

Interchangeable Preservation Format (IPF) Documentation

*By Jean Louis-Guerin (DrCoolZic)
Version 1.2 September 2013*

IPF – Interchangeable Preservation Format Documentation

Table of Content

1. Presentation	3
1.1 IPF (Interchangeable Preservation Format).....	3
1.2 IPF and Openness.....	3
1.3 Software Preservation Society.....	3
1.4 How Floppy Disks were Produced.....	3
1.5 Floppy Disk Preservation with IPF files.....	4
2. CAPS Analyzer Description Language	5
2.1 Disk Images.....	5
2.2 Format /Track Descriptors.....	5
2.3 Block/Sector Descriptors.....	6
2.4 Data Descriptors.....	6
2.4.1 Data type field.....	6
2.4.2 Some Data type.....	6
2.4.3 Data type table.....	7
2.5 Atari Example.....	8
2.5.1 Atari Format/Track descriptor.....	8
2.5.2 Atari Block & Data descriptors.....	8
3. IPF/IPX file content	10
3.1 IPF File organization.....	10
3.2 IPX File Organization.....	11
3.3 Record Header.....	11
3.4 CAPS Record.....	11
3.5 INFO Record.....	12
3.6 IMGE Record.....	13
3.7 DATA Record.....	14
3.7.1 Block Descriptor.....	14
3.7.2 Data Area.....	15
3.8 CTEI Record.....	16
3.9 CTEX Record.....	16
4. Atari Example	17
4.1 Atari FD Format short presentation.....	17
4.1.1 Atari Track Format.....	17
4.1.2 Atari Sector Format.....	17
4.1.3 Atari Sector write splices.....	18
4.1.4 Atari Track write splices.....	18
4.2 Example of Theme Park Mystery FD.....	19
4.2.1 Analysis of Track 00.0 in the Stream file.....	19
4.2.2 Analysis of track 00.0 in the IPF file.....	21
4.2.3 Copylock Protection.....	23
5. CAPS/IPF Programing	27
5.1 IPF Decoder Library.....	27
5.1.1 IPF Library Release 4.2A.....	27
5.2 IPF Dump Program.....	27
6. Glossary	28
7. References	31
8. History	31

1. Presentation

1.1 IPF (*Interchangeable Preservation Format*)

IPF stands for [Interchangeable Preservation Format](#), and is the file format defined and used by Software Preservation Society (SPS) to preserve content, that is, floppy disk images. It's naming is quite abstract for very good reasons:

- It is platform agnostic.
- It is low-level, representing the information as it would have been read by a drive head.
- It is not just for disks: Any type of digital media can be contained by the format. Currently it is used for disk images and ROM files, but in the future it may possibly be used for CDRoms and information contained in hardware “dongle” protections.
- The underlying file format has many similarities to the IFF (Interchangeable File Format) standard on the Amiga, which was a very forward-thinking format. The same techniques are used today in other file formats such as PNG.

1.2 IPF and Openness

This subject is described in [IPF and Openness subject](#) in the Kryoflux forum. After long debate the IPF source code has been made available under MAME license (see [IPF DECODER LIBRARY source code released](#)). Therefore the IPF format is now open but unfortunately completely undocumented. This is not felt as a problem by most user as the content of IPF files is supposed to be read through the IPF decoder library and **written only** by SPS people. In this document I try my best to provide information about the IPF file format so it should be possible to read and **write** IPF files directly.

 However beware that the information presented in this documents comes from different sources and inevitably must contain errors and therefore must be used with caution.

1.3 Software Preservation Society

The [Software Preservation Society \(SPS\)](#) dedicates itself to the preservation of software for the future, namely classic games. As it is, these items are no longer available from their original suppliers, and are mainly in the possession of an ever diminishing community of well willed collectors. However, just by the passage of time these games are affected by the gradual deterioration of the media that stores them. These classics risk being lost forever in the near future, a tragedy that must be prevented.

1.4 How Floppy Disks were Produced

First it is important to understand how diskettes *were* produced. Although it would have been possible to create non-protected diskettes directly on the target machine, in practice, especially to manufacture large quantity of disks, it was necessary to use dedicated *commercial floppy disk mastering machines*. Beside the capability to produce large quantity of disks these machines were also capable to write information that *could not be reproduced* on the target machine. This technique is often referred as [floppy disk copy protections](#). The copy protection method has at least two obvious qualities: first, a *key protected disk* can be simultaneously used as *protection* and *distribution* disk and second, this type of protection is very cheap but nevertheless hard to tamper with. The mastering machines were using dedicated description languages to describe pretty much any disk format, including the ones with copy protection. The duplication machines knew about complex magnetic recording theory in order to be able write perfect disks based on the description. For example the mastering machines from [Trace Products](#) were using the *Freeform* description language for that matter. The same description language could also be used to check that what has been written to a disk was correct.

1.5 Floppy Disk Preservation with IPF files

The term preservation means different things to different people. Most of the time, only the minimum required information about a floppy disk is saved so that it is possible to **emulate** correctly, for a specific platform, what would be read from the original disk. While this may be sufficient to run on an emulator a “preserved game” it should **not** be considered as preservation because a lot of information from the original master is lost. This is why SPS people have taken an innovative approach that consists in retrieving the original master description of the disk to preserve using a language equivalent to Freeform. Basic information about this language is described in the [CAPS Analyzer Description Language](#) chapter. Finding the original “Freeform” description can be compared to breaking a vault combination lock: it is a very hard process but once done you know that you have a perfect result. The description is then stored in IPF files that contain a combination of both the disk data and how this data is stored. IPF files can then be used either to master the disk images back to floppy disk or in emulation tool.


SPS have developed image creation tools (See [Kryoflux web site](#)) that can read disk information at very low level. The disk image containing the raw signal information is loaded into an Analyzer (CAPS Analyzer) where known disk formats are semi-automatically detected. The disk image is saved using the IPF format which contains the disk data, how it is stored (and hence how to write it back), and also a few other things needed to store copy protection (like how dense the data is at any part of each track).


