

EPFL

Transactions

Prof. George Candea

School of Computer & Communication Sciences

What are transactions ?

What is a transaction in the real world ?

- Two or more parties...
 - *negotiate for a while*
 - *then make a deal*
 - *write it up in a contract*
 - *all parties sign the contract*
=> transaction completes
- Implication
 - *everyone agrees*
 - *deal is binding*

Properties of real-world transactions

- Transaction is in accordance with legal protocols
 - *i.e., law governs society*
- The entire deal either takes place or not
 - *either all parties are bound by it or none are*
- Once the contract is signed, it cannot be abrogated
 - *can be amended / compensated*
- If someone engages in a different transaction doesn't affect this one



Tandem TR 81.3

**The Transaction Concept:
Virtues and Limitations**

Jim Gray
Tandem Computers Incorporated
19333 Vallco Parkway, Cupertino CA 95014

June 1981

ABSTRACT: A transaction is a transformation of state which has the properties of atomicity (all or nothing), durability (effects survive failures) and consistency (a correct transformation). The transaction concept is key to the structuring of data management applications. The concept may have applicability to programming systems in general. This paper restates the transaction concepts and attempts to put several implementation approaches in perspective. It then describes some areas which require further study: (1) the integration of the transaction concept with the notion of abstract data type, (2) some techniques to allow transactions to be composed of sub-transactions, and (3) handling transactions which last for extremely long times (days or months).

Appeared in Proceedings of Seventh International Conference on Very Large Databases, Sept. 1981. Published by Tandem Computers Incorporated.



What is a transaction in the computing world ?

- Transaction = collection of *actions* that comprise a consistent transformation of system state
 - *Actions read and transform values*
- Outcome = committed | aborted
- The only way to "correct" a committed transaction is via another (compensating) transaction
- System state may include assertions of what consistency means

What is an action in a transaction ?

- Unprotected
 - *need not be undone if txn must be aborted*
 - *need not be redone if the value needs to be reconstructed*
- Protected
 - *action can and must be undone / redone if ...*
- Real
 - *cannot be undone (once done)*
- Txn commits => all protected and real actions persist
- Txn aborts => no effects of protected and real actions are visible to other txns

How does a transaction look ?

```
DELETE FROM Orders WHERE ClientID = @DonaldTrump  
DELETE FROM Clients WHERE ClientID = @DonaldTrump
```


How does a transaction look ?

```
DELETE FROM Orders WHERE ClientID = @DonaldTrump  
DELETE FROM Clients WHERE ClientID = @DonaldTrump
```

```
BEGIN TRANSACTION  
DELETE FROM Orders WHERE ClientID = @DonaldTrump  
DELETE FROM Clients WHERE ClientID = @DonaldTrump  
COMMIT
```

