

# Gfapy

Gfapy is a Python 3 library for working with GFA files. It allows to parse, validate, edit and write GFA files.

This manual explains how to access the information in GFA files using the library. It is completed by the more technical API library, which documents each class, method and constant defined by the library.

A test suite makes sure that the functionality described by this manual also works as intended. However, if this is not the case, please report any bug using the Github issues tracked (<https://github.com/ggonnella/gfapy/issues>).

## GFA specifications

The library is based on the official GFA specifications version 1 and 2, available at <https://github.com/GFA-spec/GFA-spec>. See the Versions chapter for an overview of the differences of the two versions and methods for the conversion from one version to the other.

## The Gfa object

The main class of the library is `gfapy.Gfa`. An object of this class represents the content of a GFA file. An instance can be created directly (using the `gfapy.Gfa()` constructor) or the method `gfapy.Gfa.from_file(filename)` can be used to parse a GFA file and create an instance from it.

The `str()` method converts the Gfa instance into its textual representation. Writing all information to a GFA file can be done directly using the `to_file(filename)` method.

```
g1 = gfapy.Gfa()
g1.append("H\tVN:i:1.0")
g1.append("S\ta\t*")
g1.to_file("my.gfa")
g2 = gfapy.from_file("my.gfa")
g2 == g1 # => True
str(g1) # => "H\tVN:i:1.0\nS\ta\t*"
```

## Retrieving the lines

For many line types, iterating between all lines of the type can be done using a property which is named after the record type, in plural (`segments`, `paths`, `edges`, `links`, `containments`, `groups`, `fragments`, `comments`, `custom_lines`).

The access to the header is done using a single line, which is retrieved using the `header` property.

Some lines use identifiers: segments, gaps, edges, paths and sets. Given an identifier, the line can be retrieved using the `line(id)` method. Note that identifier are represented in Gfapy by Python strings.

The list of all identifier can be retrieved using the `names` property; for the identifiers of a single line type, use `segment_names`, `edges_names`, `gap_names`, `path_names` and `set_names`. The identifiers of external sequences in fragments are not part of the same namespace and can be retrieved using the `external_names` property.

## Segments

Segment lines are available in both GFA1 and GFA2. They they represent the pieces of molecules, whose relations to other segments are coded by other line types.

In GFA1 a segment contains a segment name and a sequence (and, eventually, optional tags). In GFA2 the syntax is slightly different, as the segment contain an additional segment length field, which represent an eventually approximate length, which can be taken as a drawing indication for segments in graphical programs.

## Relationships between segments

Segments are put in relation to each other by edges lines (E lines in GFA2, L and C Lines in GFA1), as well as gaps. Gfapy allows to convert edges lines from one spefication version to the other (subject to limitations, see the Versions chapter). Gap lines cannot be converted, as no GFA1 specification exist for them.

## Relationships to external sequences

Fragments represent relationships of segments to external sequences, i.e. sequences which are not represented in the GFA file itself. The typical application is to put contigs in relationship with the reads from which they are constructed.

The set of IDs of the external sequences may overlap the IDs of the GFA file itself (ie. the namespaces are separated). The list of external IDs referenced to by fragment lines can be retrieved using the `external_names` property of the `gfapy.Gfa` instances.

To find all fragments which refer to an external ID, the `fragments_for_external(ID)` method is used. As an external sequence can refer to different segments in different F lines, the result is always an array of F lines.

Conversely, to find all fragments for a particular segment, you may use the **fragments** property of the segment instance (see the References chapter).

## Groups

Groups are lines which combine different other lines in an ordered (paths) or unordered (sets) way. Gfapy supports both GFA1 paths and GFA2 paths and sets. Paths have a different syntax in the two specification versions. Methods are provided to edit the group components also without disconnecting the line instance (see the “References” chapter).

## Other line types

The header contain metadata in a single or multiple lines. For ease of access to the header information, all its tags are summarized in a single line instance. See the “Header” chapter for more information. All lines which start by the string **#** are comments; they are handled in the “Comments” chapter. Custom lines are lines of GFA2 files which start with a non-standard record type. Gfapy provides basic built-in support for accessing the information in custom lines, and allows to define extensions for own record types for defining more advanced functionality.

## Adding new lines

New lines can be added to a Gfa instance using the **add\_line(line)** method or its alias **append(line)**. The argument may be a string describing a line with valid GFA syntax, or an instance of the class **gfapy.Line** - if a string is added, a line instance is created and then added. A line instance can be created manually before adding it, using the **gfa.Line.from\_string(s)** method.

## Editing the lines

Accessing the information stored in the fields of a line instance is described in the “Positional fields” and “Tags” chapters.

Once a line instance has been added to a **gfapy.Gfa** instance, either directly, or using its string representation, the line is said to be *connected* to the Gfa. Reading the information in fields is always allowed, while changing the content of some fields (fields which refer to other lines) is only possible for instances which are not connected.

In some cases, methods are provided to modify the content of reference fields of connected line (see the “References” chapter).

## Removing lines

Removing a line can be done using the `rm(line)` method. The argument can be a line instance or a string (in which case the line is searched using the `line(name)` method, then eliminated). A line instance can also be disconnected using the `disconnect()` method on it. Disconnecting a line may trigger other operations, such as the disconnection of other lines (see the “References” chapter).

## Renaming lines

Lines with an identifier can be renamed. This is done simply by editing the corresponding field (such as `name` or `sid` for a segment). This field is not a reference to another line and can be freely edited also in line instances connected to a `Gfa`. All references to the line from other lines will still be up to date, as they will refer to the same instance (whose name has been changed) and their string representation will use the new name.

## Validation

Different validation levels are available. They represent different compromises between speed and warrant of validity. The validation level can be specified when the `gfapy.Gfa` object is created, using the `vlevel` parameter of the constructor and of the `gfapy.Gfa.from_file()` method. Four levels of validation are defined (0 = no validation, 1 = validation by reading, 2 = validation by reading and writing, 3 = continuous validation). The default validation level value is 1.

## Manual validation

Independently from the validation level choosen, the user can always check the value of a field calling `validate_field(fieldname)` on the line instance. If no exception is raised, the field content is valid.

To check if the entire content of the line is valid, the user can call `validate` on the line instance. This will check all fields and perform cross-field validations, such as comparing the length of the sequence of a GFA1 segment, to the value of the LN tag (if present).

It is also possible to validate the structure of the GFA, for example to check if there are unresolved references to lines. To do this, use the `validate()` method of the `gfapy.Gfa` class.

## **No validations**

If the validation is set to 0, Gfapy will try to accept any input and never raise an exception. This is not always possible, and in some cases, an exception will still be raised, if the data is invalid.

## **Validation when reading**

If the validation level is set to 1 or higher, basic validations will be performed, such as checking the number of positional fields, the presence of duplicated tags, the tag datatype of predefined tags. Additionally, all tags will be validated, either during parsing or on first access. Record-type cross-field validations will also be performed.

In other words, a validation of 1 means that Gfapy guarantees (as good as it can) that the GFA content read from a file is valid, and will raise an exception on accessing the data if not.

