

Transactions interaction and best practices for application development

Firebird Conference 2019

Berlin, 17-19 October



YOUR PREMIER SOURCE OF FIREBIRD SUPPORT

IBSurgeon



**MOSCOW
EXCHANGE**

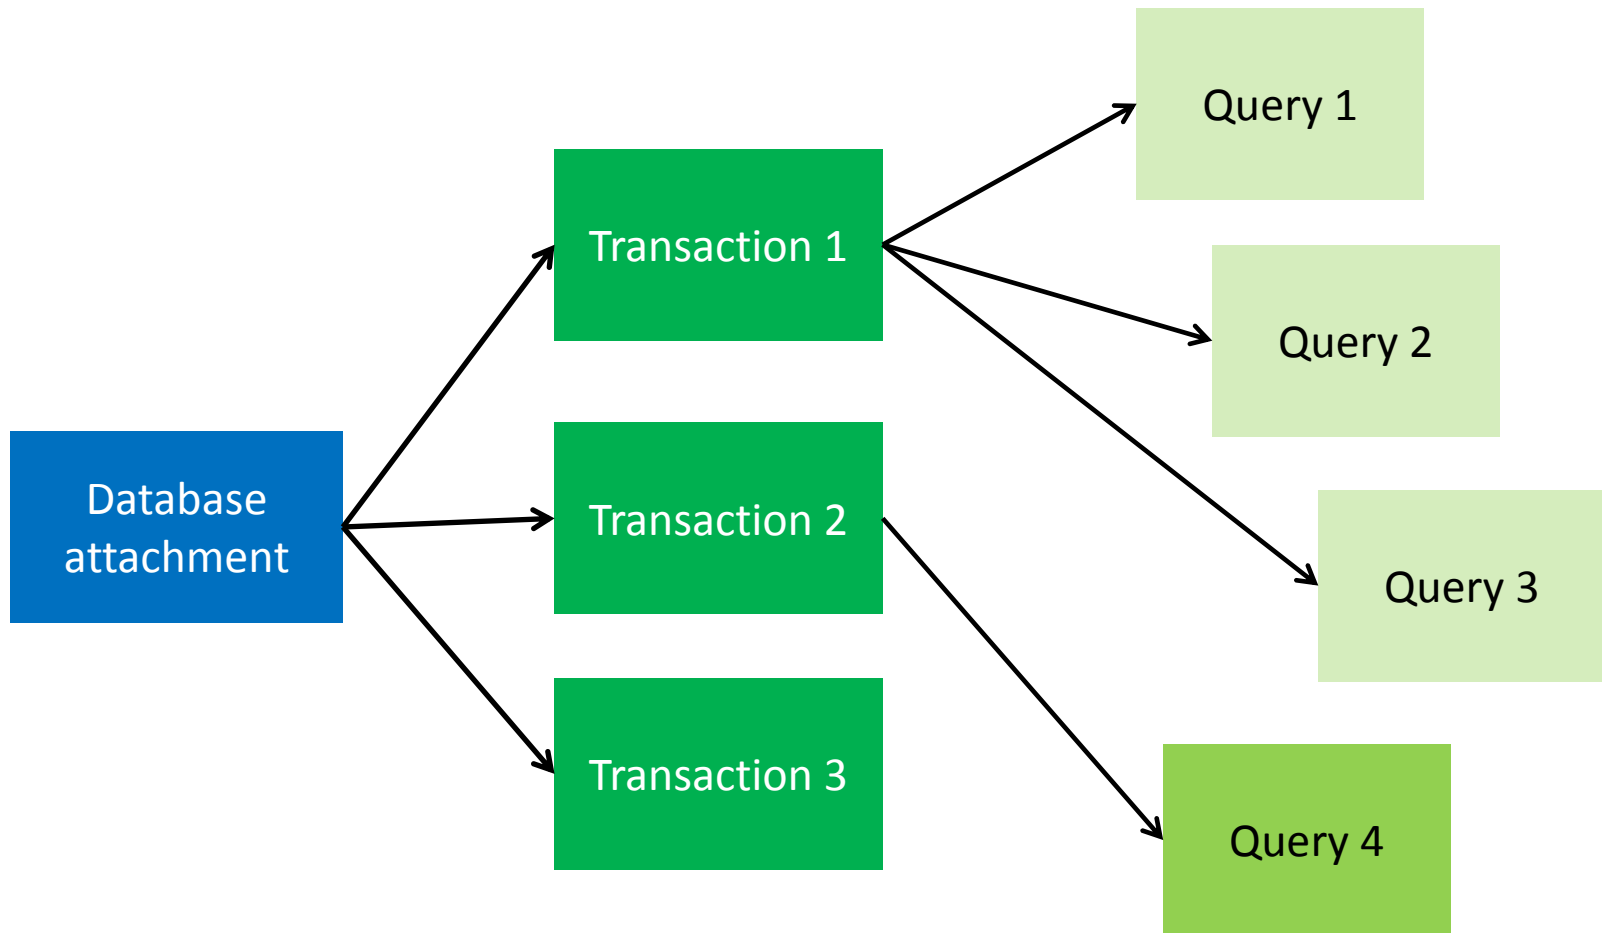


Fast Reports
Reporting must be fast!



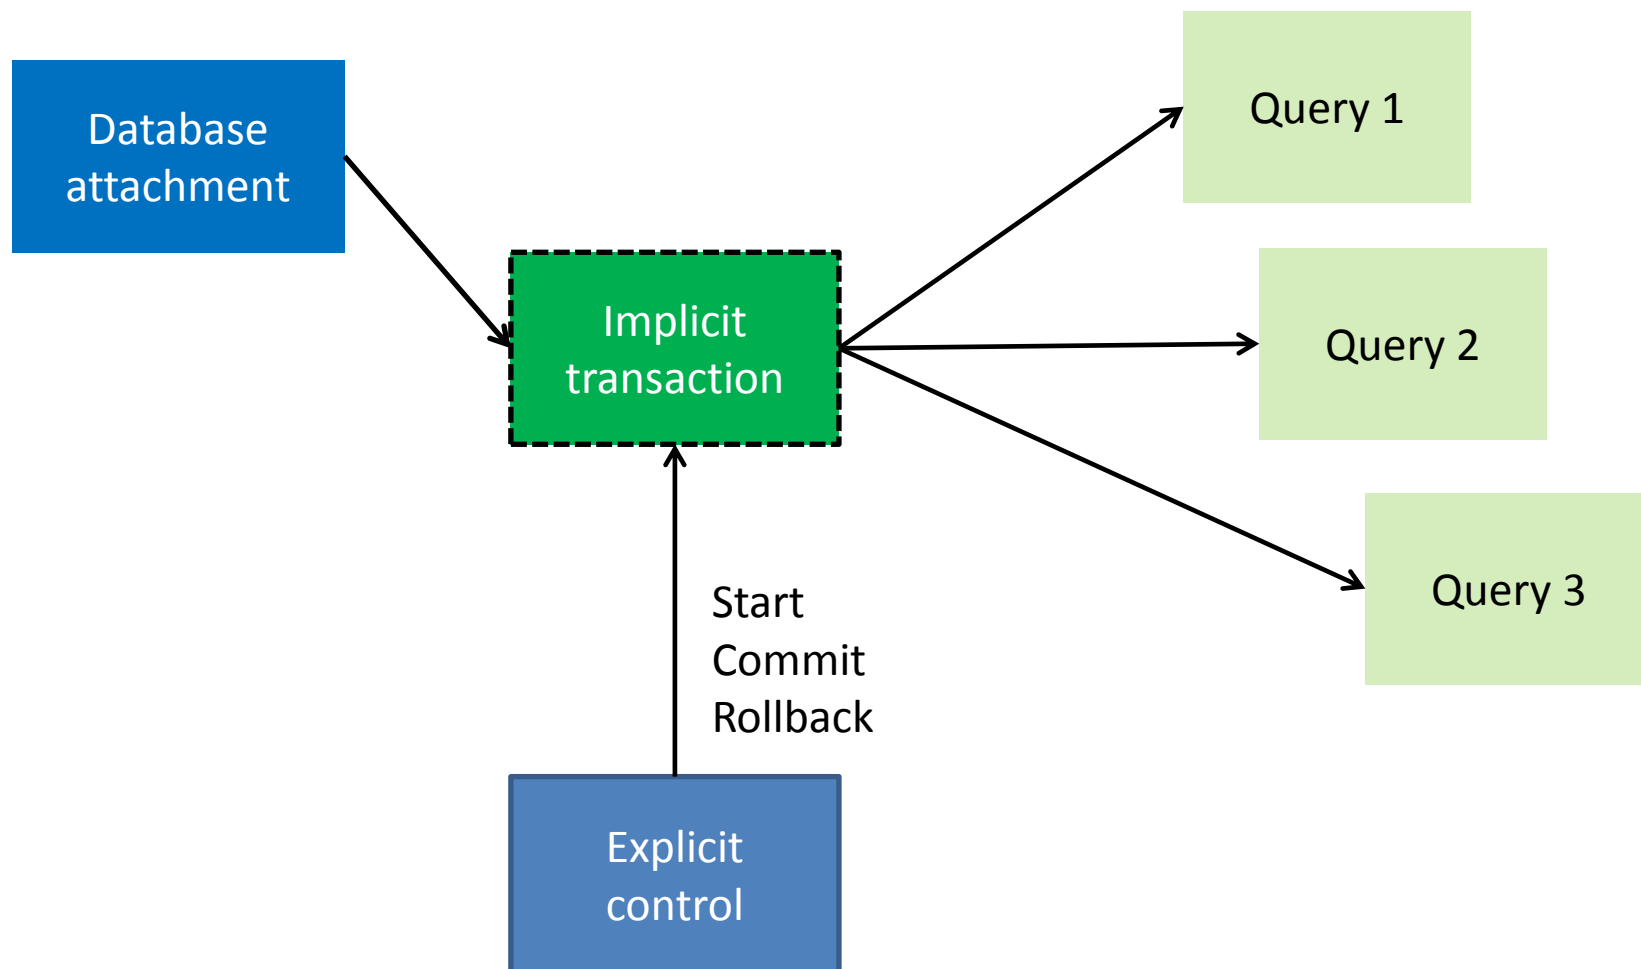
TRANSACTIONS LIFETIME IN VARIOUS DRIVERS

InterBase and Firebird API



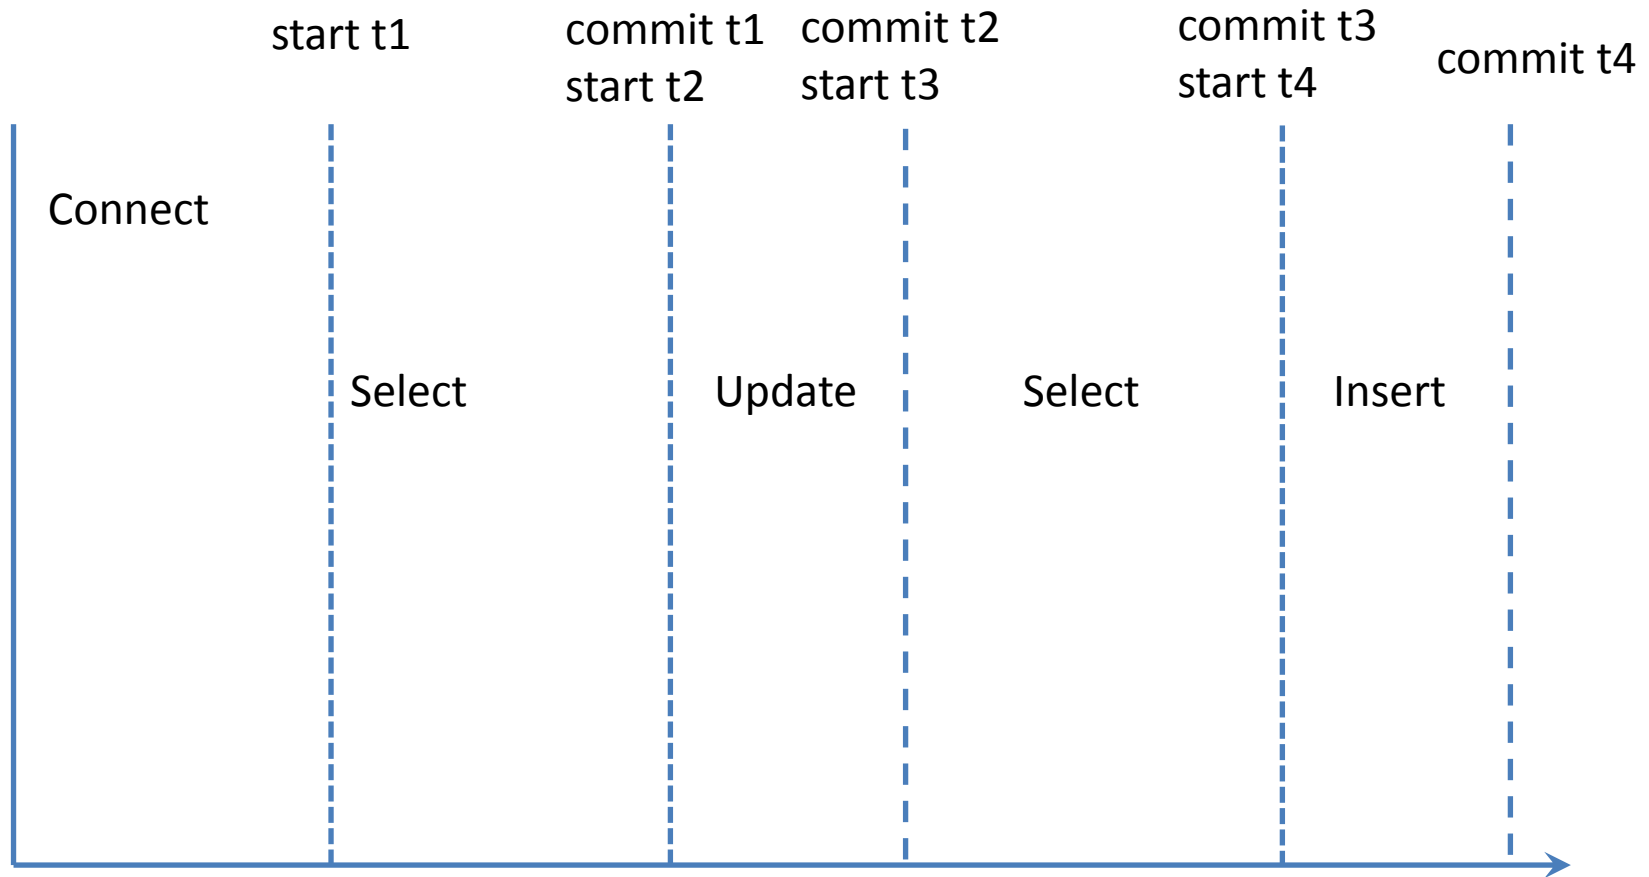
Handles!

