

Proposed Update Unicode Technical Standard #10

UNICODE COLLATION ALGORITHM

Version	6.3.0 (draft 5)
Editors	Mark Davis (markdavis@google.com), Ken Whistler (ken@unicode.org), Markus Scherer
Date	2013-03-15
This Version	http://www.unicode.org/reports/tr10/tr10-27.html
Previous Version	http://www.unicode.org/reports/tr10/tr10-26.html
Latest Version	http://www.unicode.org/reports/tr10/
Latest Proposed Update	http://www.unicode.org/reports/tr10/proposed.html
Revision	<u>27</u>

Summary

This report is the specification of the Unicode Collation Algorithm (UCA), which details how to compare two Unicode strings while remaining conformant to the requirements of the Unicode Standard. The UCA also supplies the Default Unicode Collation Element Table (DUCET) as the data specifying the default collation order for all Unicode characters.

Status

*This is a **draft** document which may be updated, replaced, or superseded by other documents at any time. Publication does not imply endorsement by the Unicode Consortium. This is not a stable document; it is inappropriate to cite this document as other than a work in progress.*

A Unicode Technical Standard (UTS) is an independent specification.
Conformance to the Unicode Standard does not imply conformance to any UTS.

Please submit corrigenda and other comments with the online reporting form [[Feedback](#)]. Related information that is useful in understanding this document is found in the [References](#). For the latest version of the Unicode Standard see [[Unicode](#)]. For a

list of current Unicode Technical Reports see [[Reports](#)]. For more information about versions of the Unicode Standard, see [[Versions](#)].

Contents

- 1 [Introduction](#)
 - 1.1 [Multi-Level Comparison](#)
 - 1.1.1 [Collation Order and Code Chart Order](#)
 - 1.2 [Canonical Equivalence](#)
 - 1.3 [Contextual Sensitivity](#)
 - 1.4 [Customization](#)
 - 1.5 [Other Applications of Collation](#)
 - 1.6 [Merging Sort Keys](#)
 - 1.7 [Performance](#)
 - 1.8 [What Collation is Not](#)
 - 1.9 [The Unicode Collation Algorithm](#)
 - 1.9.1 [Goals](#)
 - 1.9.2 [Non-Goals](#)
- 2 [Conformance](#)
- 3 [Collation Element Table](#)
 - 3.1 [Weight Levels and Notation](#)
 - 3.2 [Simple Mappings](#)
 - 3.3 [Multiple Mappings](#)
 - 3.3.1 [Expansions](#)
 - 3.3.2 [Contractions](#)
 - 3.3.3 [Many-to-Many Mappings](#)
 - 3.3.4 [Other Multiple Mappings](#)
 - 3.4 [Backward Accents](#)
 - 3.5 [Rearrangement](#)
 - 3.6 [Variable Weighting](#)
 - 3.7 [Well-Formed Collation Element Tables](#)
 - 3.8 [Default Unicode Collation Element Table](#)
 - 3.8.1 [Default Values](#)
 - 3.8.2 [Well-Formedness of the DUCET](#)
 - 3.8.3 [Stability of the DUCET](#)
 - 3.8 [Stability](#)
 - 3.8.1 [Stable Sort](#)
 - 3.8.2 [Deterministic Comparison](#)
- 4 [Main Algorithm](#)
 - 4.1 [Normalize](#)
 - 4.2 [Produce Array](#)
 - 4.3 [Form Sort Key](#)
 - 4.4 [Compare](#)
 - 4.5 [Rationale for Well-Formed Collation Element Tables](#)
- 5 [Tailoring](#)
 - 5.1 [Parametric Tailoring](#)
 - 5.2 [Tailoring Example](#)
 - 5.3 [Use of Combining Grapheme Joiner](#)
 - 5.4 [Preprocessing](#)
- 6 [Implementation Notes](#)
 - 6.1 [Reducing Sort Key Lengths](#)
 - 6.1.1 [Eliminating Level Separators](#)

- [6.1.2 L2/L3 in 8 Bits](#)
 - [6.1.3 Machine Words](#)
 - [6.1.4 Run-Length Compression](#)
 - [6.2 Large Weight Values](#)
 - [6.3 Reducing Table Sizes](#)
 - [6.3.1 Contiguous Weight Ranges](#)
 - [6.3.2 Escape Hatch](#)
 - [6.3.3 Leveraging Unicode Tables](#)
 - [6.3.4 Reducing the Repertoire](#)
 - [6.3.5 Memory Table Size](#)
 - [6.4 Avoiding Zero Bytes](#)
 - [6.5 Avoiding Normalization](#)
 - [6.6 Case Comparisons](#)
 - [6.7 Incremental Comparison](#)
 - [6.8 Catching Mismatches](#)
 - [6.9 Handling Collation Graphemes](#)
 - [6.10 Flat File Example](#)
 - [6.10.1 Collation Element Format](#)
 - [6.10.2 Sample Code](#)
 - [7 Weight Derivation](#)
 - [7.1 Derived Collation Elements](#)
 - [7.1.1 Handling Ill-Formed Code Unit Sequences](#)
 - [7.1.2 Unassigned and Other Code Points](#)
 - [7.1.3 Implicit Weights](#)
 - [7.1.4 Trailing Weights](#)
 - [7.1.5 Hangul Collation](#)
 - [7.2 Tertiary Weight Table](#)
 - [7.3 Fourth-Level Weight Assignments](#)
 - [8 Searching and Matching](#)
 - [8.1 Collation Folding](#)
 - [8.2 Asymmetric Search](#)
 - [8.2.1 Returning Results](#)
 - [9 Data Files](#)
 - [9.1 Allkeys File Format](#)
 - Appendix A: [Deterministic Sorting](#)
 - [A.1 Stable Sort](#)
 - [A.2 Deterministic Sort](#)
 - [A.3 Deterministic Comparison](#)
 - [A.4 Stable and Portable Comparison](#)
 - Appendix B: [Synchronization with ISO/IEC 14651](#)
 - [Acknowledgements](#)
 - [References](#)
 - [Migration Issues](#)
 - [Modifications](#)

1 Introduction

Collation is the general term for the process and function of determining the sorting order of strings of characters. It is a key function in computer systems; whenever a list

of strings is presented to users, they are likely to want it in a sorted order so that they can easily and reliably find individual strings. Thus it is widely used in user interfaces. It is also crucial for databases, both in sorting records and in selecting sets of records with fields within given bounds.

Collation varies according to language and culture: Germans, French and Swedes sort the same characters differently. It may also vary by specific application: even within the same language, dictionaries may sort differently than phonebooks or book indices. For non-alphabetic scripts such as East Asian ideographs, collation can be either phonetic or based on the appearance of the character. Collation can also be customized according to user preference, such as ignoring punctuation or not, putting uppercase before lowercase (or vice versa), and so on. Linguistically correct *searching* needs to use the same mechanisms: just as "v" and "w" sort as if they were the same base letter in Swedish, a loose search should pick up words with either one of them.

Collation implementations must deal with the complex linguistic conventions for ordering text in specific languages, and provide for common customizations based on user preferences. Furthermore, algorithms that allow for good performance are crucial for any collation mechanisms to be accepted in the marketplace.

Table 1 shows some examples of cases where sort order differs by language, usage, or another customization.

Table 1. Example Differences

Language	Swedish:	z < ö
	German:	ö < z
Usage	German Dictionary:	of < öf
	German Telephone:	öf < of
Customizations	Upper-First	A < a
	Lower-First	a < A

Languages vary regarding which types of comparisons to use (and in which order they are to be applied), and in what constitutes a fundamental element for sorting. For example, Swedish treats *ä* as an individual letter, sorting it after z in the alphabet; German, however, sorts it either like *ae* or like other accented forms of *a*, thus following *a*. In Slovak, the digraph *ch* sorts as if it were a separate letter after *h*. Examples from other languages and scripts abound. Languages whose writing systems use uppercase and lowercase typically ignore the differences in case, unless there are no other differences in the text.

It is important to ensure that collation meets user expectations as fully as possible. For example, in the majority of Latin languages, *ø* sorts as an accented variant of *o*, meaning that most users would expect *ø* alongside *o*. However, a few languages, such as Norwegian and Danish, sort *ø* as a unique element after z. Sorting "Søren" after "Sylt" in a long list, as would be expected in Norwegian or Danish, will cause problems if

the user expects ø as a variant of o. A user will look for "Søren" between "Sorem" and "Soret", not see it in the selection, and assume the string is missing, confused because it was sorted in a completely different location. In matching, the same can occur, which can cause significant problems for software customers; for example, in a database selection the user may not realize what records are missing. See [Section 1.5, *Other Applications of Collation*](#).

With Unicode applications widely deployed, multilingual data is the rule, not the exception. Furthermore, it is increasingly common to see users with many different sorting expectations accessing the data. For example, a French company with customers all over Europe will include names from many different languages. If a Swedish employee at this French company accesses the data from a Swedish company location, the customer names need to show up in the order that meets this employee's expectations—that is, in a Swedish order—even though there will be many different accented characters that do not normally appear in Swedish text.

For scripts and characters not used in a particular language, explicit rules may not exist. For example, Swedish and French have clearly specified, distinct rules for sorting ä (either after z or as an accented character with a secondary difference from a), but neither defines the ordering of characters such as Ж, ѱ, ♪, ∞, ⋄, or △.

1.1 Multi-Level Comparison

To address the complexities of language-sensitive sorting, a *multilevel* comparison algorithm is employed. In comparing two words, the most important feature is the identity of the base letters—for example, the difference between an *A* and a *B*. Accent differences are typically ignored, if the base letters differ. Case differences (uppercase versus lowercase), are typically ignored, if the base letters or their accents differ. Treatment of punctuation varies. In some situations a punctuation character is treated like a base letter. In other situations, it should be ignored if there are any base, accent, or case differences. There may also be a final, tie-breaking level (called an *identical* level), whereby if there are no other differences at all in the string, the (normalized) code point order is used.

Table 2. Comparison Levels

Level	Description	Examples
L1	Base characters	role < roles < rule
L2	Accents	role < r <u>ô</u> le < roles
L3	Case/Variants	role < <u>R</u> ole < r <u>ô</u> le
L4	Punctuation	role < “r <u>o</u> le” < Role
Ln	Identical	role < ro <u>□</u> le < “role”

The examples in *Table 2* are in English; the description of the levels may correspond to different writing system features in other languages. In each example, for levels L2 through Ln, the differences on that level (indicated by the underlined characters) are

swamped by the stronger-level differences (indicated by the blue text). For example, the L2 example shows that difference between an o and an accented ô is swamped by an L1 difference (the presence or absence of an s). In the last example, the □ represents a format character, which is otherwise completely ignorable.

The primary level (L1) is for the basic sorting of the text, and the non-primary levels (L2..Ln) are for adjusting string weights for other linguistic elements in the writing system that are important to users in ordering, but less important than the order of the basic sorting. In practice, fewer levels may be needed, depending on user preferences or customizations.

1.1.1 Collation Order and Code Chart Order

Many people expect the characters in their language to be in the "correct" order in the Unicode code charts. Because collation varies by language and not just by script, it is not possible to arrange the encoding for characters so that simple binary string comparison produces the desired collation order for all languages. Because multi-level sorting is a requirement, it is not even possible to arrange the encoding for characters so that simple binary string comparison produces the desired collation order for any particular language. Separate data tables are required for correct sorting order. For more information on tailorings for different languages, see [\[CLDR\]](#).

The basic principle to remember is: ***The position of characters in the Unicode code charts does not specify their sort order.***

1.2 Canonical Equivalence

There are many cases in Unicode where two sequences of characters are canonically equivalent: the sequences represent essentially the same text, but with different actual sequences. For more information, see [\[UAX15\]](#).

Sequences that are canonically equivalent must sort the same. *Table 3* gives some examples of canonically equivalent sequences. For example, the *angstrom sign* was encoded for compatibility, and is canonically equivalent to an *A-ring*. The latter is also equivalent to the decomposed sequence of *A* plus the *combining ring* character. The order of certain combining marks is also irrelevant in many cases, so such sequences must also be sorted the same, as shown in the second example. The third example shows a composed character that can be decomposed in four different ways, all of which are canonically equivalent.

Table 3. Canonical Equivalence

1	Å	U+212B ANGSTROM SIGN
	Å	U+00C5 LATIN CAPITAL LETTER A WITH RING ABOVE
	A ◌̊	U+0041 LATIN CAPITAL LETTER A, U+030A COMBINING RING ABOVE

2	x ◌ ◌	U+0078 LATIN SMALL LETTER X, U+031B COMBINING HORN, U+0323 COMBINING DOT BELOW
	x ◌ ◌	U+0078 LATIN SMALL LETTER X, U+0323 COMBINING DOT BELOW, U+031B COMBINING HORN
3	ʊ	U+1EF1 LATIN SMALL LETTER U WITH HORN AND DOT BELOW
	u◌	U+1EE5 LATIN SMALL LETTER U WITH DOT BELOW, U+031B COMBINING HORN
	u ◌ ◌	U+0075 LATIN SMALL LETTER U, U+031B COMBINING HORN, U+0323 COMBINING DOT BELOW
	ʊ ◌	U+01B0 LATIN SMALL LETTER U WITH HORN, U+0323 COMBINING DOT BELOW
	u ◌ ◌	U+0075 LATIN SMALL LETTER U, U+0323 COMBINING DOT BELOW, U+031B COMBINING HORN

1.3 Contextual Sensitivity

There are additional complications in certain languages, where the comparison is context sensitive and depends on more than just single characters compared directly against one another, as shown in *Table 4*.

The first example of such a complication consists of **contractions**, where two (or more) characters sort as if they were a single base letter. In the table below, *CH* acts like a single letter sorted after *C*.

The second example consists of **expansions**, where a single character sorts as if it were a sequence of two (or more) characters. In the table below, an *Œ* ligature sorts as if it were the sequence of *O + E*.

Both contractions and expansions can be combined: that is, two (or more) characters may sort as if they were a different sequence of two (or more) characters. In the third example, for Japanese, a length mark sorts `like` with only a tertiary difference from the vowel of the previous syllable: as an *A* after *KA* and as an *I* after *KI*.

Table 4. Context Sensitivity

Contractions	H < Z, but CH > CZ
Expansions	OE < Œ < OF

Both	カー < カイ, but
	キー > キイ
	カー < カア, but
	キー > キア

Some languages have additional oddities in the way they sort. Normally, all differences in sorting are assessed from the start to the end of the string. If all of the base letters are the same, the first accent difference determines the final order. In row 1 of *Table 5*, the first accent difference is on the *o*, so that is what determines the order. In *Canadian French* in some French dictionary ordering traditions, however, it is the *last* accent difference that determines the order, as shown in row 2.

Table 5. Backward Accent Ordering

Normal Accent Ordering	cote < coté < c ^o te < c ^o té
Backward Accent Ordering	cote < c ^o te < cot ^e < cot ^e

1.4 Customization

In practice, there are additional features of collation that users need to control. These are expressed in user-interfaces and eventually in APIs. Other customizations or user preferences include the following:

- *Language*. This is the most important feature, because it is crucial that the collation match the expectations of users of the target language community.
- *Strength*. This refers to the number of levels that are to be considered in comparison, and is another important feature. Most of the time a three-level strength is needed for comparison of strings. In some cases, a larger number of levels will be needed, while in others—especially in searching—fewer levels will be desired.
- *Case Ordering*. Some dictionaries and authors collate uppercase before lowercase while others use the reverse, so that preference needs to be customizable. Sometimes the case ordering is mandated by the government, as in Denmark. Often it is simply a customization or user preference.
- *Punctuation*. Another common option is whether to treat punctuation (including spaces) as base characters or treat such characters as only making a level 4 difference.
- *User-Defined Rules*. Such rules provide specified results for given combinations of letters. For example, in an index, an author may wish to have symbols sorted as if they were spelled out; thus "?" may sort as if it were the string "question mark".
- *Merged Tailorings*. An option may allow the merging of sets of rules for different languages. For example, someone may want Latin characters sorted as in French, and Arabic characters sorted as in Persian. In such an approach, generally one of the tailorings is designated the “master” in cases of conflicting weights for a given character.
- *Script Order*. A user may wish to specify which scripts come first. For example, in

a book index an author may want index entries in the predominant script that the book itself is written in to come ahead of entries for any other script. For example:

b < ב < β < б [Latin < Hebrew < Greek < Cyrillic] *versus*
β < b < б < ב [Greek < Latin < Cyrillic < Hebrew]

Attempting to achieve this effect by introducing an extra strength level before the first (primary) level would give incorrect ordering results for strings which mix characters of more than one script.

- *Numbers*. A customization may be desired to allow sorting numbers in numeric order. If strings including numbers are merely sorted alphabetically, the string “A-10” comes before the string “A-2”, which is often not desired. This behavior can be customized, but it is complicated by ambiguities in recognizing numbers within strings (because they may be formatted according to different language conventions). Once each number is recognized, it can be preprocessed to convert it into a format that allows for correct numeric sorting, such as a textual version of the IEEE numeric format.

Phonetic sorting of Han characters requires use of either a lookup dictionary of words or, more typically, special construction of programs or databases to maintain an associated phonetic spelling for the words in the text.

1.5 Other Applications of Collation

The same principles about collation behavior apply to realms beyond sorting. In particular, searching should behave consistently with sorting. For example, if *v* and *w* are treated as identical base letters in Swedish sorting, then they should also be treated the same for searching. The ability to set the maximal strength level is very important for searching.

Selection is the process of using the comparisons between the endpoints of a range, as when using a SELECT command in a database query. It is crucial that the range returned be correct according to the user's expectations. For example, if a German businessman making a database selection to sum up revenue in each of the cities from *O...* to *P...* for planning purposes does not realize that all cities starting with *Ö* were excluded because the query selection was using a Swedish collation, he will be one very unhappy customer.

A sequence of characters considered a unit in collation, such as *ch* in Slovak, represents a *collation grapheme cluster*. For applications of this concept, see Unicode Technical Standard #18, "Unicode Regular Expressions" [[UTS18](#)]. For more information on grapheme clusters, see Unicode Standard Annex #29, "Unicode Text Segmentation" [[UAX29](#)].

1.6 Merging Sort Keys

Sort keys may need to be merged. For example, the simplest way to sort a database according to two fields is to sort field by field, sequentially. This gives the results in column one in *Table 6*. (The examples in this table are ordered using the **Shifted** option

for handling variable collation elements such as the space character; see [Section 3.6.2 Variable Weighting](#) for details.) All the levels in Field 1 are compared first, and then all the levels in Field 2. The problem with this approach is that high-level differences in the second field are swamped by minute differences in the first field, which results in unexpected ordering for the first names.

Table 6. Merged Fields

Sequential		Weak First		Merged	
F1 _{L1} , F1 _{L2} , F1 _{L3} , F2 _{L1} , F2 _{L2} , F2 _{L3}		F1 _{L1} , F2 _{L1} , F2 _{L2} , F2 _{L3}		F1 _{L1} , F2 _{L1} , F1 _{L2} , F2 _{L2} , F1 _{L3} , F2 _{L3}	
di Silva	Fred	disílvá	Fred	di Silva	Fred
di Silva	John	diSilva	Fred	diSilva	Fred
diSilva	Fred	di Silva	Fred	disílvá	Fred
diSilva	John	di Silva	John	di Silva	John
disílvá	Fred	diSilva	John	diSilva	John
disílvá	John	disílvá	John	disílvá	John

A second way to do the sorting is to ignore all but base-level differences in the sorting of the first field. This gives the results in the second column. The first names are all in the right order, but the problem is now that the first field is not correctly ordered except by the base character level.

The correct way to sort two fields is to merge the fields, as shown in the "Merged" column. Using this technique, all differences in the fields are taken into account, and the levels are considered uniformly. Accents in all fields are ignored if there are any base character differences in any of the field, and case in all fields is ignored if there are accent or base character differences in any of the fields.

1.7 Performance

Collation is one of the most performance-critical features in a system. Consider the number of comparison operations that are involved in sorting or searching large databases, for example. Most production implementations will use a number of optimizations to speed up string comparison.

Strings are often preprocessed into sort keys, so that multiple comparisons operations are much faster. With this mechanism, a collation engine generates a *sort key* from any given string. The binary comparison of two sort keys yields the same result (less, equal, or greater) as the collation engine would return for a comparison of the original strings. Thus, for a given collation C and any two strings A and B:

$A \leq B$ according to C if and only if $\text{sortkey}(C, A) \leq \text{sortkey}(C, B)$

However, simple string comparison is faster for any individual comparison, because the generation of a sort key requires processing an entire string, while differences in most string comparisons are found before all the characters are processed. Typically, there is a considerable difference in performance, with simple string comparison being about 5 to 10 times faster than generating sort keys and then using a binary comparison.

Sort keys, on the other hand, can be much faster for multiple comparisons. Because binary comparison is much faster than string comparison, it is faster to use sort keys whenever there will be more than about 10 comparisons per string, if the system can afford the storage.

