

CS 265

*Stratos Idreos*

BIG DATA SYSTEMS

NoSQL | Neural Networks | Image AI | LLMs | Data Science

# Logistics

Systems projects and Labs are active.

## **Logistics**

Systems projects and Labs are active.

## **Action Items for now**

Readings and go to labs to start getting familiar with systems projects.

**preparation**

**Get familiar with the very basics of traditional database architectures:**

Architecture of a Database System. By J. Hellerstein, M. Stonebraker and J. Hamilton. Foundations and Trends in Databases, 2007

**Get familiar with very basics of modern database architectures:**

The Design and Implementation of Modern Column-store Database Systems. By D. Abadi, P. Boncz, S. Harizopoulos, S. Idreos, S. Madden. Foundations and Trends in Databases, 2013

**Get familiar with the very basics of modern large scale systems:**

Massively Parallel Databases and MapReduce Systems. By Shivnath Babu and Herodotos Herodotou. Foundations and Trends in Databases, 2013

first readings



## **The Periodic Table of Data Structures.**

Stratos Idreos, Kostas Zoumpatianos, Manos Athanassoulis, Niv Dayan, Brian Hentschel, Michael S. Kester, Demi Guo, Lukas Maas, Wilson Qin, Abdul Wasay, Yiyou Sun.  
IEEE Data Engineering Bull. Sep, 2018

## **Design Continuums and the Path Toward Self-Designing Key-Value Stores that Know and Learn.**

Stratos Idreos, Niv Dayan, Wilson Qin, Mali Akmanalp, Sophie Hilgard, Andrew Ross, James Lennon, Varun Jain, Harshita Gupta, David Li, and Zichen Zhu. Proceedings of CIDR Conference on Innovative Data Systems Research, 2019.



# Readings for NoSQL research & project and mathematical systems modelling

**Monkey: Optimal Navigable Key-Value Store.** Niv Dayan, Manos Athanassoulis, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2017

**Dostoevsky: Better Space-Time Trade-Offs for LSM-Tree Based Key-Value Stores via Adaptive Removal of Superfluous Merging.** Niv Dayan, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2018

**The Log-Structured Merge-Bush & the Wacky Continuum.** Niv Dayan, Stratos Idreos. In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2019




# Reading SysML research & project

**$\mu$ -TWO: 3x Faster Multi-model Training with Orchestration and Memory Optimization**, Sanket Purandare, Abdul Wasay, Animesh Jain, Stratos Idreos, Conference on Machine Learning and Systems (MLsys), 2023


**we will add 4 more in the next 2 weeks**




**fundamental** building blocks




**properties** when combined






DESIGN SPACE



COST SYNTHESIS




WHAT-IF

# NoSQL Key-value Stores



machine learning   social media  
smart homes   web browsers  
phones   web-based apps  
security   health devices  
graphs   analytics




[1,0,0,1,1,1]  
hash fun.

bloom  
filters






...




[min-max]  
/page

fence  
pointers



...



buffer

Level 1

Level 2

Level 3

...

Level N

MEMORY

DISK

pages



SSTables



tiered

leveled

sorted

## Today/This week

Advanced LSM-tree/Storage designs.

We will see that “advanced research is fundamentally simple”

## Today/This week

Advanced LSM-tree/Storage designs.

We will see that “advanced research is fundamentally simple”

Intro to Step 2 of self-designing systems: cost synthesis (math & ML).