It is interesting to note that the IPF files does not store directly detailed information often found in other preservation formats. For example the timing information measured by the preservation tool is **not** stored but the “format description” can for example indicate that specific sectors of a track use a well-known bit-cell density variation. The result is therefore perfect as the timings are not subject to measurement errors.

☞ *Note that the CAPS Analyzer does not support archival of modified disks that are based on generic MFM coding and were written on a standard machine (PC, Atari ST, etc.). The reason is that, except the Commodore Amiga that will always write a complete track at once, other machines equipped with standard controllers (e.g. Western Digital WD177x) replace single sectors only. Starting and stopping the magnetic fields of the write head will generate undefined data in GAPS as well as slight drift of the complete sector in relation to the rest of the track (also known as write splice).*

2. CAPS Analyzer Description Language

This chapter is a short presentation of the mastering description language used by the **CAPS Analyzer** (also referred as the **Analyzer** in this document) as described in the [WIP of July 18, 2002: Time for an even more generic way to describe disk formats](#) as well as some other WIP (see [references](#)). This document does not go into details (mainly because I do not know them) but nevertheless presents useful information required to better understand the IPF Format. The “CAD” Language is the SPS equivalent of the “Freeform” Language used by [Trace Machine](#).

 The information presented in this chapter is mainly coming from several WIP documents and from discussion with SPS’s developers. Therefore beware that some descriptions are speculative and may be wrong.

 Note that I am primarily knowledgeable about the Floppy Disks format used on Atari ST. Therefore all my examples are given using this format but it should be possible to transpose the information presented here to your format of choice.


We do the description from the top level of a simple disk image and we go downwards to the most basic level used by the **Analyzer** (the [Data Descriptor](#) type):

1. Disk Images
2. Format /Track Descriptors
3. Block Descriptors
4. Data Descriptors
5. Data Field Types
6. Other

We first give a generic explanation of the different descriptors followed by a detail example of an Atari format.

2.1 Disk Images

The disk images going into the **Analyzer** are floppy disks that are dumped by contributors with the [Kryoflux dumping tools](#). They normally contain 84 cylinders (0-83) if the disk drive used can safely read up that far, but the image format itself may contain any number. It is unlikely many games will use cylinders that high, but it is best to be sure.



2.2 Format /Track Descriptors

In order to recognize a track, we need to describe it. When the analysis occurs the Analyzer finds the format that best matches a track or indicates that the track is unknown. Any integrity information held by the track format is put into the track description and checked against to indicate any errors on the disk.

Most of the formats are known to be re-used several times. When we think the format was not re-used we name it after the game, otherwise we name it after the author (if known) or company. If something makes it clear it was derived from another format, it will be normally named as a variant, but it may be still be marked as such after author/company info if nothing better comes along.

Each track can contain an entirely different format to the next as [so we define disk formats on a track-by-track basis](#).

A track can contain one or more blocks of data; the exact number depends on the format. For example a standard Atari format track contains 9 blocks. Each of these blocks is described by a block descriptor.

IPF – Interchangeable Preservation Format Documentation

2.3 Block/Sector Descriptors

A track can contain one or more blocks of data (sometimes referred as sectors); the exact number depends on the format. As the name implies, a Block Descriptor is the representation of the block data format in the **Analyzer**.

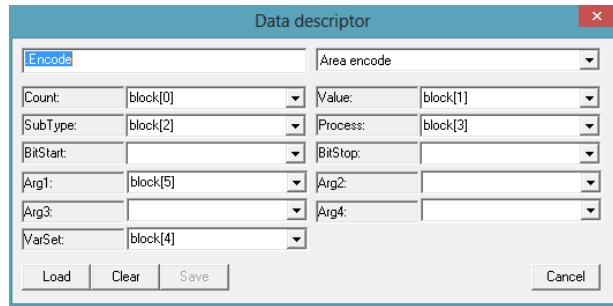
A block descriptor is described by a series of Data Descriptors and their parameters.

2.4 Data Descriptors

Format descriptors are defined by what block descriptors are used in them and whatever parameters are needed. In turn, blocks are described in terms of Data descriptors and their parameters.

Data descriptors are user-defined descriptions for a chunk of data. They indicate the “data field type” along with parameters to nail down exactly what the chunk of data is, and what it is used for. A data chunk or item is the smallest element of a format.

Block descriptor uses data descriptors as user definable functions. If you were programming, think of it as defining your functions first (in our case this would be the data descriptors) and then using them in context (in our case in the block descriptors). So basically, all the names are arbitrary. The person who writes the descriptor chooses the actual name of the format.



Data descriptor contains generic parameter fields:

- **Data Type:** The data field type to use, this is explained below.
- **Count:** Area number or data count.
- **Value:** Expected value.
- **SubType:** Depending on its context, this is information about the item, like encoder type or the basic type of data (byte, word, long word, etc.).
- **Process:** Processing requirement, like encryption.
- **Bit Start, Bit Stop:** These are values are always represented as 32 bits internally and it is possible to reference only a bit field of them so the complete value is marked as 0, 31. But any valid field selected can now be processed.
- **Arg1...4:** Generic placeholders for arguments to be passed to (for example) external processes.
- **VarSet:** Selects a variable where decoded data is placed. If more than one piece of data is being decoded by the data element, the variable is the starting point of the array being filled with decoded data.

2.4.1 Data type field

The description language is a typed system. Every data descriptor must indicate the type of data it represents; they are hard coded functions that define the “type” of each chunk of “data” on a block. Just like the C programming language where you might define a data type to be “**byte**”, “**int**” or “**long**”, here we define things as “**Mark**”, “**Sync**”, “**Data**”, etc.

2.4.2 Some Data type

Here are some examples of data type that we will be using in the [Atari example](#).

2.4.2.1 Area start/stop

Sets the beginning and the end of a logical data block. A data block is an area made of one or more data elements. After each we define the area number, which can be used to reference the area. This is any arbitrary value in the range **1..n**.

IPF – Interchangeable Preservation Format Documentation

2.4.2.2 Encode

An arbitrary number of layers can be applied to any area. The encoding layer has the following parameters:

- count: Area number where encoding applies.
- value: Physical encoding (MFM, MFM2, FM, GCR...).
- subtype: Logical encoding
- process: Any process to be applied over the data, like encryption.

