LZ4r - A New Fast Compression Algorithm for High-Speed Data Storage Systems

Rui Chen (ruichen@mathworks.com)

Lihao Xu (lihao@wayne.edu)

LZ4

- The key to success of LZ4:
- - consistently provides high throughput for universal file formats
- - creates little overhead for uncompressible data files
- - good trade-off between compression ratio and speed

LZ4r Inspiration

- Based on LZ4
- Inspired by two observation during match searching:
- - lots of 0-length literals
- - offset lengths tend to be short

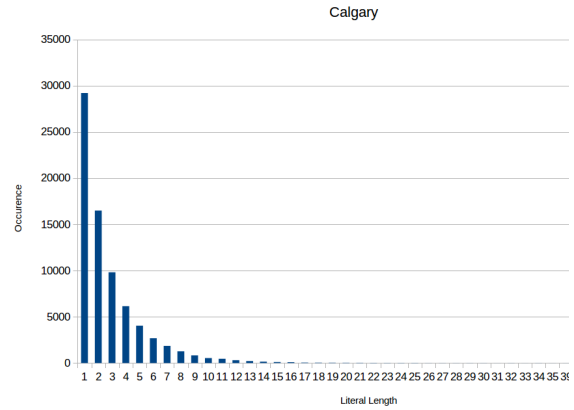

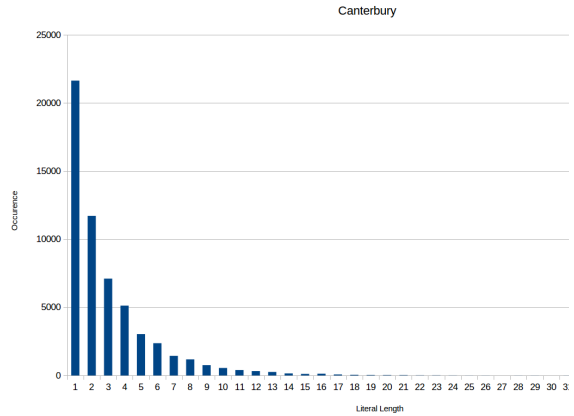

LZ4r

- Compression/decompression process is the same as LZ4
- New compression format:

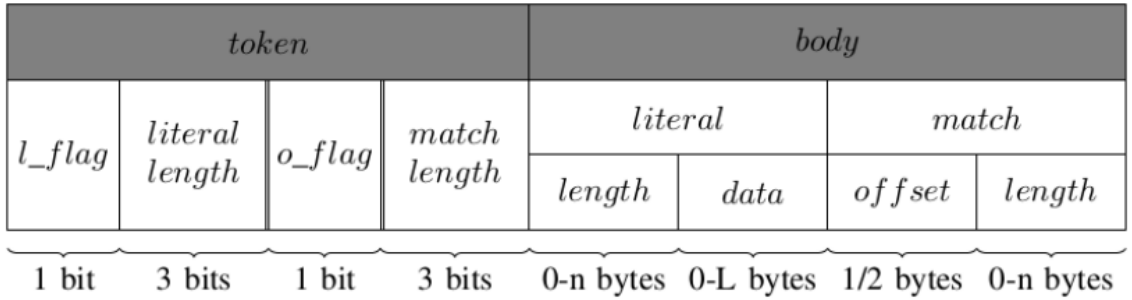

LZ4r Evaluation

- Test platforms:

Platform	CPU Model	L1 Cache	L2 Cache	L3 Cache	Memory	OS
Y700	Intel(R) Core(TM) i7-6700HQ CPU @ 2.60GHz	8 × 64KB	8 × 256KB	6MB	16GB	Ubuntu 16.04
Grid	Intel(R) Xeon(R) CPU E5-2698 v3 @ 2.30GHz	2 × 64KB	2 × 256KB	4MB	8GB	CentOS 7

- Comparison algorithms:
  - Zlib 1.2.11, Deflate 1.3, LZO1x 2.10, Snappy 1.1.4, LZ4 1.9.2
- Corpus sets:
  - Silesia, Calgary, Canterbury, enwik8

LZ4r Evaluation – Compression Ratio

- Compression ratio: 5%-10% better (LZO, Snappy, LZ4)

