

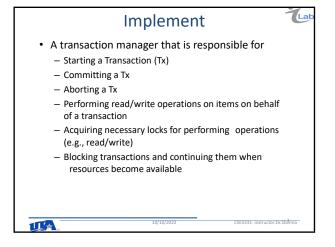
#### Note on Canvas and deadline

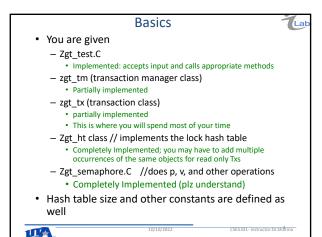


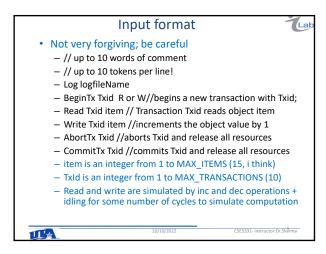
- Deadline indicated on the project description is what you should go by
- If there is an extension, I will make an announcement on Canvas explicitly indicating that

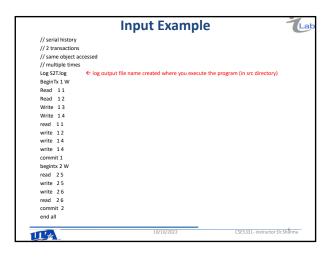
11/2

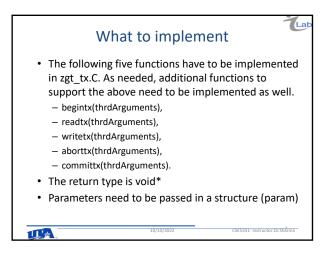
2 CSE5331- Instructor:Dr.Shail

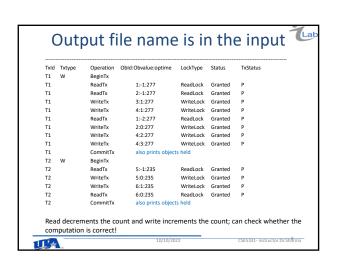


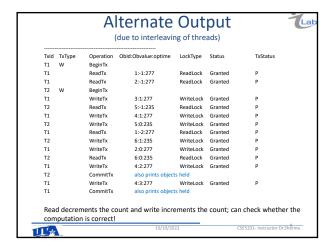


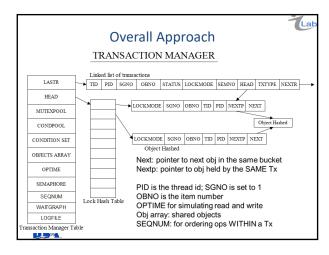


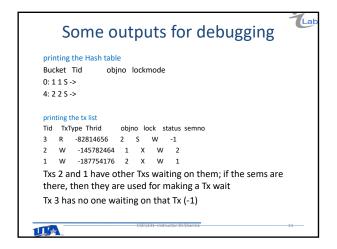


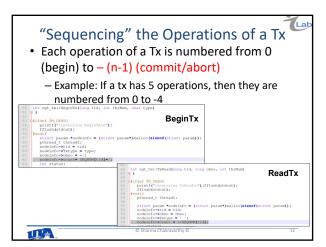


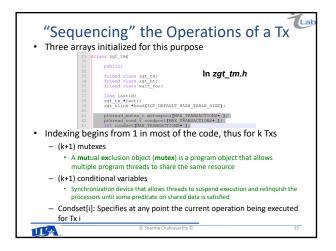


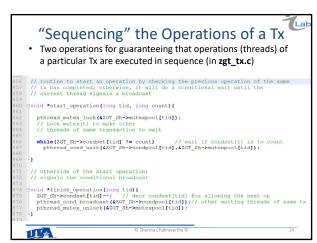




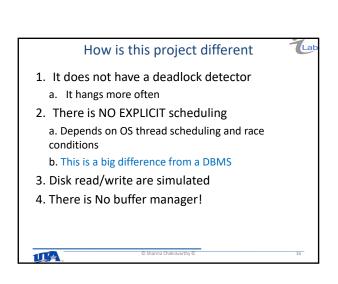








# Keep these in mind 1. Don't forget to put your team no in zgt\_tm.C 2. ipcs –s lists semaphores NOT released by you 3. ipcs\_cleanup.sh releases semaphores held by you 4. When your program hangs, it does not release resources; use ipcs\_cleanup 5. Understand the output; comes handy for debugging 6. Once everything is working, turn off DEBUGFLAGS to stop debugging output 7. You can redirect the output using >& to debug