Most standard drivers



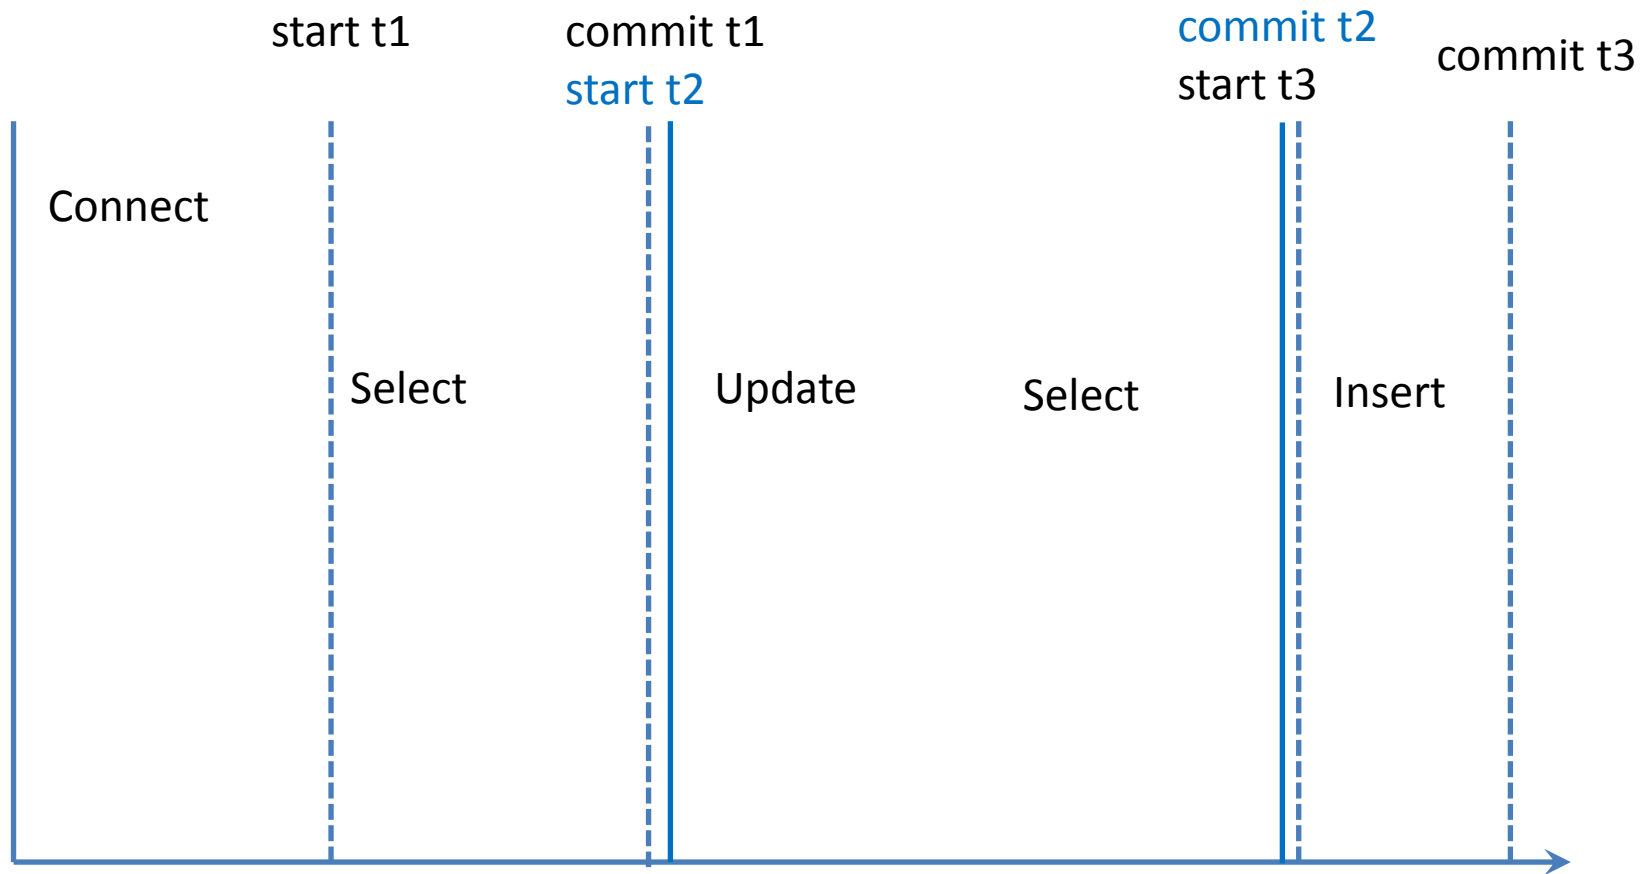
Implicit transactions

- Hidden. You do not see them, you do not control them
- Any SQL statement causes transaction start
- Autocommit mode
 - Each successful INSERT, UPDATE, DELETE and Execute Procedure causes automatic Commit. Any error causes Rollback.
 - SELECT statements may not be committed, until Insert, Update, Delete or Execute procedure.
- You cannot define set of SQL DML statements as a real transaction
- Transaction with Select statements can run forever
- Transaction may be ended by Retaining
- Connect may start transaction immediately



Explicit transactions

- You may (or not) call `StartTransaction`
- All SQL statements will execute in that transaction
- You must end transaction with explicit `Commit` or `Rollback`



One transaction per connect

- Common driver architecture
 - BDE
 - ODBC
 - JDBC
 - DBExpress
 - .Net driver
 - ...
- Implicit transactions by default
- You may start explicit transactions
- Transactions can live (be active) for a very long time

BDE

- DataSet.Open;
- Query.ExecSQL;
- Database.StartTransaction
 - Query1.ExecSQL;
 - DataSet1.Open;
- Database.Commit;

Many transactions per connect

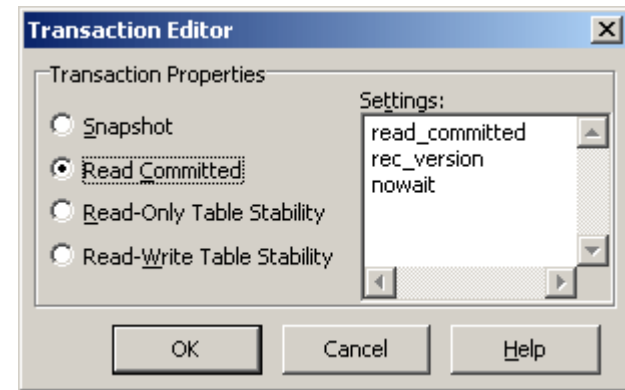
- IBX, FIBPlus, FireDAC (AnyDAC), UIB, ...
- Total control on transaction parameters
- Many transactions per connect
- Ability to use DataSets for read in one transaction, and write in another

IBX

- Transaction1.StartTransaction;
- IBQuery1.ExecSQL;

- Transaction2.StartTransaction;
- IBQuery2.ExecSQL;

- Transaction1.Commit;
- Transaction2.Rollback;



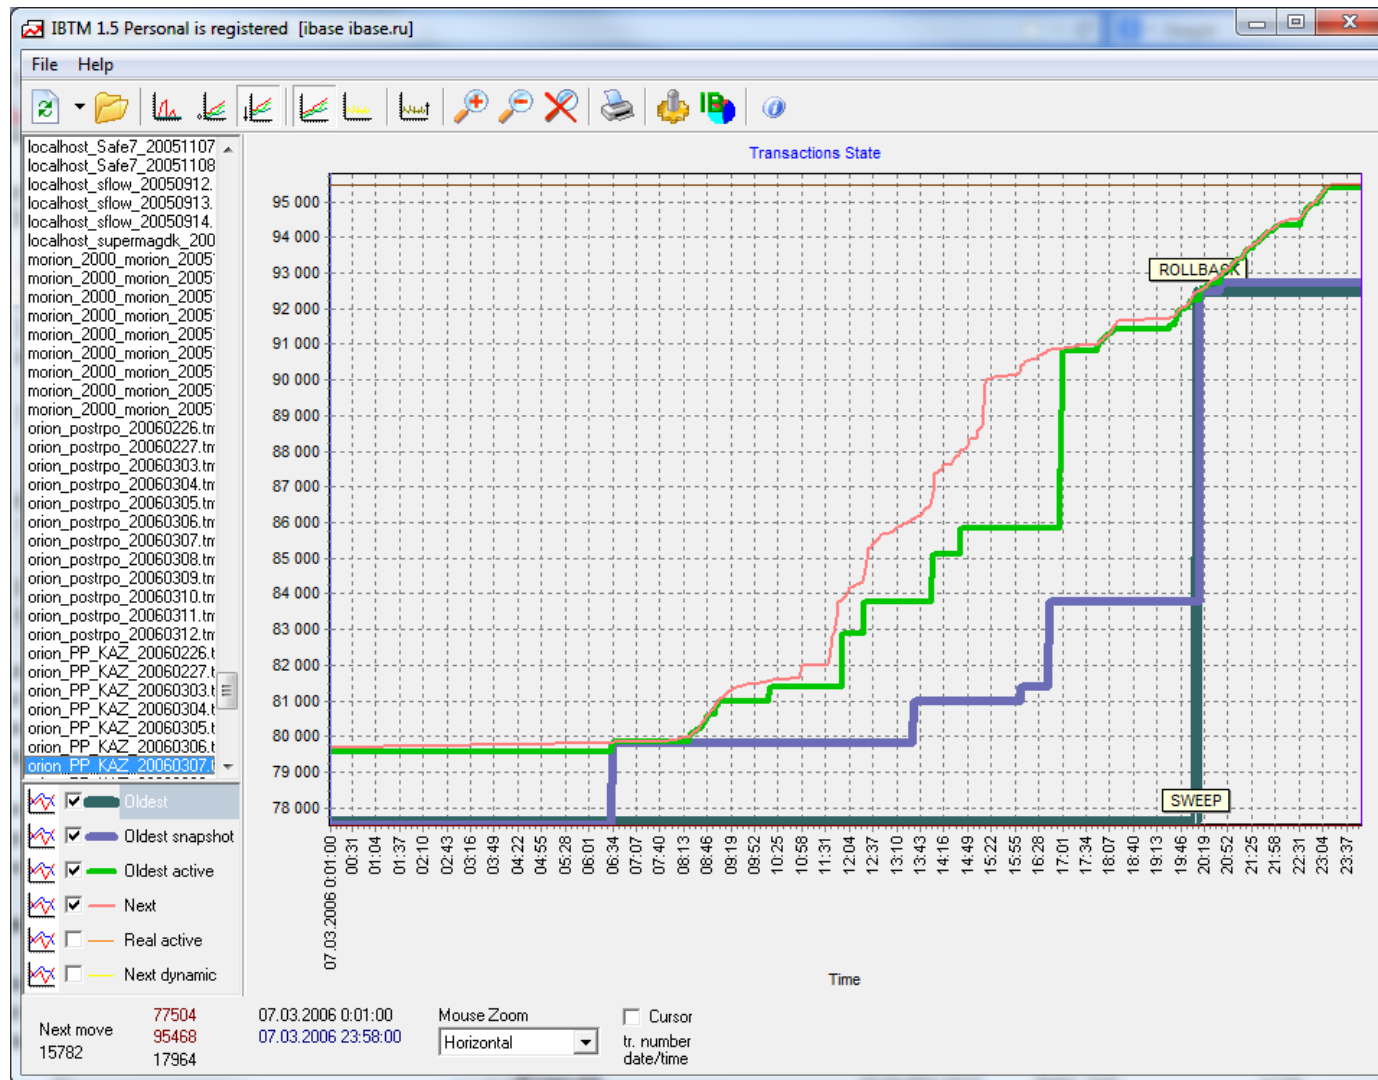
Long reading transactions

- from InterBase 6.0
- read
nowait
read_committed
rec_version
- This can be running forever

