Locality Analysis Through IBM POWER8 Hardware Sampling

Chencheng Ye
yechencheng@gmail.com

Zhizhou Zhang
zzhang72@u.rochester.edu

Ning Gu
ngu3@cs.rochester.edu

Chen Ding
cding@cs.rochester.edu

Full Trace vs. Sampled Trace

- Modern processors like IBM POWER 8 could sample the cache miss on the fly
- The overhead of sampling is low compared to collecting a full trace

Pro-grams	Ratio of full exe. time						
	1/1	1/10	1/100	1/1K	1/5K	1/10K	1/15K
BT	1.190	1.173	1.140	1.093	1.032	1.012	1.004
CG	1.168	1.193	1.176	1.151	1.151	1.108	1.009
FT	1.204	1.200	1.178	1.110	1.095	1.104	1.080
IS	1.111	1.086	1.042	1.035	1.033	1.040	1.053
LU	1.200	1.212	1.190	1.115	1.065	1.039	1.029
MG	1.205	1.204	1.125	1.098	1.051	1.050	1.053
SP	1.190	1.170	1.118	1.073	1.038	1.027	1.016
UA	1.208	1.172	1.091	1.048	1.025	1.018	1.015

Reuse Distance vs. Footprint

- Conventionally, people use reuse distance to analyze locality. A reuse involved two accesses, all accesses between them would be monitored
- Footprint is a higher order metric could be calculated in linear-time, amendable to sampling

Sampling Methods

- Reuse-based sampling is effective but costly to implement
- Miss-based sampling can be done with a negligible cost at run time. However, the downside is the loss of the ability to sample data reuse

Motivation

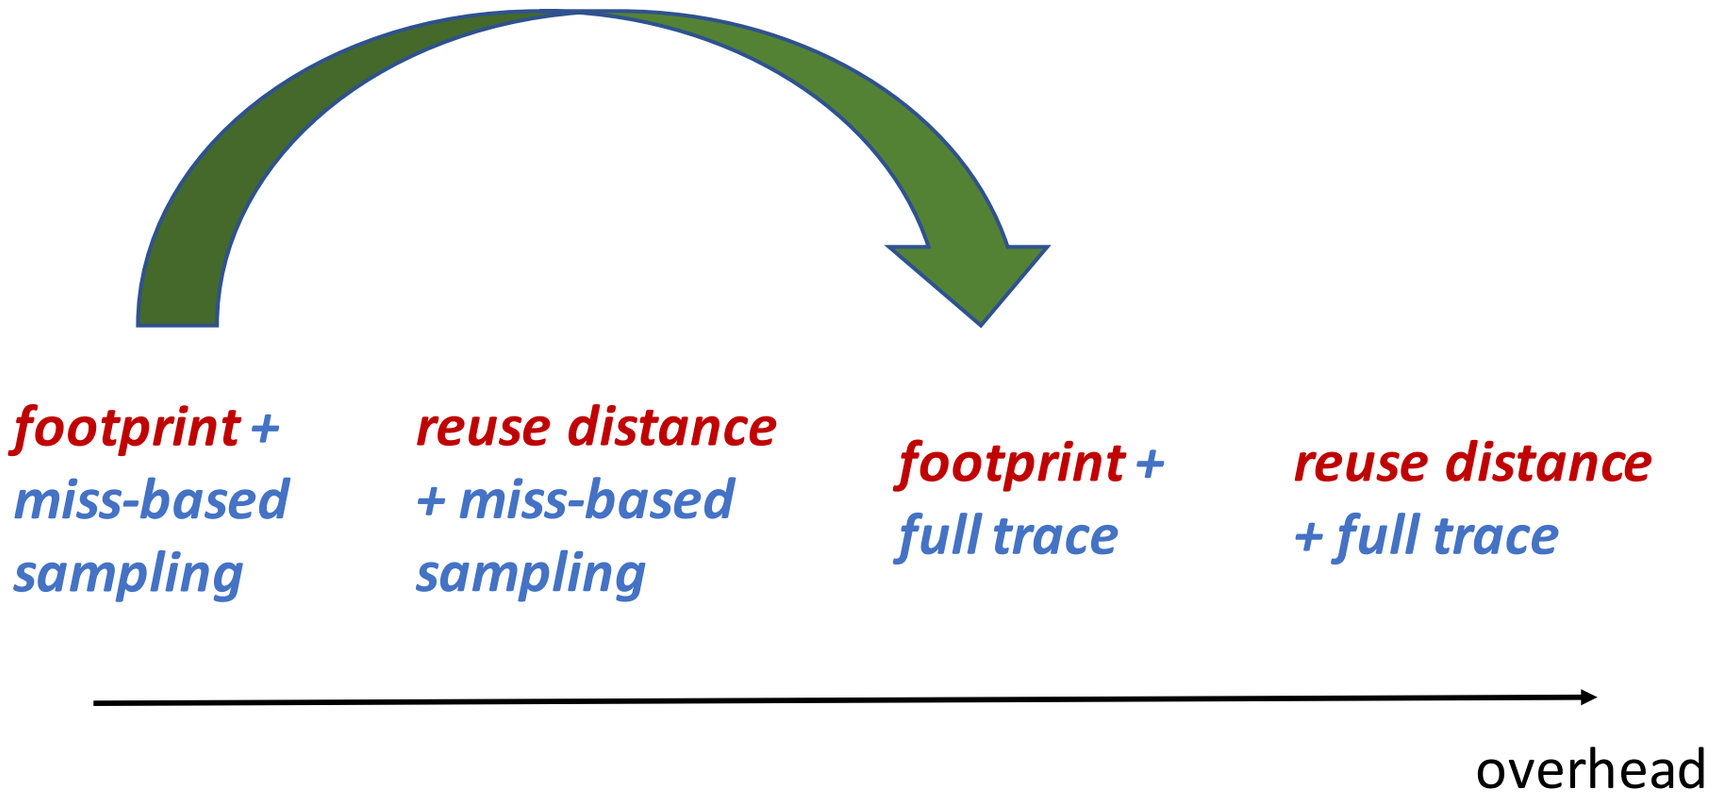
footprint + miss-based sampling
reuse distance + miss-based sampling
footprint + full trace
reuse distance + full trace
overhead