1.8 What Collation is Not

There are a number of common expectations about and misperceptions of collation. This section points out many things that collation is not and cannot be.

Collation is not aligned with character sets or repertoires of characters.

Swedish and German share most of the same characters, for example, but have very different sorting orders.

Collation is not code point (binary) order.

A simple example of this is the fact that capital Z comes before lowercase a in the code charts. As noted earlier, beginners may complain that a particular Unicode character is “not in the right place in the code chart.” That is a misunderstanding of the role of the character encoding in collation. While the Unicode Standard does not gratuitously place characters such that the binary ordering is odd, the only way to get the linguistically-correct order is to use a language-sensitive collation, not a binary ordering.

Collation is not a property of strings.

In a list of cities, with each city correctly tagged with its language, a German user will expect to see all of the cities sorted according to German order, and will not expect to see a word with ö appear after z, simply because the city has a Swedish name. As in the earlier example, it is crucially important that if a German businessman makes a database selection, such as to sum up revenue in each of the cities from O... to P... for planning purposes, cities starting with Ö *not* be excluded.

Collation order is not preserved under concatenation or substring operations, in general.

For example, the fact that x is less than y does not mean that $x + z$ is less than $y + z$, because characters may form contractions across the substring or concatenation boundaries. In summary:

$x < y$ does not imply that $xz < yz$
 $x < y$ does not imply that $zx < zy$
 $xz < yz$ does not imply that $x < y$
 $zx < zy$ does not imply that $x < y$

Collation order is not preserved when comparing sort keys generated from different collation sequences.

Remember that sort keys are a preprocessing of strings according to a given set of collation features. Different features result in different binary sequences. For example, if there are two collations, F and G, where F is a French collation, and G is a German phonebook ordering, then:

- $A \leq B$ according to F if and only if $\text{sortkey}(F, A) \leq \text{sortkey}(F, B)$, *and*
- $A \leq B$ according to G if and only if $\text{sortkey}(G, A) \leq \text{sortkey}(G, B)$
- The relation between $\text{sortkey}(F, A)$ and $\text{sortkey}(G, B)$ says nothing about whether $A \leq B$ according to F, or whether $A \leq B$ according to G.

Collation order is not a stable sort.

Stability is a property of a sort algorithm, not of a collation sequence. [For more information, see Section 3.8, *Stability*.](#)

3.8.1 Stable Sort

A *stable sort* is one where two records with a field that compares as equal will retain their order if sorted according to that field. This is a property of the sorting algorithm, *not* of the comparison mechanism. For example, a bubble sort is stable, while a Quicksort is not. This is a useful property, but cannot be accomplished by modifications to the comparison mechanism or tailorings. [See also Appendix A, *Deterministic Sorting*.](#)

3.8.2 Deterministic Comparison

A *deterministic comparison* is different. It is a comparison in which strings that are not canonical equivalents will not be judged to be equal. This is a property of the comparison, not of the sorting algorithm. This is not a particularly useful property—its implementation also requires extra processing in string comparison or an extra level in sort keys, and thus may degrade performance to little purpose. However, if a deterministic comparison is required, the specified mechanism is to append the NFD form of the original string after the sort key, in [Section 4.3, *Form Sort Key*](#). See also [Appendix A, *Deterministic Sorting*](#).

~~The fourth-level weights in the Default Collation Element Table can be used to provide an approximation of a deterministic comparison.~~

A deterministic comparison is also sometimes referred to as a *stable (or semi-stable) comparison*. Those terms are not to be preferred, because they tend

to be confused with *stable sort*.

Collation order is not fixed.

Over time, collation order will vary: there may be fixes needed as more information becomes available about languages; there may be new government or industry standards for the language that require changes; and finally, new characters added to the Unicode Standard will interleave with the previously-defined ones. This means that collations must be carefully versioned.

1.9 The Unicode Collation Algorithm

The Unicode Collation Algorithm (UCA) details how to compare two Unicode strings while remaining conformant to the requirements of the Unicode Standard. This standard includes the Default Unicode Collation Element Table (DUCET), which is data specifying the default collation order for all Unicode characters, and the CLDR root collation element table that is based on the DUCET. This table is designed so that it can be *tailored* to meet the requirements of different languages and customizations.

Briefly stated, the Unicode Collation Algorithm takes an input Unicode string and a Collation Element Table, containing mapping data for characters. It produces a sort key, which is an array of unsigned 16-bit integers. Two or more sort keys so produced can then be binary-compared to give the correct comparison between the strings for which they were generated.

The Unicode Collation Algorithm assumes multiple-level key weighting, along the lines widely implemented in IBM technology, and as described in the Canadian sorting standard [[CanStd](#)] and the International String Ordering standard [[ISO14651](#)].

By default, the algorithm makes use of three fully-customizable levels. For the Latin script, these levels correspond roughly to:

1. alphabetic ordering
2. diacritic ordering
3. case ordering.

A final level may be used for tie-breaking between strings not otherwise distinguished.

This design allows implementations to produce culturally acceptable collation, with a minimal burden on memory requirements and performance. In particular, it is possible to construct Collation Element Tables that use 32 bits of collation data for most characters.

Implementations of the Unicode Collation Algorithm are not limited to supporting only three levels. They are free to support a fully customizable 4th level (or more levels), as long as they can produce the same results as the basic algorithm, given the right Collation Element Tables. For example, an application which uses the algorithm, but which must treat some collection of special characters as ignorable at the first three levels *and* must have those specials collate in non-Unicode order (for example to emulate an existing EBCDIC-based collation), may choose to have a fully customizable 4th level. The downside of this choice is that such an application will require more

storage, both for the Collation Element Table and in constructed sort keys.

The Collation Element Table may be tailored to produce particular culturally required orderings for different languages or locales. As in the algorithm itself, the tailoring can provide full customization for three (or more) levels.

1.9.1 Goals

The algorithm is designed to satisfy the following goals:

1. A complete, unambiguous, specified ordering for all characters in Unicode.
2. A complete resolution of the handling of canonical and compatibility equivalences as relates to the default ordering.
3. A complete specification of the meaning and assignment of collation levels, including whether a character is ignorable by default in collation.
4. A complete specification of the rules for using the level weights to determine the default collation order of strings of arbitrary length.
5. Allowance for override mechanisms (*tailoring*) to create language-specific orderings. Tailoring can be provided by any well-defined syntax that takes the default ordering and produces another well-formed ordering.
6. An algorithm that can be efficiently implemented, in terms of both performance and memory requirements.

Given the standard ordering and the tailoring for any particular language, any two companies or individuals—with their own proprietary implementations—can take any arbitrary Unicode input and produce exactly the same ordering of two strings. In addition, when given an appropriate tailoring this algorithm can pass the Canadian and ISO 14651 benchmarks ([\[CanStd\]](#), [\[ISO14651\]](#)).

Note: The Default Unicode Collation Element Table does not explicitly list weights for all assigned Unicode characters. However, the algorithm is well defined over *all* Unicode code points. See [Section 7.1.2, *Unassigned and Other Code Points*](#).

1.9.2 Non-Goals

The Default Unicode Collation Element Table (DUCET) explicitly does not provide for the following features:

1. *Reversibility*: from a Collation Element one is not guaranteed to be able to recover the original character.
2. *Numeric formatting*: numbers composed of a string of digits or other numerics will not necessarily sort in *numerical order*.
3. *API*: no particular API is specified or required for the algorithm.
4. *Title sorting*: removing articles such as *a* and *the* during bibliographic sorting is not provided.
5. *Stability of binary sort key values between versions*: weights in the DUCET may change between versions. For more information, see [Section 3.8, *Stability*](#) "[Collation order is not a stable sort](#)" in [Section 1.8, *What Collation is Not*](#).

6. *Linguistic applicability*: to meet most user expectations, a linguistic tailoring is needed. For more information, see [Section 5, Tailoring](#).

The feature of linguistic applicability deserves further discussion. DUCET does not and cannot actually provide linguistically correct sorting for every language without further tailoring. That would be impossible, due to conflicting requirements for ordering different languages that share the same script. It is not even possible in the specialized cases where a script may be predominantly used by a single language, because of the limitations of the DUCET table design and because of the requirement to minimize implementation overhead for all users of DUCET.

Instead, the goal of DUCET is to provide a reasonable default ordering for all scripts that are *not* tailored. Any characters used in the language of primary interest for collation are expected to be tailored to meet all the appropriate linguistic requirements for that language. For example, for a user interested primarily in the Malayalam language, DUCET would be tailored to get all details correct for the expected Malayalam collation order, while leaving other characters (Greek, Cyrillic, Han, and so forth) in the default order, because the order of those other characters is not of primary concern. Conversely, a user interested primarily in the Greek language would use a Greek-specific tailoring, while leaving the Malayalam (and other) characters in their default order in the table.

2 Conformance

The Unicode Collation Algorithm does not restrict the many different ways in which implementations can compare strings. However, any Unicode-conformant implementation that purports to implement the Unicode Collation Algorithm must do so as described in this document.

A conformance test for the UCA is available in [\[Tests10\]](#).

The algorithm is a *logical* specification. Implementations are free to change any part of the algorithm as long as any two strings compared by the implementation are ordered the same as they would be by the algorithm as specified. Implementations may also use a different format for the data in the Collation Element Table. The sort key is a logical intermediate object: if an implementation produces the same results in comparison of strings, the sort keys can differ in format from what is specified in this document. (See [Section 6, Implementation Notes](#).)

The conformance requirements of the Unicode Collation Algorithm are as follows:

- C1** Given a well-formed Unicode Collation Element Table, a conformant implementation shall replicate the same comparisons of strings as those produced by [Section 4, Main Algorithm](#).

In particular, a conformant implementation must be able to compare any two canonical-equivalent strings as being equal, for all Unicode characters supported by that implementation.

If a conformant implementation compares strings in a legacy character set, it must provide the same results as if those strings had been transcoded to Unicode. It should specify the conversion table and transcoding mechanism.

3 Collation Element Table

A Collation Element Table contains a mapping from one (or more) characters to one (or more) *collation elements*, where a collation element is an ordered list of three or more weights (non-negative integers). (All code points not explicitly mentioned in the mapping are given an implicit weight: see *Section 7, [Weight Derivation](#)*).

Note: Implementations can produce the same result using various representations of weights. See *Section 6, [Implementation Notes](#)*.

Unless otherwise noted, all weights used in the example collation elements in this document are in hexadecimal format. The specific weight values shown are illustrative only; they may not match the weights in the latest Default Unicode Collation Element Table [[Allkeys](#)].

3.1 Weight Levels and Notation

The first weight is called the *Level 1* or *primary* weight; the second is called the *Level 2* or *secondary* weight; the third is called the *Level 3* or *tertiary* weight; the fourth is called the *Level 4* or *quaternary* weight, and so on. For a collation element X, these can be abbreviated as X_1 , X_2 , X_3 , X_4 , and so on.

Given two collation elements X and Y, this document uses the notation in *Table 7* and *Table 8*.

Table 7. Equals Notation

Notation	Reading	Meaning
$X =_1 Y$	X is primary equal to Y	$X_1 = Y_1$
$X =_2 Y$	X is secondary equal to Y	$X_2 = Y_2$ and $X =_1 Y$
$X =_3 Y$	X is tertiary equal to Y	$X_3 = Y_3$ and $X =_2 Y$
$X =_4 Y$	X is quaternary equal to Y	$X_4 = Y_4$ and $X =_3 Y$

Table 8. Less Than Notation

Notation	Reading	Meaning
$X <_1 Y$	X is primary less than Y	$X_1 < Y_1$
$X <_2 Y$	X is secondary less than Y	$X <_1 Y$ or ($X =_1 Y$ and $X_2 < Y_2$)

$X <_3 Y$	X is tertiary less than Y	$X <_2 Y$ or $(X =_2 Y$ and $X_3 < Y_3)$
$X <_4 Y$	X is quaternary less than Y	$X <_3 Y$ or $(X =_3 Y$ and $X_4 < Y_4)$

Other operations are given their customary definitions in terms of the above. That is:

- $X \leq_n Y$ if and only if $X <_n Y$ or $X =_n Y$
- $X >_n Y$ if and only if $Y <_n X$
- $X \geq_n Y$ if and only if $Y \leq_n X$

This notation for collation elements is also adapted to refer to ordering between strings, as shown in *Table 9*, where A and B refer to two strings.

Table 9. Notation for String Ordering

Notation	Meaning
$A <_2 B$	A is less than B, and there is a primary or secondary difference between them
$A <_2 B$ and $A =_1 B$	A is less than B, but there is only a secondary difference between them
$A \equiv B$	A and B are equivalent (equal at all levels) according to a given Collation Element Table
$A = B$	A and B are bit-for-bit identical

Where only plain text ASCII characters are available the fallback notation in *Table 10* may be used.

Table 10. Fallback Notation

Notation	Fallback
$X <_n Y$	$X <[n] Y$
X_n	$X[n]$
$X \leq_n Y$	$X <=[n] Y$
$A \equiv B$	$A =[a] B$

If a weight is 0000, then that collation element is *ignorable* at that level: the weight at that level is not taken into account in sorting. A *Level N ignorable* is a collation element that is ignorable at level N but not at level N+1. Thus:

D1. A *primary collation element* is a collation element that is not ignorable at Level 1.

- This is also known as a *non-ignorable*. In parametrized expressions, also known as a Level 0 ignorable.

D2. A *secondary collation element* is a collation element that is ignorable at Level 1, but not at Level 2.

- This is also known as a *Level 1 ignorable* or a *primary ignorable*.

D3. A *tertiary collation element* is ignorable at Levels 1 and 2, but not Level 3.

- This is also known as a *Level 2 ignorable* or a *secondary ignorable*.

D4. A *quaternary collation element* is ignorable at Levels 1, 2, and 3 but not Level 4.

- This is also known as a *Level 3 ignorable* or a *tertiary ignorable*.

D5. A *completely ignorable collation element* is ignorable at all levels (except the identical level).

D6. An *ignorable collation element* is ignorable at Level 1.

- It may be a secondary, tertiary, quaternary, or completely ignorable collation element. If the UCA is extended to more levels, then an ignorable collation element includes those ignorable at those levels.

For a given Collation Element Table, MIN_n is the least weight in any collation element at level n , and MAX_n is the maximum weight in any collation element at level n .

There are three kinds of collation element mappings used in the discussion below. These are defined as follows:

D7. A *simple mapping* maps one Unicode character to one collation element.

D8. An *expansion* maps one Unicode character to a sequence of collation elements.

D9. A *contraction* maps a sequence of Unicode characters to a sequence of (one or more) collation elements.

3.2 Simple Mappings

Most of the mappings in a collation element table are simple: they consist of the mapping of a single character to a single collation element.

The following list shows several simple mappings that are used in the examples illustrating the algorithm.

Character	Collation Element	Name
0300 "˘"	[.0000.0021.0002]	COMBINING GRAVE ACCENT
0061 "a"	[.06D9.0020.0002]	LATIN SMALL LETTER A

0062 "b"	[.06EE.0020.0002]	LATIN SMALL LETTER B
0063 "c"	[.0706.0020.0002]	LATIN SMALL LETTER C
0043 "C"	[.0706.0020.0008]	LATIN CAPITAL LETTER C
0064 "d"	[.0712.0020.0002]	LATIN SMALL LETTER D

3.3 Multiple Mappings

The mapping from characters to collation elements may not always be a simple mapping from one character to one collation element. In general, the mapping may be from one to many, from many to one, or from many to many.

3.3.1 Expansions

The Latin letter æ is treated as a primary equivalent to an <a e>sequence, such as in the following example:

Character	Collation Element	Name
00E6	[.15D5.0020.0004] [.0000.0139.0004] [.1632.0020.0004]	LATIN SMALL LETTER AE; "æ"

In this example, the collation element [.15D5.0020.0004] gives the primary weight for a, and the collation element [.1632.0020.0004] gives the primary weight for e.

3.3.2 Contractions

Similarly, where *ch* is treated as a single letter, as for instance in traditional Spanish, it is represented as a mapping from two characters to a single collation element, such as in the following example:

Character	Collation Element	Name
0063 0068	[.0707.0020.0002]	LATIN SMALL LETTER C, LATIN SMALL LETTER H; "ch"

In this example, the collation element [.0707.0020.0002] has a primary value one greater than the primary value for the letter *c* by itself, so that the sequence *ch* will collate after *c* and before *d*. This example shows the result of a tailoring of collation elements to weight sequences of letters as a single unit.

Characters in a contraction can be made to sort as separate characters by inserting, someplace within the contraction, a starter that maps to a completely ignorable collation element. There are two characters, *soft hyphen* and U+034F COMBINING GRAPHEME JOINER, that are particularly useful for this purpose. These can be used to separate contractions that would normally be weighted as units, such as Slovak *ch* or Danish *aa*. [Section 5.3, *Use of Combining Grapheme Joiner*](#).

Contractions that end with *non-starter* characters (those with Combining_Character_Class≠0) are known as *discontiguous contractions*. For example,

suppose that there is a contraction of **<a, combining ring above>**, as in Danish where this sorts as after "z". If the input text contains the sequence **<a, combining dot below, combining ring above>**, then the contraction still needs to be detected. This is required by the rearrangement of the combining marks:

$$\begin{aligned} &<a, combining dot below, \mathbf{combining ring above}> \\ &= \\ &<\mathbf{a, combining ring above}, combining dot below>. \end{aligned}$$

That is, discontinuous contractions must be detected in input text whenever the final sequence of non-starter characters could be rearranged so as to make a contiguous matching sequence that is canonically equivalent. In the formal algorithm this is handled by rule Rule [S2.1](#). For information on non-starters, see [\[UAX15\]](#).

3.3.3 Many-to-Many Mappings

In some cases a sequence of two or more characters is mapped to a sequence of two or more collation elements. For example, this technique is used in the Default Unicode Collation Element Table [\[Allkeys\]](#) to handle weighting of rearranged sequences of Thai or Lao left-side-vowel + consonant. See [Section 3.5, Rearrangement](#).

Both many-to-many mappings and many-to-one mappings are referred to as *contractions* in the discussion of the Unicode Collation Algorithm, even though many-to-many mappings often do not actually shorten anything. The key issue for implementations is that for both many-to-one mappings and many-to-many mappings, the weighting algorithm must first identify a sequence of characters in the input string and "contract" them together as a unit for weight lookup in the table. The identified unit may then be mapped to any number of collation elements. Contractions pose particular issues for implementations, because all eligible contraction targets must be identified first, before the application of simple mappings, so that processing for simple mappings does not bleed away the context needed to correctly identify the contractions.

3.3.4 Other Multiple Mappings

Certain characters may both expand and contract. See [Section 1.3, Contextual Sensitivity](#).

3.4 Backward Accents

In some languages (notably Canadian French) In some French dictionary ordering traditions, accents are sorted from the back of the string to the front of the string. This behavior is not marked in the Default Unicode Collation Element Table, but may occur in tailored tables. In such a case, the collation elements for the accents and their base characters are marked as being *backwards* at Level 2.

3.5 Rearrangement

Certain characters, such as the Thai vowels ɨ through ɯ (and related vowels in the Lao and Tai Viet scripts of Southeast Asia), are not represented in strings in logical order. The exact list of such characters is given by the Logical_Order_Exception property in

the Unicode Character Database [UAX44]. For collation, they are rearranged by swapping them with the following character before further processing, because logically they belong afterward. This is done by providing these sequences as many-to-many mappings in the Collation Element Table.

3.6.2 Variable Weighting

This used to be section 3.6.2.

Collation elements that are marked with an asterisk in a Unicode Collation Element Table are known as *variable collation elements*.

Character	Collation Element	Name
0020 " "	[*0209.0020.0002]	SPACE

Non-ignorable collation elements with low primary weights, usually up to and including punctuation (as in CLDR) or even symbols (as in the DUCET), are known as *variable collation elements*.

Based on the ~~setting of the variable weighting tag~~ variable-weighting setting, collation elements can be either treated as quaternary collation elements or not. When they are treated as quaternary collation elements, any sequence of ignorable collation elements that immediately follows the variable collation element is also affected.

There are ~~four~~ possible options for variable weighted characters:

1. **Non-ignorable:** Variable collation elements are not reset to be quaternary collation elements, ~~but get the weights explicitly mentioned in the file.~~ All mappings defined in the table are unchanged.

- ~~SPACE would have the value [*0209.0020.0002]~~
- ~~Capital A would be unchanged, with the value [06D9.0020.0008]~~
- ~~Ignorables are unchanged.~~

2. **Blanked:** Variable collation elements and any subsequent ignorable collation elements are reset so that all weights (except for the identical level) are zero. It is the same as the Shifted Option, except that there is no fourth level.
3. **Shifted:** Variable collation elements are reset to zero at levels one through three. In addition, a new fourth-level weight is appended, whose value depends on the type, as shown in *Table 12*. Any subsequent primary or secondary ignorables following a variable are reset so that their weights at levels one through four are zero.

