HowMessy is Your Database?

- A Robelle and Adager Tutorial
  HP World '96
  Anaheim, California
  August 5 - 8, 1996

Adager

What's Inside

How messy is your database?

- A database is messy if it takes more I/O than it should
- Unnecessary I/O is still a major limiting factor even on MPE/iX machines
- Databases are messy by nature
- Run HowMessy or DBLOADNG against your database
  - HowMessy is a bonus program for Robelle customers
  - DBLOADNG is a contributed library program

Blocks

- TurboIMAGE does all I/O operations in blocks
- A block may contain many user records
- More entries per block means fewer I/Os
- Fewer I/Os means better performance

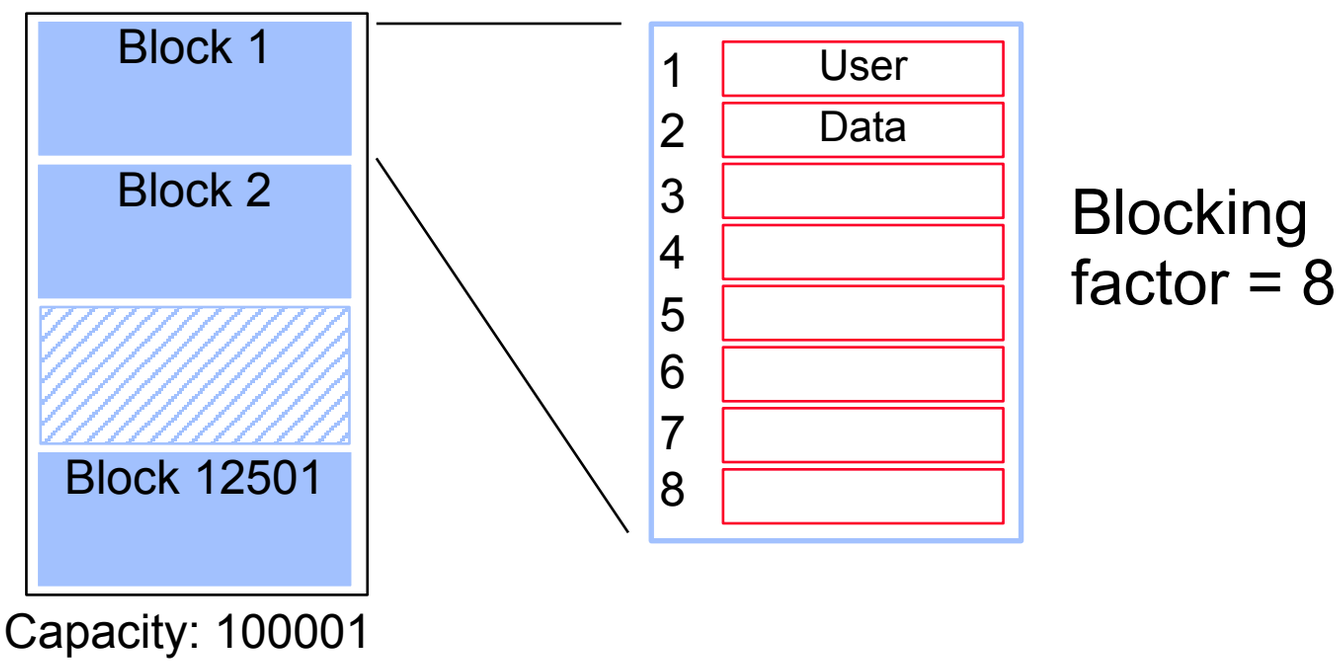

Record location in masters

- Search item values must be unique
- Location of entries is determined by a hashing algorithm or a primary address calculation
- Calculation is done on search item value to transform it into a record number between one and the capacity
- Different calculation depending on the search item type
  - X, U, Z, and P give random results
  - I, J, K, R, and E give predictable results

