

EBU

OPERATING EUROVISION AND EURORADIO

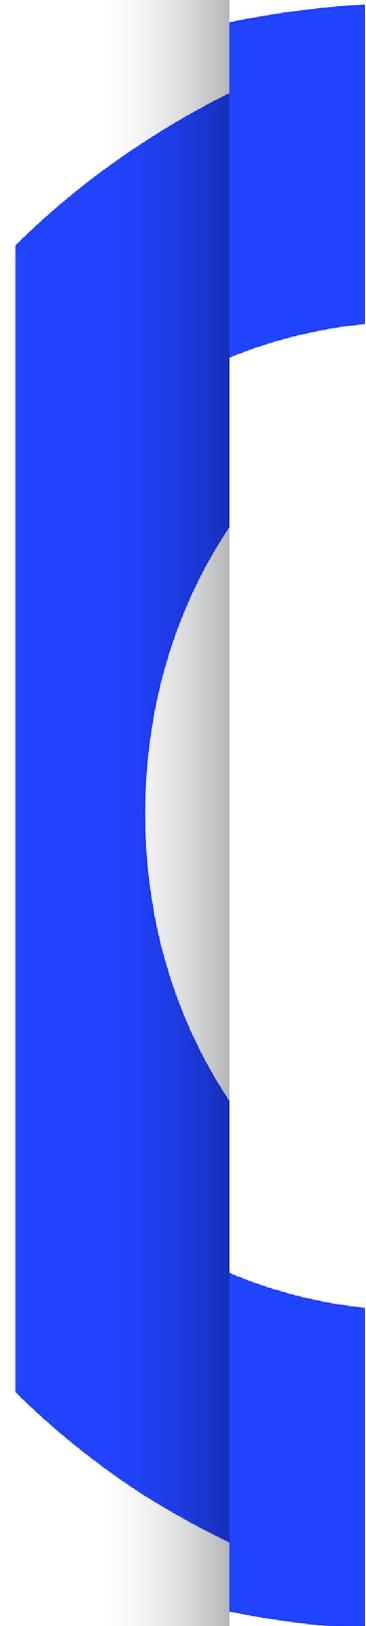
TECH 3370

EBU-TT, PART 3 LIVE SUBTITLING APPLICATIONS

SYSTEM MODEL AND CONTENT PROFILE FOR AUTHORIZING AND CONTRIBUTION

STATUS: Development Version 0.9,
For Prototype Implementation

Geneva
December 2015



Conformance Notation

This document contains both normative text and informative text.

All text is normative except for that in the Introduction, any section explicitly labelled as 'Informative' or individual paragraphs which start with 'Note:'.

Normative text describes indispensable or mandatory elements. It contains the conformance keywords 'shall', 'should' or 'may', defined as follows:

- | | |
|----------------------------|---|
| 'Shall' and 'shall not': | Indicate requirements to be followed strictly and from which no deviation is permitted in order to conform to the document. |
| 'Should' and 'should not': | Indicate that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others.
OR indicate that a certain course of action is preferred but not necessarily required.
OR indicate that (in the negative form) a certain possibility or course of action is deprecated but not prohibited. |
| 'May' and 'need not': | Indicate a course of action permissible within the limits of the document. |

Default identifies mandatory (in phrases containing "shall") or recommended (in phrases containing "should") presets that can, optionally, be overwritten by user action or supplemented with other options in advanced applications. Mandatory defaults must be supported. The support of recommended defaults is preferred, but not necessarily required.

Informative text is potentially helpful to the user, but it is not indispensable and it does not affect the normative text. Informative text does not contain any conformance keywords.

A conformant implementation is one which includes all mandatory provisions ('shall') and, if implemented, all recommended provisions ('should') as described. A conformant implementation need not implement optional provisions ('may') and need not implement them as described.

Contents

Status of this document (Informative)	7
Scope (Informative)	9
1. Introduction (Informative).....	10
1.1 EBU-TT as exchange format for live and prepared subtitles.....	10
1.2 Summary of key points	11
1.3 Example scenarios	13
1.4 Document Structure.....	13
2. Definitions and Concepts	14
2.1 Definition of terms	14
2.2 System Model	16
2.3 Timing and synchronisation	18
2.3.1 Document resolved begin and end times	18
2.3.1.1 Document resolved begin time	19
2.3.1.2 Document resolved end time	19
2.3.1.3 Behaviour when no document is active	19
2.3.1.4 Document creation and issuing strategies (informative).....	20
2.3.2 Management of delay in a live authoring environment	25
2.3.3 ebuttm:authoringDelay	26
2.3.4 Delay nodes	27
2.3.5 Reference clocks	28
2.4 Handover	28
2.4.1 Authors Group parameters	29
2.5 Describing facets of the subtitle content	31
2.5.1 Document level summaries of facets	31
2.6 Tracing and debugging.....	32
3. Document Conformance	32
3.1 Generic Constraints	32
3.1.1 Namespaces	32
3.1.2 Extensibility.....	33
3.1.3 Compatibility with TTML 1.0 timing model	33
3.2 Document Structure and Content Profile.....	33
3.2.1 Elements and attributes whose cardinality differs	34
3.2.2 Newly introduced and modified elements and attributes	34
3.2.2.1 tt:tt.....	35
3.2.2.2 ebuttm:documentMetadata	37
3.2.2.3 tt:body.....	39
3.2.2.4 tt:div	41
3.2.2.5 tt:p	42
3.2.2.6 tt:metadata	43
3.3 Datatypes	44
3.3.1 ebuttdt:delayTimingType	44

4.	Node Conformance	45
4.1	Generic Node Classes	45
4.1.1	Node	45
4.1.1.1	Processing Node	45
4.1.1.2	Passive Node	46
5.	Bibliography	46
Annex A: Overview document structure (Informative)		47
Annex B: Requirements for Carriage Specifications		57
	Core networking dependencies and connection protocols	57
	Synchronisation impacts and/or thresholds	57
	Information Security	57
	Endpoint cardinality	57
	Connection lifecycle management	58
	Channel routing	58
	Stability	58
	Interoperability	58
Annex C: Not fully defined example scenarios (Informative)		59
Annex D: Summary of document conformance differences relative to Part 1 (Informative)		61

Status of this document (Informative)

This document is a working draft and may be updated, replaced or made obsolete by other documents at any time. It is inappropriate to cite this document as other than a work in progress.

This document is part of a series of EBU-TT (EBU Timed Text) documents. The full list of published and planned EBU-TT documents is given below.

Part 1: EBU-TT Subtitling format definition (EBU Tech 3350) [EBUTT1]

Introduction to EBU-TT and definition of the XML based format.

Part 2: STL (Tech 3264) Mapping to EBU-TT (EBU Tech 3360)

How EBU-TT provides backwards compatibility with EBU STL.

Part 3: EBU-TT in Live Subtitling applications: system model and content profile for authoring and contributions

How to use EBU-TT for the production and contribution of live subtitles.

EBU-TT Annotation

How EBU-TT can be used in future scenarios for 'authoring of intent'.

EBU-TT User Guide

General guide ('How to use EBU-TT').

EBU-TT-D (EBU Tech 3380) [EBUTTD]

EBU-TT content profile for TTML that can be used for the distribution of subtitles over IP based networks.

Carriage of EBU-TT-D in ISOBMFF (EBU Tech 3381)

How EBU-TT-D can be stored using the storage format of the ISO Base Media File Format (ISO/IEC 14496-12).

EBU-TT in Live Subtitling applications: System model and content profile for authoring and contributions

<i>EBU Committee</i>	<i>First Issued</i>	<i>Revised</i>	<i>Re-issued</i>
TC	2015		

Keywords: EBU Timed Text, EBU-TT, Subtitle, STL, XML, W3C, TTML, DFXP, caption, encoder, live.

Scope (Informative)

The topic of live subtitle or caption authoring, routing and encoding is large. Organisations such as broadcasters, access service providers and other content providers face a variety of challenges ranging from the editorial, for example word accuracy and rate, to the technical, for example how the text data is routed from the author to the encoder, what format it should be in and how it can be configured and monitored. The classical way to address such a large “problem space” is to divide it up into more easily solvable constituent parts. This approach is taken here.

This document describes how EBU-TT [EBUTT1] can be used in a broadcasting environment to carry subtitles that are created in real time (“live” or from a prepared file) from an authoring station to an encoder prior to distribution, via intermediate processing units. It does this by specifying:

- a system model consisting of processing nodes that pass streams of subtitles along a chain;
- a content profile based on EBU-TT Part 1 specifying the data format of each document in a live subtitle stream;
- a mechanism by which content providers can model and potentially improve synchronisation between the subtitles and the audio to which they relate;
- a mechanism for interoperable management of the handover from one subtitler to the next, to generate a single output subtitle stream.

Future documents will address the carriage of live subtitle streams, i.e. the requirements of protocols for passing live subtitle streams between nodes and how specific protocols meet those requirements. This document defines what such a carriage specification document should consider.

Authoring conventions, for example the use of colour to identify speakers, are not directly addressed in this document; however care has been taken to ensure that the technical solutions presented here can support a variety of conventions. System setup, configuration, management, resilience, monitoring and recovery are likewise addressed indirectly by modelling the potential architectures in the abstract and designing the data format to support those architectures.

No requirement is presented here for any device used by an audience member to receive and present subtitles to behave differently in the presence of live subtitles compared to prepared subtitles: rather, it is envisaged that any processing required to manage the two cases, and the transitions between them, occurs in the broadcaster domain.

1. Introduction (Informative)

EBU-TT Live specifies how to transfer live and non-live subtitles between production systems, e.g. from an authoring station to encoder. It builds on EBU-TT part 1, but has more relaxed conformance rules. Styling and layout details can be left out for example. For home delivery, other specifications, such as EBU-TT-D, are used.

1.1 EBU-TT as exchange format for live and prepared subtitles

EBU-TT as defined in EBU-TT Part 1 [EBUTT1] is intended as a general purpose exchange format for subtitles. The workflow diagram in Part 1 is extended by this document to include live subtitle exchange.

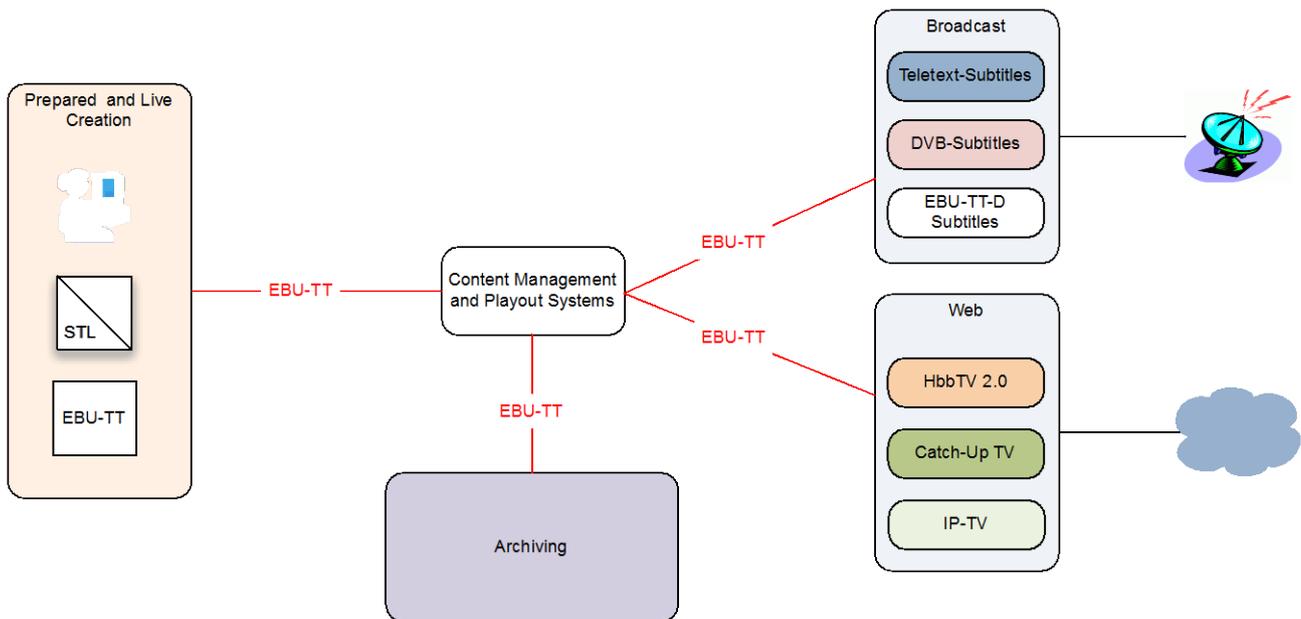


Figure 1: Subtitle workflow with EBU-TT Live

Whereas EBU-TT Part 1 focuses on document exchange for prepared subtitle documents that apply to completed programmes, EBU-TT Part 3 extends this workflow to incorporate the streaming of subtitles that may be created live in real time for programmes that could also be live. See Figure 1.

The benefits of EBU-TT Part 3 relative to other existing ways to transfer live streams of subtitles include:

- It can be used both as a source for and be generated from other EBU-TT and TTML based subtitle formats, using XML processing techniques. It is therefore possible to build a complete subtitle infrastructure based on EBU-TT throughout.
- It is an open standard.
- In general it is stateless: no prior knowledge of documents in a stream is needed to begin processing live subtitles to generate the intended output.
- It is extensible.
- It can be used in a variety of levels of complexity to suit implementation needs.
- It can be used within existing subtitle routing infrastructures, subject to constraints that could apply such as maximum data rates.

1.2 Summary of key points

The subtitles are carried in Documents, each of which is a valid W3C TTML document. Besides the subtitle's text it can contain styling, layout, timing and additional metadata information.

```
<tt xmlns:ebuttp="urn:ebu:tt:parameters" xmlns="http://www.w3.org/ns/ttml" ... xml:lang="en"
ttp:cellResolution="40 24" ttp:clockMode="local" ttp:timeBase="clock" ebuttm:authoringDelay="5s"
ebuttm:sequenceIdentifier="testSequence_1441882303" ebuttm:sequenceNumber="1858107">
  <head>
    <styling>
      <style xml:id="s1" tts:backgroundColor="black" tts:color="lime" tts:fontSize="1c 2c"/>
    </styling>
    <layout>
      <region xml:id="r0" tts:extent="37c 4c" tts:origin="3c 20c"/>
    </layout>
  </head>
  <body begin="10:29:32.36">
    <div>
      <p region="r0">
        <span style="s1"> Hello. Have you not done something <br/>
you should have done, for me? </span>
      </p>
    </div>
  </body>
</tt>
```

Figure 2: Example EBU-TT Part 3 document

One or more Documents together form a Sequence. Each Document indicates to which Sequence it belongs using a Sequence Identifier. The order is set by a Sequence Number.

```
<tt ... ebuttm:sequenceIdentifier="testSequence_1441882303" ebuttm:sequenceNumber="1858107">
...
</tt>

<tt ... ebuttm:sequenceIdentifier="testSequence_1441882303" ebuttm:sequenceNumber="1858108">
...
</tt>

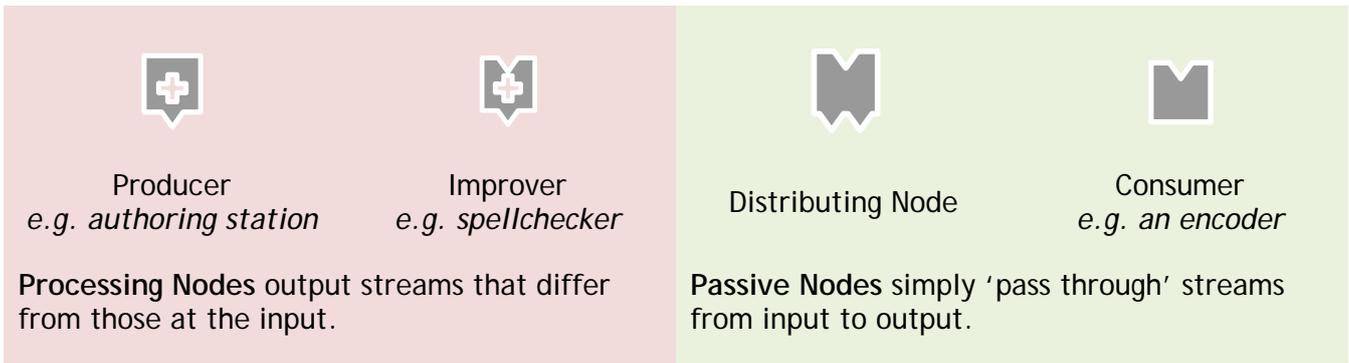
<tt ... ebuttm:sequenceIdentifier="testSequence_1441882303" ebuttm:sequenceNumber="1858109">
...
</tt>

<tt
... ebuttm:sequenceIdentifier="testSequence_1441882303" ebuttm:sequenceNumber="1858110">
```

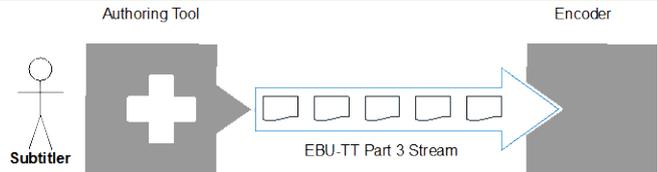
Figure 3: Example of sequence identifier and sequence numbers

The concept of sequence identification is separate to service identification. Document metadata is included to allow authors to identify the services for which the sequence is intended to carry subtitles, also known as the broadcast channel etc.

In EBU-TT Live streams are transferred between Nodes, which can consume, process and/or output Documents. Nodes can send or receive multiple Streams. Some examples are shown below.



When Nodes transfer sequences between each other, we call these transfers **Streams**.



Documents can use different types of timing. Only one Document can be active at any given time. If there is an overlap between the timing of different, the highest Sequence number 'wins'.

