

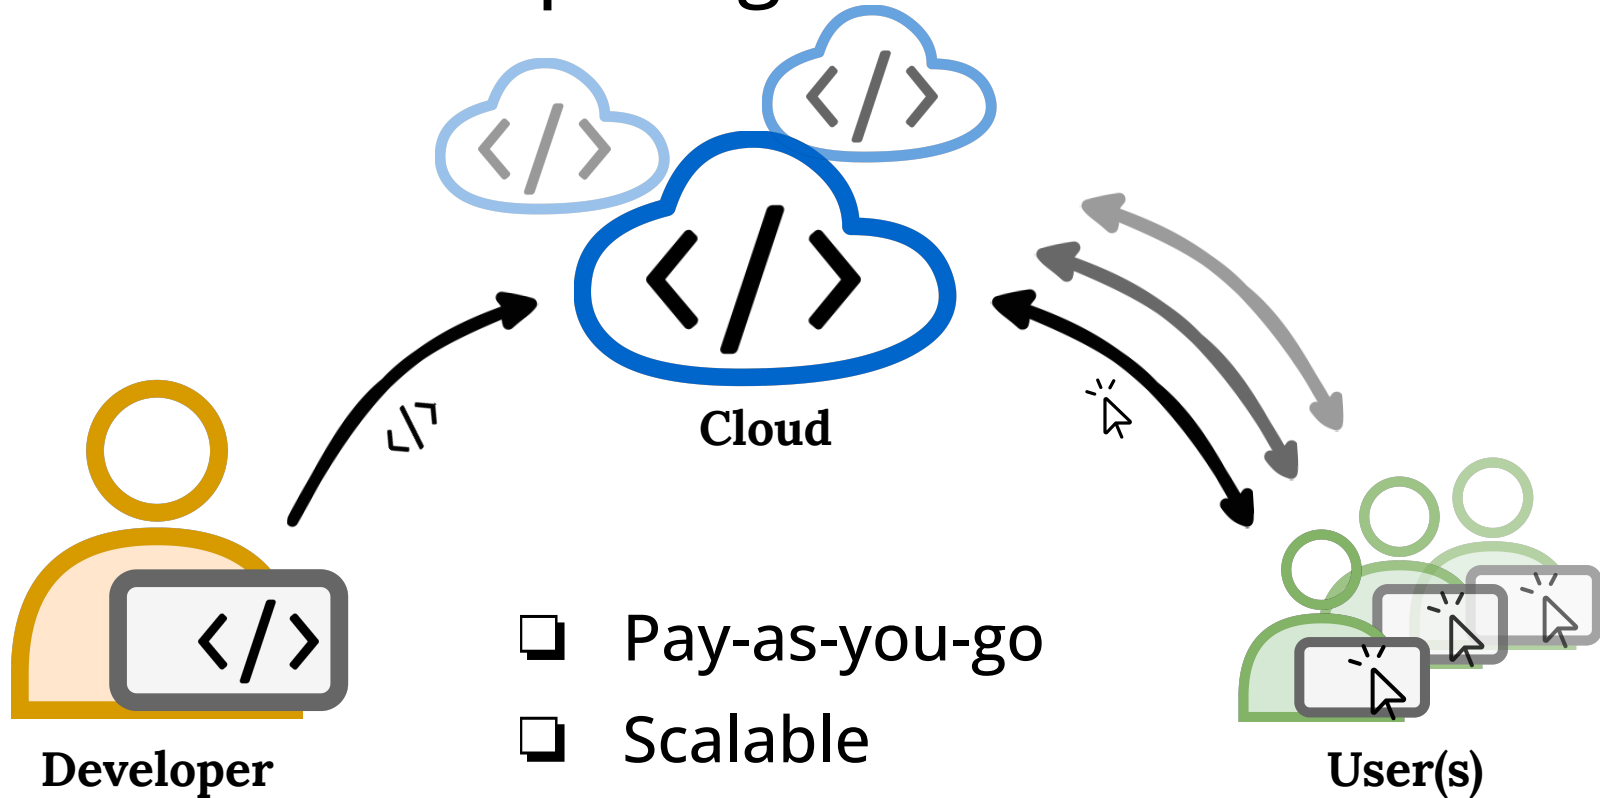
# Cosmic: Cost-Effective Support for Cloud-Assisted 3D Printing

Yuan Yao<sup>1</sup>, Chuan He<sup>2</sup>, Chinedum Okwudire<sup>2</sup>, Harsha V. Madhyastha<sup>1</sup>

