

Space Efficient Deep Packet Inspection of Compressed Web Traffic

Yehuda Afek¹, Anat Bremler-Barr^{1,1}, Yaron Koral¹,

^aBlavatnik School of Computer Sciences Tel-Aviv University, Israel

^bComputer Science Dept. Interdisciplinary Center, Herzliya, Israel

Abstract

In this paper we focus on the process of deep packet inspection of compressed web traffic. The major limiting factor in this process imposed by the compression, is the high memory requirements of 32KB per connection. This leads to the requirements of hundreds of megabytes to gigabytes of main memory on a multi-connection setting. We introduce new algorithms and techniques that drastically reduce this space requirement for such bump-in-the-wire devices like security and other content based networking tools. Our proposed scheme improves both space and time performance by almost 80% and over 40% respectively, thus making real-time compressed traffic inspection a viable option for networking devices.

Keywords: pattern matching, compressed http, network security, deep packet inspection

1. Introduction

Compressing HTTP text when transferring pages over the web is in sharp increase motivated mostly by the increase in web surfing over mobile devices. Sites such as Yahoo!, Google, MSN, YouTube, Facebook and others use HTTP compression to enhance the speed of their content download. In Section ?? we provide statistics on the percentage of top sites using HTTP Compression. Among the top 1000 most popular sites 66% use HTTP compression (see Figure ??). The standard compression method used by HTTP 1.1 is GZIP.

This sharp increase in HTTP compression presents new challenges to networking devices, such as intrusion-prevention system (IPS), content filtering and web-application firewall (WAF), that inspect the content for security hazards and balancing decisions. Those devices reside between the server and the client and perform *Deep Packet Inspection* (DPI). Upon receiving compressed traffic the networking device needs first to decompress the message in order to inspect its payload. We note that GZIP replaces repeated strings with back-references, denoted as pointers, to their prior occurrence within the last 32KB of the text. Therefore, the decompression process requires a 32KB buffer of the recent decompressed data to keep all possible bytes that might

be back-referenced by the pointers, what causes a major *space* penalty. Considering today's mid-range firewalls which are built to support 100K to 200K concurrent connections, keeping a buffer for the 32KB window for each connection occupies few gigabytes of main memory. Decompression causes also a *time* penalty but the time aspect was successfully reduced in [?].

This high memory requirement leaves the vendors and network operators with three bad options: either ignore compressed traffic, or forbid compression, or divert the compressed traffic for offline processing. Obviously neither is acceptable as they present a security hole or serious performance degradation.

The basic structure of our approach is to keep the 32KB buffer of all connections compressed, except for the data of the connection whose packet(s) is now being processed. Upon packet arrival, unpack its connection buffer and process it. One may naïvely suggest to just keep the appropriate amount of original compressed data as it was received. However this approach fails since the buffer would contain recursive pointers to data more than 32KB backwards. Our technique, called "Swap Out-of-boundary Pointers" (*SOP*), packs the buffer's connection by combining recent information from both compressed and uncompressed 32KB buffer to create the new compressed buffer that contains pointers that refer only to locations within itself. We show that by employing our technique for DPI on real life data we reduce the space requirement by a factor

*Corresponding author

¹Supported by European Research Council (ERC) Starting Grant no. 259085.

55 of 5 with a time penalty of 26%. Notice that while our
56 method modifies the compressed data locally, it is trans-
57 parent to both the client and the server.

58 We further design an algorithm that combines our
59 *SOP* technique that reduces space with the ACCH algo-
60 rithm which was presented in [?] (method that accel-
61 erates the pattern matching on compressed HTTP traf-
62 fic). The combined algorithm achieves an improvement
63 of 42% on the time and 79% on the space requirements.
64 The time-space tradeoff presented by our technique pro-
65 vides the first solution that enables DPI on compressed
66 traffic in wire speed for network devices such as IPS and
67 WAF.

68 The paper is organized as follows: A background on
69 compressed web traffic and DPI is presented in Section
70 ??. An overview on the related work appears in Section
71 ??. An overview on the challenges in performing DPI
72 on compressed traffic appears in Section ??. In Section
73 ?? we describe our *SOP* algorithm and in Section ??
74 we present the combined algorithm for the entire DPI
75 process. Section ?? describes the experimental results
76 for the above algorithms and concluding remarks appear
77 in Section ??.

78 Preliminary abstract of this paper was published in
79 the proceedings of IFIP Networking 2011 [?].

80 2. Background

81 In this section we provide background on compressed
82 HTTP and DPI and its time and space requirements.
83 This helps us in explaining the considerations behind
84 the design of our algorithm and is supported by our ex-
85 perimental results described in Section ??.

86 **Compressed HTTP:** HTTP 1.1 [?] supports the
87 usage of content-codings to allow a document to be
88 compressed. The RFC suggests three content-codings:
89 GZIP, COMPRESS and DEFLATE. In fact, GZIP uses
90 DEFLATE as its underlying compression protocol. For
91 the purpose of this paper they are considered the same.
92 Currently GZIP and DEFLATE are the common cod-
93 ings supported by current browsers and web servers.²

94 The GZIP algorithm uses a combination of the fol-
95 lowing compression techniques: first the text is com-
96 pressed with the LZ77 algorithm and then the output is
97 compressed with the Huffman coding. Let us elaborate
98 on the two algorithms:

99 *LZ77 Compression* [?]- The purpose of LZ77 is to
100 reduce the *string presentation size*, by spotting repeated

101 strings within the last 32KB of the uncompressed data.
102 The algorithm replaces a repeated string by a backward-
103 pointer consisting of a (*distance,length*) pair, where *dis-*
104 *tance* is a number in [1,32768] (32K) indicating the
105 distance in bytes of the string and *length* is a number
106 in [3,258] indicating the length of the repeated string.
107 For example, the text: ‘abcdeabc’ can be compressed
108 to: ‘abcde(5,3)’; namely, “go back 5 bytes and copy 3
109 bytes from that point”. LZ77 refers to the above pair as
110 “pointer” and to uncompressed bytes as “literals”.

111 *Huffman Coding* [?]- Recall that the second stage
112 of GZIP is the Huffman coding, that receives the LZ77
113 symbols as input. The purpose of Huffman coding is
114 to reduce the *symbol coding size* by encoding frequent
115 symbols with fewer bits. The Huffman coding method
116 builds a dictionary that assigns to symbols from a given
117 alphabet a variable-size *codeword* (coded symbol). The
118 codewords are coded such that no codeword is a prefix
119 of another so the end of each codeword can be easily
120 determined. *Dictionaries* are constructed to facilitate
121 the translation of binary codewords to bytes.

122 In the case of GZIP, Huffman encodes both literals
123 and pointers. The distance and length parameters are
124 treated as numbers, where each of those numbers is
125 coded with a separate codeword. The Huffman dictio-
126 nary which states the encoding of each symbol, is usu-
127 ally added to the beginning of the compressed file (oth-
128 erwise a predefined dictionary is selected).

129 The Huffman decoding process is relatively fast. A
130 common implementation (cf. zlib [?]) extracts the
131 dictionary, with average size of 200B, into a temporary
132 lookup-table that resides in the cache. Frequent sym-
133 bols require only one lookup-table reference, while less
134 frequent symbols require two lookup-table references.

135 **Deep packet inspection (DPI):** DPI is the process of
136 identifying signatures (patterns or regular expressions)
137 in the packet payload. Today, the performance of secu-
138 rity tools is dominated by the speed of the underlying
139 DPI algorithms [?]. The two most common algorithms
140 to perform string matching are the Aho-Corasick (AC)
141 [?] and Boyer-Moore (BM) [?] algorithms. The BM
142 algorithm does not have deterministic time complexity
143 and is prone to denial-of-service attacks using tailored
144 input as discussed in [?]. Therefore the AC algorithm
145 is the standard. The implementations need to deal with
146 thousands of signatures. For example, ClamAV [?] virus-
147 signature database contains 27 000 patterns, and the
148 popular Snort IDS [?] has 6 600 patterns; note that
149 typically the number of patterns considered by IDS sys-
150 tems grows quite rapidly over time.

