

MMA Technical Standards Board/ AMEI MIDI Committee

Letter of Agreement for Recommend Practice

Audio Clips for Mobile XMF (RP-045)

Originated By: MMA	Reference MMA Item #: 190	TSBB #: 31
Issue Date: Feb 07 2007	Approval Date: Nov 30 2007	Last Revised: _____
Related Items: Mobile XMF Content Format Specification (RP-042) ID3 Data for XMF Files (RP-047)		
Agreed to by MMA: _____ <i>Signature</i>	_____ <i>title</i>	Date: _____
Agreed to by AMEI: _____ <i>Signature</i>	_____ <i>title</i>	Date: _____

Source: MMA Mobile WG, Beatnik

Abstract:

This specification defines an extension of the Mobile XMF v1.0 Specification that allows instruments to use longer mono or stereo 'audio clip' wavetables, and enables the use of various data-compressed audio encodings beyond linear PCM.

Background:

The Mobile XMF v1.0 specification, which has been widely deployed and is referenced by at least one other standards body, defines a structured music and audio content format for use in resource-constrained mobile devices such as mobile phones. It allows custom wavetable-based Mobile DLS instrument definitions to be bundled together with a Standard MIDI File (SMF) image containing a musical performance played on a combination of the custom instruments and a standard palette of General MIDI instruments. However in Mobile XMF all wavetables must fit into the DLS synthesizer's wavetable memory, which is extremely limited (minimum of 7k or 15k bytes).

This specification allows much longer mono or stereo audio wavetables or 'audio clips' to be triggered in sync with the MIDI timeline. Audio clip instruments are intended for straight audio playback, and so do not allow most of the Mobile DLS sound shaping controls to be used. They are not polyphonic (consume 1 mono or stereo voice at most), and may not be changed in pitch, filtered, nor modulated in any way other than ordinary volume setting and, for mono wavetables rendered on stereo devices, stereo pan setting. This makes it possible to render audio clips either from within the MIDI synthesizer, or with external audio file player hardware or software.

This level of playback functionality is suitable for integrating sections of recorded music together with MIDI rendered notes, which is a common and appropriate production technique for producing full-length branded music content (as well as other content styles) at file sizes much smaller than conventionally data-compressed recorded music.

Publication Plan:

The specification will be published as a stand-alone document, or bundled with other related specification documents.

Proposal (Details):

See attached.

Audio Clips for Mobile XMF

February 7, 2007

Published By:

The MIDI Manufacturers Association

Los Angeles, CA

PREFACE

AUDIO CLIPS FOR MOBILE XMF
RP-045

Copyright © 2006 MIDI Manufacturers Association Incorporated

ALL RIGHTS RESERVED. NO PART OF THIS DOCUMENT MAY BE REPRODUCED OR TRANSMITTED IN ANY FORM OR BY ANY MEANS, ELECTRONIC OR MECHANICAL, INCLUDING INFORMATION STORAGE AND RETRIEVAL SYSTEMS, WITHOUT PERMISSION IN WRITING FROM THE MIDI MANUFACTURERS ASSOCIATION.

