

WAVE File Format

WAVE File Format is a file format for storing digital audio (waveform) data. It supports a variety of bit resolutions, sample rates, and channels of audio. This format is very popular upon IBM PC (clone) platforms, and is widely used in professional programs that process digital audio waveforms. It takes into account some peculiarities of the Intel CPU such as little endian byte order.

This format uses Microsoft's version of the Electronic Arts Interchange File Format method for storing data in "chunks".

Data Types

A C-like language will be used to describe the data structures in the file. A few extra data types that are not part of standard C, but which will be used in this document, are:

pstring Pascal-style string, a one-byte count followed by that many text bytes. The total number of bytes in this data type should be even. A pad byte can be added to the end of the text to accomplish this. This pad byte is not reflected in the count.

ID A chunk ID (ie, 4 ASCII bytes).


Also note that when you see an array with no size specification (e.g., `char ckData[];`), this indicates a variable-sized array in our C-like language. This differs from standard C arrays.

Constants

Decimal values are referred to as a string of digits, for example 123, 0, 100 are all decimal numbers. Hexadecimal values are preceded by a 0x - e.g., 0x0A, 0x1, 0x64.

Data Organization

All data is stored in 8-bit bytes, arranged in Intel 80x86 (ie, little endian) format. The bytes of multiple-byte values are stored with the low-order (ie, least significant) bytes first. Data bits are as follows (ie, shown with bit numbers on top):




File Structure


A WAVE file is a collection of a number of different types of chunks. There is a required Format ("fmt ") chunk which contains important parameters describing the waveform, such as its sample rate. The Data chunk, which contains the actual waveform data, is also required. All other chunks are optional. Among the other optional chunks are ones which define cue points, list instrument parameters, store application-specific information, etc. All of these chunks are described in detail in the following sections of this document.

All applications that use WAVE must be able to read the 2 required chunks and can choose to selectively ignore the optional chunks. A program that copies a WAVE should copy all of the chunks in the WAVE, even those it chooses not to interpret.

There are no restrictions upon the order of the chunks within a WAVE file, with the exception that the Format chunk must precede the Data chunk. Some inflexibly written programs expect the Format chunk as the first chunk (after the RIFF header) although they shouldn't because the specification doesn't require this.

Here is a graphical overview of an example, minimal WAVE file. It consists of a single WAVE containing the 2 required chunks, a Format and a Data Chunk.





A Bastardized Standard

The WAVE format is sort of a bastardized standard that was concocted by too many "cooks" who didn't properly coordinate the addition of "ingredients" to the "soup". Unlike with the AIFF standard which was mostly designed by a small, coordinated group, the WAVE format has had all manner of much-too-independent, uncoordinated aberrations inflicted upon it. The net result is that there are far too many chunks that may be found in a WAVE file -- many of them duplicating the same information found in other chunks (but in an unnecessarily different way) simply because there have been too many programmers who took too many liberties with unilaterally adding their own additions to the WAVE format without properly coming to a consensus of what everyone else needed (and therefore it encouraged an "every man for himself" attitude toward adding things to this "standard"). One example is the Instrument chunk versus the Sampler chunk. Another example is the Note versus Label chunks in an Associated Data List. I don't even want to get into the totally irresponsible proliferation of compressed formats. (ie, It seems like everyone and his pet Dachshound has come up with some compressed version of storing wave data -- like we need 100 different ways to do that). Furthermore, there are lots of inconsistencies, for example how 8-bit data is unsigned, but 16-bit data is signed.

I've attempted to document only those aspects that you're very likely to encounter in a WAVE file. I suggest that you concentrate upon these and refuse to support the work of programmers who feel the need to deviate from a standard with inconsistent, proprietary, self-serving, unnecessary extensions. Please do your part to rein in half-ass programming.


Sample Points and Sample Frames

A large part of interpreting WAVE files revolves around the two concepts of sample points and sample frames.


A sample point is a value representing a sample of a sound at a given moment in time. For waveforms with greater than 8-bit resolution, each sample point is stored as a linear, 2's-complement value which may be from 9 to 32 bits wide (as determined by the `wBitsPerSample` field in the Format Chunk, assuming PCM format -- an uncompressed format). For example, each sample point of a 16-bit waveform would be a 16-bit word (ie, two 8-bit bytes) where 32767 (0x7FFF) is the highest value and -32768 (0x8000) is the lowest value. For 8-bit (or less) waveforms, each sample point is a linear, unsigned byte where 255 is the highest value and 0 is the lowest value. Obviously, this signed/unsigned sample point discrepancy between 8-bit and larger resolution waveforms was one of those "oops" scenarios where some Microsoft employee decided to change the sign sometime after 8-bit wave files were common but 16-bit wave files hadn't yet appeared.

Because most CPU's read and write operations deal with 8-bit bytes, it was decided that a sample point should be rounded up to a size which is a multiple of 8 when stored in a WAVE. This makes the WAVE easier to read into memory. If your ADC produces a sample point from 1 to 8 bits wide, a sample point should be stored in a WAVE as an 8-bit byte (ie, unsigned char). If your ADC produces a sample point from 9 to 16 bits wide, a sample point should be stored in a WAVE as a 16-bit word (ie, signed short). If your ADC produces a sample point from 17 to 24 bits wide, a sample point should be stored in a WAVE as three bytes. If your ADC produces a sample point from 25 to 32 bits wide, a sample point should be stored in a WAVE as a 32-bit doubleword (ie, signed long). Etc.

Furthermore, the data bits should be left-justified, with any remaining (ie, pad) bits zeroed. For example, consider the case of a 12-bit sample point. It has 12 bits, so the sample point must be saved as a 16-bit word. Those 12 bits should be left-justified so that they become bits 4 to 15 inclusive, and bits 0 to 3 should be set to zero. Shown below is how a 12-bit sample point with a value of binary 10100010111 is formatted left-justified as a 16-bit word.




But note that, because the WAVE format uses Intel little endian byte order, the LSB is stored first in the wave file as so:




For multichannel sounds (for example, a stereo waveform), single sample points from each channel are interleaved. For example, assume a stereo (ie, 2 channel) waveform. Instead of storing all of the sample points for the left channel first, and then storing all of the sample points for the right channel next, you "mix" the two channels' sample points together. You would store the first sample point of the left channel. Next, you would store the first sample point of the right channel. Next, you would store the second sample point of the left channel. Next, you would store the second sample point of the right channel, and so on, alternating between storing the next sample point of each channel. This is what is meant by interleaved data; you store the next sample point of each of the channels in turn, so that the sample points that are meant to be "played" (ie, sent to a DAC) simultaneously are stored

contiguously.

The sample points that are meant to be "played" (ie, sent to a DAC) simultaneously are collectively called a **sample frame**. In the example of our stereo waveform, every two sample points makes up another sample frame. This is illustrated below for that stereo example.



For a monophonic waveform, a sample frame is merely a single sample point (ie, there's nothing to interleave). For multichannel waveforms, you should follow the conventions shown below for which order to store channels within the sample frame. (ie, Below, a single sample frame is displayed for each example of a multichannel waveform).



The sample points within a sample frame are packed together; there are no unused bytes between them. Likewise, the sample frames are packed together with no pad bytes.

Note that the above discussion outlines the format of data within an uncompressed data chunk. There are some techniques of storing compressed data in a data chunk. Obviously, that data would need to be uncompressed, and then it will adhere to the above layout.

The Format Chunk

The Format (fmt) chunk describes fundamental parameters of the waveform data such as sample rate, bit resolution, and how many channels of digital audio are stored in the WAVE.

```
#define FormatID 'fmt ' /* chunkID for Format Chunk. NOTE: There is a space at the end of this ID. */

typedef struct {
    ID          chunkID;
    long        chunkSize;

    short       wFormatTag;
    unsigned short wChannels;
    unsigned long dwSamplesPerSec;
    unsigned long dwAvgBytesPerSec;
    unsigned short wBlockAlign;
    unsigned short wBitsPerSample;

    /* Note: there may be additional fields here, depending upon wFormatTag. */
}
```

```
} FormatChunk;
```

The ID is always "fmt". The chunkSize field is the number of bytes in the chunk. This does not include the 8 bytes used by ID and Size fields. For the Format Chunk, chunkSize may vary according to what "format" of WAVE file is specified (ie, depends upon the value of wFormatTag).