151 In Section ?? we provide an algorithm that combines
152 our *SOP* technique with a Aho-Corasick based DPI al-

²Analyzing captured packets from last versions of both Internet Explorer, FireFox and Chrome browsers shows that accept only the GZIP and DEFLATE codings.

153 gorithm. The Aho-Corasick algorithm relies on an underlying
154 Deterministic Finite Automaton (DFA) to support all required
155 patterns. A DFA is represented by a “five-tuple” consisting
156 of a finite set of states, a finite set of input symbols, a
157 transition function that takes as arguments a state and an
158 input symbol and returns a state, a start state and a set of
159 accepting states. In the context of DPI, the sequence of
160 symbols in the input results in a corresponding traversal of
161 the DFA. A transition to an accepting state means that one
162 or more patterns were matched.

164 In the implementation of the traditional algorithm the
165 DFA requires dozens of megabytes and may even reach
166 gigabytes of memory. The size of the signatures databases
167 dictates not only the memory requirement but also the
168 speed, since it dictates the usage of a slower memory,
169 which is an order-of-magnitude larger DRAM, instead of
170 using a faster one, which is SRAM based. We use that
171 fact later when we compare DPI performance to that of
172 GZIP decompression. That leads to an active research
173 on reducing the memory requirement by compressing the
174 corresponding DFA [? ? ? ? ?]; however, all proposed
175 techniques suggest pure-hardware solutions, which usually
176 incur prohibitive deployment and development cost.

178 3. Related Work

179 There is an extensive research on performing pattern
180 matching on compressed files as in [? ? ? ?], but very
181 limited is on compressed traffic. Requirements posed in
182 dealing with compressed traffic are: (1) on-line scanning
183 (1-pass), (2) handling thousands of connections concurrently
184 and (3) working with LZ77 compression algorithm (as
185 oppose to most papers which deal with LZW/LZ78
186 compressions). To the best of our knowledge, [? ?]
187 are the only papers that deal with pattern matching
188 over LZ77. However, in those papers the algorithms
189 are for single pattern and require two passes over
190 the compressed text (file), which is not an option
191 in network domains that require ‘on-the-fly’ processing.

192 Klein and Shapira [?] suggest a modification to the
193 LZ77 compression algorithm, to change the backward
194 pointer into forward pointers. That modification makes
195 the pattern matching easier in files and may save some
196 of the required space by the 32KB buffer for each
197 connection. However, the suggestion is not implemented
198 in today’s HTTP.

199 The first paper to analyze the obstacles of dealing
200 with compressed traffic is [?], but it only accelerated
201 the pattern matching task on compressed traffic and did

202 not handle the space problem, and it still requires the
203 de-compression. We show in Section ?? that our paper can
204 be combined with the techniques of [?] to achieve a fast
205 pattern matching algorithm for compressed traffic, with
206 moderate space requirement. In [?] an algorithm that
207 applies the Wu-Manber [?] multi-patterns matching
208 algorithm on compressed web-traffic is presented. Al-
209 though here we combine SOP with an algorithm based
210 on Aho-Corasick, only minor modifications are required
211 to combine SOP with the algorithm of [?].

212 There are techniques developed for “in-place decom-
213 pression”, the main one is LZ0 [?]. While LZ0 claims
214 to support decompression without memory overhead it
215 works with files and assumes that the uncompressed
216 data is available. We assume decompression of thou-
217 sands of concurrent connections on-the-fly, thus what
218 is for free in LZ0 is considered overhead in our case.
219 Furthermore, while GZIP is considered the standard for
220 web traffic compression, LZ0 is not supported by any
221 web server or web browser.

222 4. Challenges in performing DPI on Compressed 223 HTTP

224 This section provides an overview of the obstacles in
225 performing deep packet inspection (DPI) in compressed
226 web traffic in a multi-connection setting.

227 While transparent to the end-user, compressed web
228 traffic needs special care by bump-in-the-wire devices
229 that reside between the server and the client and per-
230 form DPI. The device needs first to decompress the data
231 in order to inspect its payload since there is no appar-
232 ent “easy” way to perform DPI over compressed traf-
233 fic without decompressing the data in some way. This
234 is mainly because LZ77 is an *adaptive* compression
235 algorithm, namely the text represented by each symbol
236 is determined dynamically by the data. As a result,
237 the same substring is encoded differently depending on
238 its location within the text. For example the pattern
239 ‘abcdef’ can be expressed in the compressed data by
240 $abcde *^j (j + 5, 5)f$ for all possible $j < 32763$.

241 One of the main problems with the decompression is
242 its memory requirement; the straightforward approach
243 requires a 32KB sliding window for each connection.
244 Note that this requirement is difficult to avoid, since
245 the back-reference pointer can refer to any point within
246 the sliding window and the pointers may be recursive
247 (i.e., a pointer may point to an area with a pointer).
248 As opposed to compressed traffic, DPI of non-compressed
249 traffic requires storing only two or four bytes variable
250 that holds the corresponding DFA state aside of the DFA
itself,

251 which is of course stored in any case. Hence, deal- 299
252 ing with compressed traffic poses a significantly higher 300
253 memory requirement by a factor of 8 000 to 16 000. 301
254 Thus, mid-range firewall that handles 100K-200K con- 302
255 current connections (like GTA's G-800 [?], SonicWalls 303
256 Pro 3060 [?] or Stonesofts StoneGate SG-500 [? 304
257]), needs 3GB-6GB memory while a high-end firewall 305
258 that supports 500K-10M concurrent connections (like 306
259 the Juniper SRX5800 [?] or the Cisco ASA 5550 or 307
260 5580 [?]) would need 15GB-300GB memory only for 308
261 the task of decompression. This memory requirement 309
262 not only imposes high price and infeasibility of the 310
263 architecture but also implies on the capability to perform 311
264 caching or using fast memory chips such as SRAM. 312
265 Hence, reducing the space boosts the speed also because 313
266 faster memory technology is becoming a viable op- 314
267 tion, such as SRAM memory. This work deals with the 315
268 challenges imposed by that space aspect.

269 Apart from the space penalty described above, the 316
270 decompression stage also increases the overall *time* 317
271 penalty. However, we note that DPI requires signifi- 318
272 cantly more time than decompression, since decompres- 319
273 sion is based on reading consecutive memory locations 320
274 and therefore enjoys the benefit of cache block archite- 321
275 cture and has low per-byte read cost, where as DPI em- 322
276 ploys a very large data structure that is accessed by reads 323
277 to non-consecutive memory areas therefore requires ex- 324
278 pensive main memory accesses. In [?] two of us pro- 325
279 vided an algorithm that takes advantage of information 326
280 gathered by the decompression phase in order to accel- 327
281 erate the commonly used Aho-Corasick pattern match- 328
282 ing algorithm. By doing so, we significantly reduced 329
283 the time requirement of the entire DPI process on com- 330
284 pressed traffic.

285 5. SOP Packing technique

286 In this section we describe our packing technique, 336
287 which reduces the 32KB buffer space requirement to 337
288 about 5.5KB per connection. The basic idea is to keep 338
289 all active connection-buffers in their packed form and 339
290 unpack only the connection under inspection by the 340
291 DPI process. To achieve this we re-pack the buffer 341
292 after each packet-inspection completion while keeping 342
293 a valid updated-buffer as explained below. It has two 343
294 parts:

- 295 • **Packing:** Swap Out-of-boundary Pointers (*SOP*) 345
296 algorithm for buffer packing.
- 297 • **Unpacking:** Our corresponding algorithm for un-
298 packing the buffer.

Whenever a packet is received, the buffer that belongs 300
to the incoming packet connection is unpacked. After 301
the incoming packet processing is finished an updated 302
buffer is packed using the *SOP* algorithm. The next sub- 303
sections elaborate on these parts of the algorithm.

