

The Case for Cross-Entity Delta Encoding in Web Compression

Benjamin Wollmer^{1,3}, Wolfram Wingerath^{2,3}, Sophie Ferrlein³,
Fabian Panse¹, Felix Gessert³, and Norbert Ritter¹

¹ University of Hamburg, Germany
`dbis-research@uni-hamburg.de`

² University of Oldenburg, Germany
`data-science@uni-oldenburg.de`

³ Baqend, Hamburg, Germany
`research@baqend.com`

Abstract. Delta encoding and shared dictionary compression (SDC) for accelerating Web content have been studied extensively in research over the last two decades, but have only found limited adoption in the industry so far: Compression approaches that use a custom-tailored dictionary per website have all failed in practice due to lacking browser support and high overall complexity. General-purpose SDC approaches such as Brotli reduce complexity by shipping the same dictionary for all use cases, while most delta encoding approaches just consider similarities between versions of the same entity (but not between different entities). In this study, we investigate how much of the potential benefits of SDC and delta encoding are left on the table by these two simplifications. As our first contribution, we describe the idea of cross-entity delta encoding that uses cached assets from the immediate browser history for content encoding instead of a precompiled shared dictionary: This avoids the need to create a custom dictionary, but enables highly customized and efficient compression. Second, we present an experimental evaluation of compression efficiency to hold cross-entity delta encoding against state-of-the-art Web compression algorithms. We consciously compare algorithms some of which are not yet available in browsers to understand their potential value before investing resources to build them. Our results indicate that cross-entity delta encoding is over 50% more efficient for text-based resources than compression industry standards. We hope our findings motivate further research and development on this topic.

Keywords: Delta Encoding · Caching · Dictionary Compression.

1 Introduction

Every Web browser utilizes a local cache to reduce the payload of a website. But since cache entries are limited in their lifetime, they become useless in current schemes once they are stale. However, stale resources can still contain information that is useful for encoding related files efficiently. Delta encoding is

an example of such an approach, which is generally used to update one entity to its newest version by sending a diff rather than the whole asset. Most proposals revolving around this mechanism focus on the similarities between versions of the same entity (single-entity data encoding).

In this work, we argue that modern websites comprise many pages that are very similar among one another (e.g. different product pages) and therefore lend themselves to *cross-entity delta encoding* as well [10]. We use Compaz [9] to evaluate the concept of cross-entity delta encoding and provide evidence on its potential benefits for payload savings to motivate further research on the topic. Sections 2 and 3 discuss and distinguish cross-entity delta encoding from existing work. In Section 4, we present quantitative results for the potential of cross-entity delta encoding to improve compression efficiency based on real-world high traffic website traces. We discuss open challenges and conclude in Section 5.

2 Related Work

Delta Encoding. Mogul et al. proposed to use delta encoding in HTTP to update stale content [6], which is not implemented by any major browser. How well this scales depends on how much of the content changes between the two versions. They evaluated the delta calculation purely for updates of returning users, but did not consider deltas between different pages. Korn et al. proposed VCDIFF, a differencing algorithm [4]. They evaluated it similar to Mogul et al. by considering deltas between updates of the same file. Cloudflare’s Railgun uses delta encoding to update the CDN content [3]. This approach is limited between server and CDN and also only considers updates between different versions.

Shared Dictionary Compression (SDC). In a standard compression approach, like with deflate, the encoder reads the file and tries to find repeating strings from the previously read content. The previously read content is also referred to as the dictionary. Instead of just using the previously seen content, the dictionary could also be an external file. In SDC, the same dictionary is shared between multiple compression processes and can therefore improve the overall compression ratio further. Chan et al. suggest that Web pages with a similar URL path also may have similar content and Web pages may therefore be transmitted more efficiently as a differentials to previously visited Web pages [2]. They only consider HTML files and assume them to be uniquely identifiable by URL. While the approach is similar to ours, the presented results are not applicable to modern websites. First, today’s HTML files are often personalized and thus not uniquely identifiable by URL. Second, some assets are static and uniquely identifiable by URL (e.g. JavaScript or CSS), but they are not considered. Butler et al. proposed Shared Dictionary Compression over HTTP (SDCH), where the server can actively push dictionaries to the client [5]. One of the key challenges here is to find the best dictionary, since the server has to predict which dictionary would be of use for the client. This dictionary may increase the payload for the first page, since it is pushed, but maybe only used later. LinkedIn reported that generating the dictionaries took them about 7 hours per deployment and