The user is supposed to run `validate_field(fieldname)` when changing a field content to something which can be potentially invalid, or `validate()` if potentially cross-field validations could fail.

## **Validation when writing**

Setting the level to 2 will perform all validations described above, plus validate the fields content when their value is written to string.

In other words, a validation of 2 means that Gfapy guarantee (as good as it can) that the GFA content read from a file and written to a file is valid and will raise an exception on accessing the data or writing to file if not.

## **Continuous validation**

If the validation level is set to 3, all validations for lower levels described above are run, plus a validation of fields contents each time a setter method is used.

A validation of 3 means that Gfapy guarantees (as good as it can) that the GFA content is always valid.

## **Positional fields**

Most lines in GFA have positional fields (Headers are an exception). During parsing, if a line is encountered, which has too less or too many positional fields, an exception will be thrown. The correct number of positional fields is record type-specific.

Positional fields are recognized by its position in the line. Each positional field has an implicit field name and datatype associated with it.

### Field names

The field names are derived from the specification. Lower case versions of the field names are used and spaces are substituted with underscores. In some cases, the field names were changed, as they represent keywords in common programming languages (**from**, **send**).

The following tables shows the field names used in Gfapy, for each kind of line. Headers have no positional fields. Comments and custom lines follow particular rules, see the respective chapters.

#### GFA1 field names

Record Type	Field 1	Field 2	Field 3	Field 4	Field 5	Field 6
Segment	name	sequence				
Link	from_segment	from_orient	to_segment	to_orient	overlap	
Containment	from_segment	from_orient	to_segment	to_orient	pos	overlap
Path	path_name	segment_names	overlaps			

#### GFA2 field names

Record Type	Field 1	Field 2	Field 3	Field 4	Field 5	Field 6	Field 7	Field 8
Segment	sid	slen	sequence					
Edge	eid	sid1	sid2	beg1	end1	beg2	end2	alignment
Fragment	sid	external	s_beg	s_end	f_beg	f_end	alignment	
Gap	gid	sid1	d1	d2	sid2	disp	var	
Set	pid	items						
Path	pid	items						

### Datatypes

The datatype of each positional field is described in the specification and cannot be changed (differently from tags). Here is a short description of the Python classes used to represent data for different datatypes. For some complex cases, more details are found in the following chapters.

### Placeholders

The positional fields in GFA can never be empty. However, there are some fields with optional values. If a value is not specified, a placeholder character is used instead (\*). Such undefined values are represented in Gfapy by the `gfapy.Placeholder` class, which is described more in detail in the Placeholders chapter.

## Arrays

The `items` field in unordered and ordered groups and the `segment_names` and `overlaps` fields in paths are lists of objects and are represented by list instances.

```
type(set.items) # => "list"
type(gfa2_path.items) # => "list"
type(gfa1_path.segment_names) # => "list"
type(gfa1_path.overlaps) # => "list"
```

## Orientations

Orientations are represented by strings. The `gfapy.invert()` method applied to an orientation string returns the other orientation.

```
gfapy.invert("+") # => "-"
gfapy.invert("-") # => "+"
```

## Identifiers

The identifier of the line itself (available for S, P, E, G, U, O lines) can always be accessed in Gfapy using the `name` alias and is represented in Gfapy by a string. If it is optional (E, G, U, O lines) and not specified, it is represented by a Placeholder instance. The fragment identifier is also a string.

Identifiers which refer to other lines are also present in some line types (L, C, E, G, U, O, F). These are never placeholders and in stand-alone lines are represented by strings. In connected lines they are references to the Line instances to which they refer to (see the References chapter).

## Oriented identifiers

Oriented identifiers (e.g. `segment_names` in GFA1 paths) are represented by elements of the class `gfapy.OrientedLine`. The `segment` method of the oriented segments returns the segment identifier (or segment reference in connected path lines) and the `orient` method returns the orientation string. The `name` method returns the string of the segment, even if this is a reference to a segment. A new oriented line can be created using the `OL[line, orientation]` method.

Calling `invert` returns an oriented segment, with inverted orientation. To set the two attributes the methods `segment=` and `orient=` are available.

Examples:

```
p = "P\tP1\ta+,b-\t*".to_rgfa_line
p.segment_names # => [OrientedLine("a", "+"), OrientedLine("b", "-")]
p[0].segment # => "a"
p[0].name # => "a"
p[0].orient # => "+"
p[0].invert # => OrientedLine("a", "-")
p[0].orient = "-"
p[0].segment = "S\tX\t*".to_rgfa_line
p[0] # => OrientedLine(gfapy.Line("S\tX\t*"), "-")
p[0].name # => "X"
p[0] = OrientedLine(gfapy.Line("S\tY\t*"), "+")
```

## Sequences

Sequences (S field sequence) are represented by strings in Gfapy. Depending on the GFA version, the alphabet definition is more or less restrictive. The definitions are correctly applied by the validation methods.

The method `rc()` is provided to compute the reverse complement of a nucleotidic sequence. The extended IUPAC alphabet is understood by the method. Applied to non nucleotidic sequences, the results will be meaningless:

```
gfapy.rc("gcat") # => "atgc"
gfapy.rc("*") # => "*" (placeholder)
gfapy.rc("yatc") # => "gatr" (wildcards)
gfapy.rc("gCat") # => "atGc" (case remains)
gfapy.rc("ctg", rna: true) # => "cug"
```

## Integers and positions

The C lines `pos` field and the G lines `disp` and `var` fields are represented by integers. The `var` field is optional, and thus can be also a placeholder. Positions are 0-based coordinates.

The position fields of GFA2 E lines (`beg1`, `beg2`, `end1`, `end2`) and F lines (`s_beg`, `s_end`, `f_beg`, `f_end`) contain a dollar string as suffix if the position is equal to the segment length. For more information, see the Positions chapter.

## Alignments

Alignments are always optional, ie they can be placeholders. If they are specified they are CIGAR alignments or, only in GFA2, trace alignments. For more details, see the Alignments chapter.



## **GFA1 datatypes**

Datatype	Record Type	Fields
Identifier	Segment	name
	Path	path_name
	Link	from_segment, to_segment
	Containment	from_segment, to_segment
[OrientedIdentifier]	Path	segment_names
Orientation	Link	from_orient, to_orient
	Containment	from_orient, to_orient
Sequence	Segment	sequence
Alignment	Link	overlap
	Containment	overlap
[Alignment]	Path	overlaps
Position	Containment	pos

### GFA2 datatypes

Datatype	Record Type	Fields
Identifier	Segment	sid
	Fragment	sid
OrientedIdentifier	Edge	sid1, sid2
	Gap	sid1, sid2
	Fragment	external
OptionalIdentifier	Edge	eid
	Gap	gid
	U Group	oid
	O Group	uid
[Identifier]	U Group	items
[OrientedIdentifier]	O Group	items
Sequence	Segment	sequence
Alignment	Edge	alignment
	Fragment	alignment
Position	Edge	beg1, end1, beg2, end2
	Fragment	s_beg, s_end, f_beg, f_end
Integer	Gap	disp, var