5.1. Buffer Packing: Swap Out of boundary Pointers (*SOP*)

A 1st obvious attempt is to store the buffer in its origi- 304
nal compressed form as received on the channel. How- 305
ever this attempt is incorrect since the compressed form 306
of the buffer contains pointers that recursively point to 307
positions prior to the 32KB boundary. Figure ?? shows 308
an example of the original compressed traffic. Note that 309
it contains a pointer to symbols that are no longer within 310
the buffer boundaries. *Clearly, any solution must main- 311
tain a valid buffer, which is one that contains only point- 312
ers that refer to locations within itself.*

A 2nd obvious attempt is to re-compress (each time 313
from scratch) the 32KB buffer using some compression 314
algorithm such as GZIP. This solution maintains a *valid 315
buffer* as required above, since the compression is based 316
only on information within the buffer. Yet, LZ77 *com- 317
pression* consumes much more time than *decompression*, 318
usually by a factor of 20. Hence this approach has 319
a very high time penalty, making it irrelevant for real- 320
time inspection. We further show that the space saving 321
by this solution is negligible.

Our suggested solution is called Swap Out-of- 322
boundary Pointers (*SOP*). As in the first attempt, *SOP* 323
tries to preserve the original structure of the compressed 324
data as much as possible while maintaining the buffer 325
validity. To preserve buffer validity, *SOP* replaces all 326
the pointers that point to location prior the new buffer- 327
boundary with their referred literals.³ Figure ?? depicts 328
an output of the algorithm. The pointer (300,5) that 329
points outside the buffer boundary is replaced by the 330
string 'hello', where the others remain untouched. 331
332

333 Since in every stage we maintain the invariant that 334
335 pointers refer only to information within the buffer, the 336
337 buffer can be unpacked by simply using GZIP decom- 338
339 pression as discussed in [?]. *SOP* still has a good 340
341 compression ratio (as shown in Section ??); most of 342
343 the pointers are left untouched because they originally 344
345 pointed to a location within the 32KB buffer bound- 346
ary. Beside being space efficient, *SOP* is also fast as 347
it performs only single pass on the uncompressed infor- 348
mation and the compressed buffer information, taking 349

³Note that pointers can be recursive, that is pointer A points to 350
pointer B. In that case we replace pointer A by the literals that pointer 351
B points to.

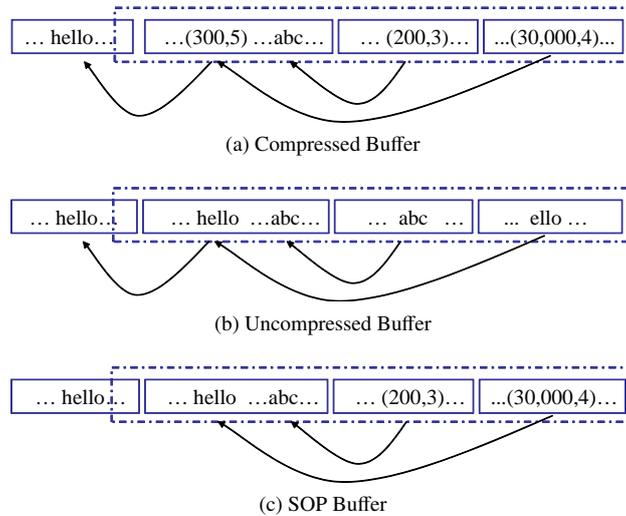


Figure 1: Sketch of the memory buffer in different scenarios. Each solid box represents a packet. Dashed frame represents the 32KB buffer. Each arrow connects between a pointer and its referred area.

346 advantage of the original GZIP traffic compression that
 347 the source (server) has compressed. The pseudocode is
 348 given in Algorithm ???. The term *old* buffer refers to
 349 the buffer that was packed and stored prior to the new
 350 packet arrival and *new* buffer refers to the compressed
 351 buffer that is preserved after the processing of the ar-
 352 rived packet. The algorithm consists of 4 parts:

- 353 1. Calculation of the arrived packet uncompressed
 354 size. This part is required for determining the
 355 boundary of the new buffer.
- 356 2. Decode the old buffer part which is *outside* the new
 357 boundary (its size equals the arrived packet uncom-
 358 pressed size).
- 359 3. Decode the old buffer part which is *within* the new
 360 boundary.
- 361 4. Decode the newly received packet.

362 The unpacking is divided into three parts (2–4) for
 363 code clarity. Recall that the LZ77 pointers and literals
 364 are compressed by the Huffman-coding stage. There-
 365 fore, the first part performs Huffman decoding of the in-
 366 coming packet to calculate its uncompressed size which
 367 in turn is used to determine the *new boundary*. The *new*
 368 *buffer* contains the arrived packet along with the *old*
 369 *buffer* minus the incoming-packet uncompressed-size.
 370 Note that pointers cannot be extracted at this point since
 371 they may refer to locations within the old packed buffer
 372 which is still compressed.

373 The other three parts consist of decoding different
 374 parts of the data, using a temporary buffer of 32KB un-
 375 compressed literals. Part 2 decodes data that would not

376 be re-packed since it is located outside of the new buffer
 377 boundary (after receiving the new packet). Parts 3 and 4
 378 decode the rest of the old packed-buffer and the input
 379 packet respectively, and prepare the the new packed-
 380 buffer along the decoding process. Throughout those
 381 parts, each pointer that back-reference to a location out-
 382 side the new buffer boundary is replaced by its referred
 383 literals.

384 An example of the different parts of the algorithm is
 385 illustrated in Figure ?? where the new arriving packet
 386 is represented by the rightmost box. The new buffer
 387 boundary, which is calculated in *Part 1*, is marked with
 388 the dashed box. Notice that part of the leftmost packet
 389 does not belong to the new buffer any more, and there-
 390 fore may not be referred by backward pointers from
 391 within the new buffer. On the following parts (2–4) the
 392 algorithm replaces each pointer that points to a position
 393 prior the new buffer left boundary with its literals, for
 394 example the (300,5) in Figure ??.

395 Notice that in some cases the output is not optimal.
 396 Figure ?? depicts a case where LZ77 algorithm (??)
 397 outperforms *SOP* (??). In that case, the original com-
 398 pression did not indicate any direct connection between
 399 the second occurrence of the ‘Hello world!’ string to the
 400 string ‘ello world’. The connection can be figured out
 401 if one follows the pointers of both strings and finds that
 402 they share common referred bytes. Since *SOP* performs
 403 a straightforward swapping without following the point-
 404 ers recursively it misses this case. However, the loss
 405 of space is limited as shown in the experimental results
 406 Section ??.

Algorithm 1 Out of Boundary Pointer Swapping

packet - input packet.

oldPacked - the old packed buffer received as input. Every cell is either a literal or a pointer.

newPacked - the new packed buffer.

unPacked - temporary array of 32K uncompressed literals.

```
1: procedure decodeBuffer(buffer)
2: for all symbols in buffer do
3:    $S \leftarrow$  next symbol in buffer after Huffman decode
4:   if  $S$  is literal then
5:     add symbol to unPacked and newPacked buffers
6:   else ▷  $S$  is Pointer
7:     store the referred literals in unPacked
8:     if pointer is out of the boundary then
9:       store the coded referred literals in newPacked
10:    else
11:      store the coded pointer in newPacked
12:    end if
13:  end if
14: end for

15: procedure handleNewPacket(packet,oldPacked)
   Part 1: calculate packet uncompressed size  $n$ 
16: for all symbols in packet do
17:    $S \leftarrow$  next symbol in packet after Huffman decode
18:   if  $S$  is literal then
19:      $n \leftarrow n + 1$ 
20:   else ▷  $S$  is Pointer
21:      $n \leftarrow n +$  pointer length
22:   end if
23: end for
   Part 2: decode out-of-boundary part of oldPacked
24: while less than  $n$  literals were unpacked do
25:    $S \leftarrow$  next symbol in oldPacked after Huffman decode
26:   if  $S$  is literal then
27:     store literal in unPacked buffer
28:   else ▷  $S$  is Pointer
29:     store pointer's referred literals in unPacked buffer
30:   end if
31: end while
   Part 3: decode oldPacked part within boundary
32:  $oldPackedForDecode \leftarrow$  oldPacked part within boundary
33: if boundary falls within a pointer in oldPacked then
34:   copy to newPacked the referred literals suffix only
35:   include boundary pointer in oldPackedForDecode
36: end if
37: decodeBuffer(oldPackedForDecode)
   Part 4: decode packet
38: decodeBuffer(packet)
```