could easily take days [7]. Therefore, they were forced to delay the generation of new dictionaries to every other week. This may be one of the reasons why SDCH was not widely adopted and removed from Chrome⁴. However, the key idea of sharing a dictionary led to Brotli, which was later developed at Google [1]. For Brotli, the shared dictionary is static and already part of the library and never has to be generated or transferred over the network. Most browsers support Brotli, but without the custom dictionary functionality which could be used for cross-entity delta encoding. Zstandard would allow the same, but has no browser support at all.

3 Cross-Entity Delta Encoding

As shown in the previous section, delta encoding has so far mostly been evaluated to compute deltas between different versions of the same file or with a shared dictionary. Calculating deltas between files that share similar data could provide similar advantages. In contrast to SDCH, this would remove the need to create and maintain dictionaries as we use the raw files as dictionaries. This has some implications. First, the compression for one asset may deliver different results for different users, since the result depends on the dictionaries available in the client cache. Second, using client cache entries requires a cache state synchronization as the server needs to know which resources can be used for content encoding.

Dictionary Scope Strategies. As a basic rule, only those assets can be used for encoding, which have already been loaded by the client. We therefore consider three different strategies for our evaluation. As the most powerful strategy, one could consider every previous asset (PA) as a potential dictionary, which was seen until the currently requested asset. This includes previously visited pages (page impressions, PIs) as well as assets from the currently requested PI. This strategy is difficult to implement in practice, because assets are typically not downloaded in sequence. We exclude the currently processed PI assets as another more practical strategy and only consider fully downloaded assets up to the previous PI (PP). As a third strategy, we exclusively consider assets of the entry page as dictionaries (EE). Due to a similar overall layout (e.g. same header), this could be a reasonable alternative with a fixed set of dictionaries.

4 Compression Efficiency on Real-World Traces

In this Section, we examine how cross-entity delta encoding could affect the transferred size within a user journey. We start by creating a dataset collected from real websites and used them to compare different compression approaches.

Creating a Dataset. The potential of cross-entity delta encoding relies on the journey taken by a user, since it defines which dictionaries are in the cache. We created artificial ones, since we have no access to real user journeys. We assume that every website offers different kind of page types of which one is

⁴ groups.google.com/a/chromium.org/d/msg/blink-dev/nQl0ORHy7sw/HNpR96sqAgAJ

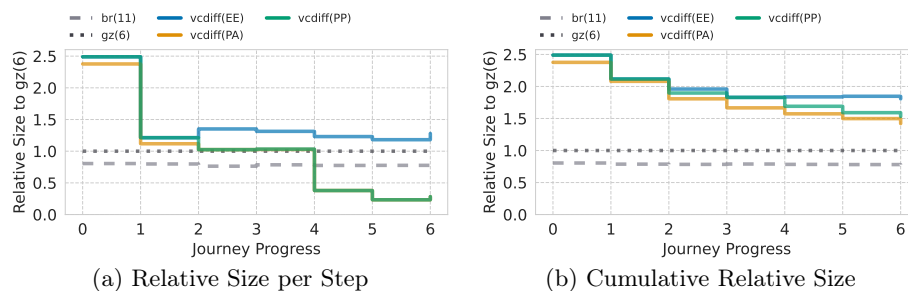


Fig. 1: The left chart shows the compression size relative to gzip (6), while the right chart shows the cumulative size up until each step to account for different page sizes. Replacing every compression with VCDIFF only pays off in the hot phase (>3), but is on average still worse than the default gzip compression.