2.4.2.3 Mark / Sync

Also known as **Sync**, renamed to be more in line with Trace terminology. Treated as data, but encoding can be modified, and of course used to spot areas on a track in the first place.

- count: number of sync
- value: sync value
- subtype: 1 (one byte)

2.4.2.4 Data

Data area, all of which support:

- count: Count of data in its un-encoded format, i.e. the real size, which completely depends on the encoding used.
- value: Expected value.
- subtype: Type of un-encoded data, 1=byte, 2=word, 3=long-word.

2.4.2.5 Gap

This is another data area, but its size may be altered via mastering - if it is allowed, and errors found after it can normally be ignored.

2.4.2.6 EDC / Checksum

Better known as a **Checksum**, but this is more in line with Trace terminology. A checksum is generated for the selected area, and stored on the spot of the command. Checksums can be nested and/or overlapped; the **Analyzer** resolves the dependencies and thus the correct order of calculating the values.

- count: Area number where the process is performed on original data. Due to various shortcuts in their code, some EDCs seem like they were originally performed on raw data but in reality they worked with un-encoded data.
- value: Starting value for the checksum process. Sometimes it is an arbitrary value, though most of the time 0.
- subtype: Type of unencoded data, 1=byte, 2=word, 3=longword.
- process: The checksum method.

2.4.3 Data type table

Function	Type	Count	Value	SubType	Process	VarSet
Encode	Area encode	BP #0	BP #1	BP #2	BP #3 (unused in example)	BP #4 (unused in example)
Area start	Area start	BP #0	n/a	n/a	n/a	n/a
Gap, byte	Gap	BP #0	BP #1	1 (i.e. byte)	n/a	n/a
Mark, byte	Mark	BP #0	BP #1	1 (i.e. byte)	BP #2	n/a
Data, byte	Data	BP #0	BP #1	1 (i.e. byte)	n/a	BP #2 (unused in example)
Data, byte f0	Data	BP #0	FP #0	1 (i.e. byte)	n/a	BP #2 (unused in example)
Sector, byte	Data	BP #0	FP #0	1 (i.e. byte)	n/a	n/a
Area stop	Area stop	BP #0	n/a	n/a	n/a	n/a
EDC, Xor16, long	EDC	BP #0	BP #1	3 (i.e. longword)	5 (i.e. CRC16 method)	n/a

IPF – Interchangeable Preservation Format Documentation

2.5 Atari Example

Now we have gone through a short description of the descriptors let's see an example of how we describe a standard Atari track format.

2.5.1 Atari Format/Track descriptor

```
.Header, Data block: 1, 512 *1 <1st>  
.Header, Data block: 2, 512 *1  
.Header, Data block: 3, 512 *1  
.Header, Data block: 4, 512 *1  
.Header, Data block: 5, 512 *1  
.Header, Data block: 6, 512 *1  
.Header, Data block: 7, 512 *1  
.Header, Data block: 8, 512 *1  
.Header, Data block: 9, 512 *1
```

Each line looks like this

```
[Block Descriptor] : [block_number], [number_bytes], *weight
```

The parameters are passed down to the block and data descriptors. These lower level descriptors have access to Format Parameter #0, Format Parameter #1 and Format Parameter #2. So now let's have a look at the data descriptors that access them.

- Format Parameter #0: set to 1...9, used by the "Sector, byte" data descriptor, i.e., we search for blocks (sector) 1 to 9.
- Format Parameter #1: set to 512, used by the "Data, byte f1" data descriptor.

You might wonder why not use fixed values. The reason is that many games use slight differences in the format descriptor.

2.5.2 Atari Block & Data descriptors

```
Encode: 1, 1  
_Area start: 1  
  _Area start: 2  
    Mark, byte: 3, $4489 *1 <esc>  
    ID, byte: 1, $FE *1 <esc>  
    Track, byte: 1 *1  
    Side, byte: 1 *1  
    Sector, byte f0: 1 *1 <esc>  
    Length, byte: 1 *1  
  Area stop: 2  
  EDC, CRC16: 2, $FFFF *2  
  GAP,byte: 34 *1  
  _Area start: 3  
    Mark, byte: 3, $4489 *1 <esc>  
    ID, byte: 1, $FB *1 <esc>  
    Data, byte f1 *1  
  Area stop: 3  
  EDC, CRC16,short: 3, $FFFF *3  
_Area stop:1  
Set GapPattern: 40 $4E  
Set GapPattern  
Set GapPattern: 12, $0  
Set GapPattern
```

Description of the statements

```
Encode: 1, 1
```

First we define Area 1 (first parameter) to be encoded as MFM (second parameter). The whole block is covered by Area 1, and therefore all data elements are set to be MFM encoded.

```
_Area start: 1
```

Start Area 1. Just marking an area to cover all the data to encode.

```
_Area start: 2
```

Next is where Area 2 starts. This is the ID/Address area of the sector.

IPF – Interchangeable Preservation Format Documentation

Mark, byte: 3, \$4489 *1 <esc>

We have 3 sync bytes with value 4489 (\$A1 with missing clock) in MFM. Weight 1. We skip (escape) to the next sync if there is no match

ID, byte: 1, \$FE *1 <esc>

1 ID byte with value \$FE, weight 1, and skip to next block if not equal

Track, byte: 1 *1

1 byte that contains the track value weight 1

Side, byte: 1 *1

1 byte that contains the side value weight 1

Sector, byte f0: 1 *1 <esc>

1 byte that contains the sector value. This value must match the value given in parameter 0 of the format descriptor (Format Parameter #0). Skip to next block if not equal

Length, byte: 1 *1

1 byte that contains the size value

_Area stop: 2

end of the ID/Address field of the sector.

EDC, CRC16: 2, \$FFFF *2

The checksum for the ID/Address field is next, which is a CRC16 type checksum on Area 2. Starting value=\$FFFF. It is weighted by a factor of 2 so the score will be worse if the checksum calculated on the header does not match the value stored here.

GAP,byte: 34 *1

There are then 34 bytes, with the value of 0 defining a gap. Weight 1x.