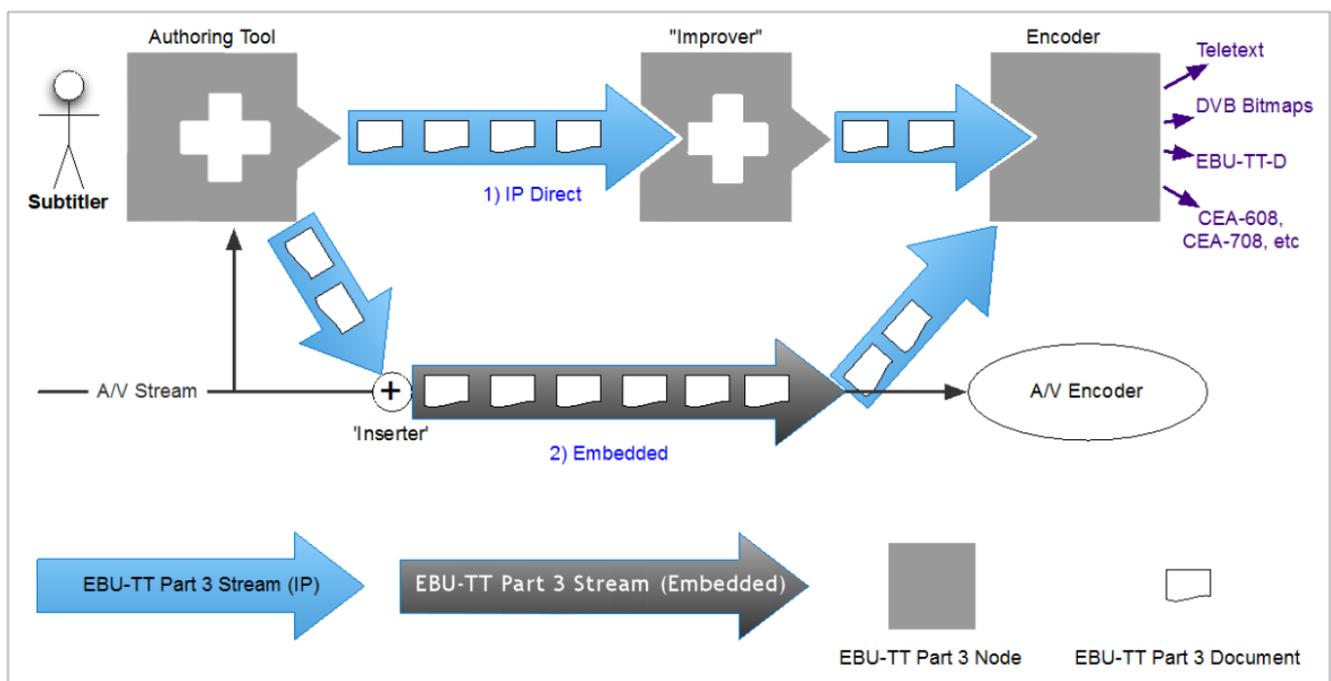
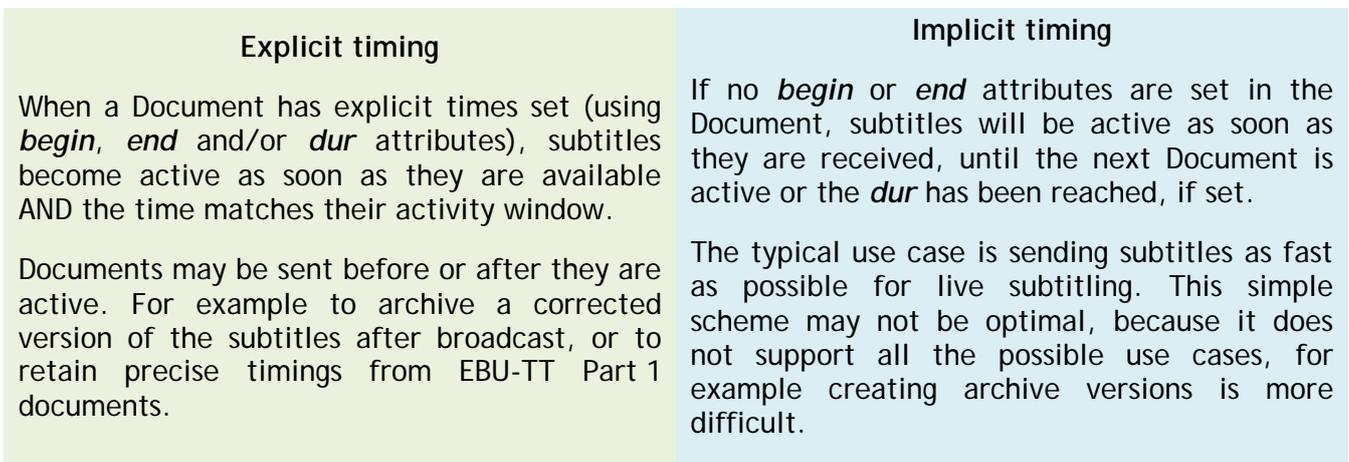


Figure 4: Schematic of simple use case showing an authoring tool generating a stream of EBU-TT Part 3 live subtitles.

Figure 4 illustrates a simple example use case in which a subtitler uses an authoring tool to create a stream of live subtitles. Those are then transferred either: 1) via a direct IP connection to an Improver and then on to an encoder; or 2) embedded into an audio/visual stream with an inserter and then de-embedded for passing to an encoder. A potential addition to this workflow would be an additional connection for example from the Improver to an Archiver to create an EBU-TT Part 1 document [EBUTT1] for later reuse.

The Improver defined in this document is a processing node that could, for example, insert a defined compensating delay, check content for correct spelling, ensure that prohibited code points are not propagated or perform other transformations to generate output suitable for encoding.

This document also provides useful options for the problem of mixing the play out of prepared subtitle documents and live subtitles, a problem that arises for all broadcast channels whose output is not either 100% pre-recorded or 100% live.

1.3 Example scenarios

The following examples represent typical real world scenarios in which documents and nodes that conform to this specification can be used, which are fully defined by this specification. Note that the scenario numbers do not increase monotonically: further example scenarios that are not fully defined by this specification are listed in Annex C, with which these form a set.

Example 1: Handover orchestration

Each subtitler in a group authors subtitles for part of a single programme; each group member takes a turn to contribute for a period of time before handing over to the next subtitler.

Each subtitler creates a distinct sequence of subtitles for their turn. Each input sequence has a different sequence identifier. The authoring stations emit the sequences as streams. As part of an externally orchestrated handover process a 'handover manager' node receives all the streams, combines them and emits a new continuous stream. This new output stream's sequence has a different sequence identifier from each of the input sequences.

Incidentally, each subtitler may subscribe to, and view the others' streams to assist with the handover orchestration.

Example 2: Author and correct

A pair of subtitlers authors and corrects live subtitles. The first subtitler creates a sequence using an authoring tool. The second subtitler receives a stream of that sequence and manipulates an Improver Node that allows the sequence to be modified and then issues a new sequence with a different sequence identifier from the input sequence, for consumption downstream.

Example 4: Timing improvement

An Improver Node receives a stream and a continuous audio track in a reference time frame. The Improver analyses the audio and subtitles and creates a new sequence whose contents are time-aligned relative to the audio track's time frame, using time stamps from a common clock source. The new sequence is issued as a stream with a new sequence identifier.

Example 10: Retrospective corrections

A subtitler authors a live subtitle stream whose sequence is archived for later reuse. On noticing an error the subtitler issues a retrospectively timed correction. The archive process uses data within the sequence to apply the correction such that the error it corrects is not apparent within the generated archive document.

1.4 Document Structure

This document is structured as follows:

Section 1 comprises this Introduction.

Section 2 defines terms and introduces core concepts.

Section 3 defines the requirements for documents that conform to this specification.

Section 4 defines the requirements for nodes that conform to this specification.

2. Definitions and Concepts

2.1 Definition of terms

Author

The term "author" describes a person or system that creates a stream of live subtitle data based on observations of some other media, for example by listening to a programme audio track.

Captions and subtitles

The term "captions" describes on screen text for use by deaf and hard of hearing audiences. Captions include indications of the speakers and relevant sound effects.

The term "subtitles" describes on screen text for translation purposes.

For easier reading only the term "subtitles" is used in this specification as the representation of captions and subtitles is identical here.

In this specification the term "captions" may be used interchangeably for the term "subtitles" (except where noted).

Document

A subtitle document conformant to this specification.

Document availability time

The time when a document becomes available for presentation.

Document resolved begin time

The time when a document becomes active during a presentation. Note: This term is used in the same sense as "resolved begin time" is used in [SMIL], when applied to a document.

Document resolved end time

The time when a document becomes inactive during a presentation. Note: This term is used in the same sense as "resolved end time" is used in [SMIL] when applied to a document.

Encoder

The term "encoder" refers to a system that receives a stream of live subtitle data and somehow encodes it into a format suitable for use downstream, for example EBU-TT-D. Note that some encoders may also package the encoded output data into other types of stream e.g. MPEG DASH.

Inserter

A unit that embeds subtitle data into an audio/visual stream. This is in common use in current subtitling architectures.

Node

A unit that creates, emits, receives or processes one or more sequences.

Node identifier

The unique identifier of a Node.

Presentation

In this document the term 'presentation' is used in the sense in which it is used in [SMIL].

Presentation Processor

This term is used as defined in [TTML1]: "A *Content Processor* which purpose is to layout, format, and render, i.e., to present, *Timed Text Markup Language* content by applying the presentation semantics defined in this specification."

Processing Context

The configuration and operating parameters of a node that processes a document.

Sequence

A set of related documents each of which shares the same sequence identifier, for example the documents that define the subtitles for a single programme.

Sequence Begin

The start of the interval in which a sequence is presented is referred to as the *sequence begin*. Equivalent to the document begin [SMIL] of the first document in the sequence.

Sequence End

The end of the interval in which a sequence is presented is referred to as the *sequence end*. Equivalent to the document end [SMIL] of the last document in the sequence.

Sequence Duration

The difference between the sequence end and the sequence begin is referred to as the *sequence duration*.

Service identifier

An identifier used to uniquely identify a broadcast service, for example the HD broadcast of the broadcaster's main channel.

Stream

The transfer of a sequence.

Logical stream

A stream offered or provided by a node to zero or more nodes; identified by the source node identifier and sequence identifier.

Physical stream

A stream provided between two nodes, identified by the source node identifier, destination node identifier and sequence identifier.

2.2 System Model

The following abstract system model is defined; systems that are conformant with this specification shall meet the requirements within the system model and the node conformance section (see § 4):

Documents

- A document is a single entity conformant to this specification.
- An implicitly timed document is one that contains no timing information and is considered active as soon as it has become available. It may also express a maximum duration after which it shall be deactivated.
- An explicitly timed document is one that contains timing information and whose activation time depends both on when it has become available and the timing information that it contains.
- When presenting a sequence of documents, at each moment in time exactly zero or one document shall be active.
- If no document is active, or if a document with no content is active, no content shall be displayed.

Sequences

- A *sequence* is a set of related documents, e.g. the documents that define the subtitles for a single programme.
- Sequences shall be considered distinct if they have had processing applied, even if the result of that processing is no change other than a known difference in state, for example if the processing has checked the spelling of the text content.
- Every distinct sequence shall have a unique *sequence identifier*.
- A sequence of implicitly timed documents may be transformed, with knowledge of the documents' associated timings, into a sequence of explicitly timed documents or a single document with explicit timings.
- A document shall be associated with exactly one sequence.
- Sequences do not have an explicit existence other than being the set of their constituent documents; the sequence identifier shall be present within every document in order to make concrete the association with the document's sequence.
- Documents with the same sequence identifier shall contain a *sequence number*. Every distinct document with the same sequence identifier shall have a different sequence number. Sequence numbers shall increase with the passage of time for each new document that is made available. Sequence numbers are used to resolve temporal overlaps: see § 2.3.1 below. Note: if sequence numbers begin at 1 then, at an example nominal mean rate of 1 document per second the maximum sequence number that will fit within a 4 byte unsigned integer corresponds to a sequence duration of over 136 years. Sequences should begin at low sequence numbers such as 1 to avoid the possibility of an unwanted integer overflow.
- Every document in a sequence shall be valid and self-contained. NOTE: in general no knowledge of other documents is required to process it¹².
- Sequences may be stored: that is, the lifetime of the sequence is unbounded.

Every document in a sequence shall have an identical timing model as defined by using the same values for the `ttp:timeBase` and `ttp:clockMode` attributes.

¹ Some specific kinds of processor may constitute exceptions, such as one that accumulates multiple documents together and combines them into a single one.

² For the purposes of efficiency, a processing node may store the preceding document for document comparison, e.g. when a subsequent document only changes the temporal validity of the document content.

Nodes and streams

- A node is an EBU-TT Part 3 aware unit or mechanism that creates, emits, receives or processes one or more sequences.
- Nodes are identified. NOTE: the format of node identifiers is not defined here.
- A node may offer a *logical stream* as the transfer of a sequence to one or more nodes. A logical stream is identified by the source node's identifier and the sequence identifier.
- A *physical stream* is the transfer of a sequence between two nodes. It is identified by the pair of identifiers of the source node and the destination node and by the sequence identifier.
- Any number of nodes may process or consume the same logical stream: by definition the delivery of those streams requires one physical stream per destination node. NOTE: a node can provide multiple physical streams of the same sequence.
- A *processing node* emits sequence(s) that are distinct³ from any of its input sequence(s). A creation node is a specialised processing node with zero input sequences.
- A *passive node* shall NOT modify input sequences and shall only emit sequences that are identical (including the sequence numbers) to the input sequence(s), for example nodes that are simply used for switching. A consumer node is a specialised passive node that does not emit any sequence.
- Streams are transient: that is, the lifetime of a stream is bounded by the period between the start of transfer (when the first document is transferred) and the end of transfer (when the last document is transferred). Streams may begin before the last document in the sequence has been generated - indeed it is envisaged that in normal operation documents within a sequence are generated dynamically and transferred in a stream with an undefined end time.
- The flow of a stream is unidirectional. Any 'back channel' reverse communication is external to the payload of the sequence⁴.
- At any moment in the presentation of a sequence by a node exactly zero or one document shall be temporally active. Note the logic defining which document is temporally active is defined in § 2.3.1 below.

Nodes are defined as abstract classes. Further detail of node behaviour is defined in § 4. See also Figure 11.

Figure 5 shows a UML class model illustrating the logical entities in this system model.

³ There may be little or no observable difference in content but the processing node may have the potential to modify the output stream; for example a spell-checker or profanity removal node may not make any changes most of the time but be called into action on an occasional basis. Nevertheless the output is considered to be different from the input because it has a logically different state, e.g. it is known not to contain any profanities, and it has a different sequence identifier.

⁴ Nodes can subscribe to multiple streams - see the Handover Manager node for example.

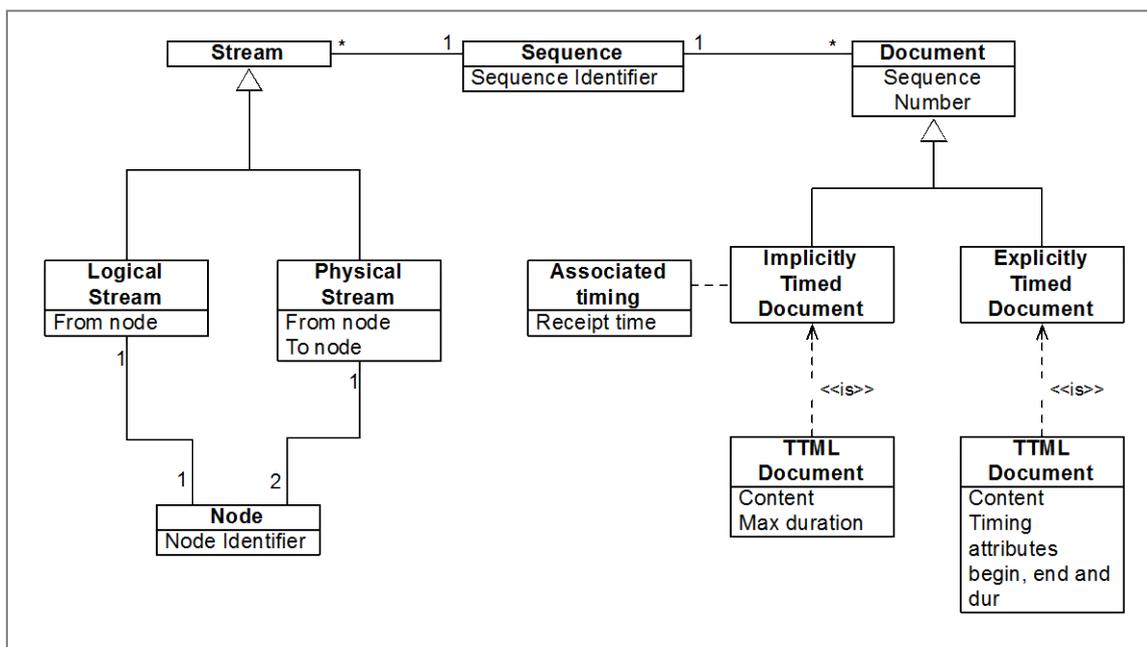


Figure 5: UML model showing system model logical entities

It is envisaged that some processing nodes will be purely software whereas others may require human intervention.

Sequence identifiers should not be used as service identifiers, such as broadcast television channel identifiers. For example consider the case of a 'simulcast' where a single video source is broadcast on more than one output service simultaneously, e.g. to support localised services or SD/HD services. A subtitler can only observe a single media source at any moment, and generates a single sequence with a single sequence identifier, which could be subsequently processed, generating new sequences with different identifiers. Those sequences could be destined to be encoded for multiple output or destination services. Destination service identifiers may be carried as metadata within documents. Similarly the observed source service identifier may be carried as metadata within documents.

2.3 Timing and synchronisation

This section defines the temporal processing of a sequence of documents within a presentation, the management of delay in a live authoring environment and the use of reference clocks.

2.3.1 Document resolved begin and end times

Every document in a sequence has a time period during which it is active within a presentation defined in TTML1 as the Root Temporal Extent. At any single moment in time during the presentation of a sequence either zero documents or one document shall be active. The period during which a document is active begins at the document resolved begin time and ends at the document resolved end time.

