


SDC
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

Generating Realistic Impressions for File-System Benchmarking

Nitin Agrawal



Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

SDC
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

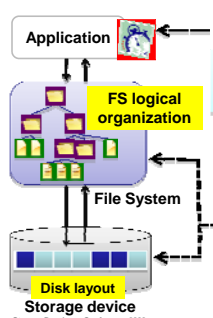
“For better or for worse, benchmarks shape a field”

David Patterson

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

SDC
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

Inputs to file-system benchmarking



Input: Postmark, FileBench, Fstress, Bonnie, IOZone, TPCC, etc etc

Input: Cold cache/warm cache

Input: Anything goes!

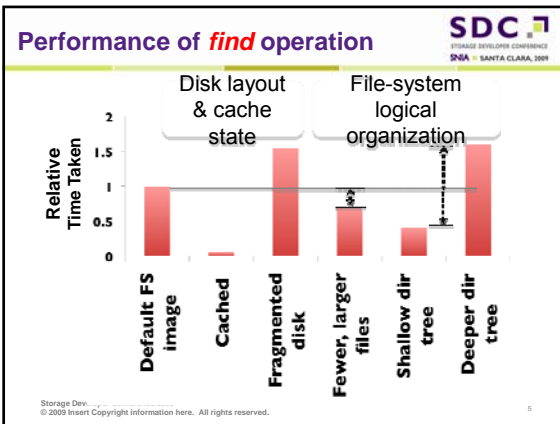
Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

FS images in past: use what is convenient

SDC
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

- Typical desktop file system w/ no description (SOSP 05)
- 5-deep tree, 5 subdirs, 10 8KB files in each (FAST 04)
- Randomly generated files of several MB (FAST 08)
- 1000 files in 10 dirs w/ random data (SOSP 03)
- 188GB and 129GB volumes in Engg dept (OSDI 99)
- 10702 files from /usr/local, size 354MB (SOSP 01)
- 1641 files, 109 dirs, 13.4 MB total size (OSDI 02)

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.



Problem scope


SDC
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

Characteristics of file-system images have strong impact on performance

We need to incorporate representative file-system images in benchmarking & design


How to create representative file-system images?

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

Requirements for creating FS images 


- Access to data on file systems and disk layout
 - Properties of file-system metadata [Satyanarayan81, Mullender84, Irlam93, Sienknecht94, Douceur99, Agrawal07]
 - Disk fragmentation [Smith97]
 - More such studies in future?
- A technique to create file-system images that is
 - **Representative**: given a set of input distributions
 - **Controllable**: supply additional user constraints
 - **Reproducible**: control & report internal parameters
 - **Easy-to-use**: for widespread adoption and

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

Introducing Impressions 

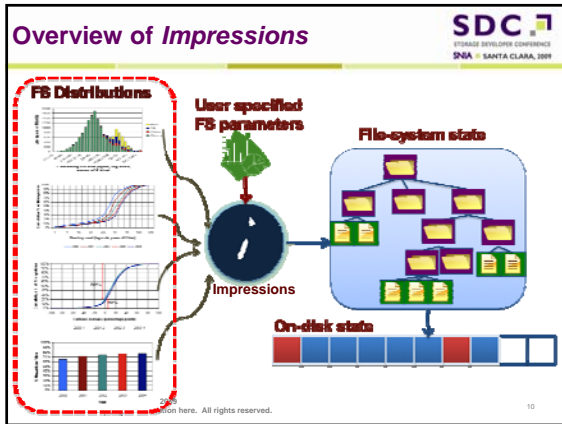
- Powerful statistical framework to generate FS images
 - Takes properties of file-system attributes as input
 - Works out underlying statistical details of the image
 - Mounted on a disk partition for real benchmarking
 - Satisfies the four design goals
- Applying Impressions gives useful insights
 - What is the impact on performance and storage size?
 - How does an application behave on a real FS image?

