# A new family of Commentz-Walter-style multiple-keyword pattern matching algorithms

Bruce W. Watson

watson@OpenFIRE.org www.OpenFIRE.org UNIVERSITY OF PRETORIA (DEPARTMENT OF COMPUTER SCIENCE) Pretoria 0002, South Africa

### 1 Introduction

In this paper, I present a new family of Commentz-Walter-style multiplekeyword string pattern matching algorithm. The problem is: given a finite set of keywords P and an input string S, find all occurrences (including overlapping ones) of a keyword (in P) as a sub string of S; register such an occurrence by the end-point within S.

The algorithm family presented in this paper is, in fact, very closely related to the original Commentz-Walter algorithm. In [3, 4], all of the (then-known) Commentz-Walter-style algorithms were presented with a common algorithmic *skeleton* — they only differed in the *shift* function used to move through the input string. In this paper, we present a significantly improved skeleton, which is able to use the same shift functions (and, therefore, all of their precomputable forms). In addition to the changed skeleton, the output function (used to register the matches) is changed — though it remains easily computed. There have previously been numerous efforts to optimize the Commentz-Walter algorithms, though most of them have focused on new shift functions, more easily precomputed shift functions, and programming-language-specific coding tricks.

I assume that the reader is familiar with the field of multiple-keyword pattern matching, and with the Commentz-Walter algorithm in particular; for more information, see [3] or the bibliography therein. The style of presentation is in Dijkstra's guarded commands [2]. In such a style, a number of seemingly redundant variables may be presented — though this is usually only to make repetition (repetition) invariants easier to express. Furthermore, some variables are assumed to have inefficient types, such as 'sub string of the input S', whereas an implementation in C, C++ or Pascal would likely use indexing into S. The mathematical prerequisites are quite simple. Function suff maps a string or a set of strings to the corresponding set of suffixes (including the empty string  $\varepsilon$ ). In this paper, we use the *reverse trie* for P; such a trie consists of state set suff(P) (though, in practice, they would be represented using integers instead of strings), and transition function  $\tau_r$ , which maps a state and an input character to the corresponding extended element (that is, state) of suff(P) or to  $\perp$  if there is no such element. The notation  $z \upharpoonright 1$  refers to the rightmost symbol of z.

## 2 The Commentz-Walter algorithm

The Commentz-Walter algorithm was derived in the mid- to late-1970s, as a multiple-keyword generalization of the highly efficient Boyer-Moore style of pattern matching algorithm. It operates by moving from left to right through input string S. At each stop in S, a match attempt (conducted by an inner repetition) proceeds from right to left, registering any matches along the way. Following the inner repetition, the algorithm may make a shift to the right of more than one character — without missing any matches. This is done based on information gathered during the match attempt.

The following rendition of the algorithm is taken directly from [3], where  $\tau_r$  is the transition function of the *reverse trie* (there are also some other some minor adjustments of notation to avoid introducing many new definitions):

#### Algorithm 2.1 (Commentz-Walter skeleton):

Note that the states of the reverse trie are elements of the set suff(P), with  $\varepsilon$  being the start state.

The expression k(l, v, r) is a place-holder for any of the shift functions given in either the original Commentz-Walter article [1] or in [3, Chapter 4]. A great many of these shift functions were developed subsequent to the original article — and some of them have proven to have extremely high performance, while being reasonably easy to precompute.

The ability of these algorithms to skip large portions of input S (without examining every single character) yields high performance — usually exceeding that of the Aho-Corasick algorithm, as is shown in [3, Chapter 13]. (This begs the question: why is the Aho-Corasick algorithm more popular in text-books and implementations? The Commentz-Walter family of algorithms have worst-case performance which is substantially below that of the (stable) Aho-Corasick variants; the Aho-Corasick algorithms have also been particularly easy to implement compared to the Commentz-Walter algorithms.)

Despite this superb performance, profiling the Commentz-Walter skeleton shows some interesting facts:

- 1. A great deal of the time is spend making transitions (using the reverse trie) in the inner repetition.
- 2. A similar amount of time is spent checking (also in the inner repetition) whether a match has been encountered.

The first item appears (at this time) unavoidable, whereas the second one is amenable to optimization (as shown in the rest of this paper). The registration of matches can be expressed as a *repetition-invariant computation*, meaning that it can be removed from the inner repetition, dramatically improving performance.

#### 3 Some observations

The inner repetition iterates (using variable v) over the set  $suff(u) \cap suff(P)$ . We can characterize this set in a way which will prove easier to implement. In the following derivation, let w be the longest element of  $suff(u) \cap suff(P)$  (there is a unique longest element, since suff(u) is linearly ordered according to the suffix relation):

$$\begin{aligned} & \operatorname{suff}(u) \cap \operatorname{suff}(P) \\ &= \{ \operatorname{split} \operatorname{suff}(u) \text{ into } \{ z \mid z \in \operatorname{suff}(u) \land |z| > |w| \} \text{ and } \operatorname{suff}(w) \} \\ & (\{ z \mid z \in \operatorname{suff}(u) \land |z| > |w| \} \cup \operatorname{suff}(w)) \cap \operatorname{suff}(P) \\ &= \{ \operatorname{distribute} \cap \operatorname{over} \cup \} \\ & (\{ z \mid z \in \operatorname{suff}(u) \land |z| > |w| \} \cap \operatorname{suff}(P)) \cup (\operatorname{suff}(w) \cap \operatorname{suff}(P)) \\ &= \{ \operatorname{by} w \text{ being the longest}, \{ z \mid z \in \operatorname{suff}(u) \land |z| > |w| \} \cap \operatorname{suff}(P) = \emptyset \} \\ & \operatorname{suff}(w) \cap \operatorname{suff}(P) \end{aligned}$$