### Reading and writing positional fields

The `positional_fieldnames` method returns the list of the names (as strings) of the positional fields of a line. The positional fields can be read using a method on the Gfapy line object, which is called as the field name. Setting the value is done with an equal sign version of the field name method (e.g. `segment.slen = 120`). In alternative, the `set(fieldname, value)` and `get(fieldname)` methods can

also be used.

```
s_gfa1.positional_fieldnames # => ["name", "sequence"]
s_gfa1.name # => "segment1"
s_gfa1.get("name") # => "segment3"
s_gfa1.name = "segment2"
s_gfa1.name # => "segment2"
s_gfa1.set("name", "segment3")
s_gfa1.name = "segment3"
```

When a field is read, the value is converted into an appropriate object. The string representation of a field can be read using the `field_to_s(fieldname)` method.

```
link.from_segment # => gfapy.line.segment.GFA1("S\tS1\t*")
link.field_to_s(from_segment) # => ("S1")
```

When setting a non-string field, the user can specify the value of a tag either as a Python non-string object, or as the string representation of the value.

```
c.pos = 1
c.pos = "1"
c.pos # => 1
c.field_to_s("pos") # => "1"
```

Note that setting the value of reference and backreferences-related fields is generally not allowed, when a line instance is connected to a Gfapy object (see the References chapter).

```
s = gfa.Line.from_string("L\tS1\t+\tS2\t-\t*")
s.from_segment = "S3"
gfa.add_line(s)
s.from_segment = "S4" # raises an exception
```

## Validation

The content of all positional fields must be a correctly formatted string according to the rules given in the GFA specifications (or a Python object whose string representation is a correctly formatted string).

Depending on the validation level, more or less checks are done automatically (see the Validation chapter). Not regarding which validation level is selected, the user can trigger a manual validation using the `validate_field(fieldname)` method for a single field, or using `validate`, which does a full validation on the whole line, including all positional fields.

```
line.validate_field("xx")
line.validate()
```

## Aliases

For some fields, aliases are defined, which can be used in all contexts where the original field name is used (i.e. as parameter of a method, and the same setter and getter methods defined for the original field name are also defined for each alias, see below).

```
gfa1_path.name == gfa1_path.path_name # True
edge.eid == edge.name # True
segment.sid == segment.name # True
containment.from_segment == containment.container # True
```

```
s = gfapy.Line.from_string("S\t1\t")
s.sid # => "1"
s.name = "a"
s.sid # => "a"
```

## Name

Different record types have an identifier field: segments (name in GFA1, sid in GFA2), paths (path\_name), edge (eid), fragment (sid), gap (gid), groups (pid).

All these fields are aliased to **name**. This allows the user for example to set the identifier of a line using the **name=(value)** method using the same syntax for different record types (segments, edges, paths, fragments, gaps and groups).

## Version-specific field names

For segments the GFA1 name and the GFA2 sid are equivalent fields. For this reason an alias **sid** is defined for GFA1 segments and **name** for GFA2 segments.

## Cryptical field names

The definition of from and to for containments is somewhat cryptical. Therefore following aliases have been defined for containments: `container[_orient]` for `from[_segment|orient]`; `contained[_orient]` for `to[_segment|orient]`.

## Placeholders

Some positional fields may contain an undefined value S: **sequence**; L/C: **overlap**; P: **overlaps**; E: **eid**, **alignment**; F: **alignment**; G: **gid**, **var**; U/O: **pid**. In GFA this value is represented by a **\***.

In Gfapy the class Placeholder represent the undefined value.

## Distinguishing placeholders

The method “`gfapy.is_placeholder()`” checks if a value is or would be represented by a placeholder in GFA (such as an empty array, or a string containing “\*”).

```
gfapy.is_placeholder("*") # => True
gfapy.is_placeholder("**") # => False
gfapy.is_placeholder([]) # => True
gfapy.is_placeholder(gfapy.Placeholder()) # => True
```

Note that, as a placeholder is `False` in boolean context, just a `if not placeholder` will also work, if placeholder is a `gfa.Placeholder()` but not if it is a string representation.

## Compatibility methods

Some methods are defined for placeholders, which allow them to respond to the same methods as defined values. This allows to write generic code.

```
placeholder.validate() # does nothing
len(placeholder) # => 0
placeholder[1] # => gfapy.Placeholder()
placeholder + anything # => gfapy.Placeholder()
```

## Position fields

The only position field in GFA1 is the `pos` field in the C lines. This represents the starting position of the contained segment in the container segment and is 0-based.

Some fields in GFA2 E lines (`beg1`, `beg2`, `end1`, `end2`) and F lines (`s_beg`, `s_end`, `f_beg`, `f_end`) are positions. According to the specification, they are 0-based and represent virtual ticks before and after each string in the sequence. Thus ranges are represented similarly to the Python range conventions: e.g. a 1-character prefix of a sequence will have begin 0 and end 1.

## GFA2 last position string

The GFA2 positions must contain an additional string (\$) appended to the integer, if (and only if) they are the last position in the segment sequence. These particular positions are represented in Gfapy as instances of the class `gfapy.LastPos`.

To create a `lastpos` instance, the constructor can be used with an integer, or the string representation (which must end with the dollar sign, otherwise an integer is returned):

```
str(gfapy.LastPos(12))    # => "12$"
gfapy.LastPos("12")       # => 12
str(gfapy.LastPos("12"))  # => "12"
gfapy.LastPos("12$")      # => gfapy.LastPos(12)
str(gfapy.LastPos("12$")) # => "12$"
```

Subtracting an integer from a `lastpos` returns a `lastpos` if 0 subtracted, an integer otherwise. This allows to do some arithmetic on positions without making them invalid.

```
gfapy.LastPos(12) - 0 # => gfapy.LastPos(12)
gfapy.LastPos(12) - 1 # => 11
```

The functions `gfapy.islastpos` and “`isfirstpos`” allow to determine if a position value is 0 (first), or the last position, using the same syntax for `lastpos` and integer instances.

```
gfapy.isfirst(0)  # True
gfapy.islast(0)   # False
gfapy.isfirst(12) # False
gfapy.islast(12)  # False
gfapy.islast(gfapy.LastPos("12")) # False
gfapy.islast(gfapy.LastPos("12$")) # True
```

## Alignments

Some fields contain alignments and lists of alignments (L/C: overlap; P: overlaps; E/F: alignment). If an alignment is not given, the placeholder symbol `*` is used instead. In GFA1 the alignments can be given as CIGAR strings, in GFA2 also as Dazzler traces.

Gfapy uses different classes for representing the two possible alignment styles (cigar strings and traces) and undefined alignments (placeholders).