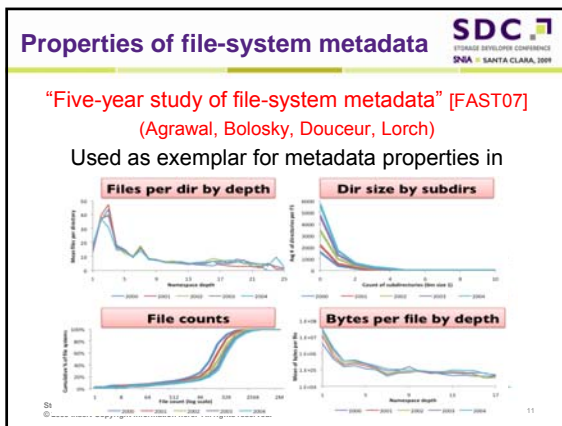
Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

Outline 

- Introduction
- **Generating realistic file-system images**
- Applying Impressions: Desktop search
- Conclusion

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.





- ### Features of Impressions
- ❑ Modes of operation for different usages
 - ❑ Basic mode: choose default settings for parameters
 - ❑ Advanced mode: several individually tunable knobs
 - ❑ Thorough statistical machinery ensures accuracy
 - ❑ Uses parameterized curve fits
 - ❑ Allows arbitrary user constraints
 - ❑ Built-in statistical tests for goodness-of-fit
 - ❑ Generates namespace, metadata, file content, and disk fragmentation using above techniques
- © 2009 Intel Corporation. All rights reserved.

Creating valid metadata

SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

File-system namespace

Generative Model processed earlier

File-system state

On-disk state

© 2009 Storage Developer Conference. All rights reserved. 13

Creating namespace

SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

Dirs by namespace

depth

Dataset

Generate

Probability of parent selection

Directory tree Monte Carlo run

Incorporates dirs by depth and dirs by subdir count

© 2009 Storage Developer Conference. All rights reserved. 14

Creating valid metadata

SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

- Creating file-system namespace
- Creating files: stepwise process
 - File size, file extension, file depth, parent directory
 - Uses statistical models & analytical approximations

Storage Developer Conference 2009
© 2009 Insert Copyright Information here. All rights reserved. 15

Example: creating realistic file sizes **SDC 7**
STORAGE DEVELOPER CONFERENCE
 SNA - SANTA CLARA, 2009

Contribution to used space

File Size
 Lognormal
 Hybrid

Containing file size (bytes, log scale)

- Pure lognormal distribution no longer good fit
- Hybrid model: lognormal body, Pareto tail
 - Fits observed data more accurately, used to recreate file sizes in Impressions

Storage Developer Conference 2009
 © 2009 Insert Copyright information here. All rights reserved. 10

Creating files **SDC 7**
STORAGE DEVELOPER CONFERENCE
 SNA - SANTA CLARA, 2009

Files by containing bytes

Fraction of bytes

File Size (bytes, log scale)

File Size Model
 Lognormal body,
 Pareto tail
Captures bimodal curve

Storage Developer Conference 2009
 © 2009 Insert Copyright information here. All rights reserved. 11

