

Reliable Transport Service


Shankar

June 10, 2014

Overview: Reliable transport service w/o Abort

Program: Reliable transport service w/o Abort

Reliable transport service with Abort



- Like reliable streaming Internet sockets
 - `addr` \leftrightarrow [ip addr, tcp port]
- User starts as **server** (accept) or **client** (connect)
- **Client-server** and **client-client** connections
- **Tx/rx data** on connection
- **Graceful closing**
- Client or server can be **first to open** // Transaction TCP
- **Non-abortable** // resend until ack rcvd

- `j.accept()` // willing to accept a conn req
 - `j` enters **server mode** (if not already so)
 - call blocks until connected or canceled
 - returns `[k]`: connected to `k`
 - overlapping `k.connect(j)`
 - returns `[]`: canceled
- `j.endServer()`
 - ends server mode at `j`
 - cancels any ongoing accept call
 - deletes any buffered conn reqs
- In server mode, incoming conn reqs are buffered
 - to be handled (by future accept call) or canceled

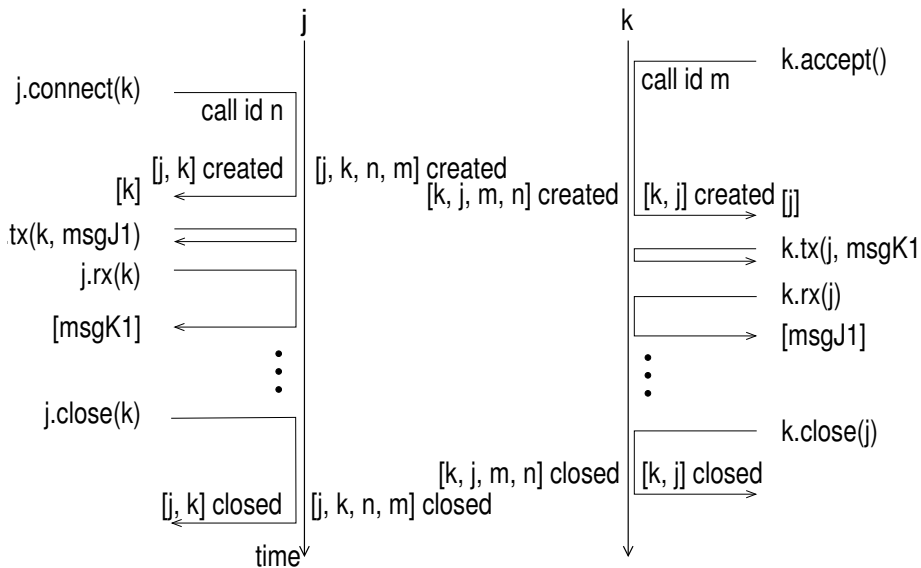
- `j.connect(k)` // request connect to k
 - blocks until connected or rejected
 - returns `[k]`: connected to k
 - returns `[]`: rejected
- `j.tx(k, msg)` // send msg to k
 - call only when j connected to k
- `j.rx(k):` // rcv msg from k
 - call only when j connected to k
 - blocking
 - returns msg or null (if remote closing and all msgs rcvd)

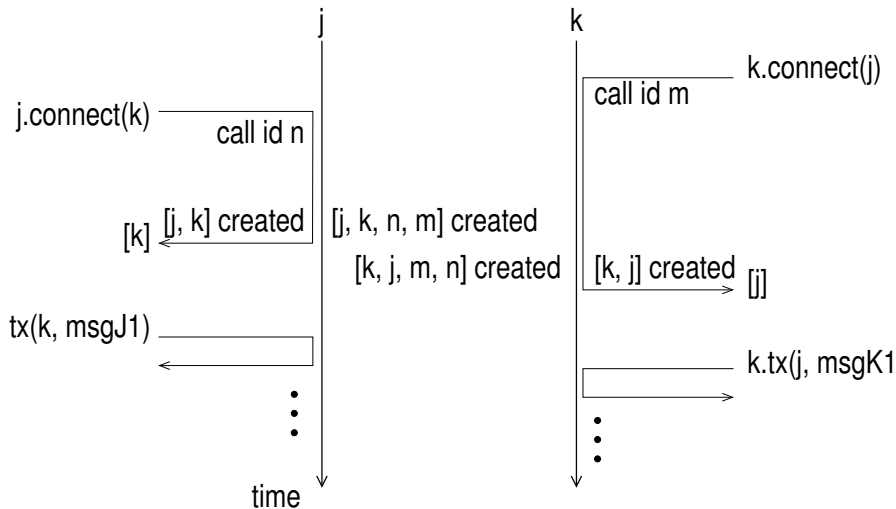
- `j.close(k)` // no further `j.tx` on connection
 - call only when `j` connected to `k`
 - returns only after // graceful closing
 - remote has called `k.close(j)`
 - all msgs sent by remote have been rcvd
 - return cancels any ongoing `j.rx(k)`