Ultimate Goal and Current work

- Predicting footprint, miss ratio curve, data size of full trace/high sampling ratio sampled trace, based on low sampling ratio sampled trace
- Data size prediction is the first step

Data Size Prediction

- Data size of a trace is defined as the number of distinct data blocks referenced in the trace. The data size is important since it defines the memory space needed
- One effect of sampling is that there is a chance for a data block not to be sampled. Therefore, the data size of a sampled trace is likely less than the data size of the full execution

Footprint and Data Size

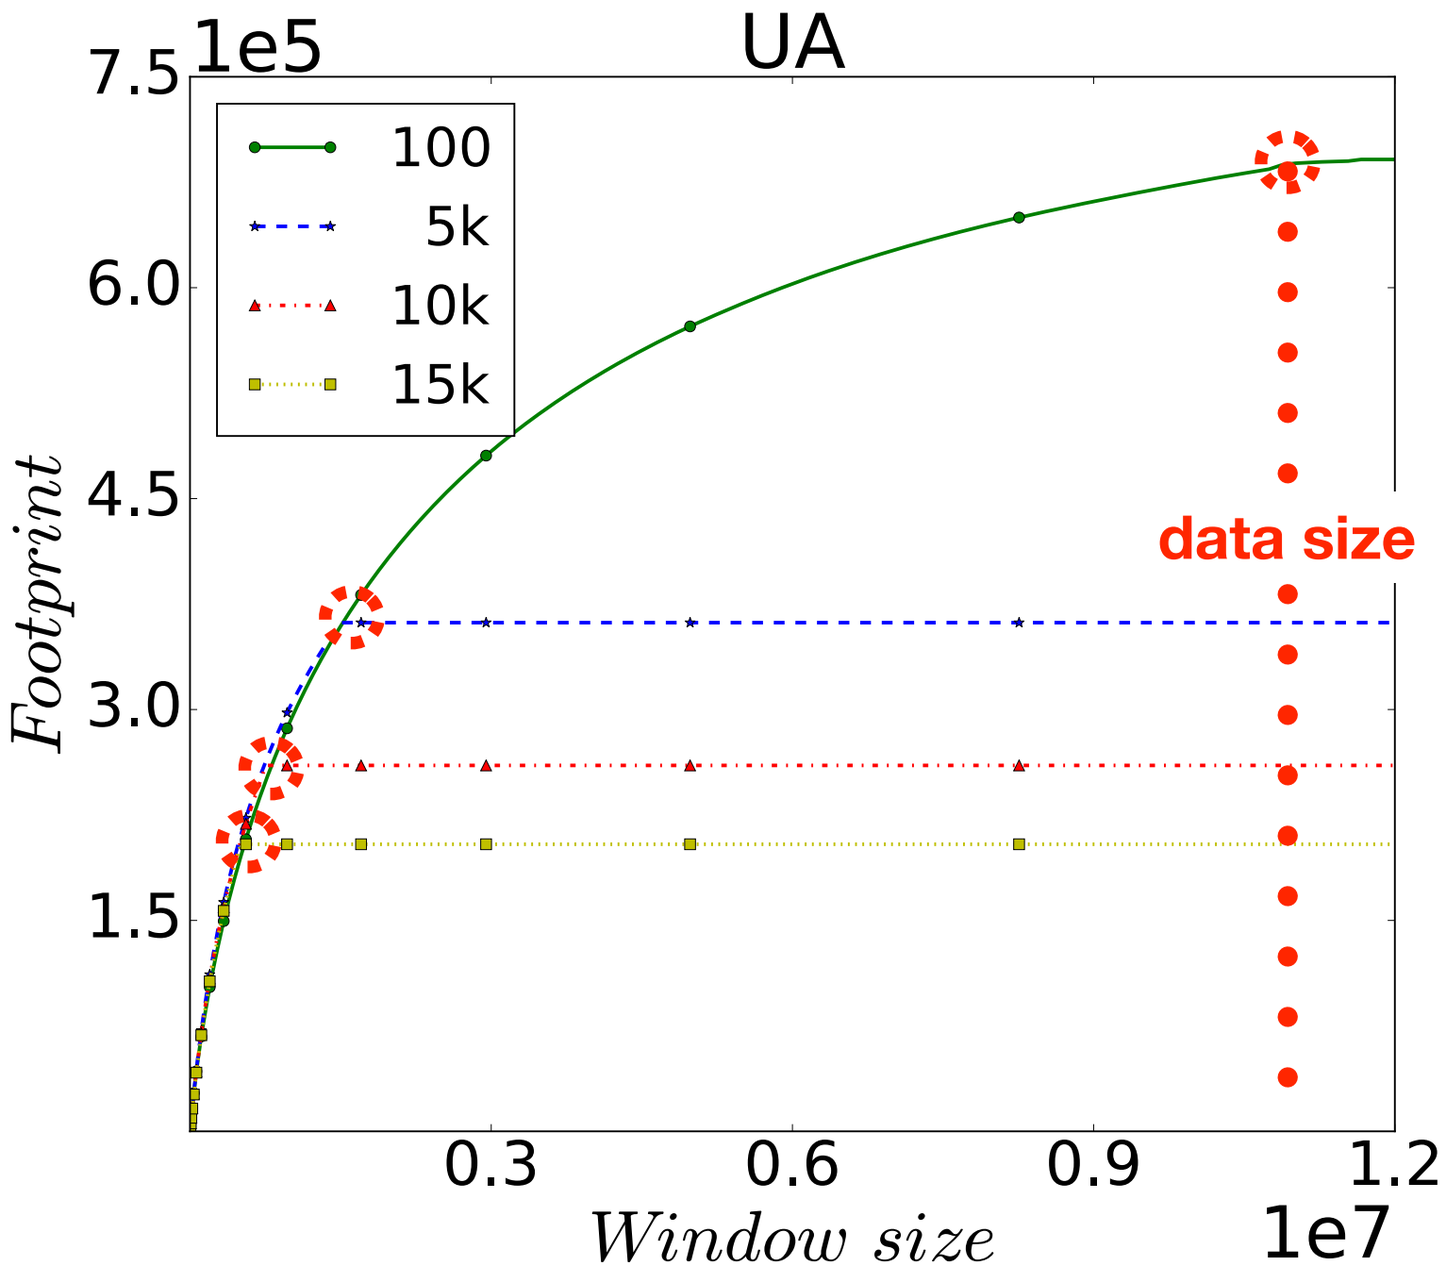

Data Size Prediction

- **Down-scaling**
  - Given a full trace, calculate data size of sampled trace based on the sampling ratio
- **Up-scaling**
  - Solve the reverse issue

