


# DISKS ARE LIKE SNOWFLAKES: NO TWO ARE ALIKE

Elie Krevat, Joseph Tucek\*, Greg Ganger (CMU, \* HP Labs)

## OVERVIEW


- Every modern disk drive is unique
  - Each has unique BW, by design
  - Even for same make and model
- Implication: perf. homogeneity not feasible
  - Varies across different disks in batch
  - Varies across different heads in each disk



Bandwidth effects from different track densities

## WHY: ADAPTIVE ZONING

- Traditional: homogeneous within make/model
  - Each disk head/surface configured identically
  - Same zoned recording density
    - Same per-track density and read/write BW
- Dilemma: waste due to process variation
  - Disk heads are ICs manufactured like CPUs
  - Some disk heads too slow → discarded
  - Some disk heads too good → throttled
- Modern: adapt to capabilities of each head
  - Average areal density maintained
  - But, each head/surface provides unique BW

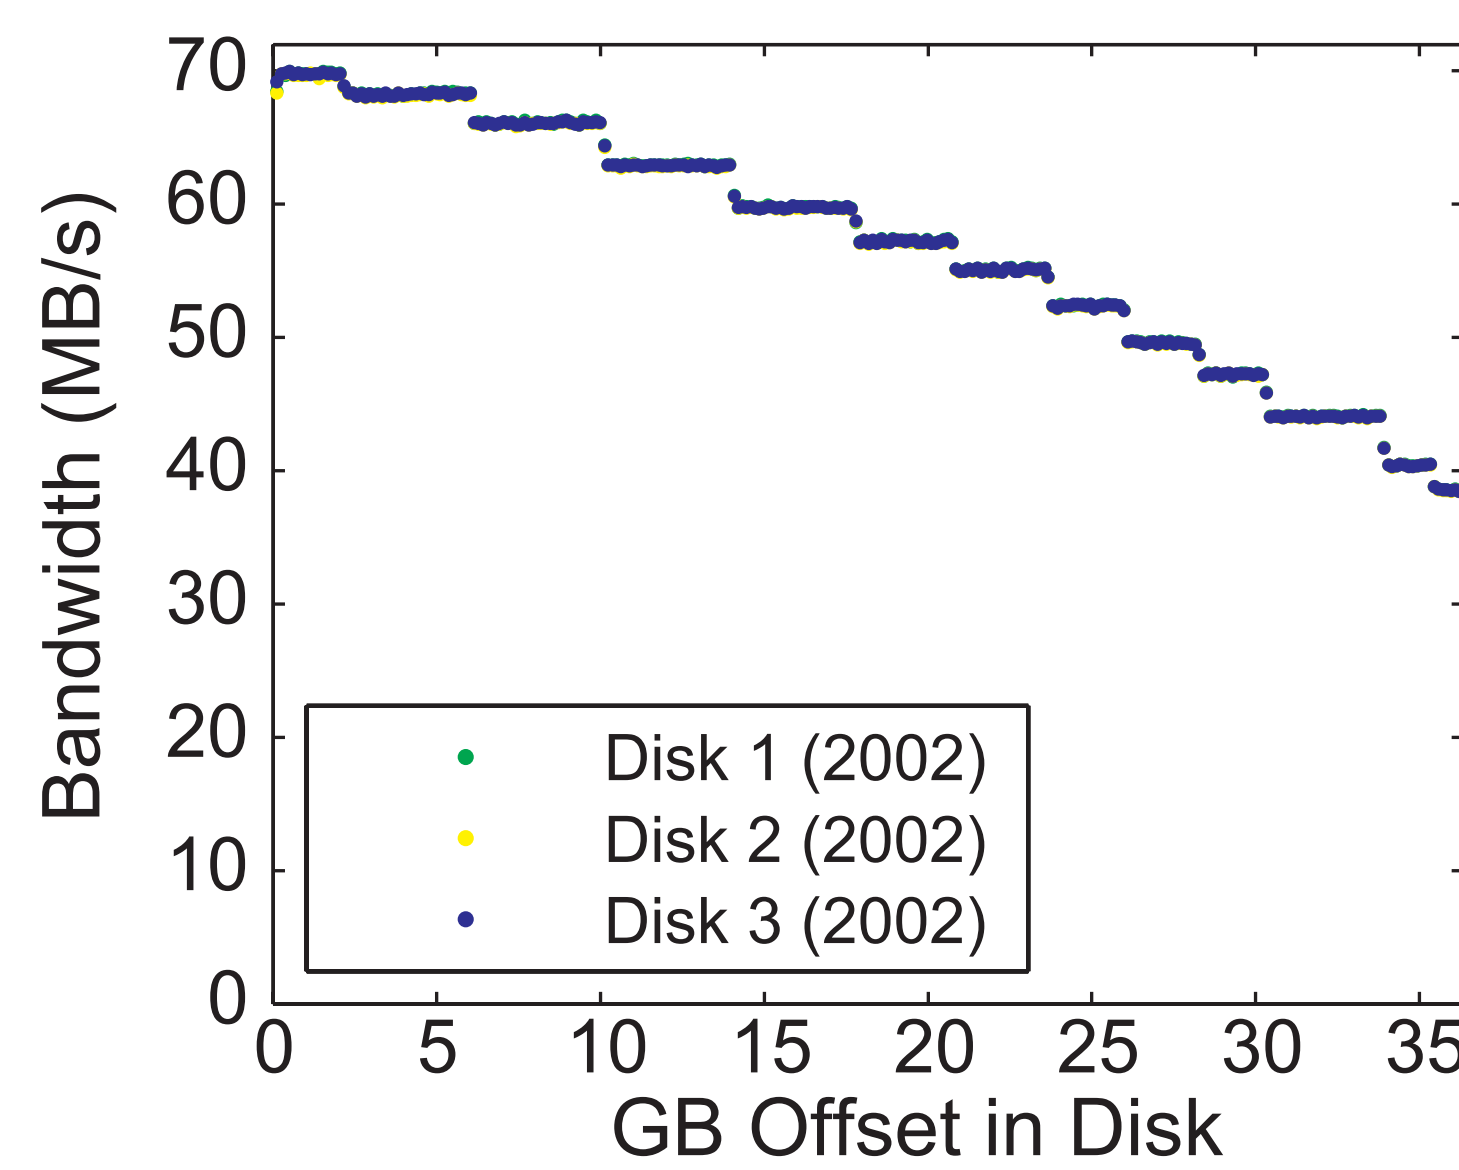


Adaptive zoning on two disk platters

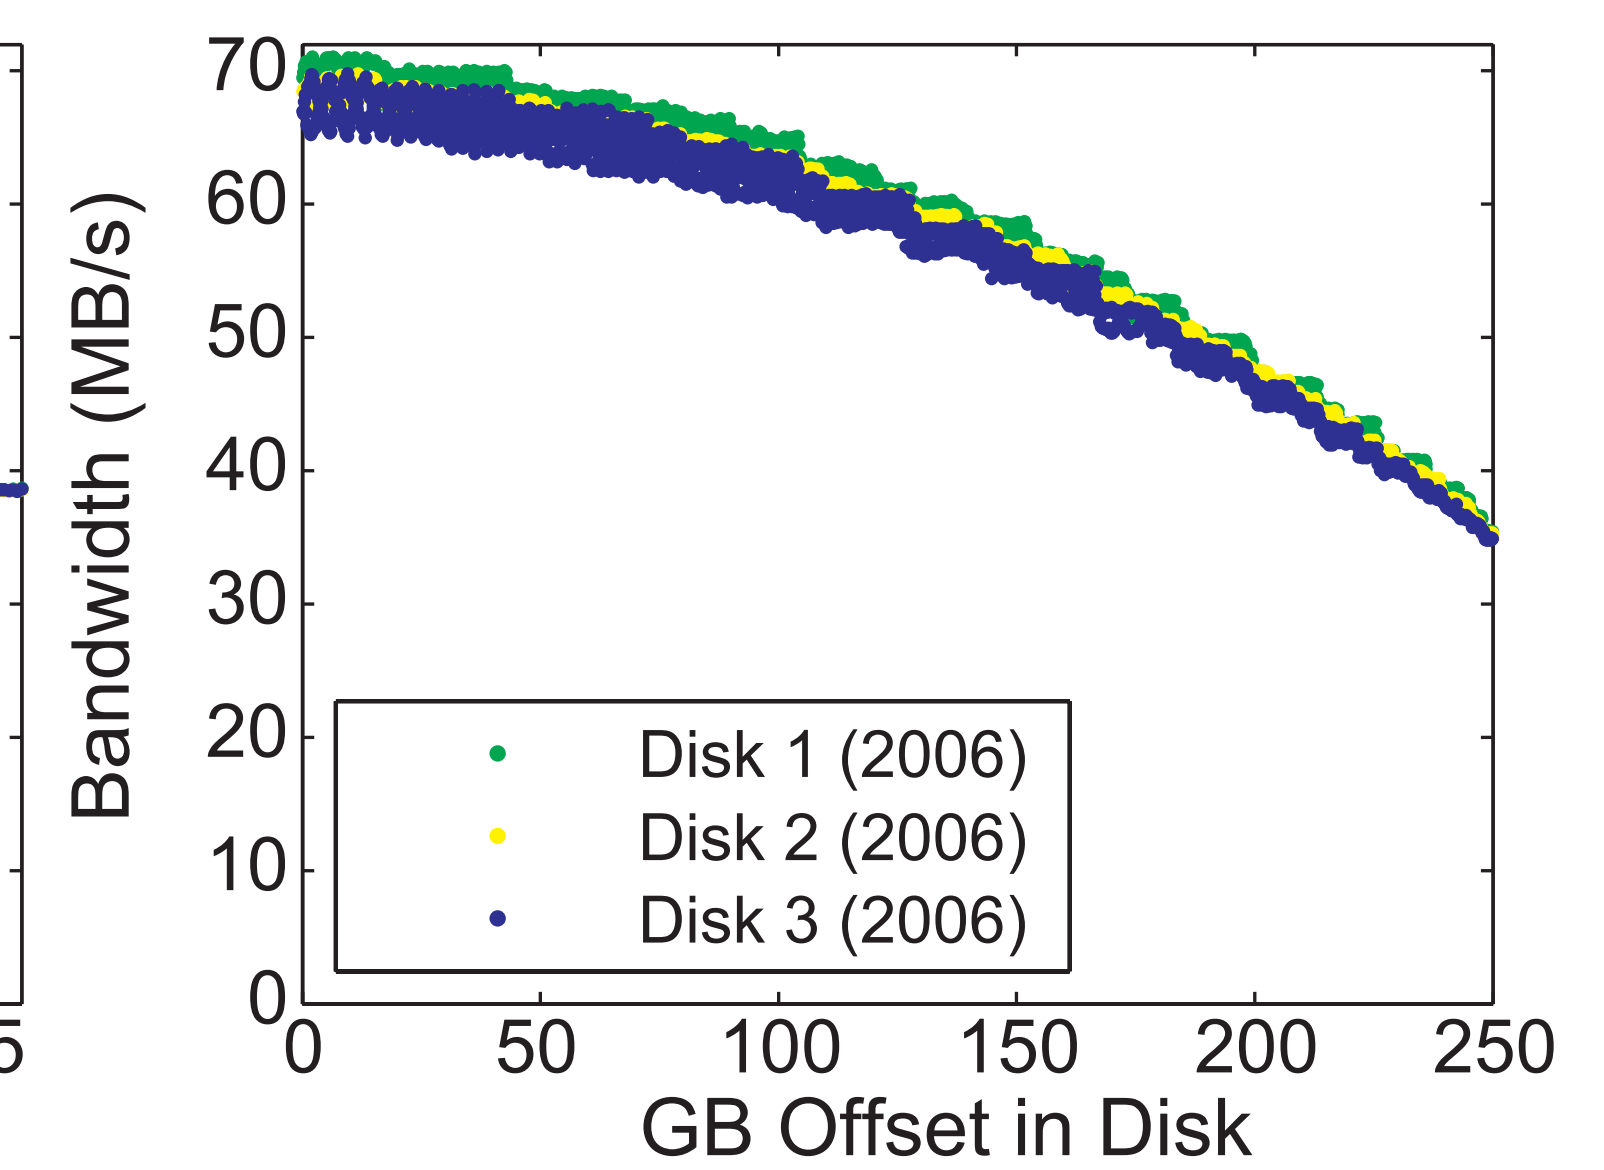
## ILLUSTRATIVE DISK MEASUREMENTS

- All data here: streaming block reads
- Trend is increasingly variable performance
  - 2002 era: Same speeds (only traditional zoning)
  - 2006-era: Variability across disks
  - 2008-era: More pronounced variability
- Reading first quarter of disk:
  - 2002-era: <1 MB/s range (0.2 stdev)
  - 2008-era: 21 MB/s range (4.4 stdev)