- A combining grave accent after a space would have the value [0000.0000.0000.0000].

- A combining grave accent after a *Capital A* would be unchanged.

4. **IgnoreSP:** This option is the same as ~~Shifted~~, except that only the variable characters that are Whitespace or Punctuation are shifted; the Symbol characters

are treated as regular (non-variable) characters.

5. **Shift-Trimmed:** This option is the same as **Shifted**, except that all trailing FFFFs are trimmed from the sort key. This could be used to emulate POSIX behavior, but is otherwise not recommended.

In UCA versions 6.1 and 6.2 another option, **IgnoreSP**, was defined. That was a variant of **Shifted** that reduced the set of variable collation elements to include only spaces and punctuation, as in CLDR.

Table 12. L4 Weights for Shifted Variables

Type	L4	Examples
L1, L2, L3 = 0	0000	null [.0000.0000.0000.0000]
L1=0, L3 ≠ 0, following a Variable	0000	combining grave [.0000.0000.0000.0000]
L1 ≠ 0, Variable	old L1	space [.0000.0000.0000.0209]
L1 = 0, L3 ≠ 0, not following a Variable	FFFF	combining grave [.0000.0035.0002.FFFF]
L1 ≠ 0, not Variable	FFFF	Capital A [.06D9.0020.0008.FFFF]

The variants of the *shifted* option provide for improved orderings when the variable collation elements are ignorable, while still only requiring three fields to be stored in memory for each collation element. Those options result in somewhat longer sort keys, although they can be compressed (see [Section 6.1, *Reducing Sort Key Lengths*](#) and [Section 6.3, *Reducing Table Sizes*](#)).

Table 13 shows the differences between orderings using the different options for variable collation elements. In this example, sample strings differ by the third character: a letter, *space*, *'-' hyphen-minus (002D)*, or *'-' hyphen (2010)*; followed by an uppercase/lowercase distinction.

Table 13. Comparison of Variable Ordering

Non-ignorable	Blanked	Shifted	IgnoreSP Shifted (CLDR)	Shift-Trimmed
de luge	death	death	death	death
de Luge	de luge	de luge	de luge	deluge
de-luge	de-luge	de-luge	de-luge	de luge
de-Luge	deluge	de-luge	de-luge	de-luge

de-luge	de-luge	deluge	deluge	de-luge
de-Luge	de Luge	de Luge	de Luge	deLuge
death	de-Luge	de-Luge	de-Luge	de Luge
deluge	deLuge	de-Luge	de-Luge	de-Luge
deLuge	de-Luge	deLuge	deLuge	de-Luge
demark	demark	demark	demark	demark
☹happy	☹happy	☹happy	☹happy	☹happy
☹sad	♥happy	♥happy	☹sad	♥happy
♥happy	☹sad	☹sad	♥happy	☹sad
♥sad	♥sad	♥sad	♥sad	♥sad

The following points out some salient features of each of the columns in Table 13.

1. **Non-ignorable.** The words with *hyphen-minus* or *hyphen* are grouped together, but before all letters in the third position. This is because they are not ignorable, and have primary values that differ from the letters. The symbols ☹ and ♥ have primary differences.
2. **Blanked.** The words with *hyphen-minus* or *hyphen* are separated by "deluge", because the letter "l" comes between them in Unicode code order. The symbols ☹ and ♥ are *ignored* on levels 1-3.
3. **Shifted.** The *hyphen-minus* and *hyphen* are grouped together, and their differences are less significant than the casing differences in the letter "l". This grouping results from the fact that they are ignorable, but their fourth level differences are according to the original primary order, which is more intuitive than Unicode order. The symbols ☹ and ♥ are *ignored* on levels 1-3.
4. **IgnoreSP, Shifted (CLDR).** The same as Shifted, except that the symbols ☹ and ♥ have primary differences.
5. **Shift-Trimmed.** Note how "deLuge" comes between the cased versions with spaces and hyphens. The symbols ☹ and ♥ are *ignored* on levels 1-3.

Primaries for variable collation elements are not *interleaved* with other primary weights. This allows for more compact storage of memory tables. Rather than using a bit per collation element to determine whether the collation element is variable, the implementation only needs to store the maximum primary value for all the variable elements. All collation elements with primary weights from 1 to that maximum are variables; all other collation elements are not.

3.7 Well-Formed Collation Element Tables

A well-formed Collation Element Table meets the following well-formedness conditions:

WF1. Except in special cases detailed in *Section 6.2, [Large Weight Values](#)*, no collation element can have a zero weight at Level N and a non-zero weight at Level N-1.

- For example, the secondary weight can only be ignorable if the primary weight is

ignorable.

- For a detailed example of what happens if the condition is not met, see *Section 4.5 [Rationale for Well-Formed Collation Element Tables](#)*.

WF2. Secondary weights of secondary collation elements must be strictly greater than secondary weights of all primary collation elements. Tertiary weights of tertiary collation elements must be strictly greater than tertiary weights of all primary and secondary collation elements.

- Given collation elements [A, B, C], [0, D, E], [0, 0, F], where the letters are non-zero weights, the following must be true:
 - $D > B$
 - $F > C$
 - $F > E$
- For a detailed example of what happens if the condition is not met, see *Section 4.5 [Rationale for Well-Formed Collation Element Tables](#)*.

WF3. No variable collation element has an ignorable primary weight.

WF4. For all variable collation elements U, V, if there is a collation element W such that $U_1 \leq W_1$ and $W_1 \leq V_1$, then W is also variable.

- This provision prevents [interleaving](#).

WF5. If a table contains a contraction consisting of a sequence of N code points, with $N > 2$ and the last code point being a non-starter, then the table must also contain a contraction consisting of the sequence of the first N-1 code points.

- For example, if "ae<umlaut>" is a contraction, then "ae" must be a contraction as well.

3.8 Default Unicode Collation Element Table

This used to be section 3.6.

The Default Unicode Collation Element Table is provided in [\[Allkeys\]](#). This table provides a mapping from characters to collation elements for all the explicitly weighted characters. The mapping lists characters in the order that they are weighted. Any code points that are not explicitly mentioned in this table are given a derived collation element, as described in *Section 7, [Weight Derivation](#)*.

The Default Unicode Collation Element Table does not aim to provide precisely correct ordering for each language and script; tailoring is required for correct language handling in almost all cases. The goal is instead to have all the *other* characters, those that are not tailored, show up in a reasonable order. This is particularly true for contractions, because contractions can result in larger tables and significant performance degradation. Contractions are required in tailorings, but their use is kept to a minimum in the Default Unicode Collation Element Table to enhance performance.

In the Default Unicode Collation Element Table, contractions are necessary where a

canonical decomposable character requires a distinct primary weight in the table, so that the canonical-equivalent character sequences are given the same weights. For example, Indic two-part vowels have primary weights as units, and their canonical-equivalent sequence of vowel parts must be given the same primary weight by means of a contraction entry in the table. The same applies to a number of precomposed Cyrillic characters with diacritic marks and to a small number of Arabic letters with *madda* or *hamza* marks.

Contractions are also entered in the table for Thai, Lao, and Tai Viet logical order exception vowels. Because these scripts all have five vowels that are represented in strings in visual order, the vowels cannot simply be weighted by their representation order in strings. One option is to preprocess relevant strings to identify and reorder all logical order exception vowels around the following consonant. That approach was used in Version 4.0 and earlier of the UCA. Starting with Version 4.1 of the UCA, contractions for the relevant combinations of vowel+consonant have been entered in the Default Unicode Collation Element Table instead.

Those are the only classes of contractions allowed in the Default Unicode Collation Element Table. Generic contractions of the sort needed to handle digraphs such as "ch" in Spanish or Czech sorting, should be dealt with in tailorings to the default table—because they often vary in ordering from language to language, and because every contraction entered into the default table has a significant implementation cost for all applications of the default table, even those which may not be particularly concerned with the affected script. See the Unicode Common Locale Data Repository [[CLDR](#)] for extensive tailorings of the DUCET for various languages, including those requiring contractions.

The Default Unicode Collation Element Table is constructed to be consistent with the Unicode Normalization algorithm, and to respect the Unicode character properties. It is not, however, merely algorithmically derivable based on considerations of canonical equivalence and an inspection of character properties, because the assignment of levels also takes into account characteristics of particular scripts. For example, the combining marks generally have *secondary collation elements*; however, the Indic combining vowels are given non-zero Level 1 weights, because they are as significant in sorting as the consonants.

Any character may have variant forms or applied accents which affect collation. Thus, for `FULL STOP` there are three compatibility variants: a fullwidth form, a compatibility form, and a small form. These get different tertiary weights accordingly. For more information on how the table was constructed, see [Section 7.2, *Tertiary Weight Table*](#).

Table 11 summarizes the overall ordering of the collation elements in the Default Unicode Collation Element Table. The collation elements are ordered by primary, secondary, tertiary, and Unicode value weights, with primary, secondary, and tertiary weights for variables blanked (replaced by "0000"). Entries in the table which contain a sequence of collation elements have a multi-level ordering applied: comparing the primary weights first, then the secondary weights, and so on. This construction of the table makes it easy to see the order in which characters would be collated.

The weightings in the table are grouped by major categories. For example, whitespace

characters come before punctuation, and symbols come before numbers. These groupings allow for programmatic reordering of scripts and other characters of interest, without table modification. For example, numbers can be reordered to be after letters instead of before. For more information, see the *Unicode Common Locale Data Repository* [CLDR].

Table 11. DUCET Ordering

Values	Type	Examples of Characters
$X_1, X_2, X_3 = 0$	completely ignorable and quaternary collation elements	<ul style="list-style-type: none"> - Control codes - Format characters - Hebrew points - Tibetan signs - Arabic tatweel ...
$X_1, X_2 = 0;$ $X_3 \neq 0$	tertiary collation elements	None in DUCET; could be in tailorings
$X_1 = 0;$ $X_2, X_3 \neq 0$	secondary collation elements	<ul style="list-style-type: none"> - Most nonspacing marks - Some letters and combining marks
$X_1, X_2, X_3 \neq 0$	primary collation elements	
	variable	<ul style="list-style-type: none"> - Whitespace (White_Space=True) - Punctuation (General_Category=Punctuation) - General symbols (General_Category=Letter Modifier or Symbol, but not Currency Symbol)
	regular	<ul style="list-style-type: none"> - General symbols (General_Category=Letter Modifier; certain characters, such as U+02D0 : MODIFIER LETTER TRIANGULAR COLON) - Currency symbols (General_Category=Currency Symbol) - Numbers (General_Category=Number) - Latin - Greek ...
	implicit	<ul style="list-style-type: none"> - CJK Unified Ideographs from the URO and CJK Compatibility blocks - CJK Extensions A, B, C, ...

		- Unassigned and others given implicit weights
	trailing	None in DUCET; could be in tailorings
	reserved	Special collation elements - U+FFFD

Note: The position of the boundary between variable and regular collation elements can be tailored.

There are a number of exceptions in the grouping of characters in DUCET, where for various reasons characters are grouped in different categories. Examples are provided below for each type of exception.

1. If the NFKD decomposition of a character starts with certain punctuation characters, it is grouped with punctuation.
 - U+2474 ⁽¹⁾ PARENTHESES DIGIT ONE
2. If the NFKD decomposition of a character starts with a character having General_Category=Number, then it is grouped with numbers.
 - U+3358 0点 IDEOGRAPHIC TELEGRAPH SYMBOL FOR HOUR ZERO
3. Many non-decimal numbers are grouped with general symbols.
 - U+2180 Ⅰ ROMAN NUMERAL ONE THOUSAND C D
4. Some numbers are grouped with the letters for particular scripts.
 - U+3280 ⊖ CIRCLED IDEOGRAPH ONE
5. Some letter modifiers are grouped with general symbols, others with their script.
 - U+3005 ㄥ IDEOGRAPHIC ITERATION MARK
6. There are a few other exceptions, such as currency signs grouped with letters because of their decompositions.
 - U+20A8 ₹ RUPEE SIGN

Note that the [\[CLDR\]](#) root collation tailors the DUCET. The differences are described in the documentation file [CollationAuxiliary.html](#) (see [Section 9, Data Files](#)). For details see [UTS #35: LDML: Collation, Section 2, Root Collation](#).

For most languages, some degree of tailoring is required to match user expectations. For more information, see [Section 5, Tailoring](#).

[Section 3.6.1 File Format](#) became [section 9.1 Allkeys File Format](#).

[Section 3.6.2 Variable Weighting](#) became [section 3.6](#).

[3.8.1](#) **Default Values**

[This used to be section 3.6.3](#).

In the Default Unicode Collation Element Table and in typical tailorings, most

unaccented letters differ in the primary weights, but have secondary weights (such as a_1) equal to MIN_2 . The secondary collation elements will have secondary weights greater than MIN_2 . Characters that are compatibility or case variants will have equal primary and secondary weights (for example, $a_1 = A_1$ and $a_2 = A_2$), but have different tertiary weights (for example, $a_3 < A_3$). The unmarked characters will a_3) equal to MIN_3 .

This use of secondary and tertiary weights does not guarantee that the meaning of a secondary or tertiary weight is uniform across tables. For example, in a tailoring a *capital A* and *katakana ta* could both have a tertiary weight of 3.

3.8.2 Well-Formedness of the DUCET

This used to be section 3.6.4.

The DUCET is *not entirely well-formed*. It does not include two contraction mappings required for [well-formedness condition 5](#):

```
0FB2 0F71 ; CE(0FB2) CE(0F71)
0FB3 0F71 ; CE(0FB3) CE(0F71)
```

However, adding just these two contractions would disturb the default sort order for Tibetan. In order to also preserve the sort order for Tibetan, the following eight contractions would have to be added as well:

```
0FB2 0F71 0F72 ; CE(0FB2) CE(0F71 0F72)
0FB2 0F73      ; CE(0FB2) CE(0F71 0F72)
0FB2 0F71 0F74 ; CE(0FB2) CE(0F71 0F74)
0FB2 0F75      ; CE(0FB2) CE(0F71 0F74)

0FB3 0F71 0F72 ; CE(0FB3) CE(0F71 0F72)
0FB3 0F73      ; CE(0FB3) CE(0F71 0F72)
0FB3 0F71 0F74 ; CE(0FB3) CE(0F71 0F74)
0FB3 0F75      ; CE(0FB3) CE(0F71 0F74)
```

The [CLDR](#) root collation adds all ten of these contractions.

3.8 Stability

The notion of *stability* in sorting often causes confusion when discussing collation.

Text of sections 3.8.1 & 3.8.2 was moved into section 1.8.

3.8.3 Stability of the DUCET

Neither a *stable sort* nor a *deterministic comparison* refers to the stability of the Default Collation Element Table itself. The contents of that table, the DUCET, will remain unchanged in any particular version of the UCA. However, the contents may change between successive versions of the UCA as new characters are added, or more information is obtained about existing characters.

Implementers should be aware that using different versions of the UCA or different

versions of the Unicode Standard could result in different collation results of their data. There are numerous ways collation data could vary across versions, for example:

1. Code points that were unassigned in a previous version of the Unicode Standard are now assigned in the current version, and will have a sorting semantic appropriate to the repertoire to which they belong. For example, the code points U+103D0..U+103DF were undefined in Unicode 3.1. Because they were assigned characters in Unicode 3.2, their sorting semantics and respective sorting weights changed as of that version.
2. Certain semantics of the Unicode standard could change between versions, such that code points are treated in a manner different than previous versions of the standard.
3. More information is gathered about a particular script, and the weight of a code point may need to be adjusted to provide a more linguistically accurate sort.

Any of these reasons could necessitate a change between versions with regards to collation weights for code points. It is therefore important that the implementers specify the version of the UCA, as well as the version of the Unicode Standard under which their data is sorted.

The policies which the UTC uses to guide decisions about the collation weight assignments made for newly assigned characters are enumerated in the [UCA Default Table Criteria for New Characters](#). In addition, there are policies which constrain the timing and type of changes which are allowed for the DUCET table between versions of the UCA. Those policies are enumerated in [Change Management for the Unicode Collation Algorithm](#).

4 Main Algorithm

The main algorithm has four steps. First is to normalize each input string, second is to produce an array of collation elements for each string, and third is to produce a sort key for each string from the collation elements. Two sort keys can then be compared with a binary comparison; the result is the ordering for the original strings.

4.1 Normalize

Step 1. Produce a normalized form of each input string, applying S1.1.

S1.1 Use the Unicode canonical algorithm to decompose characters according to the canonical mappings. That is, put the string into Normalization Form D (see [\[UAX15\]](#)).

- Conformant implementations may skip this step *in certain circumstances*: see [Section 6.5, *Avoiding Normalization*](#) for more information.

4.2 Produce Array

Step 2. The collation element array is built by sequencing through the normalized form, applying S2.1 through S2.6.

Figure 1. String to Collation Element Array

Normalized String	Collation Element Array
ca´b	[.0706.0020.0002], [.06D9.0020.0002], [.0000.0021.0002], [.06EE.0020.0002]

S2.1 Find the longest initial substring S at each point that has a match in the table.

S2.1.1 If there are any non-starters following S, process each non-starter C.

S2.1.2 If C is not blocked from S, find if S + C has a match in the table.

Note: A non-starter in a string is called *blocked* if there is another non-starter of the same canonical combining class or zero between it and the last character of canonical combining class 0.

S2.1.3 If there is a match, replace S by S + C, and remove C.

S2.2 Fetch the corresponding collation element(s) from the table if there is a match. If there is no match, synthesize a weight as described in [Section 7.1, *Derived Collation Elements*](#).

S2.3 Process collation elements according to the variable-weight setting, as described in [Section 3.6-2, *Variable Weighting*](#).

S2.4 Append the collation element(s) to the collation element array.

S2.5 Proceed to the next point in the string (past S).

S2.6 Loop until the end of the string is reached.

Note: The extra non-starter C needs to be considered in Step 2.1.1 because otherwise irrelevant characters could interfere with matches in the table. For example, suppose that the contraction $\langle a, combining_ring \rangle$ (= \acute{a}) is ordered after z. If a string consists of the three characters $\langle a, combining_ring, combining_cedilla \rangle$, then the normalized form is $\langle a, combining_cedilla, combining_ring \rangle$, which separates the *a* from the *combining_ring*. Without considering the extra non-starter, this string would compare incorrectly as after *a* and not after *z*.

If the desired ordering treats $\langle a, combining_cedilla \rangle$ as a contraction which should take precedence over $\langle a, combining_ring \rangle$, then an additional mapping for the combination $\langle a, combining_ring, combining_cedilla \rangle$ can be introduced to produce this effect.

For conformance to Unicode canonical equivalence, only unblocked non-starters are matched in Step 2.1.2. For example, $\langle a, combining_macron, combining_ring \rangle$ would compare as after *a-macron*, and not after *z*. Additional mappings can be added to customize behavior.

Also note that the Algorithm employs two distinct contraction matching methods:

- Step 2.1 “Find the longest initial substring S” is a contiguous, longest-match method. In particular, it must support matching of a contraction ABC even if there is not also a contraction AB. Thus, an implementation that incrementally matches a lengthening initial substring must be able to handle partial matches like for AB.
- Steps 2.1.1 “process each non-starter C” and 2.1.2 “find if S + C has a match in the table”, where one or more intermediate non-starters may be skipped (making it discontinuous), extends a contraction match by one code point at a time to find the next match. In particular, if C is a non-starter and if the table had a mapping for ABC but not one for AB, then a discontinuous-contraction match on text ABMC (with M being a skippable non-starter) would never be found. [Well-formedness condition 5](#) requires the presence of the prefix contraction AB.
- In either case, the prefix contraction AB cannot be added to the table automatically because it would yield the wrong order for text ABD if there is a contraction BD.

4.3 Form Sort Key

Step 3. The sort key is formed by successively appending all non-zero weights from the collation element array. The weights are appended from each level in turn, from 1 to 3. (Backwards weights are inserted in reverse order.)

Figure 2. Collation Element Array to Sort Key

Collation Element Array	Sort Key
[.0706.0020.0002], [.06D9.0020.0002], [.0000.0021.0002], [.06EE.0020.0002]	0706 06D9 06EE 0000 0020 0020 0021 0020 0000 0002 0002 0002 0002

An implementation may allow the maximum level to be set to a smaller level than the available levels in the collation element array. For example, if the maximum level is set to 2, then level 3 and higher weights are not appended to the sort key. Thus any differences at levels 3 and higher will be ignored, leveling any such differences in string comparison.

Here is a more detailed statement of the algorithm:

S3.1 For each weight level L in the collation element array from 1 to the maximum level,

S3.2 If L is not 1, append a *level separator*

Note: The level separator is zero (0000), which is guaranteed to be lower than any weight in the resulting sort key. This guarantees that when two strings of unequal length are compared, where the shorter string is a prefix of the longer string, the longer string is always sorted after the shorter—in the absence of special features like contractions. For example: "abc" < "abcX" where "X" can be any character(s).