Down-scaling

- The Probability of a data block sampled at least once:

$$P(r, f) = 1 - (1 - \frac{1}{r})^f$$

  - r: sampling frequency
  - f: The number of times referenced in the execution
- The expect data size of sampled trace for a block：

$$P(r, f) \times C(f)$$

  - $C(f)$: the number of data blocks referred exact f times

Down-scaling

- The total data size of a sampled trace:

$$D(r) = \sum_{f=1}^{n} P(r, f) \times C(f)$$

Up-scaling

- We approximate the various values $P(r, f)$ with a constant value $P(r, \bar{f})$

$$D(r) = \sum_{f=1}^{n} P(r, \bar{f}) \times C(f) = P(r, \bar{f}) \times \sum_{f=1}^{n} C(f)$$

- There is not enough information to find best $\bar{f}$, and we requires two different sampling ratios

$$\frac{D(r_0)}{D(r_1)} = \frac{P(r_0, \bar{f})}{P(r_1, \bar{f})}$$

Up-scaling

- We exhaustively check all possible values and take the best one according to measured value. we iterate from 0.01 to 10000 at step 0.01

Sample Loss

- Using *perf* collects samples

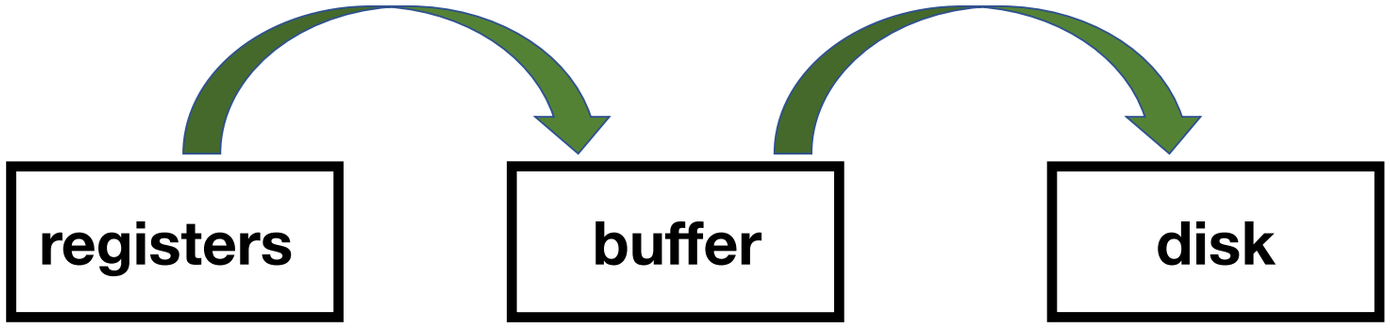

- Why sample loss?
  - The buffer would be overwritten while it is full
  - All cores are busy
  - Available bandwidth between the memory and the disk is low
  - etc.

Requested & Actual Sampling Ratio

Programs	Requested sampling ratio (10^{-5})						
	100,000 (1/1)	10,000 (1/10)	1,000 (1/100)	100 (1/1K)	20 (1/5K)	10 (1/10K)	6.7 (1/15K)
	Actual sampling ratio (10^{-5})						
BT	104.7	104.7	86.3	51.1	21.1	10.5	6.9
CG	128.4	120.0	60.4	54.3	20.1	10.0	6.7
FT	31.8	33.3	26.6	17.7	11.0	10.2	6.8
IS	361.4	318.3	105.2	35.4	19.1	10.0	6.7
LU	78.0	76.1	58.4	38.4	20.2	10.0	6.6
MG	301.6	310.6	149.1	87.3	20.6	10.1	6.7
SP	119.7	118.4	78.7	44.3	19.4	10.0	6.6
UA	403.6	332.2	145.4	72.9	19.7	10.0	6.7

Sampled Data Size on Power8

BT	1/5K	1/10k	1/15K
Data Size	1.30	1.16	1
Sampling Ratio	3	1.5	1

Pro-grams	Full Execution Data Size(MB)	Ratio of Full Data Size		
		1/5K	1/10K	1/15K
BT	172.06	0.519	0.422	0.363
CG	184.13	0.521	0.314	0.243
FT	1283.06	0.078	0.067	0.044
IS	264.17			
LU	147.60			
MG	431.87			
SP	181.35	0.319	0.231	0.153
UA	121.94	0.359	0.327	0.201

CG	1/5K	1/10k	1/15K
Data Size	2.14	1.29	1
Sampling Ratio	3	1.5	1

Data Size Variation

- Taking the maximum and the minimal as boundary, the ratio is NOT greater than 1.1 as stable
- A program is stable if it is stable under all three sampling ratios
- The stability is not monotonic to the sampling ratio due to BT, MG and UA are not stable only at the median sampling ratio

