

Data Discovery and Dissemination with DIP

Authors: Kaisen Lin, Philip Levis

Speaker: Giannakakis Spyridon

10-937-548
ETH Zurich, D-INFK

March 22, 2011

Problem

How to efficiently distribute binaries or data to remote nodes.

Problem

How to efficiently distribute binaries or data to remote nodes.

Why?

- Adjust configuration parameters.
- Reprogram nodes for extra functionality.
- Apply patches.

Limitations of wireless nodes: energy, processing power etc.

Trickle

The algorithm behind most protocols.

Operation

- Assigns keys and version numbers to data items.
- Broadcasts periodically summary of the node's data.

Trickle

The algorithm behind most protocols.

Operation

- Assigns keys and version numbers to data items.
- Broadcasts periodically summary of the node's data.
 - If neighbours agree on version number, increase the broadcast interval.
 - Else, broadcast rapidly.

Extra functionality

- Suppress message.
- Nodes can update each other.

Trickle

The algorithm behind most protocols.

Operation

- Assigns keys and version numbers to data items.
- Broadcasts periodically summary of the node's data.
 - If neighbours agree on version number, increase the broadcast interval.
 - Else, broadcast rapidly.

Extra functionality

- Suppress message.
- Nodes can update each other.

Big intervals when data consistent, rapid advertisement in case of differences.

Limitations

- Scales to many nodes but not for many data items per node.
- Detection of new data is problematic.
- Transmission costs scale linearly $O(T)$.¹

Result:

Broadcast intervals remain big, late detection of asymmetries in the network.

¹T number of data items

Performance metrics

- Latency.
- Maintenance cost.
- Identify cost.

Related to interval used.

Traditional Trickle: dynamic adjustment of interval propagates fast updates...

Performance metrics

- Latency.
- Maintenance cost.
- Identify cost.

Related to interval used.

Traditional Trickle: dynamic adjustment of interval propagates fast updates...

but does not help with detection itself!

Approaches using Trickle

① Parallel Scan

Separate Trickle for each data item. Fixed interval.

② Serial Scan

One Trickle (transmission) containing summary over a selection of data items.

		Detection		Identify	2
		Latency	Cost		
Parallel Scan		$O(1)$	$O(T)$	$O(1)$	
Serial Scan		$O(T)$	$O(1)$	$O(1)$	

²adapted from [DIP, 2008]

Cost/latency Tradeoff

Cost/latency product scales with $O(T)$.³
Administrators have to choose between speed and efficiency.

Parallel scan detects faster but costs more,
serial scan is cheaper but takes longer.

³T data items


Search Protocol

Similar to a binary search.

Main idea:

Send a hash of version numbers across H data items.

When a node hears a different hash, it sends hashes of sub-range within that hash.



Performance

Detection

When stable, top-level hash is sent \Rightarrow detects all differences.

Identification

Need to traverse the whole tree.

Protocol	Latency	Cost	Identify
Parallel Scan	$O(1)$	$O(T)$	$O(1)$
Serial Scan	$O(T)$	$O(1)$	$O(1)$
Search	$O(1)$	$O(1)$	$O(\log(T))$ ⁴

⁴[DIP, 2008]

Performance

Detection

When stable, top-level hash is sent \Rightarrow detects all differences.

Identification

Need to traverse the whole tree.

Protocol	Latency	Cost	Identify
Parallel Scan	$O(1)$	$O(T)$	$O(1)$
Serial Scan	$O(T)$	$O(1)$	$O(1)$
Search	$O(1)$	$O(1)$	$O(\log(T))$ ⁴

Problem

Identifying multiple new items simultaneously.

⁴[DIP, 2008]

DISSEMINATION PROTOCOL

Hybrid protocol

- Starts with *search* to detect a difference.
- When hashes become small enough (threshold), drop to *scan* to identify the item.

⁵[DIP, 2008]


DISSEMINATION PROTOCOL

Hybrid protocol

- Starts with *search* to detect a difference.
- When hashes become small enough (threshold), drop to *scan* to identify the item.

Goal:

Detect and identify which nodes need which updates by increasing Estimate values.



⁵[DIP, 2008]

Data stored

For each data item:

- Unique key
- Version Number
- Index (memory)
- Estimates: estimate of E means that DIP detected a difference in level E.

Estimates

$E = 0$ constant data

$E_D = \log(T)$ certainty of difference

E_O neighbour contains **o**lder information

E_N neighbour contains **n**ewer information

As we traverse down the tree \Rightarrow E grows.

Data stored

For each data item:

- Unique key
- Version Number
- Index (memory)
- Estimates: estimate of E means that DIP detected a difference in level E .

Estimates

$E = 0$ constant data

$E_D = \log(T)$ certainty of difference

E_O neighbour contains **o**lder information


E_N neighbour contains **n**ewer information

As we traverse down the tree $\Rightarrow E$ grows.

Overhead: Trickle uses 4 bytes, DIP extra 1 byte.


Messages

1 Data



Sent after detecting *and* identifying the difference, updates the node.

2 Vector



...

Provides detection of differences,
can identify the different item, can't update the node.
Scan part of DIP.

- ③ Summary:
Summary elements (branching factor), salt value

Salt s			
Begin ₁	End ₁	BloomFilter ₁	SummaryHash ₁
...			
Begin _{n}	End _{n}	BloomFilter _{n}	SummaryHash _{n}

6

Provides detection with the hash (*search part of DIP*).
Salt avoids hash collisions.
Bloom filter helps to identify with probability.

⁶[DIP, 2008]

Bloom filters

”Swiss army knife”

Usage

Test membership of element e .