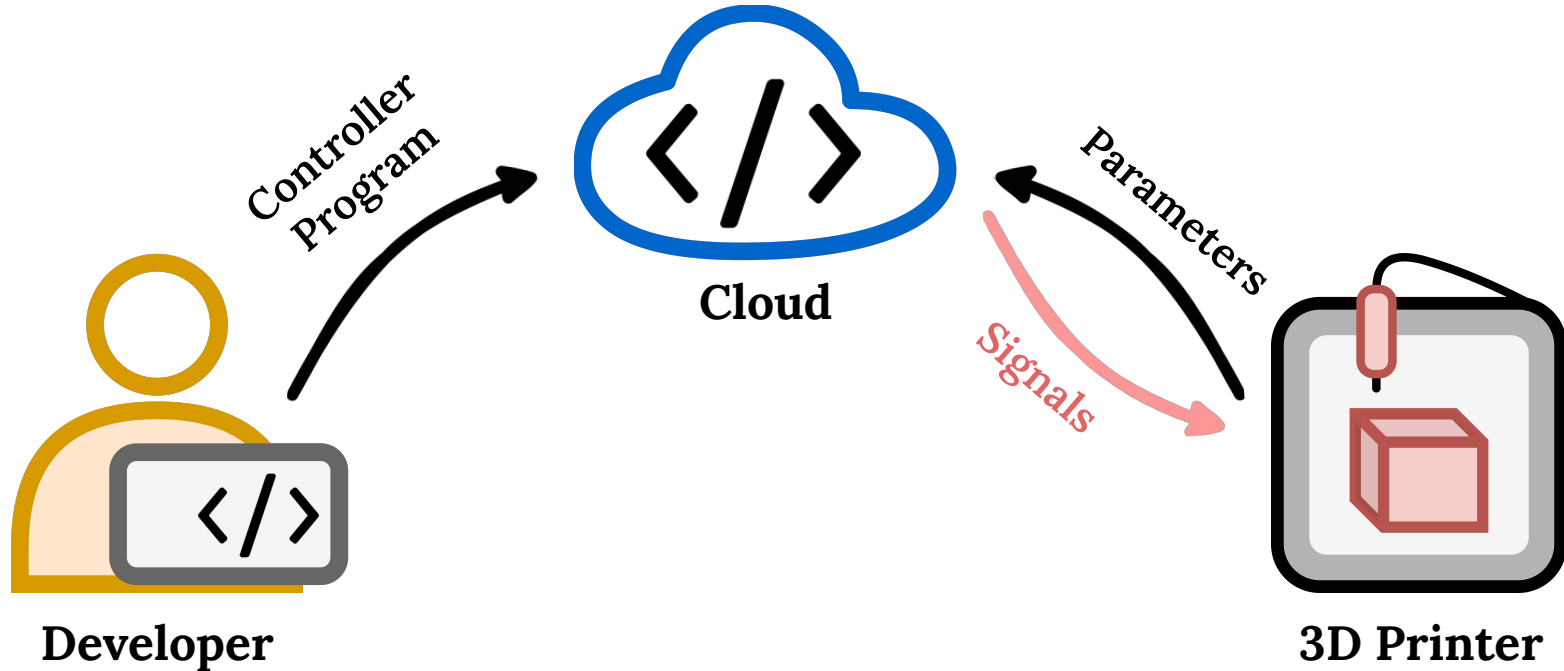
<sup>1</sup>  USC University of  
Southern California

<sup>2</sup>  UNIVERSITY OF MICHIGAN

# Serverless Computing




# Serverless Assisted 3D Printing Control




# An Example: A Metal 3D Printing Process

## Powder Bed Fusion (PBF)


- 1 Spread a thin layer of metal powder
- 2 Selectively melt powder with a laser
- 3 Repeat layer by layer until finish



