MY WEAK CONSISTENCY IS STRONG WHEN BAD THINGS DO NOT COME IN THREES

ZECHAO SHANG JEFFREY XU YU

THE UNIVERSITY OF CHICAGO 香港中文大學 The Chinese University of Hong Kong

DISCLAIMER: NOT AN OLTP TALK

HOW TO GET *ALMOST* EVERYTHING FOR NOTHING

SHARED-MEMORY SYSTEM IS BACK

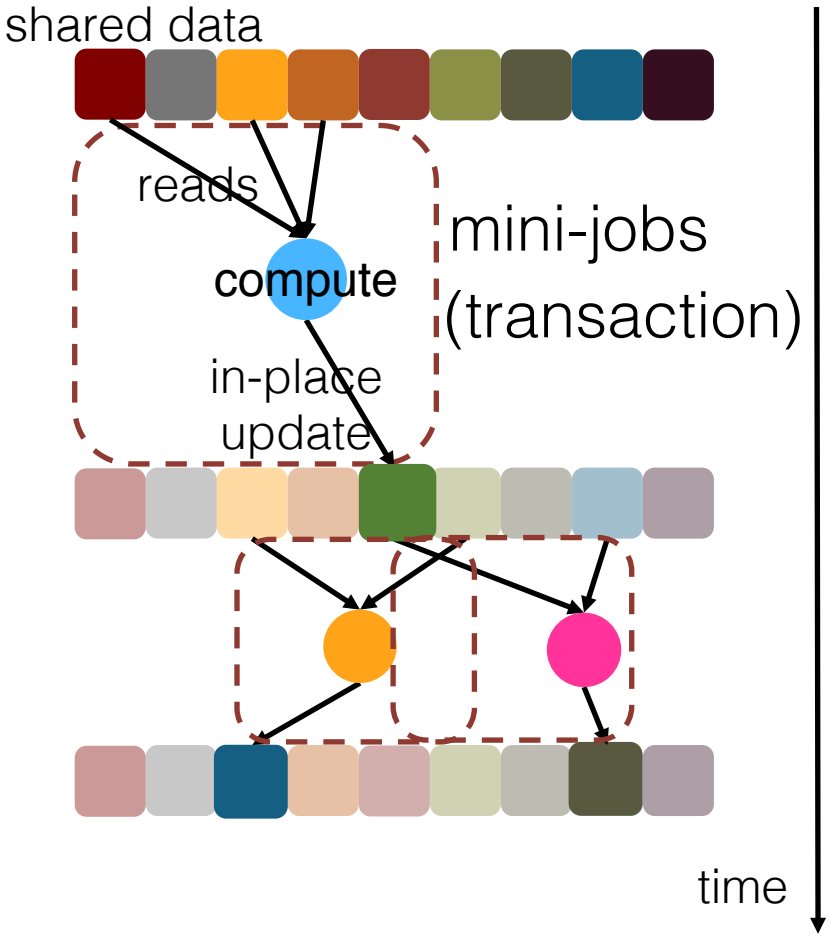

- Fine-grained mini-jobs
  - Hard to batch
- Low-latency in-place updates
- Hard to partition the data space
- Applications
  - Machine learning (SGD and others)
  - Graph computing (Vertex-centric systems)
  - Streaming (S-Store)

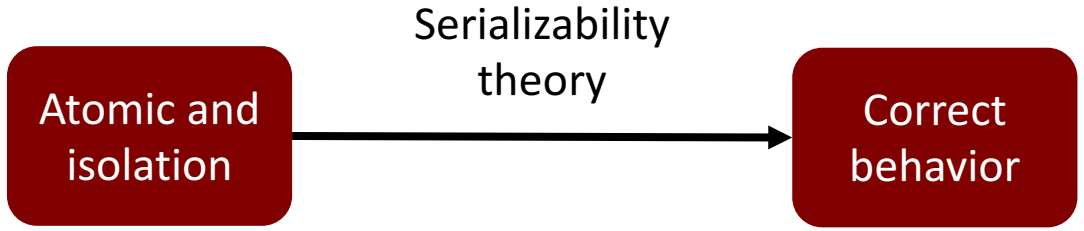


SCALABILITY
LATENCY
THROUGHPUT
DATA
CONSISTENCY
&
JOB ISOLATION
© Hank Perry/Solent News

DO WE NEED IT?