S3.3 If the collation element table is forwards at level L,

S3.4 For each collation element CE in the array

S3.5 Append CE_L to the sort key if CE_L is non-zero.

S3.6 Else the collation table is backwards at level L, so

S3.7 Form a list of all the non-zero CE_L values.

S3.8 Reverse that list

S3.9 Append the CE_L values from that list to the sort key.

S3.10 If a semi-stable sort is required, then after all the level weights have been added, append a copy of the NFD version of the original string. This strength level is called the *identical level*, and this feature is called *semi-stability*. (See also *Appendix A, Deterministic Sorting*.)

4.4 Compare

Step 4. Compare the sort keys for each of the input strings, using a binary comparison. This means that:

- Level 3 differences are ignored if there are any Level 1 or 2 differences.
- Level 2 differences are ignored if there are any Level 1 differences.
- Level 1 differences are never ignored.

Figure 3. Comparison of Sort Keys

	String	Sort Key
1	cab	<u>0706</u> 06D9 06EE 0000 0020 0020 <u>0020</u> 0000 <u>0002</u> 0002 0002
2	Cab	<u>0706</u> 06D9 06EE 0000 0020 0020 <u>0020</u> 0000 <u>0008</u> 0002 0002
3	cáb	<u>0706</u> 06D9 06EE 0000 0020 0020 <u>0021</u> 0020 0000 0002 0002 0002 0002
4	dab	<u>0712</u> 06D9 06EE 0000 0020 0020 0020 0000 0002 0002 0002

In *Figure 3*, "cab" $<_3$ "Cab" $<_2$ "cáb" $<_1$ "dab". The differences that produce the ordering are shown by the **bold underlined** items:

- For strings 1 and 2, the first difference is in 0002 versus 0008 (Level 3).
- For strings 2 and 3, the first difference is in 0020 versus 0021 (Level 2).
- For strings 3 and 4, the first difference is in 0706 versus 0712 (Level 1).

4.5 Rationale for Well-Formed Collation Element Tables

While forming sort keys, zero weights are omitted. If collation elements were not well-formed according to conditions 1 and 2, the ordering of collation elements could be incorrectly reflected in the sort key. The following examples illustrate this.

Suppose [well-formedness condition 1](#) were broken, and secondary weights of the Latin characters were zero (ignorable) and that (as normal) the primary weights of case-variants are equal: that is, $a_1 = A_1$. Then the following incorrect keys would be generated:

Order	String	Normalized	Sort Key
1	"áe"	a, acute, e	$a_1 e_1 0000 \text{ acute}_2 0000 \underline{a}_3 \text{ acute}_3 e_3 \dots$
2	"Aé"	A, e, acute	$a_1 e_1 0000 \text{ acute}_2 0000 \underline{A}_3 \text{ acute}_3 e_3 \dots$

Because the secondary weights for *a*, *A*, and *e* are lost in forming the sort key, the relative order of the acute is also lost, resulting in an incorrect ordering based solely on the case of *A* versus *a*. With well-formed weights, this does not happen, and the following correct ordering is obtained:

Order	String	Normalized	Sort Key
1	"Aé"	A, e, acute	$a_1 e_1 0000 a_2 \underline{e}_2 \text{ acute}_2 0000 a_3 \text{ acute}_3 e_3 \dots$
2	"áe"	a, acute, e	$a_1 e_1 0000 a_2 \underline{\text{acute}}_2 e_2 0000 A_3 \text{ acute}_3 e_3 \dots$

However, there are circumstances—typically in expansions—where higher-level weights in collation elements can be zeroed (resulting in ill-formed collation elements) without consequence (see [Section 6.2, Large Weight Values](#)). Implementations are free to do this as long as they produce the same result as with well-formed tables.

Suppose on the other hand, [well-formedness condition 2](#) were broken. Let there be a tailoring of 'b' as a secondary difference from 'a' resulting in the following collation elements where the one for 'b' is ill-formed.

```
0300 ; [.0000.0035.0002.0300] # (DUCET) COMBINING GRAVE ACCENT
0061 ; [.15EF.0020.0002.0061] # (DUCET) LATIN SMALL LETTER A
0062 ; [.15EF.0040.0002.0061] # (tailored) LATIN SMALL LETTER B
```

Then the following incorrect ordering would result: "aa" < "àa" < "ab" — The secondary difference on the *second* character (b) trumps the accent on the *first* character (à).

A correct tailoring would give 'b' a secondary weight lower than that of any secondary collation element, for example: (assuming the DUCET did not use secondary weight 0021 for any secondary collation element)

```
0300 ; [.0000.0035.0002.0300] # (DUCET) COMBINING GRAVE ACCENT
0061 ; [.15EF.0020.0002.0061] # (DUCET) LATIN SMALL LETTER A
0062 ; [.15EF.0021.0002.0061] # (tailored) LATIN SMALL LETTER B
```

Then the following correct ordering would result: "aa" < "ab" < "àa"

5 Tailoring

Tailoring is any well-defined syntax that takes the Default Unicode Collation Element Table and produces another well-formed Unicode Collation Element Table. This syntax

can provide linguistically-accurate collation, if desired. Such syntax will usually allow the following:

Tailoring consists of any well-defined change in the Collation Element Table and/or any well-defined change in the behavior of the algorithm. Typically, a tailoring is expressed by means of a formal syntax which allows detailed manipulation of values in a Collation Element Table, with or without an additional collection of parametric settings which modify specific aspects of the behavior of the algorithm. A tailoring can be used to provide linguistically-accurate collation, if desired. Tailorings usually specify one or more of the following kinds of changes:

1. Reordering any character (or contraction) with respect to others in the **standard default** ordering. The reordering can represent a Level 1 difference, Level 2 difference, Level 3 difference, or identity (in levels 1 to 3). Because such reordering includes sequences, arbitrary multiple mappings can be specified.
2. Removing contractions, such as the Cyrillic contractions which are not necessary for the Russian language, and the Thai/Lao reordering contractions which are not necessary for string *search*.
3. Setting the secondary level to be backwards (**for Canadian French for some French dictionary ordering traditions**) or forwards (normal).
4. Set variable weighting options.
5. Customizing the exact list of variable collation elements.
6. Allow normalization to be turned off where input is already normalized.

For best interoperability, it is recommended that tailorings for particular locales (or languages) make use of the tables provided in the Unicode Common Locale Data Repository [[CLDR](#)].

For an example of a tailoring syntax, see *Section 5.2, [Tailoring Example](#)*.

5.1 Parametric Tailoring

Parametric tailoring, if supported, is specified using a set of attribute-value pairs that specify a particular kind of behavior relative to the UCA. The standard parameter names (attributes) and their possible values are listed in *Unicode Locale Data Markup Language (LDML)* [[UTS35](#)] in the table [Collation Settings](#) (*UTS #35: LDML: Collation, Section 3.3, [Setting Options](#)*).

The default values for collation parameters specified by the UCA algorithm may differ from the LDML defaults given in the LDML table *Collation Settings*. The table indicates both default values. For example, the UCA default for alternate handling is **shifted**, while the general default in LDML is **non-ignorable**. Also, defaults in CLDR data may vary by locale. For example, **normalization** is turned off in most CLDR locales (those that don't normally use multiple accents). The default for strength in UCA is **tertiary**; it can be changed for different locales in CLDR.

When a locale or language identifier is specified for tailoring of the UCA, the identifier uses the syntax from [[UTS35](#)], *Section 3, [Unicode Language and Locale Identifiers](#)*. Unless otherwise specified, tailoring by locale uses the tables from the Unicode

Common Locale Data Repository [\[CLDR\]](#).

5.2 Tailoring Example

Unicode [\[CLDR\]](#) provides a powerful tailoring syntax in [\[UTS35\]](#), as well as tailoring data for many locales. The tailorings are based on the DUCET table, although there are some general tailorings always applied to the DUCET in CLDR 1.9. This XML format maps to a simpler representation in ICU, shown in *Table 14*. A simpler version of this syntax is also used in Java, although at the time of this writing, Java does not implement the UCA.

Table 14. ICU Tailoring Syntax

Syntax	Description
& y < x	Make x primary-greater than y
& y << x	Make x secondary-greater than y
& y <<< x	Make x tertiary-greater than y
& y = x	Make x equal to y

Either x or y in this syntax can represent more than one character, to handle contractions and expansions.

Entries for tailoring can be abbreviated in a number of ways:

- They do not need to be separated by newlines.
- Characters can be specified directly, instead of using their hexadecimal Unicode values.
- In rules of the form "x < y & y < z", "& y" can be omitted, leaving just "x < y < z".

These abbreviations can be applied successively, so the examples shown in *Table 15* are equivalent in ordering.

Table 15. Equivalent Tailorings

ICU Syntax	DUCET Syntax
a <<< A << à <<< À < b <<< B	0061 ; [0.0001.0001.0001] % a 0040 ; [0.0001.0001.0002] % A 00E0 ; [0.0001.0002.0001] % à 00C0 ; [0.0001.0002.0002] % À 0042 ; [0.0002.0001.0001] % b 0062 ; [0.0002.0001.0002] % B

The syntax has many other capabilities: for more information, see [\[UTS35\]](#) and [\[ICUCollator\]](#).

5.3 Use of Combining Grapheme Joiner

The Unicode Collation Algorithm involves the normalization of Unicode text strings

before collation weighting. U+034F COMBINING GRAPHEME JOINER (CGJ) is ordinarily ignored in collation key weighting in the UCA, but it can be used to block the reordering of combining marks in a string as described in [Unicode]. In that case, its effect can be to invert the order of secondary key weights associated with those combining marks. Because of this, the two strings would have distinct keys, making it possible to treat them distinctly in searching and sorting without having to further tailor either the combining grapheme joiner or the combining marks.

The CGJ can also be used to prevent the formation of contractions in the Unicode Collation Algorithm. Thus, for example, while *ch* is sorted as a single unit in a tailored Slovak collation, the sequence $\langle c, \text{CGJ}, h \rangle$ will sort as a *c* followed by an *h*. This can also be used in German, for example, to force *ü* to be sorted as *u + umlaut* (thus $u <_2 \ddot{u}$), even where a dictionary sort is being used (which would sort $ue <_3 \ddot{u}$). This happens without having to further tailor either the combining grapheme joiner or the sequence.

Note: As in a few other cases in the Unicode Standard, the name of the CGJ can be misleading—the usage above is in some sense the inverse of "joining".

Sequences of characters which include the combining grapheme joiner or other completely ignorable characters may also be given tailored weights. Thus the sequence $\langle c, \text{CGJ}, h \rangle$ could be weighted completely differently from either the contraction "ch" or the sequence "c" followed by "h" without the contraction. However, this application of CGJ is not recommended, because it would produce effects much different than the normal usage above, which is to simply interrupt contractions.

5.4 Preprocessing

In addition to tailoring, some implementations may choose to preprocess the text for special purposes. Once such preprocessing is done, the standard algorithm can be applied.

Examples include:

- mapping "McBeth" to "MacBeth"
- mapping "St." to "Street" or "Saint", depending on the context
- dropping articles, such as "a" or "the"
- using extra information, such as pronunciation data for Han characters

Such preprocessing is outside of the scope of this document.

6 Implementation Notes

As noted above for efficiency, implementations may vary from this logical algorithm as long as they produce the same result. The following items discuss various techniques that can be used for reducing sort key length, reducing table sizes, customizing for additional environments, searching, and other topics.

6.1 Reducing Sort Key Lengths

The following discuss methods of reducing sort key lengths. If these methods are applied to all of the sort keys produced by an implementation, they can result in significantly shorter and more efficient sort keys while retaining the same ordering.

6.1.1 Eliminating Level Separators

Level separators are not needed between two levels in the sort key, if the weights are properly chosen. For example, if all L3 weights are less than all L2 weights, then no level separator is needed between them. If there is a fourth level, then the separator before it needs to be retained.

The following example shows a sort key with these level separators removed.

String	Technique(s) Applied	Sort Key
càb	none	0706 06D9 06EE 0000 0020 0020 0021 0020 0000 0002 0002 0002 0002
càb	1	0706 06D9 06EE 0020 0020 0021 0020 0002 0002 0002 0002

While this technique is relatively easy to implement, it can interfere with other compression methods.

6.1.2 L2/L3 in 8 Bits

The L2 and L3 weights commonly are small values. Where that condition occurs for all possible values, they can then be represented as single 8-bit quantities.

The following example modifies the first example with both these changes (and grouping by bytes). Note that the separator has to remain after the primary weight when combining these techniques. If any separators are retained (such as before the fourth level), they need to have the same width as the previous level.

String	Technique(s) Applied	Sort Key
càb	none	07 06 06 D9 06 EE 00 00 00 20 00 20 00 21 00 20 00 00 00 02 00 02 00 02 00 02
càb	1, 2	07 06 06 D9 06 EE 00 00 20 20 21 20 02 02 02 02

6.1.3 Machine Words

The sort key can be represented as an array of different quantities depending on the machine architecture. For example, comparisons as arrays of unsigned 32-bit quantities may be much faster on some machines. When using arrays of unsigned 32-bit quantities, the original is to be padded with trailing (not leading) zeros as necessary.

String	Technique(s) Applied	Sort Key

càb	1, 2	07 06 06 D9 06 EE 00 00 20 20 21 20 02 02 02 02
càb	1, 2, 3	070606D9 06EE0000 20202120 02020202

6.1.4 Run-Length Compression

Generally sort keys do not differ much in the secondary or tertiary weights, which tends to result in keys with a lot of repetition. This also occurs with quaternary weights generated with the shifted parameter. By the structure of the collation element tables, there are also many weights that are never assigned at a given level in the sort key. One can take advantage of these regularities in these sequences to compact the length—while retaining the same sort sequence—by using the following technique. (There are other techniques that can also be used.)

This is a logical statement of the process; the actual implementation can be much faster and performed as the sort key is being generated.

- For each level *n*, find the most common value COMMON produced at that level by the collation element table for typical strings. For example, for the Default Unicode Collation Element Table, this is:
 - 0020 for the secondaries (corresponding to unaccented characters)
 - 0002 for tertiaries (corresponding to lowercase or unmarked letters)
 - FFFF for quaternaries (corresponding to non-ignorables with the shifted parameter)
- Reassign the weights in the collation element table at level *n* to create a gap of size GAP above COMMON. Typically for secondaries or tertiaries this is done after the values have been reduced to a byte range by the above methods. Here is a mapping that moves weights up or down to create a gap in a byte range.


```
w → w + 01 - MIN, for MIN ≤ w < COMMON
w → w + FF - MAX, for COMMON < w ≤ MAX
```
- At this point, weights go from 1 to MINTOP, and from MAXBOTTOM to MAX. These new unassigned values are used to run-length encode sequences of COMMON weights.
- When generating a sort key, look for maximal sequences of *m* COMMON values in a row. Let *W* be the weight right after the sequence.
 - If *W* < COMMON (or there is no *W*), replace the sequence by a synthetic low weight equal to (MINTOP + *m*).
 - If *W* > COMMON, replace the sequence by a synthetic high weight equal to (MAXBOTTOM - *m*).

In the example shown in *Figure 4*, the low weights are 01, 02; the high weights are FE, FF; and the common weight is 77.

Figure 4. Run-Length Compression

Original Weights	Compressed Weights
01	01
02	02

77 01	03 01
77 02	03 02
77 77 01	04 01
77 77 02	04 02
77 77 77 01	05 01
77 77 77 02	05 02
...	...
77 77 77 FE	FB FE
77 77 77 FF	FB FF
77 77 FE	FC FE
77 77 FF	FC FF
77 FE	FD FE
77 FF	FD FF
FE	FE
FF	FF

- The last step is a bit too simple, because the synthetic weights must not collide with other values having long strings of COMMON weights. This is done by using a sequence of synthetic weights, absorbing as much length into each one as possible. A value BOUND is defined between MINTOP and MAXBOTTOM. The exact value for BOUND can be chosen based on the expected frequency of synthetic low weights versus high weights for the particular collation element table.
 - If a synthetic low weight would not be less than BOUND, use a sequence of low weights of the form (BOUND-1)..(BOUND-1)(MINTOP + remainder) to express the length of the sequence.
 - Similarly, if a synthetic high weight would be less than BOUND, use a sequence of high weights of the form (BOUND)..(BOUND)(MAXBOTTOM - remainder).

This process results in keys that are never longer than the original, are generally much shorter, and result in the same comparisons.

6.2 Large Weight Values

If a collation sequence requires more than 65,535 weight values (or 65,024 values where zero bytes are avoided), this requirement can still be accommodated by assigning multiple collation elements to single characters. For example, suppose that 50,000 supplementary private-use characters are used in a particular implementation, and that these are to be sorted after a character whose primary weight is x . In such cases, the second CE ("continuation") does not have to be well formed.

Simply assign them all dual collation elements of the following form:

$[x](X+1).zzzz.www], [x]yyyy.0000.0000]$

If there is an element with the primary weight $(X+1)$, then it also needs to be converted into a dual collation element.

The private-use characters will then sort properly with respect to each other and the rest of the characters. The first collation element of this dual collation element pair is one of the instances in which ill-formed collation elements are allowed. The second collation element of each of these pairs is well-formed, and the first element only occurs in

combination with them. In this way, ordering is preserved with respect to other, non-paired collation elements.

Another example of the continuation technique appears in the DUCET:

```
2F00 ; [FB40.0020.0004] [CE00.0000.0000] # KANGXI RADICAL ONE
```

6.3 Reducing Table Sizes

The data tables required for collation of the entire Unicode repertoire can be quite sizable. This section discusses ways to significantly reduce the table size in memory. These recommendations have very important implications for implementations.

6.3.1 Contiguous Weight Ranges

Whenever collation elements have different primary weights, the ordering of their secondary weights is immaterial. Thus all of the secondaries that share a single primary can be renumbered to a contiguous range without affecting the resulting order. The same technique can be applied to tertiary weights.

6.3.2 Escape Hatch

Although the tertiary weights for the Default Unicode Collation Element Table can fit within one byte, any particular tailored table could conceivably end up with tertiary weights that exceed what can be contained in a single byte. Secondary weights in the Default Unicode Collation Element Table already far exceed what can be represented by a single byte, ranging from 0020 to 019A as of UCA Version 6.0.0. However, the same technique used for large weight values can also be used for implementations that do not want to handle more than 00FF values for a particular weight.

For example, the Java collation implementation only stored 8-bit quantities in level 2 and level 3. However, characters can be given L2 or L3 weights with greater values by using a series of two collation elements. For example, if characters require 2,000 weights at L2, then 248 characters can be given single keys, while 1,752 (2000 - 248) can be given two collation keys of the form [yyyy.00zz.00ww] [0000.00nn.0001].

The 248 in this example can be chosen to be the higher frequency characters, while there would need to be eight distinct zz values to cover the remaining characters. These zz values must only be used with dual collation elements.

6.3.3 Leveraging Unicode Tables

Because all canonically decomposable characters are decomposed in Step 1.1, no collation elements need to be supplied for them. The DUCET has over 2,000 of these, but they can all be dropped with no change to the ordering (it does omit the 11,172 Hangul syllables).

The collation elements for the Han characters (unless tailored) are algorithmically derived; no collation elements need to be stored for them either.

This means that only a small fraction of the total number of Unicode characters need to

have an explicit collation element. This can cut down the memory storage considerably.

In addition, most characters with compatibility decompositions can have collation elements computed at runtime to save space, duplicating the work that was done to compute the Default Unicode Collation Element Table. This can provide important savings in memory space. The process works as follows.

1. Derive the compatibility decomposition. For example,

```
2475 PARENTHESIZED DIGIT TWO => 0028, 0032, 0029
```

2. Look up the collation, discarding completely ignorables. For example,

```
0028 [*023D.0020.0002] % LEFT PARENTHESIS
0032 [06C8.0020.0002] % DIGIT TWO
0029 [*023E.0020.0002] % RIGHT PARENTHESIS
```

3. Set the L3 values according to the table in [Section 7.2, Tertiary Weight Table](#). If the last L3 value is 0004, and if the last element of the decomposition represents a non-alphabetic character, then reset that last L3 value to MAX (which in the default table is 001F). For example,

```
0028 [*023D.0020.0004] % LEFT PARENTHESIS
0032 [06C8.0020.0004] % DIGIT TWO
0029 [*023E.0020.0004] % RIGHT PARENTHESIS
```

4. Concatenate the result to produce the sequence of collation elements that the character maps to. For example,

```
2475 [*023D.0020.0004] [06C8.0020.0004] [*023E.0020.0004]
```

Some characters cannot be computed in this way. They must be filtered out of the default table and given specific values. For example, the *long s* has a secondary difference, not a tertiary.

```
0073 [17D9.0020.0002] # LATIN SMALL LETTER S
017F [17D9.0020.0004] [0000.013A.0004] # LATIN SMALL LETTER LONG S
```

6.3.4 Reducing the Repertoire

If characters are not fully supported by an implementation, then their code points can be treated as if they were unassigned. This allows them to be algorithmically constructed from code point values instead of including them in a table. This can significantly reduce the size of the required tables. See [Section 7.1, Derived Collation Elements](#) for more information.