An algorithm for computing the time when a document is active during a presentation has the following variables to consider:

- The document availability time, i.e. the time when each document becomes available. In a live authoring scenario it is typical for documents to become available during the presentation, soon after they have been authored. In a replay scenario all of the documents may be available prior to the presentation beginning.
- The earliest computed begin time in the document, calculated from the `begin` and `end` attribute values if present, according to the semantics of [TTML1].
- The latest computed end time in the document, calculated from the `begin` and `end` attribute values if present, according to the semantics of [TTML1].

- The value of the `dur` attribute if present.
- The document sequence number.
- Any externally specified document (de)activation times, such as the beginning of sequence presentation.

The definitions of the document resolved begin and end times below are derived from the following rules:

1. A document cannot become active until it is available, at the earliest.
2. The absence of any begin time implies that the document is active immediately;
3. The absence of any end time implies that the document remains active indefinitely;
4. The `dur` attribute on the `body` element imposes a maximum document duration relative to the point at which it became active.
5. If two documents would otherwise overlap temporally, the document with the greater sequence number supersedes the document with the lesser sequence number.

Note: the `timeContainer` attribute cannot be specified in EBU-TT Part 3 documents so `par` semantics apply as specified in [TTML1] § 10.2.4.

Note: It is not necessary in all processing contexts to resolve the document begin and end times. For example a processing node that checks text spelling only can do so without reference to the timing constructs defined in this section.

Note: In general the document time base, as specified by the `ttp:timeBase` attribute, can be transformed within a processing context to a value that can be compared with externally specified document activation times, for example the clock values when documents become available. Implementations that process documents that have `ttp:markerMode="discontinuous"` cannot assume the time values to be part of a monotonically increasing clock, but only as markers. This scenario is common with timecodes, i.e. when `ttp:timeBase="smpte"`. In this case a presentation node has to derive comparable timing events from some other source, such as time code embedded in some related media.

Note: the `dur` attribute is not permitted in documents that have `ttp:markerMode="discontinuous"`.

2.3.1.1 Document resolved begin time

The document resolved begin time shall be the later of (a) the document availability time, (b) the earliest computed begin time in the document and (c) any externally specified document activation begin time, such as the beginning of sequence presentation.

2.3.1.2 Document resolved end time

The document resolved end time shall be the earlier of (a) the earliest document resolved begin time of all available documents in the sequence with a greater sequence number, (b) the document resolved begin time plus the value of the `dur` attribute, if present, (c) the latest computed end time in the document and (d) any externally specified document deactivation time, such as the end of sequence presentation.

2.3.1.3 Behaviour when no document is active

When no document is active a presentation processor shall NOT render any content.

Note: An encoder can be considered to be a specialised type of presentation processor. More generally, this applies to any consumer node.

2.3.1.4 Document creation and issuing strategies (informative)

The above rules allow for implementations to use different strategies for creating and issuing sequences. It is beyond the scope of this document to recommend the appropriate strategy for every circumstance. This section describes some of the strategies that could be employed.

The concepts of implicitly timed and explicitly timed documents, and their potential usage, are described. Further, a selection of document issuing strategies is described; this is not intended to encompass all of the possible strategies.

2.3.1.4.1 Implicitly timed documents

Documents that do not contain any `begin` or `end` attributes are described as having implicit timing (see [SMIL] § 10.7.1). Such documents can contain a maximum duration, set as a `dur` attribute on the `body` element.

The absence of a `begin` time implies that the content is active as soon as the document becomes active. If neither an end time nor a duration is specified then the content will be active forever, or until the document becomes inactive, for example because a different document is activated in its place. See rule 5 in § 2.3.1 above.

Figure 6 shows a diagrammatic representation of the resolved begin and end times of implicitly timed documents.

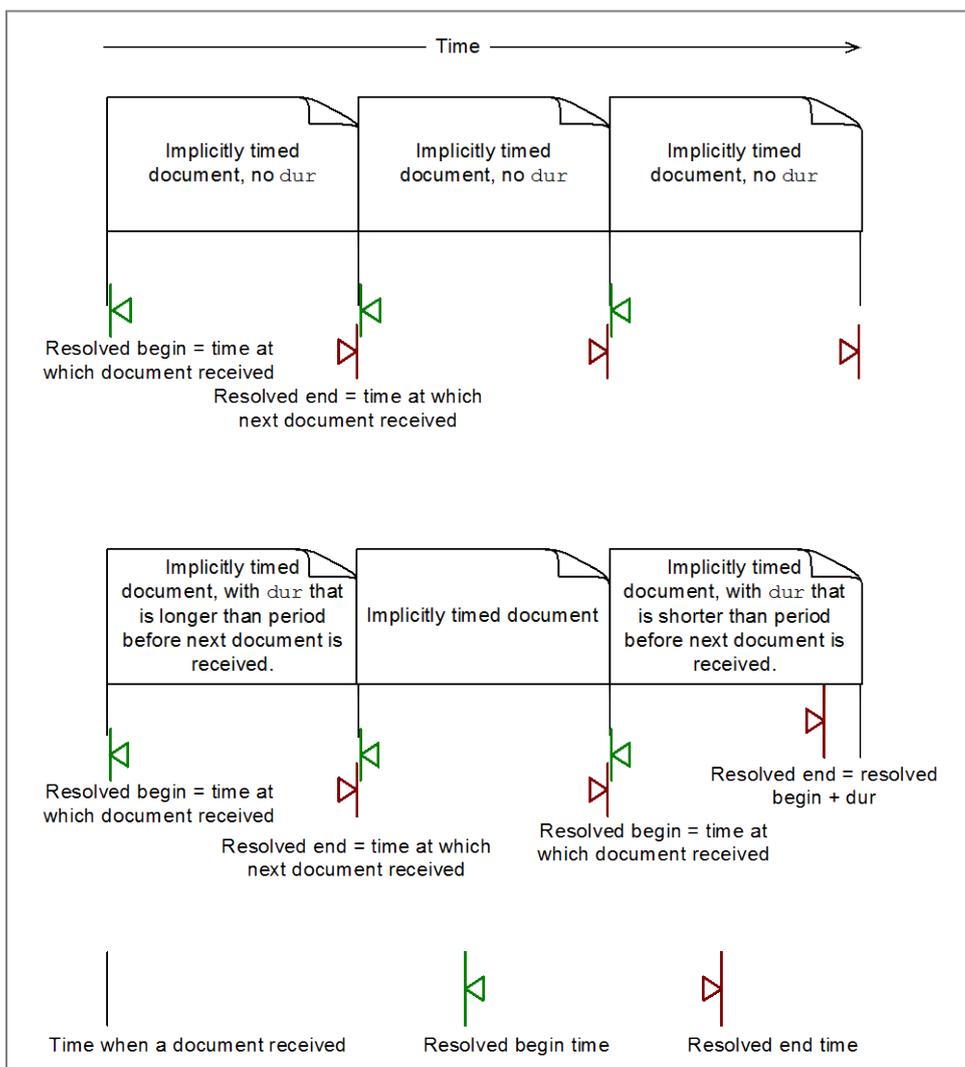


Figure 6: Diagram showing resolved begin and end times of implicitly timed documents with and without `dur` attributes

Use of documents with implicit timing restricts the available issuing strategies because it is not possible to include timed changes within them. In other words every change to the displayed content can only be contained in a new document.

It is neither possible to provide implicitly timed documents in advance of their required activation nor retrospective documents that revise earlier documents in the sequence since no time information is present within the documents to indicate such intentions.

A use case for implicitly timed documents is one where the timing for the document is defined by association with a specific frame of video or set of frames, where that association is inherent in the carriage mechanism, for example because the document is embedded in the VANC data area of an HD-SDI video stream, or in an ISOBMFF sample.

A second use case for implicitly timed documents is one where the timing does not need to be resolved with reference to a clock or other timed data at all, for example because the sequence is streamed and the recipient of the stream treats the availability of each new document in the sequence as a 'do it now' event and presents it as soon as possible. In this case the availability time could be captured from some clock source and noted alongside each document for later re-use if desired.

A document containing an empty `body` element can be used to indicate that the presentation is to be cleared while that document is active.

If implicitly timed documents are dissociated from the context that defines the timing of their presentation, the resulting sequence cannot reasonably be presented.

2.3.1.4.2 Explicitly timed documents

Documents that contain a combination of `begin` and `end` attributes are described as having *explicit timing* (see [SMIL] § 10.7.1). Such documents can additionally contain a maximum duration, set as a `dur` attribute on the `body` element.

See Figure 7 for a diagrammatic representation of the resolved begin and end times of explicitly timed documents.

Use of documents with explicit timing requires that any presentation processor that needs to resolve the timing, such as an encoder, needs to be able to compute the local real time that applies in its processing context for any given time expression. Where the marker mode indicates that an event model is being used, by the `ttp:markerMode` attribute having the value "discontinuous", the presentation processor needs to be able to derive comparable timing events from some other source, such as time code embedded in some related media.

A further implication is that the time values present in the document need to be valid and reasonable in the context of use. If for example an independent reference clock is used to match document times to encoding presentation times, then the presentation processor cannot proceed without access to that reference clock.

A use case for explicitly timed documents is one where precise timing is known at document creation time. For example the document could have been created by a synchronisation process that calculates the correct alignment of the subtitles with the audio and stores that alignment within the document.

A second use case for explicitly timed documents is one where the relationship between the media time and the subtitle time needs to be stored entirely within the documents, for example because there is no embedding context from which to derive times. If there is an embedding context that imposes times, explicitly timed documents can still be used according to the document activation rules defined above.

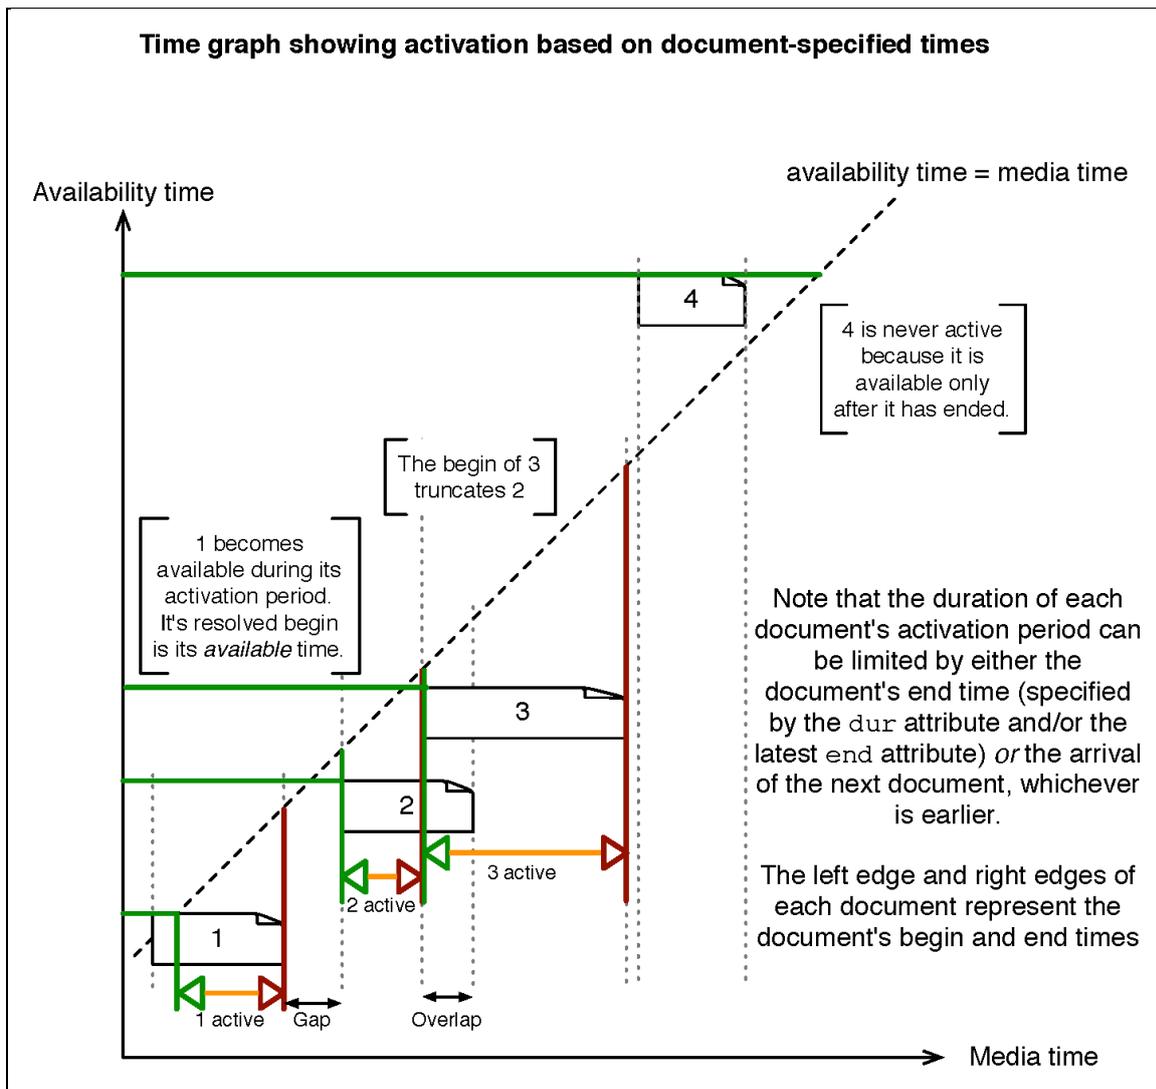


Figure 7: Diagram showing resolved begin and end times for explicitly timed documents

A third use case for explicitly timed documents is one where it is acceptable or necessary to accumulate a set of timed changes over a defined period and deliver them in a single document. For example there could be sufficient time available in the encoding and distribution system that this is considered acceptable, or if the fastest data transfer rate is limited this could be a scheme that in some cases helps to reduce the average required rate. Even if the timings are not precise absolutely, they can be precise in a relative sense within the document. For example if the time of each word issued by a subtitle author is captured in the document the relative timing per word may be precise, even though there is an overall offset for the whole document relative to the media being subtitled that may be imprecise and indeed variable.

The timings within an explicitly timed document can be used to apply retrospective changes so that the re-presentation of the sequence when all the documents are available differs from the original presentation.

The timings within an explicitly timed document can be used to signal future content if it is known. For example if a news programme contains a video package for which subtitles have been authored in advance a single document could be issued containing all of the subtitles for the package.

It is possible for explicitly timed documents to be made available before they become active. In this case it is expected that processing nodes will not discard such documents until they have become inactive, at the earliest.

It is possible to create sequences of explicitly timed documents that contain periods in which no document is active. This technique can be used to indicate that the presentation is to be cleared.

This is distinct from a sequence of implicitly timed documents where an empty document is in general required to achieve the same effect.

A second option for indicating that the presentation is to be cleared is to issue a document with appropriate `begin`, `end` or `dur` attributes on the body element that is otherwise empty of content.

2.3.1.4.2.1 Cumulative documents

It is possible to create a sequence of explicitly timed documents in which each document contains all of the content from the previous document in addition to any new content. This could be described as a cumulative document creation technique. See Figure 8, overleaf.

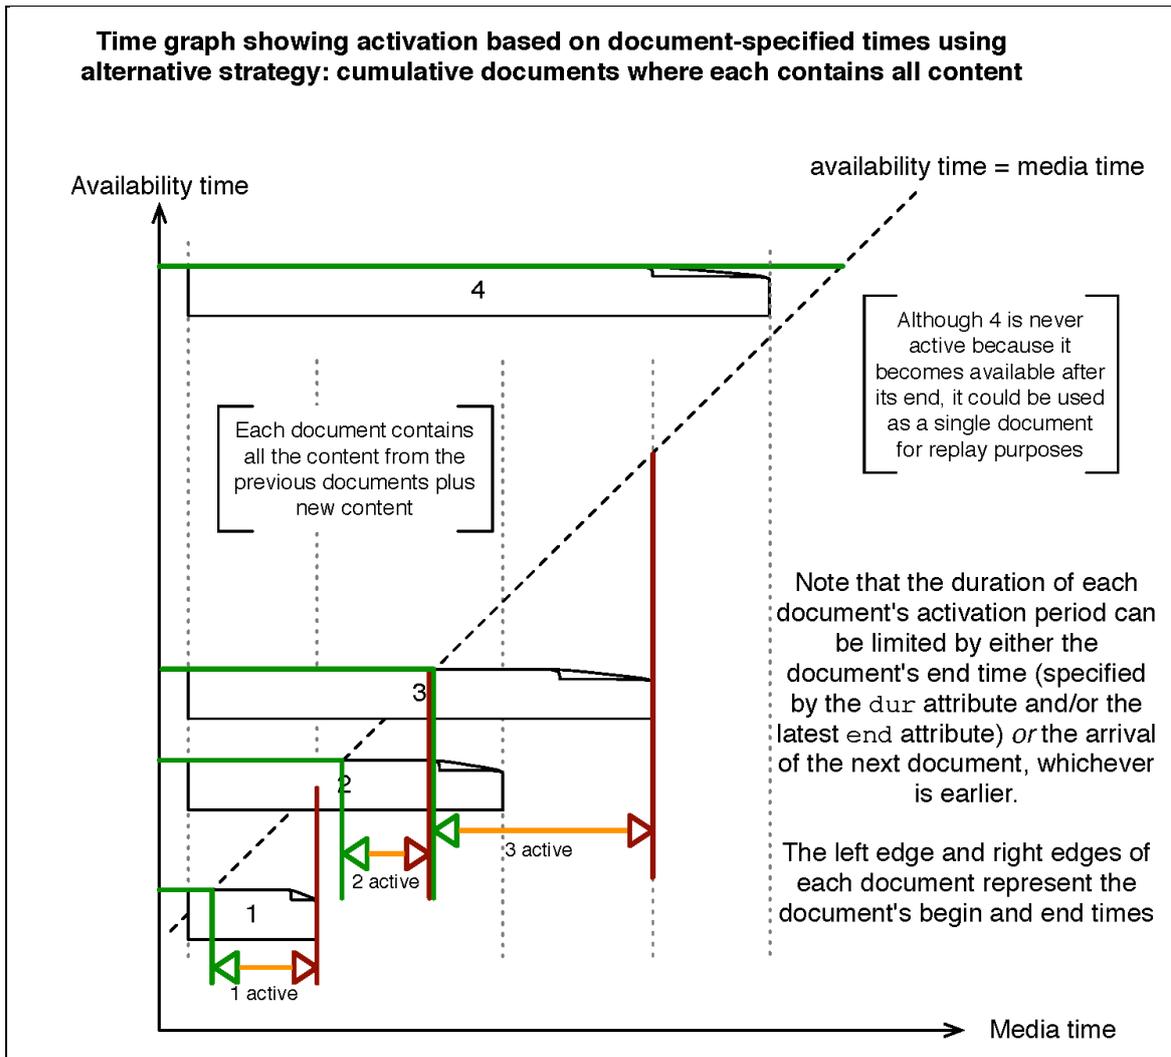


Figure 8: Timing graph showing cumulative document approach

This technique allows the entire sequence to be presented simply by activating the last document, much as though it were a prepared subtitle document. In real world implementations using this technique it could be wise to limit the duration applicable to the sequence to limit document sizes and parsing time.

