

Graphite Binary Format

Extending TrueType for Graphite

Version 1.0

*Martin Hosken and Sharon Correll,
SIL Non-Roman Script Initiative (NRSI)*

Introduction

The Graphite binary format is structured in order that a Graphite binary description may be incorporated into a TrueType font. Thus the binary format uses the TrueType table structure, identically to how it is used in a TrueType font. The only difference between using an external file containing Graphite binary information in tables, and inserting the binary information into tables in the font is that tables are considered local to their file and are considered to override those found in the font file. This allows there to be multiple, independent descriptions held in separate files. Those independent descriptions would have to be merged, in a way described in this document, if they were to be held together in the same font file or binary file.

The description consists of a set of table descriptions. The format of a file follows that of a TrueType font containing only those tables pertinent to the description (i.e. for a separate binary description, those tables listed here).

Version

This file describes version 1.0 of the Graphite tables.

Tables

This description proposes that addition of three table types to TrueType. One, “`silf`” is unique to the needs of Graphite, whilst the other two, “`gloc`” and “`glat`” provide an extended glyph attribute mechanism. (If necessary the tables could be restructured to be stored inside the single “`silf`” table.) In addition, use is made of the “`feat`” and “`name`” table types.

Glat

The `Glat` table type is used for storing glyph attributes. Each glyph may be considered to have a sparse array of, at the most, 256 16-bit signed attributes. The `Glat` table is the mechanism by which they are stored.

The `Glat` table consists of a table header and an array of `Glat_entry` items:

Type	Name	Description
FIXED	Version	Table version: 00010000
<code>Glat_entry[]</code>	Entries	Glyph attribute entries

Table 1: `Glat`

The glyph attributes associated with a particular glyph are identified by number and value. To conserve space, this storage is run-length encoded. Thus a glyph will have a series of `Glat_entries` corresponding to each non-contiguous set of attributes. The structure of a `Glat_entry` is:

Type	Name	Description
BYTE	<code>attNum</code>	Attribute number of first attribute
BYTE	<code>num</code>	Number of attributes in this run
SHORT	<code>attributes[]</code>	Array of <code>num</code> attributes

Table 2: `Glat_entry`

Notice that all glyph attributes are 16-bit signed values. If a 32-bit value is required, then two attributes should be assigned and joined together by the application.

Attribute numbers are application specific.

Gloc

The `Gloc` table is used to index the `Glat` table. It is structured identically to the `loca` table type, except that it has a header.

Type	Name	Description
FIXED	Version	Table version: 00010000
USHORT	<code>flags</code>	bit 0 = 1 for Long format, = 0 for short format bit 1 = 1 for attribute names, = 0 for stripped
USHORT	<code>numAttribs</code>	Number of attributes
USHORT/ULONG	<code>locations[]</code>	Offsets into <code>Glat</code> table for each glyph; (number of glyph IDs + 1) of these
USHORT	<code>attribIds[]</code>	Debug id for each attribute

Table 3: `Gloc`

The `flags` entry contains a bit to indicate whether the `locations` array is of type `USHORT` or `ULONG`. The `locations` array is identically structured to that of the `loca` table. There is one entry per glyph and an extra entry to identify the length of the final glyph's attribute entries. Offsets are given to a `Glat_entry` in the `Glat` table. The second bit indicates whether there is an `attribIds` array at the end of this table. If there is, then it contains name IDs for each attribute. If this bit is not set, then there is no array and the table ends after the `locations` array.

feat

The GX `feat` table will be used for storing features. This makes reference to the name table which is use for storing feature names and feature value names.

Silf

The “`silf`” table will be used for storing rules and actions for the various types of tables in a rendering description. The structure of the `silf` table is:

Type	Name	Description
FIXED	Version	Table version: 00010000
USHORT	Num	Number of SIL subtables
ULONG	Offset[]	Array of Num offsets to the subtables relative to the start of this table
SIL_Sub	Tables[]	Array of independent rendering description subtables

Table 4: Silf

Since one TrueType file may hold multiple independent rendering descriptions, each rendering description is described in a subtable. The subtable contains all that is necessary to describe the rendering of one set of writing systems.