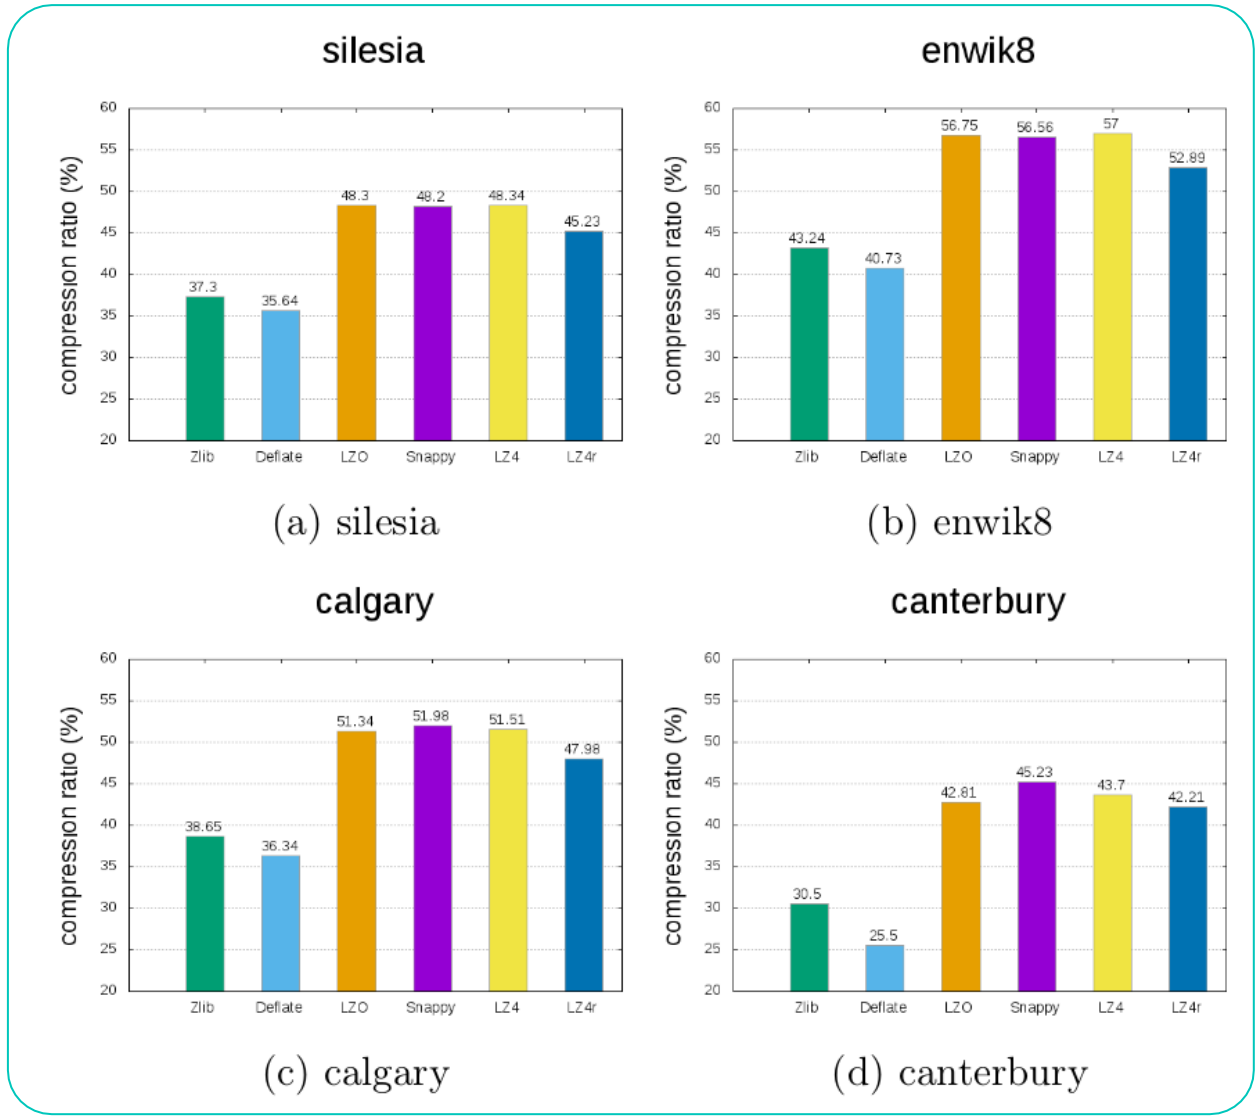

(a) silesia (b) enwik8

(c) calgary (d) canterbury

LZ4r Evaluation – Speed

- Compression speed: similar (LZO, LZ4), 5%-10% faster (Snappy)
- Decompression speed: similar (Snappy), compromised (LZ4)

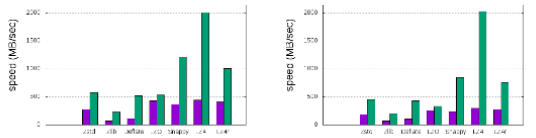

(a) silesia corpus.

(b) enwik8 corpus.

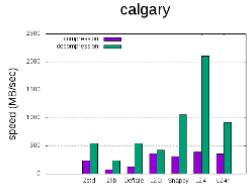

(c) calgary corpus.

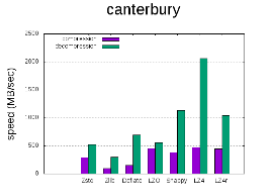

(d) canterbury corpus.

Figure 3.29: Performance on Y700 with *blocksize=64KB*.

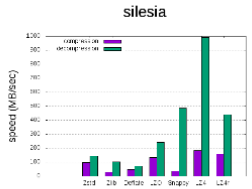

(a) silesia corpus.

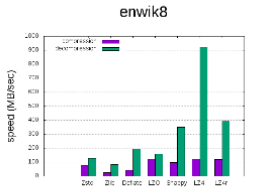

(b) enwik8 corpus.

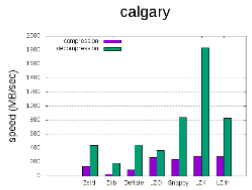

(c) calgary corpus.

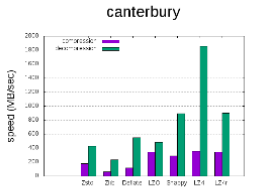

(d) canterbury corpus.

LZ4r Evaluation

- In practical systems, compression/decompression operations are at similar rate
- Calls compression/decompression to have similar speed
- LZ4r decompression speed is still much faster than compression, thus, is unlikely to become the bottleneck of the system

Thank you!

Rui Chen (ruichen@mathworks.com)

Lihao Xu (lihao@wayne.edu)