2.3.1.4.3 Mix of explicitly timed documents and implicitly timed documents

It is possible to create sequences that contain a mixture of implicitly timed and explicitly timed documents.

In this scenario implementations need to relate any external time values to the time base used within the documents. One approach is to store the document receipt timestamps using the same time base as that used within the document.

If the presentation processor does not have access to document availability times or some other context that constrains the activation period of each document then every implicit document effectively prunes all of the documents with a lower sequence number, which by definition are never activated if the sequence is re-presented.

See Figure 9 for an illustration of how mixed documents are activated.

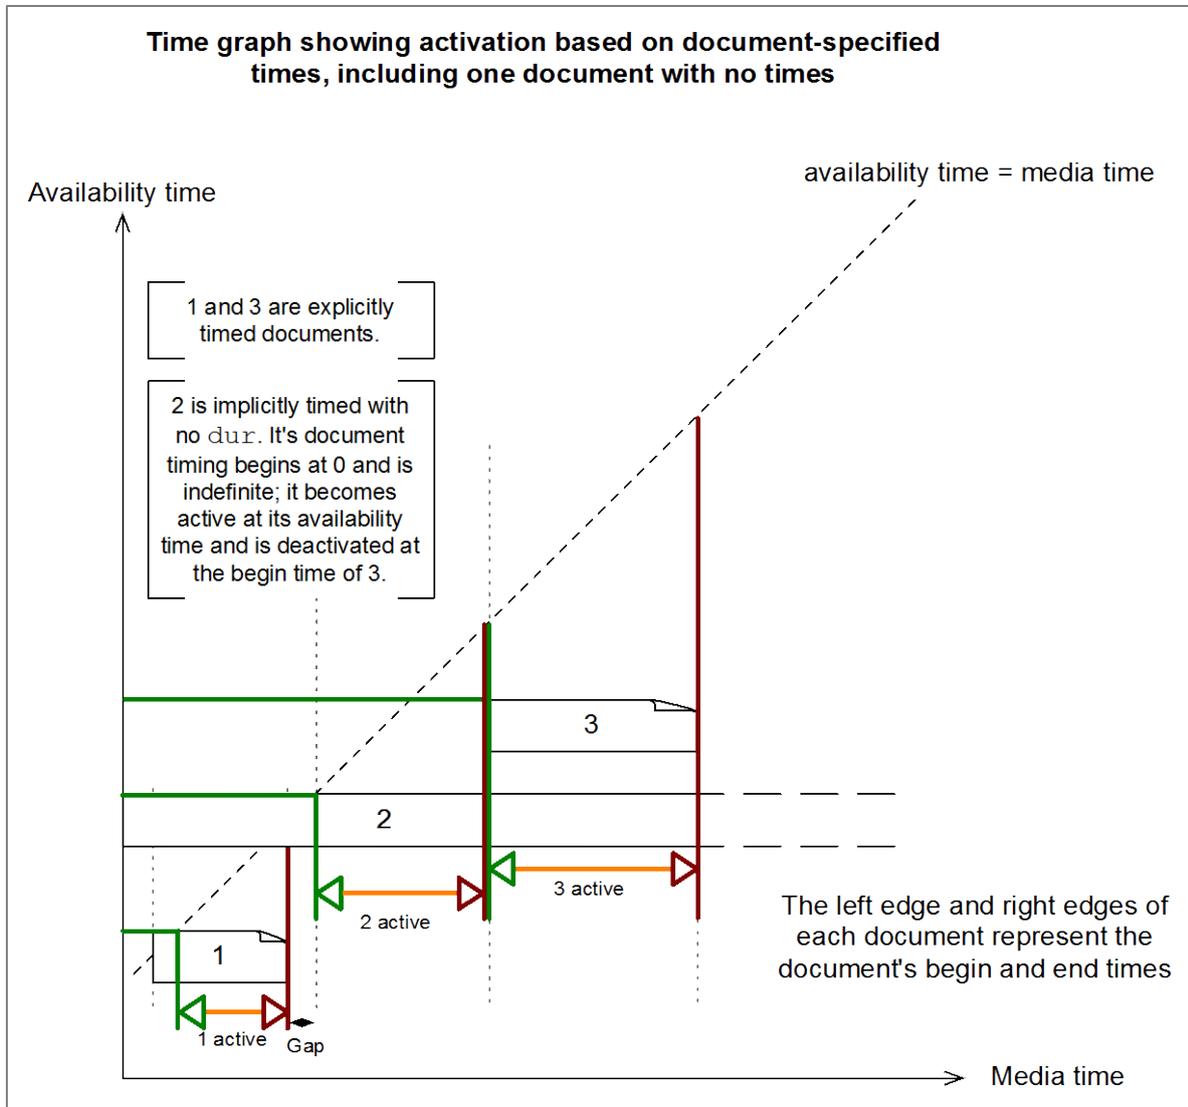


Figure 9: Timing graph showing document activation when a mix of implicitly and explicitly timed documents is in the same sequence

2.3.1.4.4 Issuing documents on every change

One strategy for issuing documents is to create a new document for every change needed to the presented content. This can allow changes to be communicated as quickly as possible without using any more data than is required. Any consumer node subscribing to the stream needs to wait until the next change before seeing any content to present; that in turn could have a negative impact on for example the recovery time if the encoder had entered a fault state before restarting its subscription.

2.3.1.4.5 Issuing documents periodically

An alternative strategy for issuing documents is to create a new document at regular intervals containing the timed content since the previous document. The interval defines the longest delay in communicating changes, additional to any other authoring delays. Conversely it also defines the maximum time until a new subscriber can begin to process content.

In extremis the interval would be the smallest unit of time available. For example a new document could be created and associated with every frame of video.

If the desired behaviour is for the presentation to change (e.g. new subtitles appearing) at times that are not coincident with the interval boundaries, then explicitly timed documents would be needed since by definition there is no way to represent such changes within implicitly timed documents.

In the case of prepared subtitle content the interval could be the duration of the prepared content.

2.3.1.4.6 Issuing documents on every change with a minimum repeat duration

A blend of the above two strategies could be useful to control data rates and set limits on recovery times in the event of source switching or other fault recovery. Such a blended strategy would be one in which a new document is created on every change, but if a predefined interval expires during which no change occurs, a new document is created regardless.

2.3.2 Management of delay in a live authoring environment

Within a typical broadcast workflow there are a variety of delays. These are imposed by the need to perform processing of content, for example transcoding, or by the need to synchronise media relative to other media or defined time structures such as frame boundaries, GOP structures etc. Many of these delays are significant, for example in an MPEG-DASH distribution model the video needs to be accumulated into segments of a predefined duration, encoded, packaged and distributed via a content delivery network.

In the context of live created subtitles, there is a variable authoring delay⁵, which is the time between the subtitler hearing the audio and the authoring tool outputting the relevant subtitles. In some environments there are unavoidable delays in the chain, for example for video processing, which are long enough to permit some re-alignment of subtitles relative to the audiovisual media, so that the combined output as experienced by the audience has reduced or even zero latency. See Figures 10 and 11.

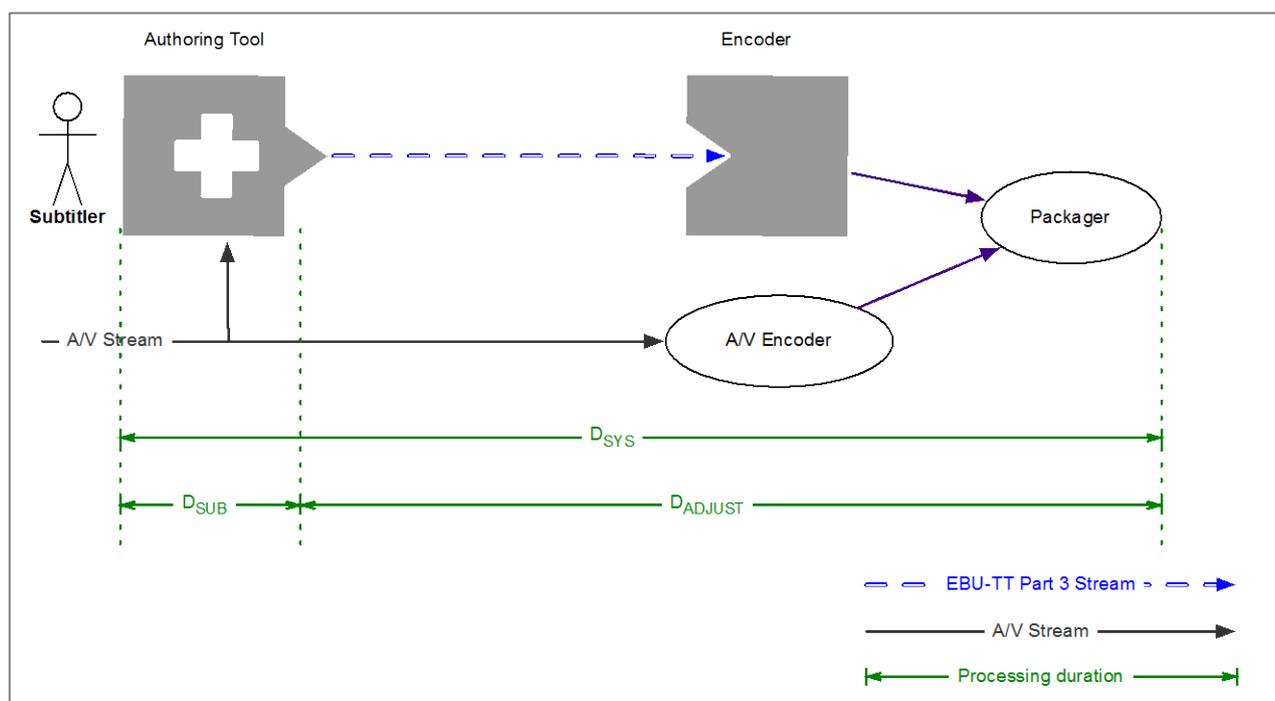


Figure 10: Use case diagram showing timing and delay concepts

⁵ Real world experiments in the UK suggest this can be around 5 - 10 seconds.

The likely size of this authoring delay depends on the technique used to create the subtitles and the author's speed. For example the delay might be longer for subtitles created using live respeaking or stenography than for subtitles prepared in advance and cued manually or automatically.

Additionally, any propagation or processing delays introduced by nodes in the chain need to be taken into account, and could in general result in the resolved document begin times being later than desired or intended.

If the total required delay is D_{SYS} , being the time between the A/V signal being observed by the subtitler and the desired moment of emission or availability of that A/V signal, and the subtitle authoring delay is D_{SUB} , then it follows that for correct timing an adjustment delay D_{ADJUST} must be imposed within the subtitle chain such that $D_{SUB} + D_{ADJUST} = D_{SYS}$.

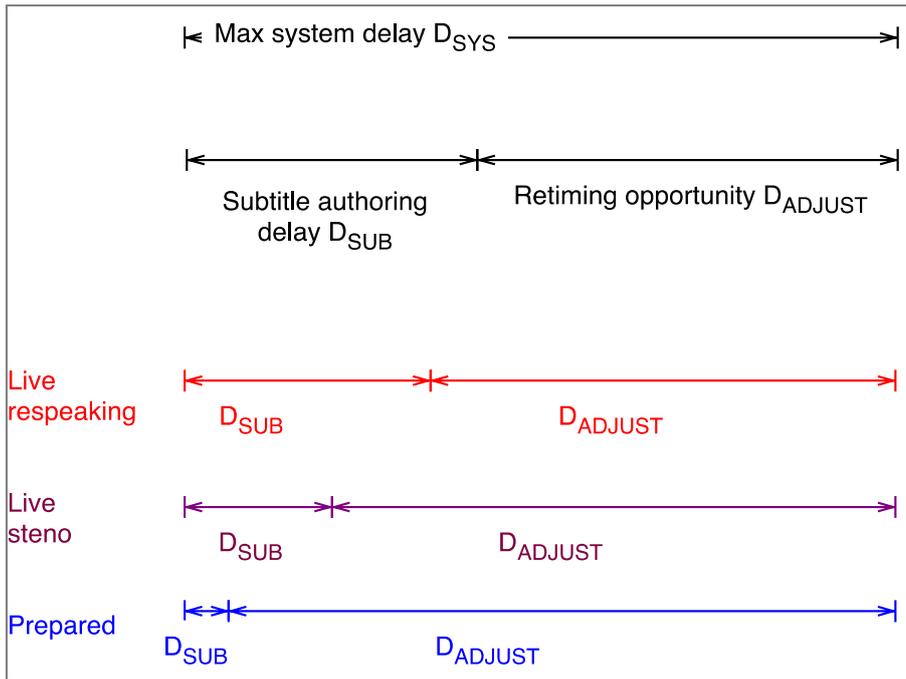


Figure 11: Opportunity for adjustable retiming of subtitles

The minimum value of D_{ADJUST} within a real world live authoring environment is 0⁶ because when D_{SUB} exceeds D_{SYS} no correction is possible. This indicates that there is a window of opportunity of duration $D_{SYS} - D_{SUB}$ within which the D_{ADJUST} may be varied to achieve correct synchronisation. See Figure 11. Note that in some environments a solution for frequently late subtitle presentation could be to add a delay to increase the value of D_{SYS} ; in scenarios where encoders are able to output accurate subtitle timings further synchronisation options could be available downstream, e.g. in client devices - such specification is out of scope of this document.

2.3.3 ebuttm:authoringDelay

The `ebuttm:authoringDelay` metadata attribute, applied to `tt:tt`, may be used to capture the equivalent of D_{SUB} , where it is known or can be estimated.

The value of `ebuttm:authoringDelay` is intended to express the latency associated with the authoring process and has no direct correlation to the timing values of a document. Specifically this means that even if a process applies a delay (e.g. a delay node), the information that is carried in `ebuttm:authoringDelay` is left unchanged.

⁶ It is conceivable that a Delay node that emits explicitly timed documents could apply a negative D_{ADJUST} offset; an encoder receiving documents that have already expired could reasonably discard them with no output; however this technique could be useful for creating correctly timed *archive* documents.

The time taken (as measured from the original spoken word) for a given subtitle to be transmitted from an Producer is a composite of the time taken to author (if the workflow is live) and the time taken for the device, application or operator to release the subtitle.

The combined authoring interval is typically:

- sub-frame for an automated prepared file play-out process;
- the operator's reaction time for live-cued subtitles;
- the 'listening and thinking' time of the subtitler plus the subtitling applications' processing time for stenography or re-speaking in live origination.

The authoring time is a variable for any human-driven process. It can vary in-session for re-speaking and stenography as a function of operator performance, media content type, language etc. The calculation or estimation of D_{SUB} can therefore facilitate an accurate 'improvement' process; it can also enable any downstream Improver Nodes to modify the timing applied to automatically or manually cued prepared subtitles.

Note: the value of D_{SUB} is likely to be higher for processes involving more human cognitive processes and lower for more automated processes; in the case that there is high certainty of upcoming subtitles they may even be sent in advance, resulting in a negative value. In the cases of a very low value Authoring Delay it is probable that any Improver will not try and improve the alignment of the subtitles - the low value signifies that the subtitles are already pretty close in time alignment.

Note: a small or negative value of `ebuttm:authoringDelay` can be considered to have greater associated confidence than a larger positive value. An Improver Node intended to resynchronise the subtitles could use this proxy for a confidence indicator as a heuristic.

2.3.4 Delay nodes

An Improver Node could apply an adjustment delay to the subtitles as one part of a solution for achieving resynchronisation. We will refer to such a node as a Delay node. The value of the delay could be derived using one of a variety of techniques, including potentially:

- heuristics based on `ebuttm:authoringDelay`, and other metadata in the stream e.g. the method used to author the subtitles etc.;
- knowledge of the typical delay within the broadcast chain;
- a comparative measurement using audio analysis to establish a varying value adjustment delay based on the actual subtitles and speech within the programme content.

Note: It is out of scope of this document to mandate the use of specific techniques; furthermore it is expected that the relative success of these techniques will depend on programme content, the level of variability in the chain and the quality of implementation of each technique.

Note: any node that receives and emits streams is likely to incur some real world processing delay; a Delay node is intended to apply a controlled adjustment delay.

From a stream and sequence perspective, the following behaviours of a Delay node are defined:

1. A Delay node is a processing node. Therefore the output sequence shall have a different sequence identifier from the input sequence.
2. A Delay node may emit explicitly timed documents within a sequence. In this case the Delay node should not delay emission of the stream. NOTE: in this scenario it is recommended that when `ttp:timeBase="clock"` the time values within the documents relate to a reference clock that is available to downstream nodes for example consumer nodes.

3. A Delay node may emit implicitly timed documents within a sequence. In this case the Delay node shall delay emission of the stream by a period equivalent to the adjustment value.
4. A Delay node shall NOT emit an output sequence with reordered subtitles. NOTE: it is possible that delay functionality is combined with other processing in a single node, for example an accumulator; hence this requirement is expressed in terms of subtitles not documents: there is for example no requirement that there is a 1:1 relationship between input and output documents from a Delay node, though such a relationship would be expected for the simplest conceivable Delay.
5. A Delay node shall NOT update the value of `ebuttm:authoringDelay`.
6. A Delay node should add an `ebuttm:trace` element to the document metadata to indicate that the delay has been added.