# ... Unoptimized scanning patterns produce uneven heat distribution



Divided into *Cells*




Successive  
ordering




**Heat-induced  
deformation**


# ... *SmartScan* generates optimized scanning patterns for PBF



Divided into *Cells*

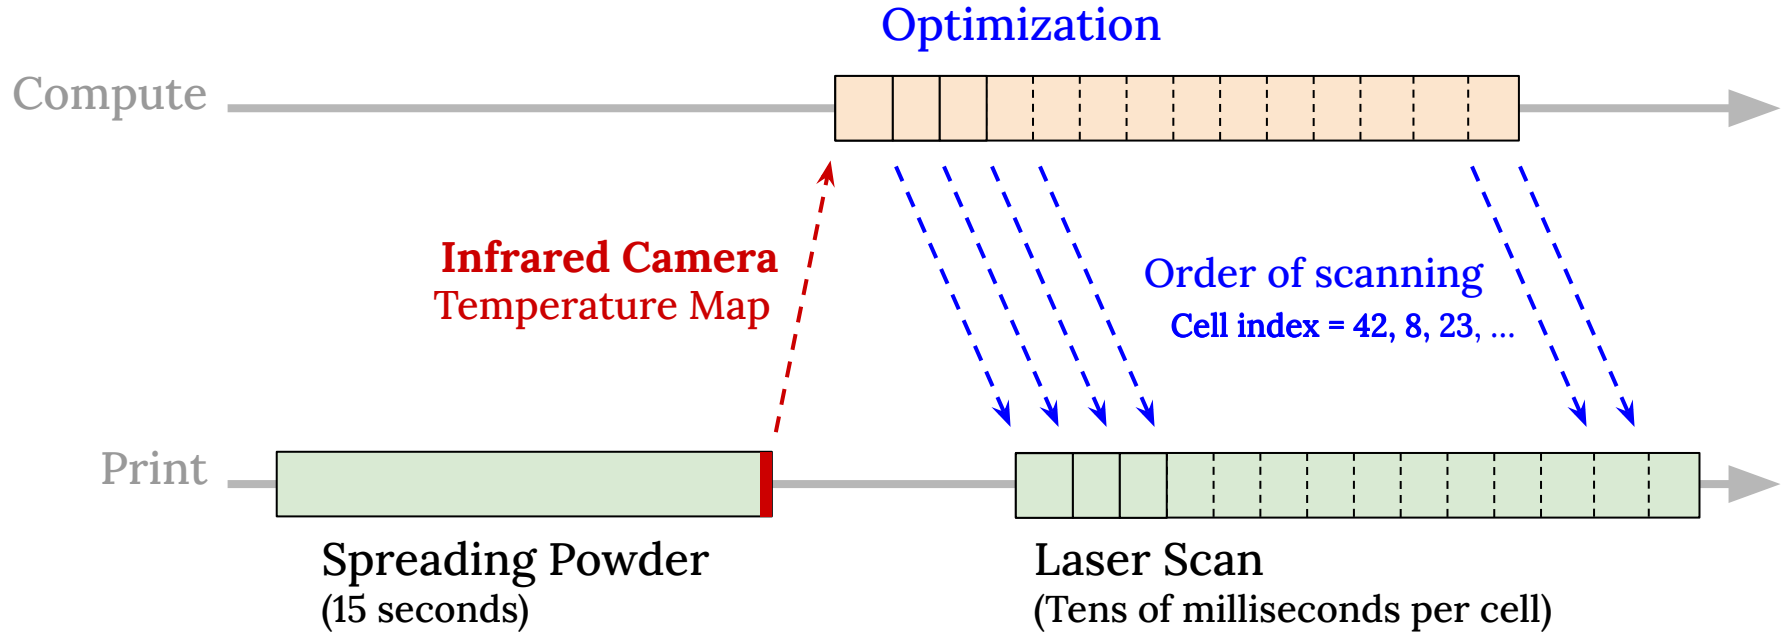


*SmartScan*  
ordering

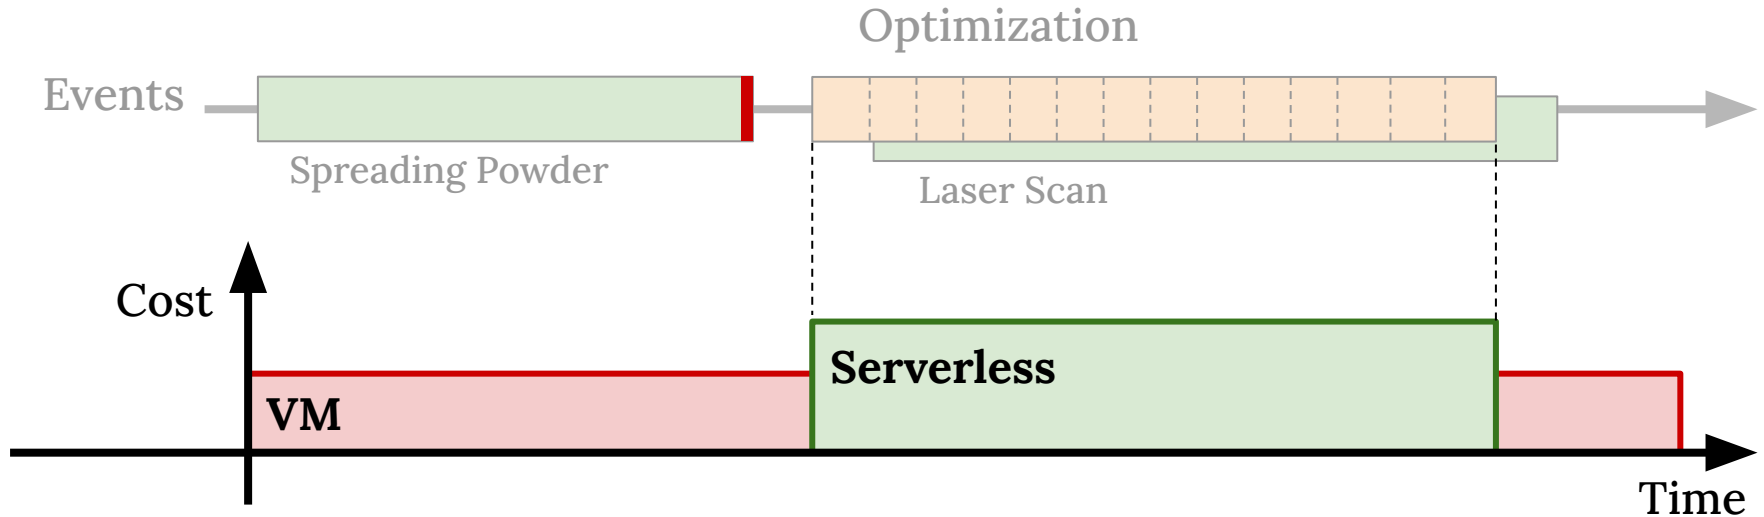


**Reduced  
deformation**