5.2. Huffman Coding Scheme

So far we discussed only the LZ77 part of GZIP and how to use it for packing each connection-buffer. However, additional space can be saved by using an appropriate symbol-coding such as Huffman coding. In fact some coding scheme must be defined since the LZ77 consists of more than 256 symbols; therefore a symbol cannot be decoded in a straight-forward manner using a single byte. We evaluate three options for symbol coding schemes in our packing algorithm:

1. Original Dictionary - Use the Huffman dictionary that was received at the beginning of the connection.
2. Global Dictionary - Use a global fixed Huffman dictionary for all connections.
3. Fixed Size Coding - Use a coding with a fixed size for all symbols, not based on the Huffman coding.

The first option uses the Huffman dictionary obtained from the *original* connection. Recall that Huffman codes frequent symbols with shorter ones. Notice that *SOP* changes the symbol frequencies therefore the original dictionary is not optimal anymore. Generally, *SOP* has more literals than in the original compressed data since it replaces few pointers with literals. We argue that determining the *SOP* symbol frequencies in order to generate a new dictionary is too expensive.

In the second option one *global Huffman dictionary* is used for all the connections. A dictionary based on those frequencies may provide better compression than the *original* dictionary. Therefore we created a common symbol frequency table using a different training dataset as described in the experimental results section.

Assume we want to avoid the complexity of coding variable size symbols as Huffman does. We suggest a third option and call it the *fixed symbol size* coding scheme. This method assigns a fixed size code for any symbol. The symbol values are based on the way GZIP defines its Huffman alphabets, that is two alphabets for symbol coding: the first alphabet represents the literals and pointer-length values and the second represents the pointer-distance values. Thus there is no need to define a special sign that distincts literals from pointers. The pointer-length value is in [3–258]. Since the literals consist of 256 different possible values a 9 bit fixed length symbol is sufficient to represent all possible values from the first alphabet. The second alphabet requires a 15 bit symbol to represent all possible distances in [1–32768]. If current symbol is a literal the decoder assumes that the next symbol is a 9 bit symbol. Otherwise the next symbol is treated as a 15 bit pointer-distance followed by a 9 bit symbol.

458 Comparison between the three coding schemes is presented in Section ???. The results show that the best
459 choice is the second option that uses the *global dictionary*. The *original* dictionary option results in a buffer
460 size which is 3% larger and the *fixed size* scheme results in a buffer size which is 7% larger as compared
461 to using the *global dictionary*. Concerning the time requirement, performing Huffman decoding is a light operation,
462 as mentioned in Section ???. The dictionary is small enough to fit in the cache and both coding and decoding
463 consist mostly of single table-lookup operation. Therefore the time differences between the above three
464 options are negligible as compared to the whole process, and result mainly from the different buffer sizes.

472 5.3. Unpacking the buffer: GZIP decompression

473 So far we introduced the SOP algorithm which is in charge of buffer packing. We now discuss the complimentary
474 algorithm for unpacking these buffers. SOP buffers are packed in a valid GZIP format, hence a standard
475 GZIP decompression can be used for unpacking.

476 Still, we should consider the fact that most of the data is decompressed more than once since it is packed and
477 unpacked whenever a connection becomes active upon a new packet arrival. According to our simulations in
478 Section ??, each byte is decompressed on average 4.2 times using SOP buffers as compared to only once in
479 the original GZIP method without buffer packing.

480 Note that upon unpacking we decompress the entire buffer prior to content inspection. One may wonder
481 if we could reduce the above number by extracting only the parts of the buffer that are actually back-
482 referenced by pointers. Since the pointers are recursive, the retrieval of literals referenced by a pointer is a recursive
483 process, which refers to several locations within the buffer. Hence, the number of decoded symbols for a
484 single pointer-extraction may be greater than the pointer length. For example: a pointer that points to another
485 pointer which in turn points to a third pointer, requires decoding three different areas where the referred pointers
486 and the literals reside.

487 This recursive behavior might lead to referring most of the 32KB buffer, making the cost of maintaining a
488 more complicated technique not worthy. To check this we designed a method for partial decompression, called
489 *SOP-Indexed*. It is defined as follows:

- 490 • Split the compressed buffer to fixed size chunks and keep indices that hold starting positions of
491 each chunk within the compressed buffer.
- 492 • Recursively extract each pointer and decode only the chunks that are referred by some pointer.

508 For example: if 256B chunks are used, referring to the 500th byte of the 32KB uncompressed buffer requires
509 decode of the second chunk that corresponds to the [256–511] byte positions within the uncompressed
510 buffer. If the pointer exceeds the chunk boundary of 511, the next chunk has to be decoded too.

511 Since compressed symbols within the buffer are coded using a variable length number of bits, a chunk
512 might end at a position which is not a multiple of 8. Hence an index to the beginning of the next chunk
513 should indicate a bit offset rather than a byte offset. In order to avoid using bit offset, which requires keeping
514 an eight times larger index, we apply zero padding that pads the end of each chunk with zero bits up to a
515 position that is a multiply of 8. Each index is coded by using 15 bits to represent offsets in [0–32K]. The chunk
516 size poses a time-space tradeoff. Smaller chunks support more precise references and result in less decoding
517 but requires more indices that have to be stored along with the buffer, and more padding for each of the smaller
518 chunks, hence increase the space waste. The analysis in Section ?? shows that a moderate time improvement of
519 36% as compared to *SOP*, is gained by *SOP-Indexed* along with an average space penalty of 300B per buffer.
520 One may choose whether to apply the *SOP-Indexed* or not according to the specific system needs considering
521 the time-space trade-off that is presented here.

522 6. Combining SOP with ACCH algorithm

523 Thus far we discussed SOP, the space efficient buffer packing and unpacking method. In this section we
524 integrate our SOP algorithm with a content inspection algorithm to support the complete DPI process. For
525 that purpose we choose the state-of-the-art ACCH algorithm, presented in [?], which is the first to provide
526 a DPI scheme that inspects GZIP compressed content. The ACCH algorithm gains up to 74% performance
527 improvement as compared to performing DPI with Aho-Corasick (i.e., three times faster). Still, it increases the
528 memory requirement by 25% (8KB per open connection). We suggest a scheme that benefits from the time
529 savings of ACCH and combines with SOP such that it reduces both the 32KB storage requirement of GZIP and
530 the 8KB of ACCH.

531 In order to understand the combined SOP-ACCH algorithm, we start with a short overview of ACCH. A
532 detailed description of the algorithm can be found in [?]. The general idea is that: Recall that GZIP compression
533 is done by compressing repeated strings. Whenever the input traffic contains a repeated string represented as an
534

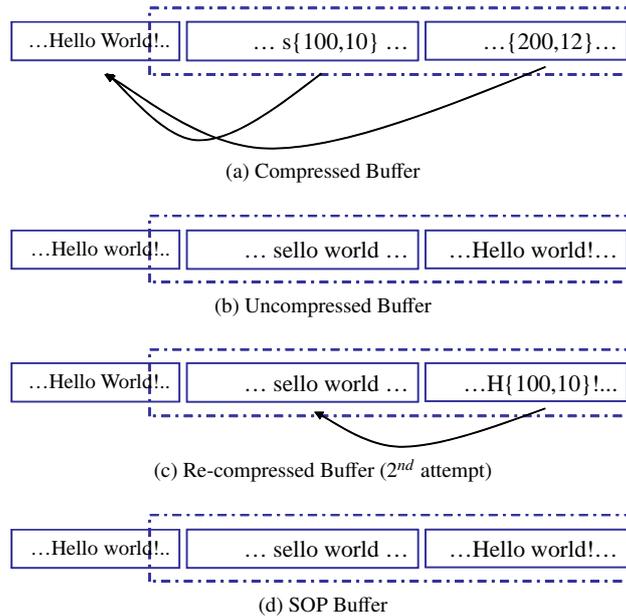


Figure 2: Sketch of the memory buffer in different scenarios.

