Implementing an Approximate Probabilistic Algorithm for Error Recovery in Concurrent Processing Systems

Dr. Silvia Heubach
Dr. Raj S. Pamula

Department of
Mathematics and Computer Science

California State University Los Angeles

Overview

- Problem Statement
- Recovery Algorithms
- Previous Work
- Implementation Issues
- Approximate Probabilistic Algorithm
- Simulation
- Comparison of Performance
- Conclusion

Problem Statement

- Design of a reliable concurrent processing system
- Transient errors detected by acceptance tests (AT) with a delay
- Consistent system states established using (global) checkpoints
- Error recovery is attempted by rolling back to previously established checkpoint and restarting the system
- What is the best method for rollback recovery?

Recovery Algorithms

- Iterative Algorithm
  - Rolls back one checkpoint (CP) at a time, starting from most recent one
  - Recovery (if possible) is achieved from the CP just prior to failure
- Probabilistic Algorithm
  - Based on distribution of latency times
  - Selects CP for recovery by comparing pairs of checkpoints for smaller <u>expected</u> cost of recovery

Previous Work

- Simulation of 4 concurrent processes
- Waiting times between events are independent exponential random variables
- Ranges for average times between
  - inter-connections: .25 - 1 (hours)
  - failures: 10 - 40 (hours)
  - acceptance tests: 1 - 2 (hours)
- Result: Probabilistic Algorithm better if checkpoint interval length < median of latency time distribution (= .45)

Implementation Issues

- Can't measure exact latency times (to get empirical distribution function)
- Theoretical derivation of latency distribution difficult due to dependencies introduced through inter-connections
- Solution: Approximate Probabilistic Algorithm

Approximate Probabilistic Algorithm (APA)

- Use iterative method to measure approximate latency times

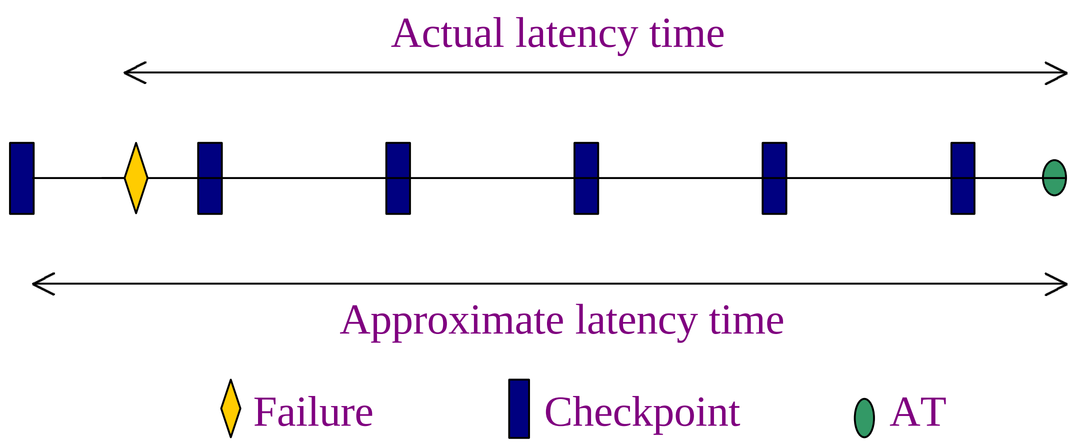

- Compute empirical distribution function of approximate latency times

- Switch to probabilistic method based on this empirical distribution function

Approximate Probabilistic Algorithm (APA)

- APA works for all distributions and allows for dependent events as well
- No parameter estimation needed (as would be necessary if one would be able to derive the theoretical distribution of latency times)
- Easy to implement
- Parameters for Implementation
  - N = number of data points collected using iterative algorithm
  - Length of checkpoint interval (=: error interval) during data collection phase

Simulation

- Two phases
  - data collection
  - performance comparison

- Phase I
  - 10 simulations, each up to 100 hours
  - 231 detected failures
  - exact and approximate latency times measured for error interval lengths 0.1, 0.15, 0.2, and 0.25