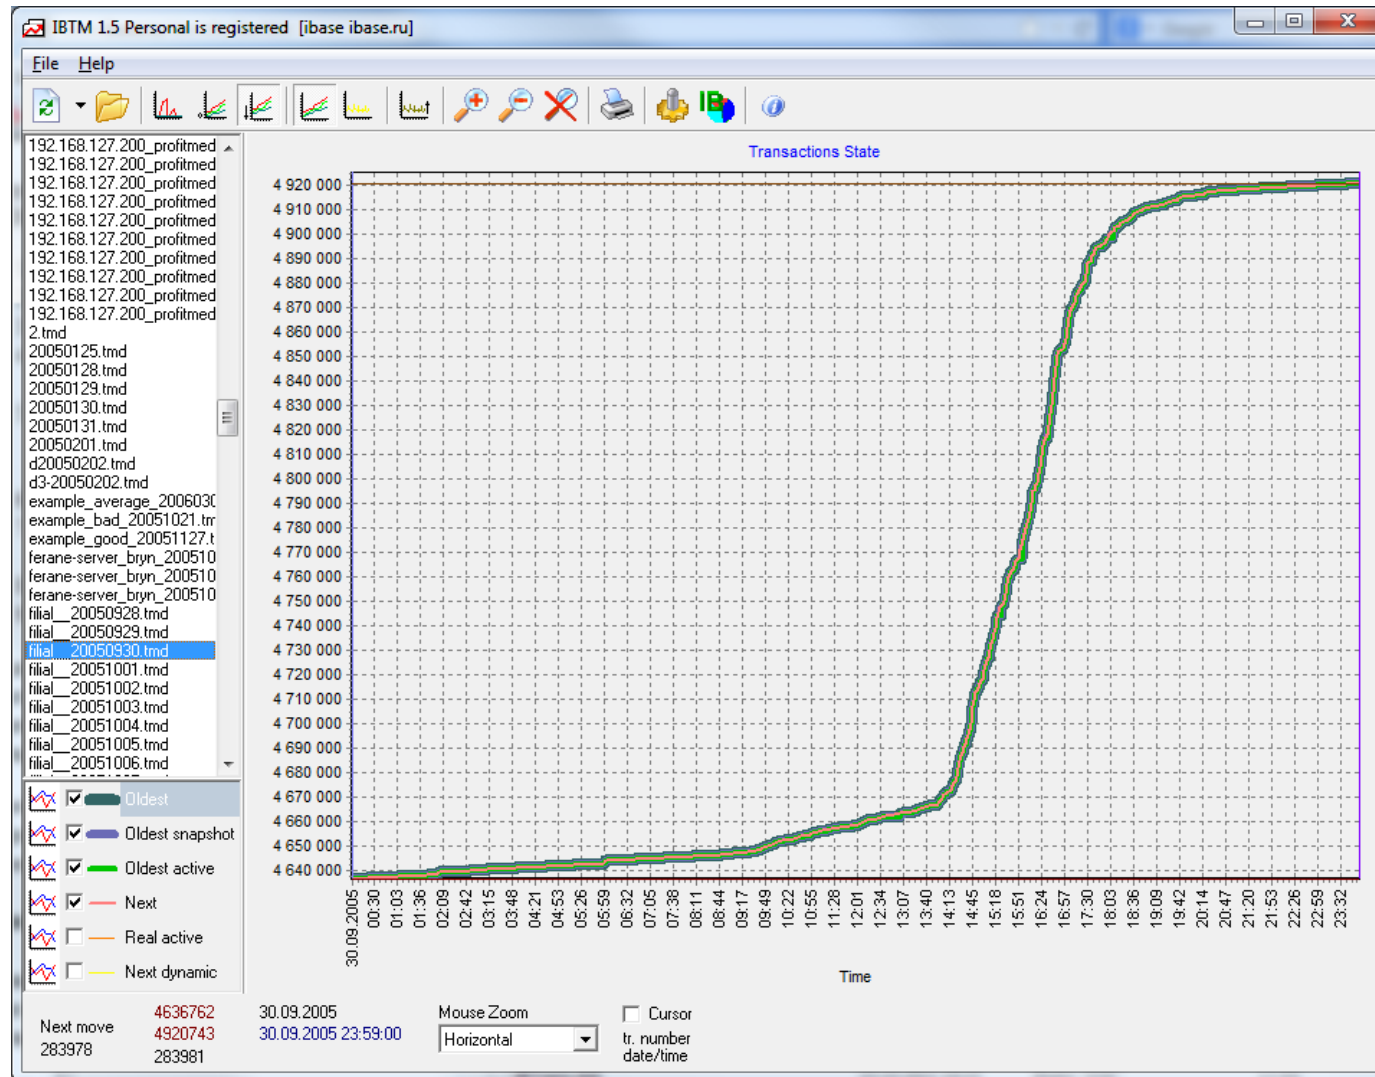
dbExpress – handles?

- Cannot set parameters
- TDBXTransaction exists, but useless
- Cannot switch between transactions
- `transaction1.BeginTransaction;`
- ...
- `transaction2.BeginTransaction;`
- ... here you can not return to `transaction1` context, you can only call it's `commit/rollback`.

Bad transaction control



Perfect transaction control



Rules for native components/drivers

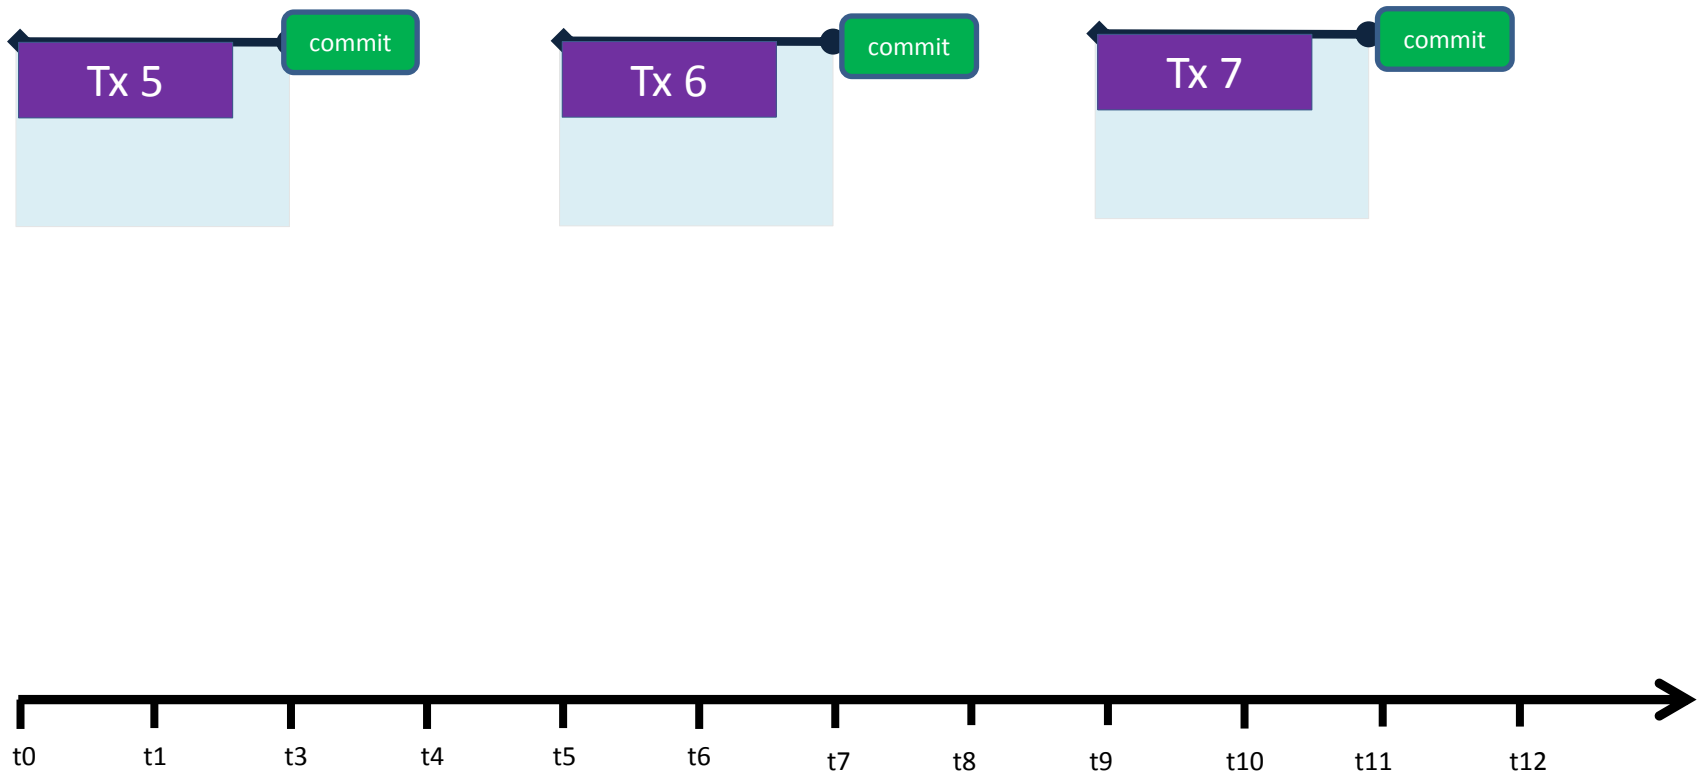
- Do not use “default” transaction. Always use explicit transaction control.
- Do not use “default transaction parameters”
 - Default may be ReadCommitted or Snapshot, you will never know
- Do not allow to live transactions for a long time.
 - Keep transactions short
 - Use read read_committed for long reading
- Do not use retaining (CommitRetaining, RollbackRetaining)
- Even if there AutoCommit option, check it not to use Retaining mode.

Rules for 'single transactional'

- Do not use these drivers/components 😊
- All you can do, is from time to time call
 - Database.StartTransaction;
 - Database.Commit;
- to end default transaction lifecycle

TRANSACTIONS INTERACTION TYPES

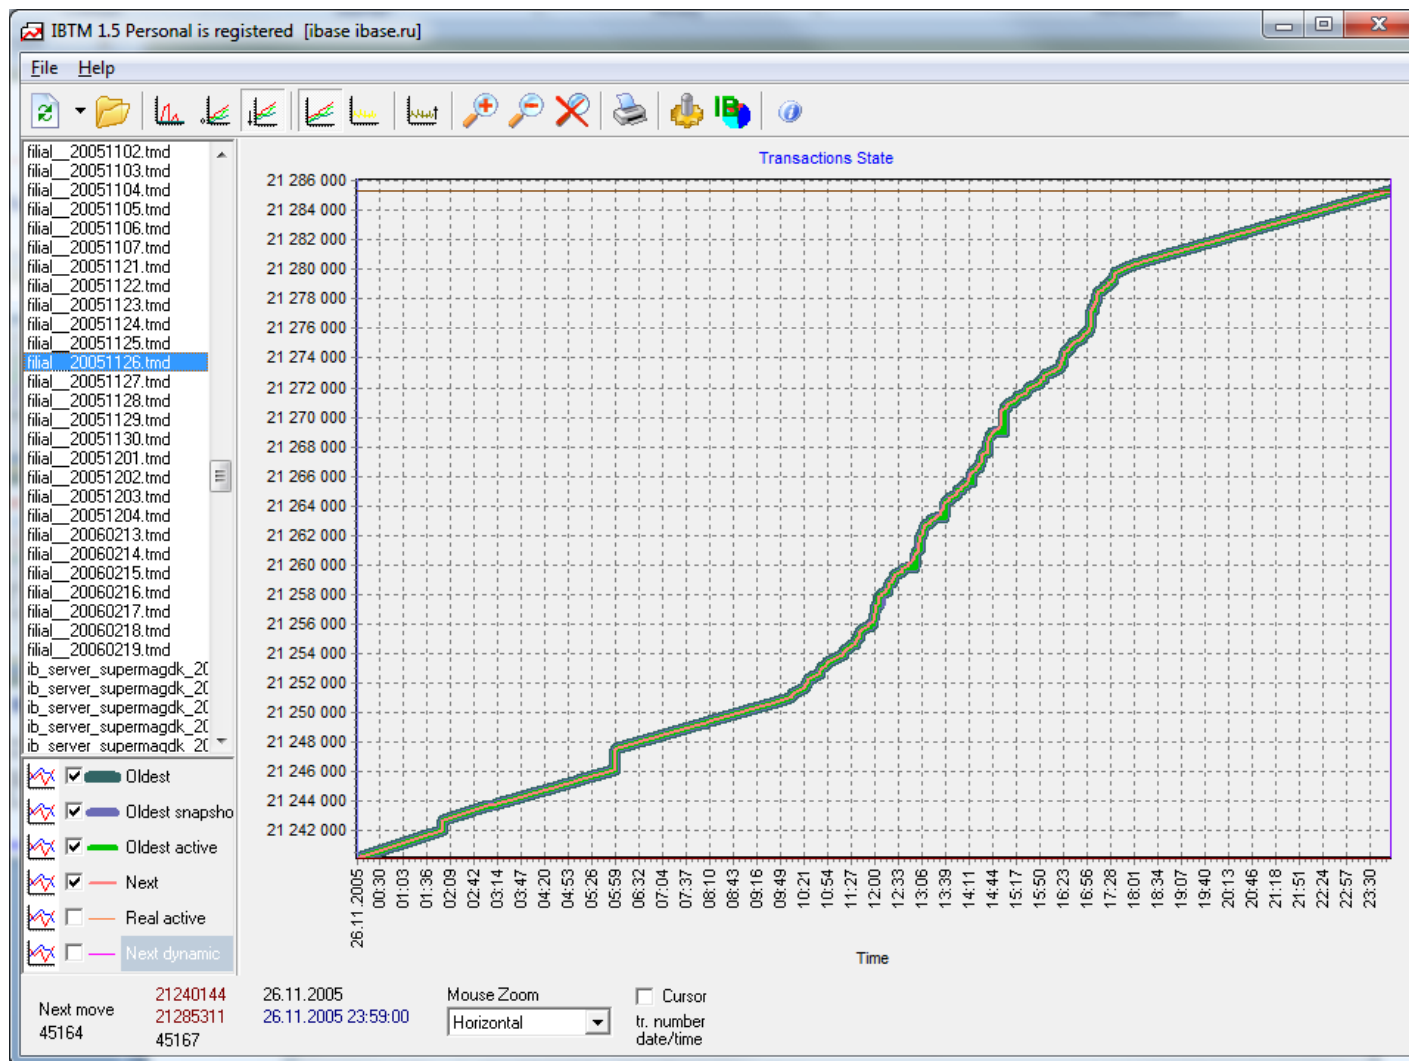
Sequential transactions



Gstat -h

- Database header page information:
 - Flags 0
 - Checksum 12345
 - Generation 112431494
 - Page size 8192
 - ODS version 11.1
 - Oldest transaction 100 x-1
 - Oldest active 101 x
 - Oldest snapshot 101 x
 - Next transaction 102 x+1
 - Bumped transaction 1
 - Sequence number 0
 - Next attachment ID 0
 - Implementation ID 16
 - Shadow count 0
 - Page buffers 256
 - Next header page 0
 - Database dialect 1
 - Creation date Jun 5, 2011 10:02:19
 - Attributes force write
- Variable header data:
 - Sweep interval: 20000
 - *END*