Hashing algorithm

- Customer number AA1000 is transformed into a record number

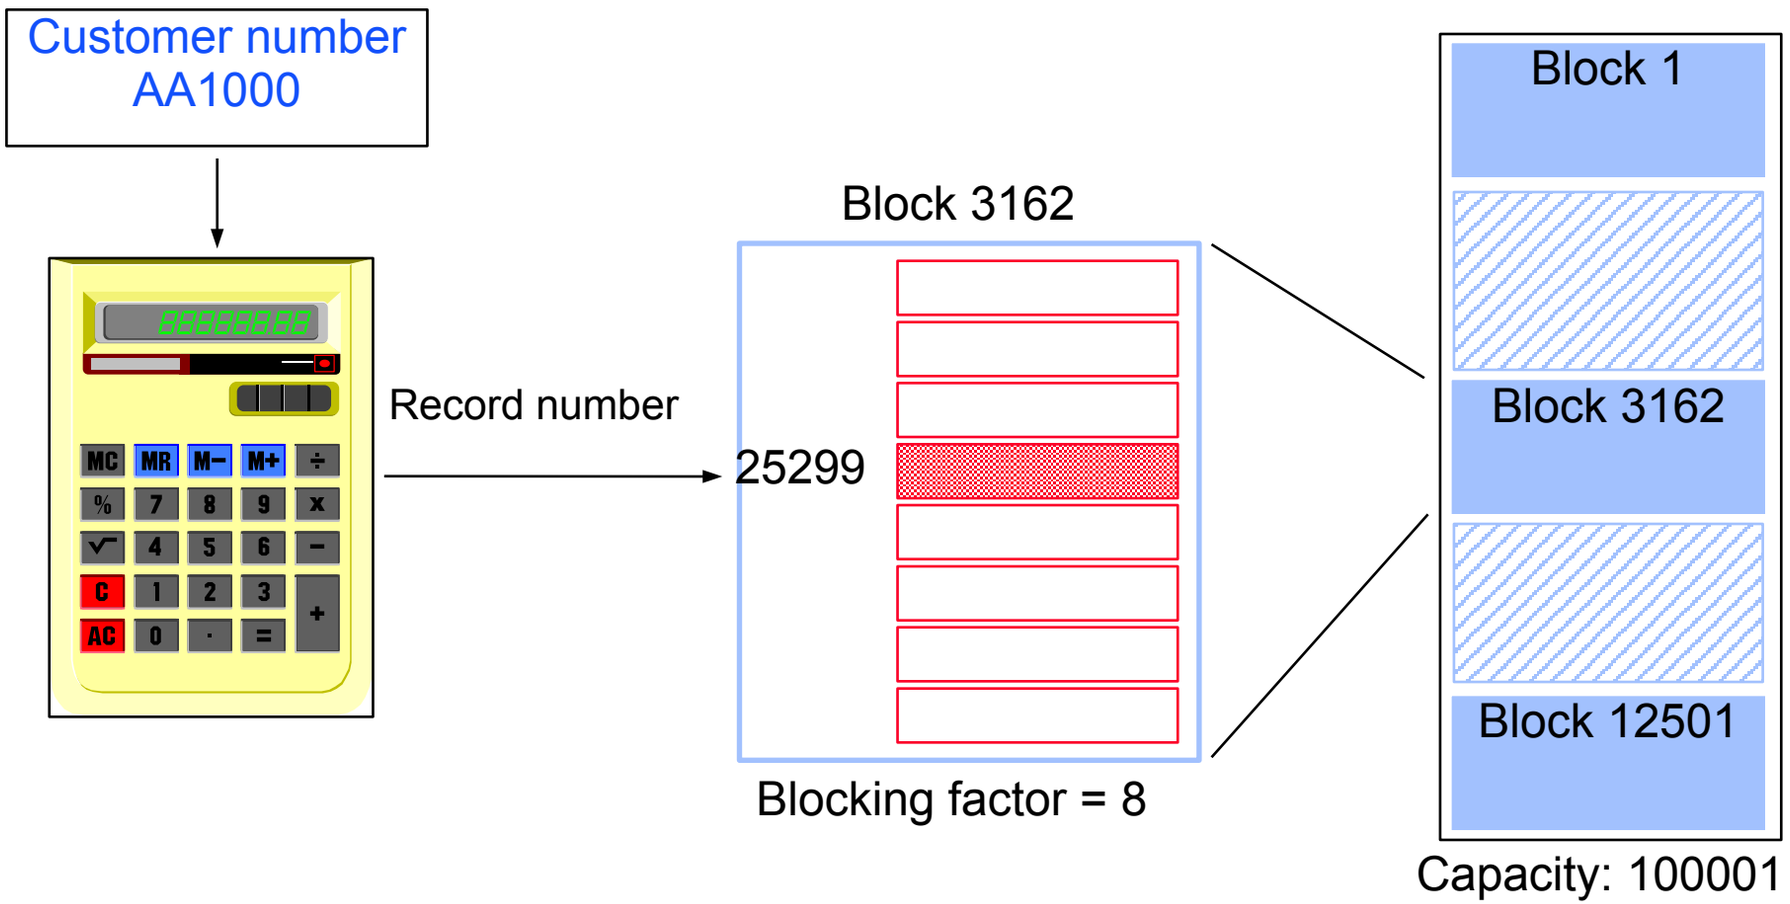

Hashing algorithm (no collision)

- Customer number BD2134 gives a different record number in a different block

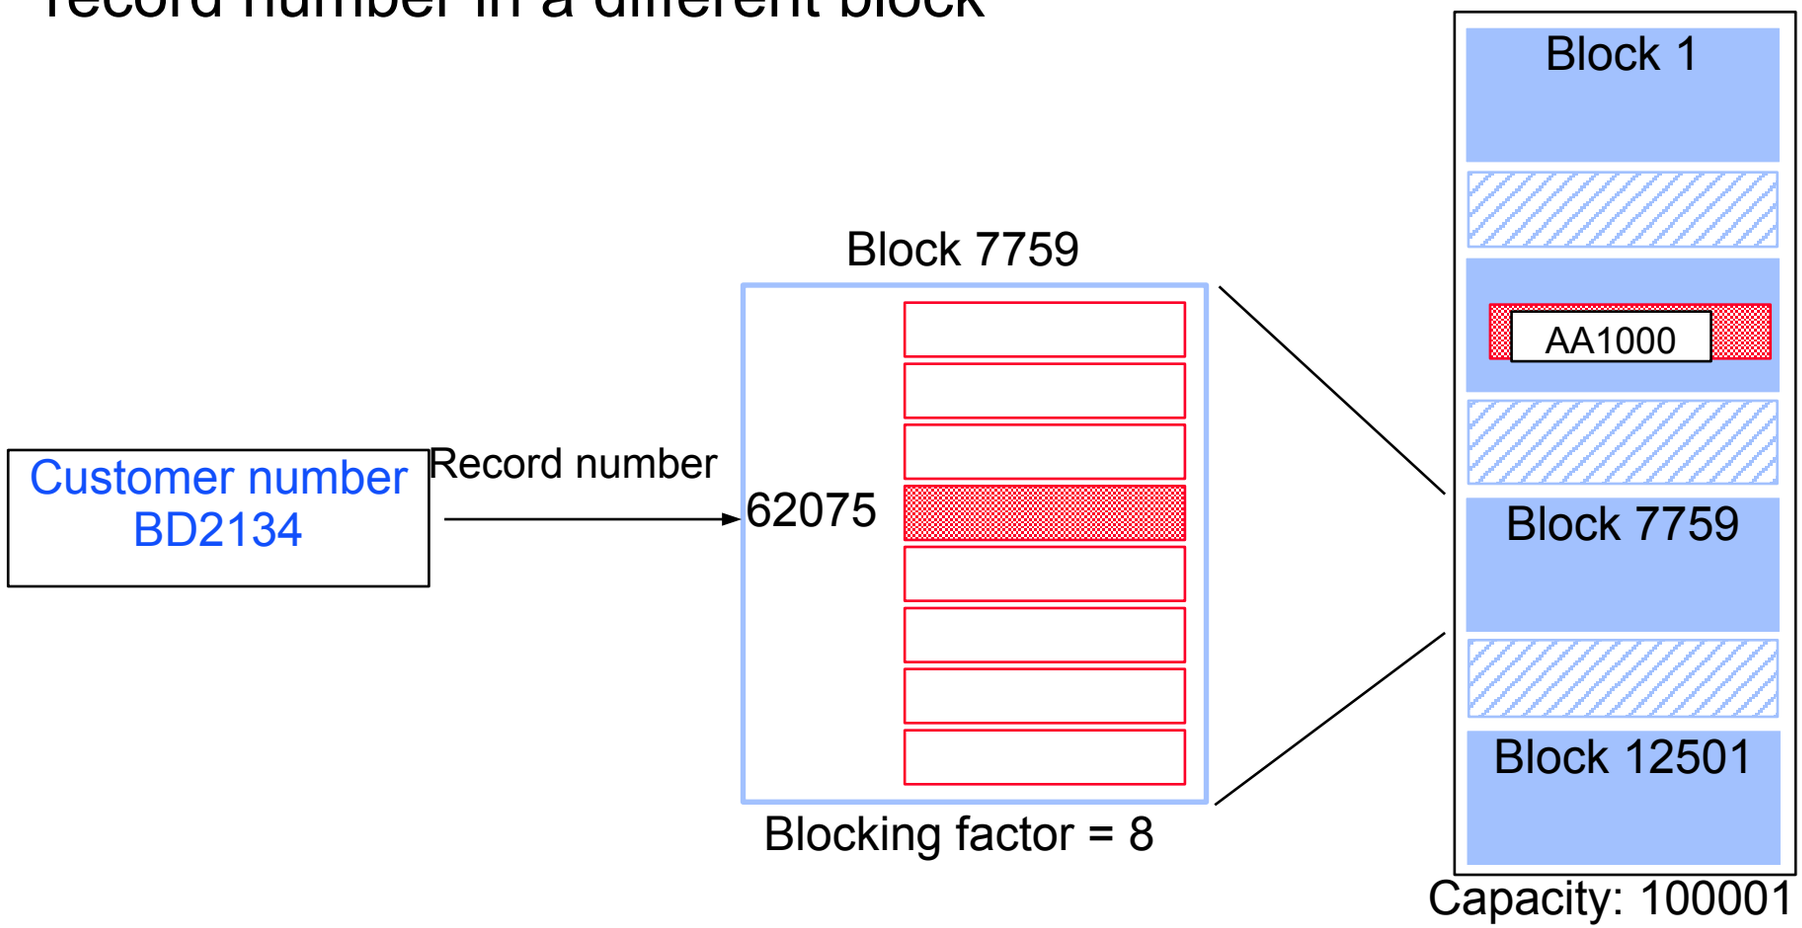

Hashing algorithm (collision - same block)

- Customer number CL1717 hashes to the same record number as AA1000 location
- TurboIMAGE tries to find an empty location in the same block. If it finds one, no additional I/O is required.
- CL1717 becomes a secondary entry. Primary and secondary entries are linked using pointers that form a chain.

