HotEmNets'08

Energy-Reliability Tradeoffs in Sensor Network Storage

Neerja Bhatnagar
Kevin Greenan
Rosie Wacha
Prof. Ethan L. Miller
Prof. Darrell D. E. Long

June 3, 2008

Storage Systems Research Center
University of California, Santa Cruz

Outline

- Data reliability in sensor networks
- Analyzing energy-reliability tradeoffs in context of
  - Choice of redundancy techniques
  - Choice of nodes
  - Frequency of integrity checks
- Optimizations
- Future experimental approach

Background: Storing Sensed Observations

- Storage operations are more energy-efficient than radio operations [Mathur '06 and Lin '06]
- Gigabyte storage is available on sensor nodes [Mitra '05]
- Storage cost is decreasing as storage gets denser [Ganesan '05]
- Want to reliably store data in local sensor network without base station

Local Sensor Network Storage

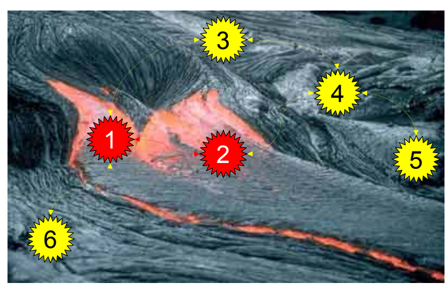

- Sensor nodes suffer from:
  - Individual failures
  - Correlated failures

Energy Tradeoffs for Reliability

- Redundancy techniques
  - Mirroring vs. erasure coding
- Choice of nodes for replication
  - Far vs. near nodes
- Frequency of remote storage verification
  - Very frequent, infrequent, or piggy-back on other traffic

Sensor Network Assumptions

- Sensor nodes limited in CPU, power, and storage
- Battery-backed RAM and NAND flash at each node
- Transmission to distant nodes consumes more energy
- Transmission cost includes retransmission cost

Redundancy Techniques

- Trade-off between processing cost and reliability
- Reliability depends on technique:
  - Irregular XOR codes tolerate **most** `j` failure sets
  - Reed-Solomon (RS) tolerates **any** `j` failures
- Details of RS and XOR
  - `(i,j)` encodes `i` data with `j` parity nodes
  - Both techniques encode same size of data
  - All encoding done by primary node
  - Data and parity chunks are distributed

RS and XOR Performance

Code Size	**Energy Expenditure (mJ)**		**Throughput (MB/s)**	
	RS	**XOR**	**RS**	**XOR**
(5, 3)	3.515	1.205	2.674	7.798
(6, 2)	3.133	0.6	3	15.654
(9, 3)	4.82	0.524	1.95	17.953
(10, 2)	3.92	0.653	2.4	14.4
(17, 3)	5.193	0.588	1.81	15.99
(18, 2)	4.36	0.589	2.156	15.972

- Experiments run on ARM9E 400MHz processor that consumes 94 mJ/s
- RS consumes 3-10 times more energy

Analyzing Reliability of Redundancy Techniques

- Mirror$_4$ – backup mirrors on 4 other nodes

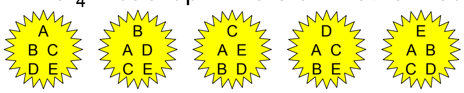

- XOR$_2$ – store two backups as XOR

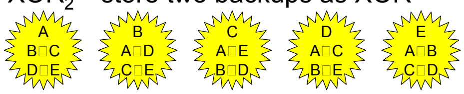

- XOR$_1$ – store one backup as XOR

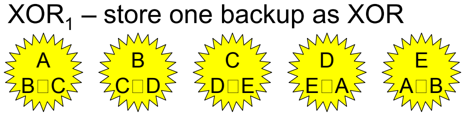

Availability Markov Model

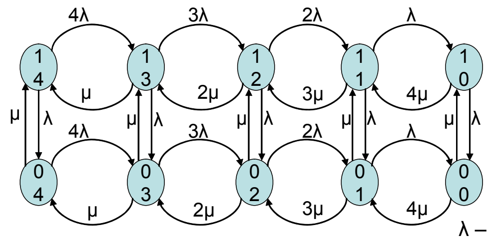

λ – failure rate
μ – repair rate

- Exponential failure and recovery rates
- Failed nodes replaced before data loss

Data Availability

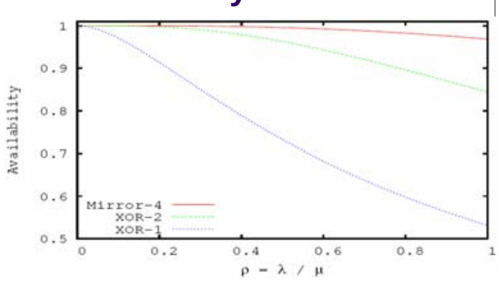

	Storage Overhead	# Failures Tolerated	Energy (mJ)	Throughput (MB/s)
Mirror$_4$	4x	4	---	---
XOR$_2$	2x	2	0.75	12.72
XOR$_1$	1x	1	0.74	12.76

Mean Time To Data Loss

	Mirror$_4$	XOR$_2$	XOR$_1$
MTTDL with repair	4.87 x 10^{11} hours	6.50 x 10^{8} hours	2.42 x 10^{6} hours
MTTDL without repair	4932 hours	2272 hours	1692 hours

- We model MTTDL using the same transition matrix that would be used modeling with differential equations.
- Assumptions
  - 5-node redundancy groups
  - $\rho = 5.56 \times 10^{-3}$: failures every 3 months and repairs take 12 hours
  - Failures and repairs exponentially distributed

Node Choice for Redundancy

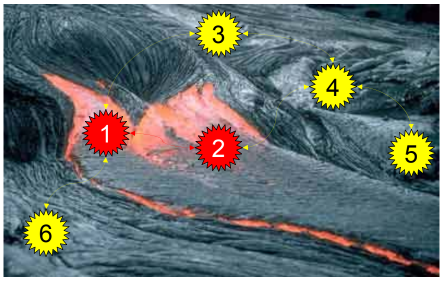

- Lava flow destroys nodes 1 and 2
- Correlated failure event can destroy data

Node Choice for Redundancy

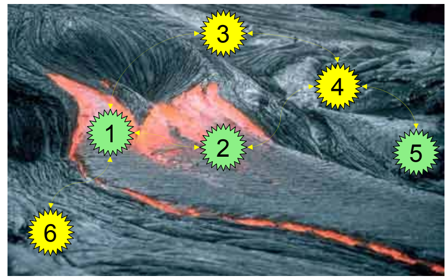

- Lava flow destroys nodes 1 and 2
- Correlated failure event can destroy data
- Choose redundancy groups with some distant nodes

Frequency of Integrity Checks

- Each node must periodically verify its backups on remote nodes
- If such integrity checks are not conducted, then overall reliability reduces
- Use algebraic signatures to detect changes in backups
- Tradeoff: frequent verifications improves reliability but consumes more energy

Optimizations

- Piggy-back integrity check messages on to other traffic
- Transmit data to a remote node and let it distribute it to its closer neighbors
- Use some intermediate nodes for redundancy

Future Experimental Approach

- Use simulation to measure total energy expenditure for reliable sensor network storage
  - Measure energy expended at originating node and each back-up node
  - Determine network protocol for establishing nodes for reliability groups
  - Data transmission costs

Summary

- Mirror to nearby nodes to guard against individual node failures
- Use erasure coding for distant nodes to guard against correlated failures
- Use XOR-based codes instead of RS codes
- Store algebraic signatures for data integrity
- Use simulation to evaluate total energy

Thank You!