_Area start: 3

Next is where Area 3 starts. This is the Data area.

Mark, byte: 3, \$4489 *1 <esc>

We have 3 sync bytes with value 4489 (\$A1 with missing clock). Weight 1. We skip (escape) to the next sync if there is no match

ID, byte: 1, \$FB *1 <esc>

1 ID (identification) byte with value \$FB, weight 1, skip to next block if not equal

Data, byte f1 *1

Then we have the data field. The length of the field is fixed by parameter 1 (sector size = 512) of the format descriptor (Format Parameter #1).

_Area stop: 3

end of the data field of the sector.

EDC, CRC16,short: 3, \$FFFF *3

The checksum for the data field is next, which is a CRC16 type checksum on Area 3. Starting value=\$FFFF. It is weighted by a factor of 3 so the score will be worse if the checksum calculated on the data does not match the value stored here, i.e. the integrity of data area is the most important.

_Area stop:1

end of area 1.

Set GapPatern: 40 \$4E

Set GapPatern

Set GapPatern: 12, \$0

Set GapPatern

These lines define the forward and backward gaps between sectors.

3. IPF/IPX file content

This section provides information about the content of IPF/IPX files, produced by CAPS Analyzer, as it would be read for example with a hexadecimal editor. This information is useful to better understand what is stored in IPF/IPX file and can be used to directly read or write IPF files.

- ✎ However it is **highly recommended** not to access the content of IPF files directly but through the [IPF/CAPS library](#) since the IPF format is complex and updated on regular basis in order to add new capabilities.
- ✎ The IPF/IPX files content description provided here is largely based on my interpretation of the [IPF/CAPS library](#) and therefore beware that some descriptions are speculative and may be wrong. So again prefer the usage of the provided library.

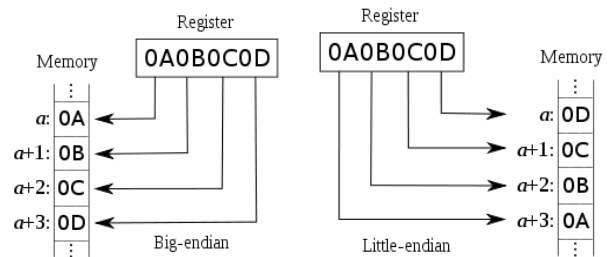
3.1 IPF File organization

The IPF file is organized in chunks of data that I call **records**. A record is composed of a fixed size **Header** block followed, by a **Type Specific** block and eventually an optional **Extra Data** block.

Therefore a generic IPF record looks like this:



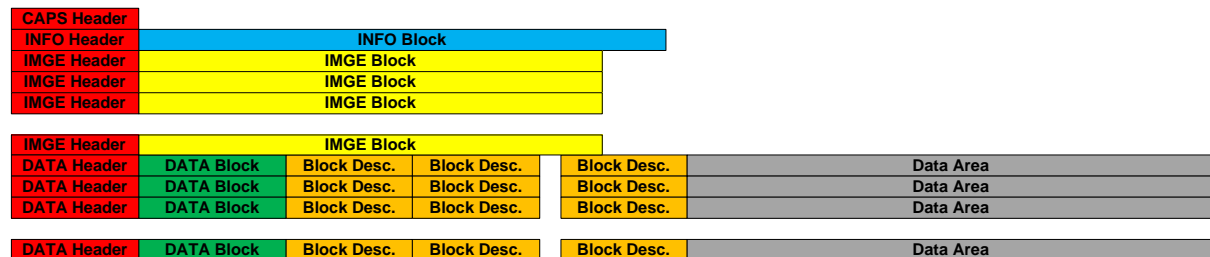
Most of the information in a block is stored as **4 bytes** (32 bits) integers using **big-endian** ordering. When reading the information on an **Intel** platform (**little-endian**) it is necessary to swap the order of the bytes to interpret the information correctly. Each block of data in an IPF file is associated with a CRC32 value that allows to verify that the block of data read has not been corrupted.



An IPF file describes the content of a floppy disk and is typically composed of the following records:

- First a **CAPS** record to identify an IPF file
- This record is followed by an **INFO** record that gives general information about the IPF file
- This record is followed by **IMGE** records: usually one per track and per side of the floppy disk.
- These records are followed by **DATA** records: one for each IMGE record.

This can be represented with the following diagram:




IPF – Interchangeable Preservation Format Documentation

3.2 IPX File Organization

The CAPS Analyzer can write an IPX file at the same time it writes an IPF file. An IPX file provides information about the analysis of a floppy disk and is composed of the following records:

- First a **CAPS** record to identify an IPF/IPX file
- This record is followed by an **INFO** record that gives general information about the IPF file
- This record is followed by a **CTEI** records: provides information about the **Analyzer**.
- These records are followed by **CTEX** records: one for each track.

This can be represented with the following diagram:



3.3 Record Header

A record header is composed of 3 fields:

- Name (4 character string): that indicates the record's type. Currently defined types:
 - ◆ **CAPS**: Used to identify an IPF file
 - ◆ **DUMP**: Unknown by me
 - ◆ **DATA**: contains the actual data from the floppy disk
 - ◆ **TRCK**: Unknown by me
 - ◆ **INFO**: contains general information about the IPF file
 - ◆ **IMGE**: contains track information
 - ◆ **CTEI**: Contains information about the Analyzer.
 - ◆ **CTEX**: Contains information about analysis of a track.
- Size (4 bytes): Size of the complete record (i.e. header followed by an optional block) in bytes.
- CRC32 (4 bytes): Store the computed CRC32 value for the complete record (i.e. the record header block plus the specific block). The CRC value is computed on the complete record with the CRC field set initially at 0x0000, and then the resulting CRC32 value is recorded in this field.

3.4 CAPS Record

CAPS Header

The content of the name field in the record header is “CAPS” in uppercase.

The CAPS record must be the first record in the file and is used to identify a CAPS/IPF file.

A CAPS record is composed of only a header block and therefore is not followed by any other block as it is the case for all other types of records.

The overall size of a CAPS record is 12 bytes.