the main page type (e.g. a product in a shop or an article of a news/blog site). The other types could be types like a category site or the homepage. We further assume the users ultimate goal to be the content of this main page type. We start the navigation at the homepage and from there we try to hit different page types for Step 1 and 2. Step 3 then navigates to the main content. We name this part of the journey from now on the *cold phase*, since we hit distinct page types and the cache is cold in terms of available dictionaries. Step 4-6 are only recommendations from the previous main content and therefore navigate over potentially similar pages. In contrast, this path is from now on referred to as the *hot phase*, since it contains possible dictionary matches. We make sure that every step in the whole journey is unique. We used 40 of the traffic-heaviest websites according to SimilarWeb⁵, providing the recommendation functionality.

Data Cleaning. This work focuses on text based content, therefore, we excluded every non-text asset, identified by the content-type header. We removed trivial cases, like the delta between two identical assets, since this is entirely preventable and would skew the results in favor of cross-entity delta encoding. Finally, we removed every third-party asset (other domain), since the provider would not have the possibility to change the compression for these kind of assets.

Calculating Deltas. The compression for the delta was done with open-VCDIFF to which we from now on will refer to as VCDIFF. We just brute-force every possible delta for a given asset and chose the smallest one. We only considered dictionaries which had at least the same type⁶, e.g. text.

Comparing the Results. The data we collected was compressed by either Brotli, gzip, or no compression. To create a baseline against which we can compare, we uncompressed every asset and compressed it with the default gzip compression level (6), which on average is slightly higher than the results we got

⁵ <https://www.similarweb.com/de/top-websites/>

⁶ Still, the best dictionary was almost every time of the same subtype, e.g. text/html.

from the server. We also compare against Brotli (11), which may be impractical due to performance reasons, but represents the currently best compression ratio.

4.1 Enforcing VCDIFF

In our first experiment, we forced VCDIFF with the different strategies on every asset. It shows that regardless of the current step within our journey, just using the entry site will be outperformed in every step by gzip (see Fig. 1a). The results for PP and PE are similar in the cold phase and on par with gzip after leaving the entry page, but significantly improve and surpass even Brotli when entering the hot phase, with as low as 28% of gzip’s size. This was expected, since we only look at previously visited page types. Still, due to the negative impact at the beginning, they cannot compete over the whole journey (see Fig. 1b).

4.2 Case-Specific VCDIFF

The previous experiment has shown that the results highly depend on the client cache. Therefore, we repeated the experiment but decided per asset to either use VCDIFF if it yields an improvement or stick with gzip (6) otherwise. Since the first PI still cannot leverage any dictionary, it mainly falls back to gzip (see Figure 2a). Just using the entry page as the dictionary source can slightly improve the compression ratio overall (~5%), but as Figure 2a indicates, this is mainly driven by similarities within the cold phase. Again, choosing the dictionary from previous assets provides the biggest impact within the hot phase and vastly outperforms Brotli. Stepping back to PP only slightly decreases the impact and can be an alternative to PA. Overall, using VCDIFF on specific assets scales well with the length of the journey, as soon as similar pages are visited. As Figure 2a and 2b show, this will converge against 25% of the payload for longer journeys.

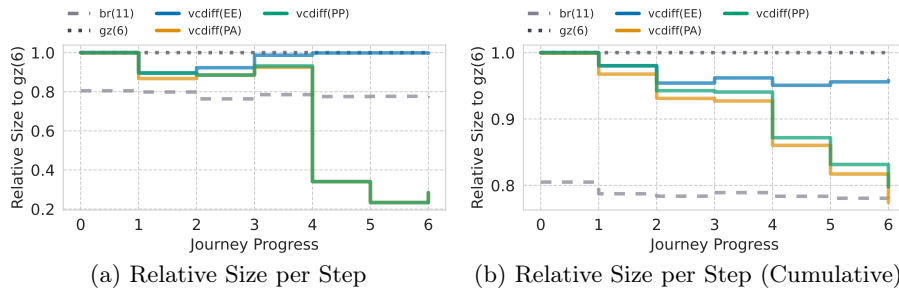


Fig. 2: Using VCDIFF only when it actually provides an uplift yields small results in the cold phase (<4), but can even further improve the hot phase and overall eventually leads to results comparable with the maximum Brotli compression.