### Creating an alignment

An alignment instance is usually created from its GFA string representation or from a list by using the `gfapy.Alignment()` constructor. If the argument is an alignment object it will be returned, so that is always safe to call the method on a variable which can contain a string or an alignment instance:

```
gfapy.Alignment("*")      # => gfapy.AlignmentPlaceholder
gfapy.Alignment("10,10,10") # => gfapy.Trace
```

```

gfapy.Alignment([10,10,10]) # => gfapy.Trace
gfapy.Alignment("30M2I30M") # => gfapy.CIGAR
gfapy.Alignment(gfapy.Alignment("*"))
gfapy.Alignment(gfapy.Alignment("10,10"))

```

## Recognizing undefined alignments

The `gfapy.is_placeholder()` method allows to understand if an alignment field contains a defined value (cigar, trace) or not (placeholder). The method works correctly with both alignment objects and their string or list representation.

```

gfapy.is_placeholder(gfapy.Alignment("30M")) # => False
gfapy.is_placeholder(gfapy.Alignment("10,10")) # => False
gfapy.is_placeholder(gfapy.Alignment("*")) # => True
gfapy.is_placeholder("*") # => True
gfapy.is_placeholder("30M") # => False
gfapy.is_placeholder("10,10") # => True
gfapy.is_placeholder([]) # => True
gfapy.is_placeholder([10,10]) # => False

```

Note that, as a placeholder is `False` in boolean context, just a `if not alignment` will also work, if alignment is an alignment object, but not if it is a string representation.

## Reading and editing CIGARs

CIGARs are represented by arrays of cigar operation objects. Each cigar operation provides the properties `length` and `code`. Length is the length of the CIGAR operation (int). Code is one of the codes allowed by the GFA specification.

```

cigar = gfapy.Alignment("30M")
isinstance(cigar, list) # => True
operation = cigar[0]
type(operation) # => "gfapy.CIGAR.Operation"
operation.code # => "M"
operation.code = "D"
operation.length # => 30
len(operation) # => 30
str(operation) # => "30D"

```

The CIGAR object can be edited using the list methods. If the array is emptied, its string representation will be `*`.

```

cigar = gfapy.Alignment("1I20M2D")
cigar[0].code = "M"
cigar.pop(1)
str(cigar) # => "1M2D"

```

```
cigar[:] = []
str(Cigar) # => "*"

```

CIGARs consider one sequence as reference and another sequence as query. The `length_on_reference` and `length_on_query` methods compute the length of the alignment on the two sequences. These methods are used by the library e.g. to convert GFA1 L lines to GFA2 E lines (which is only possible if CIGARs are provided).

```
cigar = gfapy.Alignment("30M10D20M5I10M")
cigar.length_on_reference() # => 70
cigar.length_on_query() # => 65

```

## Validation

The `validate` method checks if all operations in a cigar use valid codes and length values (which must be non-negative). The codes can be M, I, D or P. For GFA1 the other codes are formally accepted (no exception is raised), but their use is discouraged. An error is raised in GFA2 on validation, if the other codes are used.

```
cigar = gfapy.Alignment("30M10D20M5I10M")
cigar.validate() # no exception raised
cigar[1].code = "L"
cigar.validate # raises an exception
cigar = gfapy.Alignment("30M10D20M5I10M")
cigar[1].code = "X"
cigar.validate(version="gfa1") # no exception raised
cigar.validate(version="gfa2") # exception raised

```

## Reading and editing traces

Traces are arrays of non-negative integers. The values are interpreted using a trace spacing value. If traces are used, a trace spacing value must be defined in a TS integer tag, either in the header, or in the single lines which contain traces.

```
gfa.header.TS # => the global TS value
gfa.line("x").TS # => an edge's own TS tag

```

## Complement alignment

CIGARs are dependent on which sequence is taken as reference and which is taken as query. For each alignment, a complement CIGAR can be computed (using the method `complement`), which is the CIGAR obtained when the two sequences are switched. This method is used by the library e.g. to compare links, as they can be expressed in different ways, by switching the two sequences.



```
cigar = gfapy.Alignment("2M1D3M")
str(cigar.complement()) # => "3M1I2M"
```

The current version of Gfapy does not provide a way to compute the alignment, thus the trace information can be accessed and edited, but not used for this purpose. Because of this there is currently no way in Gfapy to compute a complement trace (trace obtained when the sequences are switched).

```
trace = gfapy.Alignment("1,2,3")
str(trace.complement()) # => "*"
```

The complement of a placeholder is a placeholder:

```
str(gfapy.Alignment("*").complement()) # => "*"
```

## Tags

Each record in GFA can contain tags. Tags are fields which consist in a tag name, a datatype and data. The format is NN:T:DATA where NN is a two-letter tag name, T is an one-letter datatype string and DATA is a string representing the data according to the specified datatype. Tag names must be unique for each line, i.e. each line may only contain a tag once.

```
# Examples of GFA tags of different datatypes:
"aa:i:-12"
"bb:f:1.23"
"cc:Z:this is a string"
"dd:A:X"
"ee:B:c,12,3,2"
"ff:H:122FA0"
'gg:J:["A","B"]'
```

## Custom tags

Some tags are explicitly defined in the specification (these are named *predefined tags* in Gfapy), and the user or an application can define its own custom tags.

Custom tags are user or program specific and may of course collide with the tags used by other users or programs. For this reasons, if you write scripts which employ custom tags, you should always check that the values are of the correct datatype and plausible.

```
if line.get_datatype("xx") != "i":
    raise Exception("I expected the tag xx to contain an integer!")
myvalue = line.xx
if (myvalue > 120) or (myvalue % 2 == 1):
    raise Exception("The value in the xx tag is not an even value <= 120")
# ... do something with myvalue
```

Also it is good practice to allow the user of the script to change the name of the custom tags. For example, Gfapy employs the +or+ custom tag to track the original segment from which a segment in the final graph is derived. All methods which read or write the +or+ tag allow to specify an alternative tag name to use instead of +or+, for the case that this name collides with the custom tag of another program.

```
# E.g. a method which does something with myvalue, usually stored in tag xx
# allows the user to specify an alternative name for the tag
def mymethod(line, mytag="xx"):
    myvalue = line.get(mytag)
    # ...
```

## Tag names in GFA1

According to the GFA1 specification, custom tags are lower case, while predefined tags are upper case (in both cases the second character in the name can be a number). There is a number of predefined tags in the specification, different for each kind of line.

```
"VN:Z:1.0" # VN is upcase => predefined tag
"z5:Z:1.0" # z5 first char is lowercase => custom tag

# not forbidden, but not recommended:
"zZ:Z:1.0" # => mixed case, first char lowercase => custom tag
"Zz:Z:1.0" # => mixed case, first char upcase => custom tag
"vn:Z:1.0" # => same name as predefined tag, but lowercase => custom tag
```

Besides the tags described in the specification, in GFA1 headers, the TS tag is allowed, in order to simplify the translation of GFA2 files.

## Tag names in GFA2

The GFA2 specification is currently not as strict regarding tags: anyone can use both upper and lower case tags, and no tags are predefined except for VN and TS.

However, Gfapy follows the same conventions as for GFA1: i.e. it allows the tags specified as predefined tags in GFA1 to be used also in GFA2. No other upper case tag is allowed in GFA2.