WAVE data may be stored without compression, in which case the sample points are stored as described in **Sample Points and Sample Frames**. Alternately, different forms of compression may be used when storing the sound data in the Data chunk. With compression, each sample point may take a differing number of bytes to store. The wFormatTag indicates whether compression is used when storing the data.

If compression is used (ie, WFormatTag is some value other than 1), then there will be additional fields appended to the Format chunk which give needed information for a program wishing to retrieve and decompress that stored data. The first such additional field will be an unsigned short that indicates how many more bytes have been appended (after this unsigned short). Furthermore, compressed formats must have a Fact chunk which contains an unsigned long indicating the size (in sample points) of the waveform after it has been decompressed. There are (too) many compressed formats. Details about them can be gotten from Microsoft's web site.

If no compression is used (ie, wFormatTag = 1), then there are no further fields.

The wChannels field contains the number of audio channels for the sound. A value of 1 means monophonic sound, 2 means stereo, 4 means four channel sound, etc. Any number of audio channels may be represented. For multichannel sounds, single sample points from each channel are interleaved. A set of interleaved sample points is called a sample frame.

The actual waveform data is stored in another chunk, the Data Chunk, which will be described later.

The dwSamplesPerSec field is the sample rate at which the sound is to be played back in sample frames per second (ie, Hertz). The 3 standard MPC rates are 11025, 22050, and 44100 KHz, although other rates may be used.

The dwAvgBytesPerSec field indicates how many bytes play every second. dwAvgBytesPerSec may be used by an application to estimate what size RAM buffer is needed to properly playback the WAVE without latency problems. Its value should be equal to the following formula rounded up to the next whole number:

$$\text{dwSamplesPerSec} * \text{wBlockAlign}$$

The wBlockAlign field should be equal to the following formula, rounded to the next whole number:

$$\text{wChannels} * (\text{wBitsPerSample} \% 8)$$

Essentially, wBlockAlign is the size of a sample frame, in terms of bytes. (eg, A sample frame for a 16-bit mono wave is 2 bytes. A sample frame for a 16-bit stereo wave is 4 bytes. Etc).

The wBitsPerSample field indicates the bit resolution of a sample point (ie, a 16-bit waveform would have wBitsPerSample = 16).

One, and only one, Format Chunk is required in every WAVE.

Data Chunk

The Data (data) chunk contains the actual sample frames (ie, all channels of waveform data).

```
#define DataID 'data' /* chunk ID for data Chunk */
typedef struct {
    ID          chunkID;
    long        chunkSize;

    unsigned char waveformData[];
} DataChunk;
```

The ID is always **data**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size (ie, chunkSize is the number of remaining bytes in the chunk after the chunkSize field, not counting any trailing pad byte).

Remember that the bit resolution, and other information is gotten from the Format chunk.

The following discussion assumes uncompressed data.

The waveformData array contains the actual waveform data. The data is arranged into what are called *sample frames*. For more information on the arrangement of data, see "Sample Points and Sample Frames".

You can determine how many bytes of actual waveform data there is from the Data chunk's chunkSize field. The number of sample frames in waveformData is determined by dividing this chunkSize by the Format chunk's wBlockAlign.

The Data Chunk is required. One, and only one, Data Chunk may appear in a WAVE.

Another way of storing waveform data

So, you're thinking "This WAVE format isn't that bad. It seems to make sense and there aren't all that many inconsistencies, duplications, and inefficiencies". You fool! We're just getting started with our first excursion into unnecessary inconsistencies, duplications, and inefficiency.

Sure, countless brain-damaged programmers have inflicted literally dozens of compressed data formats upon the Data chunk, but apparently someone felt that even this wasn't enough to make your life difficult in trying to support WAVE files. No, some half-wit decided that it would be a good idea to screw around with storing waveform data in something other than one Data chunk. NOOOOOOOOOOOOOOOO!!!!!!

For some god-forsaken reason, someone came up with the idea of using an imbedded IFF List inside of the WAVE file. NOOOOOOOOOOOOOOOO!!!!!! And this "Wave List" would contain multiple 'data' and 'slnt' chunks. NOOOOOOOOOOOOOOOO!!!! The Type ID for this List is 'wav!'.

I strongly suggest that you refuse to support any WAVE file that exhibits this Wave List nonsense. There's no need for it, and hopefully, the misguided programmer who conjured it up will be embarrassed into hanging his head in shame when nobody agrees to support his foolishness. Just say "NOOOOOOOOOOOOOOOO!!!!!"

Cue Chunk

The Cue chunk contains one or more "cue points" or "markers". Each cue point references a specific offset within the waveformData array, and has its own CuePoint structure within this chunk.

In conjunction with the Playlist chunk, the Cue chunk can be used to store looping information.

CuePoint Structure

```
typedef struct {
    long    dwIdentifier;
    long    dwPosition;
    ID      fccChunk;
    long    dwChunkStart;
    long    dwBlockStart;
    long    dwSampleOffset;
} CuePoint;
```

The dwIdentifier field contains a unique number (ie, different than the ID number of any other CuePoint structure). This is used to associate a CuePoint structure with other structures used in other chunks which will be described later.

The dwPosition field specifies the position of the cue point within the "play order" (as determined by the Playlist chunk. See that chunk for a discussion of the play order).

The fccChunk field specifies the chunk ID of the Data or Wave List chunk which actually contains the waveform data to which this CuePoint refers. If there is only one Data chunk in the file, then this field is set to the ID 'data'. On the other hand, if the file contains a Wave List (which can contain both 'data' and 'slnt' chunks), then fccChunk will specify 'data' or 'slnt' depending upon in which type of chunk the referenced waveform data is found.

The dwChunkStart and dwBlockStart fields are set to 0 for an uncompressed WAVE file that contains one 'data' chunk. These fields are used only for WAVE files that contain a Wave List (with multiple 'data' and 'slnt' chunks), or for a compressed file containing a 'data' chunk. (Actually, in the latter case, dwChunkStart is also set to 0, and only dwBlockStart is used). Again, I want to emphasize that you can avoid all of this unnecessary crap if you avoid hassling with compressed files, or Wave Lists, and instead stick to the sensible basics.

The dwChunkStart field specifies the byte offset of the start of the 'data' or 'slnt' chunk which actually contains the waveform data to which this CuePoint refers. This offset is relative to the start of the first chunk within the Wave List. (ie, It's the byte offset, within the Wave List, of where the 'data' or 'slnt' chunk of interest appears. The first chunk within the List would be at an offset of 0).

The dwBlockStart field specifies the byte offset of the start of the block containing the position. This offset is relative to the start of the waveform data within the 'data' or 'slnt' chunk.

The dwSampleOffset field specifies the sample offset of the cue point relative to the start of the block. In an uncompressed file, this equates to simply being the offset within the waveformData array. Unfortunately, the WAVE documentation is much too ambiguous, and doesn't define what it means by the term "sample offset". This could mean a byte offset, or it could mean counting the sample points (for example, in a 16-bit wave, every 2 bytes would be 1 sample point), or it could even mean sample frames (as the loop offsets in AIFF are specified). Who knows? The guy who conjured up the Cue chunk certainly isn't saying. I'm assuming that it's a byte offset, like the above 2 fields.

Cue Chunk

```
#define CueID 'cue' /* chunk ID for Cue Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;

    long    dwCuePoints;
    CuePoint points[];
} CueChunk;
```

The ID is always **cue**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields.

The dwCuePoints field is the number of CuePoint structures in the Cue Chunk. If dwCuePoints is not 0, it is followed by that many CuePoint structures, one after the other. Because all fields in a CuePoint structure are an even number of bytes, the length of any CuePoint will always be even. Thus, CuePoints are packed together with no unused bytes between them. The CuePoints need not be placed in any particular order.

The Cue chunk is optional. No more than one Cue chunk can appear in a WAVE.

Playlist chunk

The Playlist (plst) chunk specifies a play order for a series of cue points. The Cue chunk contains all of the cue points, but the Playlist chunk determines how those cue points are used when playing back the waveform (ie, which cue points represent looped sections, and in what order those loops are "played"). The Playlist chunk contains one or more Segment structures, each of which identifies a looped section of the waveform (in conjunction with the CuePoint structure with which it is associated).

