

PyTorch Scheduling ML Jobs at Scale

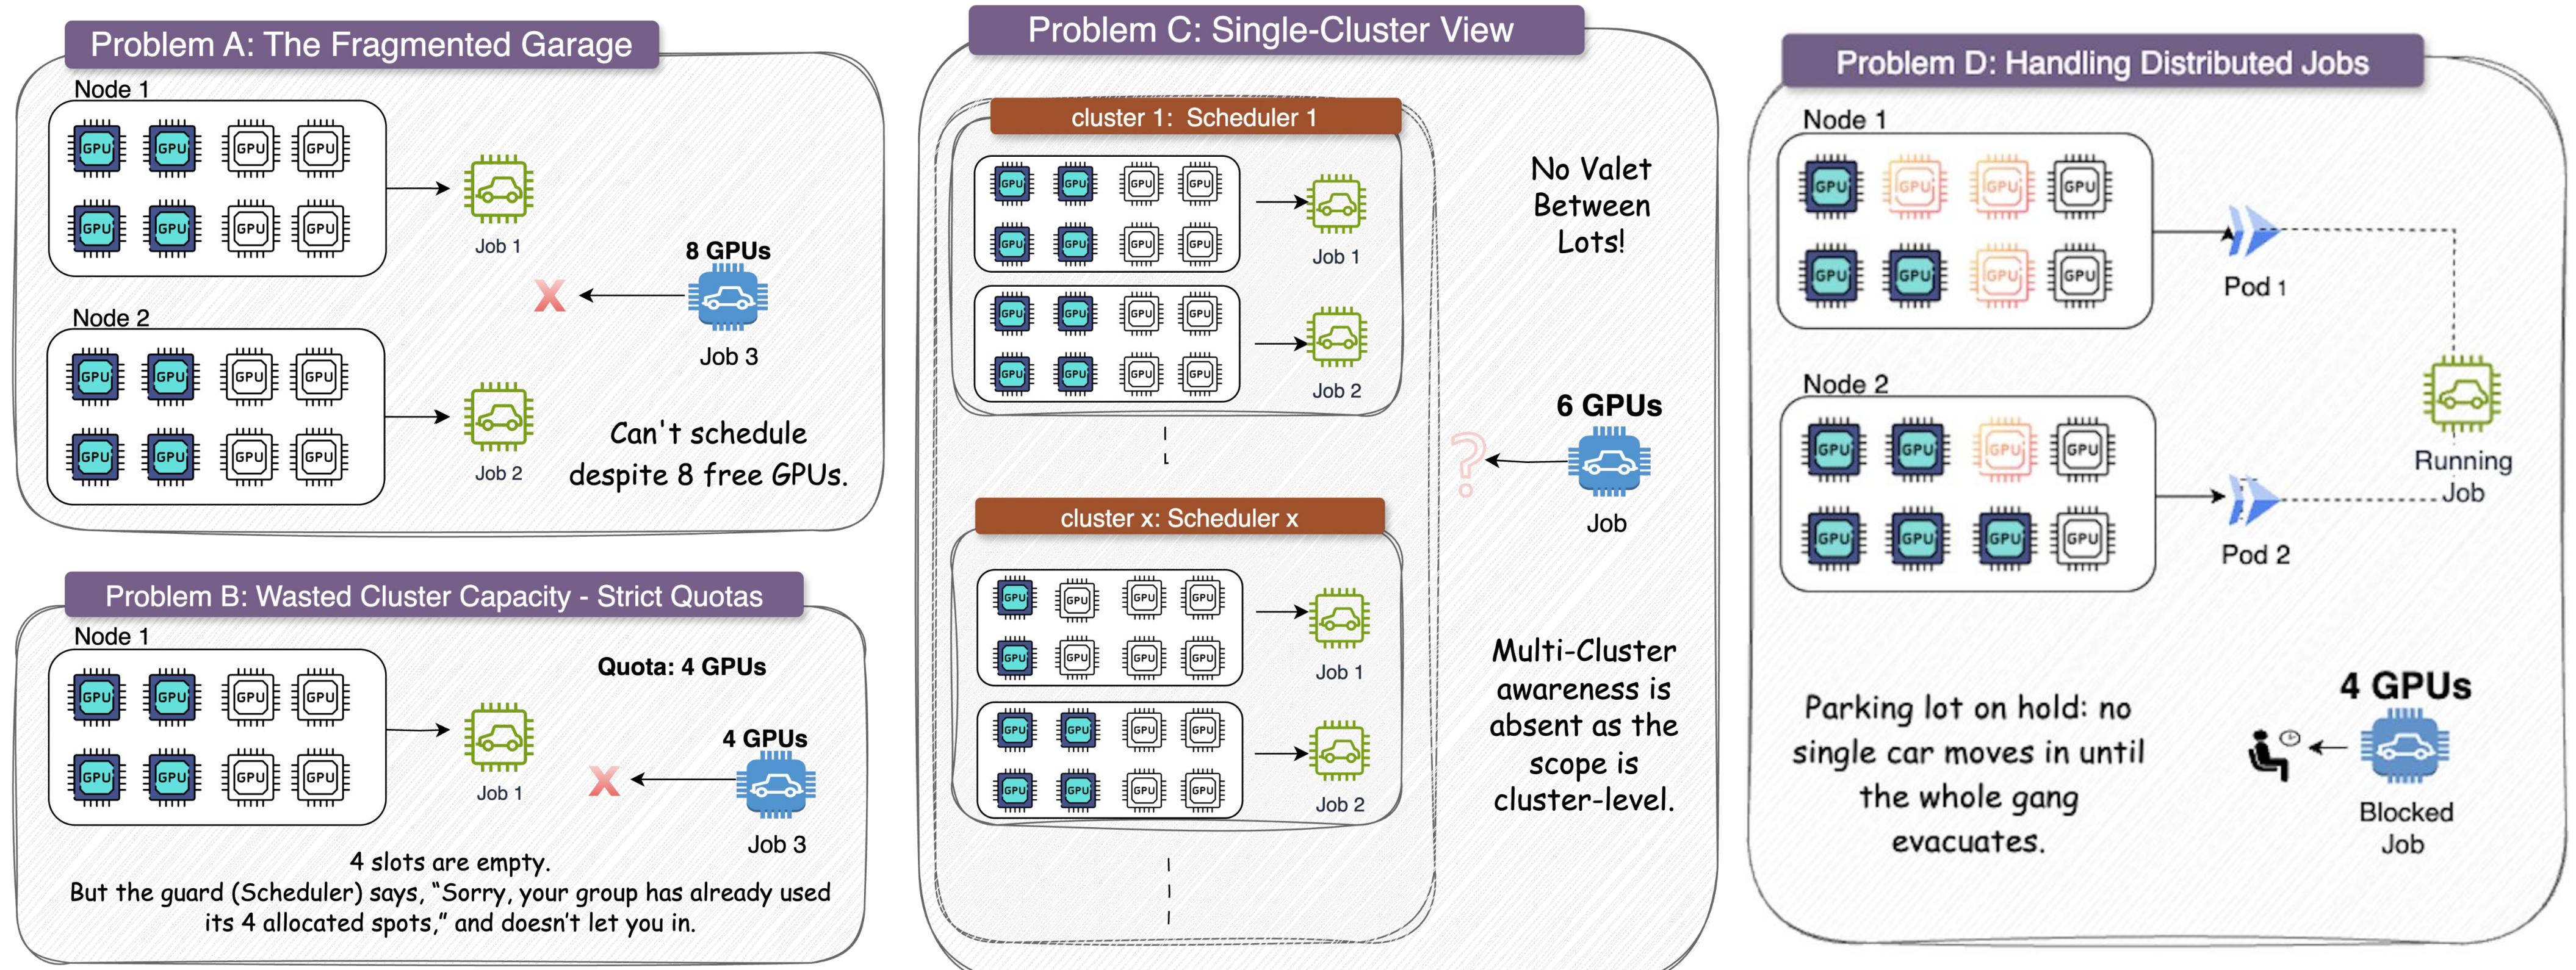
CONFERENCE 2025

Manvi Gupta
(Computer Scientist 1, Adobe Systems)