6.3.5 Memory Table Size

Applying the above techniques, an implementation can thus safely pack all of the data for a collation element into a single 32-bit quantity: 16 for the primary, 8 for the secondary and 8 for the tertiary. Then applying techniques such as the Two-Stage table

approach described in "*Multistage Tables*" in *Section 5.1, Transcoding to Other Standards* of [\[Unicode\]](#), the mapping table from characters to collation elements can both fast and small. For an example of how this can be done, see *Section 6.10, Flat File Example*.

6.4 Avoiding Zero Bytes

If the resulting sort key is to be a C-string, then zero bytes must be avoided. This can be done by:

- using the value 0101_{16} for the level separator instead of 0000
- preprocessing the weight values to avoid zero bytes, for example by remapping 16-bit weights as follows (and larger weight values in analogous ways):

$$x \rightarrow 0101_{16} + (x / 255) * 256 + (x \% 255)$$

Where the values are limited to 8-bit quantities (as discussed above), zero bytes are even more easily avoided by just using 01 as the level separator (where one is necessary), and mapping weights by:

$$x \rightarrow 01 + x$$

6.5 Avoiding Normalization

Characters with canonical decompositions do not require mappings to collation elements, because [S1.1](#) maps them to collation elements based upon their decompositions. However, they may be given mappings to collation elements anyway. The weights in those collation elements must be computed in such a way that they will sort in the same relative location as if the characters were decomposed using Normalization Form D. Including these mappings allows an implementation handling a restricted repertoire of supported characters to compare strings correctly without performing the normalization in [S1.1](#) of the algorithm. It is recommended that implementations correctly sort all strings that are in the format known as "Fast C or D form" (FCD) even if normalization is *off*, because this permits more efficient sorting for locales whose customary characters do not use multiple combining marks. For more information on FCD, see [\[UTN5\]](#).

6.6 Case Comparisons

In some languages, it is common to sort lowercase before uppercase; in other languages this is reversed. Often this is more dependent on the individual concerned, and is not standard across a single language. It is strongly recommended that implementations provide parameterization that allows uppercase to be sorted before lowercase, and provides information as to the standard (if any) for particular countries. For more information, see *Section 5.14.13, Case Parameters* in [\[UTS35\]](#).

6.7 Incremental Comparison

Implementations do not actually have to produce full sort keys. Collation elements can

be incrementally generated as needed from two strings, and compared with an algorithm that produces the same results as sort keys would have. The choice of algorithm depends on the number of comparisons between the same strings.

- Generally incremental comparison is *more* efficient than producing full sort keys if strings are only to be compared once and if they are generally dissimilar, because differences are caught in the first few characters without having to process the entire string.
- Generally incremental comparison is *less* efficient than producing full sort keys if items are to be compared multiple times.

However, it is very tricky to produce an incremental comparison that produces correct results. For example, some implementations have not even been transitive! Be sure to test any code for incremental comparison thoroughly.

6.8 Catching Mismatches

Sort keys from two different tailored collations cannot be compared, because the weights may end up being rearranged arbitrarily. To catch this case, implementations can produce a hash value from the collation data, and prepend it to the sort key. Except in extremely rare circumstances, this will distinguish the sort keys. The implementation then has the opportunity to signal an error.

6.9 Handling Collation Graphemes

A collation ordering determines a *collation grapheme cluster* (also known as a collation grapheme or collation character), which is a sequence of characters that is treated as a primary unit by the ordering. For example, *ch* is a collation grapheme for a traditional Spanish ordering. These are generally contractions, but may include additional ignorable characters.

Roughly speaking, a collation grapheme cluster is the longest substring whose corresponding collation elements start with a non-zero primary weight, and contain as few other collation elements with non-zero primary weights as possible. In some cases, collation grapheme clusters may be *degenerate*: they may have collation elements that do not contain a non-zero weight, or they may have no non-zero weights at all.

For example, consider a collation for language in which "ch" is treated as a contraction, and "à" as an expansion. The expansion for à contains collation weights corresponding to *combining-grave* + "a" (but in an unusual order). In that case, the string <`ab`ch`à> would have the following clusters:

- *combining-grave* (a degenerate case),
- "a"
- "b`"
- "ch`"
- "à" (also a degenerate case, starting with a zero primary weight).

To find the collation grapheme cluster boundaries in a string, the following algorithm can be used:

1. Set **position** to be equal to 0, and set a boundary there.
2. If **position** is at the end of the string, set a boundary there, and return.
3. Set **startPosition = position**.
4. Fetch the next collation element(s) mapped to by the character(s) at **position**, setting **position** to the end of the character(s) mapped.
 1. This fetch must collect collation elements, including discontinuous contractions, until no characters are skipped.
 2. It cannot rewrite the input string for S2.1.3 (that would invalidate the indexes).
5. If the collation element(s) contain a collation element with a non-zero primary weight, set a boundary at **startPosition**.
6. Loop to step 2.

For information on the use of collation graphemes, see [\[UTS18\]](#).

6.10 Flat File Example

This section contains a sample flat-file binary layout and sample code for collation data. It is included only for illustration. Data in the format of *Table 16* is used to generate collation elements from characters, either going forward or backward, and detect the start of a contraction. The backward generation is for searching backward or Boyer-Moore-style searching; the contraction detection is for random access.

In the file representation, ints are 32 bit values, shorts are 16 bits, bytes are 8 bits. Negatives are two's-complement. For alignment, the ends of all arrays are padded out to multiples of 32 bits. The signature determines endianness. The locale uses an ASCII representation for the Java locale: a 2-byte ISO language code, optionally followed by '_' and 2-byte ISO country code, followed optionally by a series of variant tags separated by '_'; any unused bytes are zero.

Table 16. Flat File Format

Data	Comment
int signature;	Constant <code>0x636F6C74</code> , used also for big-endian detection
int tableVersion;	Version of the table format
int dataVersion;	Version of the table data
byte[32] locale;	Target locale (if any)
int flags;	Bit01 = 1 if Backward secondary. Other values are reserved.
int limitVariable;	Every CE below this value that has a non-zero primary is variable. Because variables are not interleaved, this does not need to be stored on a per-character basis.

int maxCharsPerCE;	Maximum number of characters that are part of a contraction	
int maxCEsPerChar;	Maximum number of collation elements that are generated by an expansion	
int indexOffset;	Offset to index table	
int collationElementsOffset;	Offset to main data table	
int expansionsOffset;	Offset to expansion table	
int contractionMatchOffset;	Offset to contraction match table	
int contractionResultOffset;	Offset to contraction values table	
int nonInitialsOffset;	Offset to non-initials table. These are used for random access.	
int[10] reserved;	Reserved	
int indexLength;	Length of following table	
int[] index;	Index for high-byte (trie) table. Contains offsets into Collation Elements. Data is accessed by: <code>ce = collationElements[index[char>>8]+char&0xFF]</code>	
int collationElementsLength;	Length of following table	
int[] collationElements;	Each element is either a real collation element, an expansionsOffset, or a contractionsOffset. See below for more information.	
int expansionsLength;	Length of following table	
int[] expansions;	The expansionOffsets in the collationElements table point into sublists in this table. Each list is terminated by FFFFFFFF.	
int contractionMatchesLength;	Length of following table	
short[] contractionMatches;	The contractionOffsets in the collationElements table point into sublists in this table. Each sublist is of the following format:	
	short backwardsOffset;	When processing backward, offset to true contractions table
	short length;	Number of chars in list to search
	short[] charsToMatch;	Characters in sorted order

int contractionCEsLength;	Length of following table
int[] contractionCEs;	List of CEs. Each corresponds to a position in the contractionChars table. The one corresponding to the length in a sublist is the bail-out; this specifies what to do if a match is not found.
int nonInitialsLength;	Length of following table
short[] nonInitials;	List of characters (in sorted order) that can be non-initials in contractions. That is, if "ch" is a contraction, then "h" is in this list. If "abcd" is a contraction, then "b", "c", and "d" are in the list.

6.10.1 Collation Element Format

- 'real' collationElement
 - 16 bits primary (FFE0..FFFF not allowed)
 - 10 bits secondary
 - 6 bits tertiary
- expansionsOffset
 - 12 bits = FFF
 - 20 bits = offset (allows for 1,048,576 items)
- contractionsOffset
 - 12 bits = FFE
 - 20 bits = offset (allows for 1,048,576 items)

An alternative structure would have the offsets be byte offsets from the start of the table, instead of indexes into the arrays. That would limit the size of the table, but use fewer machine instructions.

6.10.2 Sample Code

The following is a pseudo code using this table for the required operations. Although using Java syntax in general, the code example uses arrays, which are more familiar to users of C and C++. *The code is presented for illustration only; it is not a complete statement of the algorithm.* To make it easier to follow the logic, the pseudo-code does not handle supplementary code points or discontinuous contractions.

```
char[] input;    // input buffer (i)
int inputPos;   // position in input buffer (io)
int[] output;   // output buffer (o)
int outputPos;  // position in output buffer (io)

/**
 * Reads characters from input, writes collation elements in output
 */
```

```

void getCollationElements() {
    char c = input[inputPos++];
    int ce = collationElements[index[c>>8] + c&0xFF];
    processCE(ce);
}

/**
 * Normally just returns ce. However, special forms indicate that
 * the ce is actually an expansion, or that we have to search
 * to see if the character was part of a contraction.
 * Expansions use
 */
void processCE(int ce) {
    if (ce < 0xFFE00000) {
        output[outputPos++] = ce;
    } else if (ce >= 0xFFF00000) {
        copyExpansions(ce & 0xFFFF);
    } else {
        searchContractions(ce & 0xFFFF);
    }
}

/**
 * Search through a contraction sublist to see if there is a match.
 * Because the list is sorted, we can exit if our value is too high.
 * Because we have a length, we could implement this as a
 * binary search, although we do not right now.
 * If we do find a match, we need to recurse. That's how "abc" would
 * be handled.
 * If we fail, we return the non-matching case. That can be an expansion
 * itself (it would never be a contraction).
 */
void searchContractions(int offset) {
    offset++; // skip backwardsoffset
    int goal = input[inputPos++];
    int length = contractionMatches[offset];
    int limit = offset + 1 + length;
    for (int i = offset + 1; i < limit; ++i) {
        int cc = contractionMatches[i];
        if (cc > goal) { // definitely failed
            processCE(contractionCEs[offset]);
            break;
        } else if (cc == goal) { // found match
            processCE(contractionCEs[i]);
            break;
        }
    }
}

/**
 * Copy the expansion collation elements up to the terminator.
 * Do not use 00000000 as a terminator, because that may be a valid CE.
 * These elements do not recurse.
 */
void copyExpansions (int offset) {
    int ce = expansions[offset++];
    while (ce != 0xFFFFFFFF) {
        output[outputPos++] = ce;
        ce = expansions[offset++];
    }
}

```

7 Weight Derivation

This section describes the generation of the Default Unicode Collation Element Table (DUCET), and the assignment of weights to code points that are not explicitly mentioned in that table. The assignment of weights uses information derived from the Unicode Character Database [[UAX44](#)].

7.1 Derived Collation Elements

CJK ideographs and Hangul syllables are not explicitly mentioned in the default table. CJK ideographs are mapped to collation elements that are derived from their Unicode code point value as described in *Section 7.1.3, [Implicit Weights](#)*. For a discussion of derived collation elements for Hangul syllables and other issues related to the collation of Korean, see *Section 7.1.5, [Hangul Collation](#)*.

7.1.1 Handling Ill-Formed Code Unit Sequences

Unicode strings sometimes contain ill-formed code unit sequences. Such ill-formed sequences must not be interpreted as valid Unicode characters. See *Section 3.2, Conformance Requirements* in [[Unicode](#)]. For example, expressed in UTF-32, a Unicode string might contain a 32-bit value corresponding to a surrogate code point (General_Category Cs) or an out-of-range value (< 0 or > 10FFFF), or a UTF-8 string might contain misconverted byte values that cannot be interpreted. Implementations of the Unicode Collation Algorithm may choose to treat such ill-formed code unit sequences as error conditions and respond appropriately, such as by throwing an exception.

An implementation of the Unicode Collation Algorithm may also choose not to treat ill-formed sequences as an error condition, but instead to give them explicit weights. This strategy provides for determinant comparison results for Unicode strings, even when they contain ill-formed sequences. However, to avoid security issues when using this strategy, ill-formed code sequences should not be given an ignorable `or variable` primary weight.

There are two recommended approaches, based on how these ill-formed sequences are typically handled by character set converters.

- The first approach is to weight each maximal ill-formed subsequence as if it were U+FFFD REPLACEMENT CHARACTER. (For more information about maximal ill-formed subsequences, see *Section 3.9, Unicode Encoding Forms* in [[Unicode](#)].)
- A second approach is to generate an implicit weight for any surrogate code point as if it were an unassigned code point, using the method of *Section 7.1.3, [Implicit Weights](#)*.

7.1.2 Unassigned and Other Code Points

Each unassigned code point and each other code point that is not explicitly mentioned in the table is mapped to a sequence of two collation elements as described in *Section 7.1.3, [Implicit Weights](#)*.

7.1.3 Implicit Weights

This section describes how a code point is mapped to an implicit weight. The result of this process consists of collation elements that are sorted in code point order, that do not collide with any explicit values in the table, and that can be placed anywhere (for example, at BASE) with respect to the explicit collation element mappings. By default, implicit mappings are given higher weights than all explicit collation elements (except those with decompositions to characters with implicit weights).

To derive the collation elements, the value of the code point is used to calculate two numbers, by bit shifting and bit masking. The bit operations are chosen so that the resultant numbers have the desired ranges for constructing implicit weights. The first number is calculated by taking the code point expressed as a 32-bit binary integer CP and bit shifting it right by 15 bits. Because code points range from U+0000 to U+10FFFF, the result will be a number in the range 0 to 2^{16} ($= 33_{10}$). This number is then added to the special value BASE.

$$AAAA = \text{BASE} + (\text{CP} \gg 15);$$

Now mask off the bottom 15 bits of CP. OR a 1 into bit 15, so that the resultant value is non-zero.

$$BBBB = (\text{CP} \& 0x7FFF) | 0x8000;$$

AAAA and BBBB are interpreted as unsigned 16-bit integers. The implicit weight mapping given to the code point is then constructed as:

$$[AAAA.0020.0002][BBBB.0000.0000]$$

If a fourth or higher weights are used, then the same pattern is followed for those weights. They are set to a non-zero value in the first collation element and zero in the second. (Because all distinct code points have a different **AAAA/BBBB** combination, the exact non-zero value does not matter.)

The value for BASE depends on the type of character. The first BASE value is for the core Han Unified Ideographs. The second BASE value is for all other Unified Han ideographs. In both of these cases, compatibility decomposables are excluded, because they are otherwise handled in the UCA. Unassigned code points are also excluded from these first two BASE values. The final BASE value is for all other code points, including unassigned code points.

Table 17. Values for Base

Base	Applicable Ranges
FB40	Unified_Ideograph=True AND ((Block=CJK_Unified_Ideograph) OR (Block=CJK_Compatibility_Ideographs))

	In regex notation: <code>[\\p{unified_ideograph}&\\p{Block=CJK_Unified_Ideographs}\\p{Block=CJK_Compatibility_Ideographs}]</code>
FB80	Unified_Ideograph=True AND NOT ((Block=CJK_Unified_Ideograph) OR (Block=CJK_Compatibility_Ideographs)) In regex notation: <code>[\\p{unified_ideograph}-[\\p{Block=CJK_Unified_Ideographs}\\p{Block=CJK_Compatibility_Ideographs}]]</code>
FBC0	Any other code point

These results make AAAA (in each case) larger than any explicit primary weight; thus the implicit weights will not collide with explicit weights. It is not generally necessary to tailor these values to be within the range of explicit weights. However if this is done, the explicit primary weights must be shifted so that none are between each of the BASE values and BASE + 34.

7.1.4 Trailing Weights

The range of 1,024 primary weights from FC00 to FFFC are available for use as trailing weights.

In many writing systems, the convention for collation is to order by syllables (or other units similar to syllables). In most cases a good approximation to syllabic ordering can be obtained in the UCA by weighting initial elements of syllables in the appropriate primary order, followed by medial elements (such as vowels), followed by final elements, if any. The default weights for the UCA in the DUCET are assigned according to this general principle for many scripts. This approach handles syllables within a given script fairly well, but unexpected results can occur when syllables of different lengths are adjacent to characters with higher primary weights, as illustrated in the following example:

Case 1	Case 2								
<table border="1" style="width: 100%; text-align: center;"> <tr><td style="width: 20px;">1</td><td>{G}{A}</td></tr> <tr><td>2</td><td>{G}{A}{K}</td></tr> </table>	1	{G}{A}	2	{G}{A}{K}	<table border="1" style="width: 100%; text-align: center;"> <tr><td style="width: 20px;">2</td><td>{G}{A}{K}事</td></tr> <tr><td>1</td><td>{G}{A}事</td></tr> </table>	2	{G}{A}{K}事	1	{G}{A}事
1	{G}{A}								
2	{G}{A}{K}								
2	{G}{A}{K}事								
1	{G}{A}事								

In this example, the symbols {G}, {A}, and {K} represent letters in a script where syllables (or other sequences of characters) are sorted as units. By proper choice of weights for the individual letters, the syllables can be ordered correctly. However, the weights of the following characters may cause syllables of different lengths to change

order. Thus {G}{A}{K} comes after {G}{A} in Case 1, but in Case 2, it comes *before*. That is, the order of these two syllables would be reversed when each is followed by a CJK ideograph, with a high primary weight: in this case, U+4E8B (事).

This unexpected behavior can be avoided by using trailing weights to tailor the non-initial letters in such syllables. The trailing weights, by design, have higher values than the primary weights for characters in all scripts, including the implicit weights used for CJK ideographs. Thus in the example, if {K} is tailored with a trailing weight, it would have a higher weight than any CJK ideograph, and as a result, the relative order of the two syllables {G}{A}{K} and {G}{A} would not be affected by the presence of a CJK ideograph following either syllable.

The primary weights from FFFD to FFFF are reserved for special collation elements. For example, in DUCET, U+FFFD maps to a collation element with the fixed primary weight of FFFD, thus ensuring that it is not a [variable collation element](#). This means that implementations using U+FFFD as a replacement for [ill-formed code unit sequences](#) will not have those replacement characters ignored in collation.

7.1.5 Hangul Collation

The Hangul script for Korean is in a rather unique position, because of its large number of precomposed syllable characters, and because those precomposed characters are the normal (NFC) form of interchanged text. For Hangul syllables to sort correctly, either the DUCET table must be tailored or both the UCA algorithm and the table must be tailored. The essential problem results from the fact that Hangul syllables can also be represented with a sequence of conjoining jamo characters and because syllables represented that way may be of different lengths, with or without a trailing consonant jamo. That introduces the trailing weights problem, as discussed in [Section 7.1.4, *Trailing Weights*](#). This section describes several approaches which implementations may take for tailoring to deal with the trailing weights problem for Hangul.

Note: The Unicode Technical Committee recognizes that it would be preferable if a single "best" approach could be standardized and incorporated as part of the specification of the UCA algorithm and the DUCET table. However, picking a solution requires working out a common approach to the problem with the ISO SC2 OWG-Sort group, which takes considerable time. In the meantime, implementations can choose among the various approaches discussed here, when faced with the need to order Korean data correctly.

The following discussion makes use of definitions and abbreviations from [Section 3.12, *Conjoining Jamo Behavior*](#) in [\[Unicode\]](#). In addition, a special symbol (Ⓢ) is introduced to indicate a terminator weight. For convenience in reference, these conventions are summarized here:

Description	Abbr.	Weight
Leading consonant	L	W_L
Vowel	V	W_V

Trailing consonant	T	W_T
Terminator weight	-	Ⓣ

Simple Method

The specification of the Unicode Collation Algorithm requires that Hangul syllables be decomposed. However, if the weight table is tailored so that the primary weights for Hangul jamo are adjusted, then the Hangul syllables can be left as single code points and be treated in much the same way as CJK ideographs. The adjustment is specified as follows:

1. Tailor each L to have a primary weight corresponding to the first Hangul syllable starting with that jamo.
2. Tailor all Vs and Ts to be ignorable at the primary level.

The net effect of such a tailoring is to provide a Hangul collation which is approximately equivalent to one of the more complex methods specified below. This may be sufficient in environments where individual jamo are not generally expected.