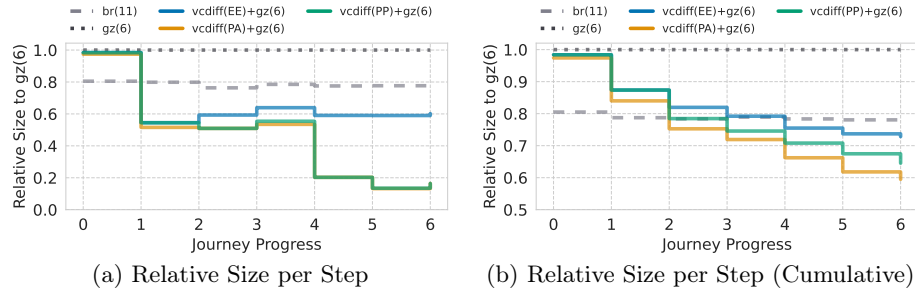


Fig. 3: While VCDIFF is still not competitive on the first page load, combining it with secondary compression outperforms Brotli on every following page load.

4.3 VCDIFF With Secondary Compression

VCDIFF files can still contain many common strings and could therefore benefit from secondary compression, which is currently not implemented in open-VCDIFF. This motivated our next experiment. We used VCDIFF as an Opt-in on every text asset, but piped the VCDIFF output through gzip on level 6. Figure 3a shows that this approach drastically improves the results, as we now outperform Brotli even before entering the hot phase. Within the hot phase, we can compress the assets as low as 14% of gzip (6). Overall the cumulative size can be reduced to 58% of our baseline (see Fig. 3b).

4.4 Impact on Different MIME-Types

We expected the HTML to gain the most benefit of cross-entity delta encoding and compared them with the other types. We grouped them by the step, as well as the subtype to make sure that the actual weight of the single assets were reflected. For the cold phase, we exclude the entry pages (step 0), since the previous experiments have already shown that cross-entity encoding is no real alternative in that step. HTML benefits the most from cross-delta entity encoding (see Figure 4) and is a safe alternative in the hot phase. We excluded




Fig. 4: Compression efficiency by MIME type: Using HTML for cross-entity delta encoding works well in almost all cases. Using other MIME types leads to mixed results in the cold phase (a), but yields high efficiency in the hot phase (b).




Fig. 5: Replacing the secondary compression yields even better results. Alternatively, Brotli and zstandard also allow a more efficient use of custom dictionaries.

XML and plain text responses, since we had too few samples. SVGs could in some cases further reduced by 10% compared to gz (6). While the other types on average can still greatly benefit, they are site-specific cases and need further investigation. Note that the hot phase had only a few CSS samples and may not be representative, since most sites do not load additional CSS files at this point.

4.5 Using Algorithms Beyond VCDIFF

Combining gzip (6) with VCDIFF already could on average reduce the journey size to 58%. But compression efficiency could be further improved by increasing the compression level and/or using different algorithms. We are aware that a higher compression level may be impractical in an actual deployment, but should act as an upper limit. While increasing gzip to level 9 had almost the same result, replacing gzip (6) with Brotli (11) reduces the cumulative journey size to 51% (see Figure 5). As mentioned earlier, one could also directly use a custom dictionary with Brotli or zstandard (which is not supported by any browser). This reached the highest compression ratio and could reduce the result to 45% for Brotli (11) and 48% for zstandard (22), compared to gzip (6)⁷

5 Conclusion

Our results show that reusing cached assets for delta encoding can significantly reduce transferred bytes in the Web, even though we simulate new users who start with cold caches. This approach should thus be seen as a complement to existing compressions rather than a replacement. But it should be noted that potential uplift is even more significant for returning user who start their journeys in the hot phase and thus directly benefit from cross-entity delta encoding.