# An Example: *SmartScan's* Timeline (Within a layer)




# Advantages of Relying on Serverless (1)




**VM: Bills total uptime including idle time**

**Serverless: Charges only during actual execution**

# Advantages of Relying on Serverless (2)




**VM: Cost scales with memory, while cores underutilized**



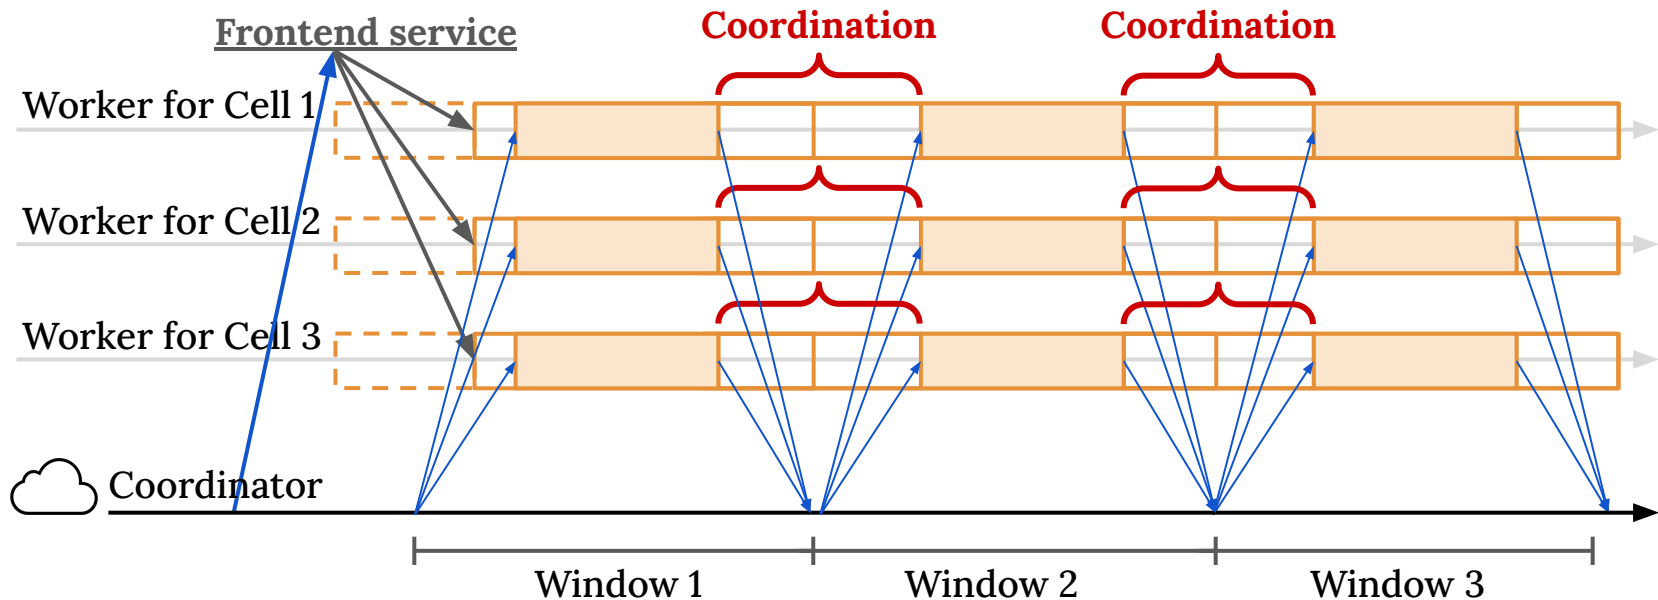
**Serverless: Data loaded in memory remains available across quick reuses, enabled by warm starts**

