

Fast Scheduling in Distributed Transactional Memory

Costas Busch Maurice Herlihy Miroslav Popovic
Gokarna Sharma

Presented by
Ramesh Adhikari
Graduate Research Assistant
School of Computer and Cyber Sciences, Augusta University

December, 2022

Table of Contents

- 1 Introduction
- 2 Graph Model
- 3 Greedy Schedule
 - Line Graph
 - Grid Graph
 - Cluster Graph
 - Star Graph
- 4 Conclusion

Table of Contents

- 1 Introduction
- 2 Graph Model
- 3 Greedy Schedule
 - Line Graph
 - Grid Graph
 - Cluster Graph
 - Star Graph
- 4 Conclusion

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions
 - Vulnerability to Failure or Delay

Introduction

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions
 - Vulnerability to Failure or Delay
- Solution (Transactional memory)

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions
 - Vulnerability to Failure or Delay
- Solution (Transactional memory)
 - Using TM, code is split into transactions

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions
 - Vulnerability to Failure or Delay
- Solution (Transactional memory)
 - Using TM, code is split into transactions
 - Sequence of Reads and Writes on shared variable execute atomically

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions
 - Vulnerability to Failure or Delay
- Solution (Transactional memory)
 - Using TM, code is split into transactions
 - Sequence of Reads and Writes on shared variable execute atomically
 - **Commit transaction**: In the absence of conflicts or failures

- Concurrent processes (threads) need to **synchronize** to avoid introducing inconsistencies in shared data objects.
- Traditional synchronization mechanisms such as **locks** and barriers have well-known **limitations**.
 - Deadlock
 - Priority inversion
 - Reliance on programmer conventions
 - Vulnerability to Failure or Delay
- Solution (Transactional memory)
 - Using TM, code is split into transactions
 - Sequence of Reads and Writes on shared variable execute atomically
 - **Commit transaction**: In the absence of conflicts or failures
 - **abort transaction**: In case of synchronization conflicts or failures

Introduction

- Paper considers a data-flow of transaction execution

Introduction

- Paper considers a data-flow of transaction execution
- Each transaction executes at a single node, but data objects are mobile.

Introduction

- Paper considers a data-flow of transaction execution
- Each transaction executes at a single node, but data objects are mobile.
- A transaction initially requests the objects it needs and executes only after it has assembled them

Introduction

- Paper considers a data-flow of transaction execution
- Each transaction executes at a single node, but data objects are mobile.
- A transaction initially requests the objects it needs and executes only after it has assembled them
- After the transaction commits, it releases its objects, possibly forwarding them to other waiting transactions

Introduction

- Paper considers a data-flow of transaction execution
- Each transaction executes at a single node, but data objects are mobile.
- A transaction initially requests the objects it needs and executes only after it has assembled them
- After the transaction commits, it releases its objects, possibly forwarding them to other waiting transactions
- In a distributed TM, execution time is dominated by the costs of moving objects from one transaction to another. The goal of a transaction scheduling algorithm (sometimes called contention management) is to minimize delays caused by data conflicts and data movement.

Table of Contents

- 1 Introduction
- 2 Graph Model**
- 3 Greedy Schedule
 - Line Graph
 - Grid Graph
 - Cluster Graph
 - Star Graph
- 4 Conclusion

- The network is modeled as a weighted graph G , where transactions reside at nodes

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.
- At any time step, a node may perform three actions:

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.
- At any time step, a node may perform three actions:
 - it may receive objects from adjacent nodes

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.
- At any time step, a node may perform three actions:
 - it may receive objects from adjacent nodes
 - it executes any transaction that has assembled its required objects

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.
- At any time step, a node may perform three actions:
 - it may receive objects from adjacent nodes
 - it executes any transaction that has assembled its required objects
 - it may forward objects to adjacent nodes.

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.
- At any time step, a node may perform three actions:
 - it may receive objects from adjacent nodes
 - it executes any transaction that has assembled its required objects
 - it may forward objects to adjacent nodes.
- A Transaction's execution terminates when it commits.

Graph Model

- The network is modeled as a weighted graph G , where transactions reside at nodes
- Edges are communication links, and edge weights are communication delays.
- At any time step, a node may perform three actions:
 - it may receive objects from adjacent nodes
 - it executes any transaction that has assembled its required objects
 - it may forward objects to adjacent nodes.
- A Transaction's execution terminates when it commits.
- Paper provided offline algorithms to compute conflict-free schedules.

Table of Contents

- 1 Introduction
- 2 Graph Model
- 3 Greedy Schedule**
 - Line Graph
 - Grid Graph
 - Cluster Graph
 - Star Graph
- 4 Conclusion

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored
- Each color is a separate time step

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored
- Each color is a separate time step
- Apply this greedy schedule to

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored
- Each color is a separate time step
- Apply this greedy schedule to
 - Line Graph

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored
- Each color is a separate time step
- Apply this greedy schedule to
 - Line Graph
 - Grid Graph

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored
- Each color is a separate time step
- Apply this greedy schedule to
 - Line Graph
 - Grid Graph
 - Cluster Graph

Greedy Schedule

- Use dependency graph
- Pick the transaction and assign valid color i.e. execution time
- Repeat until all transactions(nodes) are colored
- Each color is a separate time step
- Apply this greedy schedule to
 - Line Graph
 - Grid Graph
 - Cluster Graph
 - Star Graph

Line Graph




Fig. 1: A line graph with $n = 32$ nodes and $\ell = 8$.

- Each node has a transaction of any node of objects
- L be the longest shortest walk of any object
- Authors schedule the transactions in two phases
- **first phase** execute the transactions in $S1$ And in the **second phase** execute the transactions in $S2$.
- Total time: $\mathcal{O}(L)$

Grid Graph




Fig. 2: A grid of size 16×16 with subgrids of size 4×4 . It depicts the path of an object.

- Each transaction requests a uniformly random set of k objects (out of w)
- $n \times n$ nodes and w objects and decompose graph into sub-grids
- Execute each sub-grid one after another
- In every sub-grid apply **greedy schedule**

Cluster Graph




Fig. 3: A graph with 5 clusters where each cluster C_i is a complete graph with 6 nodes; links within clusters have weight 1, while links between clusters have weight γ .

- Each transaction requests k arbitrary objects (out of w)

Star Graph




Fig. 4: A star graph with 8 rays, each ray consisting of 7 nodes; the rings depict the set of segments V_1, V_2, V_3 .

- Each transaction requests k arbitrary objects (out of w)
- Divide star into rings
- Treat each ring as a cluster graph
- Apply the line algorithm for each ray cluster

Table of Contents

- 1 Introduction
- 2 Graph Model
- 3 Greedy Schedule
 - Line Graph
 - Grid Graph
 - Cluster Graph
 - Star Graph
- 4 Conclusion

Conclusion

- Paper discussed the data-flow model in Transactional Memory
- Paper presented efficient schedules for many network typologies

Thank You!