2002-model disks



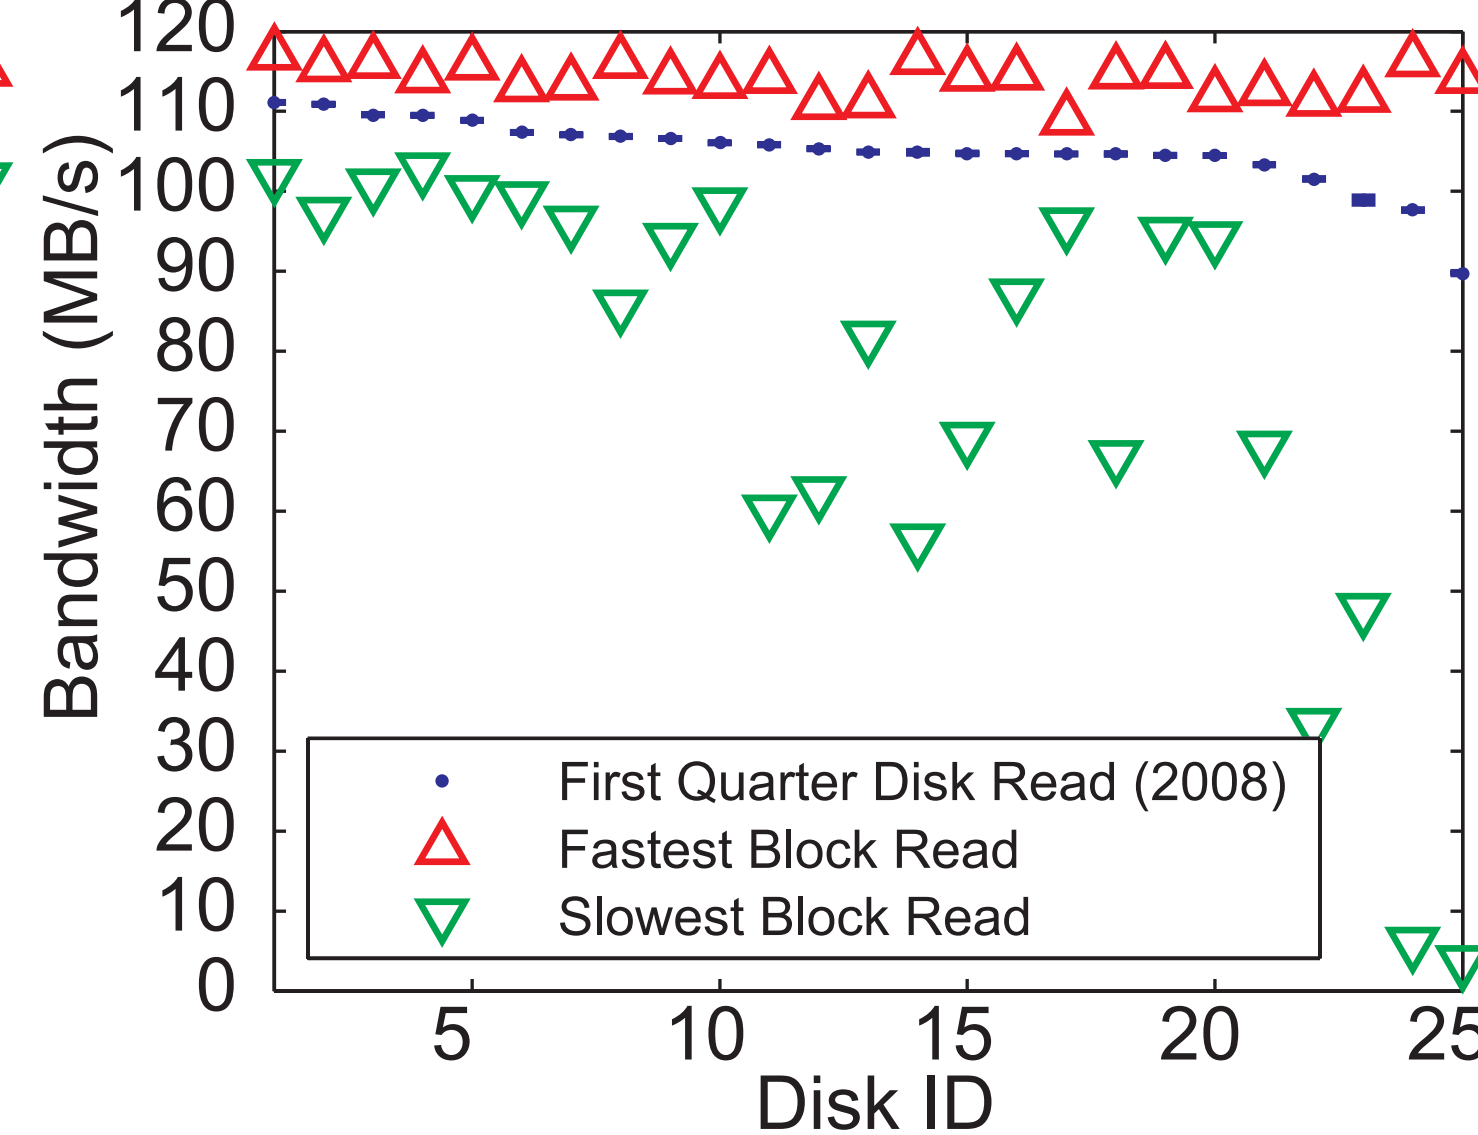
2006-model disks



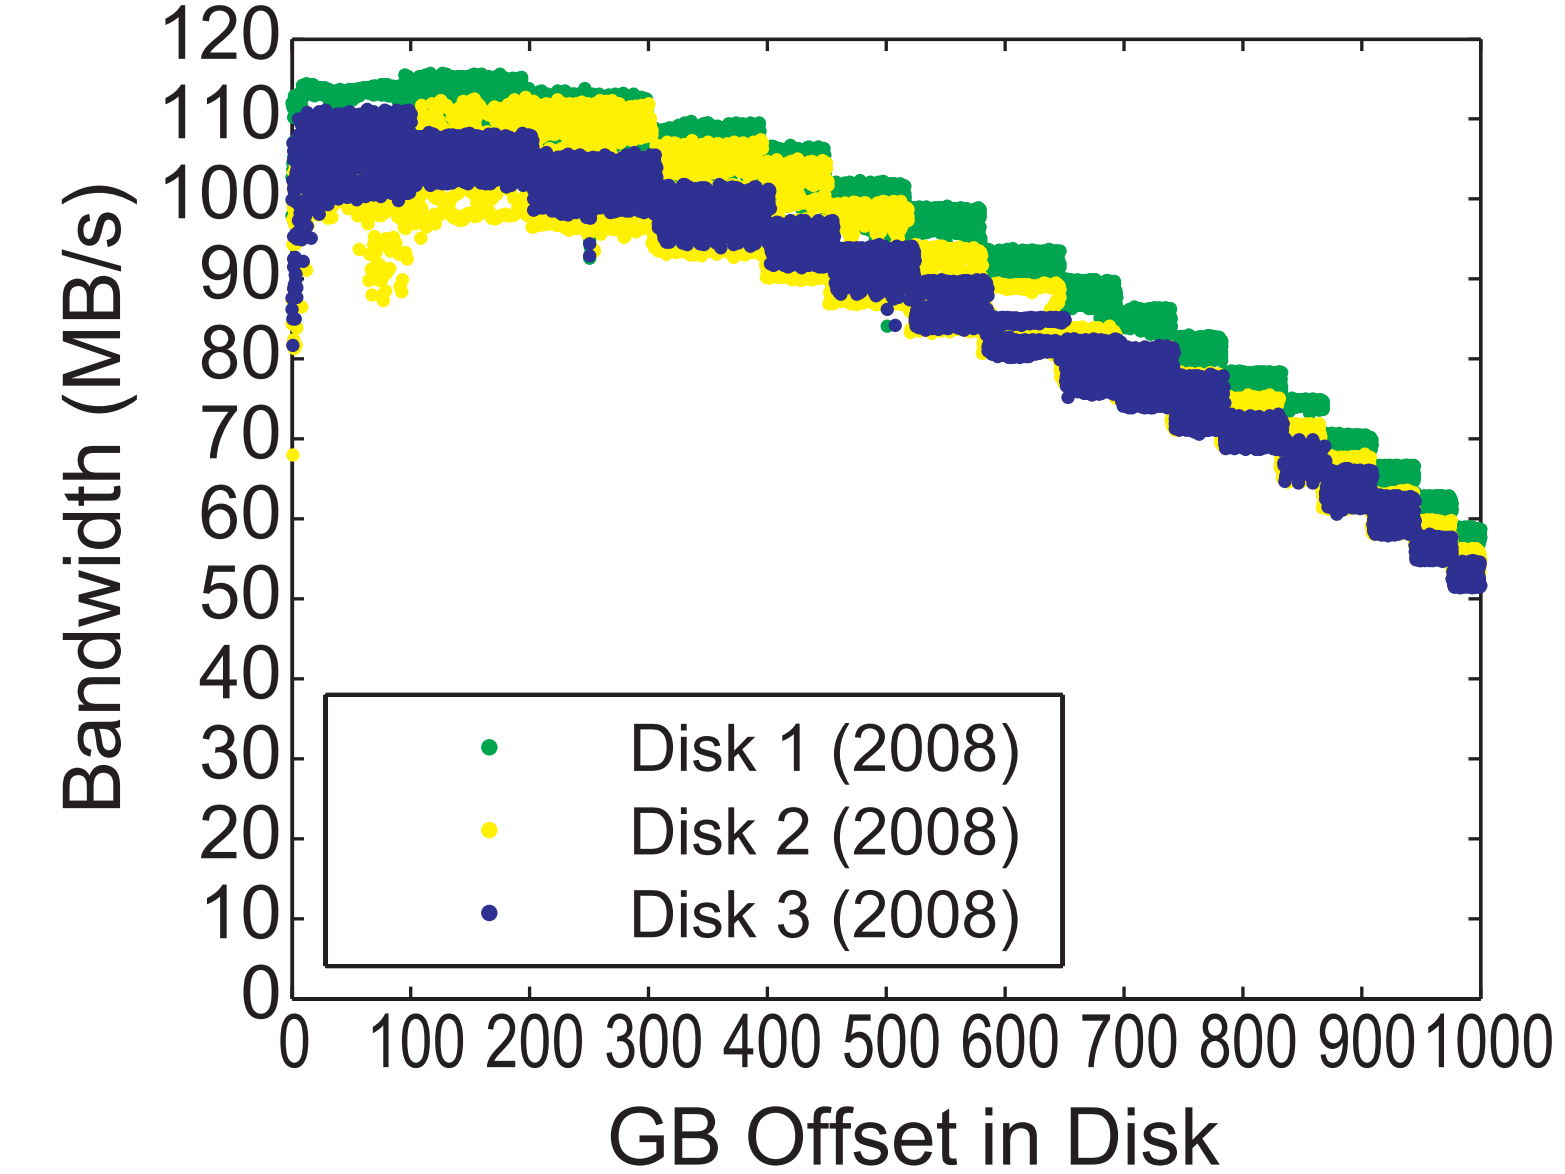