Note: if an authored sequence is intended for multiple services, care should be taken in case those services have differing D_{SYS} values, for example because one is encoded as SD video and another is HD, and their encoders have different latencies.

2.3.5 Reference clocks

Some broadcast environments do not relate time expressions to a real world clock such as UTC but to some other generic reference clock such as a studio timecode generator. When `ttp:timebase="clock"` is used and `ttp:clockMode="local"`, or when `ttp:timeBase="smpte"` is used, the `ebuttp:referenceClockIdentifier` parameter may be specified on the `tt:tt` element to identify the source of this reference clock to allow for correct synchronisation.

Note that for real time processing of EBU-TT Part 3 documents correct dereferencing of the external clock is a processing requirement, therefore the `referenceClockIdentifier` is defined as a parameter attribute in the `ebuttp:` parameter namespace. This is in contrast to an [EBUTT1] document where the `referenceClockIdentifier` attribute is defined in the `ebuttm:` metadata namespace. A Part 1 document created as an archive version of a sequence of Part 3 documents may preserve the value of the `ebuttp:referenceClockIdentifier` parameter as a `ebuttm:referenceClockIdentifier` attribute.

2.4 Handover

In a live subtitle authoring environment it is common practice for multiple subtitlers to collaborate with each other in the creation of subtitles for a single programme. From an encoder perspective, it is desirable to manage only a single stream of live subtitles. To mediate between the streams that each subtitler creates we will refer to a *Handover Manager* node. See Figure 12.

The Handover Manager subscribes to a set of sequences and selects documents from one sequence at a time, switching between sequences dependent on parameters within the documents. It then emits a new sequence of documents representing the time-interleaved combination of subtitles from each of the authors. See also Figure 13 for an example of handover sequences.

The Handover Manager node shall use a 'who claimed control most recently' algorithm for selecting the sequence, based on a control token parameter within each document.

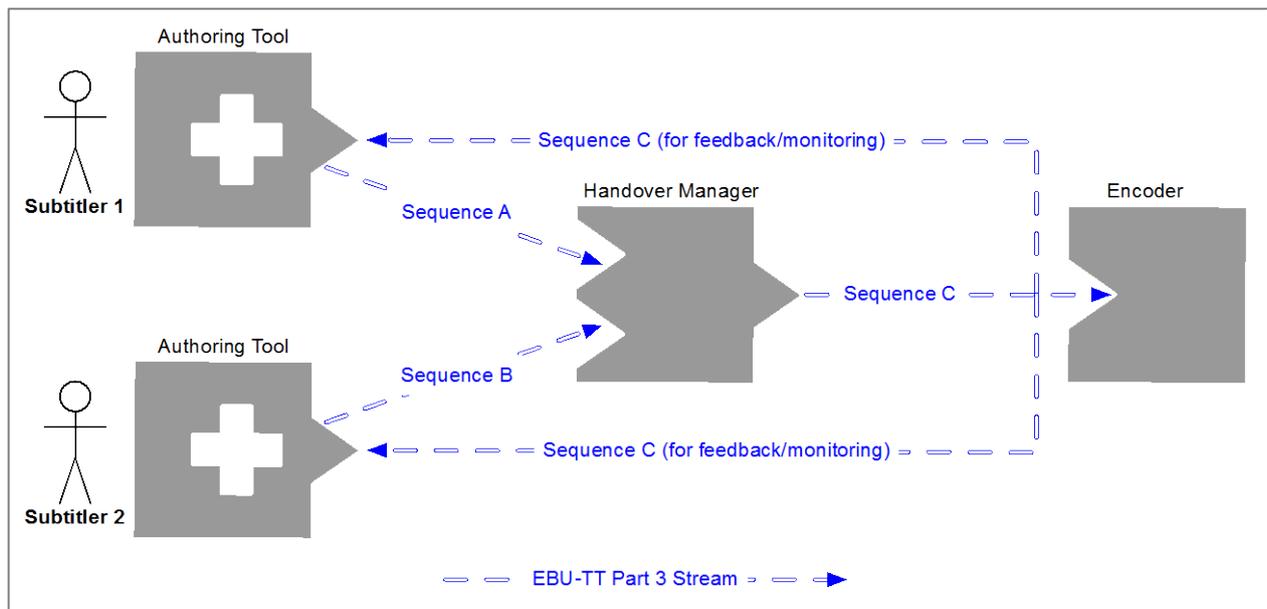


Figure 12: Use case showing a Handover Manager selecting between Sequences A and B and emitting Sequence C

Authoring tools can subscribe to the output stream from the Handover Manager; this makes the control token parameter values visible to them to permit each to direct the Handover Manager to switch to their output; it also facilitates monitoring.

NOTE: other schemes for directing handover are possible, for example the control token could be derived from a separate mediation source or even the clock.

2.4.1 Authors Group parameters

The following parameters on the `tt:tt` element are provided to facilitate handover:

The `ebuttp:authorsGroupIdentifier` is a string that identifies the group of authors from which a particular Handover Manager can choose. A Handover Manager should be configured to subscribe to all the streams whose documents have the same `ebuttp:authorsGroupIdentifier` except for its own output stream. All documents within a sequence that contain the element `ebuttp:authorsGroupIdentifier` shall have the same `ebuttp:authorsGroupIdentifier`.

The Handover Manager may include the parameters `ebuttp:authorsGroupIdentifier`, `ebuttp:authorsGroupControlToken` and `ebuttm:authorsGroupControlRequest` from the currently selected input sequence within its output sequence. This is so that each subtitle authoring station can know the current status, if subscribed to the Handover Manager's output stream.

When present in a document, the `ebuttp:authorsGroupControlToken` is a number that the Handover Manager uses to identify which sequence to select: when a document is received with a higher value `ebuttp:authorsGroupControlToken` than that most recently received in the currently selected sequence the Handover Manager shall switch to that document's sequence, i.e. it shall emit a document in its output sequence corresponding to the received document with the new control token without delay.

Note: The above definition permits the control token value to be lowered after taking control, either by setting the control token value to an explicit lower number or by omitting it in subsequent documents and relying on the initial value of "0".

Note: Care should be taken if the carriage mechanism does not guarantee delivery of every document in the sequence in case a document intended to take control is lost. One strategy for

avoiding this would be for the subtitle authoring station to observe the Handover Manager's output and verify that control has been taken before lowering the control token value. Another strategy would be to maintain the high control token value and duplicate it in each document in the sequence until the sequence switch has been verified through another mechanism.

The `ebuttm:authorsGroupControlRequest` may be used to indicate to other subtitle authors within the same group the current author's intent, i.e. if they are continuing to author, or are requesting that another author should take over control.

Note: The `ebuttm:authorsGroupControlRequest` syntax is undefined. It could include an indication that the current subtitler wishes to continue, wishes for an immediate takeover, or wishes for a takeover at some future point such as an advertising break, a sentence etc. How implementations manage a smooth handover is beyond the scope of this specification and could require out of band communication, training or other operational processes.

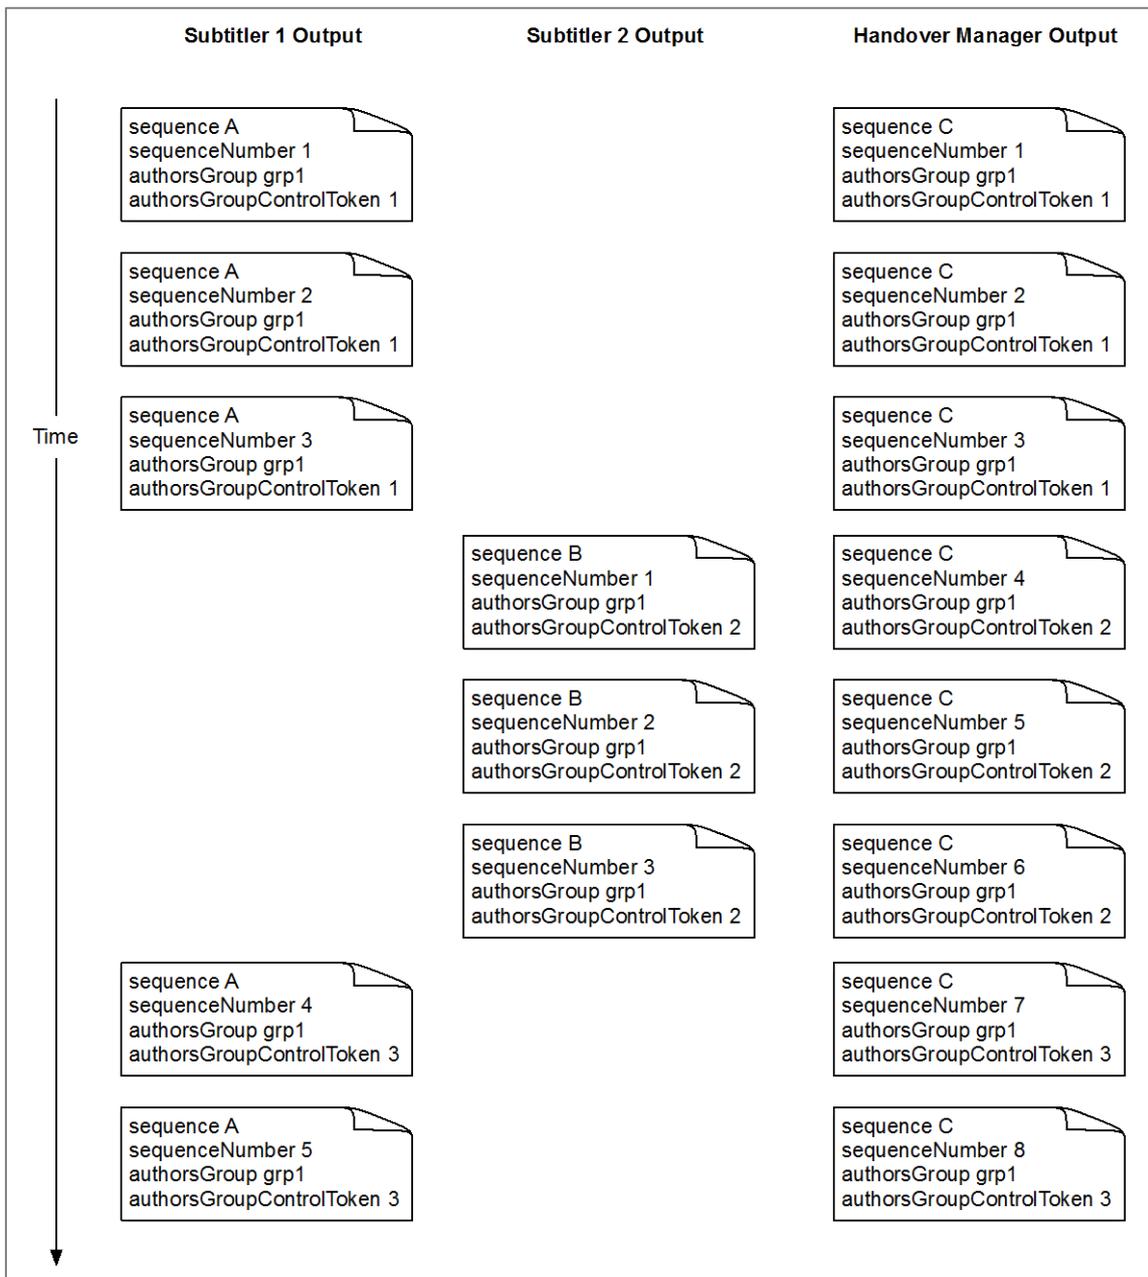


Figure 13: Sample sequences demonstrating Handover (omitting `ebuttp:authorsGroupControlRequest` for simplicity)

2.5 Describing facets of the subtitle content

In a chain of processing nodes each might modify the subtitle content to suit some particular need. For example the subtitler's priority could be to issue new documents as quickly as possible without considering spelling, grammatical correction, profanities etc. Alternatively the subtitler could be issuing a sequence that is known not to be suitable without modification for all downstream platforms: one encoder could be able to emit Unicode code points; another could be restricted to ISO 8859-x. A combination of these scenarios is possible. In order to complete the processing needed to ensure that a stream is suitable, the system model presented here proposes that a series of Improver nodes would be used to perform the necessary processing.

However from a Consumer node's perspective it is not always desirable to rely on the node configuration being correct without further information. Additionally for compliance monitoring automated processes could be used to assess conformance against some rule set without necessarily enforcing it or making modifications to the text content in the sequence.

Some knowledge of the processing applied can be obtained at a document level from the `ebuttm:trace` element (see § 2.6) or the `ebuttm:appliedProcessing` element (see [EBUTT1]) however this is coarse-grained. To indicate for a particular piece of content some aspect of its editorial or technical quality, for example if it has been spellchecked, profanity checked, had its code point set reduced for downstream compatibility, had colours applied, been positioned etc. the `ebuttm:facet` element can be applied.

The `ebuttm:facet` element may be added as one or more children of the `tt:metadata` element associated with any content, and expresses via the `expresses` attribute whether that content *has* ("has") or *does not have* ("has_not") the facet indicated by a term in a classification scheme or if it is *unknown* if the content has the facet ("unknown"). It is therefore extensible; no classification schemes are listed in this specification.

When multiple facets are applied to the same piece of content they are considered to be orthogonal and additive; it is an error to apply more than one facet with the same term identifier to the same content element, where the term identifier is the combination of the term name and the `link` attribute's value.

Facets are inherited, but may be overridden in a descendant element. For example, a document that was emitted by a spelling checker Improver node with the *has* "spell-checked" facet applied to the `tt:body` element could be edited further, resulting in an ad hoc insertion of new text, say in a `span` element. That span element could reasonably have the *has_not* "spell-checked" facet applied to override the inherited value.

This inheritance rule implies that implementations need to take care not to assume that a facet applied to a content element applies also to all of its descendants, since any one of those descendants can specify the alternate sense of the `expresses` attribute.

The `ebuttm:facet` element should contain either a text label (as the element contents) or a `link` attribute to link to a term in a classification scheme or both. A facet with an empty text label that also omits the `link` attribute has no meaning.

2.5.1 Document level summaries of facets

The `ebuttm:documentFacet` element may be added to the `ebuttm:documentMetadata` element that is a child of `tt:metadata` element in the `head` element to provide a document level summary of a particular facet.

Rather than an `expresses` attribute the `ebuttm:documentFacet` element has a `summary` attribute, that takes the values "all_has", "mixed", "all_has_not" or "unspecified". If all the content in a document *has* a facet then the summary shall be "all_has". If all the content in a document *has_not* a facet then the summary shall be "all_has_not". If there is a mix of *has* and *has_not* and *unknown* or if some of the content does not have the facet then the summary shall be "mixed". If

none of the document's content has the facet or all of the document's content has the facet described as *unknown* then the summary shall be "unspecified".

The `ebuttm:documentFacet` element is otherwise identical to the `ebuttm:facet` element.

2.6 Tracing and debugging

The model presented here allows for multiple Processing Nodes to receive and emit streams. In real world scenarios it can be useful to log the activity that generated a document for audit or debugging purposes, for example to check that the correct configurations have been applied.

The `ebuttm:trace` element permits such logging. If present, a `trace` element shall describe in text the action that generated the document, in the `action` attribute, and an identifier that performed that action, in the `generatedBy` attribute. The action can be derived from a classification scheme not specified here. The `generatedBy` node identifier is a URI and is also not further defined here.

Optionally the `ebuttm:trace` element can identify the node that supplied the source content for the action, using the `sourceId` attribute.

The `ebuttm:trace` element may contain text content providing any further logging information.

3. Document Conformance

This section defines the requirements for documents that conform to this specification.

3.1 Generic Constraints

The EBU-TT Part 3 format defines constraints for an XML document instance. A valid EBU-TT Part 3 XML document shall comply with the generic constraints in § 3.1 and the document structure defined in § 3.2.

3.1.1 Namespaces

The following external namespaces from the W3C TTML specification shall be used for the TTML elements and attributes in EBU-TT Part 3:

Name	Prefix	Value
TT	tt:	http://www.w3.org/ns/ttml
TT Parameter	ttp:	http://www.w3.org/ns/ttml#parameter
TT Style	tts:	http://www.w3.org/ns/ttml#styling
TT Metadata	ttm:	http://www.w3.org/ns/ttml#metadata

The following namespaces shall be used for the assignment of XML Schema datatypes:

Name	Prefix	Value
XML Schema	xs:	http://www.w3.org/2001/XMLSchema

The following namespaces shall be used for the EBU-TT Part 3 specific vocabulary:

Name	Prefix	Value
EBU-TT Metadata	ebuttm:	urn:ebu:tt:metadata
EBU-TT Styling	ebutts:	urn:ebu:tt:style
EBU-TT Datatypes	ebuttdt:	urn:ebu:tt:datatypes
EBU-TT Parameters	ebuttp:	urn:ebu:tt:parameters

Note: Although any prefix can be used to bind the namespaces in an XML document the use of the prefixes listed above is recommended.

If attributes in this document are defined without prefix they are not in any namespace.

3.1.2 Extensibility

EBU-TT Part 3 documents may be extended as defined by EBU TT Part 1 (Tech 3350) [EBUTT1].

Note that this document defines the `ebuttm:facet` element which is explicitly permitted as a child of the `tt:metadata` element - see § 3.2.2.6.1.

3.1.3 Compatibility with TTML 1.0 timing model

The timing model in EBU-TT Part 3 is compatible with the TTML 1.0 timing model. The additional rules concerning the temporal activation of documents defined in § 2.3.1 constitute constraints on the Document Processing Context as defined in TTML 1.0.

3.2 Document Structure and Content Profile

This specification bases document compliance on EBU-TT Part 1 (Tech 3350) [EBUTT1] with some changes. Only the changes are described here. These changes are either modifications to the cardinality of elements or attributes or additional elements and attributes. Any elements and attributes omitted from this section are defined identically to EBU-TT Part 1 (Tech 3350).

The formal definition of how the EBU-TT Part 3 specification uses EBU-TT-, TTML- and XML-vocabulary is presented in tabular form. When using this specification, the definition of the use of an element or attribute shall be interpreted relative to the position defined in the table.