# Challenge 1: Invoking functions on demand is slow



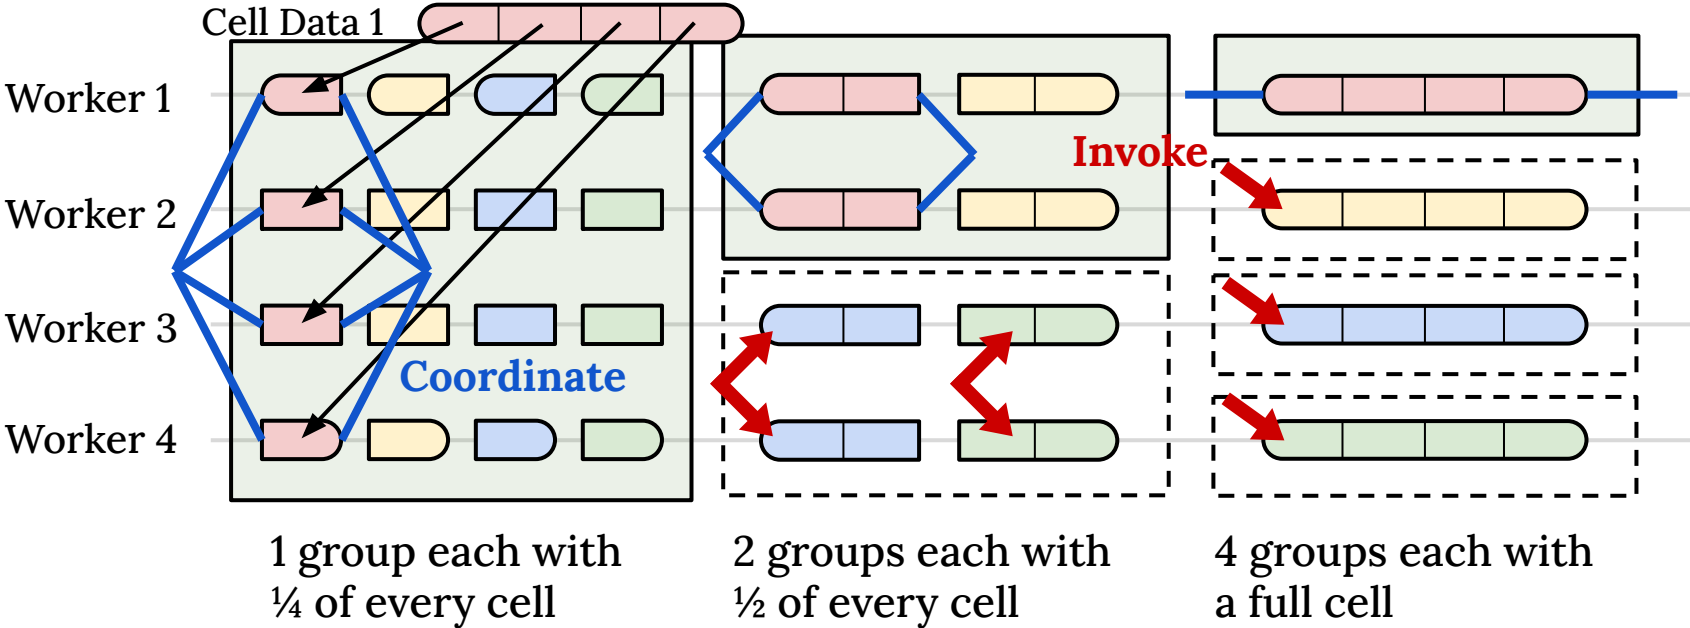
**Even with warm starts, invocation delay adds >10 ms of overhead, often causing timing violations**