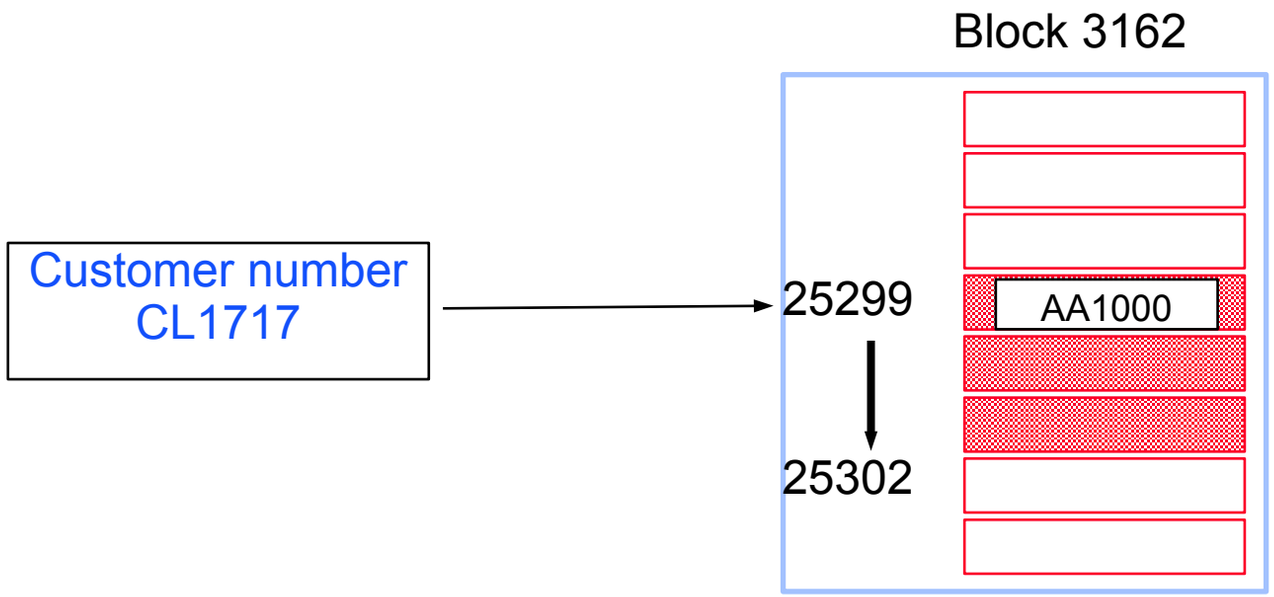

Hashing algorithm (collision - different block)

- Customer number MD4884 collides with AA1000
- No more room in this block. TurboIMAGE reads the following blocks until it finds a free record location.
- In this case, MD4884 is placed two blocks away, which requires two additional I/Os.

