Corrections to "Toward Optimal Storage Scaling via Network Coding: From Theory to Practice"

Xiaoyang Zhang, Yuchong Hu, Patrick P. C. Lee, and Pan Zhou

Last modified: April 9, 2018

Issue 1: Use of Vandermonde-based Reed-Solomon codes (Last modified: April 9, 2018)

We clarify the use of Vandermonde-based Reed-Solomon codes in NCScale. NCScale uses an $(n-k) \times k$ Vandermonde matrix as the parity-check matrix to compute parity blocks. However, it is known that systematic codes built on Vandermonde matrices generally do not preserve the MDS property under finite fields [2, 3].

Intel's Intel Storage Acceleration Library (ISA-L) also uses a Vandermonde matrix to form the parity-check matrix of systematic Reed-Solomon codes (in the function `gf_gen_rs_matrix`), and some users reported that some parameters may cause decoding operations to fail [1]. ISA-L (since version 2.19) now provides a program `gen_rs_matrix_limits` to find the valid parameters that maintain the MDS property of Reed-Solomon codes in GF(2^8). The valid parameters (for $k \leq 62$ and $n \leq 63$) include:

- $k \leq 3$;
- $k \leq 4$ and $n \leq 25$;
- $k = 5$ and $n \leq 10$;
- $k \leq 21$ and $n - k = 4$; and
- $n - k \leq 3$.

We also extended the above program and find the valid parameters in GF(2^{16}). The valid parameters (for $k \leq 62$ and $n \leq 63$) include:

- $k \leq 4$;
- $k = 5$ and $n \leq 57$;
- $k \leq 7$ and $n - k = 16$;
- $k \leq 9$ and $n - k = 14$;
- $k \leq 11$ and $n - k = 11$;
- $k \leq 14$ and $n - k = 9$;
- $k \leq 16$ and $n - k = 7$;
- $k \leq 52$ and $n - k = 5$; and
- $n - k \leq 4$.

Thus, the analysis of NCScale still holds for the above parameters. We believe that these parameters cover most erasure coding deployment settings in practical storage systems.

We apologize that the conference version of the paper at INFOCOM'18 does not properly clarify the issue, and we did not realize the issue until after we submitted the final version. This article hope to help readers better understand our work.

Issue 2: Typos (Last modified: April 9, 2018)

- Page 6, Left Column, in the "Prepare" paragraph: The notation "$X_{w \bmod n}$" is incorrect. It should read "$X_{((w-1) \bmod n)+1}$".

References

[1] corrupted fragment on decode #10. `https://github.com/01org/isa-l/issues/10`.

[2] J. Lacan and J. Fimes. Systematic MDS Erasure Codes Based on Vandermonde Matrices. *IEEE Communications Letters*, 8(9):570–572, Sep 2004.

[3] J. S. Plank and C. Huang. Tutorial: Erasure Coding for Storage Applications. Slides presented at FAST-2013: 11th Usenix Conference on File and Storage Technologies, Feb 2013.