#### Flow and Tx states

- The main thread (in zgt\_test.C) creates a transaction manager object and the needed hash table in the main or test program. There is only one transaction manager object. However, there will be one transaction object for each transaction, created by begin Tx input.
- Transaction states (reflected in the tx object)
  - TR ACTIVE (P)
  - TR WAIT (W)
  - TR\_ABORT, (A)
  - TR\_COMMIT (E)

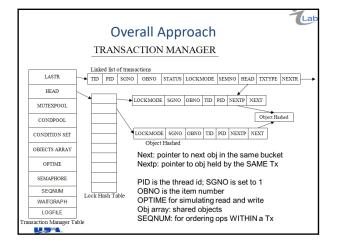


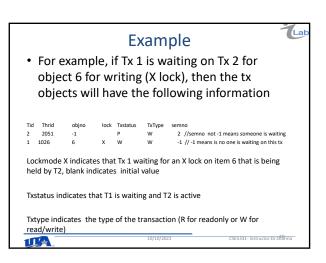
# Locking of objects

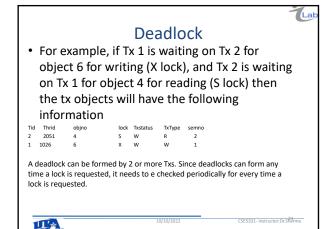


- An object is inserted into the hash table if a lock can be obtained for that object by that tx. The presence of an object in the lock table indicates that that object is being used by a tx. Lockmode in the tx object indicates the type of lock a Tx is waiting for (S or X). TxType is used to indicate the type of the tx (R or W).
- · All Txs are linked using lastr
- · Head points to the hash table
- · All objects within the same bucket are linked using next
- Head of Tx object points to the objects held by that tx as a list (using nextp of object)
- Semno in the tx object is used to make other txs wait for that tx on that semno (Tx k uses semno k)







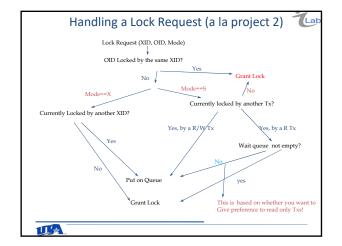


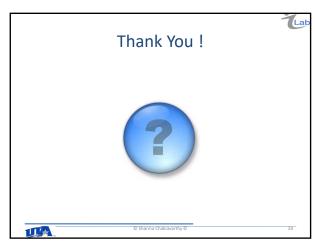
### **Important**

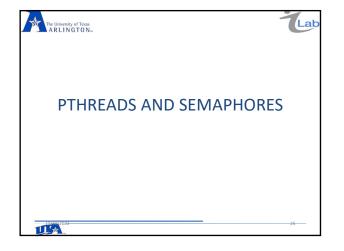
- La
- Lock table needs to be locked for every operation
  - This is done by using one semaphore for the entire table
  - In this implementation sem (an attribute of TxMgr object) is an array of locks
  - Sem 0 is used for the TM table table, sem k by Tx k
    - Sem 0 is initialized to 1 to allow first operation
    - Others sems are initialized to 0 as a p operation is done to make a

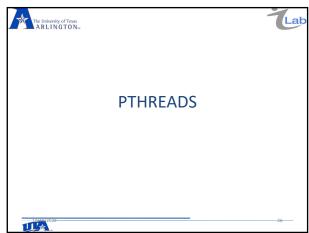
      Tx wait!
- Important: Hold a lock for the shortest duration
- Never suspend/wait holding a lock
- Make sure all p operations have a corresponding v operation (irrespective of the conditionals and flow)











#### **Pthreads**

Lab

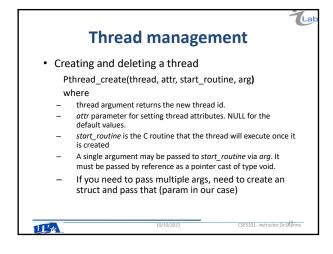
- · To take full advantage of the capabilities provided by threads, a standardized programming interface was required.
- For UNIX/Linux systems, this interface has been specified by the IEEE POSIX 1003.1c standard (1995).
- Implementations which adhere to this standard are referred to as POSIX threads, or Pthreads.