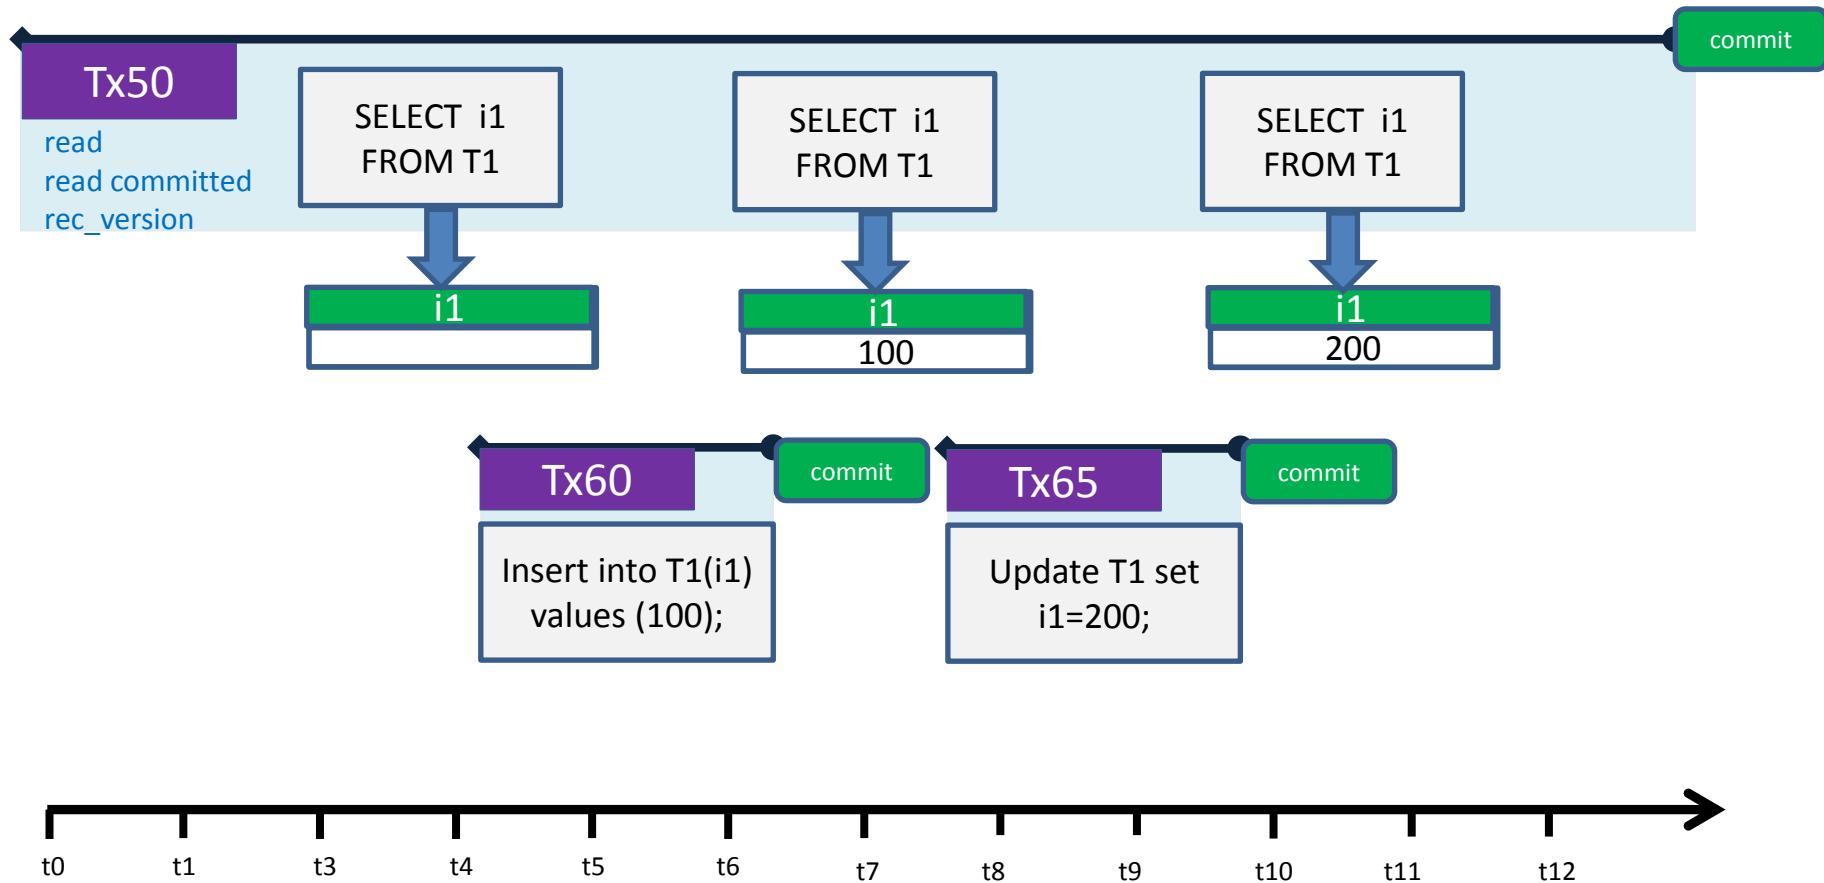
Ideal transaction control



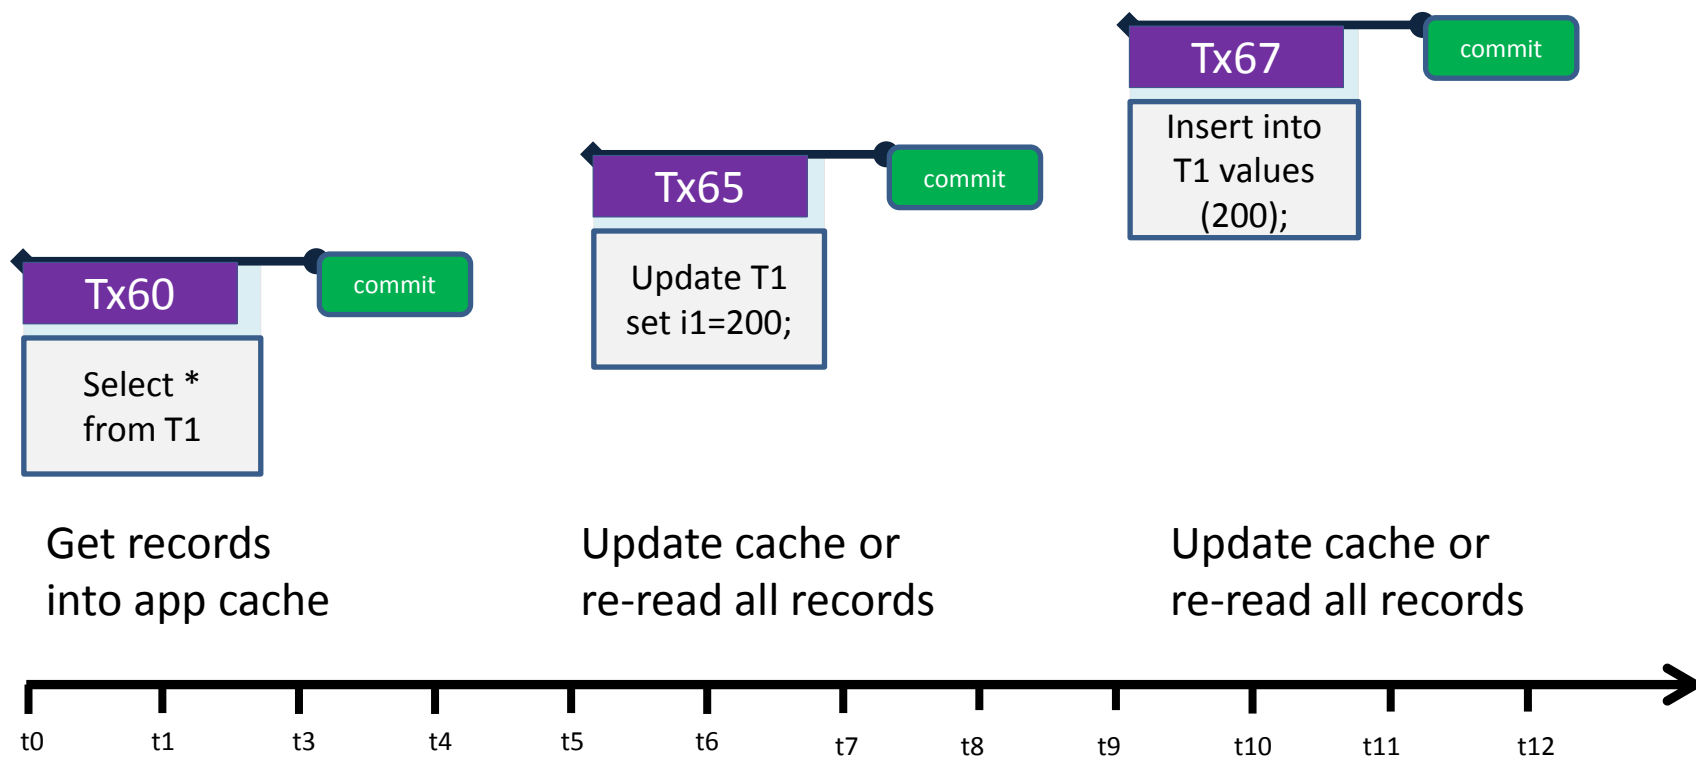
Two ways to almost ideal transaction control

1. Long read-only read-committed and short write
2. Short read and write

Long read-only RC and short write



Short read and write



Pro & Contra

Long read-only RC and short writes

- + easy to implement read and update logic
- - requires support from drivers/components (2 transactions or 2 connections)
- + more convenient for client-server
- - less convenient for multi-tier and stateless applications