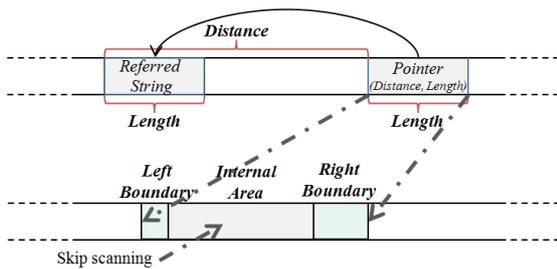


Figure 3: Illustration of *pointer* and its *referred string* with the *left boundary* and *right boundary* scan areas.

557 LZ77 pointer, ACCH uses the information about previ-
 558 ous scans of that string instead of re-scanning it. Since
 559 most of the compressed bytes are represented as point-
 560 ers, most of the byte scans can be saved.

561 In the general case, if the string that the pointer points
 562 to, denoted as the *referred string*, does not completely
 563 contain matched pattern then the pointer also contains
 564 none. In that case the algorithm may skip scanning
 565 bytes in the uncompressed data where the pointer oc-
 566 curs. Apart from that case ACCH handles three special
 567 cases where the first two refer to examining whether
 568 the string that the pointer represents (which is an identi-
 569 cal copy of the *referred string*), denoted by the *pointer*
 570 *string*, is part of a pattern and the third refers to the case
 571 where a pattern was found within the referred string. In
 572 the first case, the pattern starts prior to the pointer and

573 only its suffix is possibly in the *pointer string*. In the
 574 second case the *pointer string* contains a pattern prefix
 575 and its remaining bytes occur after the pointer. In order
 576 to detect those patterns we need to scan few bytes within
 577 the pointer starting and ending points denoted by pointer
 578 left boundary and right boundary scan areas respectively
 579 (as in Fig. ??). If no matches occurred within the refer-
 580 red string the algorithm skips the remaining bytes be-
 581 tween the pointer boundaries denoted by internal area.
 582 Note that left boundary, right boundary and internal area
 583 are parameters determined by the ACCH algorithm. In
 584 the third case a pattern ends within the referred string.
 585 The algorithm has to figure out whether the whole pat-
 586 tern is contained within the referred string or only its
 587 suffix.

588 The ACCH algorithm is based on the Aho-Corasick
 589 algorithm [?] which uses a DFA for content inspection
 590 as defined in Section ?. Handling those three cases
 591 relies strongly on using the DFA state *depth* param-
 592 eter (namely, the number of edges on the shortest simple
 593 path from the state and the DFA root). The *depth* of a
 594 state is the length of the *longest prefix* of any pattern
 595 within the pattern-set [?] at the current scan location.
 596 Using the *longest prefix* helps ACCH to determine the
 597 locations within the pointer string from where to stop or
 598 continue scanning at each one of the above cases. We
 599 elaborate on the way ACCH handles those three cases:

600 *Right Boundary (second case)*: In this case the al-
 601 gorithm determines whether some pattern starts

602 within the pointer suffix, i.e., whether a pattern pre- 654
603 fix is in the suffix of the pointer string. Here, ACCH 655
604 uses the stored information about previous scans in or- 656
605 der to determine the *longest prefix* of any pattern within 657
606 the pattern-set that ends at the last byte of the referred 658
607 string. ACCH continues its scan at the location of that 659
608 *longest prefix* within the pointer string which in turn de- 660
609 fines the right boundary area. 661

610 *Match in Referred String (third case):* This case is 662
611 handled similarly to the previous one. Using the in- 663
612 formation about previous scans, ACCH determines the 664
613 *longest prefix* of any pattern that ends at the last byte of 665
614 the matched pattern within the referred string. Know- 666
615 ing that, ACCH determines the point within the pointer 667
616 string from where to scan in order to check whether the 668
617 pointer string also contains the pattern which was found 669
618 while scanning the referred string. 670

619 *Left Boundary (first case):* In this case the algorithm 671
620 determines whether a pattern ends within the pointer left 672
621 boundary. The algorithm continue scanning the pointer 673
622 string (in the uncompressed data) until the number of 674
623 bytes scanned is equal or greater than the *depth* of the 675
624 DFA state of the current scanned symbol. At that point 676
625 the *longest prefix* of any pattern of the pattern-set is con- 677
626 tained within the pointer string, thus there cannot be any 678
627 pattern that started prior to the pattern string and the al- 679
628 gorithm may continue to one of the other two cases. 680

629 The ACCH algorithm described above increases the 681
630 memory requirement by 25% (8KB per open connec- 682
631 tion) in order to maintain the information about previ- 683
632 ous scans. In order to explain our scheme that reduces 684
633 this memory requirement we elaborate on the data struc- 685
634 ture that holds the previous scans information in ACCH, 686
635 called *Status Vector*. 687

636 Recall that ACCH uses the information about pre- 688
637 vious scans to determine the *longest prefix* in referred 689
638 strings for the *right boundary* and *match* cases. In 690
639 fact the vector does not store the exact '*longest-prefix*' 691
640 length parameter, but whether the prefix length is above 692
641 or below a certain threshold. Using that threshold al- 693
642 lows keeping only two status-indicators instead of stor- 694
643 ing all possible '*longest prefix*' lengths, which in turn 695
644 saves space. It is shown in [?] that using this thresh- 696
645 old causes only a few extra unnecessary byte scans but 697
646 does not cause misdetection of patterns. Specifically, 698
647 each vector entry represents the status of a byte within 699
648 the previously scanned uncompressed text. The entry 700
649 is 2-bit long and it stores three possible status values: 701
650 *Match*, in case that a pattern was located, and two other 702
651 states; If a prefix longer than a certain threshold was lo- 703
652 cated the status is *Check*, otherwise it is *Uncheck*. This 704
653 information suffices to get a good estimate of the *longest* 705

prefix that ends at any point within the referred string.

Evidently the 2-bit entry of the *Status Vector* causes the 25% space penalty per byte within the 32KB GZIP window (2 bits for each byte). Our suggested scheme stores the vector in a much more efficient way. Since there are long stretches with the same status, we indicate only a change of status. That is, we place a change-of-status indication after each byte at which such a change occurs. For example, in the case where the *Status Vector* contains 20 *Uncheck* status indicators in a row followed by 3 *Check* status indicators in a row, instead of using 23 pairs of bits as in the ACCH vector, we use only a single symbol that indicates a status change to *Check*, placed in the sequence after the 20 bytes of the input string. For that matter we define three new *Change-Status* symbols in the Huffman dictionary that indicate a status change to each one of the available statuses, namely *Check*, *Uncheck* and *Match*. The start status would always be *Uncheck*, hence it should not be specified. Figure ?? shows an example that demonstrates the usage of *Change-Status*. The *Change-Status* symbols S_u and S_c indicate a change of status, after the string 'hell' to *Uncheck* and after the string 'ab' to *Check* respectively.

Notice that *Change-Status* symbols are applicable only when they can be inserted between two other symbols. That does not apply in the case of a status-change within pointer bytes, since all those bytes are represented by a single *pointer-length* symbol (followed by *pointer-distance* symbol). Here we mark that ACCH is able to restore safely all pointer's statuses from the referred string except for those in the *Left Boundary* (as discussed in [?]). Thus we should only supply a solution that indicates status-changes that occur within the *Left Boundary* bytes.