- `j` can have multiple connections ongoing (to different addrs)
 - customary to identify connection to `k` by `[j, k]` // socket

- A connection should not see msgs of previous connections
 - $[j,k]$ socket should rcv msgs only from the $[k,j]$ socket to which it connected
 - and not from $[k,j]$ sockets of previous $j-k$ connections
- Conditions that ensure a connect call succeeds
 - progress, not performance
- ...

- Tag every connect and accept call with a unique **call id**
- **Service endpoint**: 4-tuple of socket and local/remote call ids
- Example
 - `j.connect(k)` with call id **n** connects to `k.accept` with call id **m**
 - endpoint at `j`: `[j,k,n,m]`
 - corresponding endpoint at `k`, **if created**: `[k,j,m,n]`
- Require `[j,k,n,m]` to rcv msgs only from `[k,j,m,n]`
- Call ids are **internal** params // internal nondeterminism





Overview: Reliable transport service w/o Abort

Program: Reliable transport service w/o Abort

Reliable transport service with Abort

- Endpoint $[j,k,n,m]$
 - **opened**: has been created
 - **closed**: opened, then `j.close(k)` called and returned
 - **open**: opened and not closed
 - **closing**: open and `j.close(k)` ongoing
- **connecting(j,k,n)**: ongoing `j.connect(k)` with cid `n`
- **accepting(j,n)**: ongoing `j.accept()` with cid `n`
- **openedTo(j,n)**: endpoint $[k,j,m,n]$ opened for some `k, m`
- **overlapped(j,k,n,m)**: `connecting(j,k,n)` overlapped with `accepting(k,m)` or with `connecting(k,j,m)`

- `connecting(j,k,n)` has **dedicated accept**:
 - while `connecting(j,k,n)`
 - `k.accept()` is ongoing (and not canceled)
 - no other client attempts to connect to `k`
 - when `connect` was called
 - no `[k,j]` socket
 - no ongoing `j.accept` // to avoid potential `k.connect(j)`
- `accepting(j,n)` has **dedicated connect**:
 - while `accept` ongoing, for some `k`
 - `k.connect(j)` is issued
 - there is no `[j,k]` socket,
 - `j.connect(k)` is not issued // to avoid `k.connect(j)`

- Parameters: ADDR
- Input fns
 - `j.accept()`, `j.endServer()`, `j.connect(k)`, `j.close(k)` // cm
 - `j.tx(k,msg)`, `j.rx(k)` // dt
- No output fns
- Main
 - `cidgen`: call-id generator, initially 0
 - `addrs` in server-mode
 - `cids` of ongoing accepts and connects
 - `cid` pairs of open sockets
 - history of `accept/connect/close` calls and returns
 - `msg txh/rxh` histories for each opened endpoint

- input fn `j.accept()`
 - ic {no ongoing `j.accept`}
 - `n ← cidgen++ ; put j in server-mode`
 - output (`rval`, remote addr `k`, remote cid `m`)
 - oc { `rval = [k]` & overlapped(`k,j,m,n`) & no open `[j,k]` &
(not openedTo(`j,n`) or opened(`k,j,m,n`))
OR
`rval = []` & ongoing `j.endServer` & not openedTo(`j,n`)
}
 - update state ; return `rval`

