INTELLIGENT TASK SCHEDULING IN HETEROGENEOUS DISTRIBUTED INFORMATION SYSTEMS: ENABLING ACADEMIC FREEDOM & DIGITAL INCLUSION

SERHII YENHALYCHEV

PHD STUDENT

THE SCHEDULING CHALLENGE IN HETEROGENEOUS ENVIRONMENTS

Heterogeneity

- Compute: CPU / GPU , on-prem HPC, and multi-cloud regions
- Storage tiers: object / block / file / archival (latency & cost vary by tier)
- Networks: intra-/inter-region latency, variable bandwidth, egress policies

Uncertainty

- Incomplete/late/incorrect metadata (queue time, throughput, spot prices)
- Prediction error (runtime, preemption risk, network congestion)
- Telemetry lag & drift (sampling intervals, stale registries)

Objective Tension

- **Performance:** throughput & tail latency
- **Economics:** cost per job, budget caps, spot volatility
- **Sustainability:** energy use & carbon intensity
- **Equity:** fair share across tenants

ACADEMIC FREEDOM & DIGITAL INCLUSION

Why this matters

- **Academic freedom:** researchers must pursue inquiry without gatekeeping by compute access or vendor lock-in.
- **Inclusion:** equitable access to GPUs/CPUs and storage so under-resourced groups aren't sidelined.
- **Risk if ignored:** compute monopolies, pay-to-play science, biased results due to skewed access to scale.

What the scheduler can guarantee

- **Fair-share policies:** minimum resource share per cohort/tenant; anti-starvation bounds.
- **Transparent decisions:** auditable logs + "why this job, why now" explanations.
- **Accessibility knobs:** budget-aware placement, data-locality compliance (e.g., EU-only), low-bandwidth paths.
- **Resilience to power asymmetries:** quotas, rate limits, and caps to prevent dominance by a single lab

RESEARCH GOAL & CONTRIBUTIONS

Design and evaluate an **uncertainty-aware, multi-objective scheduler** for heterogeneous distributed systems that **guarantees fairness and transparency** while balancing performance, cost, and energy.

Method (at a glance)

- Inputs: job DAGs, resource catalog, historical traces, budgets/SLAs, data-locality rules;
- Modeling: predictors for runtime/queue/network with confidence intervals;
- Optimizer: multi-objective (throughput, latency, cost), constraint handling for compliance (e.g., EU-only);
- Outputs: placement plan + confidence bands + human-readable decision receipt ("why this job, why now")

RESEARCH GOAL & CONTRIBUTIONS

Expected Key Contributions

- Uncertainty-aware estimation: runtime/queue/network predictors.
- Multi-objective optimization: explicit trade-offs among performance, cost, energy, and equity.
- Fairness guarantees: min-share quotas, max-wait bounds, budget-aware throttling; protection for under-resourced cohorts.
- Transparency & auditability: per-decision explanations, reproducible seeds, and queryable logs for governance.
- Robustness to bad/missing metadata: imputation, drift detection, and fallback heuristics under severe uncertainty.
- Policy/Compliance layer: data-locality rules, rate-limits against dominance/abuse.

FAIRNESS & TRANSPARENCY BY DESIGN

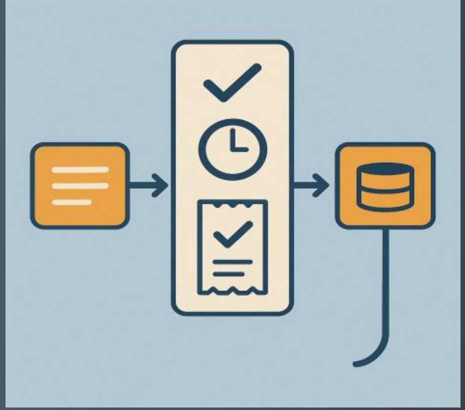

What we guarantee (policy layer)

- **Minimum-share quotas:** each cohort receives≥ ρ of compute over a sliding window WWW.
- **Max-wait bounds:** P95 queue time per cohort ≤ LLL (anti-starvation).
- **Budget-aware fairness:** respect per-tenant caps; graceful throttling vs. hard cutoffs.
- **Compliance guards:** data-locality (e.g., EU-only), export-control, carbon-aware placement.

How we measure (fairness & access)

- **Jain's index** (resource share across cohorts) → target ≥ 0.85.
- **Access share** (% of GPU hours to protected/under-resourced cohorts).
- **Tail latency** within policy bounds.
- **Gini coefficient** on compute allocation (lower is better).
- **SLO attainment** per cohort (service parity).

FROM PRINCIPLES TO PRACTICE: HOW WE GOVERN THE SCHEDULER

Governance: who decides & how

- Policy charter: scope, objectives, metrics (fairness, access), review cadence (e.g., monthly).
- Oversight roles: Product (SLOs), Research (fairness), Security (risk & access), FinOps (budget).
- Change management: versioned policies, RFC → review → staged rollout → post-mortem; emergency rollback path.

Ethics: what we optimize & protect

- **Fairness commitments:** min-share quotas, max-wait P95 bounds, anti-starvation.
- **Transparency:** per-decision "receipt" (inputs, constraints met, why-now, what-if).
- **Grievance & appeal:** ticket path for cohorts; SLA to investigate & remediate.
- **Abuse resistance:** Sybil/identity binding, rate limits, credit exhaustion rules.

Compliance: what must never break

- **Data locality & sovereignty:** region pinning (e.g., EU-only workloads), cross-border controls.
- **Privacy & DPIA:** data-minimization in telemetry, retention limits.
- **Security controls:** encryption in transit/at rest (KMS/HSM), RBAC/ABAC, key rotation, secrets hygiene.
- **Regulatory & contractual:** export-control filters, IRB/ethics for sensitive datasets, license obligations.

Operational assurance

- **Audits & logging:** signed logs, reproducible seeds, policy version tag, model hash.
- **Risk mgmt:** register with likelihood/impact, owners, mitigations, review dates.
- **Monitoring KPIs:** Jain's index, SLO attainment, % policy-compliant placements.
- **Red-teaming:** bias probes, gaming attempts (job splitting, burst floods), locality evasion tests.

CONCLUSIONS & NEXT STEPS

- We built an uncertainty-aware, multi-objective scheduler with fairness and transparency built in.
- Results indicate we can improve equitable access while staying close to performance and cost targets.
- Next steps: run a small pilot, publish reproducible artifacts, and formalize lightweight governance.
- Roadmap: expand regions/accelerators, and expose a simple “why this decision” view.

THANK YOU !

Contact email: engalichev.sergiy@hneu.net