Type	Name	Description
USHORT	maxGlyphID	Maximum valid glyph ID (including line-break & pseudo-glyphs)
SHORT	extraAscent	Em-units to be added to the font's ascent
SHORT	extraDescent	Em-units to be added to the font's descent
BYTE	numPasses	Number of rendering description passes
BYTE	iSubst	Index of first substitution pass
BYTE	iPos	Index of first Positioning pass
BYTE	iJust	Index of first Justification pass
BYTE	iBidi	Index of first pass after the bidi pass(must be <= iPos); 0xFF implies no bidi pass
BYTE	lineBreakFlag	True (1) if there is any start-, end-, or cross-line contextualization; false (0) if line boundaries are irrelevant
BYTE	maxPreContext	Max range for preceding cross-line-boundary contextualization
BYTE	maxPostContext	Max range for following cross-line-boundary contextualization
BYTE	attrPsuedo	Glyph attribute number that is used for actual glyph ID for a pseudo glyph
BYTE	attrBreakWeight	Glyph attribute number of breakweight attribute
BYTE	attrBidi	Glyph attribute number for directionality attribute
USHORT	numLigComp	Number of initial glyph attributes that represent ligature components
BYTE	numUserDefn	Number of user-defined slot attributes
BYTE	maxCompPerLig	Maximum number of components per ligature
BYTE	direction	Supported direction(s)
BYTE	reserved	

BYTE	reserved	
BYTE	reserved	
BYTE	numScriptTag	Number of scripts this subtable supports
UINT	scriptTag[]	Array of numScriptTag script tags
USHORT	lbGID	Glyph ID for line-break psuedo-glyph
ULONG	oPasses[]	Offets to passes relative to the start of this subtable; numPasses + 1 of these
USHORT	numPseudo	Number of Unicode -> pseudo-glyph mappings
USHORT	searchPseudo	(max power of 2 <= numPseudo) * sizeof(PseudoMap)
USHORT	pseudoSelector	log ₂ (max power of 2 <= numPseudo)
USHORT	pseudoShift	numPseudo - searchPseudo
PseudoMap	pMaps[]	Mappings between Unicode and pseudo-glyphs in order of Unicode
ClassMap	classes	Classes object storing replacement classes used in actions
SIL_Pass	passes[]	Array of passes

Table 5: SIL_Sub

A pseudo-glyph is a glyph which contains no font metrics (it has a GID greater than the numGlyphs entry in the maxp table) but is used in the rendering process. Each pseudo-glyph has an attribute which is the glyph ID of a real glyph which will be used to actually render the glyph. The pseudo-glyph map contains a mapping between Unicode and pseudo-glyph number:

Type	Name	Description
USHORT	uniCode	Unicode codepoint
USHORT	nPseudo	Glyph ID of pseudo-glyph

Table 6: PseudoMap

The ClassMap stores the replacement class information for the passes in this description. Replacement classes are used during substitution where a glyph id is looked up in one class and the glyph ID at the corresponding index in another class is substituted. The difficulty with the storage of such classes is in looking up a glyph ID in an arbitrarily ordered list. One approach is to use a linear search; this is very slow, but is stored very simply. Another approach is to order the glyphs in the class and to store the index against the glyph. Both approaches are supported in the ClassMap table structure:

Type	Name	Description
USHORT	numClass	Number of replacement classes
USHORT	numLinear	Number of linearly stored replacement classes
USHORT	oClass[]	Array of numClass + 1 offsets to class arrays from the beginning of the class map
USHORT	glyphs[]	Glyphs for linear classes
LookupClass	lookups[]	An array of numClass – numLinear lookups

Table 7: ClassMap

The LookupClass stores a fast lookup association between glyph ID and index. Each lookup consists of an ordered list of glyph IDs with the corresponding index for that glyph. The number

of elements in the lookup is specified by numIDs along with a search Range and shift to initialize a fast binary search engine:

Type	Name	Description
USHORT	numIDs	Number of elements in the lookup
USHORT	searchRange	(max power of 2 \leq numRange) * 4
USHORT	entrySelector	log ₂ (max power of 2 \leq numRange)
USHORT	rangeShift	numIDs*4 – searchRange
LookupPair	lookups[]	lookups; there are numIDs of these

Table 8: LookupClass

Each element in the lookup consists of a glyphId and the corresponding index in the original ordered list.

