




Go further, faster®

Paragone: What's next in block I/O trace modeling

Rukma Talwadker, Kaladhar Voruganti
Advanced Technology Group, NetApp
May 9th, 2013





Block Traces – an introduction


Timestamp	I/O Size	Offset	Operation
128166465299039000	4096	239070	Read
128166465299119000	8192	182082715136	Write
128166465299134000	16384	3154148864	Read

1. In most cases no two I/O's are identical
2. Variability exercises performance/behavioral variations on the storage systems.
3. Traces are widely used to study system behaviors for better designs.
4. Traces are difficult to store/share and manipulate.
5. Building trace based statistical models is being increasingly focused on in the storage research community.



Challenges

Model should be concise yet lossless





Why is trace modeling hard

1. I/O attributes vary per I/O. Averages are not good representatives.
2. Trace model needs to be revived over time. Attributed to change in application phases, multi-tenancy levels etc.
3. I/O parameters are not independent but are correlated and coupled with each other.
4. Ordering of I/O's is crucial.
5. Many I/O attributes to model. Dimensionality reduction is necessary.



Where do Benchmarks fail?

1. I/O sequential length preservation
 - Cannot model variable sequential lengths
 - Cannot model number of blocks to seek

2. I/O burst modeling
 - Interleaved high and low activity periods
 - Variable timing of requests during bursty and non-bursty intervals



Where do Benchmarks fail?

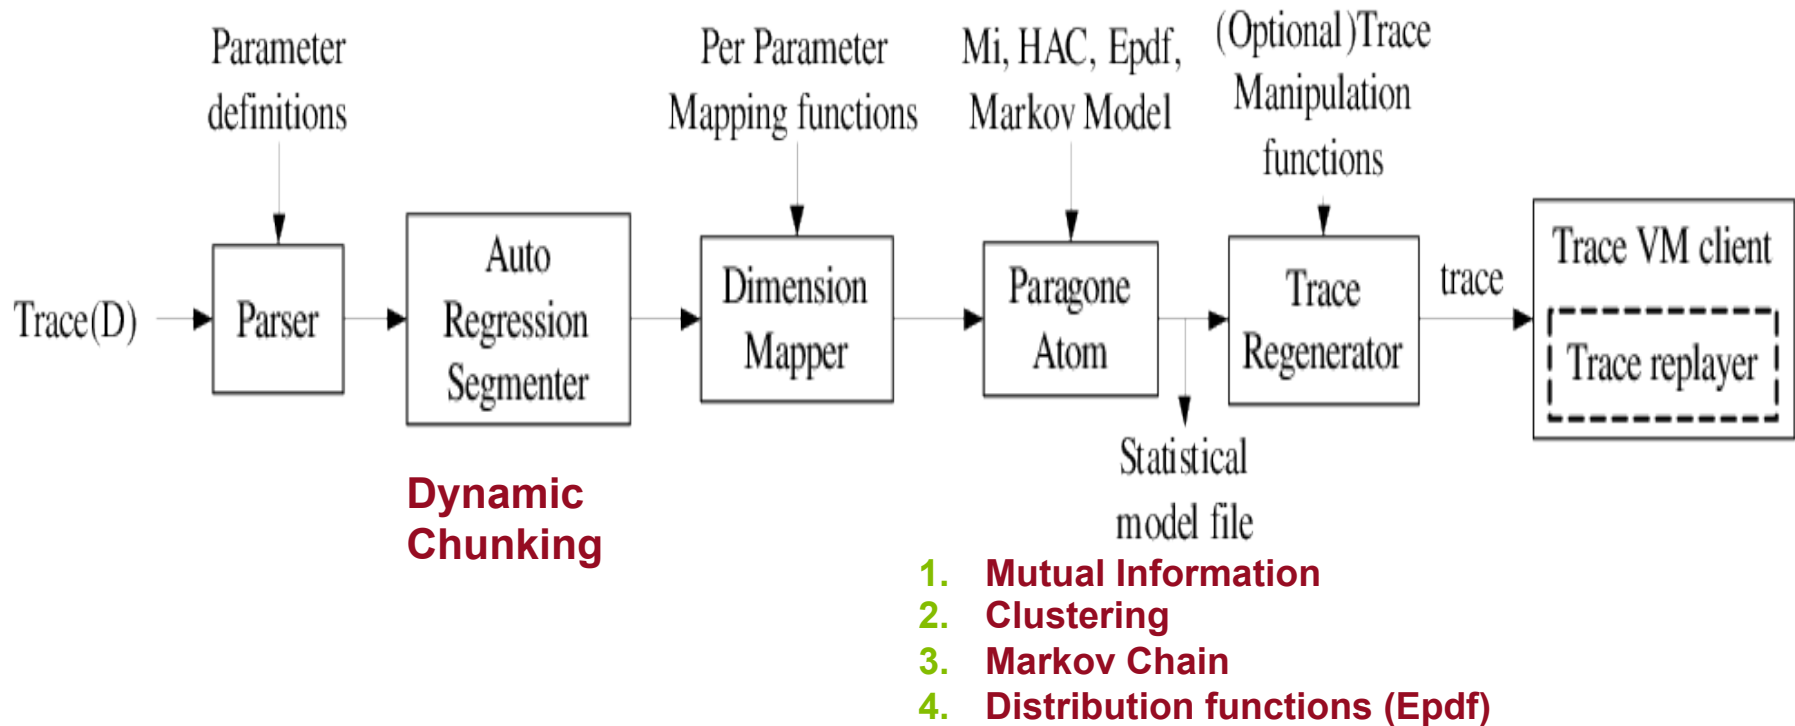
3. Scaling workloads

- Assumes workloads to be self similar
- Scales by physically creating multiple datasets.