2002-model disks



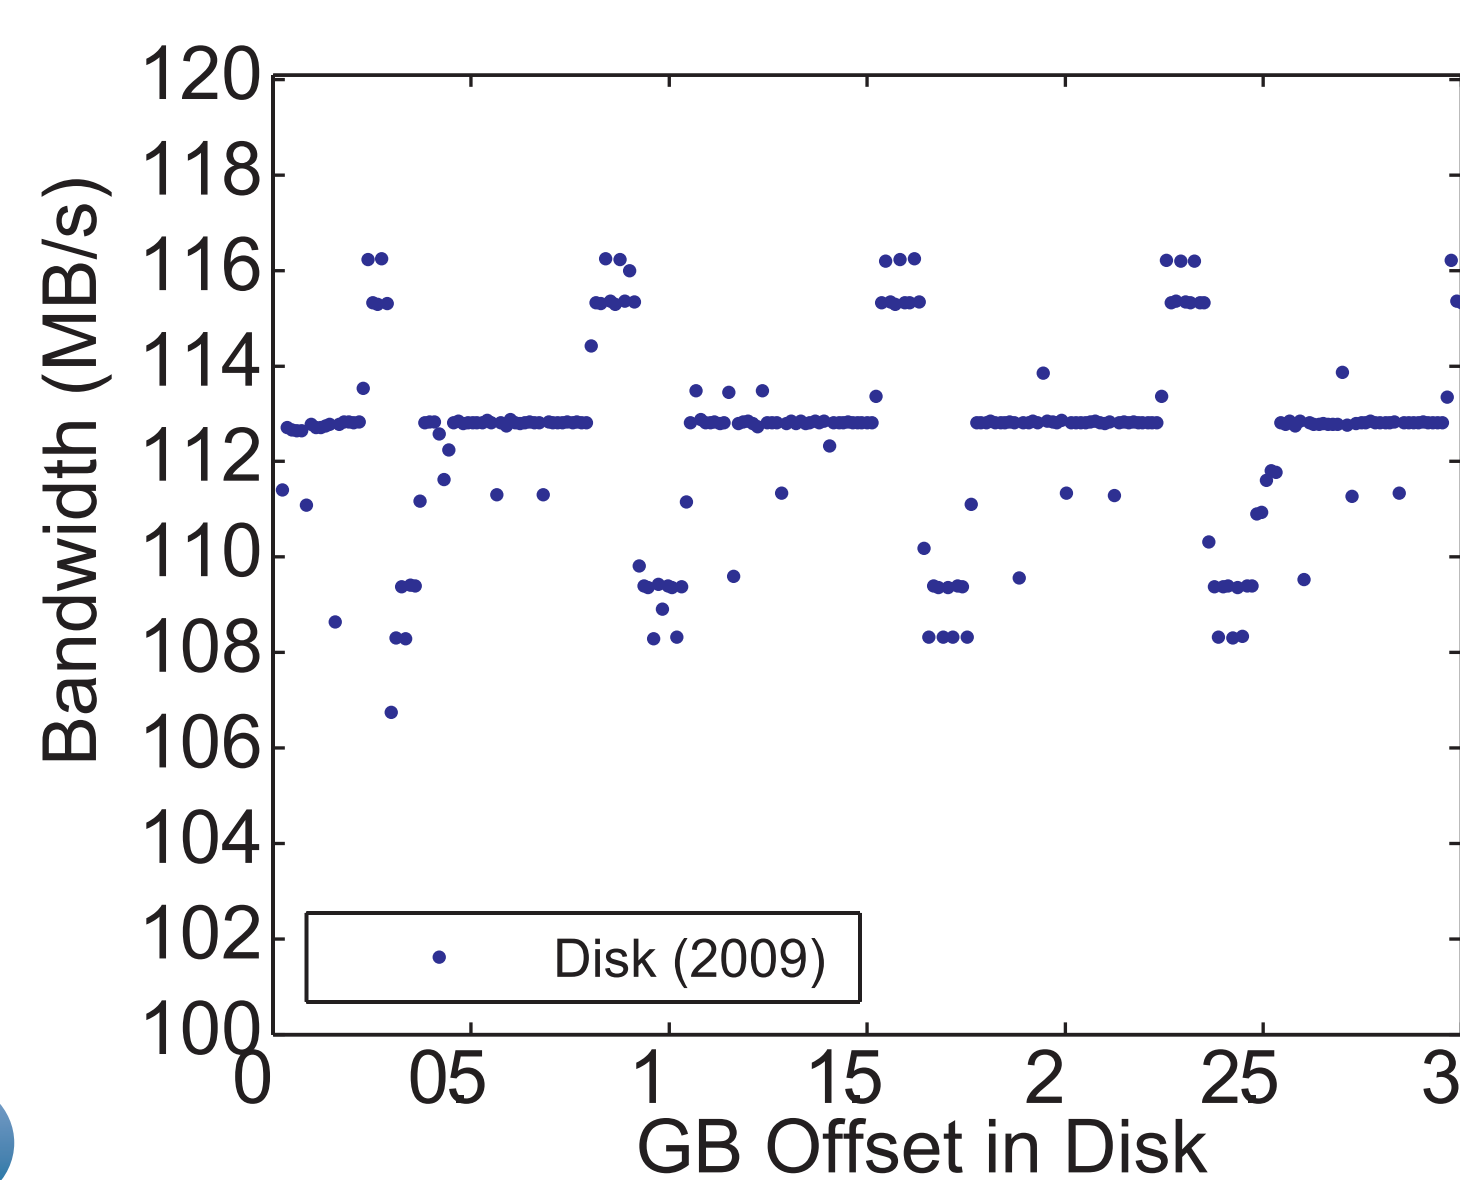
2008-model disks



2008-model disks



- Also differently varied BW within each drive



Intra-disk behavior

## IMPLICATIONS

- No more homogeneous disk-based servers
- Disk striping limited by slowest disk
- Static equal-work partitioning wasteful
- Various other low-level disk techniques broken