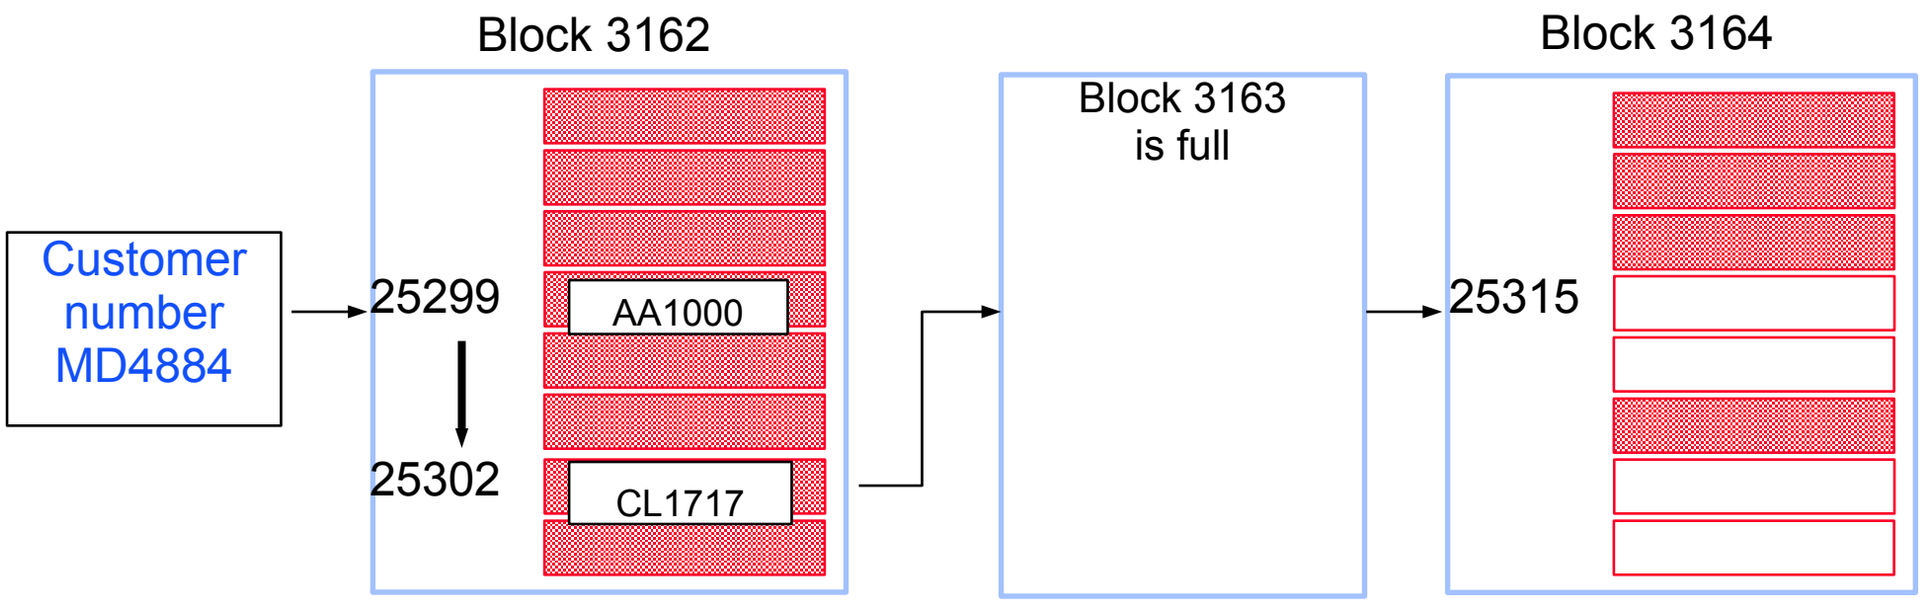

An example TurboIMAGE database

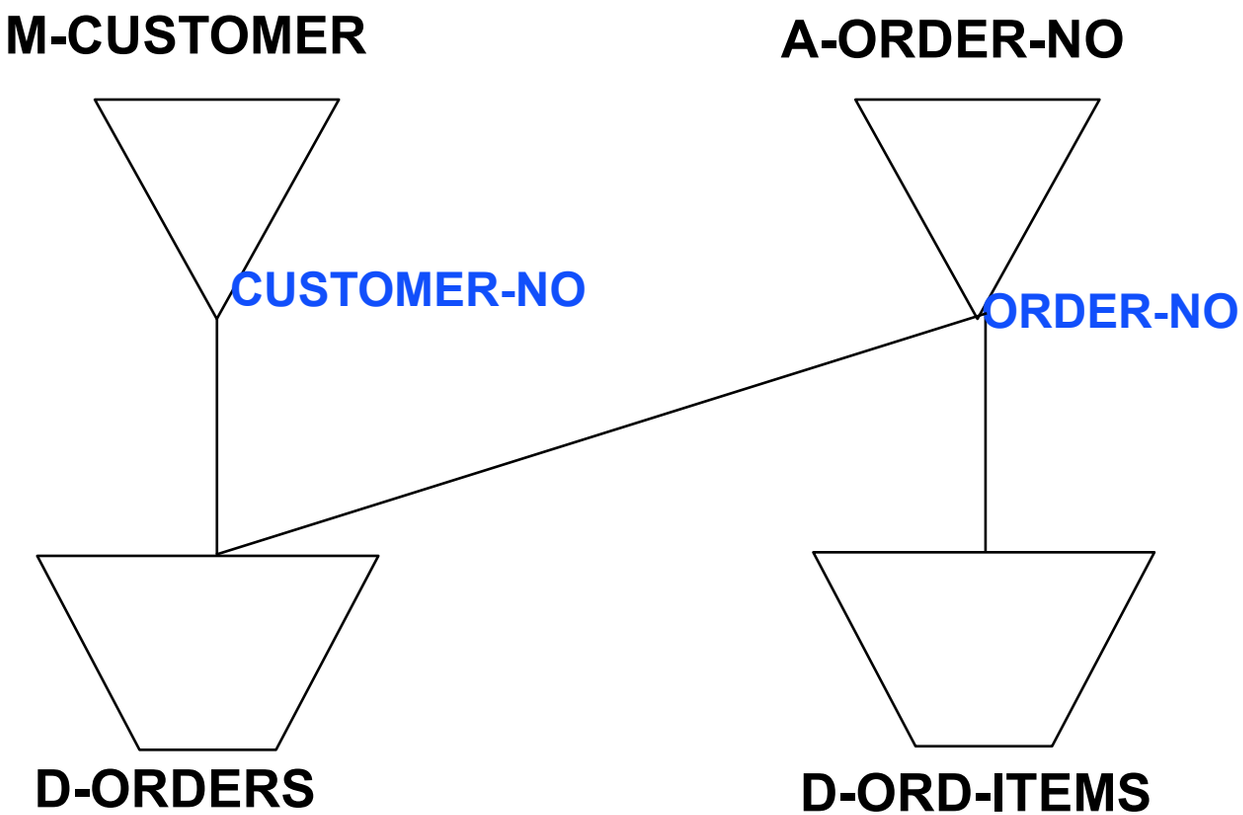

HowMessy sample report

HowMessy/XL (Version 2.2.1) Data Base: STORE.DATA.INVENT Run on: MON, JAN 9, 1996, 11:48 AM
TurboIMAGE/3000 databases By Robelle Solutions Technology Inc. Page: 1

Data Set	Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
M-Customer	Man	248113	178018	71.7%	30.5%	1496	11
A-Order-No	Ato	1266783	768556	60.7%	25.7%	1	70
D-Orders	Det	1000000	768558	76.9%	(	851445)	32
D-Ord-Items	Det	4000000	3458511	86.5%	(	3470097)	23

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
Customer-No	32	1.92	0.32	1.00	1.90	90.5%	1.90
Order-No	10	1.35	0.62	1.00	1.00	0.0%	1.00
!Order-No	1	1.00	0	1.00	1.00	0.0%	1.00
S Customer-No	80	14.34	17.76	1.75	9.20	57.2%	5.25
S !Order-No	1604	8.06	35.75	1.36	11.32	72.5%	8.34

HowMessy sample report (master dataset)

HowMessy/XL (Version 2.2.1) Data Base: STORE.DATA.INVENT Run on: MON, JAN 9, 1996, 11:48 AM
TurboIMAGE/3000 databases By Robelle Solutions Technology Inc. Page: 1

Data Set	Type	Capacity	Entries	Load Factor	Secon- daries (Highwater)	Max Blks	Blk Fact
M-Customer	Man	248113	178018	71.7%	30.5% 1496	11	
A-Order-No	Ato	1266783	768556	60.7%	25.7% 1	70	
D-Orders	Det	1000000	768558	76.9%	(851445)	32	
D-Ord-Items	Det	4000000	3458511	86.5%	(3470097)	23	

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
Customer-No	32	1.92	0.32	1.00	1.90	90.5%	1.90
Order-No	10	1.35	0.62	1.00	1.00	0.0%	1.00
!Order-No	1	1.00	0	1.00	1.00	0.0%	1.00
S Customer-No	80	14.34	17.76	1.75	9.20	57.2%	5.25
S !Order-No	1604	8.06	35.75	1.36	11.32	72.5%	8.34

Interpreting master datasets lines

- Pay attention to the following statistics:
  - High percentage of Secondaries (inefficient hashing)
  - High Maximum Blocks (clustering)
  - High Maximum and Average Chains (inefficient hashing)
  - High Inefficient Pointers (when secondaries exist)
  - High Elongation (when secondaries exist)

Report on m-customer

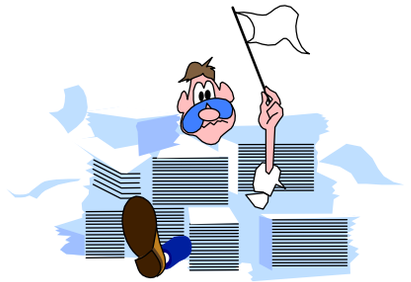

- The number of Secondaries is not unusually high
- However, there may be problems
  - Records are clustering (high Max Blks)
  - Long synonym chain
  - High percentage of Inefficient Pointers

Data Set	Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
M-CUSTOMER	Man	248113	178018	71.7%	30.5%	1496	11

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
CUSTOMER-NO	22	1.92	0.32	1.00	1.90	90.5%	1.90