Program	Sampling Ratio	Ratio of full exe.		
		Min.	Median	Max.
BT	**1/1K**	**0.687**	**0.701**	**0.746**
	1/5K	0.502	0.518	0.553
	1/15K	**0.357**	**0.364**	**0.367**
CG	**1/1K**	**0.740**	**0.751**	**0.799**
	1/5K	**0.512**	**0.522**	**0.539**
	1/15K	**0.241**	**0.243**	**0.245**
FT	**1/1K**	**0.096**	**0.099**	**0.100**
	1/5K	**0.073**	**0.074**	**0.075**
	1/15K	**0.043**	**0.045**	**0.047**
IS	1/1K	0.033	0.035	0.050
	1/5K	0.019	0.021	0.024
	1/15K	0.005	0.006	0.007
LU	1/1K	0.759	0.796	0.835
	1/5K	0.529	0.551	0.604
	1/15K	0.297	0.322	0.341
MG	**1/1K**	**0.113**	**0.117**	**0.121**
	1/5K	0.034	0.038	0.038
	1/15K	**0.012**	**0.013**	**0.013**
SP	1/1K	0.567	0.600	0.658
	1/5K	0.296	0.322	0.341
	1/15K	**0.149**	**0.153**	**0.159**
UA	**1/1K**	**0.534**	**0.539**	**0.546**
	1/5K	0.332	0.362	0.376
	1/15K	**0.202**	**0.204**	**0.209**