3.5 INFO Record

INFO Header

INFO Block

The content of the name field in the record header is “INFO” in uppercase.

The INFO record provides general information about the IPF file.

The INFO header is followed by an INFO block that contains the following information:

- Type: The type of image
 - ◆ 01 = Floppy disk (the only type supported currently).
- EncoderID: The image encoder ID. Depending on the encoder ID some fields in the IMGE records are interpreted differently. Currently this field can take the following values:
 - ◆ 01 = CAPS encoder
 - ◆ 02 = SPS encoder
- EncoderRev: The image encoder revision. For a specific EncoderID several revisions can exist. Currently only revision 1 exists for CAPS or SPS encoders.
- ReleaseID: Each IPF file has a unique ID that can be used as a unique key for database. More than one file can have the same ID for example for a game that has more than one disk.
- RevisionID: Revision of the file. Normally the revision is one.
- Origin: A reference to the original source.
- MinCylinder: The lowest cylinder (track) number of the floppy image. Usually 0
- MaxCylinder: The highest cylinder (track) number of the floppy image. Usually 83
- MinHead: The lowest head (side) number of the floppy image. Usually 0
- MaxHead: The highest head (side) number of the floppy image. Usually 1
- CreationDate: The image creation date packed into two 4 bytes integers. Specify the year, the month, the day, the hour, the minute, the second, and the tick.
- Platforms: An array of four 4 bytes integers used to store the value of the intended platforms (one image can run on several platforms). Platform can take the following values:
 - ◆ 01 = Amiga
 - ◆ 02 = Atari ST
 - ◆ 03 = PC
 - ◆ 04 = Amstrad CPC
 - ◆ 05 = Spectrum
 - ◆ 06 = Sam Coupe
 - ◆ 07 = Archimedes
 - ◆ 08 = C64
 - ◆ 09 = Atari 8-bit
- DiskNumber: Number of the disk in a multidisc release otherwise 0
- UserID: A unique user ID of the disk image creator.
- Reserved: an array of three 4 bytes integer reserved for future use.

The overall size of an INFO record is 96 bytes (12 for the header block and 84 for the Info block).

3.6 IMGE Record



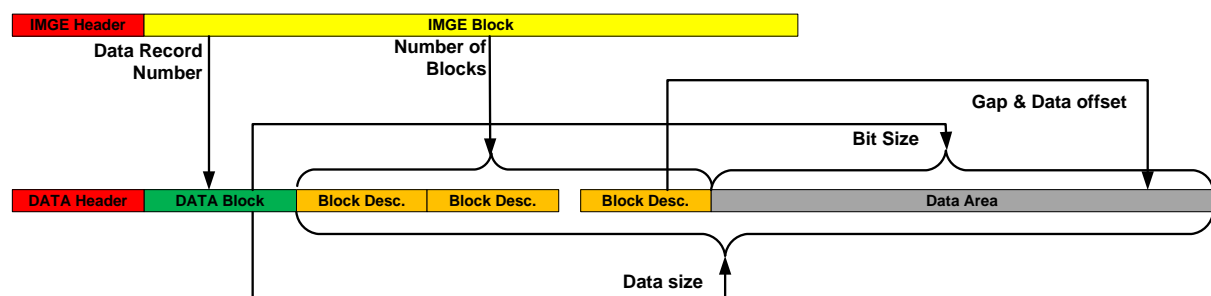
The content of the name field in the record header field is “IMGE” in uppercase.

The IMGE record provides general information for a specific track of the floppy.

The IMGE header is followed by the IMGE block that contains the following information:

- Cylinder: The cylinder (track) number. Usually in the range 0-83
- Head: The head (side) number. Usually 0 or 1
- Density Type: Indicate the density type. Currently density can take the following values:
 - ◆ 01 = Noise cell are unformatted (random size)
 - ◆ 02 = Auto automatic cell size, according to track size
 - ◆ 03 = Copylock Amiga
 - ◆ 04 = Copylock Amiga new
 - ◆ 05 = Copylock ST
 - ◆ 06 = Speedlock Amiga
 - ◆ 07 = Old Speedlock Amiga
 - ◆ 08 = Adam Brierley Amiga
 - ◆ 09 = Adam Brierley, density key Amiga
- SignalType: Signal processing type. Currently the only supported type is 1 that indicates 2μs cells
- TrackSize: Rounded decoded track size. The total number of bytes (clock + data) on the track
- StartPos: The rounded StartBits position.
- StartBits: Start position of original data from Index
- DataBits: number of decoded data bits (clock + data)
- GapBits: number of decoded gap bits (clock + data)
- TrackBits: The total number of bits on the track. Track Bits = Data Bits + Gap Bits
- BlockNumber: Number of blocks used to describe one track. Usually (but not necessarily) this is the number of sectors of the current track. This field is used to specify the number of data block descriptors in the matching [DATA record](#) (see below).
- EncoderProcess: Encoder process?
- Image Flag:
 - ◆ Bit 0 = Fuzzy
- DataID: Data Record ID. Unique ID that allow pointing to the matching [DATA record](#).

An IMGE record contains a unique identifier that allows pointing to a specific data record and it also contains the number of blocks used to describe the content of the tracks.



The overall size of an IMGE record is 80 bytes (12 for the header block and 68 for the IMGE block).

3.7 DATA Record



The content of the name field in the record header is “DATA” in uppercase.

The DATA record provides the data information for a specific track of the floppy.


The DATA header block is followed by the DATA block, and eventually by an Extra Data Block.

The DATA block contains the following information:

- ByteSize: size in bytes of the Extra Data Block following the data block (can be zero).
- BitSize: data area size in bits
- CRC32: CRC32 of the Extra Data Block
- DataID: Data Record ID. Unique ID used by [IMGE record](#) to find the corresponding DATA record.

If ByteSize is null it indicates that the data block is **not** followed by an Extra Data Block.