In Figure ?? the *Left Boundary bits* of the first pointer indicate that there is a status change to *Uncheck* at the first byte and that the *Left Boundary* ends at that point. At the second pointer the *Left Boundary bits* indicate that there was no status change at all and that *Left Boundary* ends after two bytes. The *Left Boundary bits* are encoded by adding a pair of bits per byte in the *Left Boundary*, called *Left Boundary bits*, next to each pointer. The first bit indicates whether it is the last byte of the Left boundary. The second indicates whether a status change occurs within the corresponding byte. In such a case, a third bit indicates to which one of the two the status changes. Note that no special symbol is assigned to those bits within the Huffman dictionary. Instead, we place those bits at the end of each pointer (although those bits hold information about the beginning of the pointer). The *Left Boundary* is rather

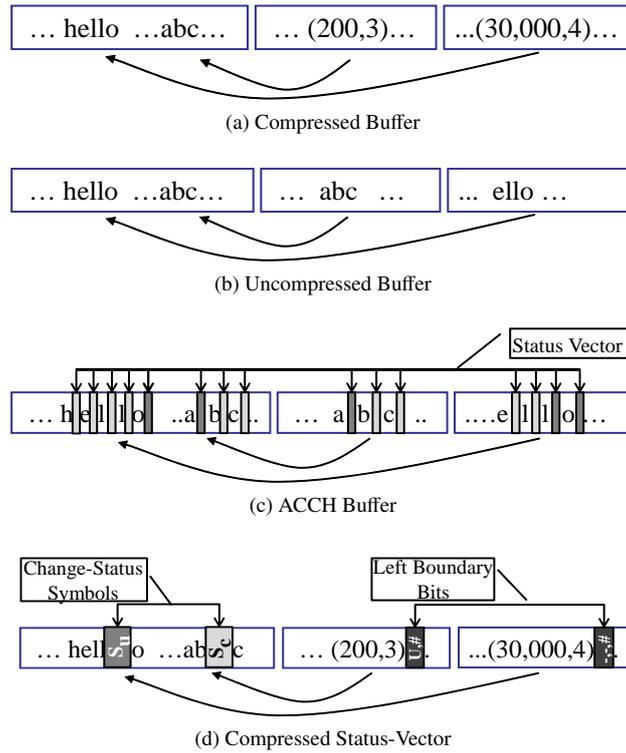


Figure 4: Sketch of the memory buffer including the *Status Vector*. (a) Original compressed traffic (b) Traffic in its uncompressed form (c) *Status Vector* packed with Uncompressed Buffer as in ACCH (d) Compressed *Status Vector*.

706 small and contains only 1.6 bytes on average, thus those
 707 *Left Boundary bits* impose a minor space penalty.

728 buffers and states from main memory. In the *realistic*
 729 *multi-connection* setup an input with interleaved pack-
 730 ets from different connections is tested, without flushing
 731 the cache between packets arrivals.

708 7. Experimental Results

709 7.1. Experimental Environment

710 Our performance results are given relative to the per-
 711 formance of a base algorithm *Plain*, which is the de-
 712 compression algorithm without packing, therefore the
 713 processor type is not an important factor for the exper-
 714 iments. The experiments were performed on a PC plat-
 715 form using Intel® Core™2 Duo CPU running at 1.8GHz
 716 using 32-bit Operating System. The system uses 1GB of
 717 Main Memory (RAM), a 256KB L2 Cache and 2×32KB
 718 write-back L1 data cache.

719 We base our implementation on the *zlib* [?] software
 720 library for the compression/decompression parts of the
 721 algorithm.

722 SOP is tested in two simulation setups: *simulated*
 723 *worst-case* setup and *realistic multi-connection* setup.
 724 In the *simulated worst-case* setup we flush the en-
 725 tire system’s cache upon each packet arrival to cre-
 726 ate the context-switch effect between multiple connec-
 727 tions. Thus every packet results in loading the required

732 7.2. Data Set

733 The data set consists of 2 308 HTML pages taken
 734 from the 500 most popular sites that use GZIP compres-
 735 sion. The web site list was constructed from the Alexa
 736 site [?], which maintains web traffic metrics and top-
 737 site lists. Total size of the uncompressed data is 359MB
 738 and in its compressed form it is 61.3MB.

739 While gathering the data-set from Alexa, we col-
 740 lected statistics about the percentage of the compressed
 741 pages among the sites as in Figure ???. The statistics
 742 shows high percentage of sites using compression, in
 743 particular among the top 1000 sites. As popularity drops
 744 the percentage slightly drops. Still, almost one out of
 745 every two sites uses compression.

746 The generation of the *global dictionary* as explained
 747 in Subsection ??? is based on the LZ77 symbol frequen-
 748 cies in a random sample of 500 compressed pages taken
 749 from a corporate firewall (a different data-set).

Packing Method	Symbol Coding Scheme	Average Buffer Size	Space Cost Ratio	Time Penalty
<i>Plain</i>	-	29.9KB	1	1
<i>OrigComp</i> (1 st attempt)	Original Huffman dict.	4.54KB	0.1518	-
<i>Re-compress</i> (2 nd attempt)	Original Huffman dict.	5.04KB	0.1576	20.77
<i>SOP</i>	Fixed size	7.33KB	0.245	3.91
<i>SOP</i>	Original Huffman dict.	6.28KB	0.211	3.89
<i>SOP</i>	Global Huffman dict.	5.17KB	0.172	3.85
<i>SOP-Indexed</i>	Global Huffman dict.	5.47KB	0.183	3.49

Table 1: Comparison of Time and Space parameters of different algorithms



Figure 5: Statistics of HTTP Compression usage among the Alexa [?] top-site lists

7.3. Space and Time Results

This subsection reports space and time results for our algorithm as shown in Table ?? . We compare SOP to several *Packing Methods* as described in Section ?? using the three symbol *Coding Schemes* as described in Section ?? . We define *Plain* as the basic algorithm that performs decompression and maintains a buffer of plain uncompressed data for each connection. The *Average Buffer Size* column indicates the average number of bytes used to store the data of the 32KB window. We are interested in the average ratio rather than in the maximum number of bytes because we try to lower the space requirement in a multi-connection setting, therefore the parameter of a single connection is of less importance. Notice that at the beginning of a connection the buffer stores less than 32KB since it stores only as much data as received. Therefore the average buffer size of *Plain* is 29.9KB which is slightly lower than the maximum value of 32KB. We use *Plain* as a reference for performance comparison to the other proposed methods and

set its time and space ratios to 1. The *Space Cost* column indicates the average buffer size of the matching scheme divided by the *Plain* average buffer size. The *Time Penalty* column indicates the time it takes to perform the DPI with the packing and unpacking operations with one method as compared to the same with the *Plain* algorithm.

We measure the size of the incoming compressed data representing the buffer and call it *OrigComp* as the 1st attempt, which is described in Section ?? . We use this method as a yard stick and a lower bound for the other methods. The average buffer size required by *OrigComp* is 4.54KB.

We define *Re-compress* (2nd attempt), as the method that compresses each buffer from scratch using GZIP. This method represents the best *practical* space result among the compared schemes, but has the worst time requirements of more than 20 times higher than *Plain*.

Next we used the *SOP* packing method with three different symbol coding schemes as described in Sec-

790 tion ??:

- 791 • *Fixed Size coding* scheme.
- 792 • *Original Huffman dictionary* as maintained from
793 the compressed connection.
- 794 • *Global Huffman dictionary*.

795 The *Global Huffman dictionary* improves the space
796 requirement by a factor of more than 3% compared with
797 the original dictionary and by more than 7% as com-
798 pared to using the *Fixed Size* coding. This underlines
799 the importance of the Huffman coding part in the space
800 reduction as the literal/pointer ratio grows.