Segment Structure

```
typedef struct {
    long    dwIdentifier;
    long    dwLength;
    long    dwRepeats;
} Segment;
```

The dwIdentifier field contains a unique number (ie, different than the ID number of any other Segment structure). This field should correspond with the dwIdentifier field of some CuePoint stored in the Cue chunk. In other words, this Segment structure contains the looping information associated with that CuePoint structure with the same ID number.

The dwLength field specifies the length of the section in samples (ie, the length of the looped section). Note that the start position of the loop would be the dwSampleOffset of the referenced CuePoint structure in the Cue chunk. (Or, you may need to hassle with the dwChunkStart and dwBlockStart fields as well if dealing with a Wave List or compressed data).

The dwRepeats field specifies the number of times to play the loop. I assume that a value of 1 means to repeat this loop once only, but the WAVE documentation is very incomplete and omits this important information. I have no idea how you would specify an infinitely repeating loop. Certainly, the person who conjured up the Playlist chunk appears to have no idea whatsoever. Due to the ambiguities, inconsistencies, inefficiencies, and omissions of the Cue and Playlist chunks, I very much recommend that you use the Sampler chunk (described later) to replace them.

Playlist chunk

```
#define PlaylistID 'plst' /* chunk ID for Playlist Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;

    long    dwSegments;
    Segment Segments[];
} PlaylistChunk;
```

The ID is always **plst**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields.

The dwSegments field is the number of Segment structures in the Playlist Chunk. If dwSegments is not 0, it is followed by that many Segment structures, one after the other. Because all fields in a Segment structure are an even number of bytes, the length of any Segment will always be even. Thus, Segments are packed together with no unused bytes between them. The Segments need not be placed in any particular order.

Associated Data List

The Associated Data List contains text "labels" or "names" that are associated with the CuePoint structures in the Cue chunk. In other words, this list contains the text labels for those CuePoints.

Again, we're talking about another imbedded IFF List within the WAVE file. NOOOOOOOOOOOOOO!!!! What's a List? A List is simply a "master chunk" that contains several "sub-chunks". Just like with any other chunk, the "master chunk" has an ID and chunkSize, but inside of this chunk are sub-chunks, each with its own ID and chunkSize. Of course, the chunkSize for the master chunk (ie, List) includes the size of all of these sub-chunks (including their ID and chunkSize fields).

The "Type ID" for the Associated Data List is "adtl". Remember that an IFF list header has 3 fields:

```
typedef struct {
    ID      listID;      /* 'list' */
    long    chunkSize;  /* includes the Type ID below */
    ID      typeID;     /* 'adtl' */
} ListHeader;
```

There are several sub-chunks that may be found inside of the Associated Data List. The ones that are important to WAVE format have IDs of "labl", "note", or "ltxt". Ignore the rest. Here are those 3 sub-chunks and their fields:

The Associated Data List is optional. The WAVE documentation doesn't specify if more than one can be contained in a WAVE file.

Label Chunk

```
#define LabelID 'labl' /* chunk ID for Label Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;

    long    dwIdentifier;
    char    dwText[];
} LabelChunk;
```

The ID is always **labl**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size (ie, chunkSize is the number of remaining bytes in the chunk after the chunkSize field, not counting any trailing pad byte).

The dwIdentifier field contains a unique number (ie, different than the ID number of any other Label chunk). This field should correspond with the dwIdentifier field of some CuePoint stored in the Cue chunk. In other words, this Label chunk contains the text label associated with that CuePoint structure with the same ID number.

The dwText array contains the text label. It should be a null-terminated string. (The null byte is included in the chunkSize, therefore the length of the string, including the null byte, is chunkSize - 4).

Note Chunk

```
#define NoteID 'note' /* chunk ID for Note Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;

    long    dwIdentifier;
    char    dwText[];
} NoteChunk;
```

The Note chunk, whose ID is **note**, is otherwise exactly the same as the Label chunk (ie, same fields). See what I mean about pointless duplication? But, in theory, a Note chunk contains a "comment" about a CuePoint, whereas the Label chunk is supposed to contain the actual CuePoint label. So, it's possible that you'll find both a Note and Label for a specific CuePoint, each containing different text.

Labeled Text Chunk

```
#define LabelTextID 'ltxt' /* chunk ID for Labeled Text Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;

    long    dwIdentifier;
    long    dwSampleLength;
    long    dwPurpose;
    short   wCountry;
    short   wLanguage;
```

```

short    wDialect;
short    wCodePage;
char     dwText[];
} LabelTextChunk;

```

The ID is always **Itxt**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size (ie, chunkSize is the number of remaining bytes in the chunk after the chunkSize field, not counting any trailing pad byte).

The dwIdentifier field is the same as the Label chunk.

The dwSampleLength field specifies the number of sample points in the segment of waveform data. In other words, a Labeled Text chunk contains a label for a **section** of the waveform data, not just a specific point, for example the looped section of a waveform.

The dwPurpose field specifies the type or purpose of the text. For example, dwPurpose can contain an ID like "scrp" for script text or "capt" for close-caption text. How is this related to waveform data? Well, it isn't really. It's just that Associated Data Lists are used in other file formats, so they contain generic fields that sometimes don't have much relevance to waveform data.

The wCountry, wLanguage, and wCodePage fields specify the country code, language/dialect, and code page for the text. An application typically queries these values from the operating system.

Sampler Chunk

The Sampler (smpl) Chunk defines basic parameters that an instrument, such as a MIDI sampler, could use to play the waveform data. Most importantly, it includes information about looping the waveform (ie, during playback, to "sustain" the waveform). Of course, as you've come to expect from the WAVE file format, it duplicates some of the information that can be found in the Cue and Playlist chunks, but fortunately, in a more sensible, consistent, better-documented way.

```

#define SamplerID 'smpl' /* chunk ID for Sampler Chunk */

typedef struct {
    ID            chunkID;
    long         chunkSize;

    long         dwManufacturer;
    long         dwProduct;
    long         dwSamplePeriod;
    long         dwMIDIUnityNote;
    long         dwMIDIPitchFraction;
    long         dwSMPTEFormat;
    long         dwSMPTEOffset;
    long         cSampleLoops;
    long         cbSamplerData;
    struct SampleLoop Loops[];
} SamplerChunk;

```

The ID is always **smpl**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size (ie, chunkSize is the number of remaining bytes in the chunk after the chunkSize field, not counting any trailing pad byte).

The dwManufacturer field contains the MMA Manufacturer code for the intended sampler. Each manufacturer of MIDI products has his own ID assigned to him by the MIDI Manufacturer's Association. See the MIDI Specification (under [System Exclusive](#)) for a listing of current Manufacturer IDs. The high byte of dwManufacturer indicates the number of low order bytes (1 or 3) that are valid for the manufacturer code. For example, this value will be 0x01000013 for Digidesign (the MMA Manufacturer code is one byte, 0x13); whereas 0x03000041 identifies Microsoft (the MMA Manufacturer code is three bytes, 0x00 0x00 0x41). If the WAVE is not intended for a specific manufacturer, then this field should be set to 0.

The dwProduct field contains the Product code (ie, model ID) of the intended sampler for the dwManufacturer. Contact the manufacturer of the sampler to ascertain the sampler's model ID. If the WAVE is not intended for a specific manufacturer's product, then this field should be set to 0.

The dwSamplePeriod field specifies the period of one sample in nanoseconds (normally 1/nSamplesPerSec from the Format chunk. But note that this field allows finer tuning than nSamplesPerSec). For example, 44.1 KHz would be specified as 22675 (0x00005893).

The dwMIDIUnityNote field is the MIDI note number at which the instrument plays back the waveform data without pitch modification (ie, at the same sample rate that was used when the waveform was created). This value ranges 0 through 127, inclusive. Middle C is 60.

The dwMIDIPitchFraction field specifies the fraction of a semitone up from the specified dwMIDIUnityNote. A value of 0x80000000 is 1/2 semitone (50 cents); a value of 0x00000000 represents no fine tuning between semitones.

The dwSMPTEFormat field specifies the SMPTE time format used in the dwSMPTEOffset field. Possible values are:


```

0 = no SMPTE offset (dwSMPTEOffset should also be 0)
24 = 24 frames per second
25 = 25 frames per second
29 = 30 frames per second with frame dropping ('30 drop')
30 = 30 frames per second

```

The dwSMPTEOffset field specifies a time offset for the sample if it is to be synchronized or calibrated according to a start time other than 0. The format of this value is 0xhhmmssff. hh is a signed Hours value [-23..23]. mm is an unsigned Minutes value [0..59]. ss is unsigned Seconds value [0..59]. ff is an unsigned value [0..(- 1)].