The Extra Data Block starts with n **Block Descriptors** (the number n is specified in the matching [IMGE record](#)) and is followed by the actual data content for all these blocks



3.7.1 Block Descriptor

Each Block Descriptor contains the following information:

- BlockBits: Decoded block size in bits
- GapBits: Decoded gap size in bits
- If the encoder type = 1 (CAPS) in the [INFO record](#) we have the following fields
 - ◆ BlockSize: decoded block size, rounded
 - ◆ GapSize: decoded gap size, rounded
- Else if the encoder type = 2 (SPS) in the [INFO record](#) we have the following fields:
 - ◆ GapOffset: offset of gap stream in data area
 - ◆ CellType: bit cell type (1 = 2 μ s cell)
- EncoderType: Block encoder type
 - ◆ 01 = MFM
 - ◆ 02 = Raw
- BockFlag: These flags are set to zero and ignored if EncoderID = 1 (CAPS) in [INFO record](#)
 - ◆ Bit 2 StreamUnit: = Unit of the Data stream sample length: 0 = bytes, 1 = bits.
Warning: only apply to data stream but **not** to gap streams (always in bits).
 - ◆ Bit 1 BackwardGap = When set it indicates an associated backward gap stream
 - ◆ Bit 0 ForwardGap = When set it indicates an associated forward gap stream
- GapValue: Default Gap value
- DataOffset: Offset value of the data stream in the extra data area

The Block Descriptors are followed by the actual data stream area. To get to a specific data or gap stream for a specific block you must use the DataOffset and/or the GapOffset information from the Block Descriptor.

3.7.2 Data Area

The Data Area of the Extra Data Block is composed of blocks of data (one for each block descriptor) that contain an optional list of data stream elements of length BlockBits, followed by an optional list of gap stream elements of length GapBits. The bit lengths indicate the length of encoded raw MFM.

3.7.2.1 Data Streams

If the data stream list exists (i.e. BlockBits is not null) the DataOffset indicates the position of the first Data Streams element from the start of the extra data block.

Each data stream element is composed of the following sequence:

- Type/Length (one byte)
 - ◆ SizeWidth (3 MSB - bits 7-5): number of bytes used to specify the sample size
 - ◆ DataType (5 LSB - bits 4-0): contain the type of data in the stream:
1 = Sync, 2 = Data, 3 = Gap, 4 = Raw, 5 = Fuzzy
- DataSize (SizeWidth bytes): The size (specified by SizeWidth bytes using big-endian ordering) of the following data stream samples. The size unit is specified by the StreamUnit flag (bit 2) in the block descriptor: 0 = size in bytes, 1 = size in bits.
- DataSample (StreamSize bytes): The actual stream samples composed of StreamSize bytes.

The list of data stream elements is terminated by a stream element with a null Type/Length.

3.7.2.2 Gap Streams


If the Gap stream list exists (i.e. GapBits is not null) the GapOffset indicates the position of the first gap streams element from the start of the extra data block.

There can be 0, 1, or 2 Gap stream list depending of the BlockFlag bits 1-0 in the [Block descriptor](#).

- Bits 1-0 = 00: No gap stream. In this case the gap is filled with the GapValue byte specified in the Block Descriptor. This byte will be repeated (including a final partial repetition, if necessary) as many times as necessary to fill the gap. The decoder fills forward and backward in the gap at the same time, hence the “write splice” or “overlap” occurs in the exact middle of the gap area.
- Bits 1-0 = 01: Forward gap stream list only. The gap stream area contains a list of gap stream elements used to fill the gap in forward direction (from the end of current data block to the beginning of the next one).
- Bits 1-0 = 10: Backward gap stream list only. The gap stream area contains a list of gap stream elements used to fill the gap in backward direction (from the beginning of the next data block to the end of the current one).
- Bits 1-0 = 11: Forward and Backward stream lists. The gap stream area contains two lists of gap stream elements used to fill the gap in forward direction (from the end of current data block to the beginning of the next one) and in backward direction (from the beginning of the next data block to the end of the current one) simultaneously.

Each gap stream element is composed of the following sequence:

- Type/Length (one byte)
 - ◆ SizeWidth (3 MSB - bits 7-5): number of bytes used to specify the sample size
 - ◆ GapType (5 LSB - bits 4-0): contain the type of gap in the stream: 1 = GapLength, 2 = Data
- GapSize (SizeWidth bytes): The size (specified by SizeWidth bytes using big-endian ordering).
 - ◆ If GapType=1 this size indicates the number of times the GapValue must be repeated.
 - ◆ if GapType=2 this size indicates the size in bits of the following GapSample
- GapSample (GapSize bits): This is a field that only exists for GapType=2 and contains the actual GapSample of GapSize bits.

 For a forward gap stream the GapSize of the final GapSample is optional. This indicates that GapSample will be repeated as many times as necessary to fill the gap.
For a backward gap stream the GapSize of the first GapSample is optional. This indicates that GapSample will be repeated as many times as necessary to fill the gap.

The list of gap stream elements is terminated by a stream element with a null Type/Length.

3.8 CTEI Record



The content of the name field in the record header field is “CTEI” in uppercase.

The CTEI record provides information about the analyzer used to create the IPF file.

The CTEI header is followed by the CTEI block that contains the following information:

- Release_CRC (4 bytes): The CRC of the associated IPF file released
- Analyzer_Rev (4 bytes): The revision of the Analyzer used
- Reserved (14 bytes): Reserved for future used usually set to 0

The overall size of a CTEI record is 76 bytes (12 for the header block and 64 for the CTEI block).

3.9 CTEX Record



The content of the name field in the record header field is “CTEX” in uppercase.

The CTEX record provides analyzer information for a specific track of the floppy.

The CTEX header is followed by the CTEX block that contains the following information:

- Cylinder (4 bytes): The cylinder (track) number. Usually in the range 0-83
- Head (4 bytes): The head (side) number. Usually 0 or 1
- Density Type (4 bytes): Indicate the density type. Currently density can take the following values:
 - ◆ 01 = Noise cell are unformatted (random size)
 - ◆ 02 = Auto automatic cell size, according to track size
 - ◆ 03 = Copylock Amiga
 - ◆ 04 = Copylock Amiga new
 - ◆ 05 = Copylock ST
 - ◆ 06 = Speedlock Amiga
 - ◆ 07 = Old Speedlock Amiga
 - ◆ 08 = Adam Brierley Amiga
 - ◆ 09 = Adam Brierley, density key Amiga
- Analyzer_Format_ID (4 bytes): Analyzer selected format
- Analyzer_Fix (4 bytes): Analyzer Format Fix method
- Analyzer_Track_Size (4 bytes): Analyzer track size
- Reserved (2 bytes): Reserved for future used usually set to 0

The overall size of a CTEX record is 44 bytes (12 for the header block and 32 for the CTEX block).

