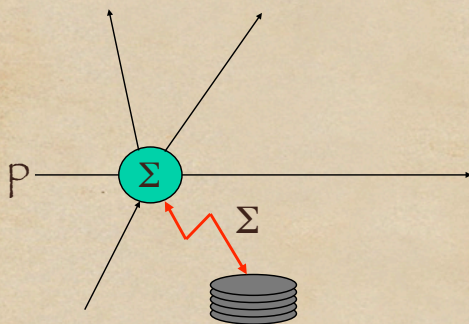


Rollback-Recovery

Uncoordinated Checkpointing



- Easy to understand
- No synchronization overhead
- Flexible
 - can choose **when** to checkpoint
- To recover from a crash:
 - go back to last checkpoint
 - restart


How (not) to take a checkpoint

- Block execution, save entire process state to stable storage
 - very high overhead during failure-free execution
 - lots of unnecessary data saved on stable storage


How to take a checkpoint

- Take checkpoints incrementally
 - save only pages modified since last checkpoint
 - use “dirty” bit to determine which pages to save
- Save only “interesting” parts of address space
 - use application hints or compiler help to avoid saving useless data (e.g. dead variables)
- Do not block application execution during recovery
 - copy-on-write


The Domino Effect




The Domino Effect




The Domino Effect




The Domino Effect




The Domino Effect




The Domino Effect




The Domino Effect




The Domino Effect



The Domino Effect



The Domino Effect



The Domino Effect



How to Avoid the Domino Effect


Coordinated Checkpointing

- No independence
- Synchronization Overhead
- Easy Garbage Collection

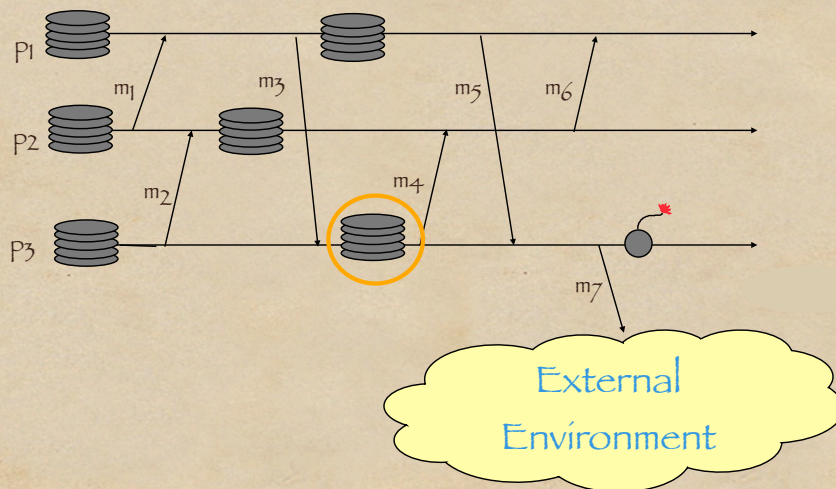
Communication Induced Checkpointing : detect dangerous communication patterns and checkpoint appropriately