The cSampleLoops field is the number (count) of SampleLoop structures that are appended to this chunk. These structures immediately follow the cbSamplerData field. This field will be 0 if there are no SampleLoop structures.

The cbSamplerData field specifies the size (in bytes) of any optional fields that an application wishes to append to this chunk. An application which needed to save additional information (ie, beyond the above fields) may append additional fields to the end of this chunk, after all of the SampleLoop structures. These additional fields are also reflected in the ChunkSize, and remember that the chunk should be padded out to an even number of bytes. The cbSamplerData field will be 0 if no additional information is appended to the chunk.

What follows the above fields are any SampleLoop structures. Each SampleLoop structure defines one loop (ie, the start and end points of the loop, and how many times it plays). What follows any SampleLoop structures are any additional, proprietary sampler information that an application chooses to store.

SampleLoop Structure

```

typedef struct {
    long dwIdentifier;
    long dwType;
    long dwStart;
    long dwEnd;
    long dwFraction;
    long dwPlayCount;
} SampleLoop;

```

The dwIdentifier field contains a unique number (ie, different than the ID number of any other SampleLoop structure). This field may correspond with the dwIdentifier field of some CuePoint stored in the Cue chunk. In other words, the CuePoint structure which has the same ID number would be considered to be describing the same loop as this SampleLoop structure. Furthermore, this field corresponds to the dwIdentifier field of any label stored in the Associated Data List. In other words, the text string (within some chunk in the Associated Data List) which has the same ID number would be considered to be this loop's "name" or "label".

The dwType field is the loop type (ie, how the loop plays back) as so:

```

0 - Loop forward (normal)
1 - Alternating loop (forward/backward)
2 - Loop backward
3-31 - reserved for future standard types
32-? - sampler specific types (manufacturer defined)

```

The dwStart field specifies the startpoint of the loop. In other words, it's the byte offset from the start of waveformData[], where an offset of 0 would be at the start of the waveformData[] array (ie, the loop start is at the very first sample point).

The dwEnd field specifies the endpoint of the loop (ie, a byte offset).

The dwFraction field allows fine-tuning for loop fractional areas between samples. Values range from 0x00000000 to 0xFFFFFFFF. A value of 0x80000000 represents 1/2 of a sample length.

The dwPlayCount field is the number of times to play the loop. A value of 0 specifies an infinite sustain loop (ie, the wave keeps looping until some external force interrupts playback, such as the musician releasing the key that triggered that wave's playback).

The Sampler Chunk is optional. I don't know as if there is any limit of one per WAVE file. I don't see why there should be such a limit, since after all, an application may need to deal with several MIDI samplers.

The Instrument Chunk Format

The Instrument Chunk contains some of the same type of information as the Sampler chunk. So what else is new?

```

#define InstrumentID 'inst' /* chunkID for Instruments Chunk */

typedef struct {
    ID chunkID;
    long chunkSize;

    unsigned char UnshiftedNote;

```

```

char      FineTune;
char      Gain;
unsigned char LowNote;
unsigned char HighNote;
unsigned char LowVelocity;
unsigned char HighVelocity;
} InstrumentChunk;

```

The ID is always **inst**. chunkSize should always be 7 since there are no fields of variable length.

The UnshiftedNote field is the same as the Sampler chunk's dwMIDIUnityNote field.

The FineTune field determines how much the instrument should alter the pitch of the sound when it is played back. Units are in cents (1/100 of a semitone) and range from -50 to +50. Negative numbers mean that the pitch of the sound should be lowered, while positive numbers mean that it should be raised. While not the same measurement is used, this field serves the same purpose as the Sampler chunk's dwFraction field.

The Gain field is the amount by which to change the gain of the sound when it is played. Units are decibels. For example, 0db means no change, 6db means double the value of each sample point (ie, every additional 6db doubles the gain), while -6db means halve the value of each sample point.

The LowNote and HighNote fields specify the suggested MIDI note range on a keyboard for playback of the waveform data. The waveform data should be played if the instrument is requested to play a note between the low and high note numbers, inclusive. The UnshiftedNote does not have to be within this range.

The LowVelocity and HighVelocity fields specify the suggested range of MIDI velocities for playback of the waveform data. The waveform data should be played if the note-on velocity is between low and high velocity, inclusive. The range is 1 (lowest velocity) through 127 (highest velocity), inclusive.

The Instrument Chunk is optional. No more than 1 Instrument Chunk can appear in one WAVE.

Audio Interchange File Format (AIFF)

Audio Interchange File Format (or AIFF) is a file format for storing digital audio (waveform) data. It supports a variety of bit resolutions, sample rates, and channels of audio. This format is very popular upon Apple platforms, and is widely used in professional programs that process digital audio waveforms.

This format uses the Electronic Arts Interchange File Format method for storing data in "chunks". You should read the article [About Interchange File Format](#) before proceeding.

Data Types

A C-like language will be used to describe the data structures in the file. A few extra data types that are not part of standard C, but which will be used in this document, are:

- extended** 80 bit IEEE Standard 754 floating point number (Standard Apple Numeric Environment [SANE] data type Extended). This would be a 10 byte field.
- pstring** Pascal-style string, a one-byte count followed by that many text bytes. The total number of bytes in this data type should be even. A pad byte can be added to the end of the text to accomplish this. This pad byte is not reflected in the count.
- ID** A chunk ID (ie, 4 ASCII bytes) as described in [About Interchange File Format](#).

Also note that when you see an array with no size specification (e.g., char ckData[];), this indicates a variable-sized array in our C-like language. This differs from standard C arrays.

Constants

Decimal values are referred to as a string of digits, for example 123, 0, 100 are all decimal numbers. Hexadecimal values are preceded by a 0x - e.g., 0x0A, 0x1, 0x64.


Data Organization

All data is stored in Motorola 68000 (ie, big endian) format. The bytes of multiple-byte values are stored with the high-order (ie, most significant) bytes first. Data bits are as follows (ie, shown with bit numbers on top):

```

      7 6 5 4 3 2 1 0
      +-----+
char: | msb                lsb |
      +-----+
15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0

```




File Structure

An Audio IFF file is a collection of a number of different types of chunks. There is a required Common Chunk which contains important parameters describing the waveform, such as its length and sample rate. The Sound Data chunk, which contains the actual waveform data, is also required if the waveform data has a length greater than 0 (ie, there actually is waveform data in the FORM). All other chunks are optional. Among the other optional chunks are ones which define markers, list instrument parameters, store application-specific information, etc. All of these chunks are described in detail in the following sections of this document.

All applications that use FORM AIFF must be able to read the 2 required chunks and can choose to selectively ignore the optional chunks. A program that copies a FORM AIFF should copy all of the chunks in the FORM AIFF, even those it chooses not to interpret.

There are no restrictions upon the order of the chunks within a FORM AIFF.

Here is a graphical overview of an example, minimal AIFF file. It consists of a single FORM AIFF containing the 2 required chunks, a Common Chunk and a Sound Data Chunk.




Sample Points and Sample Frames

A large part of interpreting Audio IFF files revolves around the two concepts of sample points and sample frames.

A sample point is a value representing a sample of a sound at a given moment in time. Each sample point is stored as a linear, 2's-complement value which may be from 1 to 32 bits wide (as determined by the sampleSize field in the Common Chunk). For example, each sample point of an 8-bit waveform would be an 8-bit byte (ie, a signed char).

Because most CPU's read and write operations deal with 8-bit bytes, it was decided that a sample point should be rounded up to a size which is a multiple of 8 when stored in an AIFF. This makes the AIFF easier to read into memory. If your ADC produces a sample point from 1 to 8 bits wide, a sample point should be stored in an AIFF as an 8-bit byte (ie, signed char). If your ADC produces a sample point from 9 to 16 bits wide, a sample point should be stored in an AIFF as a 16-bit word (ie, signed short). If your ADC produces a sample point from 17 to 24 bits wide, a sample point should be stored in an AIFF as three bytes. If your ADC produces a sample point from 25 to 32 bits wide, a sample point should be stored in an AIFF as a 32-bit doubleword (ie, signed long). Etc.