- Initialization: Set k bits $H(e) = 1$.⁷
- Query: Check if $H(e) = 1$
 - If no, e is not in the set
 - If yes, *might* be.

Purpose


Bypass thorough traversing of tree, sets directly corresponding data items to E_D .

⁷In DIP $k=1$

Upon receiving

Summary message

- Matching hash: Data is the same, Estimate = 0.




⁸[DIP, 2008]

Upon receiving

Summary message

- Matching hash: Data is the same, Estimate = 0.
- Differing hash: $E = \max(\text{currentEstimate}, \log(T) - \log(H))$




⁸[DIP, 2008]

Upon receiving

Summary message

- Matching hash: Data is the same, Estimate = 0.
- Differing hash: $E = \max(\text{currentEstimate}, \log(T) - \log(H))$
- Differing hash and bloom filter: Difference detected, E_D .



⁸[DIP, 2008]

Receiving Data or Vector messages

- Same version number: Data is the same, Estimate = 0
- Older version: Estimate = E_O unless E_N

Receiving Data or Vector messages

- Same version number: Data is the same, Estimate = 0
- Older version: Estimate = E_O unless E_N

Vector with a newer version E_N

Data with newer version install update, set E_O .

Transmitting

Decrease Estimate after transmission of a message.

Transmission

Packets with highest Estimate first.

A packet represents a single level in the tree (estimate value).

Algorithm

- if Estimate is E_O , send *data* message⁹.
- else:
 - if "cost of using summaries"¹⁰ > "cost of using vectors", send summary
 - else send vectors.

⁹forwarding has higher priority

¹⁰ $(E_D - E)$: levels until the end of the hash tree

Transmission

Packets with highest Estimate first.

A packet represents a single level in the tree (estimate value).

Algorithm

- if Estimate is E_O , send *data* message⁹.
- else:
 - if "cost of using summaries"¹⁰ > "cost of using vectors", send summary
 - else send vectors.

Basically, the scan threshold

if Estimate equals E_D or E_N , transmit a vector.

⁹forwarding has higher priority

¹⁰ $(E_D - E)$: levels until the end of the hash tree

Example¹¹

Node A

Sends *summary*,

each element of $H=8$ ↘

Node B

¹¹Assumptions: Initially both in steady state, 16 Data items, without bloom filters, vectors are sent when E_D (vectors contain one tuple), summaries contain 2 elements (branching factor of 2).

Example¹¹

Node A

Sends *summary*,
each element of $H=8$ ↘

Node B

Sets $E=1$,
↙ sends *summary*, elements $H=4$

¹¹Assumptions: Initially both in steady state, 16 Data items, without bloom filters, vectors are sent when E_D (vectors contain one tuple), summaries contain 2 elements (branching factor of 2).

Example¹¹

Node A

Sends *summary*,
each element of $H=8$ ↘

Sets $E=2$,
sends *summary* of $H=2$ ↘

Node B

Sets $E=1$,
↙ sends *summary*, elements $H=4$

¹¹Assumptions: Initially both in steady state, 16 Data items, without bloom filters, vectors are sent when E_D (vectors contain one tuple), summaries contain 2 elements (branching factor of 2).

Example¹¹

Node A

Sends *summary*,
each element of $H=8$ ↘

Sets $E=2$,
sends *summary* of $H=2$ ↘

Node B

Sets $E=1$,
↙ sends *summary*, elements $H=4$

Sets $E=3$,
↙ sends *vector* of 1 item

¹¹Assumptions: Initially both in steady state, 16 Data items, without bloom filters, vectors are sent when E_D (vectors contain one tuple), summaries contain 2 elements (branching factor of 2).

Example¹¹

Node A

Sends *summary*,
each element of $H=8$ ↘

Sets $E=2$,
sends *summary* of $H=2$ ↘

Sets $E=4(E_O, N)$,
sends *data* →

Node B

Sets $E=1$,
↙ sends *summary*, elements $H=4$

Sets $E=3$,
↙ sends *vector* of 1 item

¹¹Assumptions: Initially both in steady state, 16 Data items, without bloom filters, vectors are sent when E_D (vectors contain one tuple), summaries contain 2 elements (branching factor of 2).

Overview

Breaks the latency/cost tradeoff $O(T)$.

Overview

Breaks the latency/cost tradeoff $O(T)$.

Link layer broadcasts

Trickle timer

Hierarchical suppression

Based on local information

Overview

Breaks the latency/cost tradeoff $O(T)$.

Link layer broadcasts

Trickle timer

Hierarchical suppression

Based on local information

Contribution

- Bloom filter
- Scan after search

Evaluation

Methodology

- 1 Improvements
DIP stand alone protocol.

Evaluation

Methodology



- ① Improvements
DIP stand alone protocol.
- ② Comparisons
With a serial scan and a pure search protocol.
Simulation and testbed deployment.

DIP Implementation

TinyOS 2.0, 3K code

2 elements/summary

2 tuples/vector

(a) $L = 0$.(b) $L = 20$.

12

Figure: $T = 256$ total items, $D = 32$ nodes, L losses

Simulations





Figure: Default: $N=8$, $T=64$, $D=32$, $L=40\%$ ¹³

¹³[DIP, 2008], N new items, T total items, D nodes, L losses

Simulation and Experiment




Figure: Tx messages, simulation on TOSSIM, 15 by 15 grid, $T = 64$

Simulation and Experiment




Figure: Tx messages, simulation on TOSSIM, 15 by 15 grid, $T = 64$




Figure: Tx costs, experiment on Mirage, $T = 64$

References



Lin, Kaisen and Levis, Philip.

Data Discovery and Dissemination with DIP.

IPSN '08.



http:

`//en.wikipedia.org/wiki/Hash_tree,Bloom_filter`