Short read and writes

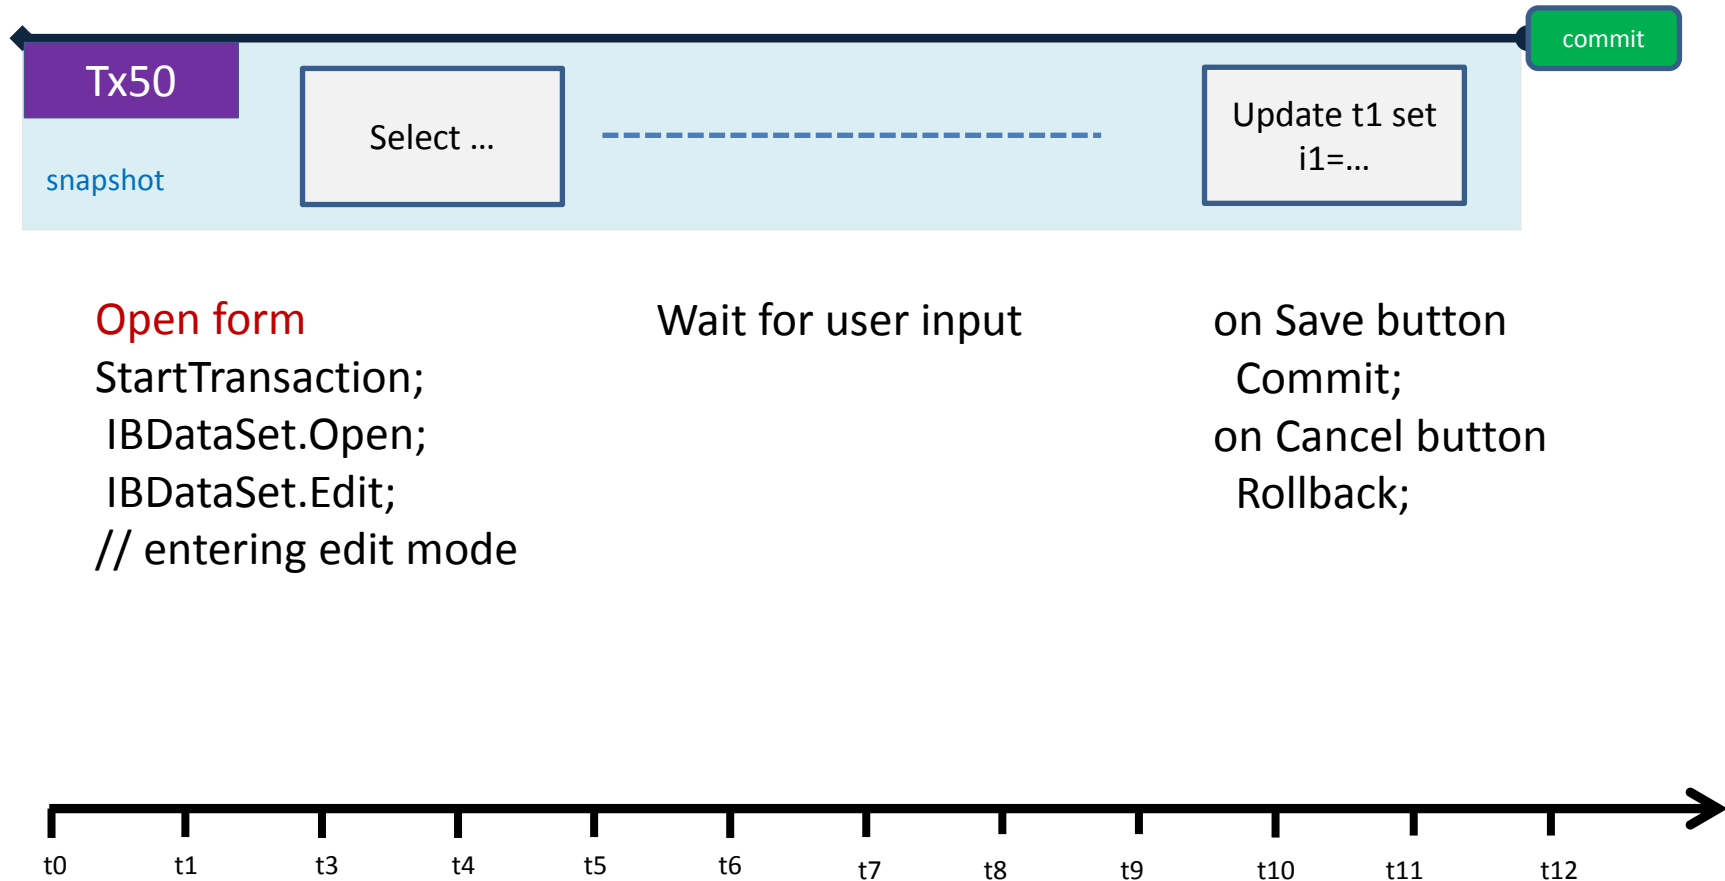
- - hard to implement sophisticated caching
- + works with any data access drivers/components
- - less convenient for client-server
- + more convenient for multi-tier and stateless applications

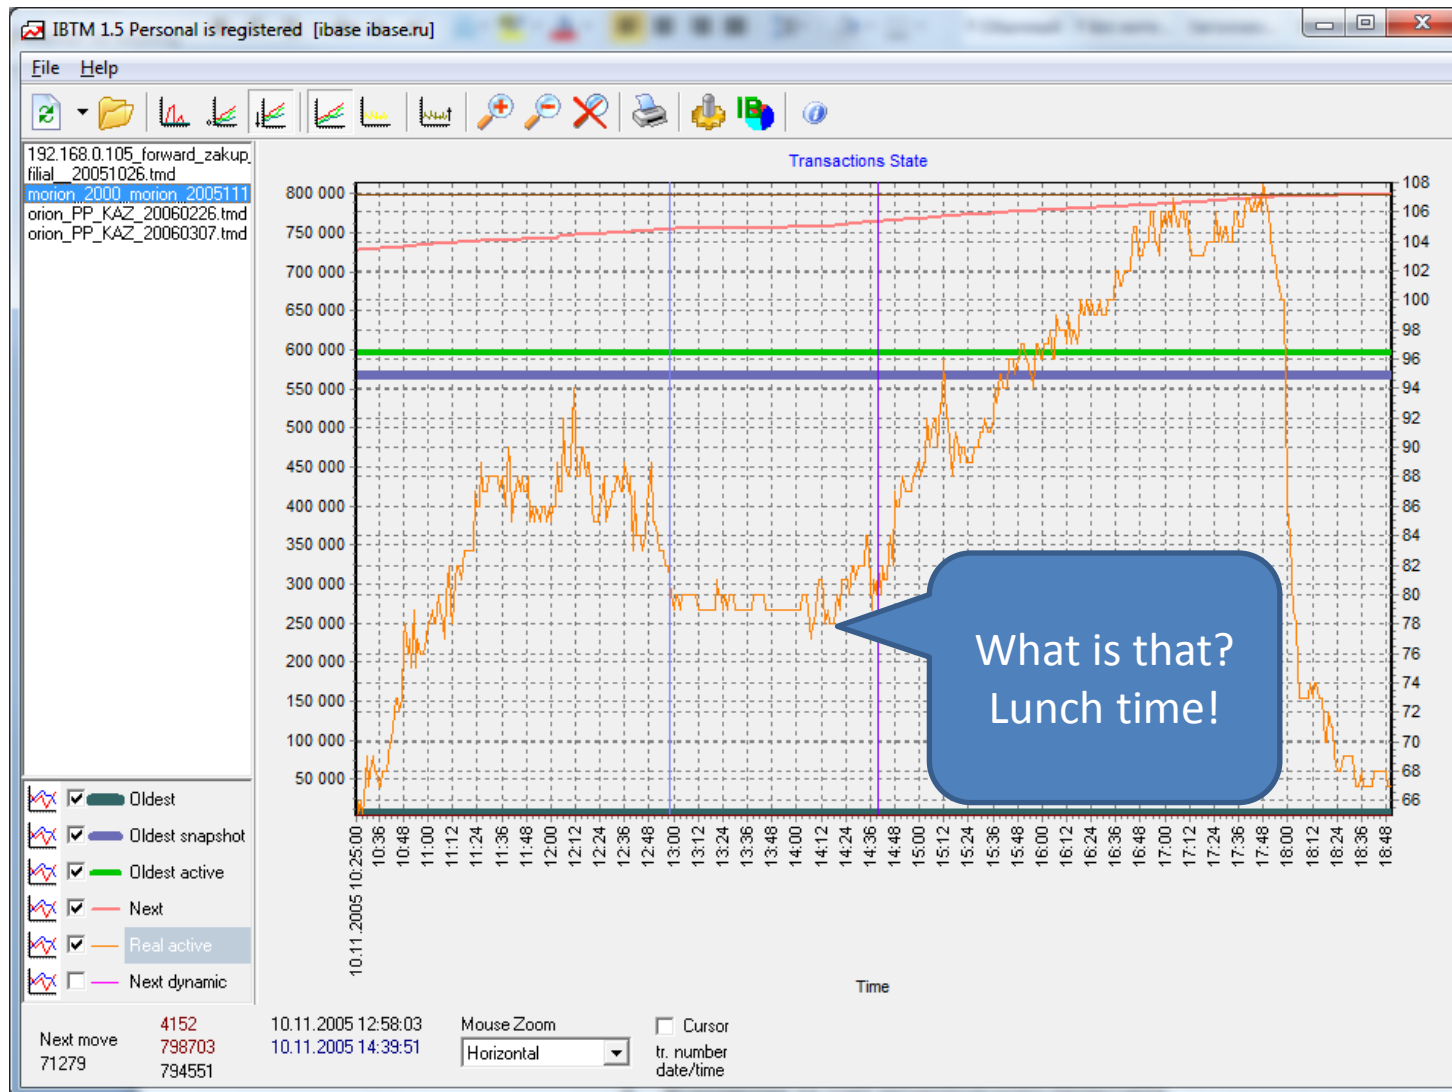
HOW TO IMPLEMENT EDIT DIALOGS IN AN EFFECTIVE AND SAFE WAY

Data editing

- Application is used by operator not in the way developer designed it
- Badly designed data editing can be a problem

Data editing: wrong scenario



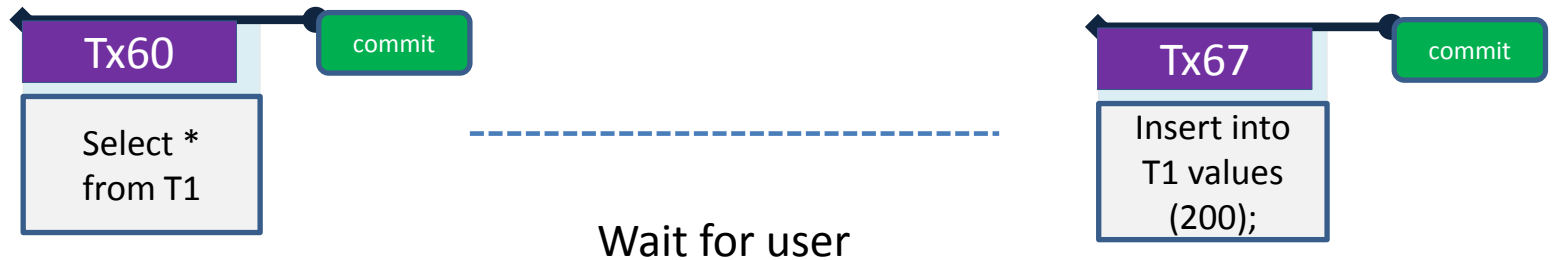


Data editing: Solution

- Open form
- StartTransaction;
- Fill controls
- Commit;
- Wait for user

- User presses Save button:
- StartTransaction;
 - IBDataSet.Edit; or IBQuery1.Prepare
- Fill data from controls
 - IBDataSet.Post; or IBQuery1.ExecSQL;
- Commit;

Data editing: Solution

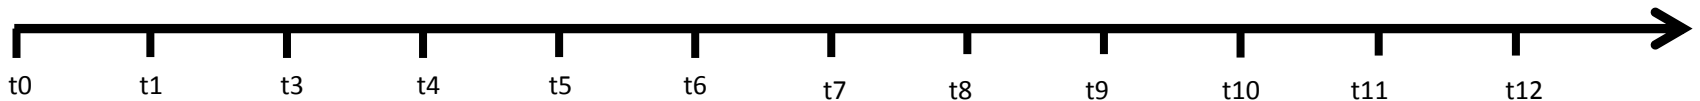


Open form

StartTransaction;
Fill controls
Commit;

on Save button:

StartTransaction;
IBDataSet.Edit; or IBQuery1.Prepare
Fill data from controls
IBDataSet.Post; or IBQuery1.ExecSQL;
Commit;



Retaining transaction context

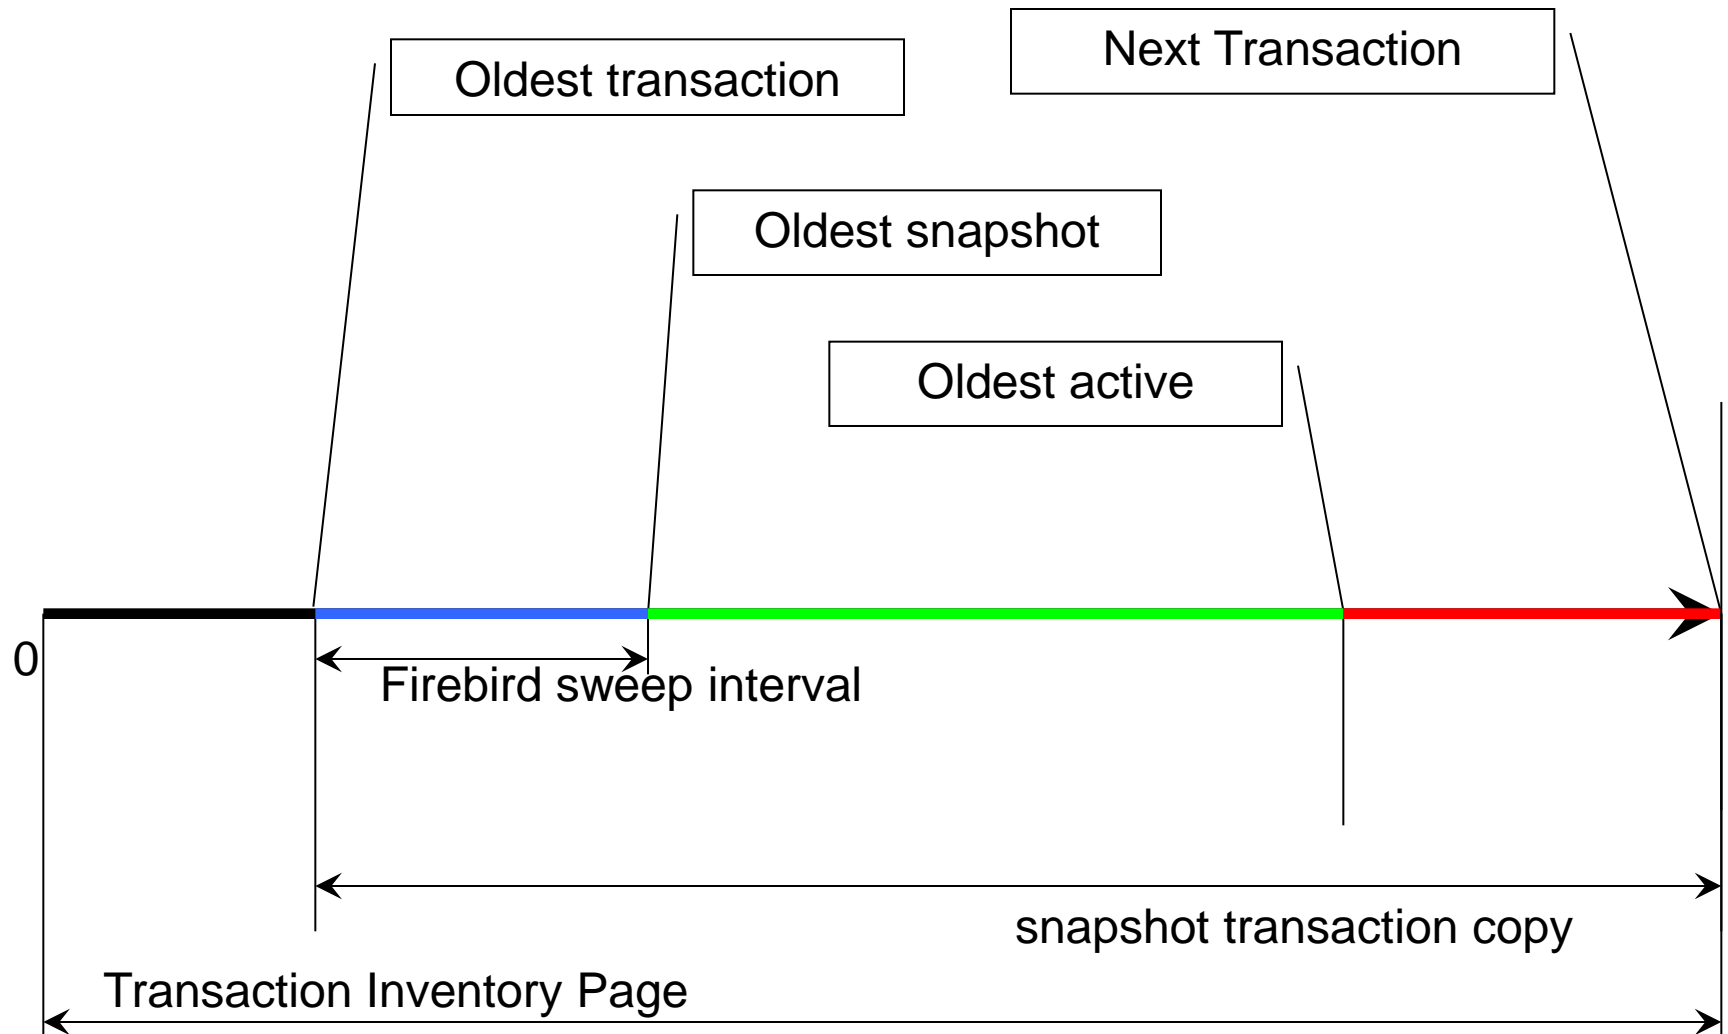
- Retaining ends transaction and starts a new one
 - Old transaction is marked in TIP as committed\rolled back
 - New transaction keeps context of old transaction
 - **Old snapshot is preserved, i.e. new transaction have the same OAT value as the old one**
 - New transaction will see changes of the old one as committed