Furthermore, the data bits should be left-justified, with any remaining (ie, pad) bits zeroed. For example, consider the case of a 12-bit sample point. It has 12 bits, so the sample point must be saved as a 16-bit word. Those 12 bits should be left-justified so that they become bits 4 to 15 inclusive, and bits 0 to 3 should be set to zero. Shown below is how a 12-bit sample point with a value of binary 101000010111 is stored left-justified as two bytes (ie, a 16-bit word).




For multichannel sounds (for example, a stereo waveform), single sample points from each channel are interleaved. For example, assume a stereo (ie, 2 channel) waveform. Instead of storing all of the sample points for the left channel first, and then storing all of the sample points for the right channel next, you "mix" the two channels' sample points together. You would store the first sample point of the left channel. Next, you would store the first sample point of the right channel. Next, you would store the second sample point of the left channel. Next, you would store the second sample point of the right channel, and so on, alternating between storing the next sample point of each channel. This is what is meant by interleaved data; you store the

next sample point of each of the channels in turn, so that the sample points that are meant to be "played" (ie, sent to a DAC) simultaneously are stored contiguously.

The sample points that are meant to be "played" (ie, sent to a DAC) simultaneously are collectively called a **sample frame**. In the example of our stereo waveform, every two sample points makes up another sample frame. This is illustrated below for that stereo example.



For a monophonic waveform, a sample frame is merely a single sample point (ie, there's nothing to interleave). For multichannel waveforms, you should follow the conventions shown below for which order to store channels within the sample frame. (ie, Below, a single sample frame is displayed for each example of a multichannel waveform).



The sample points within a sample frame are packed together; there are no unused bytes between them. Likewise, the sample frames are packed together with no pad bytes.

The Common Chunk

The Common Chunk describes fundamental parameters of the waveform data such as sample rate, bit resolution, and how many channels of digital audio are stored in the FORM AIFF.

```
#define CommonID 'COMM' /* chunkID for Common Chunk */

typedef struct {
    ID          chunkID;
    long        chunkSize;

    short       numChannels;
    unsigned long numSampleFrames;
    short       sampleSize;
    extended    sampleRate;
} CommonChunk;
```

The ID is always **COMM**. The chunkSize field is the number of bytes in the chunk. This does not include the 8 bytes used by ID and Size fields. For the Common Chunk, chunkSize should always 18 since there are no fields of variable length (but to maintain compatibility with possible future extensions, if the chunkSize is > 18, you should always treat those extra bytes as pad bytes).

The numChannels field contains the number of audio channels for the sound. A value of 1 means monophonic sound, 2 means stereo, 4 means four channel sound, etc. Any number of audio channels may be represented. For multichannel sounds, single sample points from each channel are interleaved. A set of interleaved sample points is called a sample frame.

The actual waveform data is stored in another chunk, the Sound Data Chunk, which will be described later.

The numSampleFrames field contains the number of sample frames. This is not necessarily the same as the number of bytes nor the number of sample points in the Sound Data Chunk (ie, it won't be unless you're dealing with a mono waveform). The total number of sample points in the file is numSampleFrames times numChannels.

The sampleSize is the number of bits in each sample point. It can be any number from 1 to 32.

The sampleRate field is the sample rate at which the sound is to be played back in sample frames per second.

One, and only one, Common Chunk is required in every FORM AIFF.

Sound Data Chunk

The Sound Data Chunk contains the actual sample frames (ie, all channels of waveform data).

```
#define SoundDataID 'SSND' /* chunk ID for Sound Data Chunk */
typedef struct {
    ID          chunkID;
    long        chunkSize;

    unsigned long  offset;
    unsigned long  blockSize;
    unsigned char  WaveformData[];
} SoundDataChunk;
```

The ID is always **SSND**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size (ie, chunkSize is the number of remaining bytes in the chunk after the chunkSize field, not counting any trailing pad byte).

You can determine how many bytes of actual waveform data there is by subtracting 8 from the chunkSize. Remember that the number of sample frames, bit resolution, and other information is gotten from the Common Chunk.

The offset field determines where the first sample frame in the WaveformData starts. The offset is in bytes. Most applications won't use offset and should set it to zero. Use for a non-zero offset is explained in "Block-Aligning Waveform Data".


The blockSize is used in conjunction with offset for block-aligning waveform data. It contains the size in bytes of the blocks that waveform data is aligned to. As with offset, most applications won't use blockSize and should set it to zero. More information on blockSize is in "Block-Aligning Waveform Data".

The WaveformData array contains the actual waveform data. The data is arranged into what are called *sample frames*. The number of sample frames in WaveformData is determined by the numSampleFrames field in the Common Chunk. For more information, see "Sample Points and Sample Frames".

The Sound Data Chunk is required unless the numSampleFrames field in the Common Chunk is zero. One, and only one, Sound Data Chunk may appear in a FORM AIFF.

Block-Aligning Waveform Data

There may be some applications that, to ensure real time recording and playback of audio, wish to align waveform data into fixed-size blocks. This alignment can be accomplished with the offset and blockSize parameters of the Sound Data Chunk, as shown below.



Above, the first sample frame starts at the beginning of block N. This is accomplished by skipping the first offset bytes (ie, some stored pad bytes) of the WaveformData. Note that there may also be pad bytes stored at the end of WaveformData to pad it out so that it ends upon a block boundary.

The `blockSize` specifies the size in bytes of the block to which you would align the waveform data. A `blockSize` of 0 indicates that the waveform data does not need to be block-aligned. Applications that don't care about block alignment should set the `blockSize` and `offset` to 0 when creating AIFF files. Applications that write block-aligned waveform data should set `blockSize` to the appropriate block size. Applications that modify an existing AIFF file should try to preserve alignment of the waveform data, although this is not required. If an application does not preserve alignment, it should set the `blockSize` and `offset` to 0. If an application needs to realign waveform data to a different sized block, it should update `blockSize` and `offset` accordingly.

The Marker Chunk

The Marker Chunk contains markers that point to positions in the waveform data. Markers can be used for whatever purposes an application desires. The Instrument Chunk, defined later in this document, uses markers to mark loop beginning and end points.

A marker structure is as follows:


```
typedef short  MarkerId;

typedef struct {
    MarkerID    id;
    unsigned long position;
    pstring     markerName;
} Marker;
```

The `id` is a number that uniquely identifies that marker within an AIFF. The `id` can be any positive non-zero integer, as long as no other marker within the same FORM AIFF has the same `id`.

The marker's position in the `WaveformData` is determined by the `position` field. Markers conceptually fall between two sample frames. A marker that falls before the first sample frame in the waveform data is at position 0, while a marker that falls between the first and second sample frame in the waveform data is at position 1. Therefore, the units for position are sample frames, not bytes nor sample points.

Sample Frames



The `markerName` field is a Pascal-style text string containing the name of the mark.

Note: Some "EA IFF 85" files store strings as C-strings (text bytes followed by a null terminating character) instead of Pascal-style strings. Audio IFF uses pstrings because they are more efficiently skipped over when scanning through chunks. Using pstrings, a program can skip over a string by adding the string count to the address of the first character. C strings require that each character in the string be examined for the null terminator.

Marker Chunk Format

The format for the data within a Marker Chunk is shown below.

```
#define MarkerID 'MARK' /* chunkID for Marker Chunk */

typedef struct {
    ID            chunkID;
    long          chunkSize;

    unsigned short numMarkers;
    Marker        Markers[];
} MarkerChunk;
```

The `ID` is always **MARK**. `chunkSize` is the number of bytes in the chunk, not counting the 8 bytes used by `ID` and `Size` fields.

The `numMarkers` field is the number of marker structures in the Marker Chunk. If `numMarkers` is not 0, it is followed by that many marker structures, one after the other. Because all fields in a marker structure are an even number of bytes, the length of any marker will always be even. Thus, markers are packed together with no unused bytes between them. The markers need not be placed in any particular order.

The Marker Chunk is optional. No more than one Marker Chunk can appear in a FORM AIFF.

The Instrument Chunk


The Instrument Chunk defines basic parameters that an instrument, such as a MIDI sampler, could use to play the waveform data.