# Challenge 2: Coordination can be slow and costly



**At high parallelism, coordination latency can violate timing constraints and increase cost**

# Partition Computing and Resources into Groups

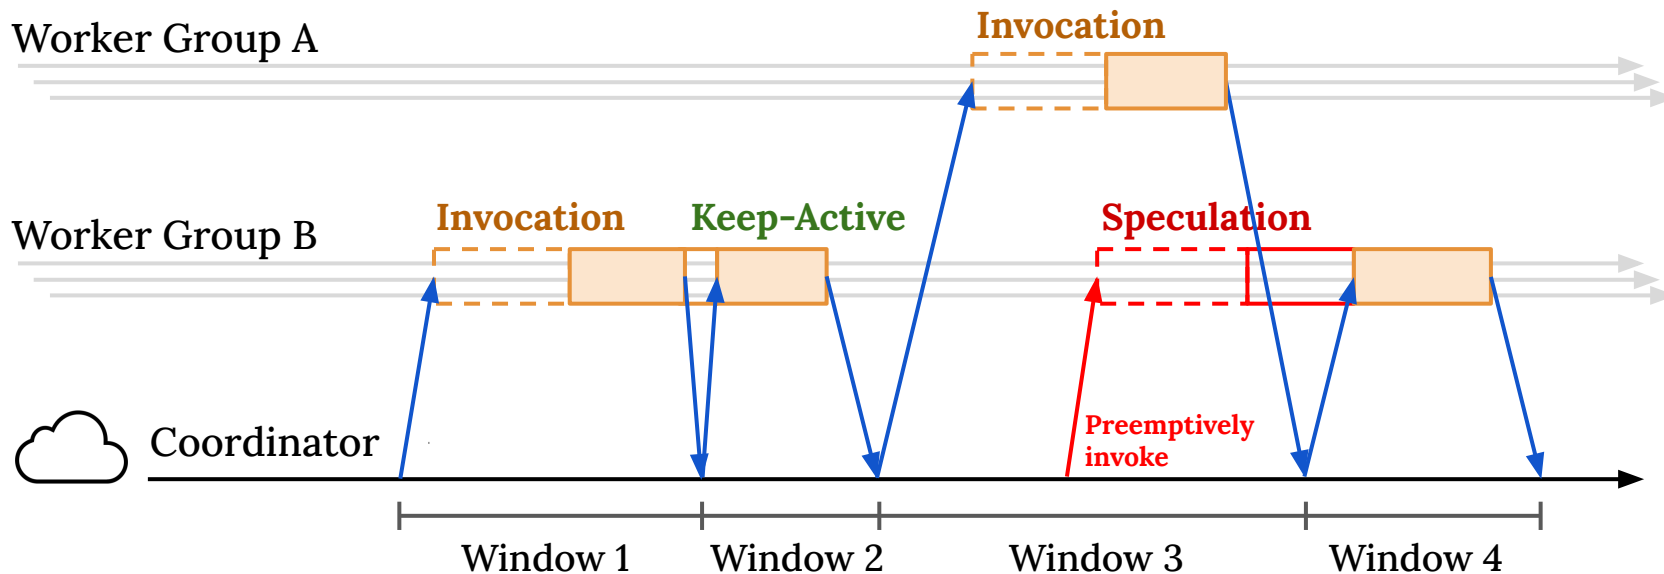


← More Coordination

More Invocation →

# Cosmic in Action

- Cell 1, in Group B
- Cell 2, in Group B again
- Cell 3, in another Group A
- Cell 4, in Group B again – correct guess