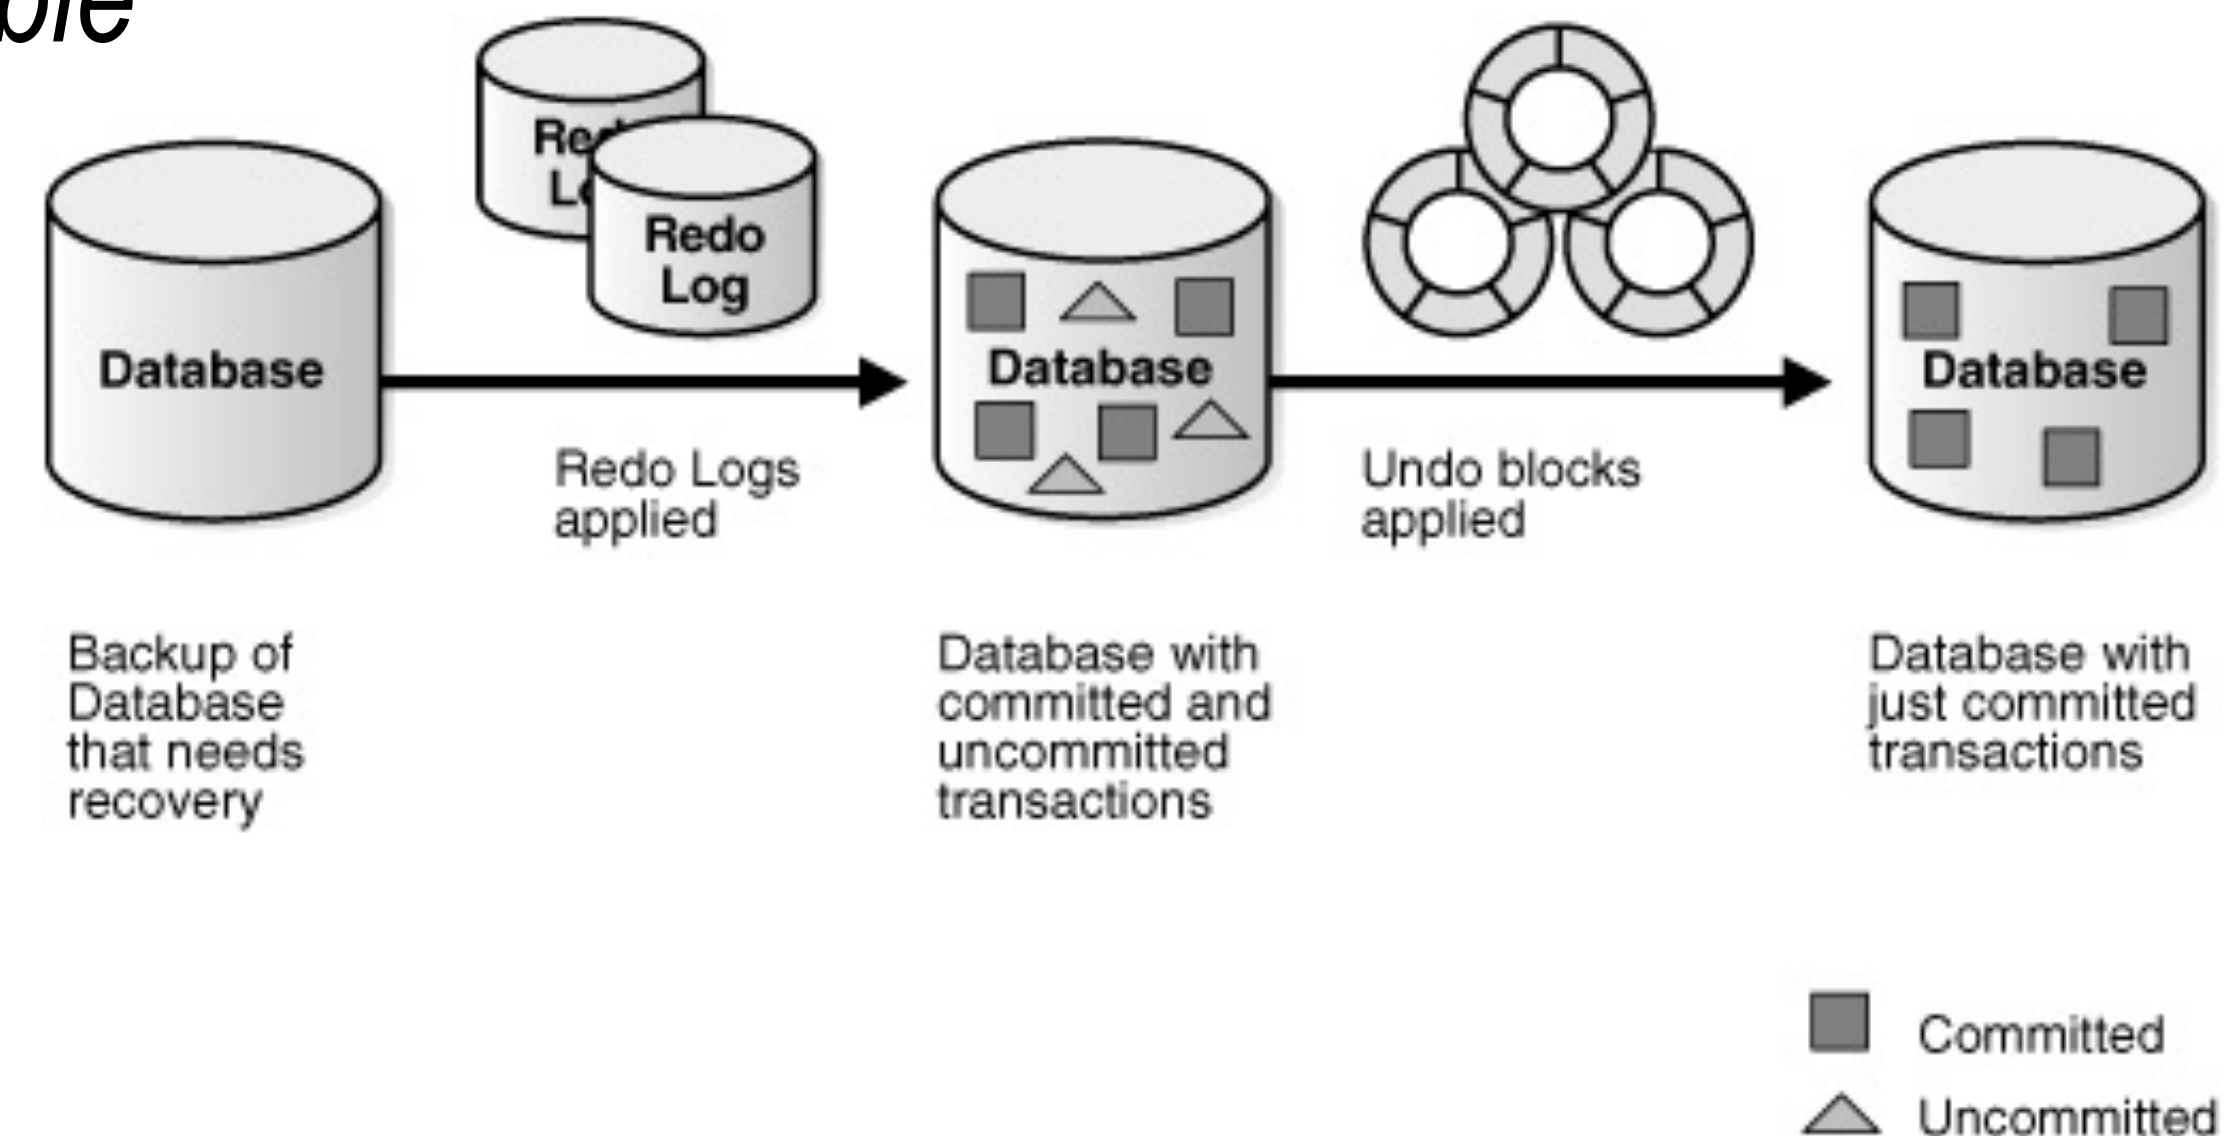
```
BEGIN TRANSACTION  
DELETE FROM Orders WHERE ClientID = @DonaldTrump  
DELETE FROM Clients WHERE ClientID = @DonaldTrump  
IF @@ROWCOUNT > 1  
    ROLLBACK  
COMMIT
```