Report on a-order-no

- Very tidy dataset
  - Number of Secondaries is acceptable
  - Max Blks, Ineff Ptrs and Elongation are at the minimum values, even if the Maximum Chain length is a bit high

Data Set	Type	Capacity	Entries	Load Factor	Secon-daries (Highwater)	Max Blks	Blk Fact
A-ORDER-NO	Ato	1266783	768556	60.7%	25.7%	1	70

Search Field	Elong-ation	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs
ORDER-NO	1.00	10	1.35	0.62	1.00	1.00	0.0%

Master dataset solutions

- Increase capacity to a higher odd number
- Increase the Blocking Factor
  - Increase block size
  - Reduce record size
- Change binary keys to type X, U, Z, or P
- Check your database early in the design
- Use HowMessy on test databases

Implementing a solution: Changing a key data-type

HowMessy/XL (Version 2.2.1) Data Base: SOMEDB.DATA.SOMEACCT Run on: MON, JAN 9, 1996, 11:48 AM
TurboIMAGE/3000 databases By Robelle Solutions Technology Inc. Page: 1

Data Set		Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
Part-Master	Ato	10000	4305	43.0%	14.2%	16	78	
Part-Loc-Master	Ato	606010	303005	50.0%	86.4%	612	67	
Parts-Detail	Det	303030	303005	100.0%	(303005)		63	

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
Part-No	3	1.17	0.40	1.00	1.67	63.3%	1.67
Part-Loc-No	4305	7.38	73.09	1.07	45.84	66.4%	42.90
Part-No	41065	70.38	666.46	1.86	3.20	3.1%	1.72
!Part-Loc-No	1	1.00	0.00	1.00	1.00	0.0%	1.00

Change key type to byte, same data values

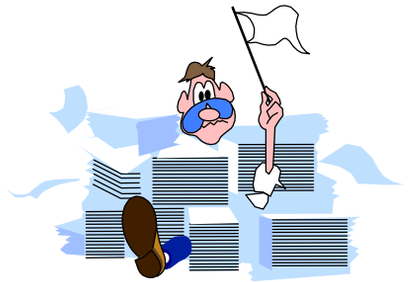

HowMessy/XL (Version 2.2.1) Data Base: SOMEDB.DATA.SOMEACCT Run on: MON, JAN 9, 1996, 11:48 AM
TurboIMAGE/3000 databases By Robelle Solutions Technology Inc. Page: 1

Data Set		Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
Part-Master	Ato	10000	4305	43.0%	14.2%	16	78	
Part-Loc-Master	Ato	606010	303005	50.0%	27.7%	95	67	
Parts-Detail	Det	303030	303005	100.0%	(303005)		63	

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
Part-No	3	1.17	0.40	1.00	1.67	63.3%	1.67
Part-Loc-No	13	1.38	0.77	1.00	1.42	29.3%	1.42
Part-No	41065	70.38	666.46	1.86	3.20	3.1%	1.72
!Part-Loc-No	1	1.00	0.00	1.00	1.00	0.0%	1.00

Change key type to byte, convert data values

HowMessy/XL (Version 2.2.1) Data Base: SOMEDB.DATA.SOMEACCT Run on: MON, JAN 9, 1996, 11:48 AM
TurboIMAGE/3000 databases By Robelle Solutions Technology Inc. Page: 1

Data Set		Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
Part-Master	Ato	10000	4305	43.0%	14.2%	16	78	
Part-Loc-Master	Ato	606010	303005	50.0%	21.4%	0	56	
Parts-Detail	Det	303054	303005	100.0%	(303005)		53	

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
Part-No	3	1.17	0.40	1.00	1.67	63.3%	1.67
Part-Loc-No	6	1.27	0.54	1.00	1.00	0.0%	1.00
Part-No	41065	70.38	666.46	2.07	3.43	3.4%	1.66
!Part-Loc-No	1	1.00	0.00	1.00	1.00	0.0%	1.00

HowMessy sample report (detail dataset)

HowMessy/XL (Version 2.2.1) for TurboIMAGE/3000 databases | Data Base: STORE.DATA.INVENT By Robelle Solutions Technology Inc. | Run on: MON, JAN 9, 1996, 11:48 AM Page: 1

Data Set	Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
M-CUSTOMER	Man	248113	178018	71.7%	30.5%	1496	1
A-ORDER-NO	Ato	126673	768556	60.7%	25.7%	1	70
D-ORDERS	Det	1000000	768556	76.9%		(851445)	12
D-ORD-ITEMS	Det	4000000	3458511	86.5%		(3470097)	23

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
Customer-No	22	1.92	0.32	1.00	1.90	90.5%	1.90
Order-No	10	1.35	0.62	1.00	1.00	0.0%	1.00
!Order-No	1	1.00	0	1.00	1.00	0.0%	1.00
S Customer-No	80	14.34	17.76	1.75	9.20	57.2%	5.25
S !Order-No	1604	8.06	35.75	1.36	11.32	72.5%	8.34

Empty detail dataset

- Records are stored in the order they are created, starting from record 1
- Records for the same customer are linked together using pointers to form a chain
- Chains are linked to the corresponding master entry