- Approach: remove data consistency controller
  - Pros: super-fast, yeeeeh!
  - Cons: could cause data consistency issues
- HogWild! & Parameter Server & others
  - Correctness proofs rely on special properties
    - Convexity
    - Lipschitz-continuity
    - Bounded staleness
- PBS: Probabilistic Bounded Staleness
  - Weak consistency actually provides strong semantics
  - Single key only
  - Probabilistic

THE DATABASE WAY

- Fewer assumptions, more applications
  - Non-convex (deep learning)
  - Discrete & combinational (graph problems)

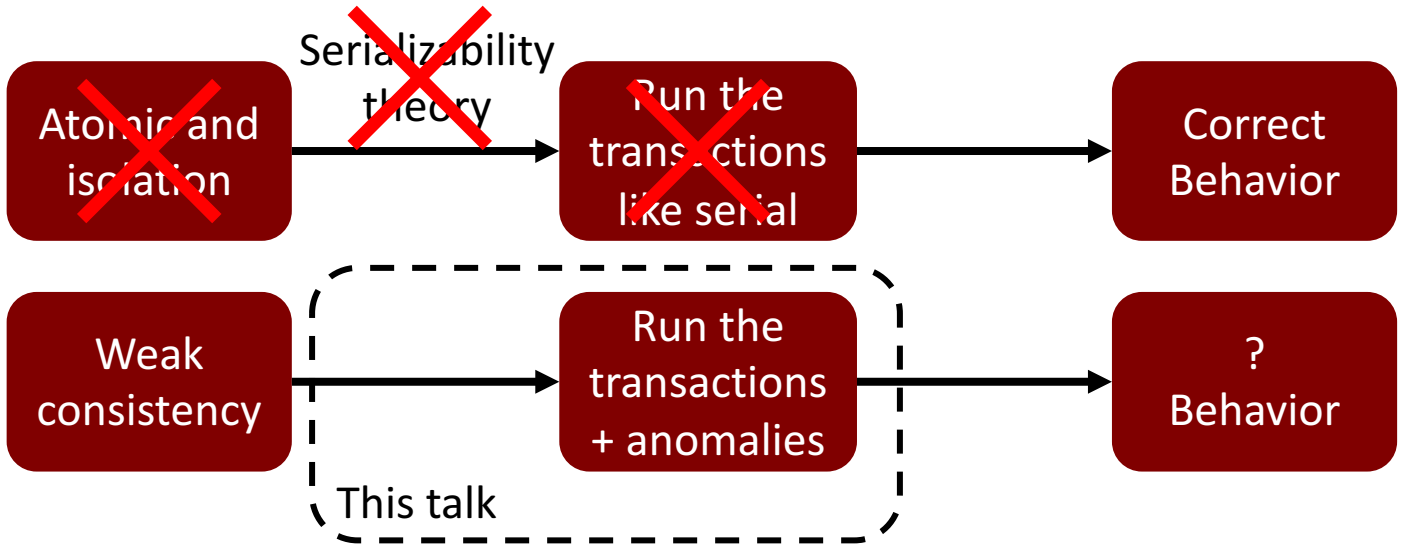

DATA CONFLICT GRAPH

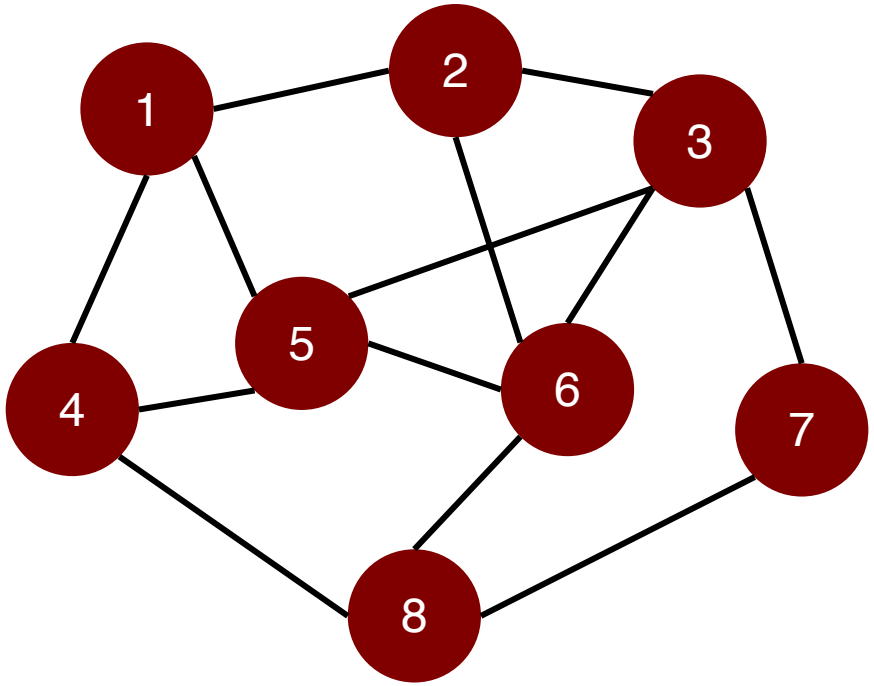

- Each vertex represents a txn
- An edge if two txns share data
  - Potential conflicts