Problem Space

The GPU Parking Problems

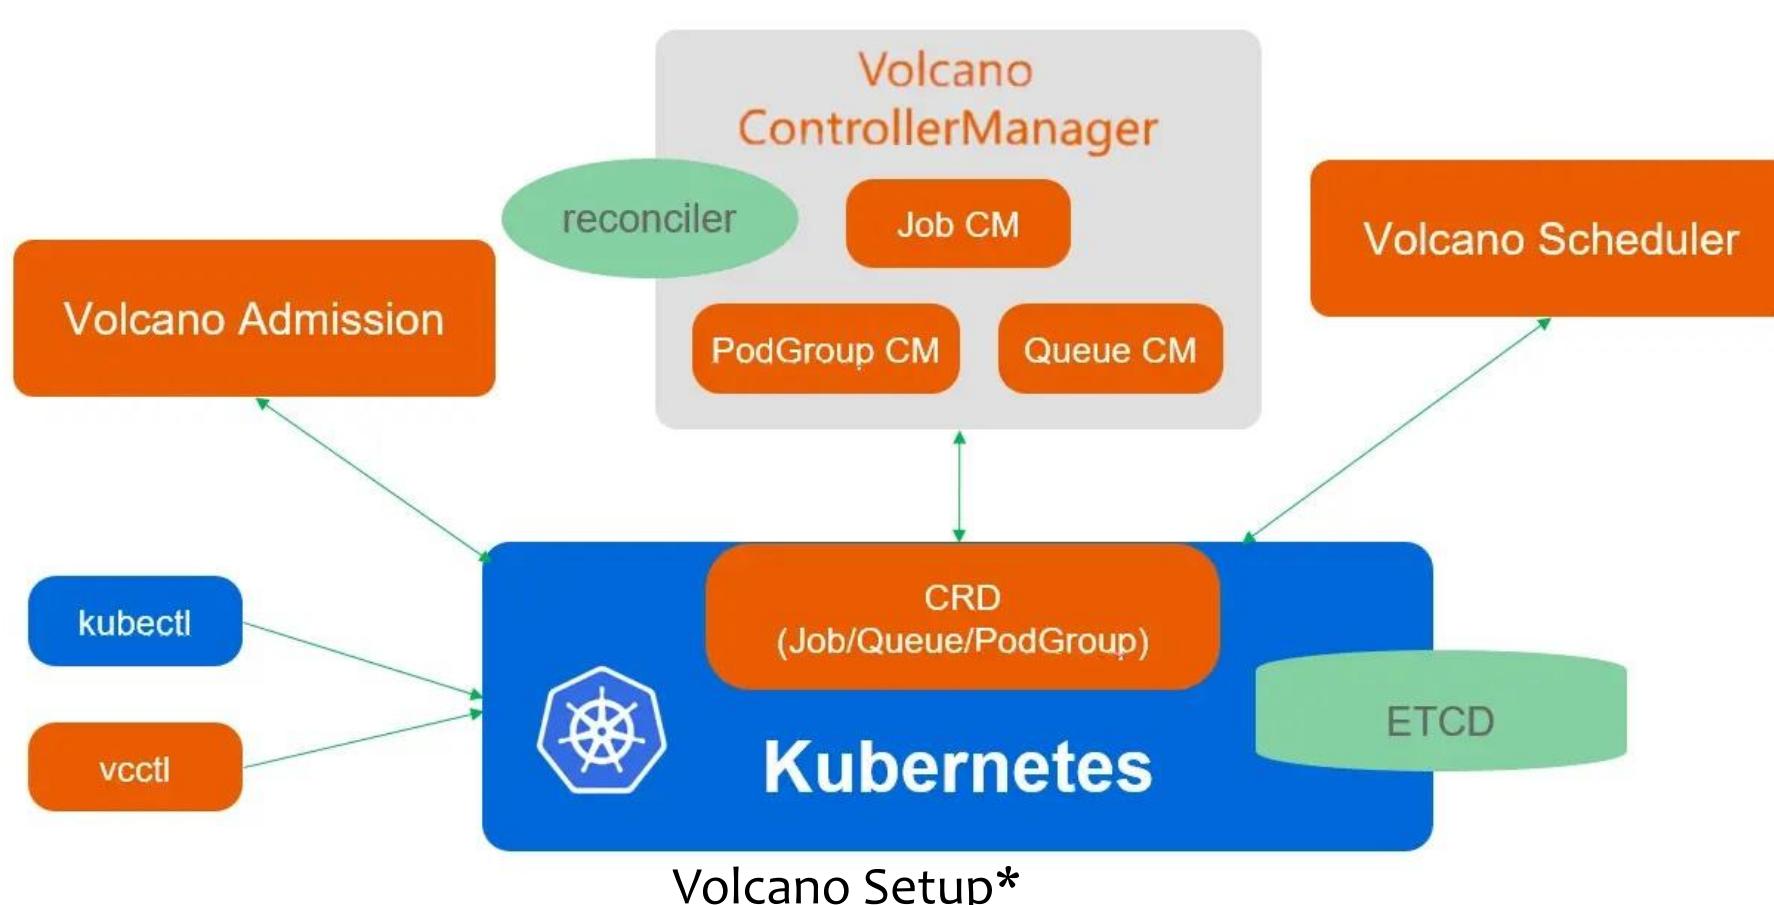
Why standard Kubernetes scheduling fails for ML workloads?



Custom Scheduling

Why Volcano?

1. Highly Extensible CNCF Design
 - a. Native capabilities implemented as plugins
2. Good core framework
 - a. Custom plugins can build a quota-based system driven by number of GPUs
3. OOTB support for PyTorch, TensorFlow workflows



How do we build our scheduler?

1. Custom Plugin
 - Optimal Over Quota Behaviour
 - Project and Job Ordering for fairness
2. New Actions
 - Custom-Reclaim
 - Over-Quota Allocate
3. Behavior tuning for Gang, Messaging, etc.
4. New CRD : ClusterProject
 - Contains project quotas
 - Read by the scheduler

Cluster Project

```
apiVersion: my.research.project/v1
kind: ClusterProject
spec:
  projectName: "ml-research-team"
  gpuQuota:
    "nvidia.com/gpu": 16
    "nvidia.com/a100": 8
  allowOverQuota: true
  status:
    overQuotaGPUs: 4
    queuedJobs: 3
```

Proposed Solution

Two-Pass Scheduling Algorithm:

- **Pass 1 (Under-Quota):** Schedule jobs within project GPU limits using fair-share ordering
- **Pass 2 (Over-Quota):** Best-effort allocation of remaining resources to pending jobs

Topology Awareness:

All pods of a job are scheduled in the same node pool