Down-scaling Prediction

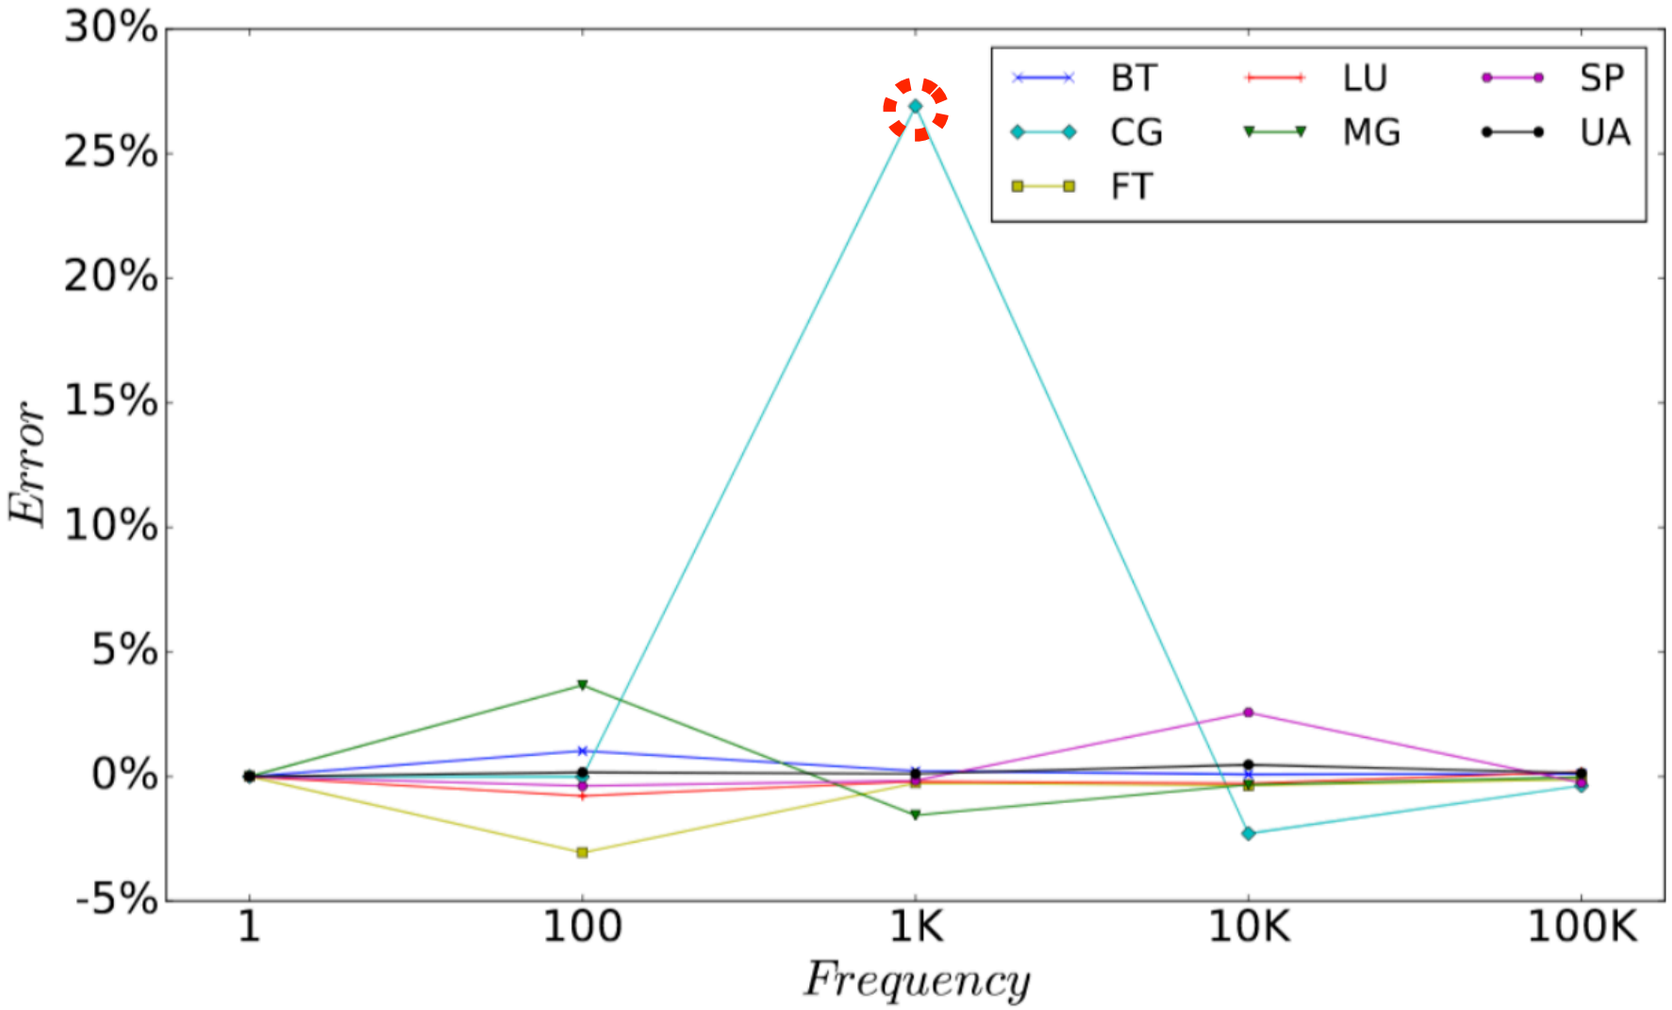

- CG at 1/1K has largest error: 26.9%
- Highest error of the rest is 3.66%, the average error is 0.71%