- Less synchronization
- Less independence
- Complex


The Output Commit Problem



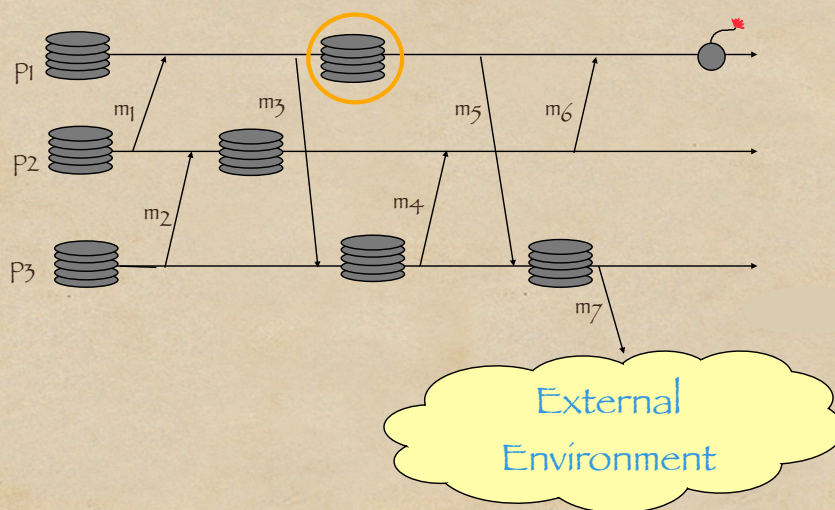
The Output Commit Problem



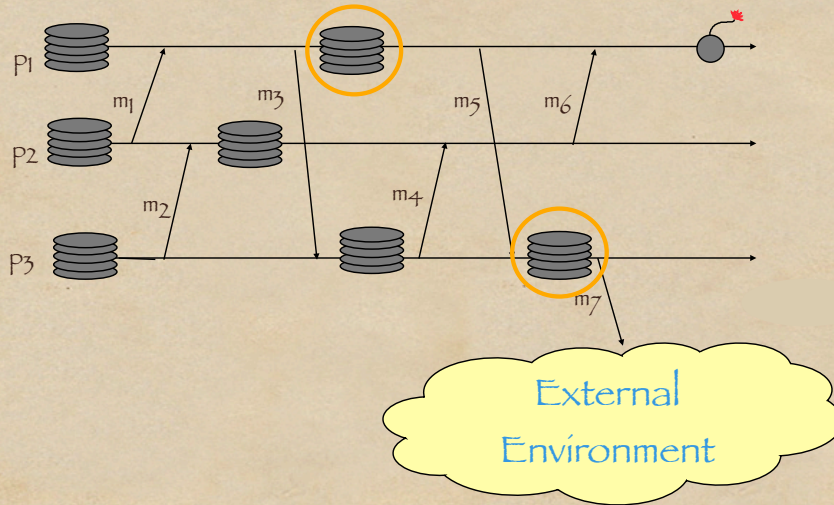
The Output Commit Problem




The Output Commit Problem



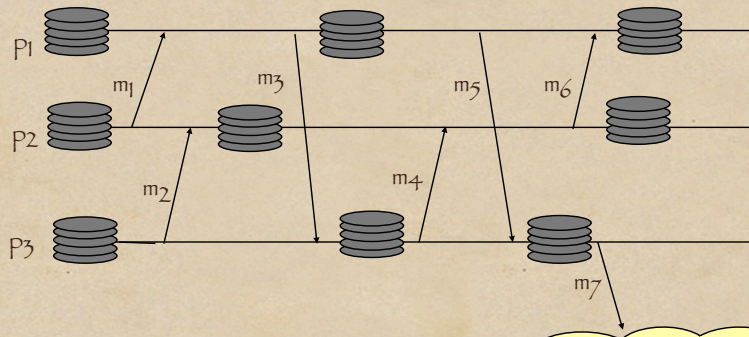
The Output Commit Problem



The Output Commit Problem




The Output Commit Problem



- Coordinated checkpoint for every output commit
- High overhead if frequent I/O with external environment

External Environment


Distributed Checkpointing at a Glance



Independent


- + Simplicity
- + Autonomy
- + Scalability
- Domino effect

at a Glance



Coordinated

- + Consistent states
- + Good performance
- + Garbage Collection
- Scalability



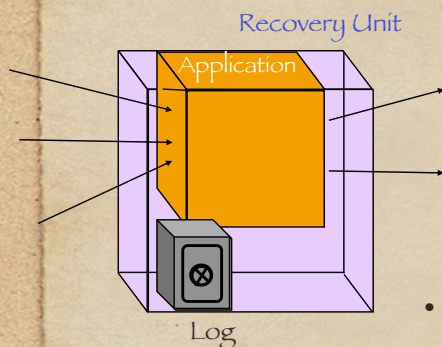
Communication-induced

- + Consistent states
- + Autonomy
- + Scalability
- None is true

Message Logging

- Can avoid domino effect
- Works with coordinated checkpoint
- Works with uncoordinated checkpoint
- Can reduce cost of output commit
- More difficult to implement

How Message Logging Works




To tolerate crash failures:

- periodically checkpoint application state;
- log on stable storage **determinants** of non-deterministic events executed after checkpointed state.
- for message delivery events:
 $\#m = (m.dest, m.rsn, m.source, m.ssn)$

Recovery:

- restore latest checkpointed state;
- replay non-deterministic events according to determinants

Pessimistic Logging



Never creates orphans

- may incur blocking
- straightforward recovery


Sender Based Logging

(Johnson and Zwaenepoel, FTCS 87)


Message log is maintained in volatile storage at the sender.

A message m is logged in two steps:

- before sending m , the sender logs its content: m is **partially logged**
- the receiver tells the sender the receive sequence number of m , and the sender adds this information to its log: m is **fully logged**.