GOOD AND BAD

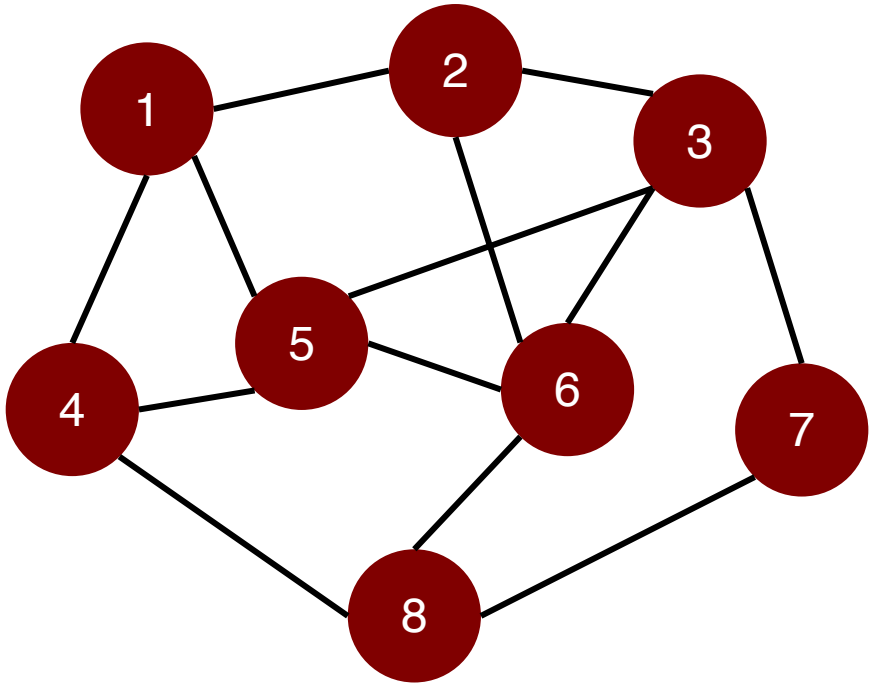

- Good No. 1: serial execution

GOOD AND BAD

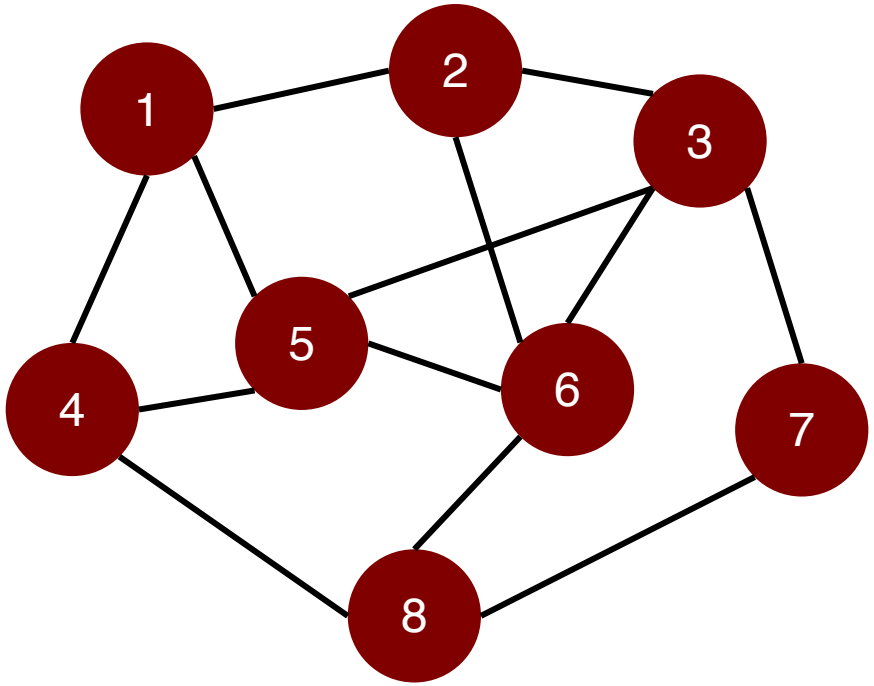

- Good No. 2: a nice scheduler
  - No direct edge in concurrent txns

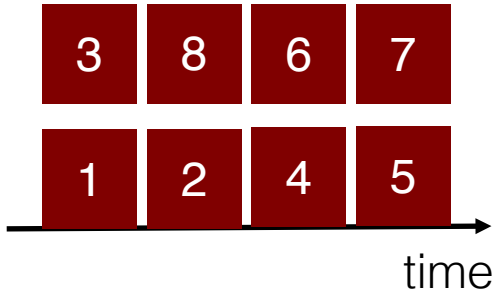

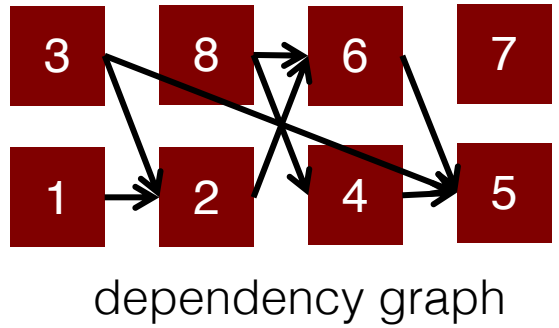

GOOD AND BAD

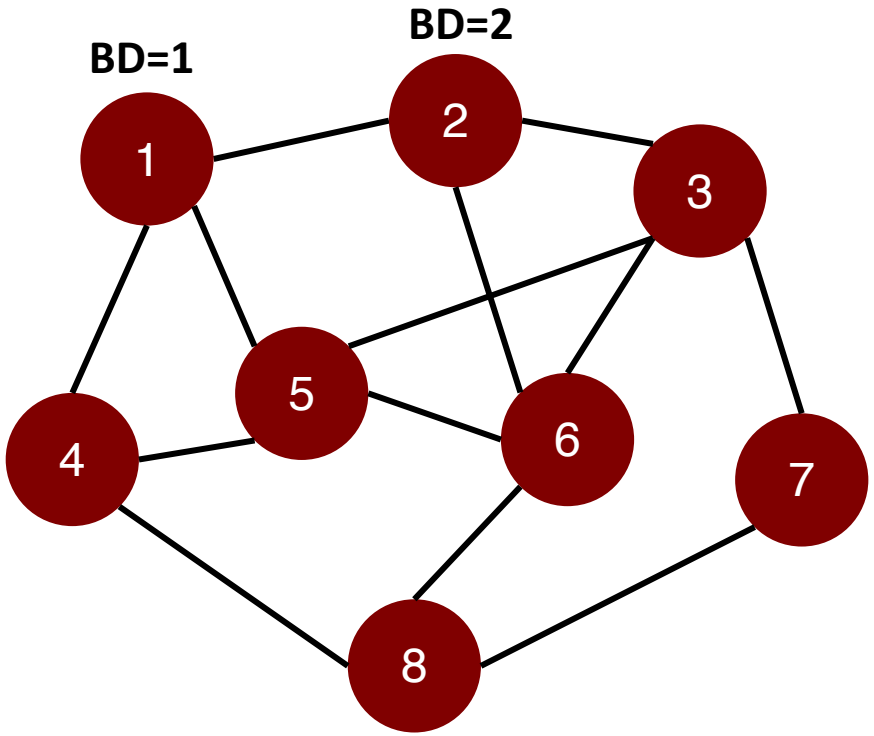