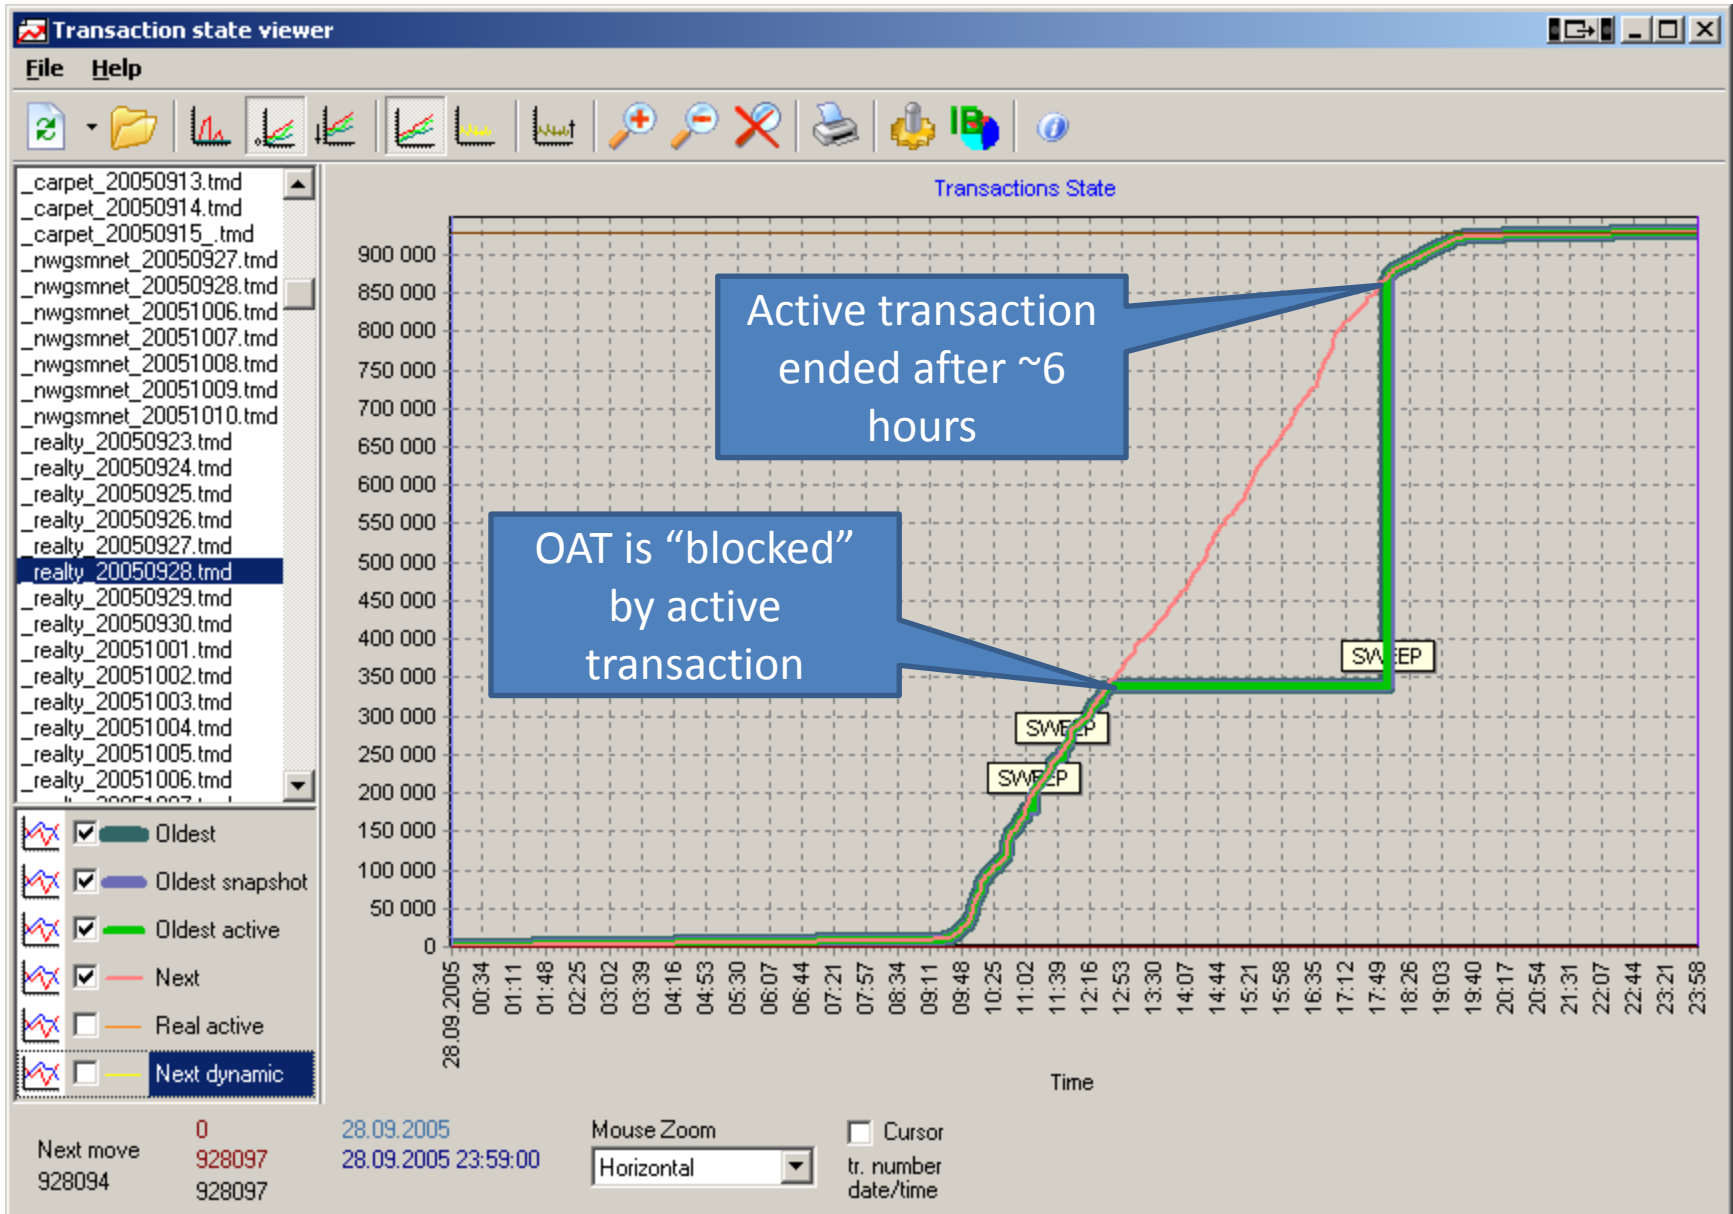
Hard commit\rollback vs retaining

- Pluses
 - One network roundtrip instead of two
 - Client recordsets survive transaction end
- Minuses
 - Open cursors are not closed
 - Temporary blobs are not released
 - Metadata locks are not released

WHEN GARBAGE COLLECTION DOES NOT WORK?

TIP markers





HOW TO IDENTIFY SWEEP

In the firebird.log !

- SRV-250 Mon May 18 21:00:01 2015
Sweep is started by SYSDBA
Database “----”
OIT 25963894, OAT 26340734, OST 26340734,
Next 27458805
- SRV-250 Mon May 18 21:46:25 2015
Sweep is finished
Database “----”
OIT 26340733, OAT 26340734, OST 26340734,
Next 27499132

What sweep could do

- Sweeping took 46 minutes (356gb database)
- OIT moved up by 376 839
- OST went up by 0
- OAT went up by 0
- Next went up by 40 327
 - 77k transactions per hour
- Next-OAT = 1 158 398
 - div 77k = oat stuck ~14 hours ago

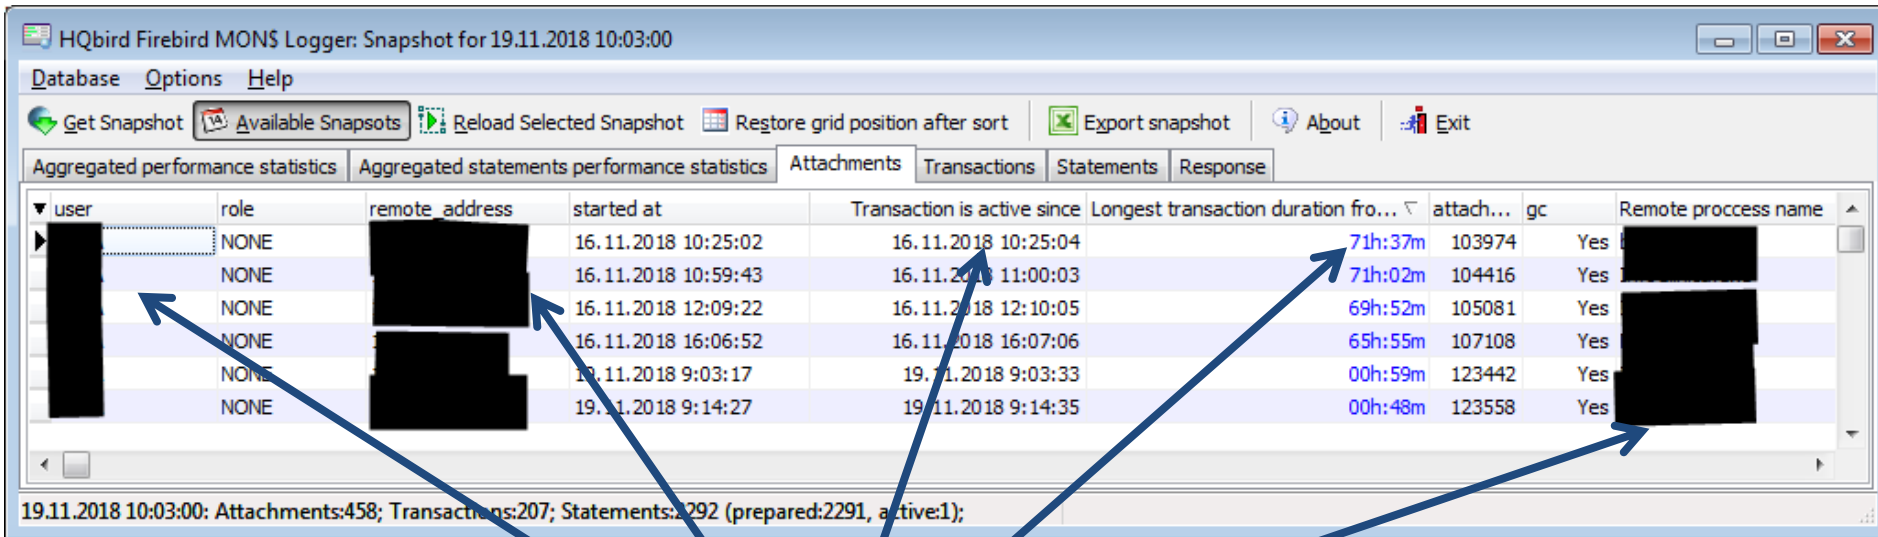
HOW TO IDENTIFY PROBLEMATIC TRANSACTIONS (TOO LONG, WRONG ISOLATION LEVEL) WITH MON\$

mon\$transactions

- MON\$TRANSACTION_ID - transaction ID
- MON\$ATTACHMENT_ID - attachment ID
- MON\$STATE - transaction state
 - 0: idle
 - 1: active
- MON\$TIMESTAMP - transaction start date/time
- MON\$TOP_TRANSACTION top transaction MON\$OLDEST_TRANSACTION - local OIT number
- MON\$OLDEST_ACTIVE - local OAT number
- MON\$ISOLATION_MODE - isolation mode
 - 0: consistency
 - 1: concurrency
 - 2: read committed record version
 - 3: read committed no record version
- MON\$LOCK_TIMEOUT - lock timeout
 - 0: no wait
 - 1: infinite wait
 - N: timeout N
- MON\$READ_ONLY - read-only flag 0/1
- MON\$AUTO_COMMIT - auto-commit flag
- MON\$AUTO_UNDO - auto-undo flag
- MON\$STAT_ID - statistics ID

