



# CS 839: Topics in Database Management Systems

## Lecture 22: Hybrid Transactional/Analytical Processing (HTAP)

Xiangyao Yu

11/20/2023

# Group Discussion from Lecture #21

---

One idea is to use a standalone cloud service to monitor a database to make tuning decisions. What would be the disadvantages of this design compared to integrating the monitor inside the database? Are there ways to overcome these disadvantages?

A large system may have millions of internal and external knobs. What characteristics make a knob suitable for auto-tuning techniques?

# Group Discussion from Lecture #21

---

## Advantages of “tuning as a service”

- Reusability
- Independent resource consumption
- Plug and play; can upgrade to better tuners

## Disadvantages of “tuning as a service”

- Longer latency and more network traffic
- External to the database therefore reactionary in nature
- Worse accuracy since it has no access to database internals
- New metrics require changes to API and tuner
- Security and privacy

## Ideas

- Hierarchical tuner with local and global tuning separated

# Group Discussion from Lecture #21

---

Characteristics making a knob suitable for auto-tuning?

- Has significant performance impact
- Impact throughput or latency and can be measured in a short period of time
- Inexpensive to tune the knob (e.g., does not require restarting the DB)
- Knobs should be independent
- Knobs that require frequent tuning

# Discussion Questions

---

In a sense, all databases in the 80s are HTAP databases. DB systems then broke into TP and AP for better workload-specific performance. What are the fundamentally new demands today that drive the unification of these two worlds again?

Submit by **11:59pm CT on Tuesday** (11/21). Title starts with “[Discussion L22]”