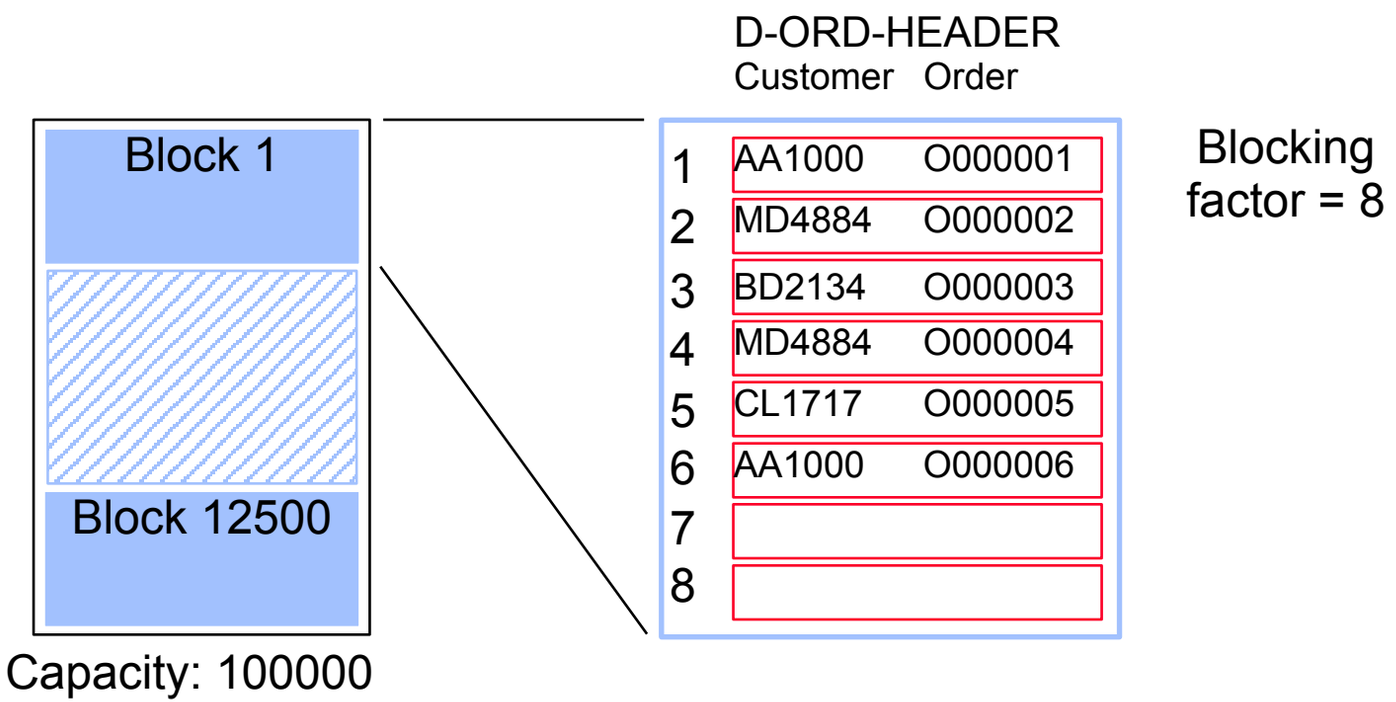

Detail chains get scattered

- Over time, records for the same customer are scattered over multiple blocks

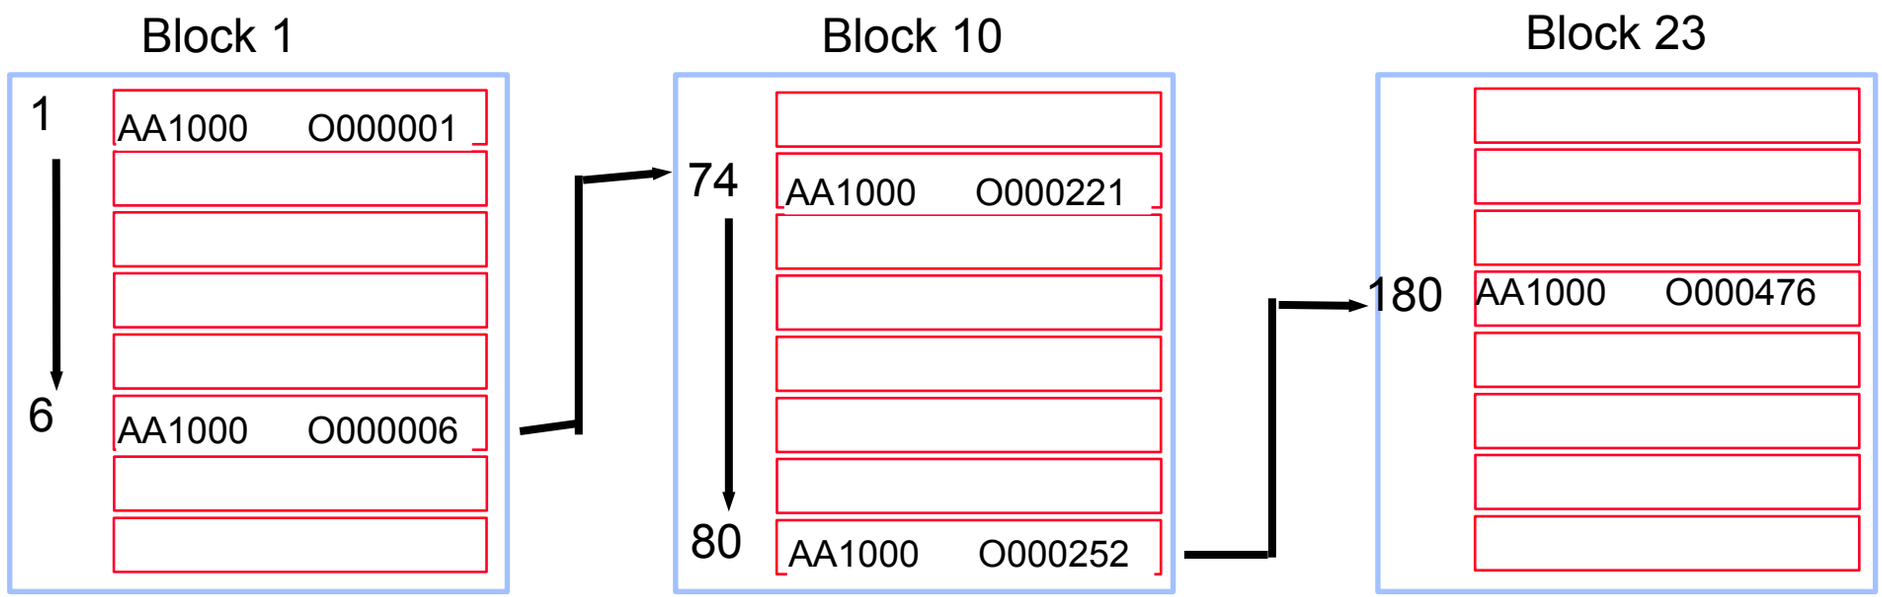

Delete chain

- Deleted records are linked together
- TurboIMAGE reuses the locations in the Delete chain, if there are any

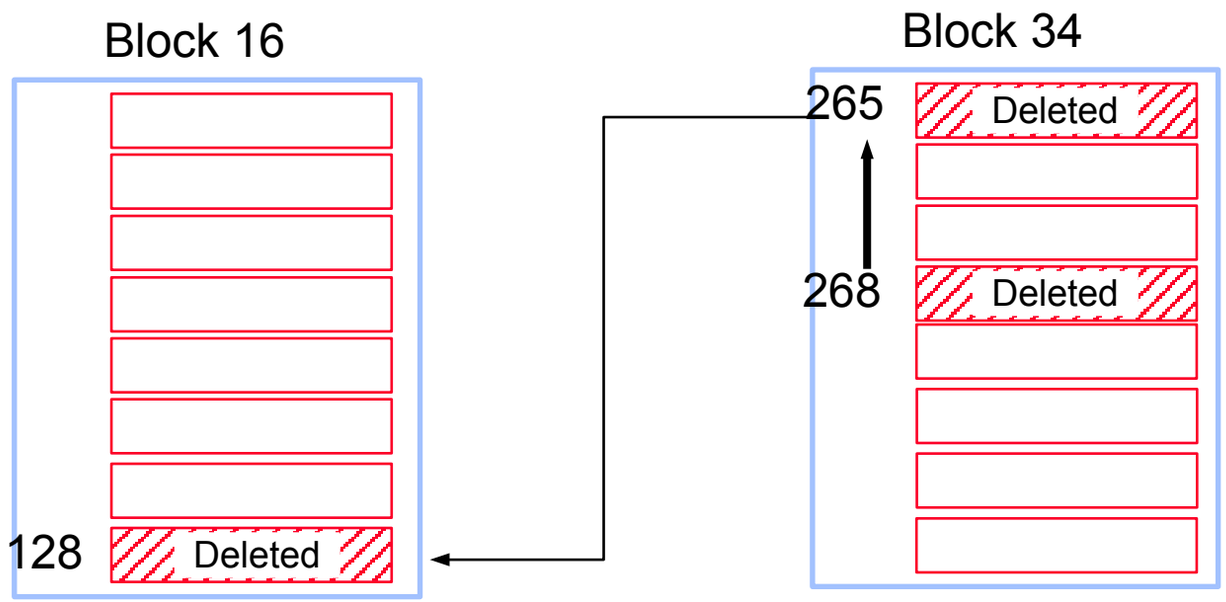

Highwater mark

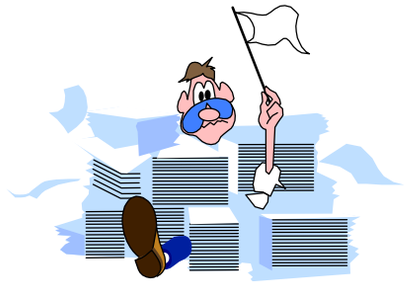

- Indicates highest record location used so far
- Serial reads scan the dataset up to the highwater mark