The inner repetition registers, as matches, only those elements  $v : v \in P$ . It follows that the set registered during a single iteration of the outer repetition is  $v \in suff(u) \cap suff(P) \cap P$ , which can be rewritten:

$$\{ v \mid v \in \mathsf{suff}(u) \cap \mathsf{suff}(P) \land v \in P \}$$

$$= \{ \text{set comprehension} \}$$

$$\{ v \mid v \in \mathsf{suff}(u) \cap \mathsf{suff}(P) \cap P \}$$

$$= \{ \text{set comprehension} \}$$

$$\mathsf{suff}(u) \cap \mathsf{suff}(P) \cap P$$

$$= \{ \text{previous derivation } \mathsf{suff}(u) \cap \mathsf{suff}(P) = \mathsf{suff}(w) \cap \mathsf{suff}(P) \}$$

$$\mathsf{suff}(w) \cap \mathsf{suff}(P) \cap P$$

$$= \{ \text{set containment: } P \subseteq \mathsf{suff}(P) \}$$

$$\mathsf{suff}(w) \cap P$$

This yields a set by which O could be updated after, instead of during, the inner repetition.

The remaining problem is to calculate w. Fortunately, the post-condition of the inner repetition is that v is the longest element of suff(u) which is still an element of suff(P) - v = w.

For conciseness, we can define function WCWOutput  $\in$  suff $(P) \longrightarrow \mathcal{P}(P)$  as

 $\mathsf{WCWOutput}(x) = \mathsf{suff}(x) \cap P$ 

## 4 The new algorithm

Given the repetition-invariance of the updates (to variable O, in which matches are registered), we can present the new algorithm skeleton:

#### Algorithm 4.1 (Watson-Commentz-Walter skeleton):

 $\begin{array}{l} u,r:=\varepsilon,S;\\ l,v:=\varepsilon,\varepsilon;\\ O:=(\{\varepsilon\}\cap P)\times\{S\};\\ \{\text{ invariant: } u=lv\wedge v\in \mathrm{suff}(P) \ \} \\ \mathbf{do} \ r\neq\varepsilon\rightarrow\\ u,r:=\mathrm{shift} \ u,r \ \mathrm{right} \ \mathrm{by} \ k(l,v,r);\\ l,v:=u,\varepsilon;\\ O:=O\cup(\{\varepsilon\}\cap P)\times\{r\};\\ \{ \ \mathrm{invariant:} \ u=lv\wedge v\in \mathrm{suff}(P) \ \} \\ \mathbf{do} \ l\neq\varepsilon\wedge\tau_r(v,l\upharpoonright 1)\neq\perp\rightarrow\\ l,v:=\mathrm{shift} \ l,v \ \mathrm{right} \ \mathrm{by} \ \mathrm{one \ character};\\ \mathbf{od};\\ O:=O\cup\mathsf{WCWOutput}(v)\times\{r\}\\ \mathbf{od}\\ \{ \ O=\{(p,r)\mid ur=S\wedge p\in \mathrm{suff}(u) \ \} \ \} \end{array}$ 

Note that we retain the invariants found in the original skeleton, meaning that exactly the same shift functions may be used. The precomputation of WCWOutput remains to be discussed. As of this writing, an efficient algorithm has not been derived, though it is clearly possible to precompute it via a brute-force method. In all likelihood, an efficient precomputation algorithm will involve a breadth-first traversal of the reverse trie (the states of which, suff(P), form the domain of WCWOutput). This will be discussed further in the final version of this paper.

### 5 Closing comments

In this section, we briefly discuss the expected<sup>1</sup> performance and the history of this algorithm variant.

Examining the generated machine code (from a compiler such as Gnu C++) reveals that the new skeleton's inner repetition consists of roughly half the number of instructions as in the original algorithm's skeleton. Since much (upwards of 80%) of the execution time is spend in the inner repetition, we can expect WCW to run in roughly 60

The derivation presented here is markedly different from the one first obtained in 1999. In the first derivation, I started with a naïve algorithm

<sup>&</sup>lt;sup>1</sup> Only the expected performance is presented here, since benchmarking is ongoing and will be presented in the final version of the paper.

which was used as a precursor to the Aho-Corasick algorithm in [3]. From that algorithm, a simple adjustment (in the step used to obtain the Aho-Corasick algorithm) yields the invariant that we need to register matches  $\operatorname{suff}(u) \cap \operatorname{suff}(P) \cap P$ . As in this paper, this set can be characterized by its longest element — eventually leading to the WCW algorithm. The alternative derivation is, in some sense, preferrable since it also allows for sets other than  $\operatorname{suff}(P)$  in the registration process. In particular,  $\operatorname{pref}(P)$ can be used (leading to the Aho-Corasick algorithm); more interestingly, the set of *factors* of P can be used — perhaps creating a new algorithm.

#### References

- 1. Commentz-Walter, B. A string matching algorithm fast on the average. (Technical Report TR 79.09.007, IBM Germany, Heidelberg Scientific Center, 1979).
- Dijkstra, E.W. A Discipline of Programming. (Prentice Hall, Englewood Cliffs, N.J., 1976).
- 3. Watson, B.W. Taxonomies and Toolkits of Regular Language Algorithms. (Ph.D dissertation, Eindhoven University of Technology, The Netherlands, 1995). See www.OpenFIRE.org
- Watson, B.W. and G. Zwaan. A taxonomy of sublinear multiple keyword pattern matching algorithms. (Science of Computer Programming 27(2), September 1996, pp. 85-118).