Looping

Waveform data can be looped, allowing a portion of the waveform to be repeated in order to lengthen the sound. The structure below describes a loop.

```
typedef struct {
    short    PlayMode;
    MarkerId beginLoop;
    MarkerId endLoop;
} Loop;
```

A loop is marked with two points, a begin position and an end position. There are two ways to play a loop, forward looping and forward/backward looping. In the case of forward looping, playback begins at the beginning of the waveform, continues past the begin position and continues to the end position, at which point playback starts again at the begin position. The segment between the begin and end positions, called the loop segment, is played repeatedly until interrupted by some action, such as a musician releasing a key on a musical controller.



With forward/backward looping, the loop segment is first played from the begin position to the end position, and then played backwards from the end position to the begin position. This flip-flop pattern is repeated over and over again until interrupted.

The playMode specifies which type of looping is to be performed:

```
#define NoLooping          0
#define ForwardLooping    1
#define ForwardBackwardLooping 2
```

If NoLooping is specified, then the loop points are ignored during playback.

The beginLoop is a marker id that marks the begin position of the loop segment.

The endLoop marks the end position of a loop. The begin position must be less than the end position. If this is not the case, then the loop segment has 0 or negative length and no looping takes place.

The Instrument Chunk Format

The format of the data within an Instrument Chunk is described below.

```
#define InstrumentID 'INST' /*chunkID for Instruments Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;

    char    baseNote;
    char    detune;
    char    lowNote;
    char    highNote;
    char    lowvelocity;
    char    highvelocity;
    short   gain;
    Loop    sustainLoop;
    Loop    releaseLoop;
} InstrumentChunk;
```

The ID is always **INST**. chunkSize should always be 20 since there are no fields of variable length.

The baseNote is the note number at which the instrument plays back the waveform data without pitch modification (ie, at the same sample rate that was used when the waveform was created). Units are MIDI note numbers, and are in the range 0 through 127. Middle C is 60.

The detune field determines how much the instrument should alter the pitch of the sound when it is played back. Units are in cents (1/100 of a semitone) and range from -50 to +50. Negative numbers mean that the pitch of the sound should be lowered, while positive numbers mean that it should be raised.

The lowNote and highNote fields specify the suggested note range on a keyboard for playback of the waveform data. The waveform data should be played if the instrument is requested to play a note between the low and high note numbers, inclusive. The base note does not have to be within this range. Units for lowNote and highNote are MIDI note values.

The lowVelocity and highVelocity fields specify the suggested range of velocities for playback of the waveform data. The waveform data should be played if the note-on velocity is between low and high velocity, inclusive. Units are MIDI velocity values, 1 (lowest velocity) through 127 (highest

velocity).

The gain is the amount by which to change the gain of the sound when it is played. Units are decibels. For example, 0db means no change, 6db means double the value of each sample point (ie, every additional 6db doubles the gain), while -6db means halve the value of each sample point.

The sustainLoop field specifies a loop that is to be played when an instrument is sustaining a sound.

The releaseLoop field specifies a loop that is to be played when an instrument is in the release phase of playing back a sound. The release phase usually occurs after a key on an instrument is released.

The Instrument Chunk is optional. No more than 1 Instrument Chunk can appear in one FORM AIFF.

The MIDI Data Chunk

The MIDI Data Chunk can be used to store MIDI data.

The primary purpose of this chunk is to store MIDI System Exclusive messages, although other types of MIDI data can be stored in the chunk as well. As more instruments come to market, they will likely have parameters that have not been included in the AIFF specification. The Sys Ex messages for these instruments may contain many parameters that are not included in the Instrument Chunk. For example, a new MIDI sampler may support more than the two loops per waveform. These loops will likely be represented in the Sys Ex message for the new sampler. This message can be stored in the MIDI Data Chunk (ie, so you have some place to store these extra loop points that may not be used by other instruments).

```
#define MIDIDataID 'MIDI' /* chunkID for MIDI Data Chunk */

typedef struct {
    ID          chunkID;
    long        chunkSize;

    unsigned char MIDIData[];
} MIDIDataChunk;
```

The ID is always **MIDI**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size.

The MIDIData field contains a stream of MIDI data. There should be as many bytes as chunkSize specifies, plus perhaps a pad byte if needed.

The MIDI Data Chunk is optional. Any number of these chunks may exist in one FORM AIFF. If MIDI System Exclusive messages for several instruments are to be stored in a FORM AIFF, it is better to use one MIDI Data Chunk per instrument than one big MIDI Data Chunk for all of the instruments.

The Audio Recording Chunk

The Audio Recording Chunk contains information pertinent to audio recording devices.

```
#define AudioRecording ID 'AESD' /* chunkID for Audio Recording Chunk. */

typedef struct {
    ID          chunkID
    long        chunkSize;

    unsigned char AESChannelStatusData[24];
} AudioRecordingChunk;
```

The ID is always **AESD**. chunkSize should always be 24 since there are no fields of variable length.

The 24 bytes of AESChannelStatusData are specified in the "AES Recommended Practice for Digital Audio Engineering - Serial Transmission Format for Linearly Represented Digital Audio Data", transmission of digital audio between audio devices. This information is duplicated in the Audio Recording Chunk for convenience. Of general interest would be bits 2, 3, and 4 of byte 0, which describe recording emphasis.

The Audio Recording Chunk is optional. No more than 1 Audio Recording Chunk may appear in one FORM AIFF.

The Application Specific Chunk

The Application Specific Chunk can be used for any purposes whatsoever by developers and application authors. For example, an application that edits sounds might want to use this chunk to store editor state parameters such as magnification levels, last cursor position, etc.

```
#define ApplicationSpecificID 'APPL' /* chunkID for Application Specific Chunk. */

typedef struct {
    ID        chunkID;
    long      chunkSize;

    char      applicationSignature[4];
    char      data[];
} ApplicationSpecificChunk;
```

The ID is always **APPL**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size.

The applicationSignature field is used by applications which run on platforms from Apple Computer, Inc. For the Apple II, this field should be set to 'pdos'. For the Mac, this field should be set to the application's four character signature as registered with Apple Technical Support.

The data field is the data specific to the application. The application determines how many bytes are stored here, and what their purpose are. A trailing pad byte must follow if that is needed in order to make the chunk an even size.

The Application Specific Chunk is optional. Any number of these chunks may exist in a one FORM AIFF.

The Comments Chunk

The Comments Chunk is used to store comments in the FORM AIFF. Standard IFF has an Annotation Chunk that can also be used for comments, but this new Comments Chunk has two fields (per comment) not found in the Standard IFF chunk. They are a time-stamp for the comment and a link to a marker.

Comment structure

A Comment structure consists of a time stamp, marker id, and a text count followed by text.

```
typedef struct {
    unsigned long    timeStamp;
    MarkerID        marker;
    unsigned short   count;
    char            text[];
} Comment;
```

The timeStamp indicates when the comment was created. On the Amiga, units are the number of seconds since January 1, 1978. On the Mac, units are the number of seconds since January 1, 1904.

A comment can be linked to a marker. This allows applications to store long descriptions of markers as a comment. If the comment is referring to a marker, then the marker field is the ID of that marker. Otherwise, marker is 0, indicating that this comment is not linked to any marker.

The count is the length of the text that makes up the comment. This is a 16-bit quantity, allowing much longer comments than would be available with a pstring. This count does not include any possible pad byte needed to make the comment an even number of bytes in length.

The text field contains the comment itself, followed by a pad byte if needed to make the text field an even number of bytes.

Comments Chunk Format

```
#define CommentID 'COMT' /* chunkID for Comments Chunk */

typedef struct {
    ID        chunkID;
    long      chunkSize;

    unsigned short   numComments;
    char            comments[];
} CommentsChunk;
```

The ID is always **COMT**. chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields.

The numComments field contains the number of Comment structures in the chunk. This is followed by the Comment structures, one after the other. Comment structures are always even numbers of bytes in length, so there is no padding needed between structures.

The Comments Chunk is optional. No more than 1 Comments Chunk may appear in one FORM AIFF.

The Text Chunks, Name, Author, Copyright, Annotation

These four optional chunks are included in the definition of every Standard IFF file.

```
#define NameID 'NAME' /* chunkID for Name Chunk */
#define NameID 'AUTH' /* chunkID for Author Chunk */
#define NameID '(c)' /* chunkID for Copyright Chunk */
#define NameID 'ANNO' /* chunkID for Annotation Chunk */

typedef struct {
    ID      chunkID;
    long    chunkSize;
    char    text[];
}TextChunk;
```

chunkSize is the number of bytes in the chunk, not counting the 8 bytes used by ID and Size fields nor any possible pad byte needed to make the chunk an even size.