Printed 2007

MMA
PO Box 3173
La Habra CA 90632-3173

CONTENTS

1. Introduction 1

 1.1 Overview 1

 1.2 Technical Goals 1

 1.3 Terminology 2

2. Content Format and Player Behavior 3

 2.1 File Layout and Supported Resource Types 3

 2.2 Interoperable SMF Content 3

 2.3 Interoperable Instrument & Audio Clip Content 3

 2.3.1 Instrument Data Format 3

 2.3.2 Constraints on Instrument Content 3

 2.3.2.1 Definition and Format of Audio Clip Instruments 3

 2.3.2.2 Player Behavior for Audio Clip Instruments 3

 2.3.2.3 Maximum Wavetable Size for Audio Clip Instruments 4

 2.3.2.4 Maximum Wavetable Data Bit Rate for Audio Clip Instruments 5

 2.3.2.5 Audio Encodings for Audio Clip Instruments 5

 2.3.2.6 Rendering Tolerances for Audio Clip Instruments 5

 2.4 How Instruments and Audio Clips are Accessed with MIDI Messages 5

 2.4.1 Accessing Audio Clip Instruments 5

 2.4.2 Playing Audio Clips 6

 2.5 Use of XMF Meta File Format Features 6

 2.6 Content Handling Behaviors 6

 2.7 Content Type Identification 6

 2.7.1 Internal Content Type Indication 6

 2.7.2 External Content Type Indication 7

 2.8 Content Description Metadata 7

3. Player Implementation Guidelines 9

4. Content Authoring Guidelines 10

 4.1 Overview 10

 4.2 Mobile DLS Content Authoring for Audio Clips 10

 4.3 SMF File Authoring for Audio Clips 11

 4.3.1 Audio Clip Polyphony 11

 4.3.2 Selecting Audio Clip Instruments 11

 4.3.3 Triggering Audio Clips 11

5. Authoring Tool Guidelines 12

6. References 13

AUDIO CLIPS FOR MOBILE XMF

1. Introduction

This specification defines an extension of Mobile XMF v1.0 [2] that allows instruments to use longer mono or stereo 'audio clip' wavetables, and enables the use of various data-compressed audio encodings beyond linear PCM (for example: AAC, aacPlus, Enhanced aacPlus, AMR-WB, Extended AMR Wideband, IMA ADPCM, MP3, etc.).

1.1 Overview

The Mobile XMF v1.0 specification, which has been widely deployed and is referenced by at least one other standards body [6], defines a structured music and audio content format for use in resource-constrained mobile devices such as mobile phones. It allows custom wavetable-based Mobile DLS [3] instrument definitions to be bundled together with a Standard MIDI File (SMF) [7] image containing a musical performance played on a combination of the custom instruments and a standard palette of General MIDI instruments. However in Mobile XMF all wavetables must fit into the DLS synthesizer's wavetable memory, which is extremely limited (minimum of 7k or 15k bytes).

This specification allows much longer mono or stereo audio wavetables or 'audio clips' to be triggered in sync with the MIDI timeline. Audio clip instruments are intended for straight audio playback, and so do not allow most of the Mobile DLS sound shaping controls to be used. They are not polyphonic (consume 1 mono or stereo voice at most), and may not be changed in pitch, filtered, nor modulated in any way other than ordinary volume setting and, for mono wavetables rendered on stereo devices, stereo pan setting. This makes it possible to render audio clips either from within the MIDI synthesizer, or with external audio file player hardware or software.

This level of playback functionality is suitable for integrating sections of recorded music together with MIDI rendered notes, which is a common and appropriate production technique for producing full-length branded music content (as well as other content styles) at file sizes much smaller than conventionally data-compressed recorded music.

1.2 Technical Goals

The following technical goals governed the development of this specification:

- 1 - Extend Mobile XMF by adding streamed, MIDI-synchronized mono or stereo linear audio clips using data-compressed audio encodings. This means short- to mid-length audio clips synced to a MIDI timeline, rather than long audio clips that would be continuously synchronized with the SMF timeline as a whole.
- 2 - Make this extended format playable on the widest possible set of playback devices.
- 3 - Easy content authoring.
- 4 - Avoid unnecessary complexity.
- 5 - Place only conservative (bandwidth) demands on the audio clip data streaming channel.
- 6 - Need a scalable polyphony approach for audio clips for broad content interoperability
- 7 - It should be possible to play this content both on players that can only handle one audio clip at a time, and on players that can play more than one audio clip at a time, and still have the content do the right thing in all cases. We don't have a typical upper limit in mind for audio clip polyphony.
- 8 - No backwards compatibility with Mobile XMF v1.0. [Note: Upon inspection this turned out to be simply too cumbersome (codec selection, memory sizes, providing fall-back instruments, & impact on data organization).]

1.3 Terminology

AMEI – Association for Musical Electronic Industry, in Japan. AMEI includes a MIDI Committee and a Mobile Division, which serve the same functions in Japan that the MMA serves in the rest of the world.

Audio Clip Instrument – Mobile DLS instrument intended for use as an audio clip; not polyphonic, self-exclusive, uses mono or stereo wavetables.

Audio Clip Polyphony – Number of simultaneously playing audio clips, or a player's maximum capacity in terms of same.

Audio Codec – Audio coder/decoder technology and associated audio data encoding format, for example linear PCM, AAC, aacPlus, Enhanced aacPlus, AMR-WB, Extended AMR Wideband, IMA ADPCM, MP3 (MPEG-2 Layer 3 Audio), etc..

Content Description Metadata (CDM) – Additional information in a Mobile XMF file describing the use of various resources per MIDI channel. Defined in section 10 of [2].