Creating files **SDC 7**
STORAGE DEVELOPER CONFERENCE
 SNA - SANTA CLARA, 2009

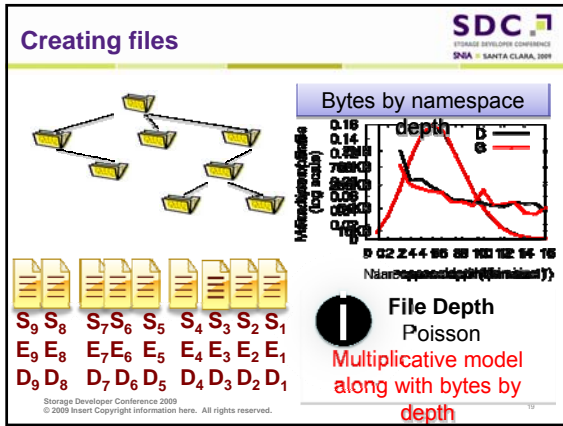
Top extensions by count

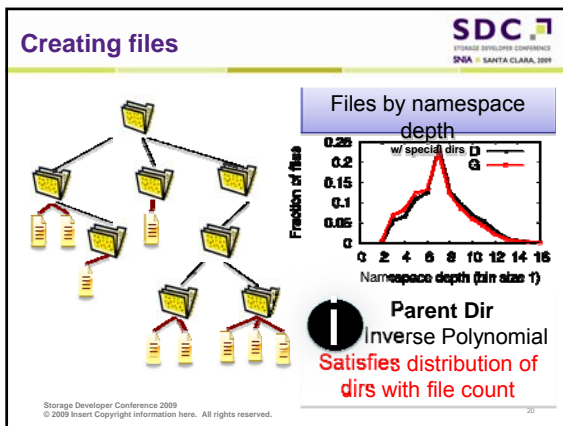
Proportion of files

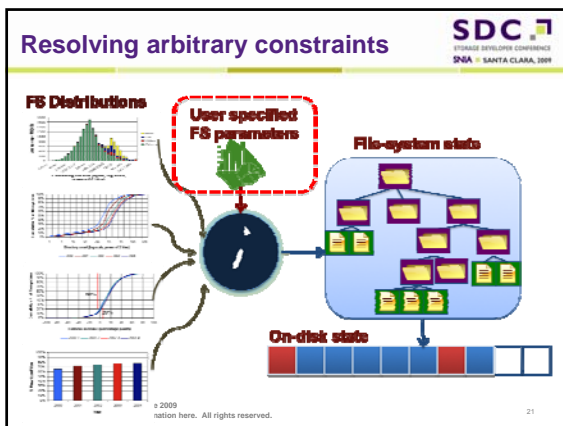
Default Generated


File Extensions
 Percentile values
Top 20 extensions account for 50% of files and bytes

Storage Developer Conference 2009
 © 2009 Insert Copyright information here. All rights reserved. 12



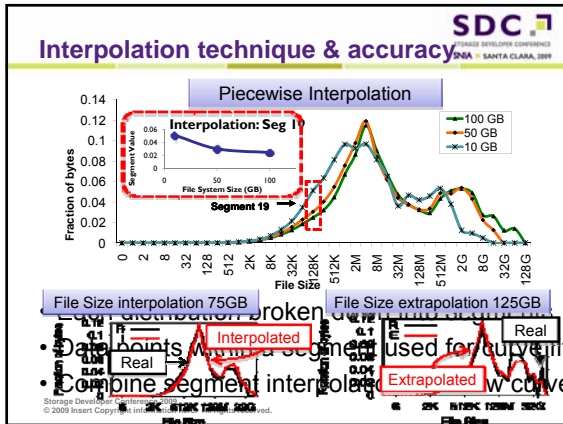




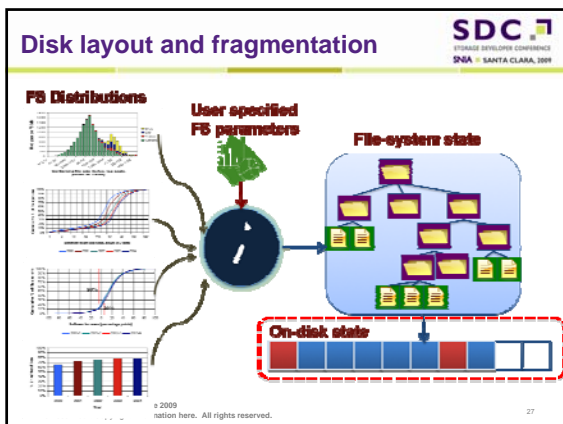


- Resolving arbitrary constraints** SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009
- ❑ Arbitrarily specified on file system parameters
 - ❑ Variant of NP-complete “Subset Sum Problem”
 - ❑ Approximation algorithm based solution (in paper)
 - ❑ **Oversampling** to get additional sample values
 - ❑ Local improvement to **iteratively converge** to the desired sum by identifying best-fit in current sample
 - ❑ While constraints are satisfied, constrained distribution also retains original characteristics
- Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.