4. Atari Example

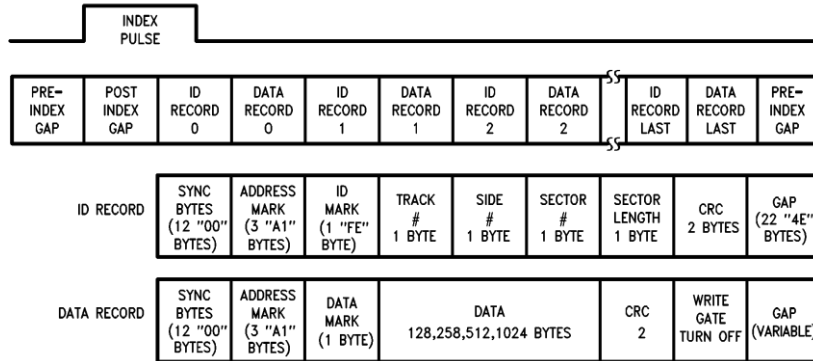
In this section we are going to analyze how the information on an Atari Floppy Disk is stored in an IPF file.

4.1 Atari FD Format short presentation


First let's quickly review the format used on the Atari machines. The Atari ST uses the Western Digital WD1772 Floppy Disc Controller (FDC) to access the 3 1/2 inch floppy diskettes.

4.1.1 Atari Track Format

Below is a description of a **Standard Atari Double Density Format** as created by the TOS.



Note: Many different conventions have been used to decompose and name the different GAPS of a track. In this document I use a GAP numbering scheme which is a combination of the IBM and ISO standards. The GAP between the ID record and the DATA record is decomposed into an ID **postamble** gap (Gap 3a) and a DATA **preamble** gap (Gap 3b). This allows a more detail description, but of course they can be recombined into a more standard inter gap (Gap3).




4.1.2 Atari Sector Format

A sector is composed of an **ID Segment** that allow identifying the sector number, and a **Data Segment** that contains the actual data. These two segments includes at their ends specific gaps.

The ID and DATA preamble gaps are used to provide enough time to lock the PLL and are normally composed of 12 \$00 bytes. The ID postamble gaps provides enough time for the WD1772 to get prepared for execution of a read or write command on the matching Data segment and is typically composed of 22 \$4E bytes. The Data postamble gaps provides enough space between sectors for the WD1172 to get prepared to work on the next sector and it also provides room for *data segment drift* during write operation. It is typically composed of 40 \$4E bytes The Atari FD format do not specifically define a sector write splice byte (where the head write current is switched on or off) but in practice it should be considered as part of the data preamble field for format and write operations, and as part of the ID postamble for read operation (see below).

One complete sector looks like this



IPF – Interchangeable Preservation Format Documentation

4.1.3 Atari Sector write splices

It is not possible with the WD1772 FDC to write an entire track in one operation. This is due to the fact that when writing a track the data in the range \$F5 to \$F7 are treated as special control bytes and therefore they cannot be written directly during a write track (format) operation. Therefore on an Atari the first operation consists in formatting a track with a **format command** then the data field of each sector is written using **write sector commands**.


Here is a simplified description of the write sector command:

Upon receipt of the Write Sector Command the FDC searches on the track an ID field that has the correct track number, correct sector number, and correct CRC. If such an ID field is found the WD1772 counts **22 bytes** from the end of the ID field CRC then it activates the WG output.

The following data are then written to the disk:

- Twelve bytes of zeroes
- Three A1 Sync
- One Data Address Mark
- The complete data field (usually 512 bytes)
- A two-byte computed CRC
- *One byte of logic ones.*

Then the WG output is deactivated.



It is easy to understand that the activation and deactivation of the Write Gate cannot be aligned precisely with the original data written during the format operation and furthermore, as the speed of the drive fluctuates, the overall size of the data segment will most likely not be exactly the same. This results in writing non-aligned (drifted) transitions, which most likely violates MFM rules, when the WG is activated **and** deactivated. These two areas are called *sector write splices*.


When data are written on a floppy disk with professional duplication equipment the complete track is written in one operation and therefore these sector write splice areas do not exist. This is one way for the Analyzer to verify if we are working on a “master floppy disk” written on professional duplication equipment and that the sector data on the FD has not been tempered.

4.1.4 Atari Track write splices

As said before when using professional duplication equipment it is possible to write a complete track in one operation and this will result in a track without any sector write splices. The writing of the complete track starts at a particular point of the track and ends up at a particular point (typically the same point). For a non-protected track these starting (WG activation) and ending (WG deactivation) points are usually located at the track index.


But even with professional equipment it is not possible to activate and deactivate the write gate at exactly the same position. This results in writing a non-aligned (drifted) transition, which most likely violates MFM rules, when the WG is deactivated. This area is called the *track write splice*.

In the case of a standard track the writing starts in the Post-index gap G1 and stops in the Pre-index gap G5 and therefore the track splice happen in this non critical area.



IPF – Interchangeable Preservation Format Documentation

However some protections are based on shifting the position of the complete track in respect with the index. For example the index might be positioned in the middle of the last data field (this protection is often referred as *data over the index*).



In such a case it is not possible to activate and deactivate the Write Gate at the position of the index because this would result in a track write splice located in the middle of the data field and therefore this will result in corrupted data.

For this kind of protection it is therefore important for the duplication equipment to activate and deactivate the write gate in a position which is still located at the boundary of the pre-index and post-index gaps, knowing that this location is **not** aligned with the index.