Optimistic Logging



Eliminates orphans during recovery

- non-blocking during failure-free executions
- rollback of correct processes
- complex recovery

Causal Logging

- ✓ No blocking in failure-free executions
- ✓ No orphans
- ✓ No additional messages
- ✓ Tolerates multiple concurrent failures
- ✓ Keeps determinant in volatile memory
- ✓ Localized output commit

Preliminary Definitions

Given a message m sent from $m.source$ to $m.dest$,

$$\text{Depend}(m): \left\{ p \in P \mid \begin{array}{l} \forall (p = m.dest) \text{ and } p \text{ delivered } m \\ \forall (\exists e_p : (deliver_{m.dest}(m) \rightarrow e_p)) \end{array} \right\}$$

$\text{Log}(m)$: set of processes with a copy of the determinant of m in their volatile memory

p orphan of a set C of crashed processes:

$$(p \notin C) \wedge \exists m : (\text{Log}(m) \subseteq C \wedge p \in \text{Depend}(m))$$

The “No-Orphans” Consistency Condition

No orphans after crash C if:

$$\forall m : (\text{Log}(m) \subseteq C) \Rightarrow (\text{Depend}(m) \subseteq C)$$

No orphans after any C if:

$$\forall m : (\text{Depend}(m) \subseteq \text{Log}(m))$$

The Consistency Condition

$$\forall m : (\neg \text{stable}(m) \Rightarrow (\text{Depend}(m) \subseteq \text{Log}(m)))$$

Optimistic and Pessimistic

No orphans after crash C if:

$$\forall m : (\text{Log}(m) \subseteq C) \Rightarrow (\text{Depend}(m) \subseteq C)$$

Optimistic weakens it to:

$$\forall m : (\text{Log}(m) \subseteq C) \Rightarrow \diamond(\text{Depend}(m) \subseteq C)$$

No orphans after any crash if:

$$\forall m : (\neg \text{stable}(m) \Rightarrow (\text{Depend}(m) \subseteq \text{Log}(m)))$$

Pessimistic strengthens it to:

$$\forall m : (\neg \text{stable}(m) \Rightarrow |\text{Depend}(m)| \leq 1)$$

Causal Message Logging

No orphans after any crash of size at most f if:

$$\forall m : (\neg \text{stable}(m) \Rightarrow (\text{Depend}(m) \subseteq \text{Log}(m)))$$

Causal strengthens it to:


$$\forall m : \left(\neg \text{stable}(m) \Rightarrow \left(\begin{array}{l} \wedge (\text{Depend}(m) \subseteq \text{Log}(m)) \\ \wedge \diamond (\text{Depend}(m) = \text{Log}(m)) \end{array} \right) \right)$$

An Example


Causal Logging:

$$\forall m : (\neg \text{stable}(m) \Rightarrow (\text{Depend}(m) \subseteq \text{Log}(m)))$$

$$\text{If } f = 1, \text{ stable}(m) \equiv |\text{Log}(m)| \geq 2$$




Recovery for $f = 1$




parents of p
Messages previously sent to p by its parents

SSN order

what is the next message from each parent?



Determinants of messages delivered by parents



RSN order
Determinants of messages delivered by p

children of p

Family-Based Logging



Each p maintains $D_p \equiv \{\#m : p \in \text{Depend}(m)\}$
in volatile memory

On sending a message m'

- adds m' to volatile send log
- piggybacks on messages to q all determinants $\#m \in D_p$ s.t.

$$|\text{Log}(m)|_p \leq f \wedge (q \notin \text{Log}(m)_p)$$

On receiving a message m'

- adds to D_p any new determinant piggybacked on m'
- adds $\#m'$ to D_p
- updates its estimate of $|\text{Log}(m)|_p$ for all determinants $\#m \in D_p$

Estimating $\text{Log}(m)$ and $|\text{Log}(m)|$

Each process p maintains estimates of

$$\text{Log}(m)_p \text{ and } |\text{Log}(m)|_p$$

p piggybacks $\#m$ on m' to q if

$$|\text{Log}(m)|_p \leq f \wedge (q \notin \text{Log}(m)_p)$$

- How can p estimate $\text{Log}(m)_p$ and $|\text{Log}(m)|_p$?
- How accurate should these estimates be?
 - inaccurate estimates cause useless piggybacking
 - keeping estimates accurate requires extra piggybacking