DLS – Downloadable Sounds, a family of MMA/AMEI content formats and player specifications for wavetable synthesizer musical instruments. See also Mobile DLS.

MIP Message – MIDI message defining the Maximum Instantaneous Polyphony per MIDI channel of a piece of scalable SMF content. Defined in the SP-MIDI specification [4].

MMA – MIDI Manufacturers Association

Mobile DLS – MMA/AMEI specification for wavetable synthesizer instruments for mobile applications. See [3].

SMF – Standard MIDI File, see [7].

SP-MIDI – Scalable Polyphony MIDI, an MMA/AMEI content format specification that adapts MIDI file playback to the available device resources. See [4] and [5].

2. Content Format and Player Behavior

This section defines the content format for Mobile XMF content that includes Audio Clips, and the required behavior of players supporting such content.

IMPORTANT: Except as detailed below, the content format and player behavior requirements of Mobile XMF content that includes Audio Clips are **identical to Mobile XMF v1.0** [2].

For easier developer reference, sections 2.1 through 2.8 of this specification correspond directly to sections 3 through 10 of the Mobile XMF v1.0 Specification [2].

2.1 File Layout and Supported Resource Types

The file layout and supported resource types for Mobile XMF content that includes Audio Clips is the same as for Mobile XMF v1.0 (see section 3 of [2]).

2.2 Interoperable SMF Content

The rules regarding interoperable SMF for Mobile XMF content that includes Audio Clips are the same as for Mobile XMF v1.0 (see section 3 of [2]). However, the operation of certain MIDI messages contained within the SMF is different when an audio clip instrument is selected, see section 2.4.

2.3 Interoperable Instrument & Audio Clip Content

With the following exceptions, the interoperable instrument (including audio clip instrument) content rules are the same as for Mobile XMF v1.0 (see section 5 of [2]).

2.3.1 *Instrument Data Format*

Except as detailed below, the bundled instrument and audio clip content must be compliant with the Mobile DLS specification v1.0 [3].

2.3.2 *Constraints on Instrument Content*

Except as detailed below, the bundled Mobile DLS instrument and audio clip content is subject to the same constraints detailed in the Mobile DLS specification v1.0 [3].

2.3.2.1 Definition and Format of Audio Clip Instruments

For Mobile XMF content that includes Audio Clips, all Mobile DLS Instrument programs in MIDI bank MSB 0x7A, LSB 0x00 are considered to be audio clip instruments. The data format for audio clip instruments is identical to ordinary Mobile DLS instruments [3], however wavetables for audio clip instruments may be encoded for different audio decoders than wavetables for ordinary Mobile DLS instruments.

Banks MSB 0x7A, LSB 0x01 through MSB 0x7A, LSB 0x7F are reserved for future MMA/AMEI definition and should not be used in Mobile XMF content that includes Audio Clips.

2.3.2.2 Player Behavior for Audio Clip Instruments

This section consists of normative requirements for implementers of players for Mobile XMF content that includes Audio Clips. See also the non-normative information at section 3. Player Implementation Guidelines.

AUDIO CLIPS FOR MOBILE XMF

The Mobile XMF player's behavior for MIDI channels on which an audio clip instrument has been selected differs from its behavior for MIDI channels on which ordinary Mobile DLS instruments have been selected in the following ways.

When handling audio clip instruments, players **MUST**:

- Play audio clip instruments without any pitch transposition other than sampling rate conversion from the input wavetable data sampling rate to the player audio output sample rate.
- Support at least all of the following audio output sample rates: 8 kHz, 11.025 kHz, 12 kHz, 16 kHz, 22.05 kHz, 24 kHz, 32 kHz, 44.1 kHz, or 48 kHz.
- Take into account the fact that audio clip instruments may use larger wavetables than ordinary Mobile DLS instruments (see section 2.3.2.3).
- Take into account the fact that audio clip instruments may include wavetables encoded for different audio decoders than ordinary Mobile DLS instruments (see section 2.3.2.5).
- Take into account the fact that audio clip instruments may include stereo wavetables, and always play stereo wavetables as stereo output whenever the Mobile XMF player is capable of producing stereo output.
- Always render audio clip instruments with a maximum polyphony of 1 and self-exclusive. Re-triggering an audio clip instrument (by sending a new MIDI Note On message) must cut off wavetable playback and restart it from the beginning of the wavetable, rather than adding a second, layered, copy.
- Never loop playback of the wavetable audio data. Any loop information included in the audio wavetable in the Mobile DLS data must be ignored.
- Always treat the Velocity field of the Note On MIDI message as a Volume control.
- Follow the MMA's standard MIDI volume response curve, which is produced by the DLS 2 Concave transform. The formula for gain in terms of -dB is given at section 1.6.5.4 of [3] as:
$$\text{gain} = -96 * (5/12) * \log_{10}(127 / (127 - (127 - \text{velocity})))$$
- Always ignore all filter and modulation connections with destinations other than Volume and, for mono wavetables rendered on stereo devices, Pan.
- For stereo source wavetables, always ignore all connections whose destination is Pan .