# Model Performance and Select The Best Configuration



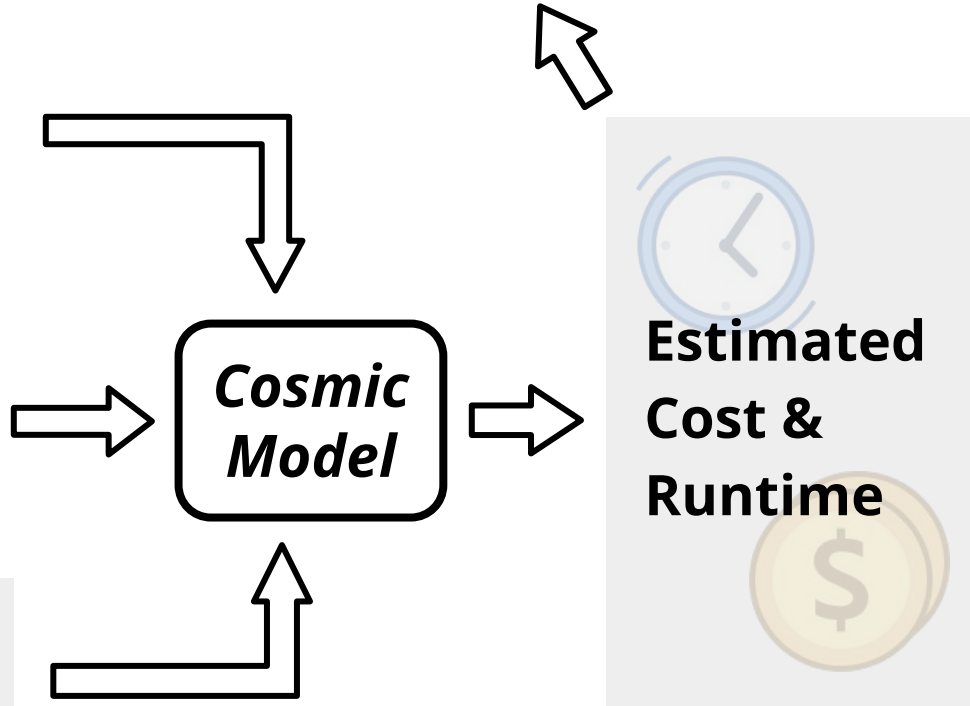
Characteristics of the printing job



Measurements of the serverless platform




Configurations to be evaluated

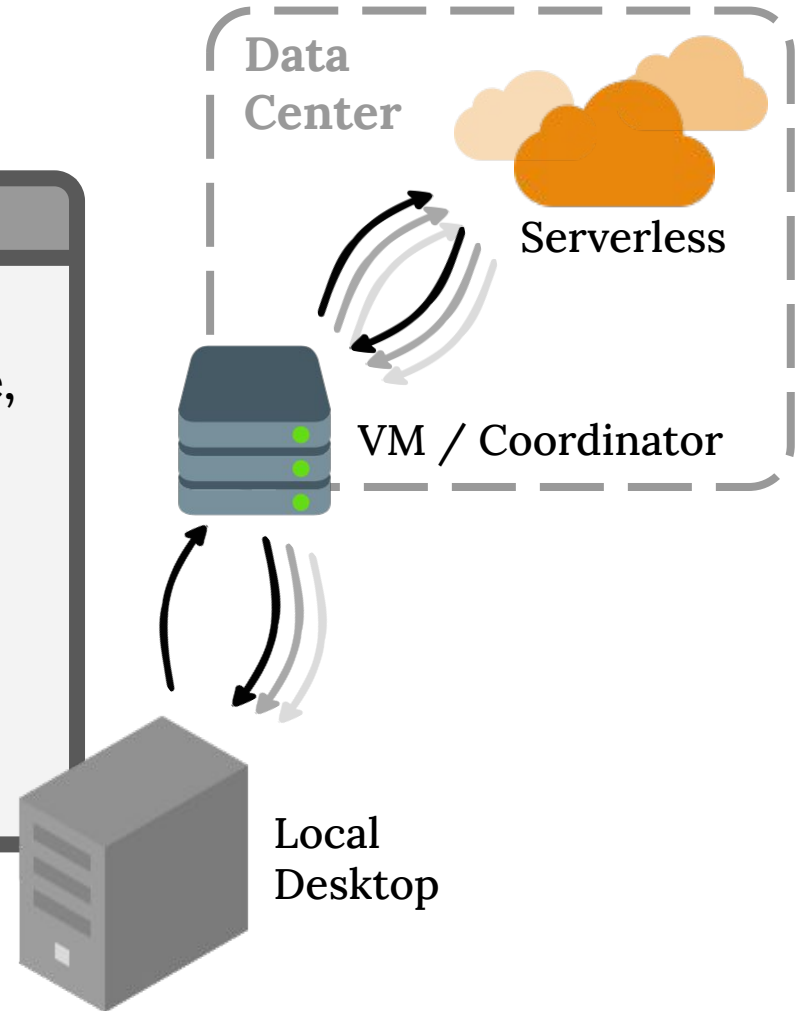


# Evaluation Setup

- > Run 3D printer controlling software
- > Mimic laser movement, temperature, etc.