Example:

The definition of the use of the `ttp:timebase` attribute in § 3.2.2.1 "tt:tt" specifies only the permitted syntax of the `dur` attribute in relation to the value of the `ttp:timebase` element; all other definition of the `ttp:timebase` element is as defined in [EBUTT1].

Definitions used within this section:

- Type:* Constraints of the Information structure of an XML element or XML attribute. The type can be further constrained through Enumerations and normative text in the description.
- Enumeration:* Enumerated values that shall be used for certain elements or attributes of type `xs:string`.
- Cardinality:* How often an element or attribute may be used inside the corresponding parent element. If the lower bound is greater than 0 (e.g. "1..1" or "1..*") the element or attribute is mandatory at this position of the document structure. If the lower bound is equal to 0 (e.g. "0..1" or "0..*") the element or attribute is optional at this position of the document structure.
- Position:* The position of an element or attribute in an EBU-TT Document as XPATH expression starting with the document root "/". A default namespace of "<http://www.w3.org/ns/ttml>" is assumed and the prefix "tt" is omitted in the XPATH expression.
- TTML:* The URL to the specific chapter in the TTML 1.0 specification where the attribute or element is defined. The normative constraints of TTML 1.0 apply unless they

are further constrained by this specification.⁷

Rows highlighted **in grey** indicate required attributes.

3.2.1 Elements and attributes whose cardinality differs

The following table shows the entities whose cardinality differs relative to [EBUTT1]. The entity names are provided in the style of an XPath [XPATH].

Entity path and name	Entity type	Cardinality		Notes
		Tech 3350	Here	
tt:tt/tt:body	Element	1..1	0..1	Body may be omitted to indicate that all content should be removed, as a 'clear' instruction in an implicit processing context.
tt:tt/tt:head/tt:styling	Element	1..1	0..1	Styling may be omitted if unknown
tt:tt/tt:head/tt:layout	Element	1..1	0..1	Layout may be omitted if unknown
/(tt:body tt:div)/@begin	Attribute	0..0	0..1	
/(tt:body tt:div)/@end	Attribute	0..0	0..1	
//tt:p/@begin	Attribute	1..1	0..1	
//tt:p/@end	Attribute	1..1	0..1	

3.2.2 Newly introduced and modified elements and attributes

Entity name	Entity type	Location
ttp:timeBase (applicable to "dur" attribute)	Attribute	tt:tt
ebuttp:sequenceIdentifier	Attribute	tt:tt
ebuttp:sequenceNumber	Attribute	tt:tt
ebuttm:authoringDelay	Attribute	tt:tt
ebuttp:authorsGroupIdentifier	Attribute	tt:tt
ebuttp:authorsGroupControlToken	Attribute	tt:tt
ebuttp:authorsGroupControlRequest	Attribute	tt:tt
ebuttp:referenceClockIdentifier	Attribute	tt:tt
dur	Attribute	tt:body
ebuttm:originalSourceServiceIdentifier	Element	ebuttm:documentMetadata
ebuttm:intendedDestinationServiceIdentifier	Element	ebuttm:documentMetadata
ebuttm:trace	Element	ebuttm:documentMetadata
ebuttm:facet	Element	tt:body//tt:metadata
ebuttm:documentFacet	Element	ebuttm:documentMetadata

The introduced entities are defined in further detail below.

⁷ 1 At time of publication the links are based on the TTML version specified at [TTML1].

3.2.2.1 tt:tt

The following attributes are defined by this specification for inclusion on the `tt:tt` element.

EBU-TT Part 3 uses the following parameters from [TTML1] to give information on how the timing information in an EBU-TT document should be interpreted. If present, these attributes shall be specified on the `tt:tt` element.

ttp:timeBase (attribute)

Type	<code>xs:string</code>
Enumeration	"smpte" "media" "clock"
Cardinality	1..1
Position	/tt
TTML	http://www.w3.org/TR/ttml1/#parameter-attribute-timeBase
Description	The <code>ttp:timeBase</code> element is as defined in [EBUTT1] with the addition that all time expressions of <code>dur</code> attributes shall denote a relative coordinate on the same timeline as the <code>begin</code> and <code>end</code> attributes.

ebuttp:sequenceIdentifier (attribute)

Type	<code>xs:string</code>
Cardinality	1..1
Position	/tt
Description	The sequence to which every document belongs shall be identified using the <code>ebuttp:sequenceIdentifier</code> attribute.

ebuttp:sequenceNumber (attribute)

Type	<code>xs:positiveInteger</code>
Cardinality	1..1
Position	/tt
Description	Every document with the same <code>ebuttp:sequenceIdentifier</code> shall be uniquely numbered using the <code>ebuttp:sequenceNumber</code> attribute.

ebuttm:authoringDelay (attribute)

Type	ebuttdt:delayTimingType
Cardinality	0..1
Position	/tt
Description	The authoring delay associated with the timed content within this document may be indicated using the <code>ebuttm:authoringDelay</code> attribute. NOTE: this delay may be estimated or measured, and is intended to represent the duration between when the subtitle authoring tool received the instantaneous media for which subtitles were authored and the moment that the authoring tool emitted those subtitles.

The parameters `ebuttp:authorsGroupIdentifier`, `ebuttp:authorsGroupControlToken` and `ebuttp:authorsGroupControlRequest` are provided to facilitate handover between subtitle authors, using semantics defined for the Handover Manager node in § 2.4.

ebuttp:authorsGroupIdentifier (attribute)

Type	xs:string
Cardinality	0..1
Position	/tt
Description	Identifies the group of authors whose sequences relate to the same content and amongst which a Handover Manager should select documents when generating its output sequence.

ebuttp:authorsGroupControlToken (attribute)

Type	xs:positiveInteger
Cardinality	0..1
Initial value	0
Position	/tt
Description	The control token used to direct a Handover Manager to select an input sequence from a particular authors group. The input sequence whose document has the greatest value <code>ebuttp:authorsGroupControlToken</code> value is selected for output.

ebuttm:authorsGroupControlRequest (attribute)

Type	xs:string
Cardinality	0..1
Position	/tt
Description	Allows an author to signal their intent to continue providing a stream, or to request that another author begins creating output with a higher value <code>ebuttp:authorsGroupControlToken</code> . NOTE: For example the value "request" might be used to request that another subtitle author takes over and the value "continuing" might be used to signal the intent to continue subtitle authoring.

ebuttp:referenceClockIdentifier (attribute)

Type	xs:anyURI
Cardinality	0..1
Position	/tt
Description	Allows the reference clock source to be identified. Permitted only when <code>ttp:timeBase="clock"</code> AND <code>ttp:clockMode="local"</code> OR when <code>ttp:timeBase="smpte"</code> .

3.2.2.2 ebuttm:documentMetadata

The following additional metadata elements are defined by this specification. If used these elements shall be the first children of the `ebuttm:documentMetadata` element in the order documented below, after any metadata elements defined by EBU-TT Part 1 Tech 3350.

3.2.2.2.1 ebuttm:originalSourceServiceIdentifier

Type	xs:string
Cardinality	0..1
Position	/tt/head/metadata/ebuttm:documentMetadata
Description	The <code>ebuttm:originalSourceServiceIdentifier</code> may be used to identify the stream of audio-visual content that was used as the source for authoring the document. Identifier of the audio-visual content that was used as the source for authoring the document.

3.2.2.2.1 ebuttm:intendedDestinationServiceIdentifier

Type	xs:string
Cardinality	0..1
Position	/tt/head/metadata/ebuttm:documentMetadata
Description	The <code>ebuttm:intendedDestinationServiceIdentifier</code> may be used to identify the streams of destination audio-visual content for which the document is intended to apply. Identifier of the audio-visual content that is the intended destination for the document.

3.2.2.2.1 ebuttm:documentFacet

Type	xs:string
Cardinality	0..*
Position	/tt/head/metadata/ebuttm:documentMetadata
Description	<p>The <code>documentFacet</code> element indicates a document level summary of a facet that applies to the document's content as defined by the <code>facet</code> element. The value may be the name of a term in a classification scheme, whose term should be identified by the <code>link</code> attribute. Each distinctly identified facet that is summarised shall have a separate <code>documentFacet</code> element. Empty <code>documentFacet</code> elements should have a <code>link</code> attribute. Documents shall NOT contain more than one <code>documentFacet</code> element referring to the same term, where the term is identified by the combination of the element contents and the value of the <code>link</code> attribute.</p>

The `ebuttm:documentFacet` element may have the following attributes:

link (attribute)

Type	xs:anyUri
Cardinality	0..1
Position	/tt/head/metadata/ebuttm:documentMetadata/ebuttm:documentFacet
Description	<p>Reference to a term in a classification scheme. The <code>link</code> attribute should be present if the <code>documentFacet</code> element is empty.</p>

summary (attribute)

Type	"all_has" "mixed" "all_has_not" "unspecified"
Cardinality	0..1
Initial	"all_has"
Position	/tt/head/metadata/ebuttm:documentMetadata/ebuttm:documentFacet
Description	<p>The summary of document content for this facet, i.e. the application of <code>facet</code> elements with the same term. The value "all_has" shall be used if all of the document's content has this facet. The value "mixed" shall be used if the document's content contains multiple values of the <code>expresses</code> attribute for this facet, or if it is not applied to some of the content. The value "all_has_not" shall be used if all of the document's content has this facet with the <code>expresses</code> attribute set to "has_not". The value "unspecified" shall be used if none of the document's content has this facet applied or if all of the document's content has this facet with the <code>expresses</code> attribute set to "unknown".</p>

3.2.2.2.1 ebuttm:trace

Type	xs:string
Cardinality	0..*
Position	/tt/head/metadata/ebuttm:documentMetadata
Description	The trace element may be used to indicate processing that has been applied during the creation process of a document; multiple trace elements may be used to indicate a chain of processing.

The `ebuttm:trace` element may have the following attributes:

action (attribute)

Type	xs:string
Cardinality	1..1
Position	/tt/head/metadata/ebuttm:documentMetadata/ebuttm:trace
Description	The action being noted by this trace element.

generatedBy (attribute)

Type	xs:anyURI
Cardinality	1..1
Position	/tt/head/metadata/ebuttm:documentMetadata/ebuttm:trace
Description	The identifier for the node that performed the action being noted by this trace element.

sourceId (attribute)

Type	xs:anyURI
Cardinality	0..1
Position	/tt/head/metadata/ebuttm:documentMetadata/ebuttm:trace
Description	The identifier for the node that supplied the source sequence from which the document was derived.

3.2.2.3 tt:body

Type	Element content
Cardinality	0..1
Position	/tt
TTML	http://www.w3.org/TR/ttml1/#document-structure-vocabulary-body
Description	<p>The <code>tt:body</code> element is defined as in [EBUTT1] with the attributes in the following tables also permitted.</p> <p>A document that contains a <code>tt:body</code> element with no content shall be treated as being active as defined by the semantics described in § 2.3.1, and shall cause no content to be presented while it is active.</p> <p>Note: [TTML1] § 9.3.2 Intermediate Synchronic Document Construction specifies that empty elements are pruned, and therefore play no part in the creation of intermediate synchronic documents; nevertheless this specification interprets the timing attributes present on <code>tt:body</code> as defining the document activation period, even if presentation of that document generates no intermediate synchronic documents by that algorithm.</p>

begin (attribute)

Type	<code>ebuttdt:smpteTimingType</code> <code>ebuttdt:mediaTimingType</code> <code>ebuttdt:clockTimingType</code>
Cardinality	0..1
Position	tt/body
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-begin
Description	<p>Start point of a temporal interval associated with a <code>tt:body</code> element.</p> <p>If the timebase is "smpte" the type shall be <code>ebuttdt:smpteTimingType</code>.</p> <p>If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code>.</p> <p>If the timebase is "media" the time expression should be the offset from a syncbase of "00:00:00.0".</p> <p>If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code>.</p>

end (attribute)

Type	<code>ebuttdt:smpteTimingType</code> <code>ebuttdt:mediaTimingType</code> <code>ebuttdt:clockTimingType</code>
Cardinality	0..1
Position	tt/body
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-end
Description	<p>End point of a temporal interval associated with a <code>tt:p</code> element.</p> <p>If the timebase is "smpte" the type shall be <code>ebuttdt:smpteTimingType</code>.</p> <p>If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code>.</p> <p>If the timebase is "media" the time expression should be the offset from a syncbase of "00:00:00.0".</p> <p>If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code>.</p>

dur (attribute)

Type	ebuttdt:mediaTimingType ebuttdt:clockTimingType
Cardinality	0..1
Position	tt/body
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-dur
Description	<p>The maximum duration of the document relative to the resolved begin time, as defined in TTML 1. See also § 2.3.1 for further details of the semantics of document activation and the resolution of begin and end times for each document in a sequence.</p> <p>The <code>dur</code> attribute shall NOT be used in documents that set <code>ttp:markerMode="discontinuous"</code>⁸.</p> <p>If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code>. If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code>.</p>

3.2.2.4 tt:div

Type	Element content
Cardinality	0..*
Position	/tt/(body div)
TTML	http://www.w3.org/TR/ttml1/#document-structure-vocabulary-div
Description	The <code>tt:div</code> element is defined as in [EBUTT1] with the attributes in the following tables also permitted.

begin (attribute)

Type	ebuttdt:smpteTimingType ebuttdt:mediaTimingType ebuttdt:clockTimingType
Cardinality	0..1
Position	tt/body//div
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-begin
Description	<p>Start point of a temporal interval associated with a <code>tt:div</code> element.</p> <p>If the timebase is "smpte" the type shall be <code>ebuttdt:smpteTimingType</code>. If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code>. If the timebase is "media" or "clock" the time expression should be the offset from a syncbase of the begin time of the closest ancestor that specifies a begin time. If no ancestor specifies a begin time and the timebase is "media" the time expression should be the offset from a syncbase of "00:00:00.0". If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code>.</p>

⁸ Note: in [EBUTT1] the required value for `ttp:markerMode` is "discontinuous" when `ttp:timeBase` is "smpte". [TTML1] specifies that `ttp:markerMode` only applies when `ttp:timeBase` is "smpte" and that the initial value of `ttp:markerMode` is "discontinuous".

end (attribute)

Type	ebuttdt:smpteTimingType ebuttdt:mediaTimingType ebuttdt:clockTimingType
Cardinality	0..1
Position	tt/body//div
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-end
Description	End point of a temporal interval associated with a <code>tt:div</code> element. If the timebase is "smpte" the type shall be <code>ebuttdt:smpteTimingType</code> . If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code> . If the timebase is "media" or "clock" the time expression should be the offset from a syncbase of the begin time of the closest ancestor that specifies a begin time. If no ancestor specifies a begin time and the timebase is "media" the time expression should be the offset from a syncbase of "00:00:00.0". If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code> .

3.2.2.5 tt:p

Type	Element content
Cardinality	0..*
Position	/tt/body//div
TTML	http://www.w3.org/TR/ttml1/#content-vocabulary-p
Description	The <code>tt:p</code> element is defined as in [EBUTT1] with adjustments to the interpretation of the attributes in the following tables to take into account the ability to nest timed content elements and to omit times.

begin (attribute)

Type	ebuttdt:smpteTimingType ebuttdt:mediaTimingType ebuttdt:clockTimingType
Cardinality	0..1
Position	tt/body//div/p
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-begin
Description	Start point of a temporal interval associated with a <code>tt:p</code> element. If the timebase is "smpte" the type shall be <code>ebuttdt:smpteTimingType</code> . If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code> . If the timebase is "media" or "clock" the time expression should be the offset from a syncbase of the begin time of the closest ancestor that specifies a begin time. If no ancestor specifies a begin time and the timebase is "media" the time expression should be the offset from a syncbase of "00:00:00.0". If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code> .

end (attribute)

Type	<code>ebuttdt:smpteTimingType</code> <code>ebuttdt:mediaTimingType</code> <code>ebuttdt:clockTimingType</code>
Cardinality	0..1
Position	<code>tt/body//div/p</code>
TTML	http://www.w3.org/TR/ttml1/#timing-attribute-end
Description	End point of a temporal interval associated with a <code>tt:p</code> element. If the timebase is "smpte" the type shall be <code>ebuttdt:smpteTimingType</code> . If the timebase is "media" the type shall be <code>ebuttdt:mediaTimingType</code> . If the timebase is "media" or "clock" the time expression should be the offset from a syncbase of the begin time of the closest ancestor that specifies a begin time. If no ancestor specifies a begin time and the timebase is "media" the time expression should be the offset from a syncbase of "00:00:00.0". If the timebase is "clock" the type shall be <code>ebuttdt:clockTimingType</code> .

3.2.2.6 tt:metadata

Type	Element content
Cardinality	0..1
Position	<code>/tt//(body div p span br)</code>
TTML	http://www.w3.org/TR/ttml1/#metadata-vocabulary-metadata
Description	The <code>tt:metadata</code> element is defined as in [EBUTT1]. The following additional elements are defined by this specification for inclusion on the <code>tt:metadata</code> element.

3.2.2.6.1 ebuttm:facet

Type	<code>xs:string</code>
Cardinality	0..*
Position	<code>/tt/body//metadata</code>
Description	The <code>facet</code> element indicates that a facet applies to the <code>metadata</code> element's parent element's content and all of its descendants that do not specify a different <code>expresses</code> attribute. The value may be the name of a term in a classification scheme, whose term should be identified by the <code>link</code> attribute. Each distinctly identified facet shall have a separate <code>facet</code> element, where facets are identified by combination of the text content and the <code>link</code> attribute. A facet can be used to describe technical or editorial aspects of the content to which it is applied, other than editorial purpose (for which the <code>ttm:role</code> attribute should be used). For example a facet could be 'checked_for_spelling', 'ISO6937_code_points_only' etc. Empty <code>facet</code> elements should have a <code>link</code> attribute. Elements shall NOT contain more than one <code>facet</code> element referring to the same term. Elements may specify a <code>facet</code> element referring to a term that is also referenced by a <code>facet</code> element applied to an ancestor or descendant element.