Notwithstanding the above differences, players must still honor the WSMP chunk's Gain Trim field.

2.3.2.3 Maximum Wavetable Size for Audio Clip Instruments

For audio clip instruments, Mobile DLS data used in Mobile XMF content that includes Audio Clips may use wavetables significantly larger than those the Mobile DLS v1.0 Specification allows. The actual maximum wavetable size depends on the context. For example, one product may allow a higher maximum than another product. Also, a document or profile that references this Audio Clips for Mobile XMF specification may set a specific maximum wavetable size.

2.3.2.4 Maximum Wavetable Data Bit Rate for Audio Clip Instruments

For audio clip instruments, Mobile DLS data used in Mobile XMF content that includes Audio Clips may have different wavetable bit rate limitations than those the Mobile DLS v1.0 Specification allows. The actual maximum wavetable data bit rate depends on the context. For example, one product may allow a higher maximum than another product. Also, a document or profile that references this Audio Clips for Mobile XMF specification may set a specific maximum wavetable data bit rate.

2.3.2.5 Audio Encodings for Audio Clip Instruments

For audio clip instruments, all required audio decoders must be indicated using the mechanism described in section 2.16.1 of [3]. In addition to defining several wFormatTags for existing audio codecs, this mechanism allows for the registration of new audio encodings via the wFormatTag value WAVE_FORMAT_EXTENSIBLE.

For audio clip instruments, Mobile DLS data used in Mobile XMF content that includes Audio Clips may use additional audio encodings not already listed in the MMA codec GUID registry referenced from section 2.16.1 of [3].

The actual set of available decoders depends on the context. For example, one product may allow a different decode, or set of decoders, than another product. Also, a document or profile that references this Audio Clips for Mobile XMF specification may specify a particular codec, or set of codecs.

2.3.2.6 Rendering Tolerances for Audio Clip Instruments

Whether audio clips are rendered inside the Mobile DLS synthesizer or outside of it, implementations must observe the following tolerances:

- Instrument audio level and gain control (CC7) response for audio clip instruments should match ordinary Mobile DLS and GM instruments as closely as possible.
- Wavetable start times of audio clip instruments must be synchronized with ordinary Mobile DLS or GM instruments. In other words, two notes with start times scheduled for the same SMF tick, where one is played on an audio clip instrument and the other one is played on an ordinary Mobile DLS or GM instrument, must be started as close to simultaneously as possible. It is the player's responsibility to automatically correct for any differences in latency between the sources.
- Drift between the MIDI and audio timelines must not exceed 1 part in 1000.

2.4 How Instruments and Audio Clips are Accessed with MIDI Messages

With the following exceptions, the rules regarding how instruments are accessed with MIDI messages for Mobile XMF content that includes Audio Clips are the same as for Mobile XMF v1.0 (see section 6 of [2]).

2.4.1 Accessing Audio Clip Instruments

Audio Clip instruments are accessed in the same way as any other Mobile DLS instrument, i.e. a MIDI channel can select the audio clip instrument by using the following sequence of MIDI messages: Bank Select MSB message, Bank Select LSB message, and finally the Program Change message. However bank MSB 0x7A, LSB 0x00 is reserved for audio clip instruments, and all audio clip instruments must be loaded only from this bank. The program number used in the Program Change message must match an audio clip instrument present in the bundled Mobile DLS content.

Banks MSB 0x7A, LSB 0x01 through MSB 0x7A, LSB 0x7F are reserved for future MMA/AMEI definition and should not be used in Mobile XMF content that includes Audio Clips.

2.4.2 *Playing Audio Clips*

Once an audio clip instrument has been selected for a MIDI channel, the audio clip can be played on that MIDI channel by sending a MIDI Note On message. The audio clip will play until it reaches the end of the wavetable data, or until a matching MIDI Note Off message (or Note On with velocity of zero) is sent on the same MIDI channel.