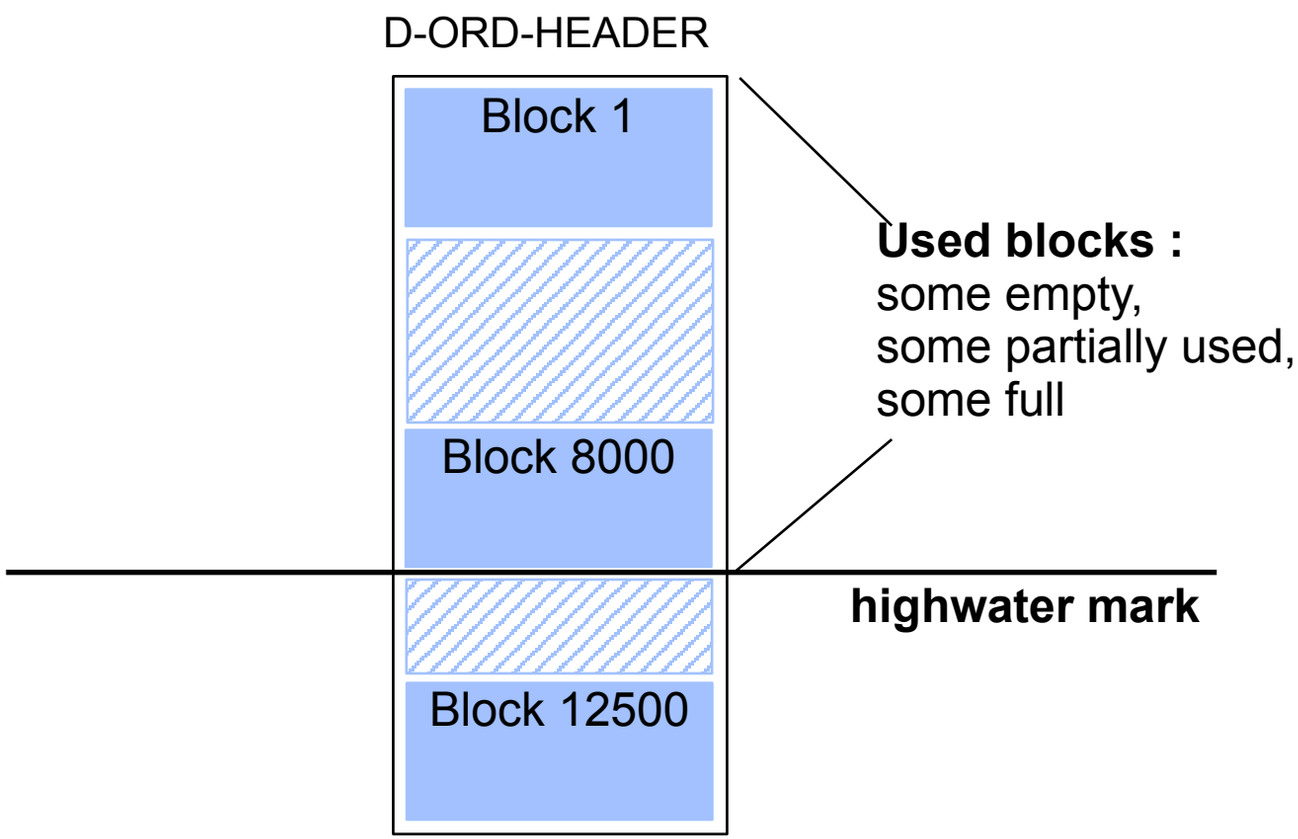

Repacking a detail dataset

- Groups records along primary path
- Removes Delete chain (no holes)
- Resets highwater mark

