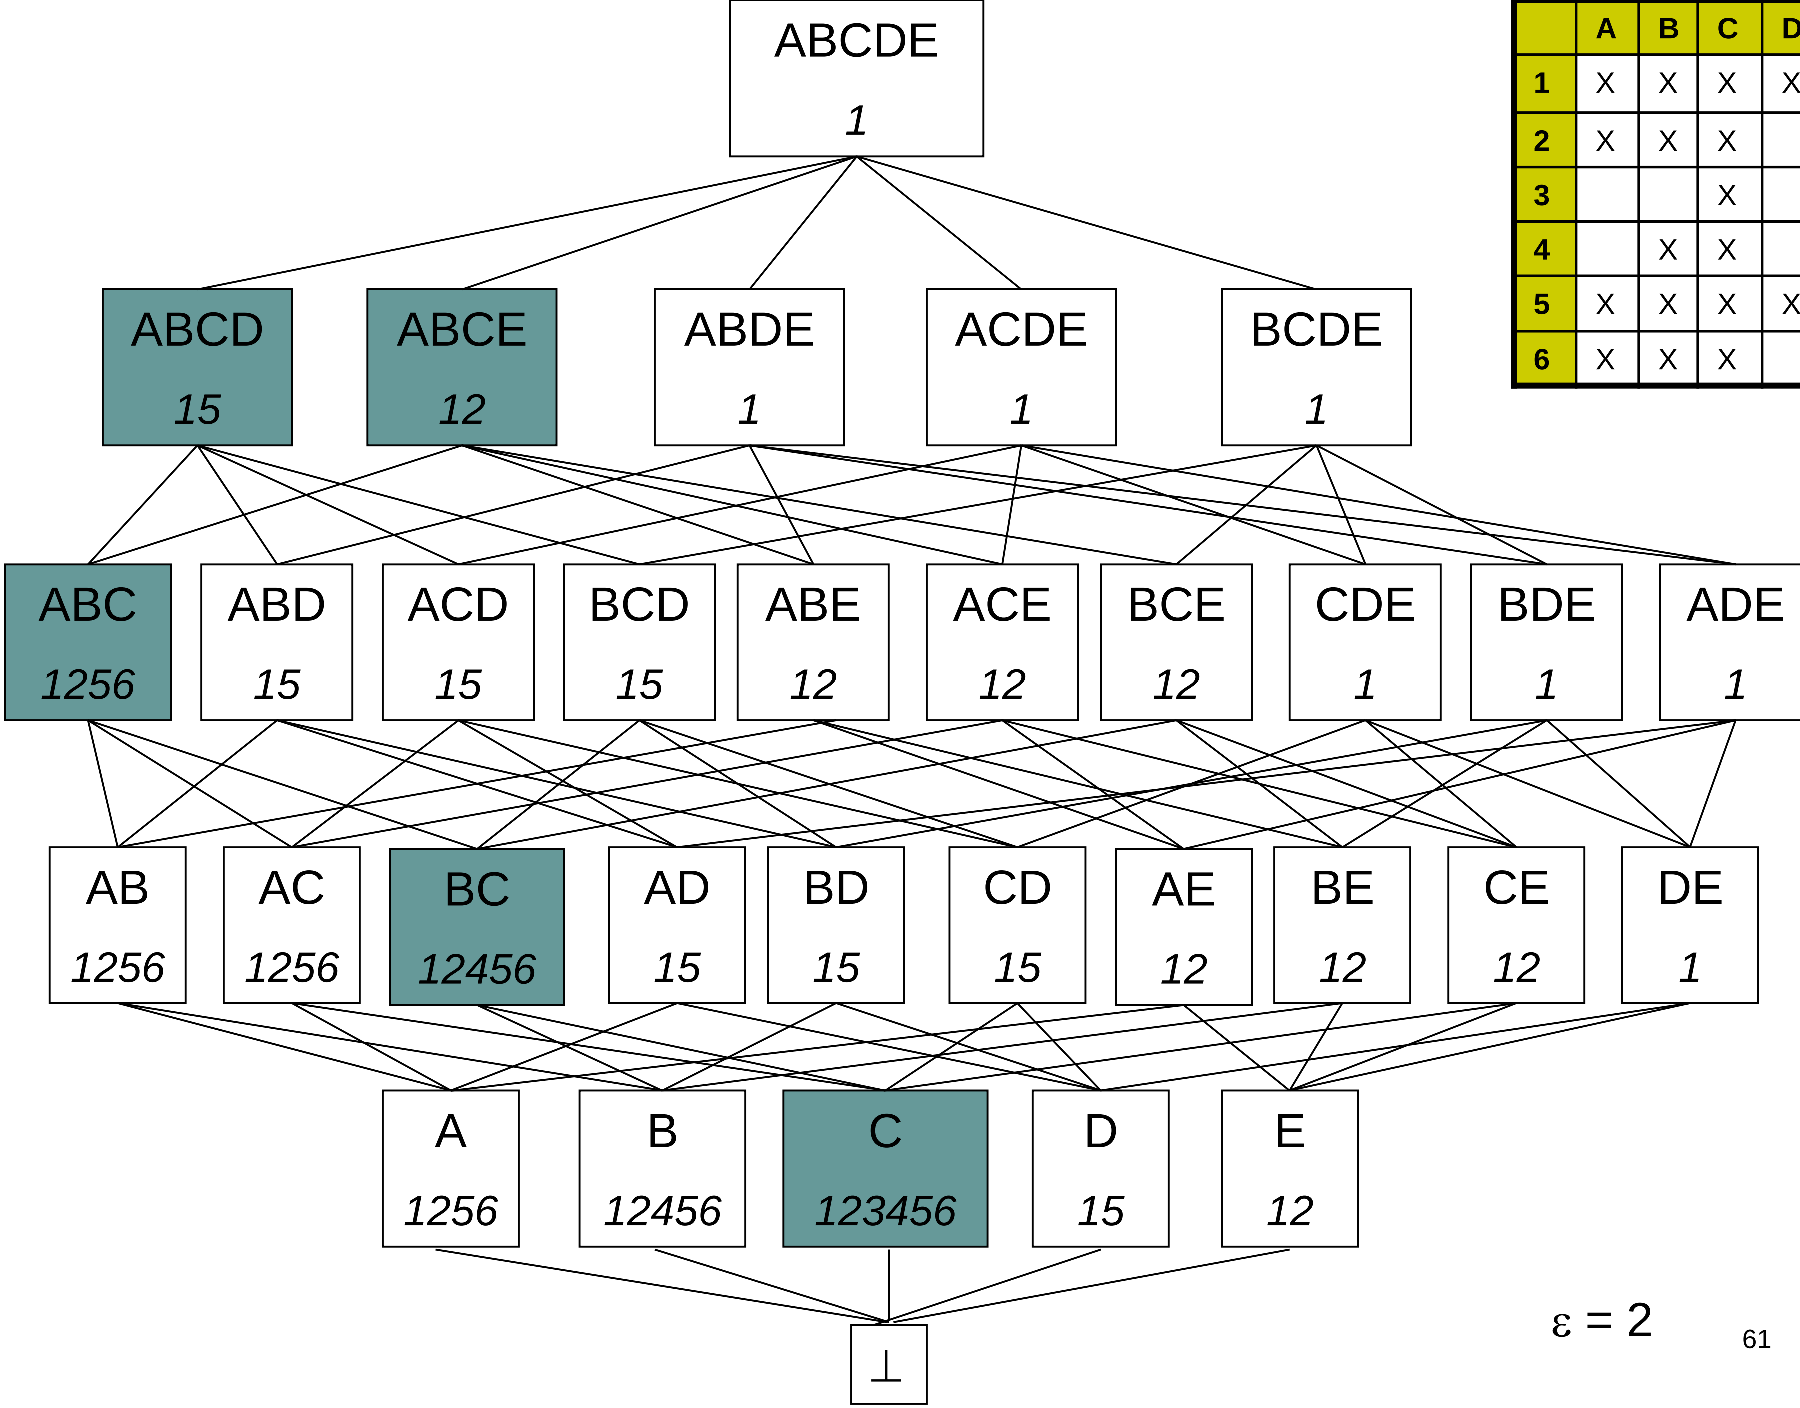

$\varepsilon = 2$

Types of Frequent Item Sets: Experiments

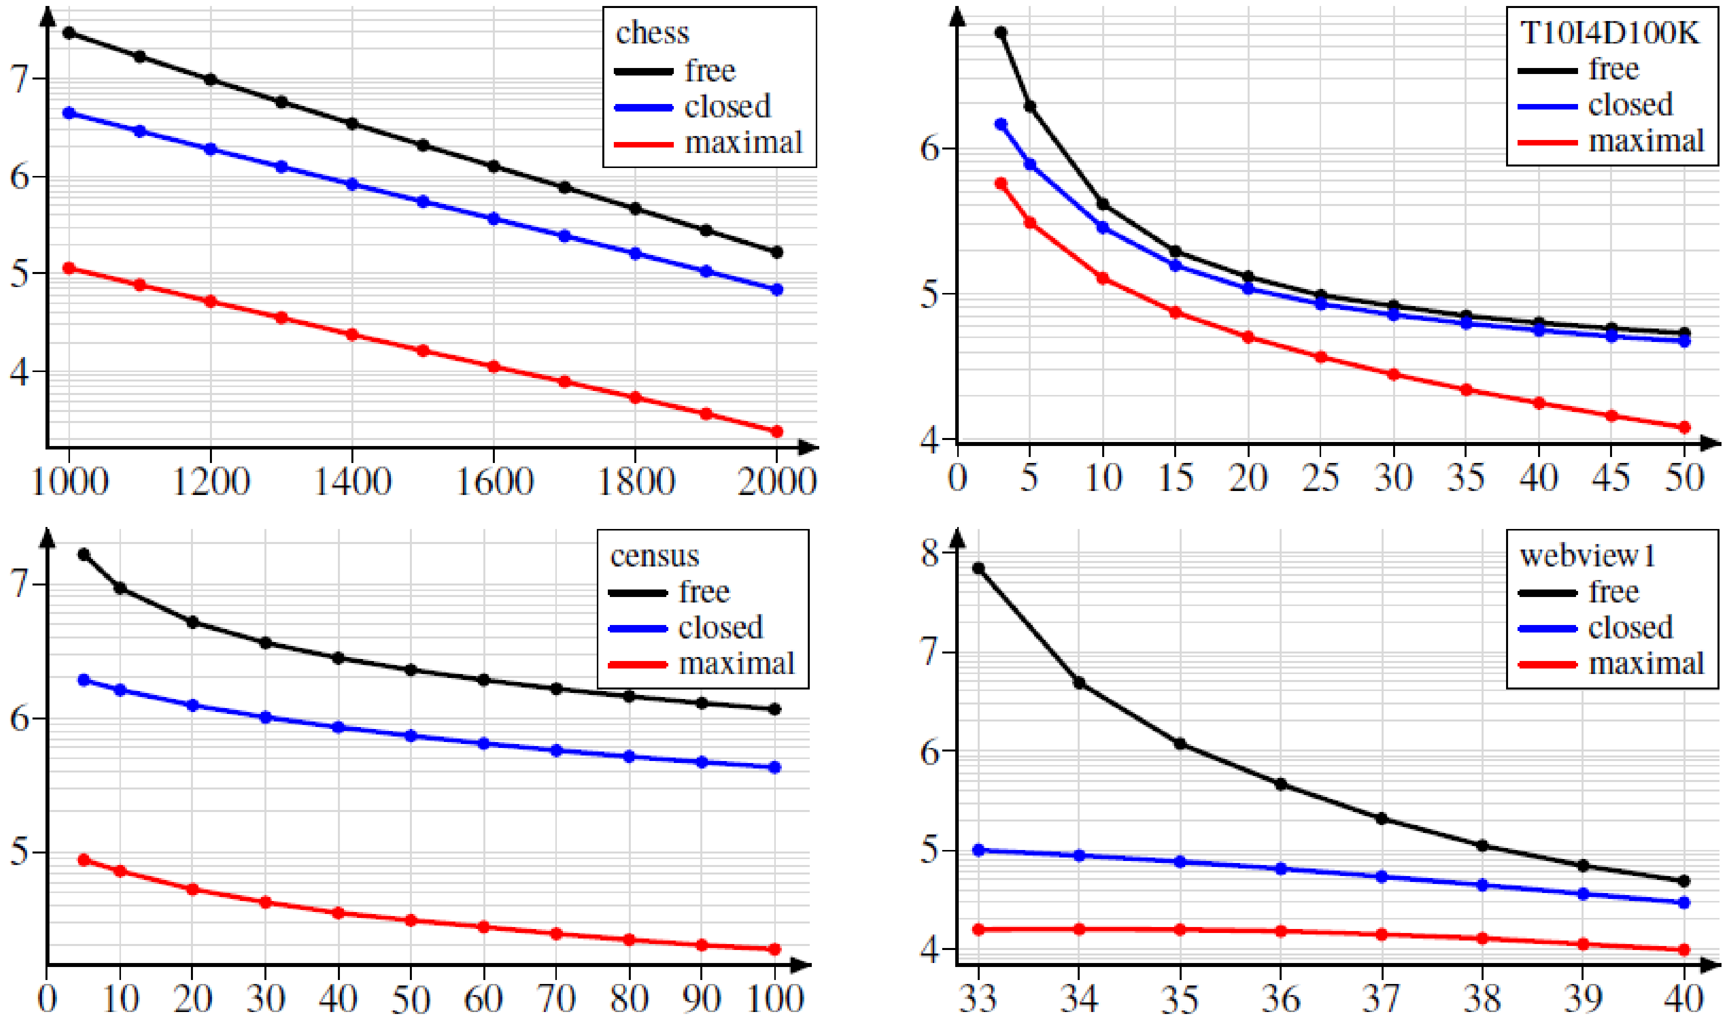

Decimal logarithm of the number of item sets over absolute minimum support.

Computing closed frequent itemsets

- Brute force *(frequent pattern base)*
  - Enumerate all the frequent patterns
  - Output only closed ones
  - *Most of the time : inefficient*
    - *Exception: if |FI| is very small*