Player response to MIDI messages when an audio clip instrument is selected is different from other Mobile DLS instruments in the following ways:

- The note number field of the MIDI Note On message will not affect audio clip pitch. The audio clip will always play at its natural pitch, even if the Mobile DLS instrument program calls for different pitch behavior.
- The note number field of the MIDI Note On message serves solely as a note ID. A MIDI Note Off message (or Note On with velocity of zero) terminates an audio clip note when its note number field matches the one used in a previous MIDI Note On message.
- The Velocity field of the Note On message will control the clip volume, according to the response specified for the Mobile DLS 1.0 default velocity-to-gain connection.
- If the selected audio clip is already playing on the indicated MIDI channel at the time of the MIDI Note On message, the audio clip will 'rewind' wavetable playback to the start of the wavetable. Audio clip instruments are always non-polyphonic and 'self-exclusive', even if the Mobile DLS instrument program calls for polyphonic behavior.
- The audio clip instrument will not respond to the MIDI Pitch Bend message, even if the Mobile DLS instrument program calls for it.
- The audio clip will respond to the MIDI Polyphonic Key Pressure or Channel Pressure messages only if the selected instrument program route them to a supported destination, i.e. clip volume, or for mono wavetables rendered on stereo devices, pan.
- In addition to MIDI Bank Select MSB & LSB, the audio clip instrument will respond to the MIDI Control Change message only if the controller number is either 7 (Volume) or 11 (Expression), or for mono wavetables rendered on a stereo device, 10 (Pan). No other control numbers will produce any audible effect, even if the selected Mobile DLS instrument program calls for it.

2.5 **Use of XMF Meta File Format Features**

The rules regarding the use of XMF Meta File Format features for Mobile XMF content that includes Audio Clips are the same as for Mobile XMF v1.0 (see section 7 of [2]) with the following exception. Use of the ID3 Metadata field (XMF Metadata Standard FieldID14, as defined in [8]) is allowed.

2.6 **Content Handling Behaviors**

The rules regarding content handling behavior for Mobile XMF content that includes Audio Clips are the same as for Mobile XMF v1.0 (see section 8 of [2]).

2.7 **Content Type Identification**

With the following exceptions, the rules regarding content type identification for Mobile XMF content that includes Audio Clips are the same as for Mobile XMF v1.0 (see section 9 of [2]).

2.7.1 *Internal Content Type Indication*

To provide content type information to XMF file parsers and generic file type recognizers, the XmfFileTypeID field in the FileHeader must indicate the proper XMF File Type number, and the XmfFileTypeRevisionID must indicate the proper spec version:

AUDIO CLIPS FOR MOBILE XMF

XMF File Type: 3

Spec Revision Level: 0

2.7.2 External Content Type Indication

Filename extension and MIME media type are the same as for Mobile XMF 1.0 content; see section 9 of [2].

2.8 Content Description Metadata

With the following exceptions, Content Description Metadata (CDM) for Mobile XMF content that includes Audio Clips is the same as for Mobile XMF 1.0 content: see section 10 of [2].

For Mobile XMF content that includes Audio Clips, CDM is extended in the following ways.

The following table extends the table in section 10.3.5 Playback Resource Group List of [2]:

PlaybackResourceGroupID	Group Name	Mutual Exclusivity	Channel Masking
3	Audio Clip Resources with Mutual Exclusivity	Yes	No
4	Audio Clip Resources without Mutual Exclusivity	No	No
5	Audio Clip Voices	Yes	Yes

Five new Standard PlaybackResourceID values are available when a Mobile XMF file includes audio clip instruments. The following table extends the table in section 10.3.4.1 Standard PlaybackResourceIDs of [2]:

PlaybackResourceID Value	Description	For Use with PlaybackResourceGroupID
5	Number of Audio Clips voices	5: Audio clip voices
6	Total encoded wavetable data (in kilobytes) for Audio Clips	3: Audio Clip Resources with Mutual Exclusivity
7	Audio clip maximum sample rate (in kilosamples per 10 seconds)	4: Audio Clip Resources without Mutual Exclusivity
8	Audio clip average bit rate (in kilobits per 10 seconds)	3: Audio Clip Resources with Mutual Exclusivity
9	Audio clip maximum bit rate (in kilobits per 10 seconds)	3: Audio Clip Resources with Mutual Exclusivity

A MIDI Channel with a non-zero entry in a MIR column whose PRL entry is { 00h, 05h } (Audio Clip Voices) is an audio clip channel. For Mobile XMF content that includes audio clips, this column **MUST** always be present. Because audio clip polyphony is limited to 1 in v1.0 of the Audio Clips for Mobile XMF specification, the maximum value in this column is 1.