801 *SOP* takes 3.85 more time than *Plain*. This is a
802 moderate time penalty as compared to *Re-compress*, the
803 space requirement of *SOP* is 5.17KB which is pretty
804 close to the 5.04KB of *Re-compress*. The small space
805 advantage gained by *Re-compress* in given its poor time
806 requirement makes it irrelevant as a solution.

807 We also examined the *SOP-Indexed* method, which
808 maintains indices to chunk offsets within the com-
809 pressed buffer to support a partial decompression of the
810 required chunks only (see Section ??). We used a 256B
811 chunks and got an average of 69% chunk accessed. The
812 time ratio is improved to 3.49 as compared to *Plain*.
813 The space penalty for maintaining the index vector and
814 chunk padding is of 0.3KB.

815 We simulated *SOP* with the *Global dictionary* using
816 our *realistic multi-connection* mode environment. First
817 we used an input with 10 interleaved connections at a
818 time. The simulations results showed time improve-
819 ment of 10.5% as compared to the *simulated worst-case*
820 mode. When we used an input with 100 interleaved con-
821 nections, we got a moderate improvement of 2.8% as
822 compared to the *simulated worst-case* mode.

823 We simulated *SOP* with the *Global dictionary* using
824 our *realistic multi-connection* setup. Two tests of the
825 realistic setup were carried out, one with 10 concurrent
826 connections and one with 100. The former performed
827 10.5% faster than the *simulated worst-case* setup and
828 the latter outperformed the worst-case by 2.8%. This
829 behavior is due to the fact that the data of a connec-
830 tion in the cache may be overwritten by the active con-
831 nection. The more concurrent connections there are the
832 higher the probability that the connections collide in the
833 cache. Thus when a connection becomes active again
834 it has to reload its data into the cache, penalizing its
835 performances. Notice that in the *simulated worst-case*
836 setup we purposely flushed the cache after the process-
837 ing of each packet.

838 7.4. Time Results Analysis

839 As explained in Section ??, *SOP* decompresses each
840 byte on average 4.2 times. Hence one would expect
841 *SOP* to take 4.2 times more than *Plain*. Still *SOP* takes
842 only 3.85 times more. This can be explained by in-
843 specting the data structures that require main memory
844 accesses by each of the algorithms (which are respon-
845 sible to most of the time it takes for the DPI process).
846 *SOP* maintains in main memory the old and new packed
847 buffers (i.e., *oldPacked* and *newPacked* in Algorithm
848 ??) which are heavily accessed during packet process-
849 ing. *Plain* on the other hand, uses parts of the 32KB
850 buffer, taken from main memory. The memory refer-
851 ences made by *Plain* are directly to the 32KB buffer in
852 memory. We measured the relative part of the buffer
853 which is accessed by *Plain* using a method similar to
854 the one in Section ?? with chunks of 32B that resem-
855 bles cache blocks and found out that an average 40.3%
856 of the buffer is accessed. Contrary to *Plain*, *SOP* ex-
857 tracts the 32KB buffer from the old and new packed
858 buffers directly to the cache. Thus, references to this
859 buffer does not cause cache-misses. Only the parts of
860 the 32KB buffer that were replaced by *SOP* should be
861 written back to main memory. Therefore the penalty of
862 the 4.2 multiple decompressions is lower than expected.
863 Specifically, the average main-memory space used for
864 *SOP* data-structures as compared to 12KB (= .4 of the
865 32KB) used by *Plain*, which is 20% higher and explains
866 why the time performance of *SOP* is better than the ex-
867 pected 4.2.

868 7.5. DPI of Compressed Traffic

869 In this subsection we analyze the performance of the
870 combined DPI process. We focus on pattern matching,
871 which is a lighter DPI task than the regular expression
872 matching. We show that the processing time taken by
873 *SOP* is minor as compared to that taken by the pattern
874 matching task. Since the regular-expression matching
875 task is more expensive than pattern matching, the pro-
876 cessing time of *SOP* is even less critical.

877 Table ?? summarizes the overall time and space
878 requirements for the different methods that imple-
879 ment pattern-matching of compressed traffic in multi-
880 connection environment. The time parameter is relative
881 to the time Aho-Corasick (AC) takes on uncompressed
882 traffic, as implemented by Snort [?]. Recall that AC
883 uses a DFA and basically performs one or two memory
884 references per scanned byte. The other pattern matching
885 algorithm we use for comparison is ACCH [?], which
886 is based on AC and takes advantage from the GZIP com-
887 pressed input, thus making the pattern matching process

Algorithms in Use			Average Time Ratio	Space per Buffer
Packing	Coding Scheme	Pattern Matching		
<i>Offline</i>	-	AC	-	170KB
<i>Plain</i>	-	AC	1.1	29.9KB
<i>Plain</i>	-	ACCH	0.36	37.4KB
<i>SOP</i>	Global Dict.	AC	1.39	5.17KB
<i>SOP</i>	Global Dict.	ACCH	0.64	6.19KB

Table 2: Overview of Pattern Matching + GZIP processing

888 faster. However, the techniques used by ACCH are inde- 928
889 pendent from the actual AC implementation. The space 929
890 per buffer parameter measures the memory required for 930
891 every connection upon context switch that happens after 931
892 packet processing is finished, i.e., in *SOP* after packing. 932

893 Current network tools that deals with compressed 933
894 traffic, construct the entire TCP connection first, then 934
895 they decompress the data and only after that they scan 935
896 the uncompressed data. We call this method the offline- 936
897 processing method (*Offline*). The space penalty is calcu- 937
898 lated by adding the average size of compressed sessions 938
899 upon TCP reconstruction and is 170KB per connection, 939
900 which is an enormous space requirement in terms of
901 mid-range security tool.

902 We use AC processing as the basic time penalty for 940
903 DPI and normalize it to 1. We measured the average 941
904 time overhead of the GZIP decompression phase, which 942
905 precedes the DPI procedure upon each packet arrival 943
906 and found out that it is 0.101, that is together with the 944
907 AC the time penalty is 1.101. We note that when per- 945
908 forming the same test on a single connection, the GZIP- 946
909 decompression time overhead is 0.035. The difference 947
910 is due to the context switch upon each packet on the 948
911 our setup that harms the decompression spatial locality 949
912 property. 950

913 As explained in Section ??, ACCH improves the time 951
914 requirement of the pattern matching process. Recall 952
915 that in order to apply the ACCH algorithm we need to 953
916 store in memory an additional data structure called *Sta-* 954
917 *tus Vector*. Applying the suggested compression algo- 955
918 rithm to the *Status Vector* as described in Section ??, 956
919 compressed it to 1.03KB. Therefore the total space re- 957
920 quirement of *SOP* combined with ACCH is 6.19KB. 958

921 The best time is achieved using *Plain* with ACCH 959
922 but the space requirement is very high as compared to 960
923 all other methods apart from *Offline*. Combining *SOP* 961
924 with ACCH achieves almost 80% space improvement 962
925 and above 40% time improvement comparing to com- 963
926 bining *Plain* with AC.

927 As we look at the greater picture that involves also

DPI, we need to refer to the space requirement applied
by its data structures. As noted before, the DPI process
itself has a large memory requirement. As opposed to
the decompression process whose space requirement is
proportional to the number of concurrent connections,
the DPI space requirements depend mainly on the num-
ber of patterns that it supports. We assume that DPI
space requirements are at the same order as those of
decompression for mid-range network tools. Thus the
80% space improvement for the decompression buffers
translates to a 40% space improvement for the overall
process.