## Datatypes

The following table summarizes the datatypes available for tags:

Symbol	Datatype	Example	Python class
Z	string	This is a string	str
i	integer	-12	int
f	float	1.2E-5	float
A	char	X	str
J	JSON	[1,{"k1":1,"k2":2},"a"]	list/dict
B	numeric array	f,1.2,13E-2,0	gfapy.NumericArray
H	byte array	FFAA01	gfapy.ByteArray

## Validation

The tag name is validated according the the rules described above: except for the upper case tags indicated in the GFA1 specification, and the TS header tag, all other tags must contain at least one lower case letter.

```
"VN:i:1" # => in header: allowed, elsewhere: error
"TS:i:1" # => allowed in headers and GFA2 Edges
"KC:i:1" # => allowed in links, containments, GFA1/GFA2 segments
"xx:i:1" # => custom tag, always allowed
"xxx:i:1" # => error: name is too long
"x:i:1" # => error: name is too short
"11:i:1" # => error: at least one letter must be present
```

The datatype must be one of the datatypes specified above. For predefined tags, Gfapy also checks that the datatype given in the specification is used.

```
"xx:X:1" # => error: datatype X is unknown
"VN:i:1" # => error: VN must be of type Z
```

The data must be a correctly formatted string for the specified datatype or a Python object whose string representation is a correctly formatted string.

```
# current value: xx:i:2
line.xx = 1 # OK
line.xx = "1" # OK, value is set to 1
line.xx = "A" # error
```

Depending on the validation level, more or less checks are done automatically (see validation chapter). Per default - validation level (1) - validation is performed only during parsing or accessing values the first time, therefore the user must perform a manual validation if he changes values to something which is not guaranteed to be correct. To trigger a manual validation, the user can call the method `validate_field(fieldname)` to validate a single tag, or `validate()` to validate the whole line, including all tags.

```
line.xx = "A"
line.validate_field("xx") # validates xx
```

```
# or, to validate the whole line, including tags:
line.validate()
```

## Reading and writing tags

Tags can be read using a property on the Gfapy line object, which is called as the tag (e.g. `line.xx`). A special version of the property prefixed by `try_get_` raises an error if the tag was not available (e.g. `line.try_get_LN`), while the tag property (e.g. `line.LN`) would return `None` in this case. Setting the value is done assigning a value to it the tag name method (e.g. `line.TS = 120`). In alternative, the `set(fieldname, value)`, `get(fieldname)` and `try_get(fieldname)` methods can also be used. To remove a tag from a line, use the `delete(fieldname)` method, or set its value to `None`.

```
# line is "H xx:i:12"
line.xx  # => 1
line.xy  # => nil
line.try_get_xx  # => 1
line.try_get_xy  # => error: xy is not defined
line.get("xx")   # => 1
line.try_get("xy") # => error, xy is not defined
line.xx = 2      # => value of xx is changed to 2
line.xx = "a"    # => error: not compatible with existing type (i)
line.xy = 2      # => xy is created and set to 2, type is auto-set to i
line.set("xy", 2) # => sets xy to 2
line.delete("xy") # => tag is eliminated
line.xx = None   # => tag is eliminated
```

The `tagnames` property of gfapy Line instances is a list of the names (as strings) of all defined tags for a line.

```
print("Line contains the following tags:")
for t in line.tagnames:
    print(t)
if "VN" in line.tagnames:
    # do something with line.VN value
```

When a tag is read, the value is converted into an appropriate object (see Python classes in the datatype table above). When setting a value, the user can specify the value of a tag either as a Python object, or as the string representation of the value.

```
# line is: H xx:i:1 xy:Z:TEXT xz:J:["a","b"]
line.xx # => 1 (Integer)
line.xy # => "TEXT" (String)
line.xz # => ["a", "b"] (Array)
```

The string representation of a tag can be read using the `field_to_s(fieldname)` method. The default is to only output the content of the field. By setting “tag: true“, the entire tag is output (name, datatype, content, separated by colons). An exception is raised if the field does not exist.

```
# line is: H xx:i:1
line.xx # => 1
line.field_to_s("xx") # => "1"
line.field_to_s("xx", tag=True) # => "xx:i:1"
```

### Datatype of custom tags

The datatype of an existing custom field (but not of predefined fields) can be changed using the `set_datatype(fieldname, datatype)` method. The current datatype specification can be read using `get_datatype(fieldname)`.

```
# line is: H xx:i:1
line.get_datatype("xx") # => "i"
line.set_datatype("xx", "Z")
```

If a new custom tag is specified, Gfapy selects the correct datatype for it: i/f for numeric values, J/B for arrays, J for hashes and Z for strings and strings. If the user wants to specify a different datatype, he may do so by setting it with `set_datatype()` (this can be done also before assigning a value, which is necessary if full validation is active).

```
# line has not tags
line.xx = "1" # => "xx:Z:1" created
line.xx      # => "1"
line.set_datatype("xy", "i")
line.xy = "1" # => "xy:i:1" created
line.xy      # => 1
```

### Arrays of numerical values

B and H tags represent array with particular constraints (e.g. they can only contain numeric values, and in some cases the values must be in predefined ranges). In order to represent them correctly and allow for validation, Python classes have been defined for both kind of tags: `gfapy.ByteArray` for H and `gfapy.NumericArray` for B fields.

Both are subclasses of list. Object of the two classes can be created by passing an existing list or the string representation to the class constructor.

```
# create a byte array instance
gfapy.ByteArray([12,3,14])
gfapy.ByteArray("A012FF")
```

```
# create a numeric array instance
gfapy.NumericArray("c,12,3,14")
gfapy.NumericArray([12,3,14])
```

Instances of the classes behave as normal lists, except that they provide a `#validate()` method, which checks the constraints, and that their string representation is the GFA string representation of the field value.

```
gfapy.ByteArray([12,1,"1x"]).validate() # error: 1x is not a valid value
str(gfapy.ByteArray([12,3,14])) # => "c,12,3,14"
```

For numeric values, the `compute_subtype()` method allows to compute the subtype which will be used for the string representation. Unsigned subtypes are used if all values are positive. The smallest possible subtype range is selected. The subtype may change when the range of the elements changes.

```
gfapy.NumericValue([12,13,14]).compute_subtype() # => "C"
```

### Special cases: custom records, headers, comments and virtual lines.

GFA2 allows custom records, introduced by record type strings other than the predefined ones. Gfapy uses a pragmatical approach for identifying tags in custom records, and tries to interpret the rightmost fields as tags, until the first field from the right raises an error; all remaining fields are treated as positional fields.

```
"X a b c xx:i:12" # => xx is tag, a, b, c are positional fields
"Y a b xx:i:12 c" # => all positional fields, as c is not a valid tag
```

For easier access, the entire header of the GFA is summarized in a single line instance. A class (`gfapy.FieldArray`) has been defined to handle the special case when multiple H lines define the same tag (see “Header” chapter for details).

Comment lines are represented by a subclass of the same class (`gfapy.Line`) as the records. However, they cannot contain tags: the entire line is taken as content of the comment. See the “Comments” chapter for more information about comments.

```
"# this is not a tag: xx:i:1" # => xx is not a tag, xx:i:1 is part of the comment
```

Virtual `gfapy.Line` instances (e.g. segment instances automatically created because of not yet resolved references found in edges) cannot be modified by the user, and tags cannot be specified for them. This includes all instances of the `gfapy::Line::Unknown` class. See the “References” chapter for more information about virtual lines.