AUDIO CLIPS FOR MOBILE XMF

NOTE: If any audio clips use any non-linear PCM codecs, these codec(s) MUST be indicated in MIR column(s) with PRL entry { 04x, <wFormatTag> } or { 05h, <GUID> }. A non-zero entry in such a MIR column indicates that that MIDI channel uses that codec. Because this column uses the exact same encodings as an ordinary Mobile DLS synthesizer voice, it cannot be used to distinguish between a Mobile DLS synthesizer voice and an audio clip voice; to distinguish between the two, use the separate { 00h, 05h } column instead.

For Mobile XMF content that includes audio clips, MIR columns with the following PRL entries SHOULD be used: { 00h, 06h }, { 00h, 07h }, { 00h, 08h }, and { 00h, 09h }.

3. Player Implementation Guidelines

This section consists of non-normative information for implementers of players for Mobile XMF content that includes Audio Clips. See also the normative requirements at section 2.3.2.2 Player Behavior for Audio Clip Instruments.

When handling audio clip instruments, players:

- May optionally render the audio clip instruments using a separate audio file playback facility (hardware or software), in other words not necessarily using the music synthesizer rendering the other Mobile DLS instruments. Players implemented in this manner must meet certain minimum performance requirements.
- May optionally play the audio clip instrument by streaming the wavetable audio data from secondary storage such as a local file system, rather than rendering directly from wavetable audio data stored wholly within the synthesizer's wavetable memory. Players offering this functionality must meet certain minimum performance requirements.

4. Content Authoring Guidelines

This section provides guidance to content authors producing Mobile XMF content that includes Audio Clip instruments.

4.1 Overview

The Audio Clips for Mobile XMF specification defines an extension of Mobile XMF v1.0 [2] that allows the use of instruments with longer mono or stereo 'audio clip' wavetables which may use various data-compressed audio encodings (for example: AAC, aacPlus, Enhanced aacPlus, AMR-WB, Extended AMR Wideband, IMA ADPCM, MP3, etc.). This means short audio snippets synced to a MIDI timeline, not very long audio playing in parallel with very long MIDI timeline. Composer control is limited to starting and stopping audio clips, plus volume and (for mono wavetables rendered on stereo devices) pan. This level of playback functionality is suitable for integrating sections of recorded music together with MIDI rendered notes, which is a common and appropriate production technique for producing full-length branded music content as well as other content styles at file sizes much smaller than conventionally data-compressed recorded music.

See also: Section 1.1. Overview.

4.2 Mobile DLS Content Authoring for Audio Clips

Wavetable content intended for use as audio clips goes into Mobile XMF files in the form of Mobile DLS instruments. Audio Clip instruments have exactly the same format as any other Mobile DLS instrument, and the player handles them exactly the same, except for the following rules which you should keep in mind when creating audio clip content:

- Audio clip instruments must appear in bank MSB 0x7A, LSB 0x00. Do not use banks MSB 0x7A, LSB 0x01 through MSB 0x7A, LSB 0x7F as they are reserved for future MMA/AMEI definition.
- Audio clip wavetables are never transposed to different pitches/notes, they are always played at the wavetable's natural pitch. The note number field of the MIDI Note On message does not affect pitch, the Pitch Bend message is ignored, and the Pitch Fine Tuning field of the WSMP chunk is ignored. So to minimize file size and avoid wasting space, don't include modulation connections with the Pitch destination.
- The Pitch Fine Tuning field of the WSMP chunk is ignored. If you need to tune the wavetable, use an audio editor program before importing the sample into the Mobile DLS file.
- Audio clip instruments can use significantly larger wavetables than ordinary Mobile DLS instruments. They don't have to fit into the synthesizer's wavetable memory, but can be streamed from secondary storage such as a local file system. See section 2.3.2.3 for size guidance.
- Audio clip wavetables may have data bit rate maximums. (This is distinct from the audio sampling rate.) See section 2.3.2.4 for bit rate guidance.
- Audio clip wavetables may use different audio encodings than ordinary Mobile DLS instruments, and therefore may require different decoders in the player. See section 2.3.2.5 for codec selection guidance.
- Audio clip wavetables can't loop. Any loop information included in the audio clip wavetable data will be ignored. So to minimize file size, don't include loop information in your wavetables.
- Audio clip instruments always treat the Velocity field of the MIDI Note On message as a Volume control. You can't prevent or override this by providing different modulation connections in the instrument definition.