8. Conclusions

With the sharp increase in cellular web surfing, HTTP
compression becomes common in today web traffic. Yet
due to its high memory requirements, most security de-
vices tend to ignore or bypass the compressed traffic
and thus introduce either a security hole or a potential
for a denial-of-service attack. This paper presents the
SOP technique, which drastically reduces this space re-
quirement by over 80% with only a slight increase in
the time overhead. It makes realtime compressed traffic
inspection a viable option for network devices. We also
present an algorithm that combines *SOP* with ACCH, a
technique that reduces the time required in DPI of com-
pressed HTTP [?]. The combined technique achieves
improvements of about 80% in space and above 40%
in the time complexity of the overall DPI processing of
compressed web-traffic. Note that the ACCH algorithm
(thus the combined algorithm) is not intrusive to the
Aho-Corasick (AC) algorithm, and it may be replaced
and thus enjoy the benefit of any DFA based algorithm
including recent improvements of AC [? ? ?].

9. Acknowledgment

The research leading to these results has received
funding from the European Research Council under

964 the European Union's Seventh Framework Programme 1025
965 (FP7/2007-2013)/ERC Grant agreement n°259085. 1026

966 References 1027

- 967 [] A. Bremler-Barr, Y. Koral, Accelerating multi-patterns match- 1032
968 ing on compressed http traffic., in: INFOCOM, IEEE, 2009, pp. 1033
969 397–405. 1034
- 970 [] Y. Afek, A. Bremler-Barr, Y. Koral, Efficient processing of 1035
971 multi-connection compressed web traffic, in: Proceedings of the 1036
972 10th international IFIP TC 6 conference on Networking - Vol- 1037
973 ume Part I, NETWORKING'11, Springer-Verlag, Berlin, Hei- 1038
974 delberg, 2011, pp. 52–65. 1039
- 975 [] Hypertext transfer protocol – http/1.1, June 1999. RFC 2616, 1040
976 <http://www.ietf.org/rfc/rfc2616.txt>. 1041
- 977 [] J. Ziv, A. Lempel, A universal algorithm for sequential data 1042
978 compression, IEEE Transactions on Information Theory 23 1043
979 (1977) 337–343. 1044
- 980 [] D. A. Huffman, A method for the construction of minimum- 1045
981 redundancy codes, Proceedings of the Institute of Radio Engi- 1046
982 neers 40 (1952) 1098–1101. 1047
- 983 [] zlib 1.2.5, April 2010. [Http://www.zlib.net](http://www.zlib.net). 1048
- 984 [] M. Fisk, G. Varghese, An analysis of fast string matching 1049
985 applied to content-based forwarding and intrusion detection, 1050
986 Technical Report CS2001-0670 (updated version) (2002). 1051
- 987 [] A. V. Aho, M. J. Corasick, Efficient string matching: an aid to 1052
988 bibliographic search, Commun. ACM 18 (1975) 333–340. 1053
- 989 [] R. S. Boyer, J. S. Moore, A fast string searching algorithm, 1054
990 Commun. ACM 20 (1977) 762–772. 1055
- 991 [] N. Tuck, T. Sherwood, B. Calder, G. Varghese, Deterministic 1056
992 memory-efficient string matching algorithms for intrusion de- 1057
993 tection, in: INFOCOM 2004. Twenty-third Annual Joint Con- 1058
994 ference of the IEEE Computer and Communications Societies, 1059
995 volume 4, pp. 2628 – 2639 vol.4. 1060
- 996 [] Clam antivirus, 2010. [Http://www.clamav.net](http://www.clamav.net) (version 0.82). 1061
- 997 [] Snort, 2010. [Http://www.snort.org](http://www.snort.org) (accessed on May 2010).
- 998 [] T. Song, W. Zhang, D. Wang, Y. Xue, A memory efficient mul-
999 tiple pattern matching architecture for network security, in: IN-
1000 FOCOM, pp. 166–170.
- 1001 [] J. van Lunteren, High-performance pattern-matching for intru-
1002 sion detection., in: INFOCOM, IEEE, 2006.
- 1003 [] V. Dimopoulos, I. Papaefstathiou, D. N. Pnevmatikatos, A
1004 memory-efficient reconfigurable aho-corasick fsm implementa-
1005 tion for intrusion detection systems., in: H. Blume, G. Gaydad-
1006 jiev, C. J. Glossner, P. M. W. Knijnenburg (Eds.), ICSAMOS,
1007 IEEE, 2007, pp. 186–193.
- 1008 [] M. Alicherry, M. Muthuprasanna, V. Kumar, High speed pat-
1009 tern matching for network ids/ips, in: Proceedings of the Pro-
1010 ceedings of the 2006 IEEE International Conference on Network
1011 Protocols, IEEE Computer Society, Washington, DC, USA,
1012 2006, pp. 187–196.
- 1013 [] L. Tan, T. Sherwood, Architectures for bit-split string scanning
1014 in intrusion detection, IEEE Micro 26 (2006) 110–117.
- 1015 [] A. Amir, G. Benson, M. Farach, Let sleeping files lie: Pattern
1016 matching in z-compressed files, Journal of Computer and Sys-
1017 tem Sciences (1996) 299–307.
- 1018 [] T. Kida, M. Takeda, A. Shinohara, S. Arikawa, Shift-and ap-
1019 proach to pattern matching in lzw compressed text, in: 10th
1020 Annual Symposium on Combinatorial Pattern Matching (CPM
1021 99).
- 1022 [] G. Navarro, M. Raffinot, A general practical approach to pat-
1023 tern matching over ziv-lempel compressed text, in: 10th Annual
1024 Symposium on Combinatorial Pattern Matching (CPM 99).
- [] G. Navarro, J. Tarhio, Boyer-moore string matching over ziv-
lempel compressed text, in: Proceedings of the 11th Annual
Symposium on Combinatorial Pattern Matching, pp. 166 – 180.
- [] M. Farach, M. Thorup, String matching in lempel-ziv com-
pressed strings, in: 27th annual ACM symposium on the theory
of computing, pp. 703–712.
- [] L. Gasieniec, M. Karpinski, W. Plandowski, W. Rytter, Ef-
ficient algorithms for lempel-ziv encoding (extended abstract),
in: In Proc. 4th Scandinavian Workshop on Algorithm Theory,
SpringerVerlag, 1996, pp. 392–403.
- [] S. Klein, D. Shapira, A new compression method for com-
pressed matching, in: Proceedings of data compression confe-
rence DCC-2000, Snowbird, Utah, pp. 400–409.
- [] A. Bremler-Barr, Y. Koral, V. Zigdon, Multi-pattern matching
in compressed communication traffic, in: Workshop on High
Performance Switching and Routing (HPSR 2011), Cartagena,
Spain.
- [] S. Wu, U. Manber, A fast algorithm for multi-pattern searching,
Technical Report TR94-17 (May 1994).
- [] M. F. Oberhumer, LZO, April 2010.
[Http://www.oberhumer.com/opensource/lzo](http://www.oberhumer.com/opensource/lzo).
- [] GB-800/GB-800e DataSheet, 2010. [Http://www.gta.com](http://www.gta.com).
- [] SonicWALL PRO 3060, 2010. [Http://www.sonicwall.com](http://www.sonicwall.com).
- [] StoneGate FW/VPN Appliances, 2010.
[Http://www.stonesoft.cn](http://www.stonesoft.cn).
- [] SRX5800 specification, 2010. [Http://www.juniper.net](http://www.juniper.net).
- [] Cisco ASA 5500 series, 2010. [Http://www.cisco.com](http://www.cisco.com).
- [] P. Deutsch, Gzip file format specification, May 1996. RFC 1952,
<http://www.ietf.org/rfc/rfc1952.txt>.
- [] T. H. Cormen, C. Stein, R. L. Rivest, C. E. Leiserson, Introduc-
tion to Algorithms, McGraw-Hill Higher Education, 2nd edi-
tion, 2001.
- [] Top sites, July 2010. [Http://www.alexa.com/topsites](http://www.alexa.com/topsites).
- [] W. Lin, B. Liu, Pipelined parallel ac-based approach for multi-
string matching, in: Proceedings of the 2008 14th IEEE Inter-
national Conference on Parallel and Distributed Systems, IEEE
Computer Society, Washington, DC, USA, 2008, pp. 665–672.