## References

Some fields in GFA lines contain identifiers or lists of identifiers (sometimes followed by orientation strings), which reference other lines of the GFA file. In Gfapy it is possible to follow these references and traverse the graph.

### Connecting a line to a Gfa object

In stand-alone line instances, the identifiers which reference other lines are either strings containing the line name, pairs of strings (name and orientation) in a `gfapy.OrientedLine` object, or lists of lines names or `gfapy.OrientedLine` objects.

Using the `add_line(line)` (alias: `append(line)`) method of the `gfapy.Gfa` object, or the equivalent `connect(gfa)` method of the `gfapy.Line` instance, a line is added to a Gfa instance (this is done automatically when a GFA file is parsed). All strings expressing references are then changed into references to the corresponding line objects. The method `is_connected()` allows to determine if a line is connected to an `gfapy` instance. The read-only property `gfa` contains the `gfapy.Gfa` instance to which the line is connected.

```
link.is_connected() # => False
link.gfa            # => None
link.from_segment   # => "A"
link.connect(gfa)   # or gfa.add_line(link); or gfa.append(link)
link.is_connected() # => True
link.gfa            # => gfapy.Gfa(...)
link.from_segment   # => gfapy.Segment("S\tA\t*", ...)
```

### References for each record type

The following tables describes the references contained in each record type. The notation `[]` represent lists.

#### GFA1

Record type	Fields	Type of reference
Link	from, to	Segment
Containment	from, to	Segment
Path	segment_names, links (1)	[OrientedLine(Segment)] [OrientedLine(Link)]

(1): paths contain information in the fields `segment_names` and `overlaps`, which

allow to find the identify from which they depend; these links can be retrieved using `links` (which is not a field).

## GFA2

Record type	Fields	Type of reference
Edge	sid1, sid2	Segment
Gap	sid1, sid2	Segment
Fragment	sid	Segment
Set	items	[Edge/Set/Path/Segment]
Path	items	[OrientedLine(Edge/Set/Segment)]

## Backreferences for each record type

When a line containing a reference to another line is connected to a Gfa object, backreferences to it are created in the targeted line.

For each backreference collection a read-only property exist, which is named as the collection (e.g. `dovetails_L` for segments). Note that the reference list returned by these arrays are read-only and editing the references is done using other methods (see the section “Editing reference fields” below).

`segment.dovetails_L # => [gfapy.line.edge.Link(...), ...]`

The following tables describe the backreferences collections for each record type.

## GFA1

Record type	Backreferences
Segment	dovetails_L dovetails_R edges_to_contained edges_to_containers paths
Link	paths

## GFA2

Record type	Backreferences	Type
Segment	dovetails_L	E
	dovetails_R	E
	edges_to_contained	E



Record type	Backreferences	Type
Edge	edges_to_containers	E
	internals	E
	gaps_L	G
	gaps_R	G
	fragments	F
	paths	O
	sets	U
	paths	O
	sets	U
	paths	O
O Group	sets	U
U Group	sets	U

### Segment backreference convenience methods

For segments, additional methods are available which combine in different way the backreferences information. The `dovetails_of_end(end)` and `gaps_of_end(end)` methods take an argument “L” or “R” and return the dovetails overlaps (or gaps) of the left or, respectively, right end of the segment sequence are returned (equivalent to `dovetails_L/dovetails_R` and `gaps_L/gaps_R`).

The segment `containments` methods returns both containments where the segment is the container or the contained segment. The segment `edges` property is a list of all edges (dovetails, containments and internals) with a reference to the segment.

Other methods directly compute list of segments from the edges lists mentioned above. The `neighbours_L`, `neighbours_R` properties and the “`neighbours(end)`” method computes the set of segment instances which are connected by dovetails to the segment. The `segmentcontainersandcontained` properties similarly compute the set of segment instances which, respectively, contains the segment, or are contained in the segment.

```
s.dovetails_of_end("L") # => [gfapy.line.edge.Link(...), ...]
s.dovetails_L == segment.dovetails_of_end("L") # => True
s.gaps_of_end("R") # => []
s.edges # => [gfapy.line.edge.Link(...), ...]
s.neighbours_L # => [gfapy.line.segment.GFA1(...), ...]
s.containers # => [gfapy.line.segment.GFA1(...), ...]
```

## Multiline group definitions

The GFA2 specification opens the possibility (experimental) to define groups on multiple lines, by using the same ID for each line defining the group. This is supported by gfapy.

This means that if multiple `gfapy.line.group.Ordered` or `gfapy.line.group.Unordered` instances connected to a Gfa object have the same `gid`, they are merged into a single instance (technically the last one getting added to the graph object). The items list are merged.

The tags of multiple line defining a group shall not contradict each other (i.e. either are the tag names on different lines defining the group all different, or, if the same tag is present on different lines, the value and datatype must be the same, in which case the multiple definition will be ignored).

```
gfa.add_line("U\tu1\tts1 s2 s3")
[s.name for s in gfa.sets[-1].items] # => ["s1", "s2", "s3"]
gfa.add_line("U\tu1\tt4 5")
[s.name for s in gfa.sets[-1].items] # => ["s1", "s2", "s3", "s4", "s5"]
```

## Induced set and captured path

The item list in GFA2 sets and paths may not contain elements which are implicitly involved. For example a path may contain segments, without specifying the edges connecting them, if there is only one such edge. Alternatively a path may contain edges, without explicitly indicating the segments. Similarly a set may contain edges, but not the segments referred to in them, or contain segments which are connected by edges, without the edges themselves. Furthermore groups may refer to other groups (set to sets or paths, paths to paths only), which then indirectly contain references to segments and edges.

Gfapy provides methods for the computation of the sets of segments and edges which are implied by an ordered or unordered group. Thereby all references to subgroups are resolved and implicit elements are added, as described in the specification. The computation can, therefore, only be applied to connected lines. For unordered groups, this computation is provided by the method `induced_set()`, which returns an array of segment and edge instances. For ordered group, the computation is provided by the method `captured_path()`, which returns a list of `gfapy.OrientedLine` instances, alternating segment and edge instances (and starting and ending in segments).

The methods `induced_segments_set()`, `induced_edges_set()`, `captured_segments()` and `captured_edges()` return, respectively, the list of only segments or edges, in ordered or unordered groups.

```
gfa.add_line("U\tu1\tts1 s2 s3")
u = gfa.sets[-1]
```

```
u.induced_edges_set # => [gfapy.line.edge.GFA2("E\te1\ts1+\ts2-...", ...)]
[l.name for l in u.induced_set ] # => ["s1", "s2", "s3", "e1"]
```

## Disconnecting a line from a Gfa object

Lines can be disconnected using the `rm(line)` method of the `gfapy.Gfa` object or the `disconnect()` method of the line instance.

```
line = gfa.segment("sA")
gfa.rm(line)
# or equivalent:
line.disconnect()
```

Disconnecting a line affects other lines as well. Lines which are dependent on the disconnected line are disconnected as well. Any other reference to disconnected lines is removed as well. In the disconnected line, references to lines are transformed back to strings and backreferences are deleted.