Three more complex and complete methods are spelled out below. First the nature of the tailoring is described. Then each method is exemplified, showing the implications for the relative weighting of jamo and illustrating how each method produces correct results.

Each of these three methods can correctly represent the ordering of all Hangul syllables, both for modern Korean and for Old Korean. However, there are implementation trade-offs between them. These trade-offs can have a significant impact on the acceptability of a particular implementation. For example, substantially longer sort keys will cause serious performance degradations and database index bloat. Some of the pros and cons of each method are mentioned in the discussion of each example. Note that if the repertoire of supported Hangul syllables is limited to those required for modern Korean (those of the form LV or LVT), then each of these methods becomes simpler to implement.

Data Method

1. Tailor the Vs and Ts to be Trailing Weights, with the ordering $T < V$
2. Tailor each sequence of multiple L's that occurs in the repertoire as a contraction, with an independent primary weight after any prefix's weight.

For example, if L_1 has a primary weight of 555, and L_2 has a primary weight of 559, then the sequence L_1L_2 would be treated as a contraction and be given a primary weight chosen from the range 556 to 558.

Terminator Method

1. Add an internal terminator primary weight (Ⓣ).
2. Tailor all jamo so that $Ⓣ < T < V < L$

3. Algorithmically add the terminator primary weight (Ⓣ) to the end of every standard Korean syllable block.

The details of the algorithm for parsing Hangeul data into standard Korean syllable blocks can be found in *Section 8, Hangeul Syllable Boundary Determination* of [\[UAX29\]](#)

Interleaving Method

The interleaving method requires tailoring both the DUCET table and the way the algorithm handles Korean text.

Generate a tailored weight table by assigned an explicit primary weight to each precomposed Hangeul syllable character, with a 1-weight gap between each one. (See *Section 6.2, Large Weight Values.*)

Separately define a small, internal table of jamo weights. This internal table of jamo weights is separate from the tailored weight table, and is only used when processing standard Korean syllable blocks. Define this table as follows:

1. Give each jamo a 1-byte weight.
2. Add an internal terminator 1-byte weight (Ⓣ).
3. Assign these values so that: $\text{Ⓣ} < T < V < L$.

When processing a string to assign collation weights, whenever a substring of jamo and/or precomposed Hangeul syllables is encountered, break it into standard Korean syllable blocks. For each syllable identified, assign a weight as follows:

1. If a syllable is canonically equivalent to one of the precomposed Hangeul syllable characters, then assign the weight based on the tailored weight table.
2. If a syllable is not canonically equivalent to one of the precomposed Hangeul syllable characters, then assign a weight sequence by the following steps:
 - a. Find the greatest precomposed Hangeul syllable that the parsed standard Korean syllable block is greater than. Call that the "base syllable".
 - b. Take the weight of the base syllable from the tailored weight table and increment by one. This will correspond to the gap weight in the table.
 - c. Concatenate a weight sequence consisting of the gap weight, followed by a byte weight for each of the jamo in the decomposed representation of the standard Korean syllable block, followed by the byte for the terminator weight.

Data Method Example

The data method provides for the following order of weights, where the X_b are all the scripts sorted before Hangeul, and the X_a are all those sorted after.

X_b	L	X_a	T	V
-------	---	-------	---	---

This ordering gives the right results among the following:

Chars	Weights				Comments
$L_1 V_1 X_a$	W_{L1}	W_{V1}	W_{Xa}		
$L_1 V_1 L_{n...}$	W_{L1}	W_{V1}	W_{Ln}	...	
$L_1 V_1 X_b$	W_{L1}	W_{V1}	W_{Xb}		
$L_1 V_1 T_1$	W_{L1}	W_{V1}	W_{T1}		Works because $W_T >$ all W_X and W_L
$L_1 V_1 V_2$	W_{L1}	W_{V1}	W_{V2}		Works because $W_V >$ all W_T
$L_1 L_2 V_1$	W_{L1L2}	W_{V1}			Works if $L_1 L_2$ is a contraction

The disadvantages of the data method are that the weights for T and V are separated from those of L, which can cause problems for sort key compression, and that a combination of LL that is outside the contraction table will not sort properly.

Terminator Method Example

The terminator method would assign the following weights:

⊕	X_b	T	V	L	X_a
---	-------	---	---	---	-------

This ordering gives the right results among the following:


Chars	Weights				Comments
$L_1 V_1 X_a$	W_{L1}	W_{V1}	⊕	W_{Xa}	
$L_1 V_1 L_{n...}$	W_{L1}	W_{V1}	⊕	W_{Ln}	...
$L_1 V_1 X_b$	W_{L1}	W_{V1}	⊕	W_{Xb}	
$L_1 V_1 T_1$	W_{L1}	W_{V1}	W_{T1}	⊕	Works because $W_T >$ all W_X and ⊕
$L_1 V_1 V_2$	W_{L1}	W_{V1}	W_{V2}	⊕	Works because $W_V >$ all W_T
$L_1 L_2 V_1$	W_{L1}	W_{L2}	W_{V1}	⊕	Works because $W_L >$ all W_V

The disadvantages of the terminator method are that an extra weight is added to all Hangul syllables, increasing the length of sort keys by roughly 40%, and the fact that the terminator weight is non-contiguous can disable sort key compression.

Interleaving Method Example

The interleaving method provides for the following assignment of weights. W_n represents the weight of a Hangul syllable, and W'_n is the weight of the gap right after it.

The L, V, T weights will only occur after a W, and thus can be considered part of an entire weight.



byte weights:



This ordering gives the right results among the following:

Chars	Weights	Comments
$L_1 V_1 X_a$	W_n X_a	
$L_1 V_1 L_n \dots$	W_n $W_k \dots$	The L_n will start another syllable
$L_1 V_1 X_b$	W_n X_b	
$L_1 V_1 T_1$	W_m	Works because $W_m > W_n$
$L_1 V_1 V_2$	$W_{m'L1V1V2\text{Ⓣ}}$	Works because $W_{m'} > W_m$
$L_1 L_2 V_1$	$W_{m'L1L2V1\text{Ⓣ}}$	Works because the byte weight for $L_2 > \text{all } v$

The interleaving method is somewhat more complex than the others, but produces the shortest sort keys for all of the precomposed Hangul syllables, so for normal text it will have the shortest sort keys. If there were a large percentage of ancient Hangul syllables, the sort keys would be longer than other methods.

7.2 Tertiary Weight Table

Characters are given tertiary weights according to *Table 18*. The Decomposition Type is from the Unicode Character Database [UAX44]. The Case or Kana Subtype entry refers either to a case distinction or to a specific list of characters. The weights are from MIN = 2 to MAX = $1F_{16}$, excluding 7, which is not used for historical reasons. The MAX value $1F$ was used for some trailing collation elements. This usage began with UCA version 9 (Unicode 3.1.1) and continued until UCA version 6.2. It is no longer used in DUCET.

The Samples show some minimal values that are distinguished by the different weights. All values are distinguished. The samples have empty cells when there are no (visible) values showing a distinction.

Table 18. Tertiary Weight Assignments

Decomposition Type	Case or Kana Subtype	Weight	Samples					
NONE		0x0002	i	ı)	mw	1/2	X

<wide>		0x0003	i					
<compat>		0x0004	i					
		0x0005	i					
<circle>		0x0006	⓪					
!unused!		0x0007						
NONE	Uppercase	0x0008	l			MW		
<wide>	Uppercase	0x0009	l)			
<compat>	Uppercase	0x000A	l					
	Uppercase	0x000B	ŷ					
<circle>	Uppercase	0x000C	Ⓛ					
<small>	small hiragana (3041, 3043, ...)	0x000D						ぁ
NONE	normal hiragana (3042, 3044, ...)	0x000E						あ
<small>	small katakana (30A1, 30A3, ...)	0x000F)			ァ
<narrow>	small narrow katakana (FF67..FF6F)	0x0010						ヾ
NONE	normal katakana (30A2, 30A4, ...)	0x0011						ア
<narrow>	narrow katakana (FF71..FF9D), narrow hangul (FFA0..FFDF)	0x0012						ヾ
<circle>	circled katakana (32D0..32FE)	0x0013						㊦
<super>		0x0014)			
<sub>		0x0015)			
<vertical>		0x0016)			
<initial>		0x0017			!			
<medial>		0x0018			!			
<final>		0x0019			!			
<isolated>		0x001A			!			

<noBreak>		0x001B						
<square>		0x001C			mW			
<square>, <super>, <sub>	Uppercase	0x001D			MW			
<fraction>		0x001E				½		
n/a	(MAX value)	0x001F						

The <compat> weight 0x0004 is given to characters that do not have more specific decomposition types. It includes superscripted and subscripted combining letters, for example U+0365 COMBINING LATIN SMALL LETTER I and U+1DCA COMBINING LATIN SMALL LETTER R BELOW. These combining letters occur in abbreviations in Medieval manuscript traditions.

The MAX weight is used to prevent certain overlaps between collation elements. In particular, if it were not used there would be some unexpected cases where $X \neq \frac{Y}{3}$, but $XY = \frac{YX}{3}$. For example, consider:

2024 ; [0273.0020.0004] # ONE DOT LEADER
 2025 ; [0273.0020.0004][0273.0020.001F] # TWO DOT LEADER

If the last GE were not given MAX ($1F_{16}$) as a tertiary, then

- ONE DOT LEADER $\neq \frac{TWO DOT LEADER}{3}$, but
- ONE DOT LEADER + TWO DOT LEADER = $\frac{TWO DOT LEADER + ONE DOT LEADER}{3}$

7.3 Fourth-Level Weight Assignments

The assignment of fourth-level weights in the Default Unicode Collation Element Table is done as follows:

1. If a character is ignorable at the first 3 levels, and is a control code (gc=Cc), format control (gc=Cf), or variation selector, then assign "0000" for the fourth level.
2. If a character is weighted as an expansion based on a compatibility decomposition or a synthetic expansion, then assign the code point of the character itself as the fourth-level weight for each element of the expansion.
3. If a character is weighted as an expansion based on a canonical decomposition, then assign the code point of each character in the decomposition as the fourth-level weight for the corresponding element of the expansion.
4. Otherwise, assign the code point of the character itself as the fourth-level weight.

Conditions 1 to 3 are disjoint. Condition 4 is the default assignment of fourth-level weights; it applies in any case which does not meet one of the first three conditions.

8 Searching and Matching

Language-sensitive searching and matching are closely related to collation. Strings that compare as equal at some strength level should be matched when doing language-sensitive matching. For example, at a primary strength, "ß" would match against "ss" according to the UCA, and "aa" would match "å" in a Danish tailoring of the UCA. The main difference from the collation comparison operation is that the ordering is not important. Thus for matching it does not matter that "å" would sort after "z" in a Danish tailoring—the only relevant information is that they do not match.

The basic operation is matching: determining whether string X matches string Y. Other operations are built on this:

- Y contains X when there is some substring of Y that matches X
- A search for a string X in a string Y succeeds if Y contains X.
- Y starts with X when some initial substring of Y matches X
- Y ends with X when some final substring of Y matches X

The collation settings determine the results of the matching operation (see [Section 5.1, Parametric Tailoring](#)). Thus users of searching and matching need to be able to modify parameters such as locale or comparison strength. For example, setting the strength to exclude differences at Level 3 has the effect of ignoring case and compatibility format distinctions between letters when matching. Excluding differences at Level 2 has the effect of also ignoring accentual distinctions when matching.

Conceptually, a string matches some target where a substring of the target has the same sort key, but there are a number of complications:

1. The lengths of matching strings may differ: "aa" and "å" would match in Danish.
2. Because of ignorables (at different levels), there are different possible positions where a string matches, depending on the attribute settings of the collation. For example, if hyphens are ignorable for a certain collation, then "abc" will match "abc", "ab-c", "abc-", "-abc-", and so on.
3. Suppose that the collator has contractions, and that a contraction spans the boundary of the match. Whether it is considered a match may depend on user settings, just as users are given a "Whole Words" option in searching. So in a language where "ch" is a contraction with a different primary from "c", "bac" would not match in "bach" (given the proper user setting).
4. Similarly, combining character sequences may need to be taken into account. Users may not want a search for "abc" to match in "...abç..." (with a cedilla on the c). However, this may also depend on language and user customization. In particular, a useful technique is discussed in [Section 8.2, Asymmetric Search](#).
5. The above two conditions can be considered part of a general condition: "Whole Characters Only"; very similar to the common "Whole Words Only" checkbox that is included in most search dialog boxes. (For more information on grapheme clusters and searching, see [\[UAX29\]](#) and [\[UTS18\]](#).)
6. If the matching does not check for "Whole Characters Only," then some other complications may occur. For example, suppose that P is "x^", and Q is "x ^, ". Because the cedilla and circumflex can be written in arbitrary order and still be equivalent, in most cases one would expect to find a match for P in Q. A

canonically-equivalent matching process requires special processing at the boundaries to check for situations like this. (It does not require such special processing within the P or the substring of Q because collation is defined to observe canonical equivalence.)

The following are used to provide a clear definition of searching and matching that deal with the above complications:

DS1. Define $S[start,end]$ to be the substring of S that includes the character after the offset *start* up to the character before offset *end*. For example, if S is "abcd", then $S[1,3]$ is "bc". Thus $S = S[0,length(S)]$.

DS1a. A boundary condition is a test imposed on an offset within a string. An example includes Whole Word Search, as defined in [\[UAX29\]](#).

The tailoring parameter *match-boundaries* specifies constraints on matching (see [Section 5.1, Parametric Tailoring](#)). The parameter *match-boundaries=whole-character* requires that the start and end of a match each be on a grapheme boundary. The value *match-boundaries=whole-character* further requires that the start and end of a match each be on a word boundary as well. For more information on the specification of these boundaries, see [\[UAX29\]](#).

By using grapheme-complete conditions, contractions and combining sequences are not interrupted except in edge cases. This also avoids the need to present visually discontinuous selections to the user (except for BIDI text).

Suppose there is a collation C, a pattern string P and a target string Q, and a boundary condition B. C has some particular set of attributes, such as a strength setting, and choice of variable weighting.

DS2. The pattern string P *has a match at* $Q[s,e]$ *according to collation C* if C generates the same sort key for P as for $Q[s,e]$, and the offsets *s* and *e* meet the boundary condition B. One can also say P has a match in Q according to C.

DS3. The pattern string P has a *canonical match at* $Q[s,e]$ *according to collation C* if there is some Q' that is canonically equivalent to $Q[s,e]$, and P has a match in Q'.

For example, suppose that P is "À", and Q is "...A◌◌...". There would not be a match for P in Q, but there would be a canonical match, because P does have a match in "A◌◌", which is canonically equivalent to "A◌◌". However, it is not commonly necessary to use canonical matches, so this definition is only supplied for completeness.

Each of the following definitions is a qualification of DS2 or DS3:

DS3a. The match is *grapheme-complete* if B requires that the offset be at a grapheme cluster boundary. Note that Whole Word Search as defined in [\[UAX29\]](#) is grapheme complete.

DS4. The match is *minimal* if there is no match at $Q[s+i,e-j]$ for any *i* and *j* such that $i \geq 0$, $j \geq 0$, and $i + j > 0$. In such a case, one can also say that P has a *minimal match at*

Q[s,e].

DS4a. A *medial* match is determined in the following way:

1. Determine the minimal match for P at Q[s,e]
2. Determine the "minimal" pattern P[m,n], by finding:
 1. the largest m such that P[m,len(P)] matches P, then
 2. the smallest n such that P[m,n] matches P.
3. Find the smallest $s' \leq s$ such that Q[s',e] is canonically equivalent to P[m',n] for some m'.
4. Find the largest $e' \geq e$ such that Q[s',e'] is canonically equivalent to P[m', n'] for some n'.
5. The medial match is Q[s', e'].

DS4b. The match is *maximal* if there is no match at Q[s-i,e+j] for any i and j such that $i \geq 0, j \geq 0$, and $i + j > 0$. In such a case, one can also say that P has a *maximal* match at Q[s,e].

Figure 5 illustrates the differences between these type of matches, where the collation strength is set to ignore punctuation and case, and **format** indicates the match.

Figure 5. Minimal, Medial, and Maximal Matches

	Text	Description
Pattern	*!abc!*	Notice that the *! and !* are ignored in matching.
Target Text	def\$!Abc%\$ghi	
Minimal Match	def\$!Abc%\$ghi	The minimal match is the tightest one, because !\$ and %\$ are ignored in the target.
Medial Match	def\$!Abc%\$ghi	The medial one includes those characters that are binary equal.
Maximal Match	def\$!Abc%\$ghi	The maximal match is the loosest one, including the surrounding ignored characters.

By using minimal, maximal, or medial matches, the issue with ignorables is avoided. Medial matches tend to match user expectations the best.

When an additional condition is set on the match, the types (minimal, maximal, medial) are based on the matches *that meet that condition*. Consider the example in Figure 6.

Figure 6. Alternate End Points for Matches

	Value	Notes
Pattern	abc	

Strength	primary	thus ignoring combining marks, punctuation
Text	abc _, -°d	two combining marks, cedilla and ring
Matches	abc _, [°] d	four possible end points, indicated by

If, for example, the condition is Whole Grapheme, then the matches are restricted to "abc_,|[°]|d", thus discarding match positions that would not be on a grapheme cluster boundary. In this case the minimal match would be "abc_,|[°]d"

DS6. The *first forward match* for P in Q starting at *b* is the least offset *s* greater than or equal to *b* such that for some *e*, P matches within Q[*s*,*e*].

DS7. The *first backward match* for P in Q starting at *b* is the greatest offset *s* less than or equal to *b* such that for some *e*, P matches within Q[*s*,*e*].

In DS6 and DS7, matches can be minimal, medial, or maximal; the only requirement is that the combination in use in DS6 and DS7 be specified. Of course, a possible match can also be rejected on the basis of other conditions, such as being grapheme-complete or applying Whole Word Search, as described in [\[UAX29\]](#).

The choice of medial or minimal matches for the "starts with" or "ends with" operations only affects the positioning information for the end of the match or start of the match, respectively.

Special Cases. Ideally, the UCA at a secondary level would be compatible with the standard Unicode case folding and removal of compatibility differences, especially for the purpose of matching. For the vast majority of characters, it is compatible, but there are the following exceptions:

1. The UCA maintains compatibility with the DIN standard for sorting German by having the German *sharp-s* (U+00DF (ß) LATIN SMALL LETTER SHARP S) sort as a secondary difference with "SS", instead of having ß and SS match at the secondary level.
2. Compatibility normalization (NFKC) folds stand-alone accents to a combination of space + combining accent. This was not the best approach, but for backwards compatibility cannot be changed in NFKC. UCA takes a better approach to weighting stand-alone accents, but as a result does not weight them exactly the same as their compatibility decompositions.
3. Case folding maps *iota-subscript* (U+0345 (ι) COMBINING GREEK YPOGEGRAMMENI) to an iota, due to the special behavior of *iota-subscript*, while the UCA treats *iota-subscript* as a regular combining mark (secondary ignorable).
4. When compared to their case and compatibility folded values, UCA compares the following as different at a secondary level, whereas other compatibility differences are at a tertiary level.
 - U+017F (f) LATIN SMALL LETTER LONG S (and precomposed characters containing it)
 - U+1D4C (ʒ) MODIFIER LETTER SMALL TURNED OPEN E
 - U+2D6F (̆) TIFINAGH MODIFIER LETTER LABIALIZATION MARK

In practice, most of these differences are not important for modern text, with one exception: the German ß. Implementations should consider tailoring ß to have a tertiary difference from SS, at least when collation tables are used for matching. Where full compatibility with case and compatibility folding are required, either the text can be preprocessed, or the UCA tables can be tailored to handle the outlying cases.

8.1 Collation Folding

Matching can be done by using the collation elements, directly, as discussed above. However, because matching does not use any of the ordering information, the same result can be achieved by a folding. That is, two strings would fold to the same string if and only if they would match according to the (tailored) collation. For example, a folding for a Danish collation would map both "Gård" and "gaard" to the same value. A folding for a primary-strength folding would map "Resume" and "résumé" to the same value. That folded value is typically a lowercase string, such as "resume".

A comparison between folded strings cannot be used for an ordering of strings, but it can be applied to searching and matching quite effectively. The data for the folding can be smaller, because the ordering information does not need to be included. The folded strings are typically much shorter than a sort key, and are human-readable, unlike the sort key. The processing necessary to produce the folding string can also be faster than that used to create the sort key.

The following is an example of the mappings used for such a folding using to the [\[CLDR\]](#) tailoring of UCA:

Parameters:

```
{locale=da_DK, strength=secondary, alternate=shifted}
```

Mapping:

...

^a → a Map compatibility (tertiary) equivalents, such as full-width and superscript characters, to representative character(s)

a → a

A → a

A → a

^a → a

...

å → aa Map contractions (a + ring above) to equivalent values

Å → aa

...