# **BITS PER ENTRY IN FILTERS: OPTIMIZED OUT**

**M**onkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017






# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017













# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017

bits per entry:  
fixed per run







# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017

worst lookup cost:  
sum of false positive rates

bits per entry:  
fixed per run





# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT






Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017

*the same memory budget  
is more impactful at smaller levels*



bits per entry:  
fixed per run






# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017

*the same memory budget  
is more impactful at smaller levels*









# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017

*the same memory budget  
is more impactful at smaller levels*




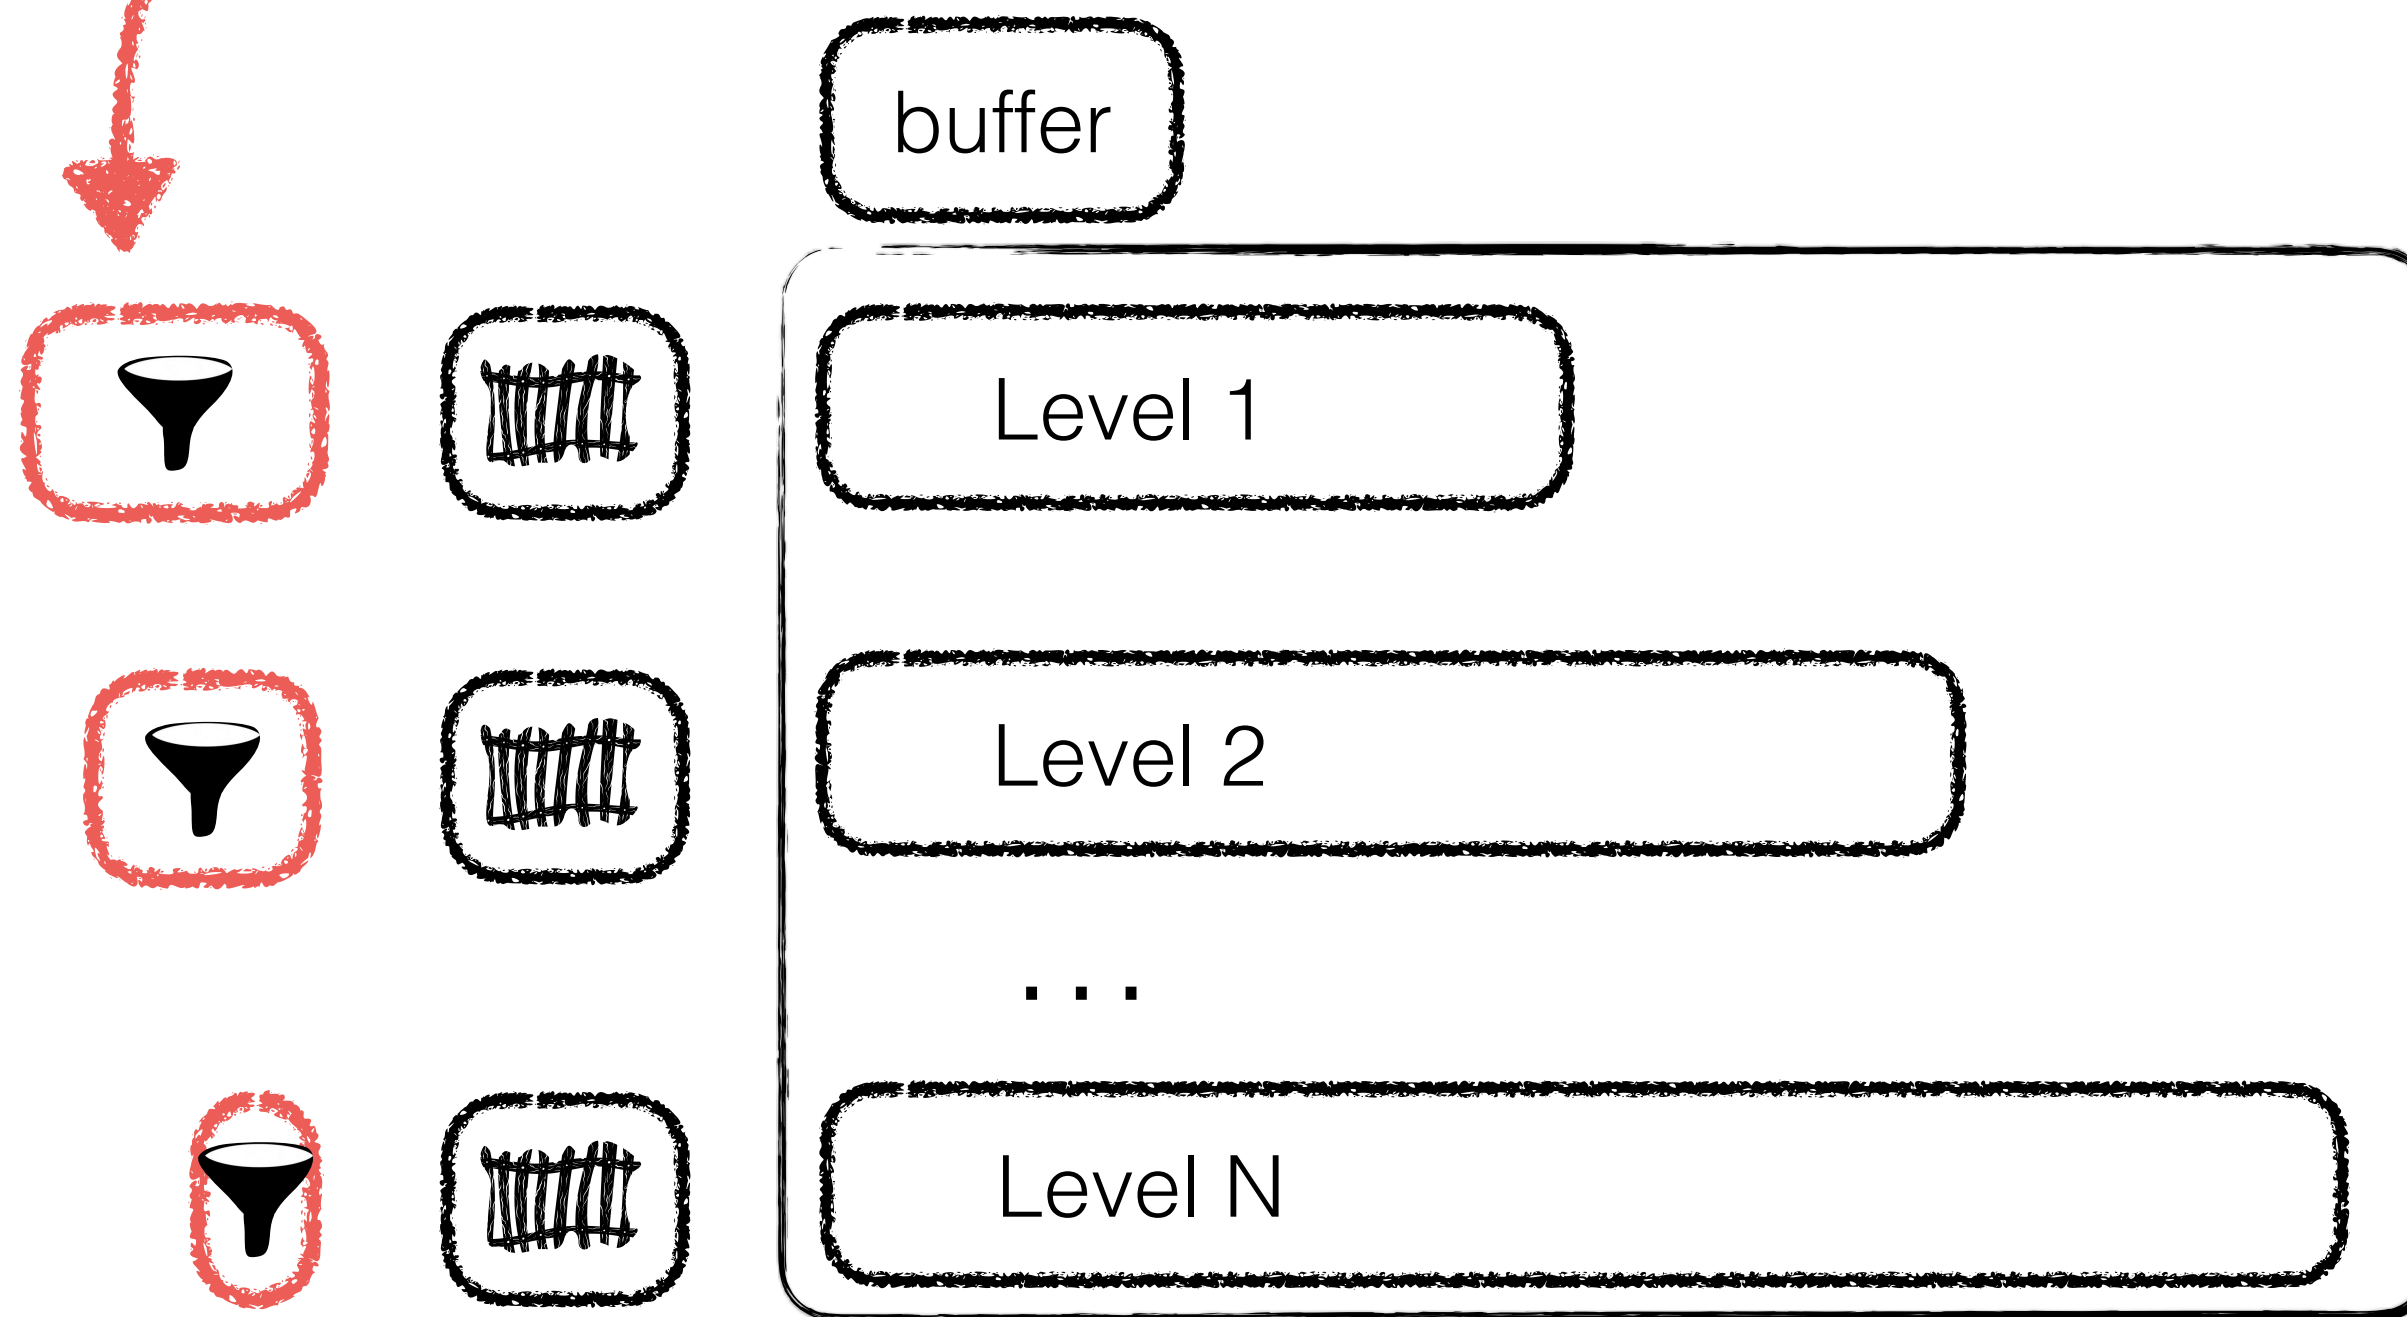


# BITS PER ENTRY IN FILTERS: OPTIMIZED OUT

Monkey: **O**ptimal **N**avigable **K**ey-Value Store

@SIGMOD2017

*the same memory budget  
is more impactful at smaller levels*



uniform, zero result, point queries, entry size=1KB

CS 265

*Stratos Idreos*

BIG DATA SYSTEMS

NoSQL | Neural Networks | Image AI | LLMs | Data Science