What is ACID

A = Atomicity

“All or nothing”

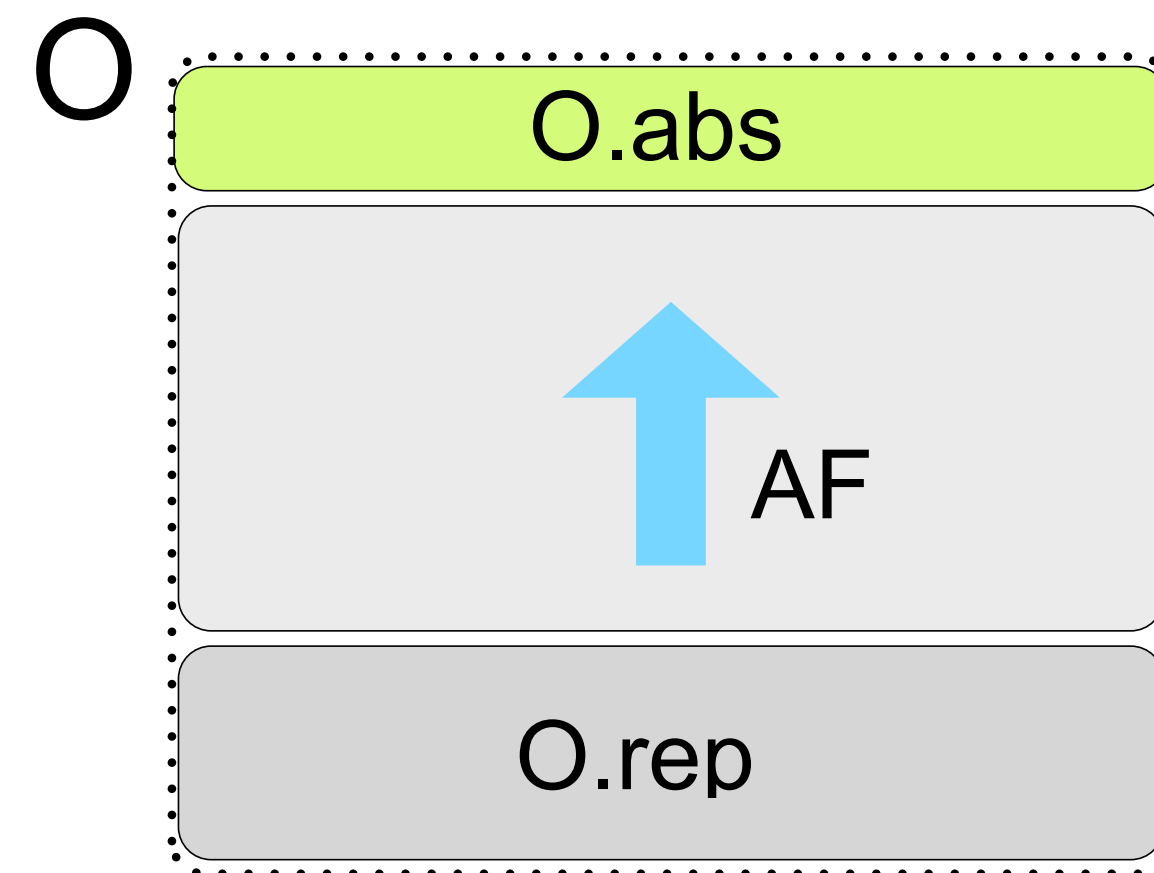
- Either all protected and real actions are visible or none
- Key = how txn looks “from the outside”
 - *expressed in terms of abstract state*
 - *partial results ok, as long as not visible*



C = Consistency

“Obey legal protocols”

- Txn transitions system from one valid state to another
- *intermediate states are not visible*