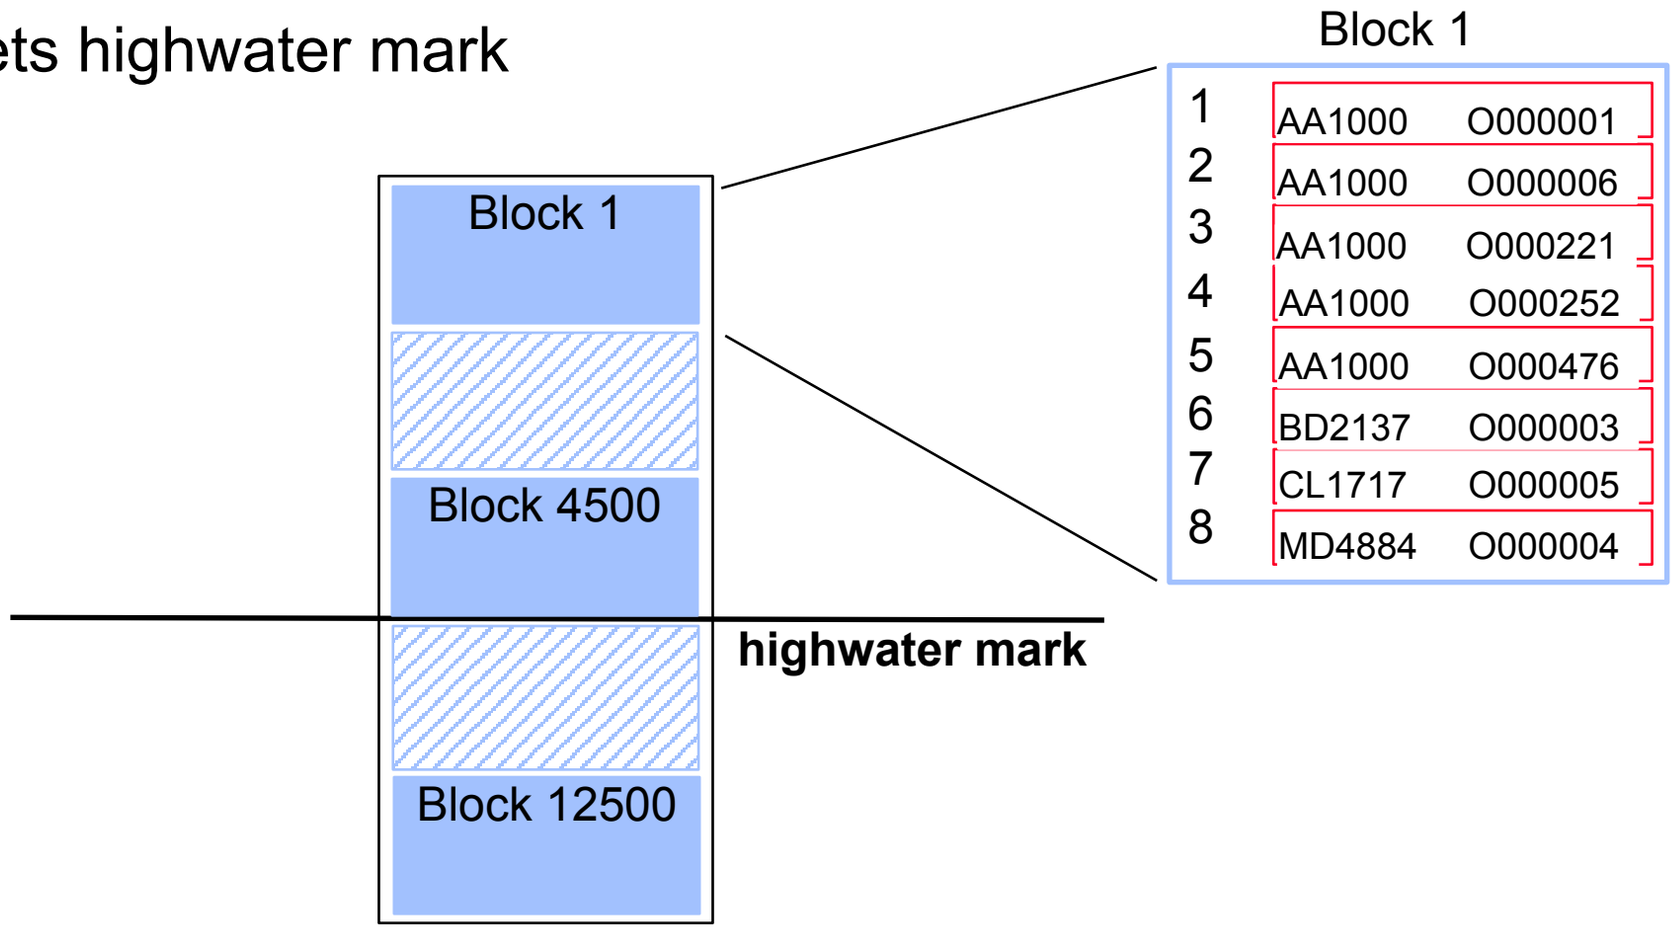

Interpreting detail dataset lines

- Pay attention to the following statistics:
  - Load Factor approaching 100% (dataset full)
  - Primary path (large Average Chain and often accessed)
  - High Average Chain and low Standard Deviation, especially with a sorted path (Is path really needed?)
  - High Inefficient Pointers (entries in chain not consecutive)
  - High Elongation (entries in chain not consecutive)

Report on d-orders

- Primary path should be on customer-no, not on order-no
- Highwater mark is high
- Repack along new primary path regularly

Data Set	Type	Capacity	Entries	Load Factor	Secon-Max daries Blks (Highwater)	Blk Fact
D-ORDERS	Det	1000000	768556	76.9%	(851445)	12

Search Field		Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong-ation
!ORDER-NO		1	1.00	0	1.00	1.00	0.0%	1.00
S 5.25	CUSTOMER-NO	80	14.34		17.76	1.75	9.20	57.2%

Report on d-ord-items

- Inefficient Pointers and Elongation are high
- Highwater mark is fairly high
- Repack the dataset regularly
- Is the sorted path really needed?

Data Set	Type	Capacity	Entries	Load Factor	Secon- daries	Max Blks (Highwater)	Blk Fact
D-ORD-ITEMS	Det	4000000	3458511	86.5%	(	3470097)	23

Search Field	Max Chain	Ave Chain	Std Dev	Expd Blocks	Avg Blocks	Ineff Ptrs	Elong- ation
S !ORDER-NO	1604	8.06	35.75	1.36	11.32	72.5	8.34

Detail dataset solutions

- Assign the primary path correctly; select a search item with an Average Chain length > 1 that is accessed most often
- Repack datasets along the primary path regularly
- Increase the Blocking Factor
  - Increase block size
  - Reduce record size
- Understand sorted paths
- Check your databases early in the design; use HowMessy on test databases

Minimum number of disc I/Os

Intrinsic	Disc I/Os
DBGET	1
DBFIND	1
DBBEGIN	1
DBEND	1
DBUPDATE	1 (non-critical item)
DBUPDATE	13 (critical item)
DBPUT	3 [+ (4 x #paths, if detail)]
DBDELETE	2 [+ (4 x #paths, if detail)]

Serial reads:

Master	Capacity / Blocking factor
Detail	# entries / Blocking factor

Estimating response time

- Deleting 100,000 records from a detail dataset with two paths would take:
  - 2 + (4 x 2 paths) = 10 I/Os per record
  - 100,000 records x 10 I/Os per record = 1,000,000 I/Os
- Classic: around 25 I/Os per second
  - 1,000,000 I/Os / 25 = 40,000 seconds
  - 40,000 seconds / 3600 = 11.1 hours
- iX: around 40 I/Os per second
  - 1,000,000 I/Os / 40 = 25,000 seconds
  - 25,000 seconds / 3600 = 6.9 hours

Automating HowMessy analysis

- Recent version of HowMessy creates a self-describing file with these statistics
- Process the file with generic tools (Suprtool, AskPlus) or custom programs (COBOL, 4GL), and produce custom reports
- Send messages to database administrators
- Write “smart” job to fix databases without user intervention

Processing Loadfile with Suprtool

- Datasets more than 80% full

```
>input  loadfile
>if loadfactor > 80
>ext database, dataset, datasettype, loadfactor
>list standard
```

- Only one address per customer

```
>input  loadfile
>if  dataset = "D-ADDRESSES" and &
  maxchain > 1
```

References

- The TurboIMAGE/3000 Handbook (Chapter 23)

- Available for $ 49.95 from:

WORDWARE
P.O. Box 14300
Seattle, WA 98114

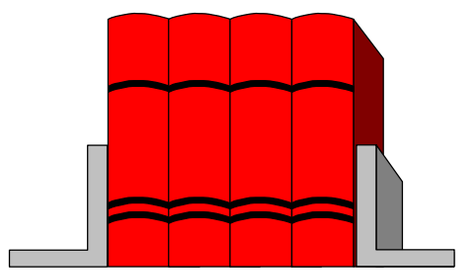

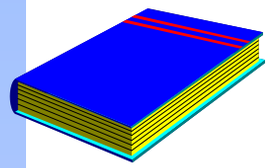

Summary

- TurboIMAGE databases become messy over time, especially if they are active
- HowMessy and DBLOADNG let you analyze the database's efficiency
- You should have some knowledge of the internal workings of TurboIMAGE
- Monitor your databases regularly