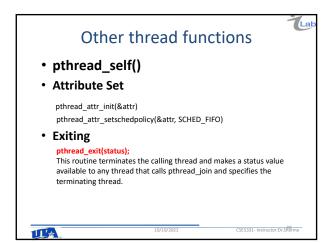
11/\*

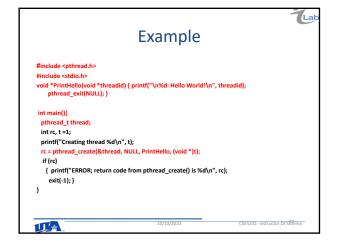
#### **Thread Basics**

- Multiple threads can be created within a process.
- Threads use process resources and exist within a process (different from a real DBMS)
- Scheduled by the operating system (you have some control over its scheduling, can specify FIFO, etc.)
- Run as independent entities within a process.
- If the main program blocks, all the threads will block.

Lab









# Synchronization primitives

- Semaphores
- Mutexes
- Condition variables
- We will be using all of the above and I want you to understand clearly why!



# Synchronization primitives



- Semaphores
- A locking mechanism
- Any thread can acquire and release
- Generalization of mutex.
- A semaphore restricts the number of simultaneous users of a shared resource up to a maximum number
- Operations: p and v (Dijkstra)
- Think of 4 toilets with 4 keys. 4 people can be using the resource at the same time!
- We use this for the lock table



# semaphores

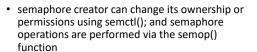
- Array of semaphores are generated by semid = semget(key, nsems, semflg) where nsems = 0 to no of transactions.
- Semaphore O(SHARED\_MEM\_AVAIL) is for locking the transaction manager.
- Semaphores: 1 to no\_of\_transactions are used for threads to wait when objects are locked by other transactions.

11/2

10/10/2022

CSE5331- instructor: Dr. Sharma

## semaphore



- Semaphore 0 is initialized to 1 i.e., holds one resource (transaction manager). Do 'p'(zgt\_p) operation to obtain the resource and 'v'(zgt\_v) to release the resource.
- Rest of semaphores are initialized to 0 i.e., hold no resources. Hence on the first p operation, the thread/process will wait till a v operation is done on the semaphore.

10/10/2022

CSE5331- instructor:Dr.SRarma

# Synchronization primitives (2)

- Mutexes: Deals with synchronization, which is an abbreviation for "mutual exclusion"
  - A semaphore with count as 1
  - A signaling mechanism
  - There is ownership with mutex
- Only the owner can release the lock
- Used for exclusive access to a shared resource (critical section)
- Operations: lock, unlock
- This is like the key to the door of the bathroom!
  Only one person can use at a time!



10/10/2022

CCCC221 Instructor Dr CRism

# Synchronization primitives (3)



- Condition variables (CV): Condition variables provide yet another way for threads to synchronize.
  - While mutexes implement synchronization by controlling thread access to data, condition variables allow threads to synchronize based upon the actual value of data/condition.
  - A thread can wait on a CV and then the resource producer can signal or broadcast the variable
  - Tied to a mutex for mutual exclusion
  - Wait for event and signal or broadcast
    - Signal if any thread can proceed
    - Broadcast if you have to select a thread based on Cv value!!
    - We use this for sequencing operations of a Tx. Using conset and SEONUM

111/2

10/10/2022

CSE5331- instructor:Dr.Sharma

#### **Condition Variable**

- To synchronize thread A and B
  - Declare and initialize global data/variables for synchronization. e.g:condset[tid] =0
  - Declare and initialize a condition variable object.
    - pthread\_cond\_init (condition,attr)
  - Create and initialize associated mutex.
    - pthread\_mutex\_init (mutex,attr)
  - Create threads A and B to do work.



10/10/2022

CSE5331- Instructor: Dr.SHarma

#### Thread A

- Lock associated mutex
- Change the value of the variable (If condset[tid] =0, set it to −1)
- ...... operations .....
- Set the global variable condset[tid]= 0, for thread B to continue
- pthread\_cond\_signal(condition) or pthread\_cond\_broadcast(condition)
- Unlock mutex
- Continue

#### Thread B

- Lock associated mutex and check value of a variable(condset[tid]=0)
- Call pthread\_cond\_wait to perform a blocking wait if condset[tid]!= 0. Note that a call to pthread\_cond\_wait automatically and atomically unlocks the associated mutex variable so that it can be used.
- When signalled, wake up. Mutex is automatically and atomically
  looked.
- Explicitly unlock mutex after completion of operation.
- Continue



CSE5331- instructor:Dr.Sharma

\_\_\_\_\_ a 40

# **Compilation Details**

- Pthreads are defined as a set of C language programming types and procedure calls, implemented with a pthread.h
- a thread library 'pthread' has to be linked. ie. -lpthread



10/10/2022

41