- Interpolation and extrapolation** SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009
- ❑ Why don't we just use available data values?
 - ❑ Limited to empirical data in input dataset
 - ❑ “What-if” analysis beyond available dataset
 - ❑ Efficient to maintain compact curve fits and use interpolation/extrapolation instead of all data
 - ❑ **Technique: Piecewise interpolation**
- Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.



- File content**
- Files having natural language content
 - Word-popularity model (heavy tailed)
 - Word-length frequency model (for the long tail)
 - Content for other files (mp3, gif, mpeg etc)
 - Impressions generates valid header/footer
 - Uses third-party libraries and software
- Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved.



Disk layout and fragmentation SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

- Simplistic technique
 - **Layout Score** for degree of fragmentation [Smith97]
 - Pairs of file create/delete operations till desired layout score is achieved
- In future more nuanced ways are possible

Access to file in specific interface
 1 non-contiguous block (out of 8) All blocks contiguous
 File Layout Score = 7/8 File Layout Score = 1

- Perhaps a tool complementary to Impressions

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved. 28

Outline SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009


- Introduction
- Generating realistic file-system images
- **Applying Impressions: Desktop search**
- Conclusion

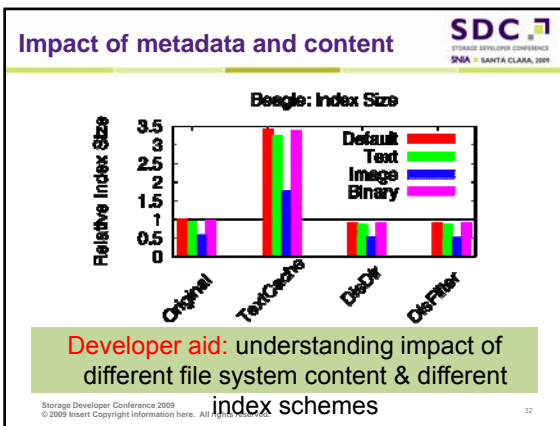
Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved. 29

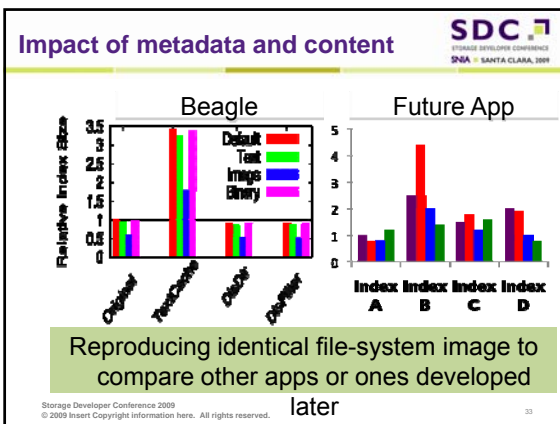
Applying Impressions SDC 7
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

- Case study: desktop search
 - Google desktop for linux (GDL) and Beagle
 - Metrics of interest:
 - Size of index, time to build initial search index
 - Identifying application bugs and policies
 - **GDL doesn't index content beyond 10-deep**
- Computing realistic rules of thumb
 - **Overhead of metadata replication?**

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved. 30







Conclusion **SDC 7**
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

- Impressions framework for realistic FS images
 - **Representative, controllable, reproducible, easy to use**
 - Includes most file system params of interest
- Extensible architecture
 - Plug in new statistical constructs, new models for metadata and content generation
- Powerful utility for file-system benchmarking
 - Available publicly

<http://www.cs.wisc.edu/adsl/Software/Impression>

Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved. 34

Questions? **SDC 7**
STORAGE DEVELOPER CONFERENCE
SNA - SANTA CLARA, 2009

Nitin Agrawal
<http://www.cs.wisc.edu/~nitina>

Impressions download
<http://www.cs.wisc.edu/adsl/Software/Impressions>



Storage Developer Conference 2009
© 2009 Insert Copyright information here. All rights reserved. 35