- Bad: potential conflict
- Bad degree (for a transaction)
  - # of potential conflict transactions
    - Concurrent
    - Share same data (adjacent in graph)

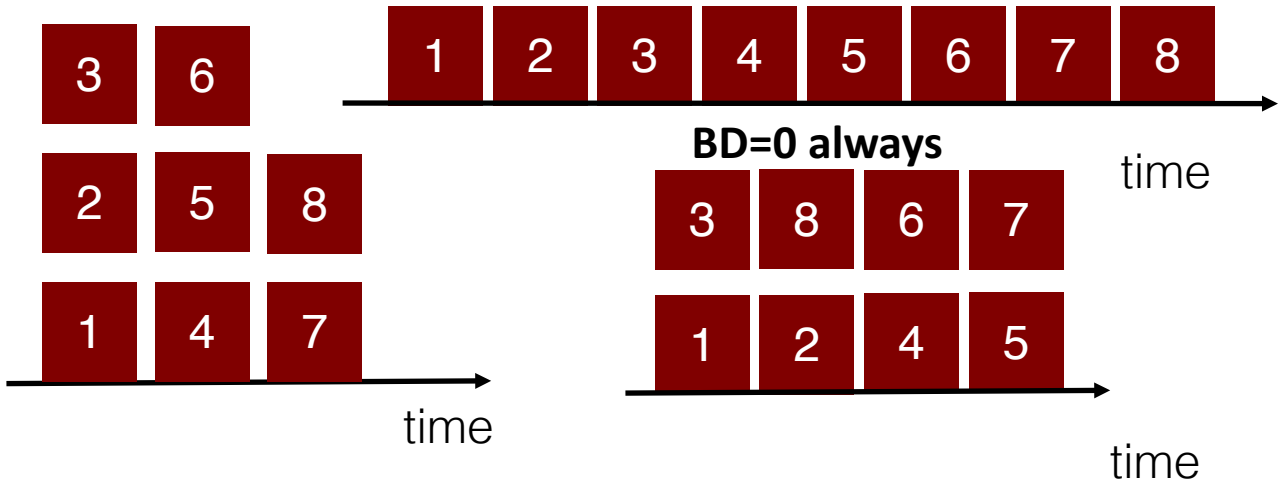

BAD DEGREE AND CORRECTNESSES

MAX BAD DEGREE	CONCURRENCY CONTROL	TXN SEMANTICS	RESULTS ACCURACY
0	NO	SERIALIZABILITY	CORRECT
>0	NO	NO	DON'T KNOW
	YES	SERIALIZABILITY	CORRECT