O exposes correct abstraction

↑

$O.abs = AF(O.rep)$

↑

$RI(O.rep)=true$

$$RI(O.rep) = \bigwedge_{r_i \in O.rep} RI(r_i) + \text{composition of } r_i$$

Integrity Constraints

```
CREATE TABLE Clients(  
  Id int NOT NULL PRIMARY KEY,  
  ...  
)
```

```
CREATE TABLE Orders (  
  OrderId int NOT NULL PRIMARY KEY,  
  ...  
  ClientId int FOREIGN KEY REFERENCES Clients(Id)  
)
```

```
CREATE TABLE Orders (  
  OrderId int NOT NULL PRIMARY KEY,  
  ...  
  ClientId int FOREIGN KEY REFERENCES Clients(Id) ON DELETE CASCADE  
)
```

```

public class Foo {
  public static int BinarySearch(int[] a, int key)
    requires forall{ int i in (0: a.Length), int j in (i: a.Length); a[i] <= a[j] };
    ensures 0 <= result ==> a[result] == key;
    ensures result < 0 ==> forall{int i in (0: a.Length); a[i] != key};
  {
    int low = 0;
    int high = a.Length;

    while (low < high)
      invariant 0 <= low && high <= a.Length;
      invariant forall{ int i in (0: low); a[i] != key };
      invariant forall{ int i in (high: a.Length); a[i] != key };
      {
        int mid = low + (high - low) / 2;
        int midVal = a[mid];

        if (key < midVal) {
          low = mid + 1;
        } else if (midVal < key) {
          high = mid;
        } else {
          return mid; // key found
        }
      }
    }
    return -low - 1; // key not found
  }
}