- input fn `j.connect(k)`
 - ic {no open or connecting [j,k]}
`n ← cidgen++ ; update state`
 - output (`rval`, remote cid `m`)
oc { `rval = [k]` & overlapped(`j,k,n,m`) & no open [j,k] &
(not openedTo(`j,n`) or opened(`k,j,m,n`))
OR
`rval = []` & no dedicated accept & not openedTo(`j,n`)
}
`update state ; return rval`

- input fn `j.endServer()`
 - ic {no ongoing `j.endServer`}
remove `j` from server-mode
 - oc {no ongoing `j.accept`}
update state ; return
- input fn `j.tx(k, msg)`
 - ic {[`j,k`] open and not closing & no ongoing `j.tx(k,.)`}
update state
 - oc {true}
return

- input fn `j.endServer()`
 - ic {no ongoing `j.endServer`}
remove `j` from server-mode
 - oc {no ongoing `j.accept`}
update state ; return

- input fn `j.tx(k, msg)`
 - ic {[`j,k`] open and not closing & no ongoing `j.tx(k,.)`}
update state
 - oc {true}
return

- input fn $j.rx(k)$
 - ic $\{[j,k] \text{ open} \ \& \ \text{no ongoing } j.rx(k)\}$
 - output $(rval, msg)$
 - let $[n,m]$ be cid-pair of open $[j,k]$
 - oc $\{rval = [0, msg] \ \& \ (rxh[n,m] \circ [msg] \text{ prefix-of } txh[m,n])$
OR
 $rval = [-1] \ \& \ rxh[n,m] = txh[m,n] \ \&$
 $\text{closing}(k, j, m, n) \text{ or } \text{closed}(k, j, m, n)$
- update state ; return $rval$

- input fn `j.close(k)`
 - ic `{[j,k] open & no ongoing j.close(k)}`
 - let `[n,m]` be cid-pair of open `[j,k]`
 - oc `{no ongoing j.tx(k,.) or j.rx(k) &`
`rxh[n,m] = txh[m,n] &`
`closing(k,j,m,n) or closed(k,j,m,n)`
 - }
 - update state ; return

- Progress assumption
 - $\text{accepting}(j,n) \ \& \ (\text{dedicated connect or ending server})$
leads-to $\text{not accepting}(j,n)$
 - $\text{connecting}(j,k,n)$ *leads-to* $\text{not connecting}(j,k,n)$
 - $\text{opened}(j,k,n,m)$ *leads-to* $\text{opened}(k,j,m,n)$
 - $\text{ongoing } j.\text{endServer}$ *leads-to* $\text{no ongoing } j.\text{endServer}$
 - $\text{ongoing } j.\text{tx}(k,.)$ *leads-to* $\text{no ongoing } j.\text{tx}(k,.)$
 - $\text{ongoing } j.\text{endServer}$ *leads-to* $\text{no ongoing } j.\text{endServer}$

// continued

// continued

■ Progress assumption

- ongoing $j.tx(k,.)$ *leads-to* no ongoing $j.tx(k,.)$
- ongoing $j.rx(k)$ & open(j,k,n,m) & ($rxh[n,m] \neq txh[m,n]$ or closing(k,j) or closed(k,j,m,n))
leads-to no ongoing $j.rx(k)$
- closing(j,k,n,m) & (ongoing $k.close(j)$ or closed(k,j,m,n))
leads-to (no ongoing $j.close(k)$ or no ongoing $j.rx(k)$)

Overview: Reliable transport service w/o Abort

Program: Reliable transport service w/o Abort

Reliable transport service with Abort

- **Abortless** service is implementable over an LRD channel
 - need to resend a “to-be-acked” msg until the ack is rcvd
 - impractical in context of failures
- **Abortable** service: abort after specified number of resends
- Abortable service has many more acceptable evolutions
 - endpoint [j,k,n,m] opens but endpoint [k,j,m,n] never opens
 - j.close(k) returns before remote is closing or all data rcvd
 - j.connect(k) returns unsuccessfully even if it is not rejected