Type	Name	Description
USHORT	glyphId	glyph id to be compared
USHORT	index	index corresponding to this glyph id in ordered list

Table 9:LookupPair

Pass

Each processing pass consists of a finite state machine description for rule finding, and the actions that are executed when a rule is matched.

Type	Name	Description
BYTE	flags	bit 0 – This pass makes no change to the slot stream
BYTE	maxRuleLoop	MaxRuleLoop for this pass
BYTE	maxRuleContext	Number of slots of input needed to run this pass
BYTE	maxBackup	Number of slots by which the following pass needs to trail this pass (ie, the maximum this pass is allowed to back up)
USHORT	numRules	Number of action code blocks
ULONG	oCode	Offset to start of constraint code from start of subtable (*constraints[0]*)
ULONG	aCode	Offset to start of action code relative to start of subtables(*actions[0]*)
ULONG	oDebug	Offset to debug arrays (*dActions[0]*); equals 0 if debug stripped
USHORT	numRows	Number of FSM states
USHORT	numTransitional	Number of transitional states in the FSM (length of *states* matrix)
USHORT	numSuccess	Number of success states in the FSM (size of *oRuleMap* array)
USHORT	numColumns	Number of FSM columns
USHORT	numRange	Number of contiguous glyph ID ranges which map to columns
USHORT	searchRange	(maximum power of 2 \leq numRange)*sizeof(Pass_Range)
USHORT	entrySelector	log ₂ (maximum power of 2 \leq numRange)
USHORT	rangeShift	numRange*sizeof(Pass_Range)-searchRange
Pass_Range	ranges[]	Ranges of glyph IDs for this FSM; *numColumns* of these
USHORT	oRuleMap[]	Maps from success state to offset into ruleMap array from start of array. First item corresponds to state # (numRows – numSuccess); ie, non-success states are

		omitted. [0xFFFF implies rule number is equal to state number (i.e. no entry in ruleMap) – NOT IMPLEMENTED]
USHORT	ruleMap[]	Array of rule numbers corresponding to an success state number
BYTE	minRulePreContext	Minimum number of items in any rule's context before the first modified rule item
BYTE	maxRulePreContext	Maximum number of items in any rule's context before the first modified rule item
SHORT	startStates[]	Array of size (maxRulePreContext – minRulePreContext + 1), indicating the start state in the state machine based on how many pre-context items a rule has
USHORT	ruleSortKeys[]	Array of *numRules* sort keys, indicating precedence of rules
BYTE	rulePreContext[]	Array of *numRules* items indicating the number of items in the context before the first modified item, one for each rule
USHORT	oConstraints[]	numRules + 1 offsets to constraint code blocks relative to *oCode* and start of subtable
USHORT	oActions[]	numRules + 1 offsets to action code blocks relative to *aCode* and start of subtable
SHORT	stateTrans[][]	array of *numTransitional* rows of *numColumns* state transitions.
BYTE	constraints[]	sequences of constraint code
BYTE	actions[]	sequences of action code
USHORT	dActions[]	name index for each action for documentation purposes. 0 = stripped ¹ . numRules of these
USHORT	dStates[]	name index for each intermediateFSM row/state for debugging. 0 = stripped. Corresponds to the last numRows – numRules
USHORT	dCols[]	name index for each state (numRows of these)

Table 10: SIL_Pass

Notice that the ranges array has fast lookup information on the front to allow for the quick identification of which range a particular glyph id is in. Each range consists of the first and last glyph id in the range.

Type	Name	Description
USHORT	firstId	First Glyph id in the range
USHORT	lastId	Last Glyph id in the range
USHORT	colId	Column index for this range

Table 11: Pass_Range

Pass Contents

A pass contains a Finite State Machine (FSM) which is used to match input strings to rules. It also contains constraints for further testing whether a matched string should fire, and it contains the action code to execute against the matched string.

The FSM consists of a set of states. A state consists of a row of transitions between that state and another state dependent upon the next glyph in the input stream. Each state may be an acceptance state, in which case it corresponds to a rule match, or a transition state, in which case the state is on the way to matching a rule, or both. A null state transition is one in which the occurrence of

¹ Should debug tables go at the end, and be marked via a flag as per Gloc?

this particular class of the following glyph, will result in no extension of a rule match anywhere, just fail on all further searching. A final state is one in which all its transitions are null transitions.