The following tables show which dependent lines are disconnected if they refer to a line which is being disconnected.

### GFA1

Record type	Dependent lines
Segment	links (+ paths), containments
Link	paths

### GFA2

Record type	Dependent lines
Segment	edges, gaps, fragments, sets, paths
Edge	sets, paths
Sets	sets, paths

## Editing reference fields

In connected line instances, it is not allowed to directly change the content of fields containing references to other lines, as this would make the state of the `Gfa` object invalid.

Besides the fields containing references, some other fields are read-only in connected lines. Changing some of the fields would require moving the backreferences to other collections (position fields of edges and gaps, `from_orient` and

`to_orient` of links). The `overlaps` field of connected links is read-only as it may be necessary to identify the link in paths.

### Renaming an element

The name field of a line (e.g. segment `name/sid`) is not a reference and thus can be edited also in connected lines. When the name of the line is changed, no manual editing of references (e.g. `from/to` fields in links) is necessary, as all lines which refer to the line will still refer to the same instance. The references to the instance in the Gfa lines collections will be automatically updated. Also, the new name will be correctly used when converting to string, such as when the Gfa instance is written to a GFA file.

Renaming a line to a name which already exists has the same effect of adding a line with that name. That is, in most cases, `gfapy.NotUniqueError` is raised. An exception are GFA2 sets and paths: in this case the line will be appended to the existing line with the same name (as described in “Multiline group definitions”).

### Adding and removing group elements

Elements of GFA2 groups can be added and removed from both connected and non-connected lines, using the following methods.

To add an item to or remove an item from an unordered group, use the methods `add_item(item)` and `rm_item(item)`, which take as argument either a string (identifier) or a line instance.

To append or prepend an item to an ordered group, use the methods `append_item(item)` and `prepend_item(item)`. To remove the first or the last item of an ordered group use the methods `rm_first_item()` and `rm_last_item()`.

### Editing read-only fields of connected lines

Editing the read-only information of edges, gaps, links, containments, fragments and paths is more complicated. These lines shall be disconnected before the edit and connected again to the Gfa object after it. Before disconnecting a line, you should check if there are other lines dependent on it (see tables above). If so, you will have to disconnect these lines first, eventually update their fields and reconnect them at the end of the operation.

### Virtual lines

The order of the lines in GFA is not prescribed. Therefore, during parsing, or constructing a Gfa in memory, it is possible that a line is referenced to, before it

is added to the Gfa instance. Whenever this happens, Gfapy creates a “virtual” line instance.

Users do not have to handle with virtual lines, if they work with complete and valid GFA files.

Virtual lines are similar to normal line instances, with some limitations (they contain only limited information and it is not allowed to add tags to them). To check if a line is a virtual line, one can use the `is_virtual()` method of the line.

As soon as the parser finds the real line corresponding to a previously introduced virtual line, the virtual line is exchanged with the real line and all references are corrected to point to the real line.

```
g = gfapy.Gfa()
g.add_line("S\t1\t*")
g.add_line("L\t1\t+\t2\t+\t*")
l = g.dovetails[-1]
g.segment("1").is_virtual() # => False
g.segment("2").is_virtual() # => True
l.to_segment == g.segment("2") # => True
g.segment("2").dovetails = [l] # => True
g.add_line("S\t2\t*")
g.segment("2").is_virtual() # => False
l.to_segment == g.segment("2") # => True
g.segment("2").dovetails = [l] # => True
```

## The Header

GFA files may contain one or multiple header lines (record type: “H”). These lines may be present in any part of the file, not necessarily at the beginning.

Although the header may consist of multiple lines, its content refers to the whole file. Therefore in Gfapy the header is accessed using a single line instance (accessible by the `header` method). Header lines contain only tags. If not header line is present in the Gfa, then the header line object will be empty (i.e. contain no tags).

Note that header lines cannot be connected to the Gfa as other lines (i.e. calling `connect` on them raises an exception). Instead they must be merged to the existing Gfa header, using `add_line(line)` on the gfa instance.

```
gfapy.Line.from_string("H\tnn:f:1.0").connect(gfa) # exception
gfa.add_line("H\tnn:f:1.0") # this works!
gfa.header.nn # => 1.0
```

## Multiple definitions of the predefined header tags

For the predefined tags (VN and TS), the presence of multiple values in different lines is an error, unless the value is the same in each instance (in which case the repeated definitions are ignored).

```
gfa.add_line("H\tVN:Z:1.0")
gfa.add_line("H\tVN:Z:1.0") # ignored
gfa.add_line("H\tVN:Z:2.0") # exception!
```

## Multiple definitions of custom header tags

If the tags are present only once in the header in its entirety, the access to the tags is the same as for any other line (see Tags chapter).

However, the specification does not forbid custom tags to be defined with different values in different header lines (which we name “multi-definition tags”). This particular case is handled in the next sections.

## Reading multi-definitions tags

Reading, validating and setting the datatype of multi-definition tags is done using the same methods as for all other lines (see Tags chapter). However, if a tag is defined multiple times on multiple H lines, reading the tag will return a list of the values on the lines. This array is an instance of the subclass `gfapy.FieldArray` of list.

```
gfa.add_line("H\txx:i:1")
gfa.add_line("H\txx:i:2")
gfa.add_line("H\txx:i:3")
gfa.header.xx # => gfapy.FieldArray("i", [1,2,3])
```

## Setting tags

There are two possibilities to set a tag for the header. The first is the normal tag interface (using `set` or the tag name property). The second is to use `add`. The latter supports multi-definition tags, i.e. it adds the value to the previous ones (if any), instead of overwriting them.

```
gfa.header.xx # => None
gfa.header.add("xx", 1)
gfa.header.xx # => 1
gfa.header.add("xx", 2)
gfa.header.xx # => gfapy.FieldArray("i", [1,2])
gfa.header.set("xx", 3)
gfa.header.xx # => 3
```

## Modifying field array values

Field arrays can be modified directly (e.g. adding new values or removing some values). After modification, the user may check if the array values remain compatible with the datatype of the tag using the `validate_field` method.

```
gfa.header.xx # => gfapy.FieldArray([1,2,3])
gfa.header.validate_field("xx") # => True
gfa.header.xx.append("X")
gfa.header.validate_field("xx") # => False
```

If the field array is modified using array methods which return a list or data of any other type, a field array must be constructed, setting its datatype to the value returned by calling `get_datatype(tagname)` on the header.

```
gfa.header.xx # => gfapy.FieldArray([1,2,3])
gfa.header.xx = gfa.FieldArray(gfa.header.get_datatype("xx"),
                               map(lambda x: x+1, gfa.header.xx))
gfa.header.xx # => gfapy.FieldArray([2,3,4])
```

## String representation of the header

For consistency with other line types, the string representation of the header is a single-line string, eventually non standard-compliant, if it contains multiple instances of the tag. (and when calling `field_to_s(tag)` for a tag present multiple times, the output string will contain the instances of the tag, separated by tabs).