Chunk Precedence

Several of the local chunks for FORM AIFF may contain duplicate information. For example, the Instrument Chunk defines loop points and some MIDI Sys Ex data in the MIDI Data Chunk may also define loop points. What happens if these loop points are different? How is an application supposed to loop the sound? Such conflicts are resolved by defining a precedence for chunks. This precedence is illustrated below.



The Common Chunk has the highest precedence, while the Application Specific Chunk has the lowest. Information in the Common Chunk always takes precedence over conflicting information in any other chunk. The Application Specific Chunk always loses in conflicts with other chunks. By looking at the chunk hierarchy, for example, one sees that the loop points in the Instrument Chunk take precedence over conflicting loop points found in the MIDI Data Chunk.

It is the responsibility of applications that write data into the lower precedence chunks to make sure that the higher precedence chunks are updated accordingly (ie, so that conflicts tend not to exist).

Errata

The Apple IIGS Sampled Instrument Format also defines a chunk with ID of "INST," which is not the same as the AIFF Instrument Chunk. A good way to tell the two chunks apart in generic IFF-style readers is by the chunkSize fields. The AIFF Instrument Chunk's chunkSize field is always 20, whereas the Apple IIGS Sampled Instrument Format Instrument Chunk's chunkSize field, for structural reasons, can never be 20.

Storage of AIFF on Apple and Other Platforms

On a Macintosh, the FORM AIFF, is stored in the data fork of an Audio IFF file. The Macintosh file type of an Audio IFF file is 'AIFF'. This is the same as the formType of the FORM AIFF. Macintosh applications should not store any information in Audio IFF file's resource fork, as this information may not be preserved by all applications. Applications can use the Application Specific Chunk, defined later in this document, to store extra information specific to their application.

Audio IFF files may be identified in other Apple file systems as well. On a Macintosh under MFS or HFS, the FORM AIFF is stored in the data fork of a file with file type "AIFF." This is the same as the formType of the FORM AIFF.

On an operating system such as MS-DOS or UNIX, where it is customary to use a file name extension, it is recommended that Audio IFF file names use ".AIF" for the extension.

Referring to Audio IFF

The official name is "Audio Interchange File Format". If an application needs to present the name of this format to a user, such as in a "Save As..." dialog box, the name can be abbreviated to Audio IFF. Referring to Audio IFF files by a four-letter abbreviation (ie, "AIFF") at the user-level should be avoided.

Converting Extended data to a unsigned long

The sample rate field in a Common Chunk is expressed as an 80 bit IEEE Standard 754 floating point number. This isn't a very useful format for computer software and audio hardware that can't directly deal with such floating point values. For this reason, you may wish to use the following ConvertFloat() hack to convert the floating point value to an unsigned long. This function doesn't handle very large floating point values, but certainly larger than you would ever expect a typical sampling rate to be. This function assumes that the passed *buffer* arg is a pointer to the 10 byte array which already contains the 80-bit floating point value. It returns the value (ie, sample rate in Hertz) as an unsigned long. Note that the FlipLong() function is only of use to Intel CPU's which have to deal with AIFF's Big Endian order. If not compiling for an Intel CPU, comment out the INTEL_CPU define.

```
#define INTEL_CPU

#ifdef INTEL_CPU
/* ***** FlipLong() *****
 * Converts a long in "Big Endian" format (ie, Motorola 68000) to Intel
 * reverse-byte format, or vice versa if originally in Big Endian.
 * ***** */

void FlipLong(unsigned char * ptr)
{
    register unsigned char val;

    /* Swap 1st and 4th bytes */
    val = *(ptr);
    *(ptr) = *(ptr+3);
    *(ptr+3) = val;

    /* Swap 2nd and 3rd bytes */
    ptr += 1;
    val = *(ptr);
    *(ptr) = *(ptr+1);
    *(ptr+1) = val;
}
#endif

/* ***** FetchLong() *****
 * Fools the compiler into fetching a long from a char array.
 * ***** */

unsigned long FetchLong(unsigned long * ptr)
{
    return(*ptr);
}

/* ***** ConvertFloat() *****
 * Converts an 80 bit IEEE Standard 754 floating point number to an unsigned
 * long.
 * ***** */

unsigned long ConvertFloat(unsigned char * buffer)
{
    unsigned long mantissa;
    unsigned long last = 0;
    unsigned char exp;

#ifdef INTEL_CPU
    FlipLong((unsigned long *) (buffer+2));
#endif
}

```

```

mantissa = FetchLong((unsigned long *) (buffer+2));
exp = 30 - *(buffer+1);
while (exp--)
{
    last = mantissa;
    mantissa >>= 1;
}
if (last & 0x00000001) mantissa++;
return(mantissa);
}

```

Of course, you may need a complementary routine to take a sample rate as an unsigned long, and put it into a buffer formatted as that 80-bit floating point value.

```

/* ***** StoreLong() *****
 * Fools the compiler into storing a long into a char array.
 * ***** */

void StoreLong(unsigned long val, unsigned long * ptr)
{
    *ptr = val;
}

/* ***** StoreFloat() *****
 * Converts an unsigned long to 80 bit IEEE Standard 754 floating point
 * number.
 * ***** */

void StoreFloat(unsigned char * buffer, unsigned long value)
{
    unsigned long exp;
    unsigned char i;

    memset(buffer, 0, 10);

    exp = value;
    exp >>= 1;
    for (i=0; i<32; i++) { exp>>= 1;
        if (!exp) break;
    }
    *(buffer+1) = i;

    for (i=32; i; i--)
    {
        if (value & 0x80000000) break;
        value <<= 1; } StoreLong(value, buffer+2); #ifdef INTEL_CPU FlipLong((unsigned long *) (buffer+2)); #endif }

```

Multi-sampling

Many MIDI samplers allow splitting up the MIDI note range into smaller ranges (for example by octaves) and assigning a different waveform to play over each range. If you wanted to store all of those waveforms into a single data file, what you would do is create an IFF LIST (or perhaps CAT if you wanted to store FORMs other than AIFF in it) and then include an embedded FORM AIFF for each one of the waveforms. (ie, Each waveform would be in a separate FORM AIFF, and all of these FORM AIFFs would be in a single LIST file). See [About Interchange File Format](#) for details about LISTS.

```

Content-Type: text/html; charset=iso-8859-1; name="New WAVE RIFF Chunks.htm"
Content-Disposition: inline; filename="New WAVE RIFF Chunks.htm"
Content-Base: "file:///D:/Source/Sound/Format/New%20WAVE%20RIFF%20Chunks.htm"

```

X-MIME-Autoconverted: from 8bit to quoted-printable by skynet.usb.ve id TAA05017



New WAVE RIFF Chunks

Added: 05/01/92
 Author: Microsoft, IBM

Most of the information in this section comes directly from the IBM/Microsoft RIFF standard document.

The WAVE form is defined as follows. Programs must expect (and ignore) any unknown chunks encountered, as with all RIFF forms. However, **<fmt-ck>** must always occur before , and both of these chunks are mandatory in a WAVE file.

```
→
RIFF( 'WAVE'
<'fmt'-ck> // Format
[] // Fact chunk
[] // Cue points
[] // Playlist
[] // Associated data list
) // Wave data
```

The WAVE chunks are described in the following sections.

Fact Chunk

The stores file dependent information about the contents of the WAVE file. This chunk is defined as follows:

```
→ fact( )
```

<dwSampleLength> represents the length of the data in samples. The **<nSamplesPerSec>** field from the wave format header is used in conjunction with the **<dwSampleLength>** field to determine the length of the data in seconds.

The fact chunk is required for all new WAVE formats. The chunk is not required for the standard WAVE_FORMAT_PCM files.

The fact chunk will be expanded to include any other information required by future WAVE formats. Added fields will appear following the field. Applications can use the chunk size field to determine which fields are present.