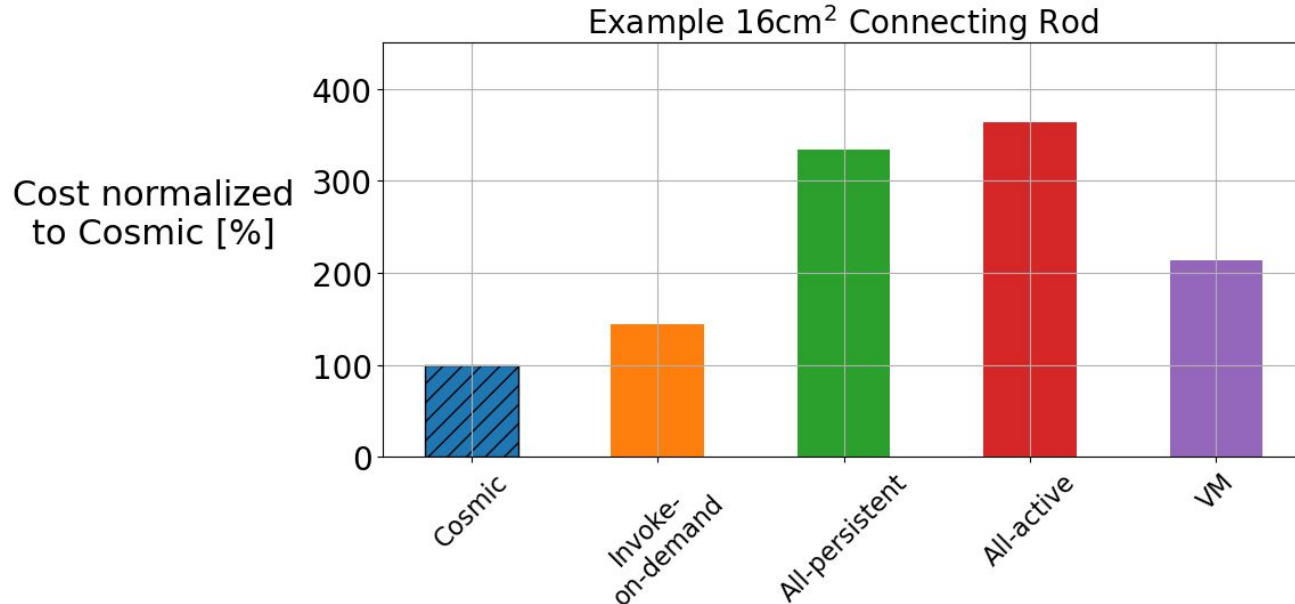


... 30 Mimicked printing jobs  
of different shapes & sizes



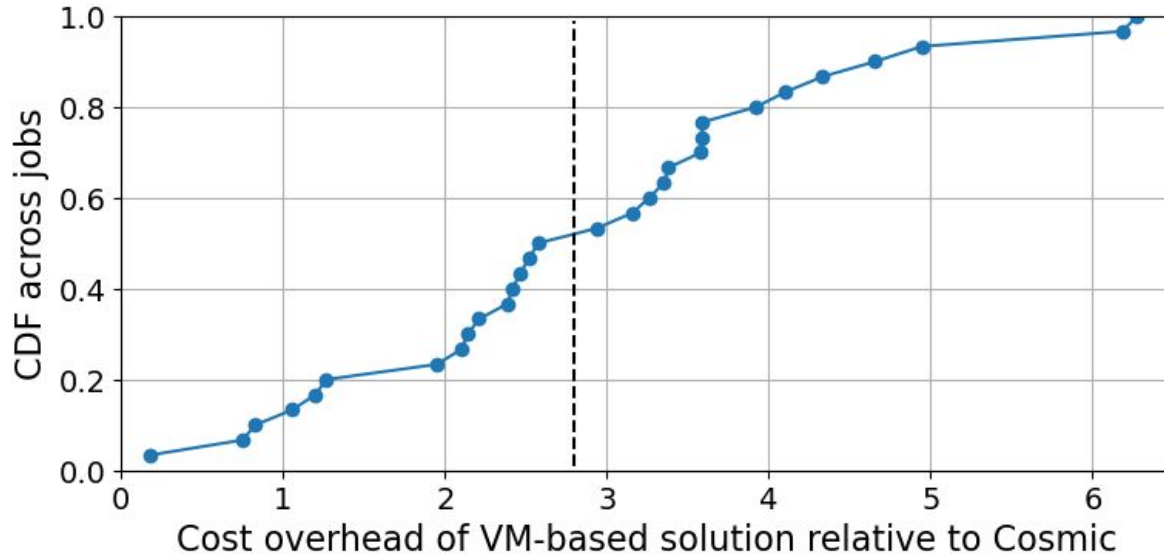
# Meets Real-Time Goals and Outperforms Baselines

- Cosmic's execution of SmartScan **meets timing in all 30 print jobs**
- Cosmic outperforms all 4 baselines in 26 jobs in terms of cost



# Reduces Cost Compared to VMs

- Cosmic achieves 2.8× median cost savings vs VMs across all 30 printing jobs



# Conclusion

**We explored a new serverless workload: cloud-based 3D printer control with strict timing constraints**

- Identified millisecond-scale platform delays that hinder timely execution

**Cosmic uses function grouping and timely speculation**

- Ensures timely execution across 30 print jobs
- Reduces cost by 2.8× at the median relative to the VM approach