The `ebuttm:facet` element may have the following attributes:

link (attribute)

Type	<code>xs:anyUri</code>
Cardinality	0..1
Position	<code>tt/body//metadata/ebuttm:facet</code>
Description	Reference to a term in a classification scheme. The <code>link</code> attribute should be present if the <code>facet</code> element is empty.

expresses (attribute)

Type	<code>"has" "has_not" "unknown"</code>
Cardinality	0..1
Initial	<code>"has"</code>
Position	<code>tt/body//metadata/ebuttm:facet</code>
Description	The <code>expresses</code> attribute indicates if the content has the specified facet or does not have it. The value <code>"has"</code> implies that the element's content has this facet. The value <code>"has_not"</code> implies that the element's content explicitly does not have this facet. The value <code>"unknown"</code> implies that it is not known whether or not the element's content has this facet.

3.3 Datatypes

EBU-TT Part 3 defines specific datatypes to restrict the content of attributes or textual Element content.

EBU-TT Part 3 also uses any datatypes referenced and defined in [EBUTT1].

Note: If a datatype is applied to an attribute that was taken from [TTML1] the restriction of the datatype is equal to the definition in TTML 1.0 or it is a further restriction of the content as defined in TTML 1.0. Therefore all values that conform to the EBU-TT Part 3 datatypes also conform to the values allowed in TTML 1.0. However it is possible to create a value that conforms to the TTML 1.0 definitions but does not conform to the EBU-TT Part 3 datatypes.

3.3.1 `ebuttdt:delayTimingType`

The content shall be constrained to a signed (positive or negative) number with an optional decimal fraction, followed by a time metric being one of:

- "h" (hours)
- "m" (minutes)
- "s" (seconds)
- "ms" (milliseconds)

4. Node Conformance

This section defines the requirements for nodes that conform to this specification.

Node conformance is defined in terms of their behaviours only. The nodes defined here are not the only permitted nodes; the purpose of defining them is to define minimal expectations to support interoperability of node implementations.

Nodes are defined as abstract classes. See Figure 14 for a UML representation of the abstract node class structure.

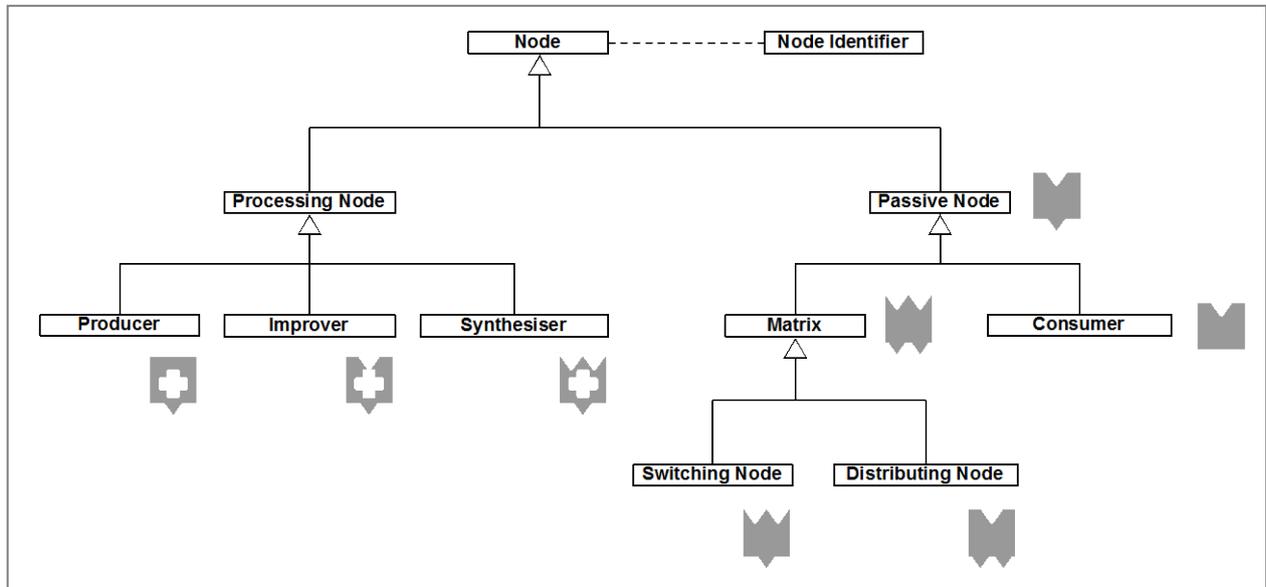


Figure 14: UML model diagram showing logical nodes and their relationships

4.1 Generic Node Classes

This section defines the behaviours of nodes. The nodes are specified as abstract classes that inherit behaviour from their parent. This section is structured according to the inheritance tree, so a subsection inherits the definitions of its parent section within the document. For example a Processing Node *is* a Node.

4.1.1 Node

A node is a logical processing unit that consumes or emits a sequence.

Note: physical implementations could combine the actions of multiple nodes.

4.1.1.1 Processing Node

A processing node emits a sequence that can differ from any of its inputs.

A processing node may consume one or more sequences. The sequence that it emits shall be differently identified to all of its input sequences.

4.1.1.1.1 Producer Node

A Producer consumes no sequences.

Note: An authoring station is a concrete example of a Producer node.

4.1.1.1.2 Improver Node

An Improver consumes exactly one sequence.

4.1.1.1.2.1 *Delay*

A Delay imposes a specified delay between its input sequence and its output sequence according to the semantics defined in section § 2.3.4.

4.1.1.1.3 **Synthesiser Node**

A Synthesiser consumes multiple sequences and outputs a new sequence in some way derived from those input sequences.

4.1.1.1.3.1 *Handover Manager*

A Handover Manager synthesises an output sequence based on a combination of its input sequences, according to the semantics defined in section § 2.4.

4.1.1.2 **Passive Node**

A Passive Node receives one or more sequences. Any sequences that it emits are identical to the sequence(s) that it consumes.

4.1.1.2.1 **Matrix**

A matrix receives multiple sequences and emits multiple sequences.

4.1.1.2.1.1 *Switching Node*

A switching node receives multiple sequences and emits one of the received sequences.

4.1.1.2.1.2 *Distributing Node*

A distributing node receives one sequence and emits it as multiple streams.

4.1.1.2.4 **Consumer**

A consumer receives one sequence and emits zero sequences.

Note: An encoder is a concrete example of a Consumer node.

5. Bibliography

[EBUTT1]	EBU Tech 3350	EBU-TT Part 1 v1.1 Subtitling format definition https://tech.ebu.ch/publications/Tech 3350
[EBUTTD]	EBU Tech 3380	EBU-TT-D Subtitling Distribution Format https://tech.ebu.ch/publications/Tech 3380
[TTML1]	TTML 1.0	Timed Text Markup Language (TTML) 1.0 (Second Edition), W3C Recommendation http://www.w3.org/TR/2013/REC-ttml1-20130924/
[UAX9]	UAX9	Mark Davis. Unicode Standard Annex #9. Unicode Bidirectional Algorithm. http://unicode.org/reports/tr9/
[XML1]	XML 1.0	Tim Bray, et al. Extensible Markup Language (XML) 1.0 (Fifth Edition), W3C Recommendation, 26 November 2008. http://www.w3.org/TR/2008/REC-xml-20081126/
[SMIL]	SMIL 2.1	Synchronized Multimedia Integration Language (SMIL 2.1), W3C Recommendation. http://www.w3.org/TR/2005/REC-SMIL2-20051213/
[XPATHT]	XPATHT 2.0	XPath 2.0 (Second Edition) http://www.w3.org/TR/xpath20/
[OPENSTD]	Open Stand	Open Stand Principles https://open-stand.org/about-us/principles/

Annex A: Overview document structure (Informative)

The following is a syntactic representation of the EBU-TT Part 3 document model. It is derived from the syntactic representation of TTML 1.0 and the definition of the reduced XML Infoset in TTML 1.0 and is aligned where appropriate with the equivalent section of [EBUTT1].

ELEMENT INFORMATION ITEMS

```

<tt:tt
  xml:space = ("default"|"preserve")>
  ttp:timeBase = ( "media" | "smpte" | "clock") #REQUIRED
  ttp:frameRate = <digit>* <digitGreaterZero> <digit>*
  ttp:frameRateMultiplier = <ebuttdt:frameRateMultiplierType>
  ttp:markerMode = ("discontinuous")
  ttp:dropMode = ("nonDrop" | "dropNTSC" | "dropPAL")
  ttp:clockMode = ( "local" | "gps" | "utc")
  xml:lang = (" " | <xs:language>) #REQUIRED
  ttp:cellResolution = <ebuttdt:cellResolutionType>
  tts:extent = <ebuttdt:extentType> /* Restricted to length metric "px" */
  ebuttp:sequenceIdentifier = xs:string #REQUIRED
  ebuttp:sequenceNumber = xs:positiveInteger #REQUIRED
  ebuttm:authoringDelay = <ebuttdt:delayTimingType>
  ebuttp:authorsGroupIdentifier = xs:string
  ebuttp:authorsGroupControlToken = xs:positiveInteger
  ebuttp:authorsGroupControlRequest = xs:string
  ebuttp:referenceClockIdentifier = xs:anyURI
  >
  Content: tt:head, tt:body?
</tt:tt>

<tt:head>
  Content: tt:metadata?, ttm:copyright?, tt:styling?, tt:layout?
</tt:head>

<ttm:copyright>
  Content: <xs:string>
</ttm:copyright>

<tt:metadata>
  Content:
    ebuttm:documentMetadata?, /* Only permitted and required in
      '/tt/head/metadata' */
    ebuttm:binaryData*, /* Only permitted in '/tt/head/metadata' and
      '/tt/body//div/metadata' */
    ebuttm:authoringTechnique*, /* Only permitted in 'tt:body/tt:metadata',
      'tt:div/tt:metadata' and 'tt:p/tt:metadata' */
    ebuttm:transitionStyle*, /* Only permitted in 'tt:body/tt:metadata',
      'tt:div/tt:metadata' and 'tt:p/tt:metadata' */
    ebuttm:font*, /* Only permitted in 'tt:styling/tt:metadata' */
    ebuttm:facet*, /* Only permitted in 'tt:body//tt:metadata' */
    ttm:title?,
    ttm:desc?,
    ttm:agent* /* Only permitted in 'tt:head/metadata' */
  {Any element not in a TTML namespace and not in an EBU-TT namespace except where
  explicitly permitted in this specification}
</tt:metadata>

```

```

<ebuttm:documentMetadata>
  Content:
    ebuttm:conformsToStandard*,
    ebuttm:documentEbutttVersion?,
    ebuttm:documentIdentifier?,
    ebuttm:documentOriginatingSystem?,
    ebuttm:documentCopyright?,
    ebuttm:documentReadingSpeed?,
    ebuttm:documentTargetAspectRatio?,
    ebuttm:documentTargetActiveFormatDescriptor?,
    ebuttm:documentIntendedTargetBarData?,
    ebuttm:documentIntendedTargetFormat?,
    ebuttm:documentCreationMode?,
    ebuttm:documentContentType?,
    ebuttm:sourceMediaIdentifier*,
    ebuttm:relatedMediaIdentifier?,
    ebuttm:relatedObjectIdentifier*,
    ebuttm:appliedProcessing*,
    ebuttm:relatedMediaDuration?,
    ebuttm:documentBeginDate?,
    ebuttm:localTimeOffset?,
    ebuttm:referenceClockIdentifier?,
    ebuttm:broadcastServiceIdentifier*,
    ebuttm:documentTransitionStyle*,
    ebuttm:documentOriginalProgrammeTitle?,
    ebuttm:documentOriginalEpisodeTitle?,
    ebuttm:documentTranslatedProgrammeTitle?,
    ebuttm:documentTranslatedEpisodeTitle?,
    ebuttm:documentTranslatorsName?,
    ebuttm:documentTranslatorsContactDetails?,
    ebuttm:documentSubtitleListReferenceCode?,
    ebuttm:documentCreationDate?,
    ebuttm:documentRevisionDate?,
    ebuttm:documentRevisionNumber?,
    ebuttm:documentTotalNumberOfSubtitles?,
    ebuttm:documentMaximumNumberOfDisplayableCharacterInAnyRow?,
    ebuttm:documentStartOfProgramme?,
    ebuttm:documentCountryOfOrigin?,
    ebuttm:documentPublisher?,
    ebuttm:documentEditorsName?,
    ebuttm:documentEditorsContactDetails?,
    ebuttm:documentUserDefinedArea?,
    ebuttm:stlCreationDate?,
    ebuttm:stlRevisionDate?,
    ebuttm:stlRevisionNumber?
    ebuttm:subtitleZero?
    ebuttm:originalSourceServiceIdentifier?,
    ebuttm:intendedDestinationServiceIdentifier*,
    ebuttm:trace*,
    ebuttm:documentFacet*
  >
</ebuttm:documentMetadata>

<ebuttm:conformsToStandard>
  Content: <xs:anyURI>
</ebuttm:conformsToStandard>

<ebuttm:documentEbutttVersion>
  Content: "v1.0"
</ebuttm:documentEbutttVersion>

<ebuttm:documentIdentifier>
  Content: <xs:string>
</ebuttm:documentIdentifier>

```

```

<ebuttm:documentOriginatingSystem>
  Content: <xs:string>
</ebuttm:documentOriginatingSystem>

<ebuttm:documentCopyright>
  Content: <xs:string>
</ebuttm:documentCopyright>

<ebuttm:documentReadingSpeed>
  Content: <xs:positiveInteger>
</ebuttm:documentReadingSpeed>

<ebuttm:documentTargetAspectRatio>
  Content: <xs:string>
</ebuttm:documentTargetAspectRatio>

<ebuttm:documentTargetActiveFormatDescriptor>
  Content: <xs:string>
</ebuttm:documentTargetActiveFormatDescriptor>

<ebuttm:documentIntendedTargetBarData
  position = ("topBottom" | "leftRight") #REQUIRED
  lineNumberEndOfTopBar = <xs:nonNegativeInteger>
  lineNumberStartOfBottomBar = <xs:nonNegativeInteger>
  pixelNumberEndOfLeftBar = <xs:nonNegativeInteger>
  pixelNumberStartOfRightBar = <xs:nonNegativeInteger> >
  Content: <xs:string>
</ebuttm:documentIntendedTargetBarData>

<ebuttm:documentIntendedTargetFormat
  link = <xs:anyURI> >
  Content: <xs:string>
</ebuttm:documentIntendedTargetFormat>

<ebuttm:documentCreationMode>
  Content: ("live" | "prepared")
</ebuttm:documentCreationMode>

<ebuttm:documentContentType
  link = <xs:anyURI> >
  Content: <xs:string>
</ebuttm:documentContentType>

<ebuttm:sourceMediaIdentifier
  type = <xs:string> >
  Content: <xs:string>
</ebuttm:sourceMediaIdentifier>

<ebuttm:relatedMediaIdentifier>
  Content: <xs:string>
</ebuttm:relatedMediaIdentifier>

<ebuttm:relatedObjectIdentifier
  type = <xs:string> >
  Content: <xs:string>
</ebuttm:relatedObjectIdentifier>

<ebuttm:appliedProcessing
  appliedDateTime = <xs:dateTime> >
  Content (Mixed)
</ebuttm:appliedProcessing>

```

```

<ebuttm:relatedMediaDuration>
  Content: <ebuttdt:mediaTimingType>
</ebuttm:relatedMediaDuration>

<ebuttm:documentBeginDate>
  Content: <ebuttdt:noTimezoneDateType>
</ebuttm:documentBeginDate>

<ebuttm:localTimeOffset>
  Content: <xs:string>
</ebuttm:localTimeOffset>

<ebuttm:referenceClockIdentifier>
  Content: <xs:string>
</ebuttm:referenceClockIdentifier>

<ebuttm:broadcastServiceIdentifier
  serviceBegin = <xs:dateTime>
  serviceEnd = <xs:dateTime> >
  Content: <xs:string>
</ebuttm:broadcastServiceIdentifier>

<ebuttm:documentTransitionStyle
  inUnit = ("block" | "line" | "word" | "partOfWord" | "groupOfWords") # REQUIRED
  outUnit = ("block" | "line" | "word" | "partOfWord" | "groupOfWords") # REQUIRED
/>

<ebuttm:documentOriginalProgrammeTitle>
  Content: <xs:string>
</ebuttm:documentOriginalProgrammeTitle>

<ebuttm:documentOriginalEpisodeTitle>
  Content: <xs:string>
</ebuttm:documentOriginalEpisodeTitle>

<ebuttm:documentTranslatedProgrammeTitle>
  Content: <xs:string>
</ebuttm:documentTranslatedProgrammeTitle>

<ebuttm:documentTranslatedEpisodeTitle>
  Content: <xs:string>
</ebuttm:documentTranslatedEpisodeTitle>

<ebuttm:documentTranslatorsName>
  Content: <xs:string>
</ebuttm:documentTranslatorsName>

<ebuttm:documentTranslatorsContactDetails>
  Content: <xs:string>
</ebuttm:documentTranslatorsContactDetails>

<ebuttm:documentSubtitleListReferenceCode>
  Content: <xs:string>
</ebuttm:documentSubtitleListReferenceCode>

<ebuttm:documentCreationDate>
  Content: <xs:date>
</ebuttm:documentCreationDate>

<ebuttm:documentRevisionDate>
  Content: <xs:date>
</ebuttm:documentRevisionDate>