The Idea

Because $\forall m : (\neg \text{stable}(m) \Rightarrow (\text{Depend}(m) \subseteq \text{Log}(m)))$

we can approximate $\text{Log}(m)$ from below with:

and then use vector clocks to track $\text{Depend}(m)$!

$$\text{Log}(m) = \begin{cases} \text{Depend}(m) & \text{if } |\text{Depend}(m)| \leq f \\ \text{Any set } S : |S| > f & \text{otherwise} \end{cases}$$

Dependency Vectors

Dependency Vector (DV): vector clock that tracks causal dependencies between message delivery events.

$$\text{deliver}_p(m) \rightarrow \text{deliver}_q(m') \equiv \\ DV_p(\text{deliver}_p(m))[p] \leq DV_q(\text{deliver}_q(m'))[p]$$

Weak Dependency Vectors

Weak Dependency Vector (WDV):

track causal dependencies on $deliver(m)$ as long as

$$(|Depend(m)| \leq f)$$

$$(deliver_p(m) \rightarrow deliver_q(m')) \wedge (|Depend(m)| \leq f) \Rightarrow WDV_p(deliver_p(m))[p] \leq WDV_q(deliver_q(m'))[p]$$

$$WDV_p(deliver_p(m))[p] \leq WDV_q(deliver_q(m'))[p] \Rightarrow deliver_p(m) \rightarrow deliver_q(m')$$

Dependency Matrix

Use WDV to determine if $p \in \text{Log}(m)$:

$$p \in Depend(m) \wedge |Depend(m)| \leq f \Rightarrow WDV_p[m.dest] \geq m.rsn$$

$$WDV_p[m.dest] \geq m.rsn \Rightarrow p \in Depend(m)$$


Each p keeps a **Dependency Matrix** (DM_p)

Given $\#m = \langle u, \text{source } \$, \text{des } 14, \text{ssn } 15 \rangle$,

$$DM_p = \begin{matrix} p \\ q \\ r \\ s \\ t \\ u \end{matrix} \begin{bmatrix} & s & \\ & 21 & \\ & 16 & \\ & 8 & \\ & 21 & \\ & 12 & \\ & 7 & \end{bmatrix}$$

$$\text{Log}(m)_p = \{p, q, s\}$$

Message Logging at a Glance




Pessimistic
 + No orphans
 + Easy recovery
 - Blocks

Optimistic
 + Non-blocking
 - Orphans
 - Complex recovery

Causal
 + Non-blocking
 + No orphans
 - Complex recovery

Communication Induced Checkpointing



- + Consistent states
- + Autonomy
- + Scalability
- + No useless checkpoints

Really?

CIC Protocols

- Independent **local** checkpoints
- **Forced** checkpoints before processing some messages
- Piggyback information about checkpoints on application messages

Always a consistent set of checkpoints without

- explicit coordination
- protocol-specific messages

CIC Protocol Families

Index-Based


- Each checkpoint has an index
- Indices piggybacked on application messages
- Checkpoints with same index are consistent

Pattern-Based

- Detect communication patterns
- Checkpoint to prevent dangerous patterns
- Avoid useless checkpoints

They are equivalent

Example of Index Based




Local checkpoint

Forced checkpoint

After
Briatico, Ciuffoletti & Simoncini 84

Z-Paths




A Z-Path exists between C_{xi} and C_{yj} iff [Netzer & Xu 95]:

$i < j$ and $x = y$ or There exists a path $[m_0, m_1, \dots, m_n]$ such that:


- $C_{xi} \rightarrow send_x(m_0)$
- $\forall l < n$ either $deliver_k(m_l) \rightarrow send_k(m_{l+1})$ or $send_k(m_{l+1}, deliver_k(m_l))$ are in the same checkpoint interval
- $deliver_y(m_n) \rightarrow C_{yj}$

Z-Cycles




A Z-Cycle is a Z-path that begins and ends at the same checkpoint

Z-Cycles & Useless Checkpoints



A checkpoint in a Z-cycle can never be part of a consistent state

Example of Pattern-Based



The forced checkpoint breaks the Z-cycle, preventing the local checkpoint from becoming useless

(Baldoni, Quaglia & Ciciani 98)

Experiment Goals

- How to implement CIC protocols?
- What is the performance?
- How do they scale?
- Which is better, index-based or pattern-based?