BAD THINGS DO NOT COME IN 3 (BN3)

- BN3: bad degree ≤1 for all transactions

MAX BAD DEGREE	CONCURRENCY CONTROL	TXN SEMANTICS	RESULTS ACCURACY
0	NO	SERIALIZABILITY	CORRECT
1 (BN3)			
>1	NO	NO	DON'T KNOW
	YES	SERIALIZABILITY	CORRECT

IS BN3 TRUE?

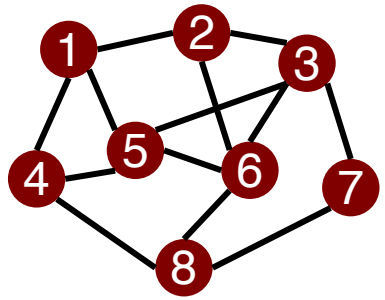

- Depends on
  - Data conflict severity: the density of data conflict graph $\frac{|E|}{\binom{|V|}{2}}$
  - Job type
  - Access pattern

	GRAPH NAME	IVI (in 10^6)	IEI (in 10^6)	DENSITY (in 10^{-4})
Web Graphs	*uk-2007-05*	106	3,739	**4.2**
	uk-2014	787	47,614	**4.7**
	eu-2015	1,070	91,792	**5.8**
	claw-2012	3,563	128,736	**1.4**
Social Networks	*wise*	59	265	**4.0**
	friendster	66	1,806	**0.7**
TPC-C	*New Order*			**>1000**

BAD DEGREE DISTRIBUTION

0BD (bad degree = 0)

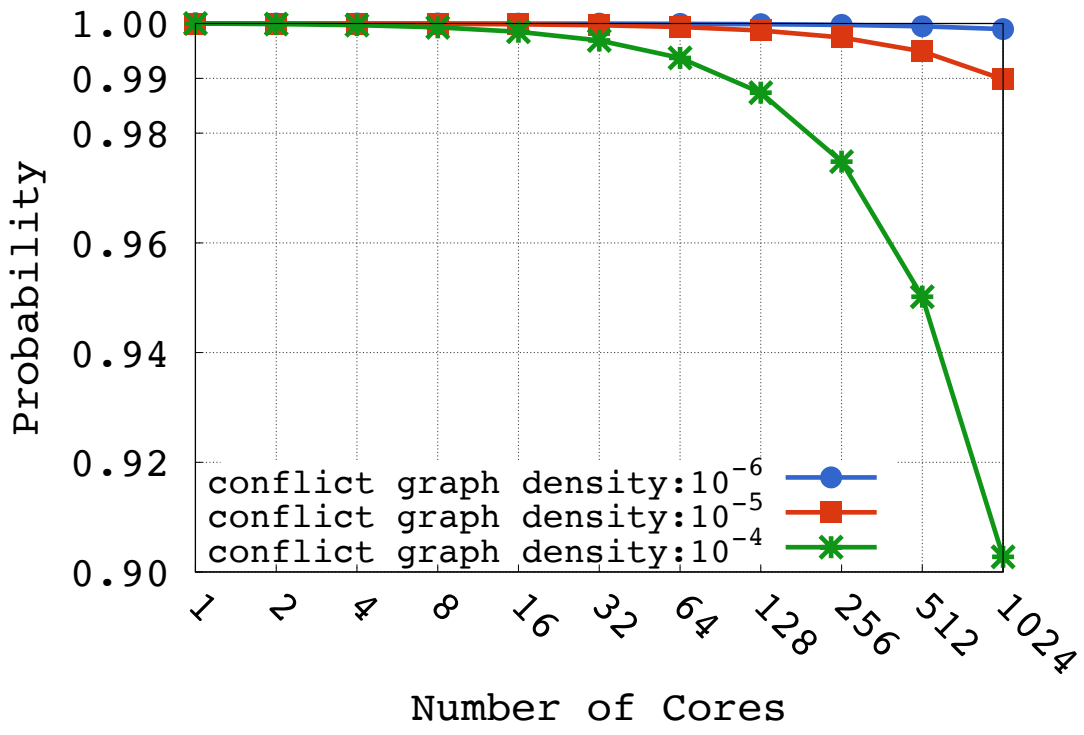

BN3 (bad degree ≦ 1)