- `select * from mon$transactions
order by mon$timestamp desc`
- `select a.*, t.*
from mon$attachments a, mon$transactions t
where a.mon$attachment_id =
t.mon$attachment_id
order by t.mon$timestamp desc`

MonLogger



HQbird Firebird MONS Logger: Snapshot for 19.11.2018 10:03:00

Database Options Help

Get Snapshot Available Snapshots Reload Selected Snapshot Restore grid position after sort Export snapshot About Exit

Aggregated performance statistics Aggregated statements performance statistics Attachments Transactions Statements Response

user	role	remote address	started at	Transaction is active since	Longest transaction duration from...	attach...	gc	Remote process name
██████████	NONE	██████████	16.11.2018 10:25:02	16.11.2018 10:25:04	71h:37m	103974	Yes	██████████
██████████	NONE	██████████	16.11.2018 10:59:43	16.11.2018 11:00:03	71h:02m	104416	Yes	██████████
██████████	NONE	██████████	16.11.2018 12:09:22	16.11.2018 12:10:05	69h:52m	105081	Yes	██████████
██████████	NONE	██████████	16.11.2018 16:06:52	16.11.2018 16:07:06	65h:55m	107108	Yes	██████████
██████████	NONE	██████████	19.11.2018 9:03:17	19.11.2018 9:03:33	00h:59m	123442	Yes	██████████
██████████	NONE	██████████	19.11.2018 9:14:27	19.11.2018 9:14:35	00h:48m	123558	Yes	██████████

19.11.2018 10:03:00: Attachments:458; Transactions:207; Statements:292 (prepared:2291, active:1);

who, where, when, how long, what application

read/write, RC & snapshot, wait

HQbird Firebird MONS Logger: Snapshot for 22.10.2013 10:37:12

Database Options Help

Get Snapshot Available Snapshots Reload Selected Snapshot Restore grid position after sort Export snapshot About Exit

Aggregated performance statistics Aggregated statements performance statistics Attachments Transactions Statements Response

Link to selected attachment: MON\$USER=SYSDBA; MON\$ROLE=NONE; MON\$REMOTE_ADDRESS=192.168.1.52; MON\$TIMESTAMP=21.10.2013 11:43:39; Ok Remove readonly and readcommitted

started at	transacti...	attachm...	state	isolation_mode	lock_timeout	read_only	auto_co...	auto_undo	record_seq_reads	record_idx_reads
22.10.2013 10:34:04	457093	205815	active	read committed record version	no wait	read write	No	Yes	7	2022
22.10.2013 10:34:11	457094	204831	active	read committed record version	no wait	read write	No	Yes	3	30947
22.10.2013 10:34:32	457142	205845	active	read committed record version	no wait	read write	No	Yes	41280	239894
22.10.2013 10:34:38	457135	205847	active	concurrency	infinite wait	read write	No	Yes	10	1829
22.10.2013 10:35:08	457186	202910	active	read committed record version	no wait	read write	No	Yes	5	29
22.10.2013 10:35:11	457188	202910	idle	read committed record version	no wait	read only	No	Yes	0	0
22.10.2013 10:35:11	457190	205878	active	concurrency	infinite wait	read write	No	Yes	2142	2503
22.10.2013 10:35:27	457208	205260	idle	read committed record version	no wait	read only	No	Yes	0	76
22.10.2013 10:35:34	457316	204320	active	read committed record version	no wait	read write	No	Yes	866893	1429907
22.10.2013 10:35:38	457334	204772	active	read committed record version	no wait	read write	No	Yes	58390	175359
22.10.2013 10:35:40	457224	203673	active	read committed record version	no wait	read write	No	Yes	608825	44091
22.10.2013 10:35:41	457225	204772	idle	read committed record version	no wait	read only	No	Yes	0	68
22.10.2013 10:35:45	457227	166995	active	read committed record version	no wait	read write	No	Yes	64484	296581
22.10.2013 10:35:54	457236	202949	active	read committed record version	no wait	read write	No	Yes	57686	415080
22.10.2013 10:36:00	457247	203414	active	read committed record version	no wait	read write	No	Yes	44685	334843
22.10.2013 10:36:09	457259	203848	active	read committed record version	no wait	read write	No	Yes	7635	51256
22.10.2013 10:36:13	457272	205928	active	read committed record version	no wait	read only	No	Yes	987	13923
22.10.2013 10:36:14	457290	205928	active	read committed record version	no wait	read write	No	Yes	30897	329501
22.10.2013 10:36:15	457277	205929	active	read committed record version	no wait	read only	No	Yes	987	13883
22.10.2013 10:36:15	457331	205929	active	read committed record version	no wait	read write	No	Yes	31071	174112

22.10.2013 10:37:12: Attachments:78; Transactions:61; Statements:3545 (prepared:3540, active:5);

How to track conflicts and deadlocks

- Set trace config
- Run trace session
- Analyze trace log
- <https://ib-aid.com/en/how-to-track-deadlocks-in-firebird/>

LEGACY APPLICATIONS: WORKAROUNDS FOR ERROR IN TRANSACTIONS MANAGEMENT

Oops...

- If you do not have sources, or you do not understand sources, or
The only way is to **terminate** these applications by at/cron schedule, for example, each hour or two.
- Active transactions in these applications decrease performance by accumulating record versions, blocking sweep, etc.
- ! rewrite these applications

BEST PRACTICES

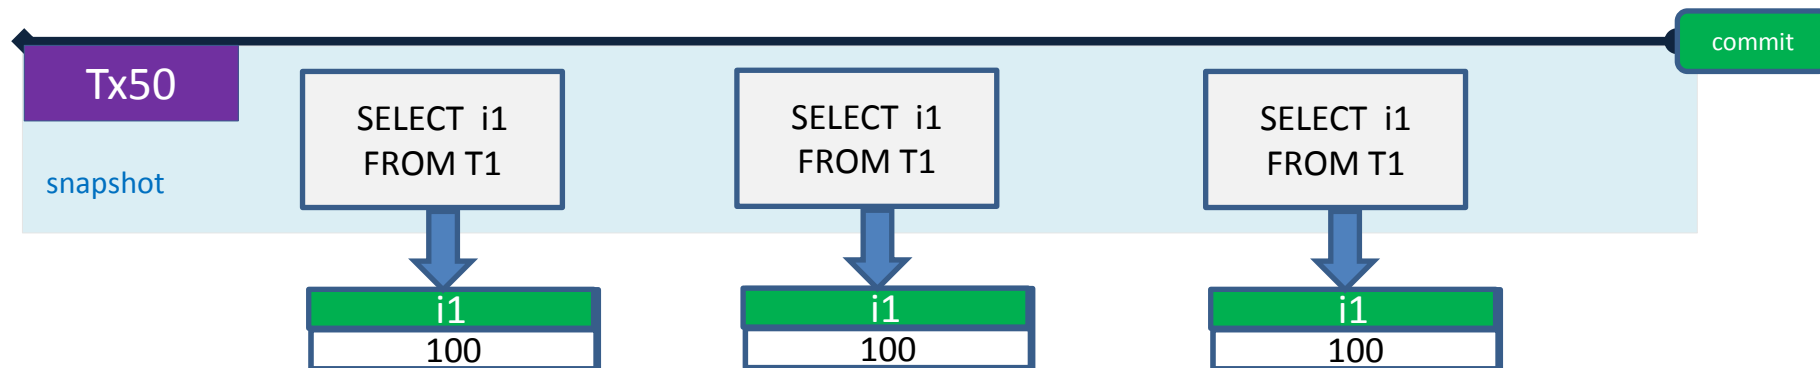
Exceptions from ideal transaction control

- Reports
- Goods balance
- Explicit record locking
- Robots

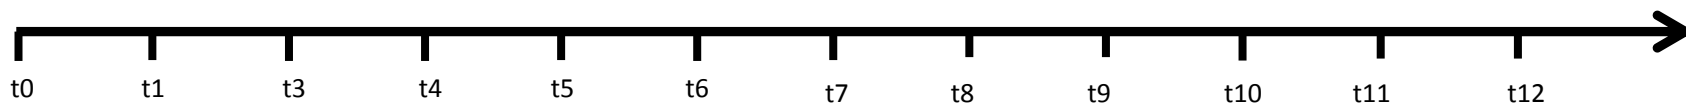
Reports

- Need data consistency
- Long queries
- Complex reports read the same data several times

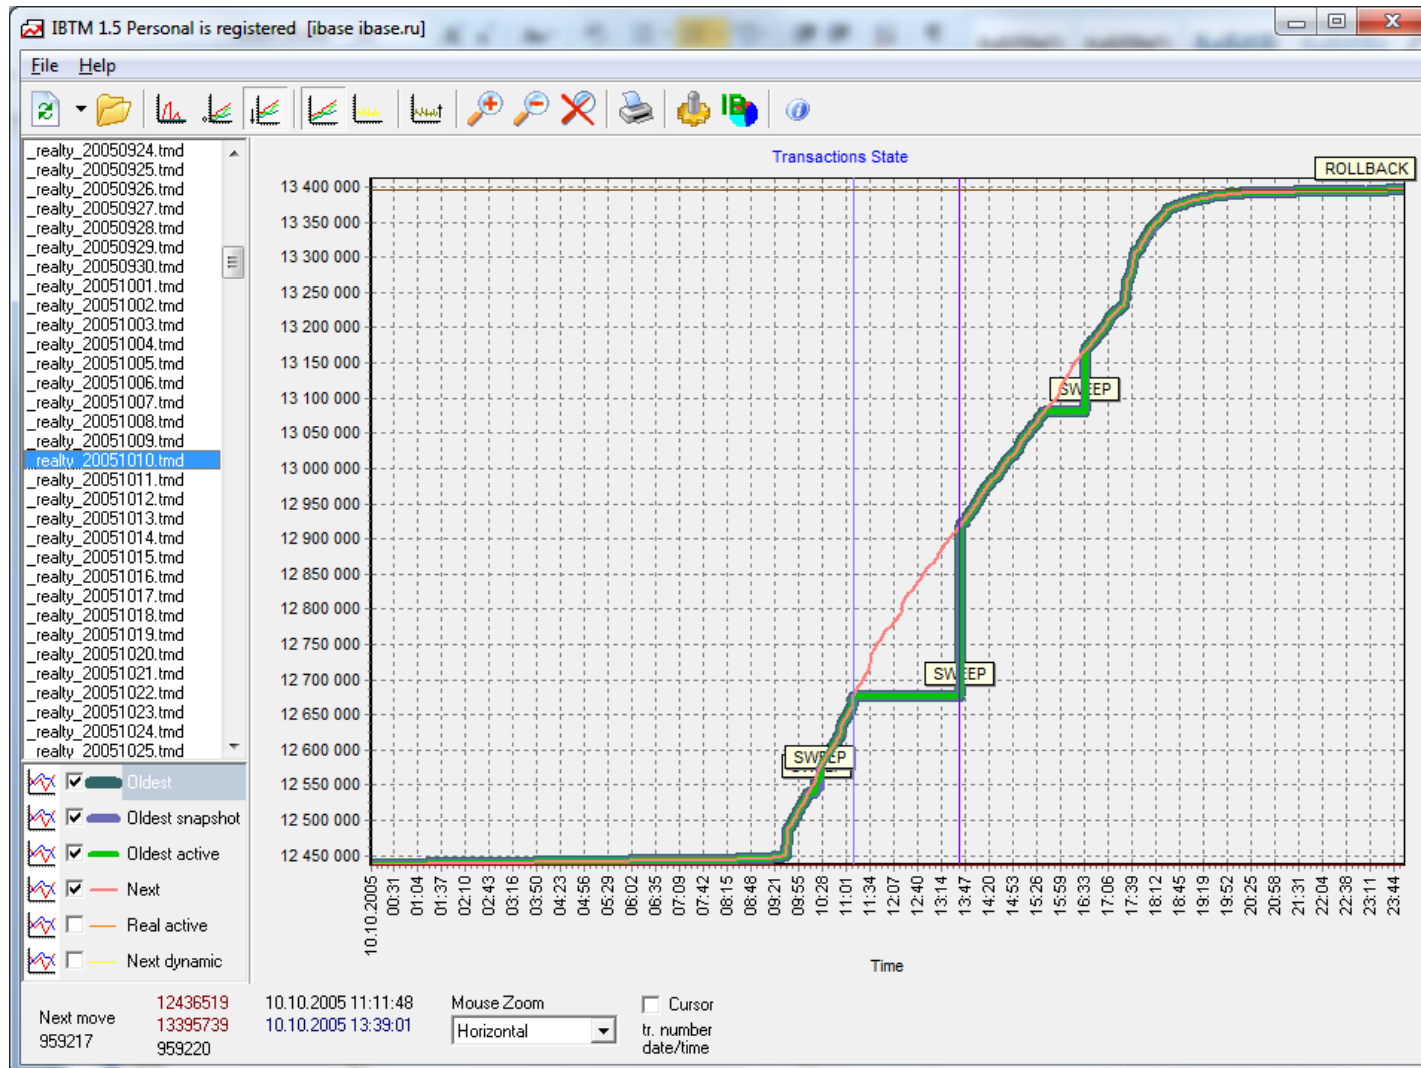
Reports - snapshot



No difference – wait/nowait (except concurrency), read/write



Heavy report example – OAT stuck



How to workaround long report problem

- Most reports does not need real-time data
- Change logic of data processing
- Scaling
 - Replication
 - Transferring data to another DB with Execute Statement On External
 - Nbackup