```

C = Consistency

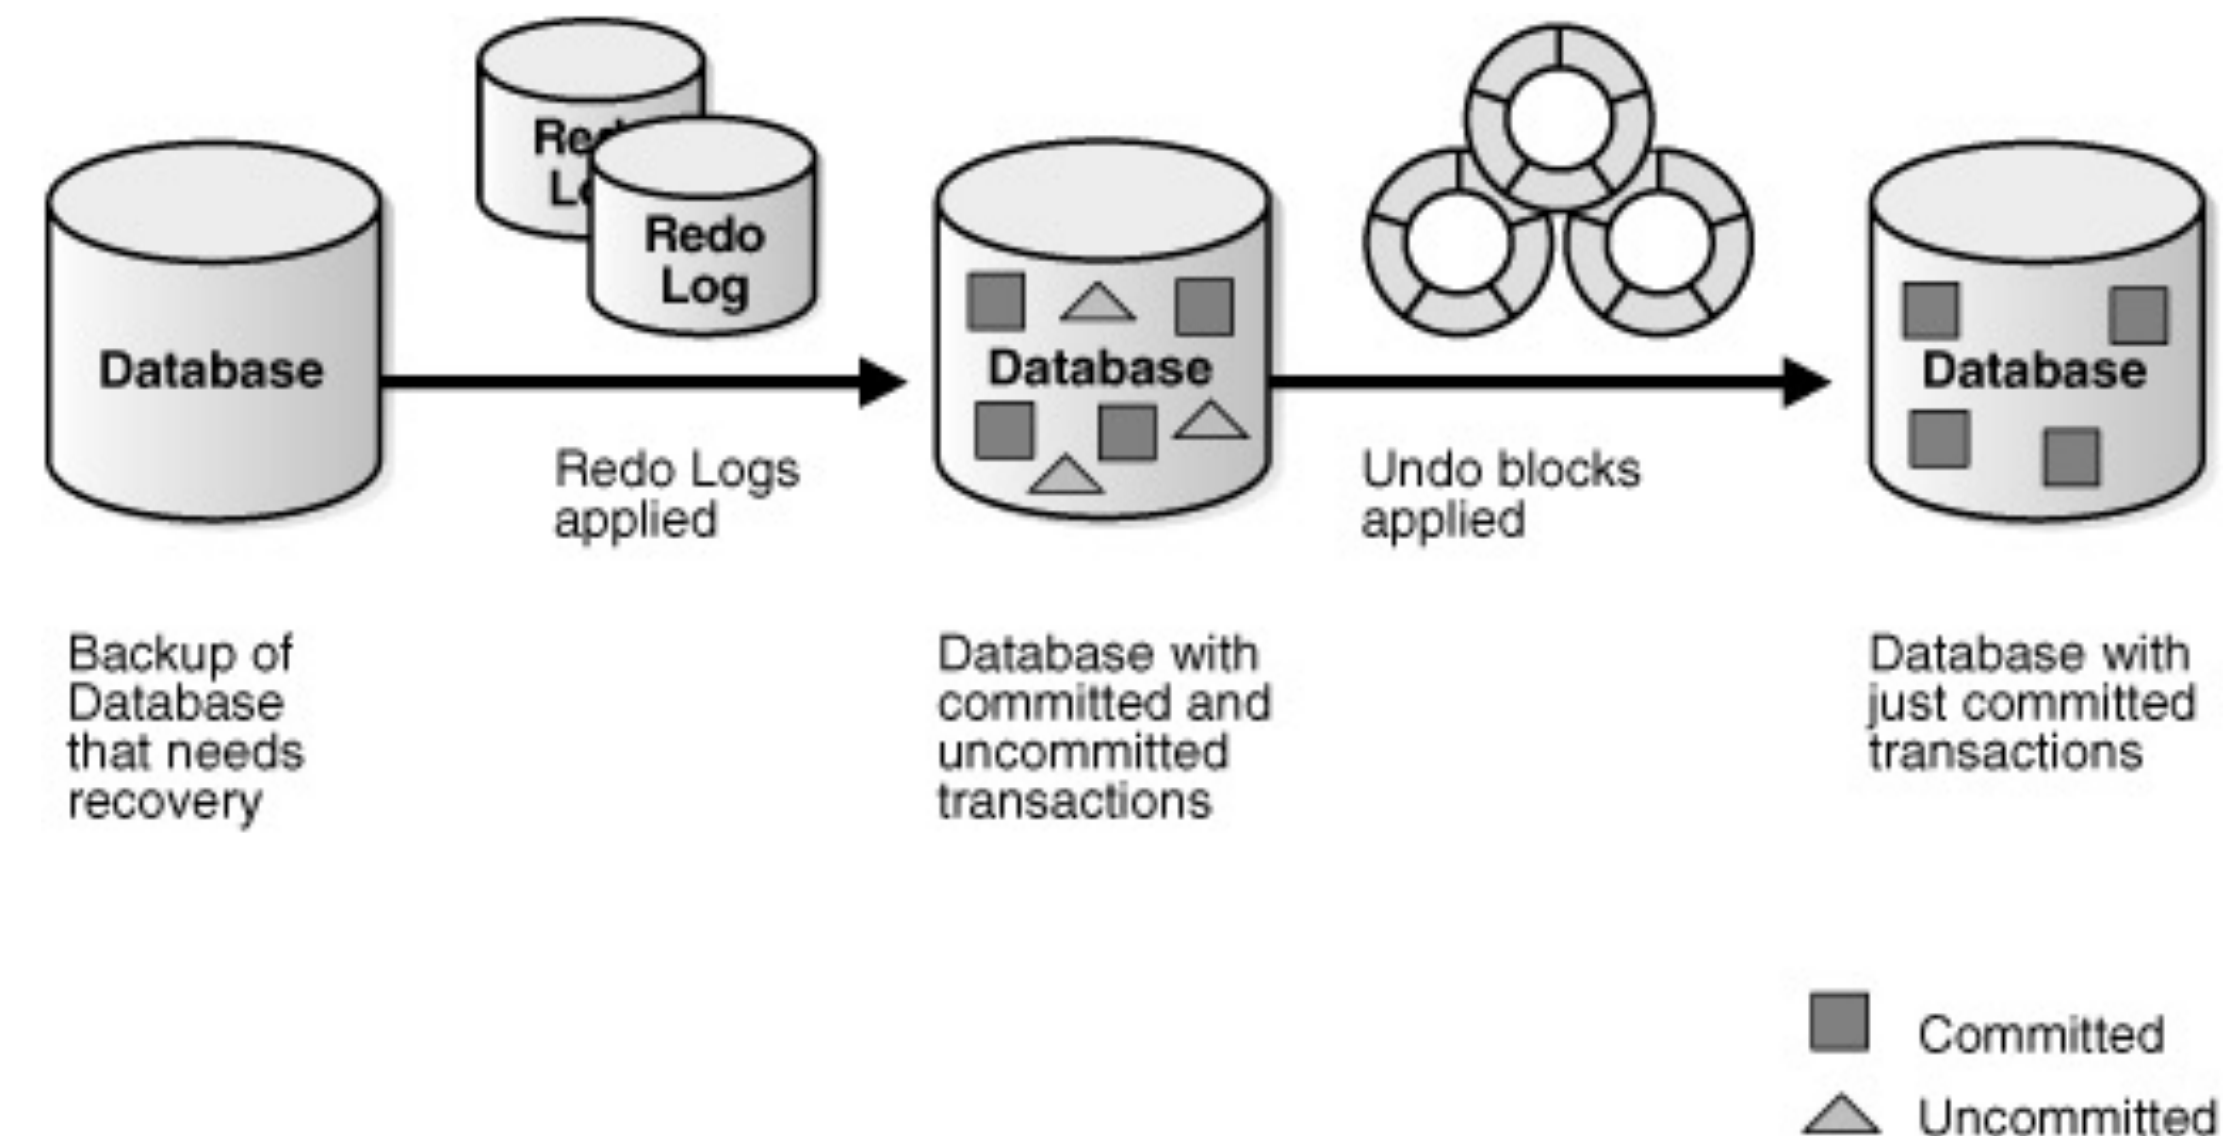
“Obey legal protocols”