Up-scaling Prediction

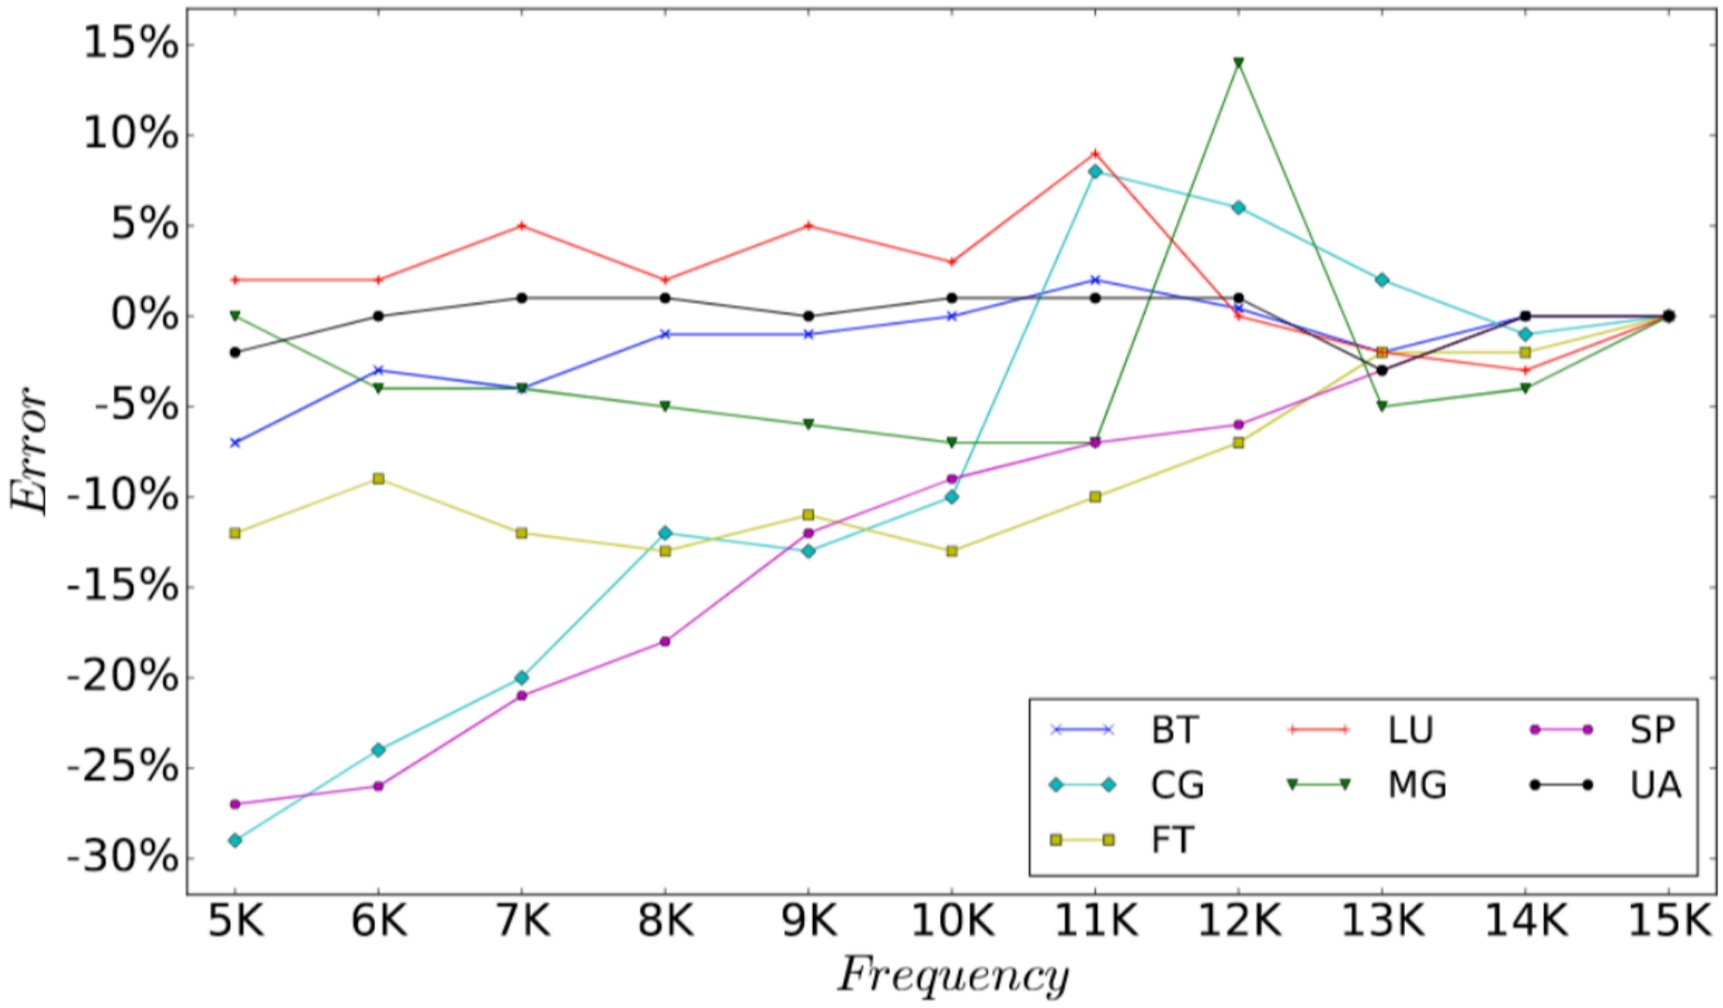

- $\frac{Prediction}{Measurement} - 1$
- IS is not shown since it has limited data blocks
- Prediction error: 3 of them within 10%. 1 of them within 20%.

Thank you for listening

Any questions?