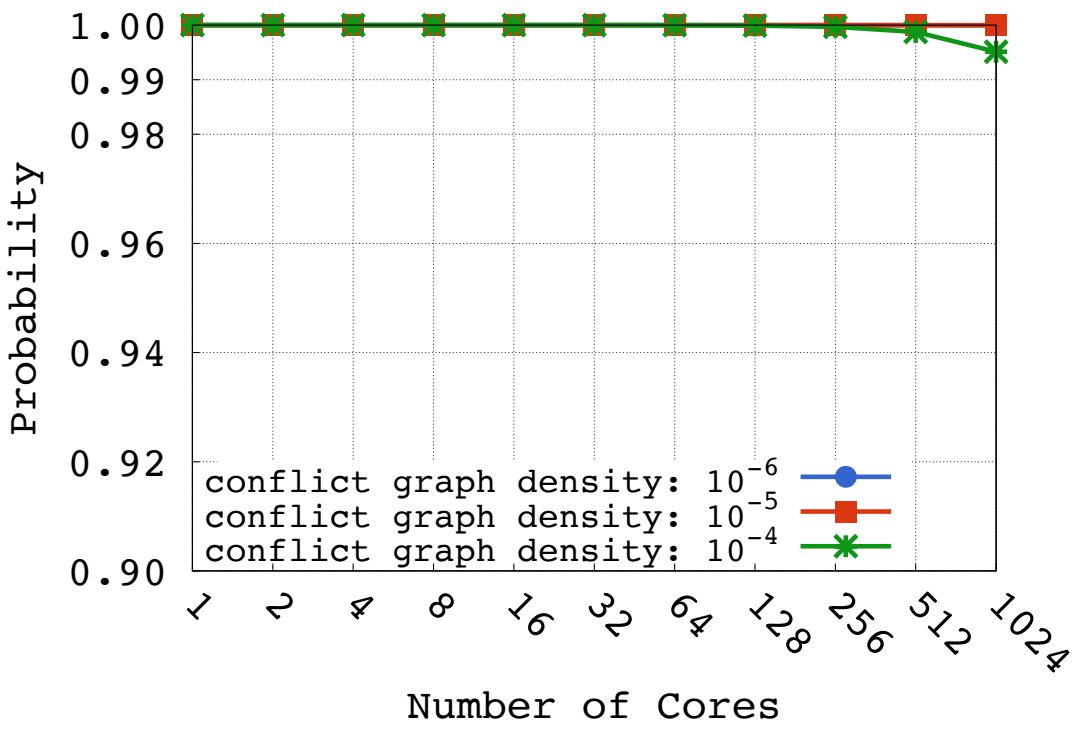

THE UNIVERSITY OF CHICAGO 香港中文大學 The Chinese University of Hong Kong

WHAT GOOD IS BN3?

THE TRANSACTIONS EXECUTED
WITHOUT ANY CONSISTENCY MECHANISM
IS UNDER
SNAPSHOT ISOLATION (SI)

PROOF: A TWO-STEP APPROACH

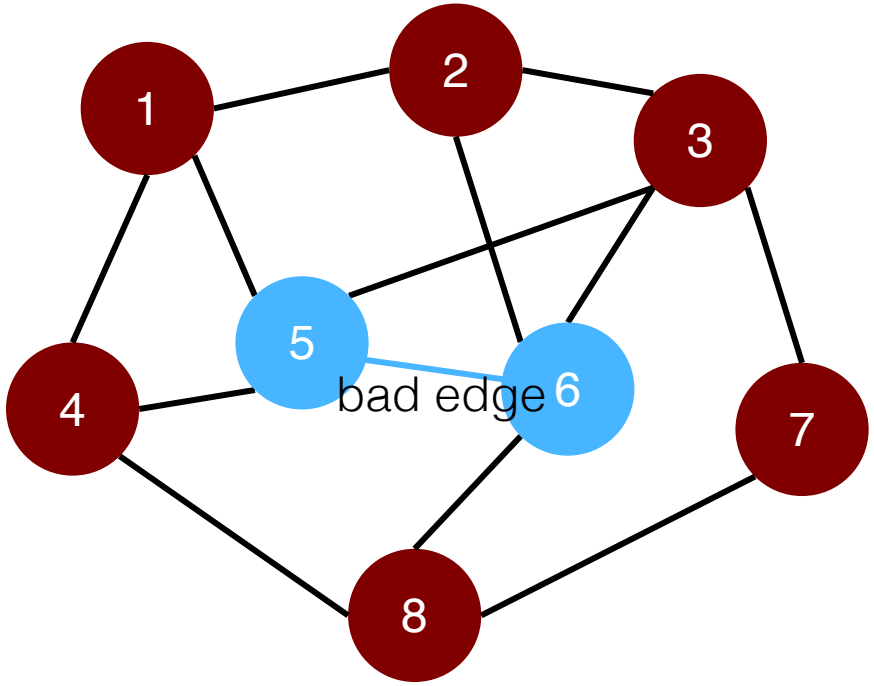

0. BN3 restricts the size of “mafia”
   - Two crews (vertices) at most
1. Only two bad transactions case
   - Proof by enumerating the type of edges
2. Other good transactions
   - Does not cause more cycles
     - Adjacent (non-bad) vertices: behind or after
     - Non-adjacent vertices: none of their business