Note that the `stateTrans` array only needs to represent transitional states, not final states. Similarly, the `oRuleMap` array only needs entries for acceptance states (whether final or transitional). For this reason the FSM is set up (conceptually) in the following order: transitional non-accepting states first, followed by transitional accepting states, followed by final (accepting) states.

Note also that because there may be more than one matched rule for a given state, `oRuleMap` indicates a list of rule indices in the `ruleMap` array; `oRuleMap[i+1] - oRuleMap[i]` indicates how many there are for state `i`.

Normally the start state for an FSM is zero. But for each pass there is the idea of a “pre-context,” that is, there are slots that need to be taken into consideration in the rule-matching process that are before the current position of the input stream. If we are very near the beginning of the input, we may need to adjust by skipping some states, which corresponds to skipping the “pre-context” slots that not present due to being prior to the beginning of the input. This is what the `maxRulePreContext`, `minRulePreContext`, and `startStates` items are used for. Specifically, we need to skip the number of transitions equal to the difference between the `maxRulePreContext` and the current stream position, if greater than zero. The `startStates` array indicates what the adjusted start state should be. If the current input position is less than `minRulePreContext`, no rule will match at all.

Rules are matched in order of length, so that longest rules are given precedence over shorter rules. However, the length of some rules may have been adjusted to allow for a consistent “pre-context” for all rules, so the number of matched states in the FSM may not correspond to the actual number of matched items in the rule. For this reason, it is not adequate to simply order rules based on the number of traversed states in the FSM. Rather, rules are given sort keys indicating their precedence, which is based primarily on the length of the rule and secondarily on its original position within the source code.

The FSM engine keeps track of all the acceptance states it passes through on its path to a final state. This results in a list of rules matched by the string sorted by precedence. The engine takes the first rule index off the list and looks up the offset to some constraint code. This code is executed and if the constraint passes, then the action code associated with that offset is executed and the FSM restarts at the returned slot position. If the constraint fails, then the FSM considers the next-preferred rule, tests that constraint, and so forth. If no accepting state is found or all rules fail their constraints, then no rule applies, in which case a single glyph is put into the output stream and the current position advances by one slot.

The action strings are simply byte strings of actions, much like hinting code, but using a completely different language. (See “Stack Machine Commands.doc”.)

Sild

This table holds the debug strings for debugging purposes. Since the strings are only used for debugging, they are held somewhat optimised for space over speed and are not considered to be multilingual. Thus strings are considered to be 7-bit ASCII, with a possible extension to UTF-8 at a later stage. The table consists of a sequence of strings each preceded by a length byte. The first string is `id 0` and so on to the end of the table.

Multiple Descriptions

In the case where multiple descriptions are to be stored in the same set of tables, the following unifications need to occur:

- The feature sets must be unified, thus limiting two features with the same name to having the same settings and corresponding values.
- The glyph attributes must be unified. This can be done by using different attribute number ranges, or by examining for identical attribute mappings or for non-intersecting attribute mappings.
- The use of the name table must be unified to ensure that two features or feature settings do not refer to the same entry in the name table.

Notice that the requirement that any tables declared in an external binary description override the corresponding font table in the font, means that a name table in an external binary description must be complete, including all the strings from the original font.

Changes

0.02 22 March 99

- Added debug information
- Changed version information from USHORT to FIXED to conform with TrueType conventions

0.03 25 June 99

- Fill in details of state table
- Sort out Action classes
- Add special glyph attribute indices
- Add Sild table
- 9 Dec 99—trivial changes

0.04 18 January 2000

- Modifications based on actual implementation
- Added a LinearClass structure

0.05 11 September 2000

- Changed WinRend to Graphite

0.06 22 November 2000

- Added constants and arrays for finite state machine rework (SIL_Pass table: min/maxRulePreContext, startStates, ruleSortKeys)
- Added extraAscent and extraDescent (SIL_Sub table)

0.07 4 February 2002

- Reworked explanation of FSM behavior.
- Made a few trivial changes in the Silf table.

0.08 18 March 2003

- Fixed problem with version numbers.