Change logic of data processing – Stored aggregates

- 1 order - ~10 goods
- 100 orders per day
- $100 * 10 = 1000$ records per day
- 365000 records per year

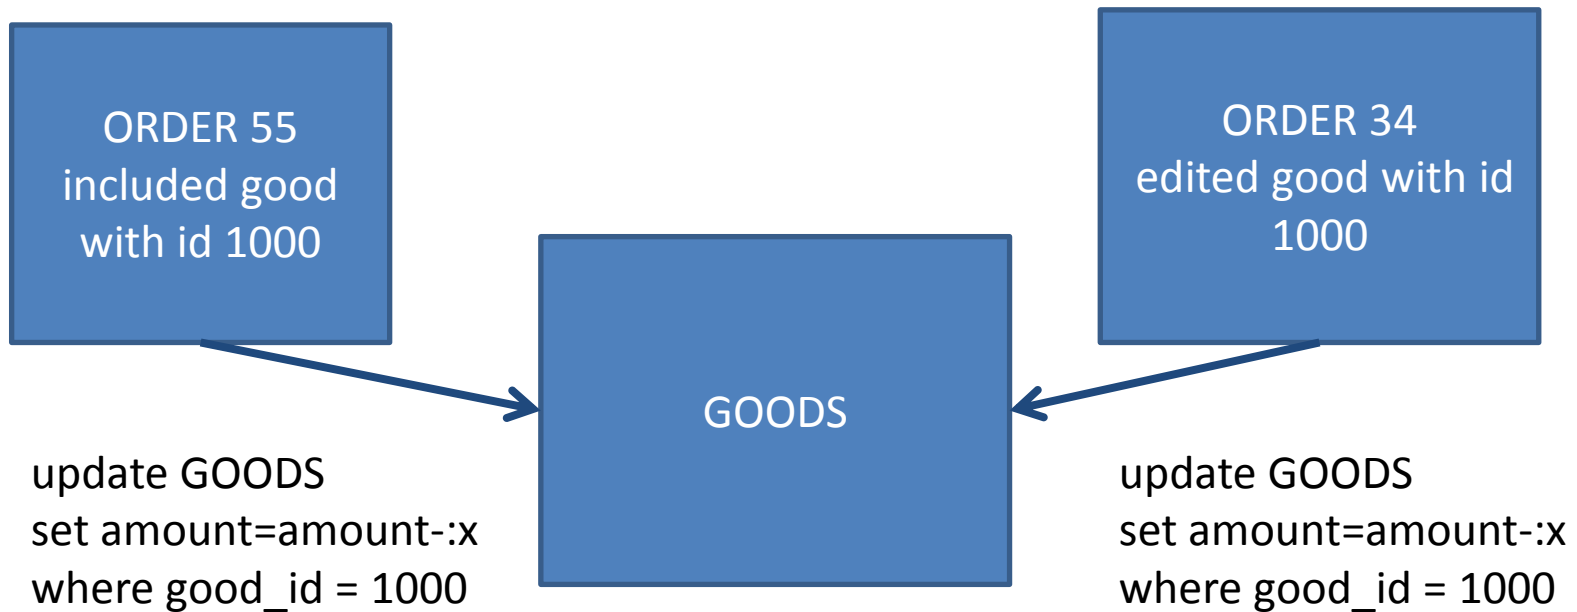
- Store “order_total” in ORDERS table – 10 times less records
- **Pro**: less records, faster queries
 - No update conflicts if there is no concurrent order editing
- **Con**: additional field in ORDERS

If you want to go further...

- To store sum by day, month, ...
- Updates by triggers “in place” won’t work – too high possibility of lock conflicts
- Solution? Routine updates
 - Routine procedure must be run in exclusive mode
 - Using generator
 - Using consistency isolation mode
 - By schedule (at night)

Goods balances – update locks

- Change goods AMOUNT while order is processed
-
- Insert - set $AMOUNT = AMOUNT - new.INORDER$
- Delete - set $AMOUNT = AMOUNT + new.INORDER$
- Update - set $AMOUNT = AMOUNT + new.INORDER - old.INORDER$
- There may be conflicts when 2 people sell same `good_id`
 - Long transaction will lock all concurrent order processing
 - Short transactions have less chances to get update conflict, and may be retried



Editing one order by 2 users is a rare case, but using same item is not rare

Goods balances- solution

- CREATE TABLE MOVEMENTS(
GOOD INTEGER NOT NULL REFERENCES GOODS, AMOUNT
INTEGER NOT NULL)
- CREATE TABLE GOODS_AMOUNTS_AGG(
GOOD INTEGER NOT NULL REFERENCES GOODS, AMOUNT
INTEGER NOT NULL)
- On insert update and delete MOVEMENTS do
- **INSERT** INTO GOODS_AMOUNT_AGG
(GOOD, AMOUNT) VALUES
 - (NEW.GOOD, NEW.AMOUNT);
 - (NEW.GOOD, NEW.AMOUNT-OLD.AMOUNT);
 - (OLD.GOOD, -OLD.AMOUNT);

- CREATE VIEW GOODS_AMOUNT
(GOOD, AMOUNT) AS
 SELECT GOOD, SUM(AMOUNT)
 FROM GOODS_AMOUNT_AGG
 GROUP BY GOOD
- CREATE PROCEDURE GOODS_AMOUNT_ROLL_UP AS
 DECLARE GOOD INTEGER;
 DECLARE TOTAL INTEGER;
 BEGIN
 FOR SELECT GOOD, SUM(AMOUNT)
 FROM GOODS_AMOUNT_AGG
 GROUP BY GOOD
 HAVING COUNT(*)>1 – *interested of 2 or more records*
 INTO :GOOD, :TOTAL
 DO
 BEGIN
 DELETE FROM GOODS_AMOUNT_AGG
 WHERE GOOD=:GOOD;
 INSERT INTO GOODS_AMOUNT_AGG
 (GOOD, AMOUNT) VALUES(:GOOD, :TOTAL);
 END
 END
- Run procedure in concurrency (or consistency)

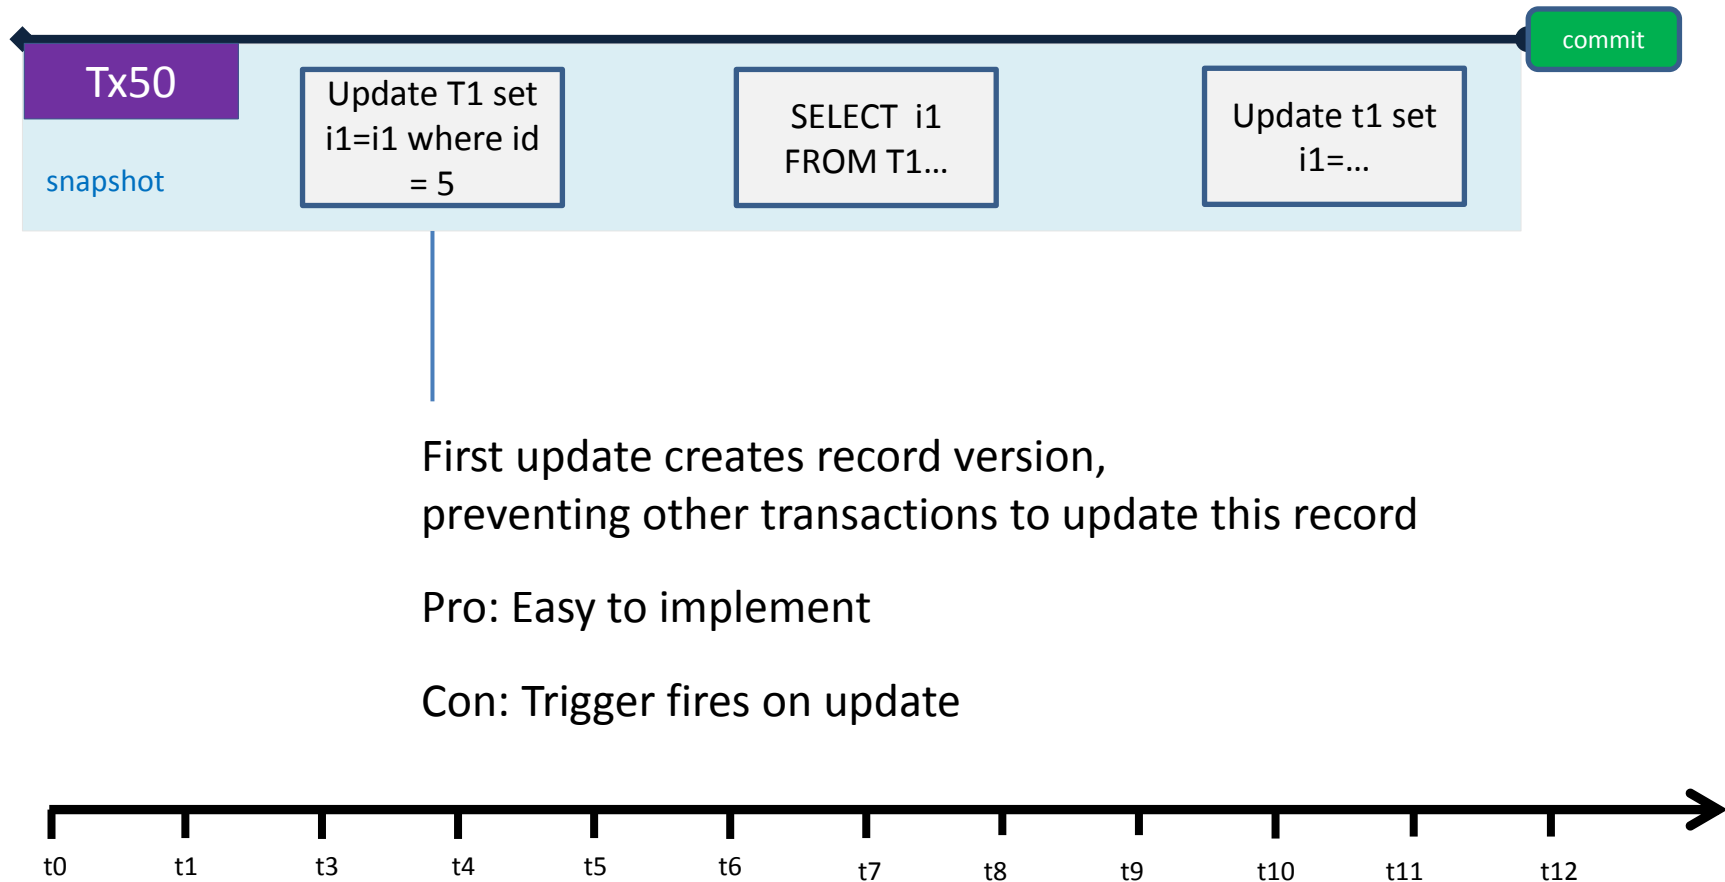
Exclusive document editing

- Goal – implement exclusive changes
- Rollbacks are not welcome
- Need explicit record locking

How to implement explicit record locking

- Blank update in long transaction
 - Or `SELECT ... FOR UPDATE WITH LOCK`
- Flags at business logic level

Blank update



SELECT ... FOR UPDATE WITH LOCK

- Same as blank update
- Can lock several records
- Locks record on fetch
 - Result returns one record per one fetch (no buffering)
- Useless for aggregates (SUM, AVG, COUNT, ...)

- Locking in the versioning server is not normal
- It maybe not enough to choose appropriate transaction isolation level

Flags at business logic level

- Add User and TimeStamp fields, or create additional table
- When you want to “lock”, write USER and CURRENT_TIMESTAMP in short transaction
- if user <> myself then
 - if TimeStamp is far then
 - UPDATE set User, TimeStamp
 - else Fail(“locked by user User at TimeStamp”)
- else
 - UPDATE set TimeStamp
- Additional table need to be cleared (disconnected apps)

Robot rules

Reading robots

- Use read-only
ReadCommitted
- Try to do work in one transaction, if possible
- Multi-tier - connection and transaction pooling
- Goals
 - Do not stuck OAT
 - Do not advance Next too much

Writing robots

- Do not keep attachment open
 - attach, do work, close;
- Keep transactions short
- Try to do work in one transaction, if possible
- Goals
 - Do not stuck OAT

- Thank you!
- www.firebirdsql.org
- www.ib-aid.com
- support@ib-aid.com