- Closure base
  - Compute only closed patterns with closure operations
  - *Can be very efficient*

Efficient computation

- First algorithms (Closet, Charm,…)
  - Candidate-based method
  - **Try** to compute as many non-closed frequent itemsets as possible
  - OR Closure Extension: **add an item to an existing closed frequent itemset, and take closure**
  - Keep in memory all closed frequent itemsets found so far
  - → *Need a lot of memory during execution*

- Reverse search (LCM)
  - Depth First Search algorithm so no global memory needed
  - *Fast computation time, Little memory usage*

Closure Extension of Itemset

- Usual backtracking does not work for closed itemsets, because there are possibly big gap between closed itemsets

- On the other hand, any closed itemset is obtained from another by **"add an item and take closure (maximal)"**

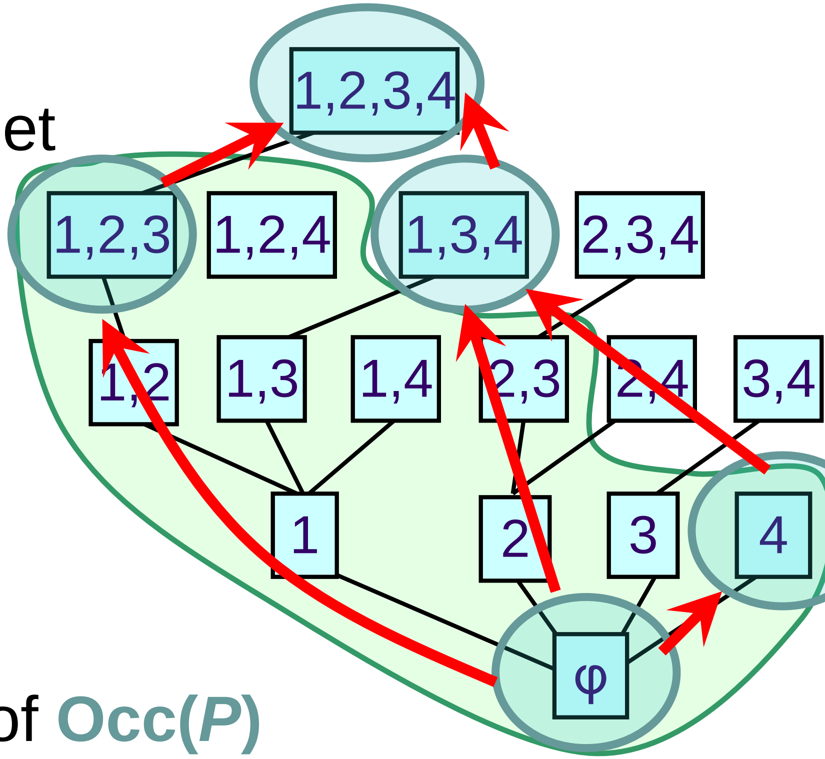

- closure of ***P*** is the closed itemset having the same denotation to ***P***, and computed by taking intersection of **Occ(*P*)**

This is an adjacency structure defined on closed itemsets, thus we can perform graph search on it, with using memory

Reverse Search

- Uno and Arimura found that the closed frequent itemsets are organized in a **directed spanning tree**