- guarantee is simply as strong as the defined rules
- *If application-level code translates all its semantics into such constraints, then an ACID system guarantees application-level consistency*
- Is a txn-level property, restricting what the transaction itself can do

D = Durability

“Data is forever”


- A committed transaction cannot be undone by any failure
- What is the price of accomplishing this?
- How do you choose how much to do/pay?



I = Isolation

“Each transaction runs alone”

- Txns run concurrently —> it's as if each one runs on its own
 - *each txn commits before a new one starts*
- Strict isolation:
 - *Txn T has inputs I and outputs O*
 - *Other txns can read I but cannot read or write O*
- Serializable execution & serialization points
- Can sacrifice serializability for performance
 - *Hard to do ACID at scale*
 - *Introduces complexity in applications*



A = Atomicity
C = Consistency
I = Isolation
D = Durability

Nested Transactions

Nested transactions

- Customer calls the travel agent giving destination and travel dates.
Agent negotiates with airlines for flights.
Agent negotiates with car rental companies for cars.
Agent negotiates with hotels for rooms.
Agent receives tickets and reservations.
Agent gives customer tickets and gets credit card number.
Agent bills credit card.
Customer uses tickets.
- Each step is a transaction and an action at the same time

Redefining the transaction

- Transaction = collection of
 - *Unprotected actions (don't require redo/undo)*
 - *Protected actions (need to be undoable/redoable)*
 - *Real actions (may be deferred but not undone)*
 - *Nested transactions which may be undone by invoking compensating transactions*
- Nested txns != protected actions
 - *effects are visible to the outside world prior to the commit of the parent transaction*
- Nested txn returns the name and params of the compensating txn
 - *keep in log of the parent txn*

Transactional Memory

Transactional memory: Overview

- concurrency control mechanism
- provide ACI but no D
- can be implemented in HW or SW

```
atomic {  
    ...  
}
```

```
atomic (numElements > 0) {  
    ...  
}
```

Intel “Transactional Synchronization Extensions” (TSX)

- Available in Intel’s Skylake and ARM
- RTM = Restricted Transactional Memory
- Three new instructions:
 - XBEGIN = start txnal execution
 - XEND = end txnal execution
 - XABORT = abort txnal execution