➢ Latency Time Distributions

exact

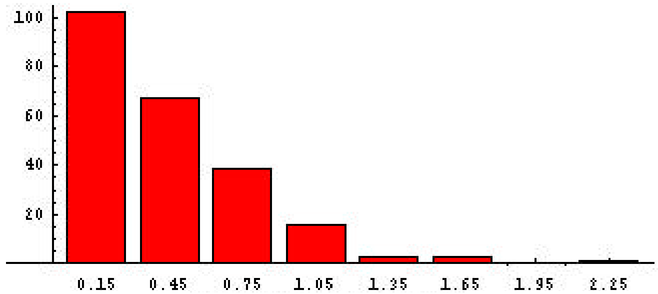

0.1

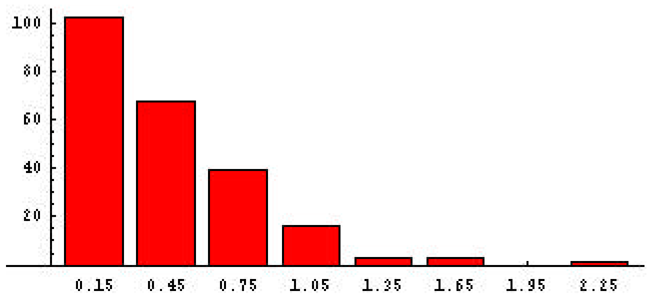

0.15

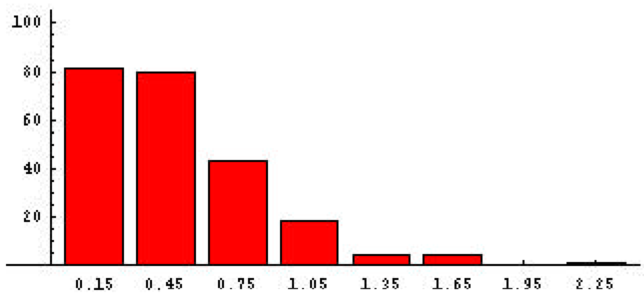

0.2

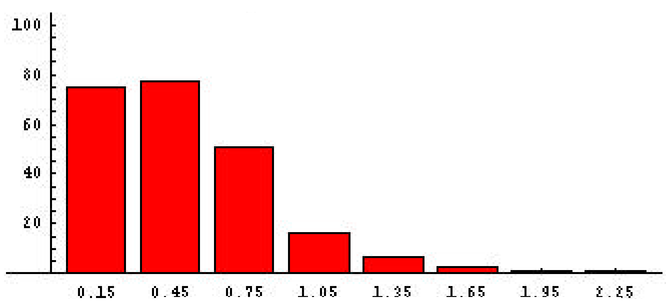

0.25

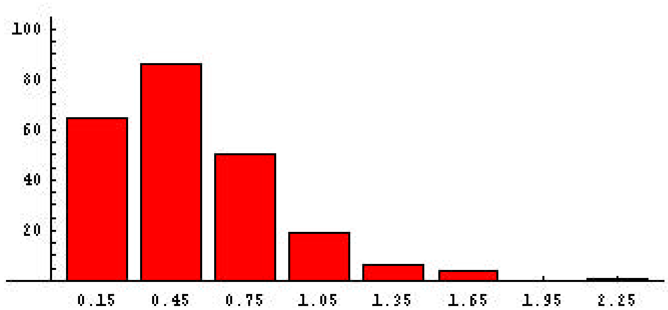

Simulation (Phase II)

- 30 simulations (up to time 100) for each combination of
  - CP interval length C = 0.1, 0.15,....., 0.45
  - Recovery level 90% and 95%
  - N = 25, 50, 75, 100, 150 and 200
- Compute average cost for each case
- Performance measure: Cost of APA as percentage of cost incurred with iterative algorithm

Comparison of Performance

Exact Probabilistic

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.10	56.5	50.5	46.7	49.6	54.2	49.1
	0.15	71.7	65.5	63.8	64.0	65.7	63.4
	0.20	82.5	75.5	72.5	74.2	73.5	76.2
C	**0.25**	88.8	83.1	82.4	81.3	80.7	80.4
	0.30	94.4	88.9	87.2	86.5	89.1	87.6
	0.35	93.6	95.8	95.2	91.7	94.3	96.5
	0.40	97.0	97.9	94.7	99.3	97.6	96.7
	0.45	99.2	99.2	97.3	99.2	96.9	97.1

≤**59.99** **60 - 69.99** **70 -79.99** **80 - 89.99** **>100.5**

- **probabilistic method better**
- **for C = .40 and .45, methods almost identical**
- **if N = 25 is disregarded, get nice efficiency brackets**

Comparison of Performance

Exact Probabilistic

		N					
		25	50	75	100	150	200
	0.10	56.5	50.5	46.7	49.6	54.2	49.1
	0.15	71.7	65.5	63.8	64.0	65.7	63.4
	0.20	82.5	75.5	72.5	74.2	73.5	76.2
C	0.25	88.8	83.1	82.4	81.3	80.7	80.4
	0.30	94.4	88.9	87.2	86.5	89.1	87.6
	0.35	93.6	95.8	95.2	91.7	94.3	96.5
	0.40	97.0	97.9	94.7	99.3	97.6	96.7
	0.45	99.2	99.2	97.3	99.2	96.9	97.1

Error Interval 0.1

		N					
		25	50	75	100	150	200
	0.1	54.9	53.3	53.1	52.1	57.2	51.6
	0.15	79.6	69.8	65.1	67.2	69.7	66.9
	0.2	85.2	76.6	72.8	78.4	76.1	75.8
C	0.25	93.6	80.8	85.1	88.9	84.3	81.7
	0.3	97.5	97.4	87.6	87.1	98.8	94.0
	0.35	**100.8**	95.1	94.4	97.7	93.7	93.5
	0.4	99.7	95.6	95.7	94.6	95.9	96.5
	0.45	**100.7**	**108.9**	99.1	99.5	**107.4**	**109.6**