Open Challenges. In our experiments, we made several simplifying assumptions that do not hold in a real-world setting. First, we employ the perfect dictionary selection via a brute-force approach, but a more efficient heuristic would be required for practical use. Also, cache state synchronization remains challenging:

⁷ Due to limited space, we only present a few selected alternatives here and refer to <https://icwe.compaz.info> for an extensive overview.

The server does not only have to select the ideal dictionary for encoding, but also one that is already present in the client cache to enable decoding. Another open challenge is the lack of browser support for different aspects of delta encoding and shared dictionary compression. While all major browsers support generic Brotli and gzip, VCDIFF and Brotli with a custom dictionary are currently not supported by any of them. While using HTML files as dictionaries was most effective in our evaluation, content that is generated per user makes it infeasible to keep all dictionaries (HTMLs) in the server. Reducing this complexity would require some kind of normalization to strip personalized content for encoding and decoding (cf. Dynamic Blocks [8]), akin to app shells in single-page applications.

Closing Thoughts. Despite a host of literature on delta encoding and shared dictionary compression from more than two decades of research, there is still a lot of untapped potential in existing compression technologies. Our results indicate that using the client cache as a dictionary for delta encoding can reduce the text payload by up to 86% for single pages and by 55-80% for user journeys over recommended content. But there is still further research needed in areas like dictionary selection and cache state synchronization. Lacking browser support for cross-entity delta encoding algorithms is another practical barrier, but could be added in platform-independent fashion with a service worker implementation. However, performance depends on the client device and is likely not comparable with native compression algorithms. Without native browser support, delta encoding only seems viable for scenarios where network efficiency is critical (e.g. for mobile users in data saving mode).

References

1. Alakuijala, J., Farruggia, A., Ferragina, P., Kliuchnikov, E., Obyrk, R., Szabadka, Z., Vandevenne, L.: Brotli: A general-purpose data compressor. *ACM TOI* **37**(1)
2. Chan, M.C., Woo, T.: Cache-Based Compaction: A New Technique for Optimizing Web transfer. In: *IEEE INFOCOM '99. Conference on Computer Communications*.
3. Knecht, D.O., Graham-Cumming, J., Prince, M.B.: Method and apparatus for reducing network resource transmission size using delta compression
4. Korn, D.G., Vo, K.P.: Engineering a differencing and compression data format. In: *USENIX annual technical conference, general track*. pp. 219–228 (2002)
5. McQuade, B., Mixter, K., Lee, W.H., Butler, J.: A proposal for shared dictionary compression over http (2016)
6. Mogul, J.C., Douglass, F., Feldmann, A., Krishnamurthy, B.: Potential Benefits of Delta Encoding and Data Compression for HTTP. *SIGCOMM CCR* (1997)
7. Shapira, O.: SDCH at LinkedIn. (2015), <https://engineering.linkedin.com/shared-dictionary-compression-http-linkedin>, accessed: 2022-01-20
8. Wingerath, W., Gessert, F., Witt, E., Kuhlmann, H., Bücklers, F., Wollmer, B., Ritter, N.: Speed Kit: A Polyglot & GDPR-Compliant Approach For Caching Personalized Content. In: *ICDE, Dallas, Texas* (2020)
9. Wollmer, B., Wingerath, W., Ferrlein, S., Gessert, F., Ritter, N.: Compaz: Exploring the Potentials of Shared Dictionary Compression on the Web. In: *22th International Conference on Web Engineering, ICWE* (2022)
10. Wollmer, B., Wingerath, W., Ritter, N.: Context-Aware Encoding & Delivery in the Web. In: *20th International Conference on Web Engineering, ICWE* (2020)