Intel TSX: XBEGIN and XEND

- Operand provides a relative offset to the fallback instruction address
 - *If the RTM region could not be successfully executed transactionally, jumps there*
 - *Post-abort, architectural state corresponds to that just before XBEGIN (eax contains abort status)*
- XBEGIN instruction does not have fencing semantics
 - *but, upon abort, all memory updates inside RTM region are invisible*
 - *same semantics as LOCK-prefixed instructions but without the cost*
- Intel provides no guarantee that the RTM region will eventually commit

Intel TSX: XABORT

- abort the execution of an RTM region explicitly
- takes an 8-bit immediate argument for status code (goes into `eax`)

```

__inline unsigned int _xbegin() {
    unsigned status;
    __asm {
        move    eax, 0xFFFFFFFF    // put _XBEGIN_STARTED in eax
        xbegin _txnL1
        _txnL1:
        move    status, eax
    }
    return status;
}

```

Q: Can we pass to xbegin the address of some fallback code other than the instruction immediately following xbegin?

A: In principle yes, but keep in mind that, upon reaching that code, the registers and memory are restored to their state just prior to executing xbegin. The easiest is to transfer control as in the example here; if control is transferred elsewhere, then you will have to explicitly handle the discrepancies between the actual and expected state at that point. If you don't write the machine code directly, then the compiler will have, e.g., allocated variables to registers in a way that getting to that fallback code with the register and memory state of xbegin will confuse the program and exhibit undefined behavior.

```

while (1) {
    unsigned status = _xbegin(); // start transaction
    if (status == _XBEGIN_STARTED) {
        (*g)++; // non atomic increment of shared global variable
        /*... do more stuff ...*/
        _xend();
        break; // break on success
    } else if (status == _XABORT_RETRY) {
        // try again
    } else {
        // fallback path
        LOG("couldn't update global variable");
    }
}

```

```

__inline unsigned int _xbegin() {
    unsigned status;
    __asm {
        move    eax, _XBEGIN_STARTED
        xbegin _txnL1
        _txnL1:
        move    status, eax
    }
    return status;
}

```

```
while (1) {
    unsigned status = _xbegin(); // start transaction
    if (status == _XBEGIN_STARTED) {
        (*g)++; // non atomic increment of shared global variable
        /... do more stuff .../
        _xend();
        break; // break on success
    } else if (status == _XABORT_RETRY) {
        // try again
    } else {
        // fallback path
        LOG("couldn't update global variable");
    }
}

__inline unsigned int _xbegin() {
    unsigned status;
    __asm {
        move    eax, _XBEGIN_STARTED
        xbegin _txnL1
        _txnL1:
        move    status, eax
    }
    return status;
}
```

The diagram illustrates the execution flow between the `_xbegin()` function and the `while` loop. A green line, labeled with a circled '1', originates from the `_xbegin()` call in the `while` loop and points to the function definition. A red line, labeled with a circled '2', originates from the `return status;` line in the `_xbegin()` function and points to the `unsigned status = _xbegin();` assignment in the `while` loop.

Transactional memory: Virtues

- No deadlocks => even bad programmers can write concurrent code
- No lost locks on crash of a thread
- Can get additional parallelism (as long as contention is low)
- Composition of transactions is smoother
 - *it's relegated to runtime (i.e., when you put two lock-based correct pieces of code together, they don't work well anymore, but with atomic{} they do)*

Transactional memory: Limitations

- Inherent in the tension between high / low levels of abstraction
- Long-running txns are more likely to abort
- Limits to what you can do within a txn in systems code ("real" actions)
- Interacting with legacy/non-txnal code
 - *If uses locks, shared memory, etc. —> then what?*
 - *Transaction abstraction is hard to deploy "incrementally"*
- Hard to debug due to non-determinism of aborted transaction
- Performance

Recap

- Transactions in real life => transaction abstraction
- True transactions = ACID
- Can nest transactions (but not trivially)
- Transactional memory updates