Once the table of such mappings is generated, the folding process is a simple

longest-first match-and-replace: a string to be folded is first converted to NFD, then at each point in the string, the longest match from the table is replaced by the corresponding result.

However, ignorable characters need special handling. Characters that are fully ignorable at a given strength level normally map to the empty string. For example, at *strength=quaternary*, most controls and format characters map to the empty string; at *strength=primary*, most combining marks also map to the empty string. In some contexts, however, fully ignorable characters may have an effect on comparison, or characters that are not ignorable at the given strength level may be treated as ignorable.

1. Any discontinuous contractions need to be detected in the process of folding and handled according to Rule [S2.1](#). For more information about discontinuous contractions, see [Section 3.3.2, *Contractions*](#).
2. An ignorable character may interrupt what would otherwise be a contraction. For example, suppose that "ch" is a contraction sorting after "h", as in Slovak. In the absence of special tailoring, a CGJ or SHY between the "c" and the "h" prevents the contraction from being formed, and causes "c<CGJ>h" to not compare as equal to "ch". If the CGJ is simply folded away, they would incorrectly compare as equal. See also [Section 5.3, *Use of Combining Grapheme Joiner*](#).
3. With the parameter values *alternate=shifted* or *alternate=blanked*, any (partially) ignorable characters after variable collation elements have their weights reset to zero at levels 1 to 3, and may thus become fully ignorable. In that context, they would also be mapped to the empty string. For more information, see [Section 3.6.2, *Variable Weighting*](#).

8.2 Asymmetric Search

Users often find *asymmetric searching* to be a useful option. When doing an asymmetric search, a character (or grapheme cluster) in the query that is *unmarked* at the secondary and/or tertiary levels will match a character in the target that is either marked or unmarked at the same levels, but a character in the query that is *marked* at the secondary and/or tertiary levels will only match a character in the target that is marked in the same way.

At a given level, a character is unmarked if it has the lowest collation weight for that level. For the tertiary level, a plain lowercase 'r' would normally be treated as unmarked, while the uppercase, fullwidth, and circled characters 'R', ' r ', 'Ⓐ' would be treated as marked. There is an exception for *kana* characters, where the "normal" form is unmarked: 0x000E for *hiragana* and 0x0011 for *katakana*.

For the secondary level, an unaccented 'e' would be treated as unmarked, while the accented letters 'é', 'è' would (in English) be treated as marked. Thus in the following examples, a lowercase query character matches that character or the uppercase version of that character even if *strength* is set to tertiary, and an unaccented query character matches that character or any accented version of that character even if *strength* is set to secondary.

Asymmetric search with strength = tertiary

Query	Target Matches
resume	resume, Resume, RESUME, résumé, rèsumè, Résumé, RÉSUMÉ, ...
Resume	Resume, RESUME, Résumé, RÉSUMÉ, ...
résumé	résumé, Résumé, RÉSUMÉ, ...
Résumé	Résumé, RÉSUMÉ, ...
けんこ	けんこ, げんこ, けんご, げんご, ...
げんご	げんご, ...

Asymmetric search with strength = secondary

Query	Target Matches
resume	resume, Resume, RESUME, résumé, rèsumè, Résumé, RÉSUMÉ, ...
Resume	resume, Resume, RESUME, résumé, rèsumè, Résumé, RÉSUMÉ, ...
résumé	résumé, Résumé, RÉSUMÉ, ...
Résumé	résumé, Résumé, RÉSUMÉ, ...
けんこ	けんこ, ケンコ, げんこ, けんご, ゲンコ, ケンゴ, げんご, ゲンゴ, ...
げんご	げんご, ゲンゴ, ...

8.2.1 Returning Results

When doing an asymmetric search, there are many ways in which results might be returned:

1. Return the next single match in the text.
2. Return an unranked set of all the matches in the text, which could be used for highlighting all of the matches on a page.
3. Return a set of matches in which each match is ranked or ordered based on the closeness of the match. The closeness might be determined as follows:
 - The closest matches are those in which there is no secondary difference between the query and target; the closeness is based on the number of tertiary differences.
 - These are followed by matches in which there is a secondary difference between query and target, ranked first by number of secondary differences, and then by number of tertiary differences.

9 Data Files

The data files for each version of UCA are located in versioned subdirectories in [\[Data10\]](#). The main data file with the DUCET data for each version is **allkeys.txt** [\[Allkeys\]](#).

Starting with Version 3.1.1 of UCA, the data directory also contains **CollationTest.zip**, a zipped file containing conformance test files. The documentation file **CollationTest.html** describes the format and use of those test files. See also [\[Tests10\]](#).

~~Starting with Version 6.0.0 of UCA, the data directory also contains **CollationAuxiliary.zip**, a zipped file containing auxiliary data files used for CLDR, and for comparison across versions. The documentation file **CollationAuxiliary.html** describes the format and use of those data files.~~

Starting with Version 6.2.0 of UCA, the data directory also contains **decomps.txt**. This file lists the decompositions used when generating the DUCET. These decompositions are loosely based on the normative decomposition mappings defined in the Unicode Character Database, [often mirroring the NFKD form](#). However, those decomposition mappings are adjusted as part of the input to the generation of DUCET, in order to produce default weights more appropriate for collation. For more details and a description of the file format, see the header of the **decomps.txt** file.

9.1 Allkeys File Format

~~This used to be section 3.6.1.~~

~~Each of the files~~ The **allkeys.txt** file consists of a version line followed by [optional lines for parametric attributes](#), and a series of entries, all separated by newlines. A '#' or '%' and any following characters on a line are comments. Whitespace between literals is ignored. The following is an extended BNF description of the format, where "x+" indicates one or more x's, "x*" indicates zero or more x's, "x?" indicates zero or one x, <char> is a hexadecimal Unicode code point value, and <weight> is a hexadecimal collation weight value.

```
<collationElementTable> := <version>
                           <entry>+
```

The version line is of the form:

```
<version> := '@version' <major>.<minor>.<variant> <eol>
```

Each entry is a mapping from character(s) to collation element(s), and is of the following form:

```
<entry>          := <charList> ';' <collElement>+ <eol>
<charList>      := <char>+
<collElement>   := "[" <alt> <weight> "." <weight> "." <weight> ("." <weight>)? "]"
<alt>           := "*" | "."
```

Collation elements marked with a "*" are [variable](#).

Every collation element in the table should have the same number of fields.

Here are some selected entries taken from a particular version of the data file. (It may not match the actual values in the current data file.)

```
0020 ; [*0209.0020.0002-0020] % SPACE
02DA ; [*0209.002B.0002-02DA] % RING ABOVE
```

```

0041 ; [06D9.0020.0008.0041] % LATIN CAPITAL LETTER A
3373 ; [06D9.0020.0017.0041] [08C0.0020.0017.0055] % SQUARE AU
00C5 ; [06D9.002B.0008.00C5] % LATIN CAPITAL LETTER A WITH RING ABOVE
212B ; [06D9.002B.0008.212B] % ANGSTROM SIGN
0042 ; [06EE.0020.0008.0042] % LATIN CAPITAL LETTER B
0043 ; [0706.0020.0008.0043] % LATIN CAPITAL LETTER C
0106 ; [0706.0022.0008.0106] % LATIN CAPITAL LETTER C WITH ACUTE
0044 ; [0712.0020.0008.0044] % LATIN CAPITAL LETTER D

```

Although this document describes collation elements as three levels, the DUCET contains a fourth level (as in [0712.0020.0008.0044]), which can be used to approximate Step [S3.10](#). Note, however, that this approximation is not very good, the UCA does not use this fourth level of data, and this data is not related to the fourth level introduced by variable handling and thus leads to confusion. So implementations should only use it at their own risk.

This DUCET fourth level is computable: In most cases the fourth level is equal to the code point itself. For composite characters which map to a sequence of collation elements, the fourth level for each collation element is based on the decomposition of the character. For further details see [Section 7.3](#), [Fourth-Level Weight Assignments](#). For more information on the use of the fourth level and stable sorts, see [Section 3.8](#), [Stability](#).

Implementations can also add more customizable levels, as discussed in [Section 2](#), [Conformance](#). For example, an implementation might want to handle the standard Unicode Collation, but also be capable of emulating an EBCDIC multi-level ordering (having a fourth-level EBCDIC binary order).

Review Note: All digits map to collation elements with the default secondary weight (0020 in the DUCET, "<BASE>" in the ISO 14651 CTT) rather than script-specific secondary weights for non-ASCII digits. Such secondary distinctions between sets of digits appear unnecessary: They only distinguish strings that have no primary differences, but have digits of the same value and different scripts.

A possible alternative is to retain a secondary distinction between ASCII digits and all other digits. ASCII digits would map to collation elements with the default secondary weight, and all other digits would map to collation elements with a higher secondary weight, using the same secondary weight for all non-ASCII digits.

Appendix A: Deterministic Sorting

There is often a good deal of confusion about what is meant by the terms "stable" or "deterministic" when applied to sorting or comparison. This confusion in terms often leads people to make mistakes in their software architecture, or make choices of language-sensitive comparison options that have significant impact in terms of performance and footprint, and yet do not give the results that users expect.

A.1 Stable Sort

A stable sort is one where two records will retain their order when sorted according to a particular field, even when the two fields have the same contents. Thus those two records come out in the same relative order that they were in before sorting, although

their positions relative to other records may change. Importantly, this is a property of the sort algorithm, *not* the comparison mechanism.

Two examples of differing sort algorithms are Quicksort and Merge sort. Quicksort is not stable while Merge sort is stable. (A Bubble sort, as typically implemented, is also stable.)

- For background on the names and characteristics of different sorting methods, see [[SortAlg](#)]
- For a definition of stable sorting, see [[Unstable](#)]

Assume the following records:

Original Records

Record	Last_Name	First_Name
1	Davis	John
2	Davis	Mark
3	Curtner	Fred

The results of a Merge sort on the Last_Name field only are:

Merge Sort Results

Record	Last_Name	First_Name
3	Curtner	Fred
1	Davis	John
2	Davis	Mark

The results of a Quicksort on the Last_Name field only are:

Quicksort Results

Record	Last_Name	First_Name
3	Curtner	Fred
2	Davis	Mark
1	Davis	John

As is apparent, the Quicksort algorithm is not stable; records 1 and 2 are not in the same order they were in before sorting.

A stable sort is often desirable—for one thing, it allows records to be successively

sorted according to different fields, and to retain the correct lexicographic order. Thus, with a stable sort, one could sort all the records by First_Name, and then sort them again by Last_Name, giving the desired results: that all records would be ordered by Last_Name, and in the case where the Last_Name values are the same, be further subordered by First_Name.

A.2 Deterministic Sort

A *deterministic* sort is a very different beast. This is a sort algorithm that returns the same results each time. On the face of it, it would seem odd for any sort algorithm to *not* be deterministic, but there are examples of real-world sort algorithms that are not.

The key concept is that these sort algorithms *are* deterministic when two records have unequal fields, but they may return different results at different times when two records have equal fields.

For example, a classic Quicksort algorithm works recursively on ranges of records. For any given range of records, it takes the first element as the *pivot element*. However, that algorithm performs badly with input data that happens to be already sorted (or mostly sorted). A randomized Quicksort, which picks a random element as the pivot, can on average be faster. Because of this random selection, different outputs can result from *exactly* the same input: the algorithm is not deterministic.

Enhanced Quicksort Results

Record	Last_Name	First_Name
3	Curtner	Fred
2	Davis	John
1	Davis	Mark

or

Record	Last_Name	First_Name
3	Curtner	Fred
1	Davis	Mark
2	Davis	John

As another example, multiprocessor sort algorithms can be non-deterministic. The work of sorting different blocks of data is farmed out to different processors and then merged back together. The ordering of records with equal fields might be different according to when different processors finish different tasks.

Note that a deterministic sort is weaker than a stable sort. A stable sort is always deterministic, but not vice versa. Typically, when people say they want a deterministic sort, they really mean that they want a stable sort.

A.3 Deterministic Comparison

A *deterministic comparison* is different than either a stable sort or a deterministic sort; it is a property of a comparison function, not a sort algorithm. This is a comparison where strings that do not have identical binary contents (optionally, after some process of normalization) will compare as unequal. A deterministic comparison is sometimes called a *stable* (or *semi-stable*) *comparison*.

There are many people who confuse a deterministic comparison with a deterministic (or stable) sort, but this ignores the fundamental difference between a comparison and a sort. A comparison is used by a sort algorithm to determine the relative ordering of two fields, such as strings. Using a deterministic comparison cannot cause a sort to be deterministic, nor to be stable. Whether a sort is deterministic or stable is a property of the sort algorithm, not the comparison function, as the prior examples show.

A.3.1 Forcing Deterministic Comparisons

One can produce a deterministic comparison function from a non-deterministic one, in the following way (in pseudo-code):

```
int new_compare (String a, String b) {
    int result = old_compare(a, b);
    if (result == 0) {
        result = binary_compare(a, b);
    }
    return result;
}
```

Programs typically also provide the facility to generate a *sort key*, which is a sequences of bytes generated from a string in alignment with a comparison function. Two sort keys will binary-compare in the same order as their original strings. The simplest means to create a deterministic sort key that aligns with the above `new_compare` is to append a copy of the original string to the sort key. This will force the comparison to be deterministic.

```
byteSequence new_sort_key (String a) {
    return old_sort_key(a) + SEPARATOR + toByteSequence(a);
}
```

Because sort keys and comparisons must be aligned, a sort key generator is deterministic if and only if a comparison is.

A.3.2 Best Practice

A deterministic comparison is generally not best practice. First, it has a certain performance cost in comparison, and a quite substantial impact on sort key size. (For example, ICU language-sensitive sort keys are generally about the size of the original string, so appending a copy of the original string generally doubles the size of the sort key.) A database using these sort keys can have substantially increased disk footprint and memory footprint, and consequently reduced performance.

More importantly, a deterministic comparison function does not actually achieve the effect people think it will have. Look at the sorted examples above. Whether a deterministic comparison is used or not, there will be *no* effect on the Quicksort example, because the two records in question have identical `Last_Name` fields. It does not make a non-deterministic sort into a deterministic one, nor does it make a non-stable sort into a stable one.

Thirdly, a deterministic comparison is often not what is wanted, when people look closely at the implications. Look at the example again, and suppose that this time the

user is sorting first by last name, then by first name.

Original Records

Record	Last_Name	First_Name
1	Davis	John
2	Davis	Mark
3	Curtner	Fred

The desired results are the following, which should result whether the sort algorithm is stable or not, because it uses both fields.

First then Last

Record	Last_Name	First_Name
3	Curtner	Fred
1	Davis	John
2	Davis	Mark

Now suppose that in record 2, the source for the data caused the last name to contain a format control character, such as a ZWJ (used to request ligatures on display). In this case there is no visible distinction in the forms, because the font does not have any ligatures for these sequences of Latin letters. The default UCA collation weighting causes the ZWJ to be—correctly—ignored in comparison, since it should only affect rendering. However, if that comparison is changed to be deterministic (by appending the binary values for the original string), then unexpected results will occur.

First then Last (Deterministic)

Record	Last_Name	First_Name
3	Curtner	Fred
2	Davis	Mark
1	Davis	John

A.3.3 Forcing Stable Sorts

Typically, what people really want when they say they want a deterministic comparison is actually a stable sort.

One can force a non-stable sort algorithm to produce stable results by how one does the comparison. However, this has literally nothing to do with making the comparison deterministic or not. Forcing stable results can be done by appending the *current record*

number to the strings to be compared. (The implementation may not actually append the number; it may use some other mechanism, but the effect would be the same.)

If such a modified comparison is used, for example, it forces Quicksort to get the same results as a Merge sort. In that case, the irrelevant character ZWJ does not affect the outcome. The correct results occur, as illustrated below.

First then Last (Forced Stable Results)

Record	Last_Name	First_Name
3	Curtner	Fred
1	Davis	John
2	Davis	Mark

If anything, this then is what users want when they say they want a deterministic comparison. See also *Section 1.6, [Merging Sort Keys](#)*.

A.4 Stable and Portable Comparison

There are a few other terms worth mentioning, simply because they are also subject to considerable confusion. Any or all of the following terms may be easily confused with the discussion above.

A *stable comparison* is one that does not change over successive software versions. That is, as one uses successive versions of an API, with the same "settings" (such as locale), one gets the same results.

A *stable sort key generator* is one that generates the same binary sequence over successive software versions.

Warning: If the sort key generator is stable, then the associated comparison will perform the same. However, the reverse is not guaranteed. To take a trivial example, suppose the new version of the software always adds an 0xFF byte at the front of every sort key. The results of any comparison of any two new keys would be identical to the results of the comparison of any two corresponding old keys. However, the bytes have changed, and the comparison of old and new keys would give different results. Thus one can have a stable comparison, yet an associated non-stable sort key generator.

A *portable comparison* is where corresponding APIs for comparison produce the same results across different platforms. That is, if one uses the same "settings" (such as locale), one gets the same results.

A *portable sort key generator* is where corresponding sort key APIs produce exactly the same sequence of bytes across different platforms.

Warning: As above, a comparison may be portable without the associated sort

key generator being portable.

Ideally, all products would have the same string comparison and sort key generation for, say Swedish, and thus be portable. For historical reasons, this is not the case. Even if the main letters sort the same, there will be differences in the handling of other letters, or of symbols, punctuation, and other characters. There are some libraries that offer portable comparison, such as [\[ICUCollator\]](#), but in general the results of comparison or sort key generation may vary significantly between different platforms.

In a closed system, or in simple scenarios, portability may not matter. Where someone has a given set of data to present to a user, and just wants the output to be reasonably appropriate for Swedish, the exact order on the screen may not matter.

In other circumstances, differences can lead to data corruption. For example, suppose that two implementations do a database SELECT for records between a pair of strings. If the collation is different in the least way, they can get different data results. Financial data might be different, for example, if a city is included in one SELECT on one platform and excluded from the same SELECT on another platform.

Appendix B: Synchronization with ISO/IEC 14651

The Unicode Collation Algorithm is maintained in synchronization with the International Standard, ISO/IEC 14651 [\[ISO14651\]](#). Although the presentation and text of the two standards are rather distinct, the approach toward the architecture of multi-level collation weighting and string comparison is closely aligned. In particular, the synchronization between the two standards is built around the data tables which define the default (or tailorable) weights. The UCA adds many additional specifications, implementation guidelines, and test cases, over and above the synchronized weight tables. This relationship between the two standards is similar to that maintained between the Unicode Standard and ISO/IEC 10646.

For each version of the UCA, the Default Unicode Collation Element Table (DUCET) [\[Allkeys\]](#) is constructed based on the repertoire of the corresponding version of the Unicode Standard. The synchronized version of ISO/IEC 14651 has a Common Tailorable Template (CTT) table built for the same repertoire and ordering. The two tables are constructed with a common tool, to guarantee identical default (or tailorable) weight assignments. The CTT table for ISO/IEC 14651 is constructed using only symbols, rather than explicit integral weights, and with the Shift-Trimmed option for variable weighting.

The detailed synchronization points between versions of UCA and published editions (or amendments) of ISO/IEC 14651 are shown in [Table 19](#).

Table 19. UCA and ISO/IEC 14651

UCA Version	UTS #10 Date	DUCET File Date	ISO/IEC 14651 Reference
6.3.0	TBD	TBD	---
6.2.0	2012-08-30	2012-08-14	---

6.1.0	2012-02-01	2011-12-06	14561:2011 Amd 1
6.0.0	2010-10-08	2010-08-26	14561:2011 (3rd ed.)
5.2.0	2009-10-08	2009-09-22	---
5.1.0	2008-03-28	2008-03-04	14561:2007 Amd 1
5.0.0	2006-07-10	2006-07-14	14561:2007 (2nd ed.)
4.1.0	2005-05-05	2005-05-02	14561:2001 Amd 3
4.0.0	2004-01-08	2003-11-01	14561:2001 Amd 2
9.0 (= 3.1.1)	2002-07-16	2002-07-17	14561:2001 Amd 1
8.0 (= 3.0.1)	2001-03-23	2001-03-29	14561:2001
6.0 (= 2.1.9)	2000-08-31	2000-04-18	---
5.0 (= 2.1.9)	1999-11-22	2000-04-18	---

Acknowledgements

Mark Davis authored most of the original text of this document. Mark Davis and Ken Whistler together have added to and continue to maintain the text.

Thanks to Bernard Desgraupes, Richard Gillam, Kent Karlsson, York Karsunke, Michael Kay, Åke Persson, Roozbeh Pournader, Markus Scherer, Javier Sola, Otto Stolz, Ienup Sung, Yoshito Umaoka, Andrea Vine, Vladimir Weinstein, Sergiusz Wolicki, and Richard Wordingham for their feedback on previous versions of this document, to Jianping Yang and Claire Ho for their contributions on matching, and to Cathy Wissink for her many contributions to the text. Julie Allen helped in copyediting of the text.