BAD DEGREE AND CORRECTNESSES

MAX BAD DEGREE	CONCURRENCY CONTROL	TXN SEMANTICS	RESULTS ACCURACY
0	NO	SERIALIZABILITY	CORRECT
1 (BN3)	NO	SNAPSHOT ISOLATION	WRITE-SKEW
ANY	NO	NO	DON'T KNOW
	YES	SERIALIZABILITY	CORRECT

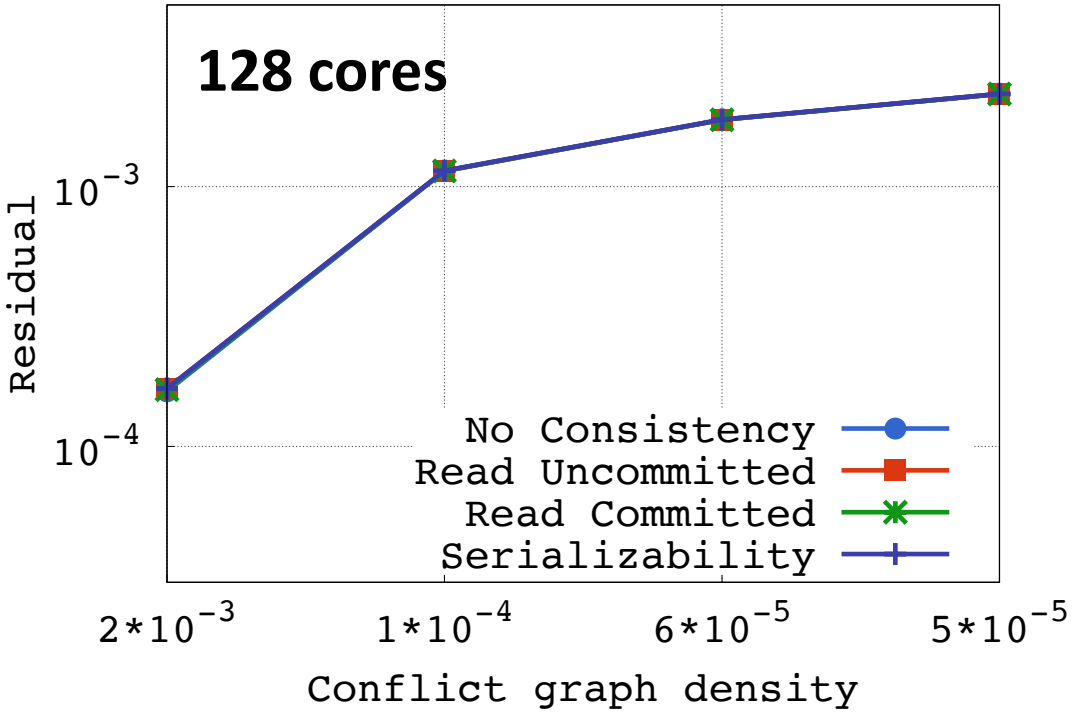

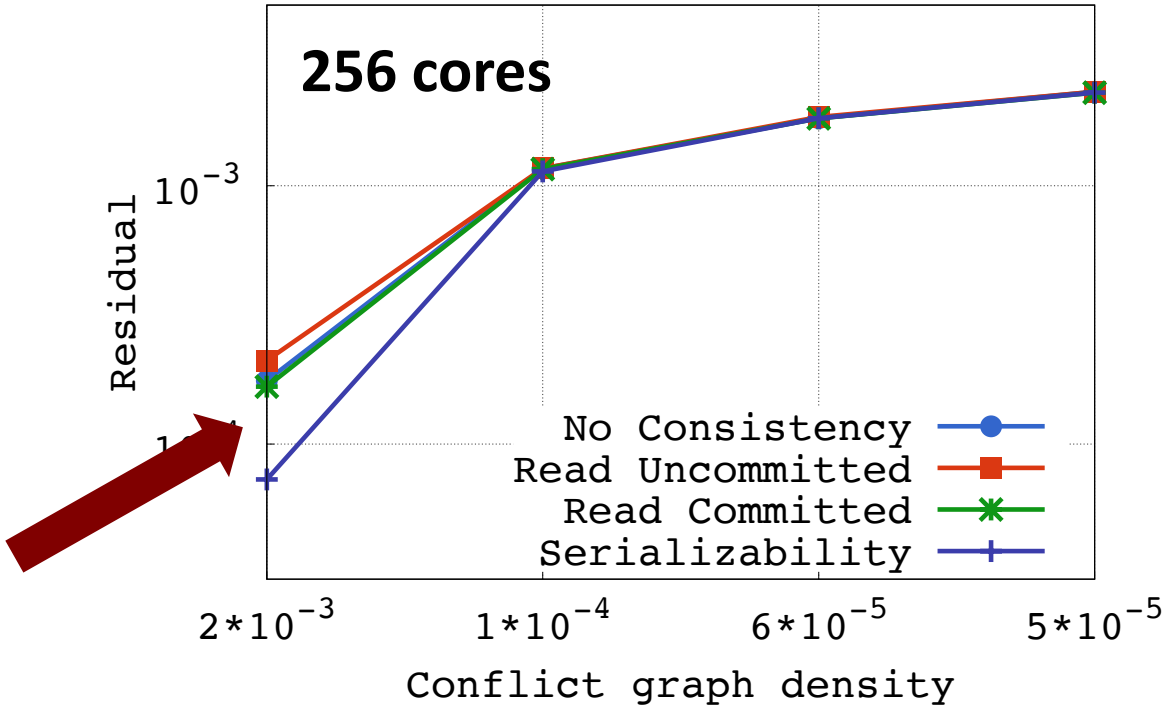

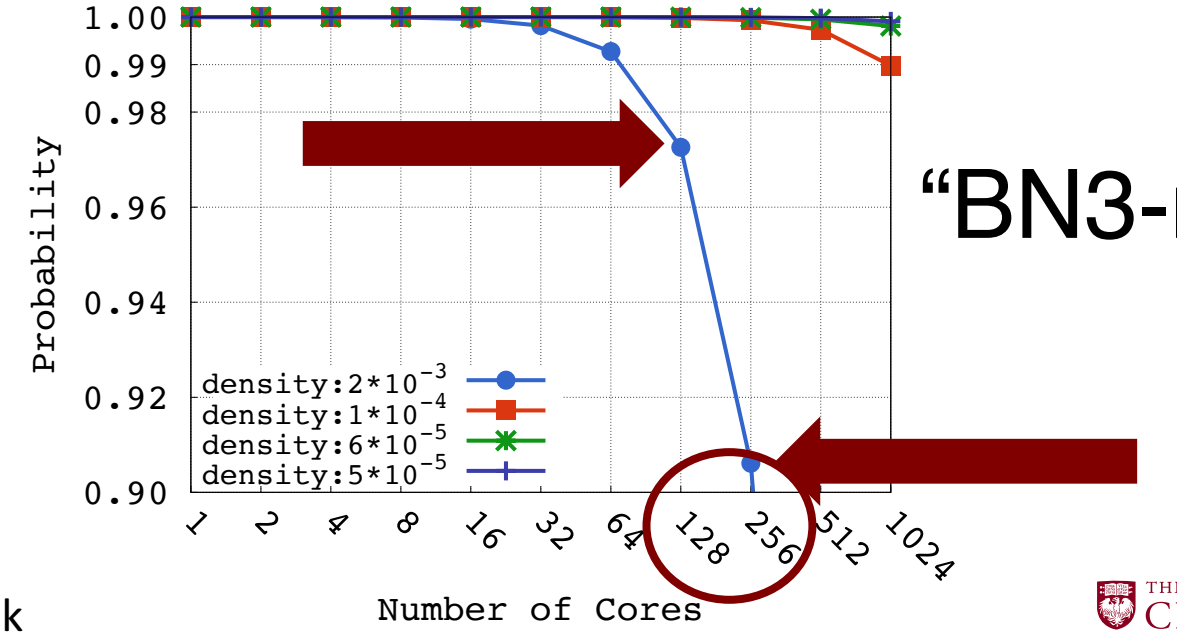

"BN3-ness"

x: vary the conflict graph density
lines: vary isolation levels
y: residual after 50 iterations of Page Rank

THE UNIVERSITY OF CHICAGO
香港中文大學 The Chinese University of Hong Kong

TAKE HOME MESSAGES

- Life is not just all-or-nothing
- Flawlessness costs a lot
- It is possible to have almost everything for free
- BN3: realistic assumption, practical conclusion
- Some future works
  - Runtime: monitor the BN3-ness
  - BN3 as a new consistency level
  - Mixed concurrency control

Thank you

EXPERIMENTAL STUDIES (THROUGHPUT)