Outline

- Implemented 3 CIC protocols in Egida
- Used NASA NPB 2.3 benchmark applications

Appl.	Communication Rate		Communication Pattern	Exec. Time (sec)
	Mess/sec	Size(KB)		
bt	6	50.7	All processes	1530
cg	20	60.7	Two neighbors	1516
lu	62	3.7	Two neighbors	975
sp	22	44.4	All processes	1222

- For most experiments, direct measures
- Simulation to extrapolate for scale
 - Used implementation to validate simulator

The Three Protocols

Index-Based:

- Briatico, Ciuffoletti & Simoncini '84,
 - BCS, $O(1)$ /message
- H elary, Mostefaoui, Netzer & Raynal '97,
 - HMNR, $O(n)$ /message

Pattern-based:


- Baldoni, Quaglia & Ciciani '98,
 - BQC, $O(n^2)$ /message

Autonomy?

Processes take independent checkpoints


But:

- Selecting a checkpointing placement policy is hard
- A process has no control over forced checkpoints



No Useless Checkpoints?

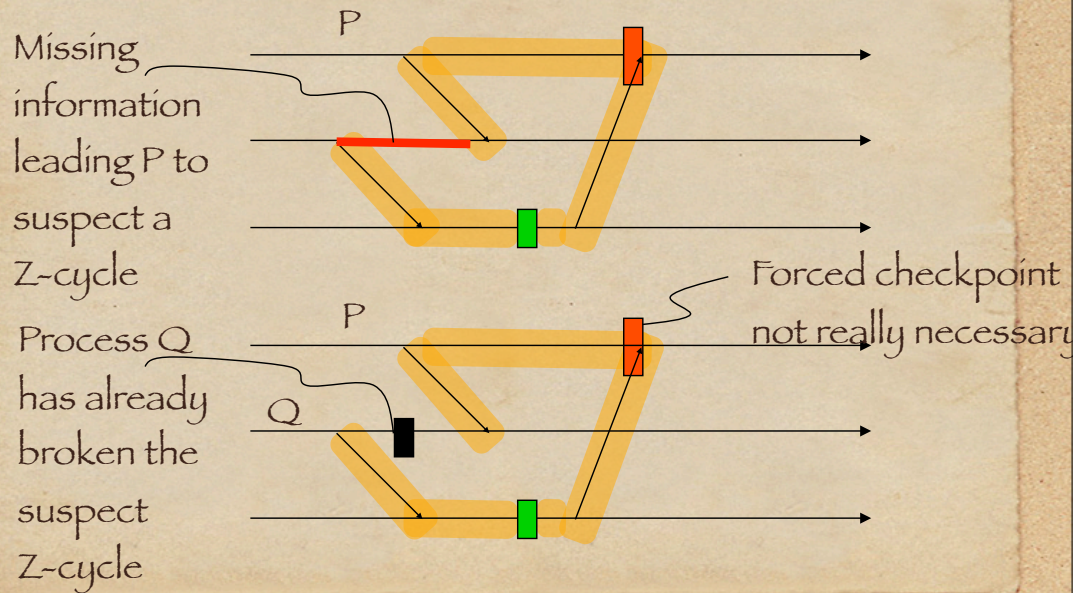
- Yes, but only if checkpoints are blocking!



Checkpoint (6) of p₃ can become useless

p₁ may run garbage collection and discard checkpoint (6)

Analysis of BQC's



Non-Blocking Checkpointing & CIC

Solutions:

Blocking checkpoints, as assumed by protocol

- Performance problems

Be optimistic and don't block

- Violates protocol's invariant

we take non-blocking checkpoints


Execution Overhead



Local checkpoints taken with exponential distribution, mean 6 minutes

- 4 Pentium-based workstations, 300 MHz
- Lightly-loaded 100Mb/s Ethernet
- Stable storage on local disk

Forced vs. Local



Local checkpoints taken with exponential distribution, mean 6 minutes

Scalability



Simulation with 119 local ckp's/proc
Low comm. load of 10 msg/ckp, random pattern

Scalability (cont'd)



Simulation with 118 local ckp's/proc
High comm. load of 500 msg/ckp, uniform pattern

Summary

- Scalability? Not exactly...
- Autonomy in checkpointing? Not exactly...
 - # of forced ckp's is often greater than twice the # of local ones
 - adaptation necessary for good performance
- Unpredictable behavior:
 - Difficult to plan resources, decide on local ckpts, or estimate overhead
- Performs well for random pattern, low-load