AUDIO CLIPS FOR MOBILE XMF

- To trim an instrument's gain, use the WSMP chunk's Gain Trim field.
- Audio clip instruments never use the lowpass filter. So don't create instruments that depend on the filter. And to minimize file size, don't include any connections that set or modulate the filter parameters.
- Audio clip instruments will ignore any modulation connections with destinations other than Volume, and if the wavetable is mono and the device is stereo, Pan. So to minimize file size, don't include modulation connections with any destination other than Volume and Pan, and don't use Pan if the instrument wavetable is stereo.

4.3 SMF File Authoring for Audio Clips

Audio clips will not be played unless the SMF inside the Mobile XMF file includes the necessary MIDI messages. To play an audio clip, you must devote a MIDI channel to it, select the Mobile DLS audio clip instrument program that references the desired audio clip wavetable, and then send MIDI Note On and Note Off messages on that MIDI channel to start and stop the clip.

4.3.1 Audio Clip Polyphony

Audio clip polyphony is limited to 1. When an audio clip instrument is playing, any subsequently started audio clip will replace the first one.

4.3.2 Selecting Audio Clip Instruments

See section 2.4.1.

4.3.3 Triggering Audio Clips

See section 2.4.2.

5. Authoring Tool Guidelines

This section provides guidance to Mobile XMF authoring tool developers who wish to support content that includes Mobile DLS Audio Clip instruments.

- All Mobile DLS audio clip instruments must appear in bank MSB 0x7A, LSB 0x00. Do not use banks MSB 0x7A, LSB 0x01 through MSB 0x7A, LSB 0x7F as they are reserved for future MMA/AMEI definition.
- All wavetables used in audio clip instruments use the WAVE_FORMAT_EXTENSIBLE mechanism to identify their audio data encoding, and consequently the required decoder. However it is expected that some application domains for Audio Clips for Mobile XMF (i.e. format or standards specifications that refer to this specification but define a unique decoder set) may specify support for different sets of audio decoders for audio clips. Authoring tools should therefore guide content creators away from the use of audio decoders not appropriate for the intended end deployment of the content.
- Audio clip instruments only respond to the Volume connection, and (if the wavetable is mono and the device is stereo) the Pan connection, so to save space any connections with other destinations should be omitted (perhaps stripped) from the final Mobile DLS and/or Mobile XMF content.
- Volume and Pan have defined default connections, so if the content creator only uses default values, there is no need to include these connections in the final Mobile DLS and/or Mobile XMF content either.

6. References

- [1] “*MIDI 1.0 Detailed Specification, Document Version 4.2*”, in “The Complete MIDI 1.0 Detailed Specification, Document Version 96.1”, The MIDI Manufacturers Association Inc., Los Angeles, CA, USA, February 1996
- [2] “*Mobile XMF Content Format Specification*”, RP-042, MIDI Manufacturers Association, Los Angeles, CA, USA, 2004
- [3] “*Mobile DLS Specification*”, RP-041, MIDI Manufacturers Association, Los Angeles, CA, USA, 2004
- [4] “*Scalable Polyphony MIDI Specification, Version 1.0*” (SP-MIDI), RP-034, MIDI Manufacturers Association, Los Angeles, CA, USA, February 2002
- [5] “*Scalable Polyphony MIDI Device 5–24 Note Profile for 3GPP, Version 1.0*”, RP-035, MIDI Manufacturers Association, Los Angeles, CA, USA, February 2002
- [6] “3GPP Technical Specification 26.234 for Release 6 : *End-to-end transparent streaming service; Protocols and codecs*”, http://www.3gpp.org/ftp/Specs/archive/26_series/26.234/26234-680.zip. The 3rd Generation Partnership Project, Sophia Antipolis, France, 2004
- [7] “*Standard MIDI Files*”, RP-001, in The Complete MIDI 1.0 Detailed Specification, Document Version 96.1, MIDI Manufacturers Association, Los Angeles, CA, USA, 1996
- [8] “*ID3 Metadata for XMF Files*”, RP-047, MIDI Manufacturers Association, Los Angeles, CA, USA, 2007