What is Paragone

We combine known techniques in a novel way to build Paragone





Traces Used

Workload Type	duration	Trace size	Model size	% size reduction
1. MSR Project Directory	2 Weeks	1.6 GB	21 MB	98%
2. MSR Webserver	2 Weeks	276 MB	10 MB	96%
3. UMASS - Financial	12 Hours	155 MB	8 MB	95%
4. Oil workload	20 Hours	137 MB	6.3 MB	95%
5. Oil workload	1 Week	256 MB	8.1 MB	97%
6. Animation workload	20 Hours	201 MB	7.2 MB	96%
7. Animation workload	1 Week	197 MB	5 MB	97%
8. Hadoop workload	40 Hours	387 MB	9.7 MB	97%
9. Support logs workload	1.3 Weeks	256 MB	7.9 MB	96%

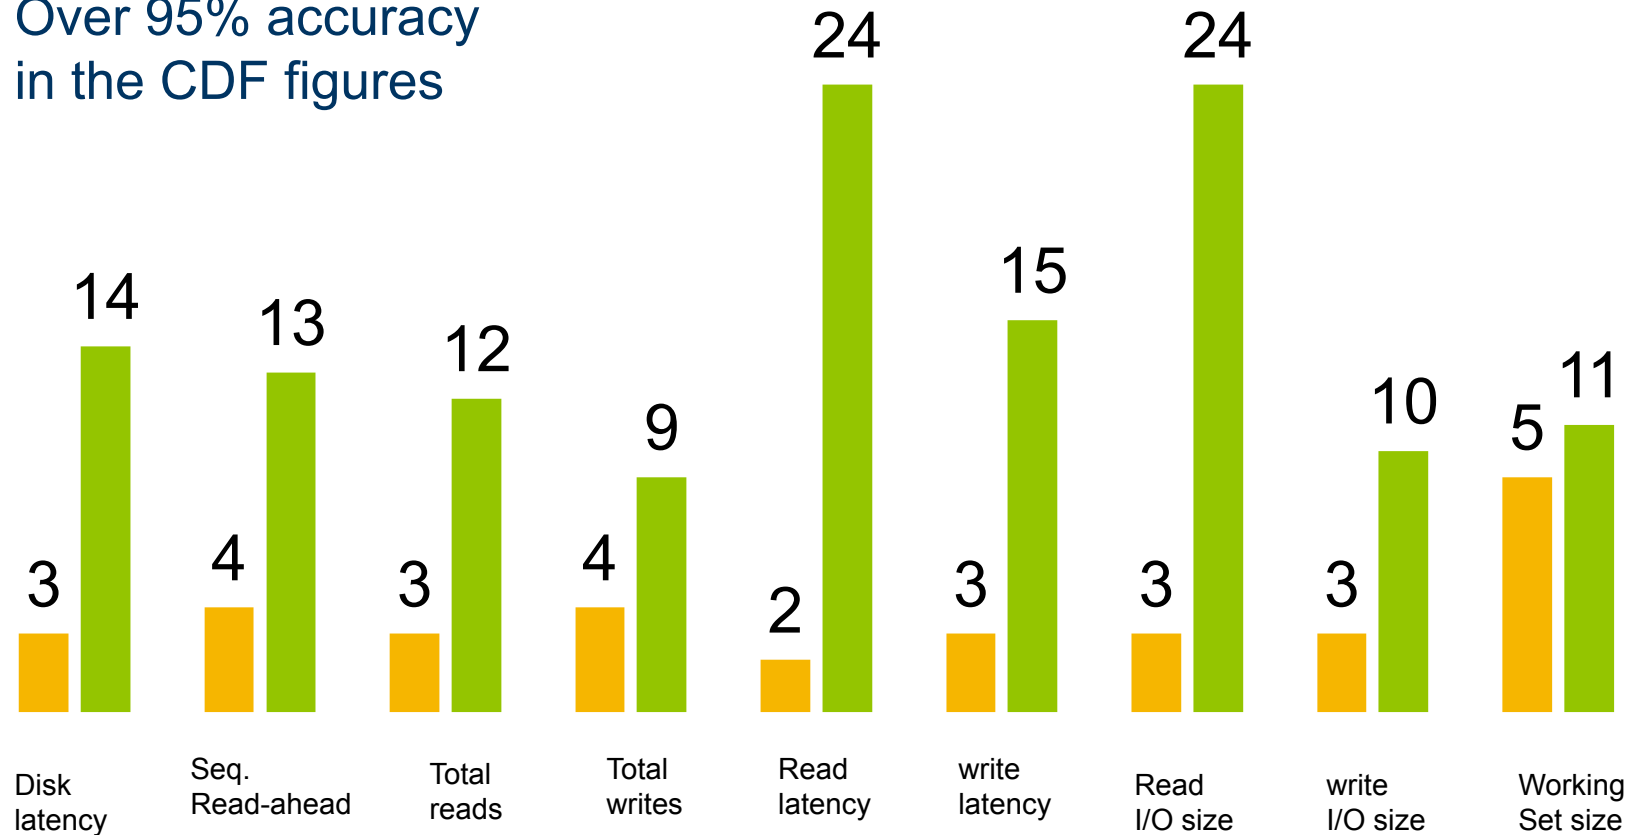


CDF errors for UMASS Finance workload

■ Paragone

■ Without Clustering+Epdf

Over 95% accuracy
in the CDF figures





Thank you