```

```
<ebuttm:documentRevisionNumber>
  Content: <xs:nonNegativeInteger>
</ebuttm:documentRevisionNumber>

<ebuttm:documentTotalNumberOfSubtitles>
  Content: <xs:nonNegativeInteger>
</ebuttm:documentTotalNumberOfSubtitles>

<ebuttm:documentMaximumNumberOfDisplayableCharacterInAnyRow>
  Content: <xs:nonNegativeInteger>
</ebuttm:documentMaximumNumberOfDisplayableCharacterInAnyRow>

<ebuttm:documentStartOfProgramme>
  Content: ( <ebuttdt:smpTeTimingType> | <ebuttdt:clockTimingType> )
</ebuttm:documentStartOfProgramme>

<ebuttm:documentCountryOfOrigin>
  Content: <xs:string>
</ebuttm:documentCountryOfOrigin>

<ebuttm:documentPublisher>
  Content: <xs:string>
</ebuttm:documentPublisher>

<ebuttm:documentEditorsName>
  Content: <xs:string>
</ebuttm:documentEditorsName>

<ebuttm:documentEditorsContactDetails>
  Content: <xs:string>
</ebuttm:documentEditorsContactDetails>

<ebuttm:documentUserDefinedArea>
  Content: <xs:string>
</ebuttm:documentUserDefinedArea>

<ebuttm:stlCreationDate>
  Content: <xs:date>
</ebuttm:stlCreationDate>

<ebuttm:stlRevisionDate>
  Content: <xs:date>
</ebuttm:stlRevisionDate>

<ebuttm:stlRevisionNumber>
  Content: <xs:nonNegativeInteger>
</ebuttm:stlRevisionNumber>

<ebuttm:authoringTechnique
  link = <xs:anyURI> >
  Content: <xs:string>
</ebuttm:authoringTechnique>

<ebuttm:subtitleZero>
  Content: <xs:string>
</ebuttm:subtitleZero>
```

```

<ebuttm:binaryData
  textEncoding = "BASE64" #REQUIRED
  binaryDataType = <xs:string> #REQUIRED
  fileName = <xs:string>
  creationDate = <xs:date>
  revisionDate = <xs:date>
  revisionNumber = <xs:nonNegativeInteger> >
  Content: <xs:string>
</ebuttm:binaryData>

<ebuttm:originalSourceServiceIdentifier>
  Content: <xs:string>
</ebuttm:originalSourceServiceIdentifier>

<ebuttm:intendedDestinationServiceIdentifier>
  Content: <xs:string>
</ebuttm:intendedDestinationServiceIdentifier>

<ebuttm:documentFacet
  link = <xs:anyURI>,
  summary = ("all_has" | "mixed" | "all_has_not" | "unspecified") >
  Content: <xs:string>
</ebuttm:documentFacet>

<ebuttm:trace
  action = <xs:string> #REQUIRED
  generatedBy = <xs:anyURI> #REQUIRED
  sourceId = <xs:anyURI> >
  Content: <xs:string>
</ebuttm:trace>

<ebuttm:font
  fontFamilyName = <ebuttdt:fontFamilyType> #REQUIRED
  src = <xs:anyUri> #REQUIRED
  fontStyle = ("italic" | "normal")
  fontWeight = ("bold" | "normal")
  fontSize = <ebuttdt:fontSizeType> > /* The metric '%' shall not be used */
/>

<ebuttm:facet
  link = <xs:anyURI>,
  expresses = ("has" | "has_not" | "unknown") >
  Content: <xs:string>
</ebuttm:documentFacet>

<tt:styling>
  Content: tt:metadata?, tt:style*
</tt:styling>

```

```

<tt:style
  xml:id = <xs:ID> #REQUIRED
  style = <xs:IDREFS>
  tts:direction = ( "ltr" | "rtl" )
  tts:fontFamily = <ebuttdt:fontFamilyType>
  tts:fontSize = <ebuttdt:fontSizeType>
  tts:lineHeight = ( "normal" | <ebuttdt:lengthType> )
  tts:textAlign = ( "left" | "center" | "right" | "start" | "end" )
  tts:color = <ebuttdt:colorType>
  tts:backgroundColor = <ebuttdt:colorType>
  tts:fontStyle = ( "normal" | "italic" )
  tts:fontWeight = ( "normal" | "bold" )
  tts:textDecoration = ( "none" | "underline" )
  tts:unicodeBidi = ( "normal" | "embed" | "bidiOverride" )
  tts:wrapOption = ( "wrap" | "noWrap" )
  tts:padding = <ebuttdt:paddingType>
  ebutts:multiRowAlign = ( "start" | "center" | "end" | "auto" )
  ebutts:linePadding = <ebuttdt:linePaddingType> >
  Content: tt:metadata?
</tt:style>

<tt:layout>
  Content: tt:metadata?, tt:region*
</tt:layout>

<tt:region
  xml:id = <xs:ID> #REQUIRED
  tts:origin = <ebuttd:originType> #REQUIRED
  tts:extent = <ebuttd:extentType> #REQUIRED
  style = <xs:IDREFS>
  tts:displayAlign = ( "before" | "center" | "after" )
  tts:padding = <ebuttdt:paddingType>
  tts:writingMode = ( "lrbt" | "rltb" | "tblr" | "tblr" | "lr" | "rl" | "tb" )
  tts:showBackground = ( "always" | "whenActive" )
  tts:overflow = ( "visible" | "hidden" ) >
  Content: tt:metadata?
</tt:region>

<tt:body
  style = <xs:IDREFS>
  begin = ( <ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType> )
  end = ( <ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType> )
  dur = ( <ebuttdt:mediaTimingType> | <ebuttdt:clockTimingType> )
  ttm:agent = <xs:IDREFS>
  ttm:role = As defined in TTML 1.0 [TTML1], section 12.2.2>
  Content: tt:metadata?, tt:div+
</tt:body>

<tt:div
  xml:id = <xs:ID>
  region = <xs:IDREF>
  style = <xs:IDREFS>
  begin = ( <ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType> )
  end = ( <ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType> ) >
  xml:lang = ( "" | <xs:language> )
  ttm:agent = <xs:IDREFS>
  ttm:role = As defined in TTML 1.0 [TTML1], section 12.2.2>
  Content: tt:metadata?, tt:div*, tt:p*
</tt:div>

```

```

<tt:p
  xml:id = <xs:ID> #REQUIRED
  xml:space = ("default" | "preserve")
  xml:lang = (" " | <xs:language>)
  region = <xs:IDREF>
  style = <xs:IDREFS>
  begin = (<ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType>)
  end = (<ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType>)
  ttm:agent = <xs:IDREFS>
  ttm:role = As defined in TTML 1.0 [TTML1], section 12.2.2
  Content (Mixed): tt:metadata?, (tt:span | tt:br)*
</tt:p>

<tt:span
  xml:id = <xs:ID>
  xml:space = ("default" | "preserve")
  xml:lang = (" " | <xs:language>)
  style = <xs:IDREFS>
  begin = (<ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType>)
  end = (<ebuttd:smpteTimingType> | <ebuttdt:mediaTimingType> |
<ebuttdt:clockTimingType>)
  ttm:agent = <xs:IDREFS>
  ttm:role = As defined in TTML 1.0 [TTML1], section 12.2.2
  Content (Mixed): tt:metadata?, tt:span*, tt:br*
</tt:span>

<tt:br
  ttm:agent = <xs:IDREFS>
  ttm:role = As defined in TTML 1.0 [TTML1], section 12.2.2
  Content: tt:metadata?
</tt:br>

```

EXPRESSIONS

```

<ebuttdt:CellResolutionType>
  : <columns> <whiteSpace> <rows>

<columns> | <rows>
  : <digit>* <digitGreaterZero> <digit>*

<ebuttdt:colorType>
  : <color> as defined in TTML 1

<ebuttdt:extentType>
  : <width> <whiteSpace> <height>

<width> | <height>
  : <ebuttdt:lengthType>

<ebuttdt:fontFamilyType>
  : <familyName> | <genericFamilyName> as defined in TTML 1

<ebuttdt:fontSizeType>
  : <ebuttdt:lengthType> <whiteSpace> <ebuttdt:lengthType>?

<ebuttdt:frameRateMultiplierType>
  : <numerator> <whiteSpace> <denominator>

<numerator> | <denominator>
  : <digit>* <digitGreaterZero> <digit>*

```

```

<ebuttdt:lengthType>
  : <scalar>
  | <percentage>

<scalar>
  : <number> <units>

<percentage>
  : <number> "%"

<number>
  : <sign>? <non-negative-number>

<sign>
  : "+" | "-"

<non-negative-number>
  : <non-negative-integer>
  | <non-negative-real>

<non-negative-integer>
  : <digit>+

<non-negative-real>
  : <digit>* "." <digit>+

<units>
  : "px" /* abbreviation of "pixel" */
  | "c" /* abbreviation of "cell" */

<ebuttdt:lineHeightType>
  : "normal" | <ebuttdt:lengthType> /* length >= 0 */

<ebuttdt:originType>
  : <x-coord> <whiteSpace> <y-coord>

<x-coord> | <y-coord>
  : <ebuttdt:lengthType>

<ebuttdt:paddingType>
  : <all-edges>
  | <beforeAndAfter> <whiteSpace> <startAndEnd>
  | <before> <whiteSpace> <startAndEnd> <whiteSpace> <after>
  | <before> <whiteSpace> <end> <whiteSpace> <after> <whiteSpace> <start>

<all-edges> | <before> | <end> | <after> | <start> | <beforeAndAfter> |
<startAndEnd>
  : <ebuttdt:lengthType>

<ebuttdt:linePaddingType>
  : <non-negative-number> "c"

<ebuttdt:noTimezoneDateType>
  : <xs:date> with no timezone specified

<ebuttdt:smpteTimingType>
  : <hh> ":" <mm> ":" <ss> ":" <ff>

<hh> | <mm> | <ss> | <ff>
  : <digit> <digit>

```

```
<ebuttdt:mediaTimingType>
  : <full-clock-value>
  | <timecount-value>

<ebuttdt:clockTimingType>
  : <full-clock-value>
  | <timecount-value>

<ebuttdt:delayTimingType>
  : <sign> <timecount-value>

<full-clock-value>
  : <hh> ":" <mm> ":" <ss> ( "." <digit>+ )?

<timecount-value>
  : <digit>+ ( "." <digit>+ )? <metric>

<metric>
  : "h" /* hours */
  | "m" /* minutes */
  | "s" /* seconds */
  | "ms" /* milliseconds */

<digit>
  : "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"

<digitGreaterZero>
  : "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"

<whiteSpace> /*(space, carriage return, line feed, tab)*/
  : (#x20 | #x9 | #xD | #xA)+
```

Annex B: Requirements for Carriage Specifications

This document does not define carriage mechanisms for streaming sequences between nodes.

A specification described as an "EBU-TT Part 3 Carriage Specification" is conformant if it complies with the requirements described in this annex.

Authors of such carriage specification are encouraged to inform EBU for example by email to subtitling@ebu.ch.

Core networking dependencies and connection protocols

Carriage specification documents shall describe the core networking dependencies and protocols, i.e. if the mechanism is based on IP with TCP, IP with UDP, VBI in SDI, VANC in HD-SDI etc. and shall reference any related standards.

It is recommended that carriage specifications should NOT depend on non-standard connection protocols, i.e. those that do not conform to common definitions of open standards such as those developed to conform to the principles of [OPENSTD].

Where the networking dependencies impose constraints, those constraints shall be described. For example a carriage specification for EBU-TT Live over VBI embedded in SDI would impose a maximum data rate constraint. Another example could be the need for, and impact of packetisation if used by a networking protocol.

Synchronisation impacts and/or thresholds

Carriage specification documents shall describe any interactions between the carriage mechanism and the synchronisation of live subtitles. For example an embedded mechanism such as VANC in HD-SDI could maintain frame-based synchronisation wherever it is routed, meaning that the synchronisation is defined when the subtitles are embedded rather than when they are received further downstream.

Information Security

Carriage specification documents shall describe:

- what provision is made for supporting information security requirements including but not limited to authentication, encryption and error checking mechanisms;
- under what circumstances the mechanisms described can be considered acceptable for crossing organisational boundaries;
- any known impacts or dependencies on other information security technologies, for example is the mechanism transparent to firewalls or does it need special configuration, is there a framework for allowing future authentication and encryption mechanisms to be used, etc.

Endpoint cardinality

Carriage specification documents shall describe whether they natively support point to point delivery, or point to multipoint delivery, or both.

If a carriage specification makes use of a duplex reverse channel for delivery monitoring or fault identification then that mechanism shall be described including the technique used, e.g. acknowledgement messages and their format, the time impact of such techniques, and the expected behaviour in case a fault is identified. Note that reverse channels are not a requirement in general.

Connection lifecycle management

Carriage specification documents shall either a) define the lifecycle of any connections, or b) reference standards that define the lifecycle of any connections. The lifecycle here refers to the initiation, establishment, ongoing maintenance, planned closure and error handling of the connections.

Note: If carriage specifications do not need use any form of connection this requirement is relaxed.

Channel routing

Carriage specifications may define registries or other ways to associate a) streams with the services, channels, programmes, languages etc. for which they are intended and b) nodes that can provide particular streams. This information could be used operationally to automate the routing of subtitle streams from authors to encoders. If such mechanisms are included the metadata model used should be described or referenced, and any impact on synchronisation should be described.

Note: the channel routing mechanisms could include a description of how to switch between input streams in a content dependent way.

Stability

Carriage specifications shall describe or reference the expected level of connection stability, and operational approaches for maintaining that stability. For example how the mechanism handles dropped connection fault conditions, whether data delivery is guaranteed or some documents may be lost as a matter of course, how the latency characteristics of the mechanism may vary, how version compatibility issues are managed, what statistics and monitoring are available etc.

Note: A TCP socket based protocol would exhibit guaranteed data delivery behaviour at the possible expense of delivery timing; conversely a UDP based protocol would exhibit minimal delivery timing at the possible expense of data loss. These characteristics affect the stability of the operational connections, and are need to be described so that stable systems can be engineered.

Interoperability

It is recommended that carriage mechanisms should be defined with interoperability in mind. Interfaces should not need to be hardware or vendor specific, should ideally publish details of available services on request, should self-describe their external interfaces, and should permit any relevant configuration of the services available to be published. Where carriage mechanisms multiplex EBU-TT Part 3 data with other data, e.g. by embedding in an HD-SDI signal, the mechanism shall be transparent to devices that are unable to process EBU-TT Part 3 data, i.e. the presence of EBU-TT Part 3 data shall NOT cause a fault in compliant devices, for example by extending beyond fixed size data windows, including bytes with special meaning, mis-representing the EBU-TT Part 3 data as an incompatible format etc.

Serial data communication protocols such as RS232 and RS422 are considered unsuitable for meeting this need and are deprecated.

Annex C: Not fully defined example scenarios (Informative)

The following examples represent typical real world scenarios in which documents and nodes that conform to this specification can be used, which this specification does not fully define. Note that the scenario numbers form a set with those in the § 1.1 and therefore do not increase monotonically.

Example 3: Prepared subtitle playback – not defined in this specification

A broadcast playout provider is tasked with playing out a prepared EBU-TT Part 1 document alongside a video programme, and inserting the subtitles into the video stream for downstream routing. A playout automation system initiates playback of the EBU-TT Part 1 document alongside the video content and a 'playback' node transforms the EBU-TT Part 1 document into a sequence of EBU-TT Part 3 subtitle documents, one per video frame⁹, for insertion into video, or other emission as a stream.

Example 5: Speaker based styling – not fully defined in this specification

An Improver Node receives a sequence in which each document has no styling data. The Improver analyses speaker ("agent") metadata within each document, and uses that in combination with a predefined style set to add the appropriate styles and style references to each document before issuing it as a new sequence with a new sequence identifier.

Example 6: Transcoding for distribution – not fully defined in this specification

An encoder receives a stream for a specific service, having been configured to subscribe to the node that emits a stream whose documents includes that service's destination service identifier. It accumulates documents in the sequence from the stream and transcodes them to the required output encoding, perhaps time-bounded samples of EBU-TT-D for packaging and distribution in MPEG-DASH, or DVB bitmaps for multiplexing in an MPEG-2 transport stream.

Example 7: Creation of an archive document – not fully defined in this specification

An 'archiver' node creates archive documents for each programme on a service. It receives a stream for a specific service, having been configured to subscribe to the node that emits a stream whose documents includes that service's destination service identifier. It accumulates those documents from the sequence that have the same broadcast programme identifier and combines them into a single EBU-TT Part 1 document for each programme, for storage in an archive.

Example 8: Resilient failover – not fully defined in this specification

A playout facility receives multiple physical streams (for redundancy) from different nodes, for a specific service, each with the same sequence identifier, since they originated from a single node. The facility will normally be configured to accept one of these streams for insertion into the video output, but on failure of that stream will select in turn an alternate stream in an order of preference.

⁹ NOTE: A similar alternative is to insert one EBU-TT Part 3 document per change in subtitles, or to insert an EBU-TT Part 3 document when a maximum period of time has elapsed since the previous document: this is an implementation choice that could be determined by the maximum permitted recovery time for the system, for example.

Example 9: Subtitle localisation for opt-outs – not fully defined in this specification

A single broadcast channel has multiple, local downstream 'opt-outs'¹⁰ for localisation purposes. A single subtitler authors the main stream, and other subtitlers author the 'opt-out' streams. Local downstream switching nodes facilitate the switching of the subtitle streams in time with the switch in video and audio streams.

¹⁰ An 'opt-out' is where a service overwrites part of another service's output for a limited period of time, for example to deliver local news for a particular region.

Annex D: Summary of document conformance differences relative to Part 1 (Informative)

The format for EBU-TT Part 3 is based on EBU-TT Part 1 with some relaxations and some additions. Processors need to be able to receive and parse those documents, and possibly to create new documents and stream those.

The conformance rules for documents defined in Part 3 are in § 3. Here's a quick summary of the differences compared to Part 1:

- The parameter attributes `sequenceIdentifier` and `sequenceNumber` are required on the `tt` element.
- `body` can be omitted to signal 'display nothing while this document is active'.
- `styling` and `layout` can be omitted and then added later.
- Timing (`begin` and `end` attributes) can be used on `body` and `div`, and don't have to be on `p`. Also, the `dur` attribute can be used on `body`.
- In case a studio or other clock is used, it can be identified with a `referenceClockIdentifier` parameter.
- The authoring delay between the subtitle hearing the programme audio and the subtitles being generated can be signalled if known using the `authoringDelay` metadata attribute.
- A set of parameters is available to support handover between different subtitlers working on the same programme output. They are `authorsGroupIdentifier`, `authorsGroupControlToken` and `authorsGroupControlRequest`.
- The source service from which the subtitles were generated can be signalled using the `originalSourceServiceIdentifier` metadata element.
- The service(s) for which the subtitles are intended can be signalled using the `intendedDestinationServiceIdentifier` metadata element.
- The processing history for a document can be captured for debugging purposes using the `trace` element.
- The `trace` element can be used to identify processing that occurred in the chain used to create the element, for example for debugging purposes.
- The `facet` element can be used to tag content with characteristics that apply or do not apply to it. These can be summarised at the document level using the `documentFacet` element.