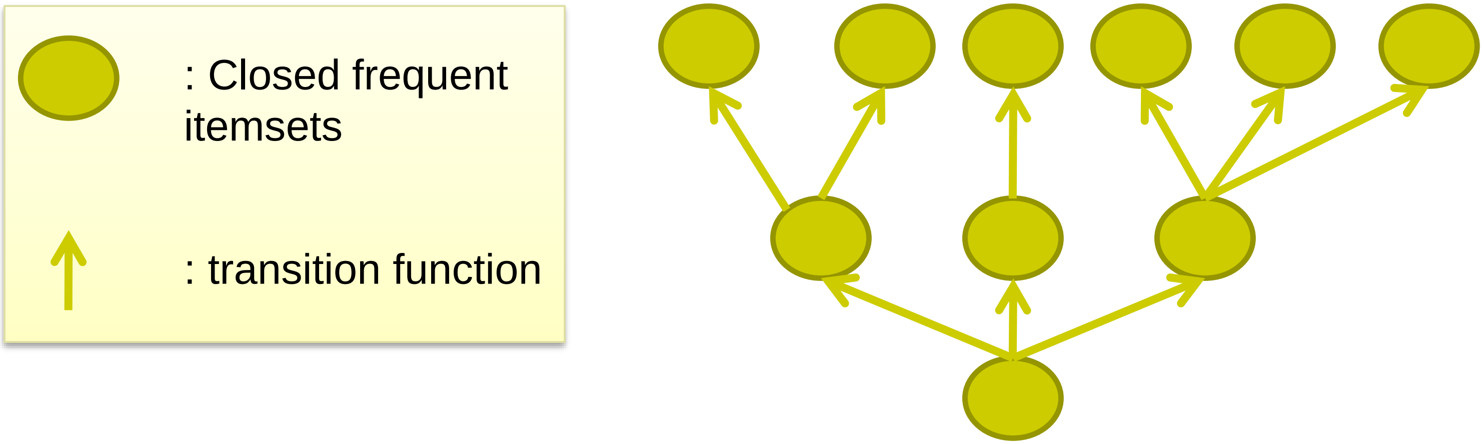

- $\Rightarrow$ they can be visited by DFS
- $\Rightarrow$ from a node of the tree, need of a *transition function* to compute its children

Tree of closed frequent itemsets

- Search space of CFIS = lattice = DAG

- DAG → Tree : impose order of exploration
- Order need to:
  - follow enumeration strategy
  - be inexpensive to enforce

- Order of Arimura and Uno
  - CFIS P, Q
  - Q children of P if all items of Q < maxitem(P)

Pseudo-code

Algorithm 1: LCM

Data: dataset D, minimum support threshold ε
Result: Outputs all frequent closed itemsets in $\mathcal{D}$

begin
    $\perp_{closed} \leftarrow \bigcap_{T \in \mathcal{D}} T$
    output $\perp_{closed}$
    foreach $i \in \mathcal{I} \mid i \notin \perp_{closed}$ **do**
        $expand(\perp_{closed}, i, \mathcal{D}, \varepsilon)$

Function $expand(I, i, \mathcal{D}_I, \varepsilon)$
    Data: Closed frequent itemset I, extension item i, reduced dataset D_I, minimum support threshold ε
    Result: Outputs all closed itemsets containing $\{i\} \cup I$
    begin
        if $support_{\mathcal{D}_I}(\{i\}) \geq \varepsilon$ **then** // Frequency test
            $I_{ext} \leftarrow \bigcap_{T \in \mathcal{D}_I[\{i\}]} T$ // Closure computation
            if $maxItem(I_{ext}) = i$ **then** // 1^{st} parent test
                $J \leftarrow I \cup I_{ext}$
                output $(J, support_{\mathcal{D}_I}(\{i\}))$
                $D_J = \{T \setminus J \mid T \in \mathcal{D}_I[\{i\}]\}$
                foreach $j \in \mathcal{I} \setminus J \mid j < i$ **do** // Augmentations
                    $expand(J, j, \mathcal{D}_J, \varepsilon)$

Database Reductions

Conditional database is to reduce database by unnecessary items and transactions, for deeper levels

1,3,5
1,3,5
1,3,5
1,2,5,6
1,4,6
1,2,6

$\theta = 3$

filtering

Remove infrequent items, items included in all

3,5
3,5
3,5
5,6
6
6

filtering

Unify same transactions

3,5 ×3
5,6
6 ×2

Linear time

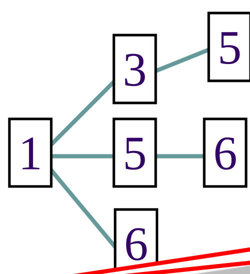

FP-tree, prefix tree

O(||D||log ||D||) time

Compact if database is dense and large

Remove infrequent items, automatically unified

Prize for the Award

Prize is {beer, diapers}

“Most Frequent Itemset”