However, when the Gfa is output to file or string, the header is splitted into multiple H lines with single tags, so that standard-compliant GFA is output. The splitted header can be retrieved using the `headers` method on the Gfa instance.

```
gfa.header.field_to_s("xx") # => "xx:i:1\txx:i:2"
str(gfa.header) # => "H\tVN:Z:1.0\txx:i:1\txx:i:2"
[str(h) for h in gfa.headers] # => ["H\tVN:Z:1.0", "H\txx:i:1", "H\txx:i:2"]
str(gfa) # => """
    H VN:Z:1.0
    H xx:i:1
    H xx:i:2
    """
```

## Custom records

According to the GFA2 specification, each line which starts with a non-standard record type shall be considered an user- or program-specific record.

Gfapy allows to retrieve custom records and access their data using a similar interface to that for the predefined record types. It assumes that custom records consist of tab-separated fields and that the first field is the record type.

Validation of custom records is very limited; therefore, if you work with custom records, you may define your own validation method and call it when you read or write custom record contents.

### Retrieving, adding and deleting custom records

The custom records of a Gfa instance can be retrieved using its `custom_records` property. This returns a list of all custom records, regardless of the record type.

To retrieve only the custom records of a given type use the method `custom_records_of_type(record_type)`.

```
gfa.custom_records
gfa.custom_records_of_type("X")
```

Adding custom records to and removing them from a Gfa instance is similar to any other line. So to delete a custom record, `disconnect()` is called on the instance. To add a custom record line, the instance or its string representation is added using `add_line` on the Gfa instance.

```
gfa.add_line("X\ta\tb")
gfa.custom_records("X")[-1].disconnect()
```

### Tags

As Gfapy cannot know how many positional fields are present when parsing custom records, an heuristic approach is followed, to identify tags. A field resembles a tag if it starts with `tn:d:` where `tn` is a valid tag name and `d` a valid tag datatype (see Tags chapter). The fields are parsed from the last to the first. As soon as a field is found which does not resemble a tag, all remaining fields are considered positionals (even if another field parsed later resembles a tag).

```
gfa.add_line("X\ta\tb\tcc:i:10\tdd:i:100")
x1 = gfa.custom_records("X")[-1]
x1.cc # => 10
x1.dd # => 100
gfa.add_line("X\ta\tb\tcc:i:10\tdd:i:100\te")
x2 = gfa.custom_records("X")[-1]
x1.cc # => None
x1.dd # => None
```

This parsing heuristics has some consequences on validations. Tags with an invalid tag name (such as starting with a number, or with a wrong number of



letters), or an invalid tag datatype (wrong letter, or wrong number of letters) are considered positional fields. The only validation available for custom records tags is thus the validation of the content of the tag, which must be valid according to the datatype.

```
gfa.add_line("X\ta\tb\tcc:i:10\tddd:i:100")
x = gfa.custom_records("X")[-1]
x.cc # => None
# (as ddd:i:100) is considered a positional field
```

## Positional fields

The positional fields in a custom record are called "field1", "field2", .... The user can iterate over the positional field names using the array obtained by calling `positional_fieldnames` on the line.

Positional fields are allowed to contain any character (including non-printable characters and spacing characters), except tabs and newlines (as they are structural elements of the line).

Due to the parsing heuristics mentioned in the Tags section above, invalid tags are sometimes wrongly taken as positional fields. Therefore, the user is responsible of validating the number of positional fields.

```
gfa.add_line("X\ta\tb\tcc:i:10\tdd:i:100")
x = gfa.custom_records("X")[-1]
len(x.positional_fieldnames) # => 2
x.positional_fieldnames # => ["a", "b"]
```

## Extensions

The support for custom fields is limited, as Gfapy does not know which and how many fields are there and how shall they be validated. It is possible to create an extension of Gfapy, which defines new record types: this will allow to use these record types in a similar way to the built-in types. However, extending the library requires slightly more advanced programming than just using the predefined record types.

The manual for writing extensions is provided as Supplementary Information to the manuscript describing Gfapy.

## Comments

GFA lines starting with a # symbol are considered comments. In Gfapy comments are represented by instances of `gfapy.line.Comment`. They have a similar

interface to other line instances, with some differences, e.g. they do not support tags.

### Accessing the comments

Adding a comment to a `gfapy.Gfa` instance is done similarly to other lines, by using the `add_line(line)` method. The comments of a `Gfa` object can be accessed using the `comments` method. This returns a list of comment line instances. To remove a comment from the `Gfa`, you need to find the instance in the list, and call `disconnect()` on it.

```
g.add_line("# this is a comment")
[str(c) for c in g.comments] # => ["# this is a comment"]
g.comments[0].disconnect()
g.comments # => []
```

### Accessing the comment content

The content of the comment line, excluding the initial `+#` and eventual initial spacing characters, is included in the field `+content+`.

The initial spacing characters can be read/changed using the `+spacer+` field. The default value is a single space.

```
g.add_line("# this is a comment")
c = g.comments[-1]
g.content # => "this is a comment"
g.spacer # => " "
```

Tags are not supported by comment lines. If the line contains tags, these are not parsed, but included in the `+content+` field. Trying to set tags values raises exceptions.

```
c = gfapy.Line.from_string("# this is not a tag\txx:i:1")
c.content # => "this is not a tag\txx:i:1"
c.xx # => None
c.xx = 1 # raises an exception
```

## Errors

All exception raised in the library are subclasses of `gfapy.Error`. This means that `except gfapy.Error` catches all library errors.

Different types of errors are defined and are summarized in the following table:

Error	Description	Examples
Version	An unknown or wrong version is specified or implied	“GFA0”; or GFA1 in GFA2 context
Value	The value of an object is invalid	a negative position is used
Type	The wrong type has been used or specified	Z instead of i used for VN tag; Hash for an i tag
Format	The format of an object is wrong	a line does not contain the expected number of fields
NotUnique	Something should be unique but is not	duplicated tag name or line identifier
Inconsistency	Pieces of information collide with each other	length of sequence and LN tag do not match
Runtime	The user tried to do something which is not allowed	editing from/to field in connected links
Argument	Problem with the arguments of a method	wrong number of arguments in dynamically created method
Assertion	Something unexpected happened	there is a bug in the library

Some error types are generic (such as `RuntimeError` and `ArgumentError`), and their definition may overlap that of more specific errors (such as `ArgumentError`, which overlaps `ValueError` and `TypeError`). The user should not rely on the type of error alone, but rather take it as an indication. The error message tries to be informative and for this reason often prints information on the internal state of the relevant variables.

Assertion errors are reserved for those situation where something is implied by the programmer (e.g. a value is implied to be positive at a certain point of the code). If the checks fails, an assertion error is raised. The user may report the problem, as this may indicate a bug (unless the user did something he was not supposed to do, such as calling an API private method).

## Graph operations

Graph operations such as linear paths merging, multiplication of segments and other are provided. These operations are similar to those provided by the RGFA library. A description of these operation can be found in the RGFA paper (Gonnella and Kurtz, 2016).