≤**59.99** **60 - 69.99** **70 -79.99** **80 - 89.99** **>100.5**

Comparison of Performance

Exact Probabilistic

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.10	56.5	50.5	46.7	49.6	54.2	49.1
	0.15	71.7	65.5	63.8	64.0	65.7	63.4
	0.20	82.5	75.5	72.5	74.2	73.5	76.2
C	**0.25**	88.8	83.1	82.4	81.3	80.7	80.4
	0.30	94.4	88.9	87.2	86.5	89.1	87.6
	0.35	93.6	95.8	95.2	91.7	94.3	96.5
	0.40	97.0	97.9	94.7	99.3	97.6	96.7
	0.45	99.2	99.2	97.3	99.2	96.9	97.1

Error Interval 0.15

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.1	60.5	50.6	50.3	52.0	57.1	51.6
	0.15	79.6	68.1	64.9	65.3	70.9	67.8
	0.2	87.7	77.0	72.9	82.1	80.2	76.1
C	**0.25**	95.9	83.9	85.5	89.0	86.8	82.5
	0.3	**100.6**	91.4	88.0	87.9	99.2	94.1
	0.35	**102.6**	96.0	96.4	94.0	94.9	95.3
	0.4	**102.6**	95.7	96.5	95.7	96.1	97.0
	0.45	**107.0**	97.9	100.0	99.6	**109.1**	**110.2**

≤**59.99** **60 - 69.99** **70 -79.99** **80 - 89.99** **>100.5**

Comparison of Performance

Exact Probabilistic

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.10	56.5	50.5	46.7	49.6	54.2	49.1
	0.15	71.7	65.5	63.8	64.0	65.7	63.4
	0.20	82.5	75.5	72.5	74.2	73.5	76.2
C	**0.25**	88.8	83.1	82.4	81.3	80.7	80.4
	0.30	94.4	88.9	87.2	86.5	89.1	87.6
	0.35	93.6	95.8	95.2	91.7	94.3	96.5
	0.40	97.0	97.9	94.7	99.3	97.6	96.7
	0.45	99.2	99.2	97.3	99.2	96.9	97.1

Error Interval 0.2

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.1	62.0	52.3	52.4	57.4	58.1	52.7
	0.15	81.1	70.8	67.5	70.1	71.1	68.2
	0.2	90.5	77.2	75.0	83.5	78.5	76.1
C	**0.25**	94.0	83.8	84.9	89.6	86.4	82.6
	0.3	**102.0**	97.9	88.0	88.5	98.1	93.3
	0.35	**103.9**	92.2	97.7	99.3	95.2	94.0
	0.4	**101.7**	93.7	96.3	93.9	96.5	97.4
	0.45	**101.8**	**109.4**	100.1	**101.4**	**109.5**	**110.1**

≤**59.99** **60 - 69.99** **70 -79.99** **80 - 89.99** **>100.5**

Comparison of Performance

Exact Probabilistic

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.10	56.5	50.5	46.7	49.6	54.2	49.1
	0.15	71.7	65.5	63.8	64.0	65.7	63.4
	0.20	82.5	75.5	72.5	74.2	73.5	76.2
C	**0.25**	88.8	83.1	82.4	81.3	80.7	80.4
	0.30	94.4	88.9	87.2	86.5	89.1	87.6
	0.35	93.6	95.8	95.2	91.7	94.3	96.5
	0.40	97.0	97.9	94.7	99.3	97.6	96.7
	0.45	99.2	99.2	97.3	99.2	96.9	97.1

Error Interval 0.25

		N					
		25	**50**	**75**	**100**	**150**	**200**
	0.1	65.5	58.5	52.5	54.0	62.6	53.3
	0.15	88.9	77.5	68.5	68.7	75.8	69.9
	0.2	98.9	85.1	75.1	83.5	85.1	78.2
C	**0.25**	**100.9**	89.4	89.1	89.2	93.2	83.5
	0.3	**108.5**	97.6	92.2	88.5	**101.3**	94.1
	0.35	**107.1**	100.1	98.2	99.7	94.4	93.2
	0.4	**109.3**	94.3	96.6	94.4	96.1	97.2
	0.45	**107.3**	**108.6**	99.7	100.1	**109.6**	**109.6**

≤**59.99**   **60 - 69.99**   **70 -79.99**   **80 - 89.99**   **>100.5**

Conclusion

- APA performs very similar to "exact" probabilistic method for error interval length 0.1 and N > 25
    - At least 40% for C = 0.1
    - At least 30% for C = 0.15
    - At least 20% for C = 0.2
    - At least 10% for C = 0.25
- For error interval lengths 0.15 and 0.2, APA achieves comparable efficiency brackets
- APA universally applicable, independent of assumptions on waiting times between events

Future Work

- Simulation
  - uniform selection of data points for all error intervals
  - other distributions
- Theoretical results
  - number of data points necessary as function of error interval length
- Author Info:

  - sheubac@calstatela.edu
  - rpamula@calstatela.edu
  - www.calstatela.edu/faculty/sheubac

Author Info:

- Email:
  - sheubac@calstatela.edu
  - rpamula@calstatela.edu

- Paper and Presentation
  - www.calstatela.edu/faculty/sheubac