Cue Points Chunk

The **<cue-ck>** cue-points chunk identifies a series of positions in the waveform data stream. The is defined as follows:

```
→ cue( // Count of cue points
... ) // Cue-point table

→ struct {
DWORD dwName;
DWORD dwPosition;
FOURCC fccChunk;
DWORD dwChunkStart;
DWORD dwBlockStart;
DWORD dwSampleOffset;
}
```

The **<cue-point>** fields are as follows:

Field	Description
dwName	Specifies the cue point name. Each record must have a unique dwName field.
dwPosition	Specifies the sample position of the cue point. This is the sequential sample number within the play order. See "Playlist Chunk," later in this document, for a discussion of the play order.
fccChunk	Specifies the name or chunk ID of the chunk containing the cue point.
dwChunkStart	Specifies the position of the start of the data chunk containing the cue point. This should be zero if there is only one chunk containing data (as is currently always the case).
dwBlockStart	Specifies the position of the start of the block containing the position. This is the byte offset from the start of the data section of the chunk, not the chunk's FOURCC.
dwSampleOffset	Specifies the sample offset of the cue point relative to the start of the block.

Examples of File Position Values

The following table describes the field values for a WAVE file containing a single 'data' chunk:

Cue Point Location	Field	Value
Within PCM data	fccChunk	FOURCC value 'data'.
	dwChunkStart	Zero value.
	dwBlockStart	File position of the sample (nBlockAlign aligned bytes) relative to the start of the data section of the 'data' chunk (not the FOURCC).
	dwSampleOffset	Sample position of the cue point relative to the start of the 'data' chunk.
In all other 'data' chunks	fccChunk	FOURCC value 'data'.
	dwChunkStart	Zero value.
	dwBlockStart	File position of the enclosing block relative to the start of the data section of the 'data' chunk (not the FOURCC). The software can begin the decompression at this point.
	dwSampleOffset	Sample position of the cue point relative to the start of the block.

Playlist Chunk

The playlist chunk specifies a play order for a series of cue points. The is defined as follows:

```
@ plst(
// Count of play segments
... ) // Play-segment table
```

```
@ struct {
DWORD dwName;
DWORD dwLength;
DWORD dwLoops;
}
```

The <play-segment> fields are as follows:

Field	Description
dwName	Specifies the cue point name. This value must match one of the names listed in the cue-point table.
dwLength	Specifies the length of the section in samples.
dwLoops	Specifies the number of times to play the section.

Associated Data Chunk

The associated data list provides the ability to attach information like labels to sections of the waveform data stream. The is defined as follows:

```
→ LIST( 'adt1'
// Label
// Note
} // Text with data length
```

```

→ labl(
)

→ note(
)

→ ltxt(
... )

```

Label and Note Information

The 'labl' and 'note' chunks have similar fields. The 'labl' chunk contains a label, or title, to associate with a cue point. The 'note' chunk contains comment text for a cue point. The fields are as follows:

Field	Description
dwName	Specifies the cue point name. This value must match one of the names listed in the cue-point table.
data	Specifies a NULL-terminated string containing a text label (for the 'labl' chunk) or comment text (for the 'note' chunk).

Text with Data Length Information

The "ltxt" chunk contains text that is associated with a data segment of specific length. The chunk fields are as follows:

Field	Description
dwName	Specifies the cue point name. This value must match one of the names listed in the cue-point table.
dwSampleLength	Specifies the number of samples in the segment of waveform data.
dwPurpose	Specifies the type or purpose of the text. For example, can specify a FOURCC code like 'scrp' for script text or 'capt' for close-caption text.
wCountry	Specifies the country code for the text. See "Country Codes" for a current list of country codes.
wLanguage, wDialect	Specify the language and dialect codes for the text. See "Language and Dialect Codes" for a current list of language and dialect codes.
wCodePage	Specifies the code page for the text.

inst (Instrument) Chunk

Added: 12/29/92
Author: IBM
Defined for: WAVE form

The WAVE form is NEARLY the perfect file format for storing a sampled sound synthesizer's samples. Bits per sample, sample rate, number of channels, and complex looping can be specified with current WAVE subchunks, but a sample's pitch and its desired volume relative to other samples cannot. The optional instrument subchunk defined below fills in these needed parameters:

```
|| → inst(
)
```

bUnshiftedNote	the MIDI note number that corresponds to the unshifted pitch of the sample. Valid values range from 0 to 127.
chFineTune	the pitch shift adjustment in cents (or 100ths of a semitone) needed to hit bUnshiftedNote value exactly. chFineTune can be used to compensate for tuning errors in the sampling process. Valid values range from -50 to 50.
chGain	the suggested volume setting for the sample in decibels. A value of zero decibels suggests no change in the volume. A value of -6 decibels suggests reducing the amplitude of the sample by two.
bLowNote and bHigh Note	the suggested usable MIDI note number range of the sample. Valid values range from 0 to 127.
bLowVelocity and bHighVelocity	the suggested usable MIDI velocity range of the sample. Valid values range from 0 to 127.

smpl (Sample) Chunk

Added: 11/09/93
Author: Digidesign, Sonic Foundary, Turtle Beach
Defined for: WAVE form

The sampled instrument chunk describes the minimum necessary information needed to allow a sampling keyboard to use a WAVE file as an instrument. Samplers which require more information can save their extended information in the sampler specific data section. The is defined as follows:

```
|| → smpl(
)
```

```
struct
{
    DWORD dwIdentifier;
    DWORD dwType;
    DWORD dwStart;
    DWORD dwEnd;
    DWORD dwFraction;
    DWORD dwPlayCount;
}
```

The chunk:

dwManufacturer	Specifies the MMA Manufacturer code for the intended target device. The high byte indicates the number of low order bytes (1 or 3) that are valid for the manufacturer code. For example, this value will be 0x01000013 for Digidesign (the MMA Manufacturer code is one byte, 0x13); whereas 0x03000041 identifies Microsoft (the MMA Manufacturer code is three bytes, 0x00 0x00 0x41). If the sample is not intended for a specific manufacturer, then this field should be set to zero.
dwProduct	Specifies the Product code of the intended target device for the dwManufacturer. If the sample is not intended for a specific manufacturer's product, then this field should be set to zero.
dwSamplePeriod	Specifies the period of one sample in nanoseconds (normally 1/nSamplesPerSec from the WAVEFORMAT structure for the RIFF WAVE file--however, this field allows fine tuning). For example, 44.1 kHz would be specified as 22675 (0x00005893).
dwMIDIUnityNote	Specifies the MIDI note which will replay the sample at original pitch. This value ranges from 0 to 127 (a value of 60 represents Middle C as defined by the MMA).
dwMIDIPitchFraction	Specifies the fraction of a semitone up from the specified dwMIDIUnityNote. A value of 0x80000000 is 1/2 semitone (50 cents); a value of 0x00000000 represents no fine tuning between semitones.
dwSMPTEFormat	Specifies the SMPTE time format used in the dwSMPTEOffset field. Possible values are (unrecognized formats should be ignored): 0 - specifies no SMPTE offset (dwSMPTEOffset should also be zero) 24 - specifies 24 frames per second 25 - specifies 25 frames per second 29 - specifies 30 frames per second with frame dropping ('30 drop') 30 - specifies 30 frames per second
dwSMPTEOffset	Specifies a time offset for the sample if it is to be synchronized or calibrated according to a start time other than 0. The format of this value is 0xhhmmssff. hh is a <i>signed</i> Hours value [-23..23]. mm is an unsigned Minutes value [0..59]. ss is unsigned Seconds value [0..59]. ff is an unsigned value [0..(- 1)].
cSampleLoops	Specifies the number (count) of records that are contained in the chunk. The records are stored immediately following the cbSamplerData field.
cbSamplerData	Specifies the size in bytes of the optional . Sampler specific data is stored immediately following the records. The cbSamplerData field will be zero if no extended sampler specific information is stored in the chunk.

The structure:

dwIdentifier	Identifies the unique 'name' of the loop. This field may correspond with a name stored in the chunk. The name data is stored in the chunk.
dwType	Specifies the loop type: 0 - Loop forward (normal) 1 - Alternating loop (forward/backward) 2 - Loop backward 3-31 - reserved for future standard types 32-? - sampler specific types (manufacturer defined)
dwStart	Specifies the startpoint of the loop in samples.
dwEnd	Specifies the endpoint of the loop in samples (this sample will also be played).
dwFraction	Allows fine-tuning for loop fractional areas between samples. Values range from 0x00000000 to 0xFFFFFFFF. A value of 0x80000000 represents 1/2 of a sample length.
dwPlayCount	Specifies the number of times to play the loop. A value of 0 specifies an infinite sustain loop.