References

- [Allkeys] Default Unicode Collation Element Table (DUCET)
For the latest version, see:
<http://www.unicode.org/Public/UCA/latest/allkeys.txt>
For the 6.3.0 version, see:
<http://www.unicode.org/Public/UCA/6.3.0/allkeys.txt>
- [CanStd] CAN/CSA Z243.4.1. For availability see
<http://www.shopcsa.ca/>
- [CLDR] Common Locale Data Repository
<http://unicode.org/cldr/>
- [Data10] For all UCA implementation and test data

For the latest version, see:

<http://www.unicode.org/Public/UCA/latest/>

For the 6.3.0 version, see:

<http://www.unicode.org/Public/UCA/6.3.0/>

For ftp access, see:

<ftp://www.unicode.org/Public/UCA/>

- [FAQ] Unicode Frequently Asked Questions
<http://www.unicode.org/faq/>
For answers to common questions on technical issues.
- [Feedback] Reporting Errors and Requesting Information Online
<http://www.unicode.org/reporting.html>
- [Glossary] Unicode Glossary
<http://www.unicode.org/glossary/>
For explanations of terminology used in this and other documents.
- [ICUCollator] ICU User Guide: Collation Introduction
<http://userguide.icu-project.org/collation>
- [ISO14651] International Organization for Standardization. Information Technology—International String ordering and comparison —Method for comparing character strings and description of the common template tailorable ordering. (ISO/IEC 14651:2011). For availability see <http://www.iso.org>
- [JavaCollator] <http://docs.oracle.com/javase/6/docs/api/java/text/Collator.html>,
<http://docs.oracle.com/javase/6/docs/api/java/text/RuleBasedCollator.html>
- [Reports] Unicode Technical Reports
<http://www.unicode.org/reports/>
For information on the status and development process for technical reports, and for a list of technical reports.
- [Sample] <http://www.unicode.org/reports/tr10/Sample/>

- [SortAlg] For background on the names and characteristics of different sorting methods, see http://en.wikipedia.org/wiki/Sorting_algorithm
- [Tests10] Conformance Test and Documentation
For the latest version, see:
<http://www.unicode.org/Public/UCA/latest/CollationTest.html>
<http://www.unicode.org/Public/UCA/latest/CollationTest.zip>
For the 6.3.0 version, see:
<http://www.unicode.org/Public/UCA/6.3.0/CollationTest.html>
<http://www.unicode.org/Public/UCA/6.3.0/CollationTest.zip>
- [UAX15] UAX #15: Unicode Normalization Forms
<http://www.unicode.org/reports/tr15/>
- [UAX29] UAX #29: Unicode Text Segmentation
<http://www.unicode.org/reports/tr29/>
- [UAX44] UAX #44: Unicode Character Database
<http://www.unicode.org/reports/tr44/>
- [Unicode] The Unicode Consortium. The Unicode Standard, Version 6.3.0 (Mountain View, CA: The Unicode Consortium, 2013. ISBN 978-1-936213-XX-X)
<http://www.unicode.org/versions/Unicode6.3.0/>
- [Unstable] For a definition of stable sorting, see <http://planetmath.org/encyclopedia/UnstableSortingAlgorithm.html>
- [UTN5] UTN #5: Canonical Equivalence in Applications
<http://www.unicode.org/notes/tn5/>
- [UTS18] UTS #18: Unicode Regular Expressions
<http://www.unicode.org/reports/tr18/>
- [UTS35] UTS #35: Unicode Locale Data Markup Language (LDML)
<http://www.unicode.org/reports/tr35/>
- [Versions] Versions of the Unicode Standard

<http://www.unicode.org/versions/>

For details on the precise contents of each version of the Unicode Standard, and how to cite them.

Migration Issues

This section summarizes important migration issues which may impact implementations of the Unicode Collation Algorithm when they are updated to a new version.

UCA 6.3.0 from UCA 6.2.0 (or earlier)

- The [IgnoreSP](#) option for variable weighted characters has been removed. Implementers of this option may instead refer to CLDR Shifted behavior.
- U+FFFD is mapped to a collation element with a very high primary weight. This changes the behavior of ill-formed code unit sequences, if they are weighted as if they were U+FFFD. When using the Shifted option, ill-formed code unit are no longer ignored.
- Fourth-level weights have been removed from the DUCET. Parsers of allkeys.txt may need to be modified. If an implementation relies on the fourth-level weights, then they can be computed according to the derivation described in UCA version 6.2.
- CLDR root collation data files have been moved from the UCA data directory (where they were combined into a CollationAuxiliary.zip) to the CLDR repository. See *UTS #35: LDML: Collation, Section 2.1, [Root Collation Data Files](#)*.

UCA 6.2.0 from UCA 6.1.0 (or earlier)

- There are a number of clarifications to the text that people should revisit, to make sure that their understanding is correct. These are listed in the modifications section.
- Users of the conformance test data files need to adjust their test code. For details see the CollationTest.html documentation file.

UCA 6.1.0 from UCA 6.0.0 (or earlier)

- A new [IgnoreSP](#) option for variable weighted characters has been added. Implementations may need to be updated to support this additional option.
- Another option for parametric tailoring, reorder, has been added. Although parametric tailoring is not a required feature of UCA, it is used by [\[UTS35\]](#), and implementers should be aware of its implications.

UCA 6.0.0 from UCA 5.2.0 (or earlier)

- Ill-formed code unit sequences are no longer required to be mapped to [0000.0000.0000] when not treated as an error; instead, implementations are strongly encouraged not to give them ignorable primary weights, for security

reasons.

- Noncharacter code points are also no longer required to be mapped to [0000.0000.0000], but are given implicit weights instead.
- The addition of a new range of CJK unified ideographs (Extension D) means that some implementations may need to change hard-coded ranges for ideographs.

UCA 5.2.0 from UCA 5.1.0 (or earlier)

- The clarification of implicit weight BASE values in *Section 7.1.3, [Implicit Weights](#)* means that any implementation which weighted unassigned code points in a CJK unified ideograph block as if they were CJK unified ideographs will need to change.
- The addition of a new range of CJK unified ideographs (Extension C) means that some implementations may need to change hard-coded ranges for ideographs.

Modifications

The following summarizes modifications from the previous revisions of this document.

Revision 27 [MS, KW]

- **Proposed Update** for Unicode 6.3.0.
- [Changes in draft 5](#)
- *Section 5, [Tailoring](#)*: Clarified the definition of "Tailoring".
- [Changes in draft 4](#)
- Reordered some sections for better flow.
[Before final publication, this should probably go to the end of the revision-27 modifications section because other modification notes refer to the old section numbers — or else we renumber the sections in the following revision-27 modification notes](#)
 - *Section 3.6, [Default Unicode Collation Element Table](#)* became section 3.8.
 - *Section 3.6.1, [File Format](#)* became section 9.1.
 - *Section 3.6.2, [Variable Weighting](#)* became section 3.6.
 - *Section 3.6.3, [Default Values](#)* became section 3.8.1.
 - *Section 3.6.4, [Well-Formedness of the DUCET](#)* became section 3.8.2.
 - *Section 3.8, [Stability](#)* was removed after moving its subsections.
 - The text of *Section 3.8.1, [Stable Sort](#)* and *Section 3.8.2, [Deterministic Comparison](#)* was moved into *Section 1.8, [What Collation is Not](#)* under "Collation order is not a stable sort".
- [Changes in draft 3](#)
- Changed the wording about where backwards-secondary ordering is used. This practice is associated with major French dictionary ordering traditions, rather than with Canadian locales.
- [Changes in drafts 1 & 2](#)
- *Section 3.6.2, [Variable Weighting](#)*: Removed option [IgnoreSP](#).
- *Section 7.1.4, [Trailing Weights](#)*: Weights FFFD..FFFF are reserved for special

collation elements. U+FFFD is mapped to a collation element with a very high primary weight (0xFFFD).

- *Section 7.2, [Tertiary Weight Table](#)*: Trailing collation elements use regular tertiary weights rather than MAX = 1F. The MAX tertiary weight is not used any more in the DUCET.
- Removed *Section 7.3, Fourth-Level Weight Assignments*: Fourth-level weights have been removed from the DUCET. They were intended for an approximation of a deterministic comparison, but this approximation was not very good, the UCA did not use this fourth level of data, and this data was not related to the fourth level introduced by variable handling and thus led to confusion.
- In *Section 9, [Data Files](#)*, added a brief description of decomp.txt.
- CLDR root collation data files have been moved from the UCA data directory (where they were combined into a CollationAuxiliary.zip) to the CLDR repository. See *UTS #35: LDML: Collation, Section 2.1, [Root Collation Data Files](#)*.

Revision 26 [MD, KW, MS]

- **Reissued** for Unicode 6.2.0.
- Used "identical level" consistently.
- Changed *Section 1.6, Interleaved Levels* to [Merging Sort Keys](#), to avoid collision with other uses of 'interleaving'.
- *Section 3.1, [Weight Levels and Notation](#)*: Added definitions of primary, secondary, tertiary, quaternary collation elements, for clarity.
- *Section 3.3.2, [Contractions](#)*: Clarified which characters prevent contractions.
- *Section 3.6, [Default Unicode Collation Element Table](#)*: Description of differences between DUCET and CLDR root collation moved out of this document and merged with existing text in the CollationAuxiliary.html documentation file.
- *Section 3.6.1, [File Format](#)* documentation bug fixes.
- *Section 3.6.2, [Variable Weighting](#)*: Added text and rearrangements for clarity.
- Added *Section 3.6.4, [Well-Formedness of the DUCET](#)* about the DUCET not being entirely well-formed, including the contractions that would need to be added.
- *Section 3.7, [Well-Formed Collation Element Tables](#)*: Narrowed and clarified [well-formedness condition 2](#). Added new [well-formedness condition 5](#) on contractions.
- *Section 4.5, [Well-Formedness Examples](#)*: Created section with existing example, added second example.
- *Section 5.1, [Parametric Tailoring](#)*: Removed Table 14, incorporating material into other sections and/or LDML. Renumbered tables 15-20 to 14-19.
- Moved and merged *Section 6.5.2, Compatibility Decompositions* into *Section 6.3.3, [Leveraging Unicode Tables](#)*.
- *Section 6.9, [Handling Collation Graphemes](#)*: Added algorithm steps 4.1 and 4.2 for handling discontinuous contractions.
- *Section 6.10.2, [Sample Code](#)*: Corrected bitmasks and rewrote the implementation of searchContractions().
- Narrowed backward accents to Canadian French as the one known locale

requiring this option.

- CollationAuxiliary.html: Added a description of the implicit weight generation (CJK and Unassigned characters), a description of the context syntax, and a note about additional Tibetan contractions.
- CollationTest.html: The conformance test data now uses the standard tie-breaker ([S3.10](#)).
- Many minor clarifications and wording changes.

Revision 25 being a proposed update, only changes between revisions 26 and 24 are noted here.

Revision 24 [MD, PE, KW]

- **Reissued** for Unicode 6.1.0.
- Described the new reorder parameter in *Table 14* (by reference to [UTS35](#)).
- Corrected duplicate anchor for "Stable Sort".
- Updated text in *Section 3.8*, [Stability](#) regarding "semi-stable collation" to use term "deterministic comparison" for consistency with Appendix A.
- Moved position of *Table 12* in [Variable Weighting](#) for better text flow and presentation.
- Added listing of migration issues for this version.
- Added subheads to *Section 3.8*, [Stability](#) and reference links to the UCA change management policy pages.
- Documented use of U+FFFF and U+FFFE in CLDR, in [Table 11](#).
- Added additional FFFF example for clarity, to [Table 12](#).
- Added examples of symbols to [Table 13](#).
- Documented the new zipped files and .html files better in [Data Files](#).
- Updated references list.
- Moved definitions of Simple, Expansion, and Contraction ahead of their first use in *Section 3.2*, [Simple Mappings](#).
- Consolidated discussion of derived weights for Hangul syllables into *Section 7.1.5*, [Hangul Collation](#) and did an extensive rewrite of that section.
- Added new *Section 7.3*, [Fourth-Level Weight Assignments](#).
- Added subheads for Appendix A to table of contents.
- Added new *Appendix B*, [Synchronization with ISO/IEC 14651](#).
- Described major revision to the ordering of variable characters into groups, separating punctuation and symbols.
- Added option [IgnoreSP](#).
- Fixed statement about soft hyphen.
- Fixed section on contiguous weights
- Fixed section on finding collation grapheme clusters.
- Added new *Section 8.2*, [Asymmetric Search](#).

Revision 23 being a proposed update, only changes between revisions 24 and 22 are noted here.

Revision 22 [KW]

- **Reissued** for Unicode 6.0.0.
- Updated text of Summary at top of document.
- Added Migration Issues section after References.
- Reorganized and renumbered several sections for better text flow.
- Provided numbers and anchors for tables, and updated table and caption formats to match current Technical Report style. Added captions for tables or figures that did not have them. Removed unneeded color backgrounds from tables.
- Updated several obsolete links in the References section.
- Reorganized the References section and updated style of references.
- Added *Section 9 [Data Files](#)*.
- Significant editorial corrections throughout.
- Completely rewrote the discussion of "illegal" and "legal" code points to bring it up to date with the Unicode Standard. See *Section 7.1.1 [Handling Ill-Formed Code Unit Sequences](#)*.
- Split *Section 7.1.5 [Hangul Collation](#)* from the discussion of trailing weights.
- Corrected order of first names in Sequential column of the *[Interleaved Levels Table](#)* and added explanation of the option used for variable collation elements in the table.
- Updated the *[Tailoring Example](#)* to use the ICU syntax instead of Java. [MD]

Revision 21 being a proposed update, only changes between revisions 22 and 20 are noted here.

Revision 20

- **Reissued** for Unicode 5.2.0.
- In *Section 7.1.3 [Implicit Weights](#)*, clarified the calculation of implicit weights.
- Made it clear that the BASE value does not include unassigned code points.
- Clarified why some sample cells are empty in the first table.
- General: updated references to UAX/UTS's
- Removed reference to UTR #30
- Better aligned the options with the 3 values for variableChoice.
- Clarified the computation of the fourth level in *Section 3.2.1, [File Format](#)*. [KW]
- Changed bit layout in *Section 6.10.1 [Collation Element Format](#)* for a real collation element, to account for the fact that the DUCET secondary values number more than 255, so no longer fit in 8 bits. [KW]
- Made small editorial clarifications regarding variable weighting in *Section 3.2.2, [Variable Weighting](#)*. [KW]
- Updated reference to SC22 WG20 to SC2 OWG-SORT in *Section 7.1.4.1*. [KW]
- Made a minor wording clarification in *Section 7.3 [Compatibility Decompositions](#)*. [KW]
- Small editorial updates through for formatting consistency. [KW]
- Updated Modifications section to current conventions for handling proposed

update drafts. [KW]

Revision 19 being a proposed update, only changes between revisions 20 and 18 are noted here.

Revision 18

- **Reissued** for Unicode 5.1.0.
- Disallowed skipping 2.1.1 through 2.1.3 (*Section 4.2, [Produce Array](#)*).
- Clarified use of contractions in the DUCET in *Section 3.2, [Default Unicode Collation Element Table](#)* and *Section 3.1.1.2, [Contractions](#)*.
- Added information about the use of parameterization (*Section 5.1, [Parametric Tailoring](#)*) and a new conformance clause [C6](#).
- In *Section 8, [Searching and Matching](#)*, added new introduction and explained special cases; clarified language in definitions.
- Added *Section 8.1, [Collation Folding](#)*.
- Fixed a number of reported typos.

Revision 17 being a proposed update, only changes between revisions 18 and 16 are noted here.

Revision 16

- **Reissued** for Unicode 5.0.0.
- Replaced "combining mark" by "non-starter" where necessary.
- Updated reference to Unicode 5.0 with the ISBN number.
- Added UTN#9 text in informative appendix as [Appendix A: Deterministic Sorting](#).

Revision 15 being a proposed update, only changes between revisions 16 and 14 are noted here.

Revision 14

- **Reissued** for Unicode 4.1.0.
- Expanded use of 0x1D in *Section 7.3.1, [Tertiary Weight Table](#)*.
- Removed DS5, added DS1a, DS2a, explanations of interactions with other conditions, such as Whole Word or Whole Grapheme.
- Added conformance clause C5 for searching and matching.
- Many minor edits.
- Removed S1.3, so that fully ignorable characters will interrupt contractions (that do not explicitly contain them).
- Added related *Section 3.1.6, [Combining Grapheme Joiner](#)*.
- Removed S1.2 for Thai, and a paragraph in 1.3.
- Added more detail about Hangul to *Section 7.1.4, [Trailing Weights](#)*, including a description of the Interleaving method.
- Fixed dangling reference to base standard in [C4](#).
- Added definitions and clarifications to *Section 8, [Searching and Matching](#)*.

- Added more information on user expectations to *Section 1, [Introduction](#)*.

Data tables for 4.1.0 contain the following changes:

1. The additions of weights for all the new Unicode 4.1.0 characters.
2. The change of weights for characters Æ, Æ̇, Æ̈́; Đ, Đ̇; Ħ, Ħ̇, Ħ̈́; and Ø, Ø̇ (and their lowercase and accented forms) to have secondary (accent) differences from AE; D; H; L; and O, respectively. This is to provide a much better default for languages in which those characters are not tailored. See also the section on user expectations.
3. Change in weights for U+0600 ARABIC NUMBER SIGN and U+2062 INVISIBLE TIMES and like characters (U+0600..U+0603, U+06DD, U+2061..U+2063) to be not completely ignorable, because their effect on the interpretation of the text can be substantial.
4. The addition of about 150 contractions for Thai. This is synchronized with the removal of S1.2. The result produces the same results for well-formed Thai data, while substantially reducing the complexity of implementations in searching and matching. Other changes for Thai include:
 - a. After U+0E44 ᨿ THAI CHARACTER SARA AI MAIMALAI
Insertion of the character U+0E45 ᨾ THAI CHARACTER LAKKHANGYAO
 - b. Before U+0E47 ᨾ THAI CHARACTER MAITAIKHU
Insertion of the character U+0E4E ᨾ THAI CHARACTER YAMAKKAN
 - c. After U+0E4B ᨾ THAI CHARACTER MAI CHATTAWA
Insertion of the character U+0E4C ᨾ THAI CHARACTER THANTHAKHAT
Then the character U+0E4D ᨾ THAI CHARACTER NIKHAHIT
5. Changed the ordering of U+03FA GREEK CAPITAL LETTER SAN and U+03FB GREEK SMALL LETTER SAN.

Revisions 12 and 13 being proposed updates, only changes between revisions 14 and 11 are noted here.

Revision 11

- Changed the version to synchronize with versions of the Unicode Standard, so that the repertoire of characters is the same. This affects the header and [C4](#). This revision is synchronized with Unicode 4.0.0.
- Location of data files changed to <http://www.unicode.org/Public/UCA/>
- Added new [Introduction](#). This covers concepts in *Section 5.17, "Sorting and Searching"*, in *The Unicode Standard, Version 3.0*, but is completely reworked. The Scope section has been recast and is now at the end of the introduction.
- In *Section 6.9, [Tailoring Example: Java](#)*, added informative reference to LDML; moved informative reference to ICU.
- Added explanation of different ways that the Hangul problem can be solved in *Section 7.1.4, [Trailing Weights](#)*.
- Copied sentence from Scope up to Summary, for more visibility.

Revision 10 being a proposed update, only changes between revisions 11 and 9 are

noted here.

Revision 9

- Added [C4](#).
- Added more conditions in *Section 3.3*, [Well-Formed Collation Element Tables](#).
- Added S1.3.
- Added treatment of ignorables after variables in *Section 3.2.2*, [Variable Weighting](#).
- Added *Section 3.4*, [Stability](#).
- Modified and reorganized *Section 7*, [Weight Derivation](#). In particular, CJK characters and unassigned characters are given different weights. Added MAX to *Section 7.3*.
- Added references.
- Minor editing.
- Clarified noncharacter code points in *Section 7.1.1*, *Illegal code points*.
- Modified S1.2 and *Section 3.1.3*, [Rearrangement](#) to use the Logical_Order_Exception property, and removed *rearrange* from the file syntax in *Section 3.2.1*, [File Format](#), and from *Section 5*, [Tailoring](#).
- Incorporated Cathy Wissink's notes on linguistic applicability.
- Updated links for [\[Test\]](#).

Copyright © 1998-2012 Unicode, Inc. All Rights Reserved. The Unicode Consortium makes no expressed or implied warranty of any kind, and assumes no liability for errors or omissions. No liability is assumed for incidental and consequential damages in connection with or arising out of the use of the information or programs contained or accompanying this technical report. The Unicode [Terms of Use](#) apply.

Unicode and the Unicode logo are trademarks of Unicode, Inc., and are registered in some jurisdictions.