4.2 Example of Theme Park Mystery FD

It is interesting to use a real Atari FD example to illustrate the information stored in a stream file and an IPF file. For that matter we will take the Atari's **Theme Park Mystery** game. We first analyze the Stream files of this game using my KFAalyze program (for more information refer to [KFAalyze documentation](#) as well as my document of [Atari copy protection](#)) then we will analyze the content of the IPF file as produced by the CAPS Analyzer.

4.2.1 Analysis of Track 00.0 in the Stream file

Let's first look at the [results of analysis](#) with the KFAalyze program of the flux transitions as provided by the Kryoflux board on track 00 head 0 of the game.

Several protections are detected by KFAalyze on this track:

- A No Flux Transition Area (NFA)
- Sectors with Fuzzy Data (FZD)
- Data Over Index (DOI)
- Sector within sector (SWS)
- Sector with bad data (SBD) - CRC Errors

As well as timing violations, timing in ambiguous area, and sync errors.


The picture below shows the overall layout of the track as detected by KFAalyze:



We can see that the last two sectors are detected as continuing beyond the end of the track (DOI) that they overlap (SWS) and as they contain Fuzzy data (FZD displayed in orange). But we can also see that sector 1 starts at a normal position and therefore this indicates that sectors 11 & 12 are in fact special short sectors.

In order to see the no flux transition area (NFA) we look at the detail graphic for sector 11.

IPF - Interchangeable Preservation Format Documentation



Here we can see that around the position 194500µs till 200000µs (the end of track) there are **no flux transitions**. This NFA is directly located at the end of the track inside the sector 11 and 12. This alone explains why these two sectors read with fuzzy bits due to multiple timing errors.

One last thing to mention: The sector 12 is located inside sector 11. But if we look at the data inside sector 11 **we do not see** any sequence of 3 \$A1 sync mark followed by an \$FE Index Address Mark!

```

Detail buffer content for sector 11 with 554 bytes
= ID=11 7 bytes @190229 length=225.91 T=0 H=0 S=11 Z=512 OK TMV=0 BRD=0 BS=0
 0000 190229 4000 fe 00 00 0b 02 25 a4 .....%.
+ GAP3 31 bytes @190455 us length=990.74 us - TMV=0 BRD=0 BS=0 IDG=0 SS=0
 0007 190455 4000 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e NNNNNNNNNNNNNNNNN
 0017 190968 4000 4e 4e 4e 4e 4e 00 00 00 00 00 a1 a1 NNNNNNN.....
= DATA ID=11 515 bytes @191446 us length=16411.99 CRC BAD CLK=3.98 TMV=2 BRD=1 DOI=250 BS=6
*** Fuzzy Sector *** starting at byte position 265
 0026 191446 4000 fb 00 00 00 00 00 00 00 00 a1 a1 a1 fe 00 00 0c .....
 0036 191961 4000 02 bc 33 4e 4e 4e 4e 4e 4e 4e 4e d9 23 76 c5 e6 ..3NNNNNNNNN.#v..
 0046 192474 4000 d3 31 b2 4e 4e 4e 4e 4e 4e 4e 4e 4e ff ff ff ff ff .1.NNNNNNNNN.....
 0056 192982 3968 fe 14 14 14 00 ff ff ff ff ff ff ff ff ff ff ff .....
 0066 193492 4000 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff .....
    
```

This was done on purpose to render the detection of a somewhat hidden sector 12 inside sector 11 more difficult. But if we also ask KFAalyze to display the clock at the same time as the data we have:

```

Detail buffer content for sector 11 with 515 bytes
= DATA ID=11 515 bytes @191446 us length=16411.99 CRC BAD CLK=3.98 TMV=2 BRD=1 DOI=250 BS=6
*** Fuzzy Sector *** starting at byte position 265
 0000 191446 4000 fb 00 00 00 00 00 00 00 00 a1 a1 a1 fe 00 00 0c .....
                00 7f ff ff ff ff ff ff ff 0a 0a 0a 00 ff ff f1 .□.....
 0010 191961 4000 02 bc 33 4e 4e 4e 4e 4e 4e 4e 4e d9 23 76 c5 e6 ..3NNNNNNNNN.#v..
                fc 01 c4 10 90 90 90 90 90 90 90 90 90 02 4c 00 18 08 .....L...
 0020 192474 4000 d3 31 b2 4e 4e 4e 4e 4e 4e 4e 4e ff ff ff ff ff .1.NNNNNNNNN.....
                04 46 04 90 90 90 90 90 90 90 90 90 00 00 00 00 .F.....
 0030 192982 3968 fe 14 14 14 00 ff ff ff ff ff ff ff ff ff ff ff .....
                00 a1 a1 a1 fb 00 00 00 00 00 00 00 00 00 00 00 00 .....
 0040 193492 4000 ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff .....
                00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
    
```

Here we can clearly see the 3 \$A1 followed by the \$FB hidden **inside the clock** information. This indicates that the data for sector 12 are **shifted by one half bit** from the data in sector 11. This will be important when we will look at the dump of the IPF file.

IPF – Interchangeable Preservation Format Documentation

4.2.2 Analysis of track 00.0 in the IPF file

Analysis of track 00.0 indicates that it uses a **Copylock short** format. This format has 10 normal sectors (as already described in [Atari Example](#)) followed by a **.Copylock short sgn block** in sector 11, followed by a **.Copylock short key block** in sector 12.

The complete track format is therefore:

```
.Header, Data block: 1, 512 *1 <1st>
.Header, Data block: 2, 512 *1
.Header, Data block: 3, 512 *1
.Header, Data block: 4, 512 *1
.Header, Data block: 5, 512 *1
.Header, Data block: 6, 512 *1
.Header, Data block: 7, 512 *1
.Header, Data block: 8, 512 *1
.Header, Data block: 9, 512 *1
.Copylock short sgn block: 11, *1 <->
.Copylock short key block: 12 *1
```

We will detail the last two Block Descriptors later during the description of the content of the IPF file for track 00.0.

The IPF file starts with a “**CAPS**” and “**INFO**” records that contain the following information:

```
DrCoolZic - IPF Info v0.3a - Working on file ThemeParkMystery.ipf
Record CAPS length=12 crc=1cd573ba (0 - 12)
Record INFO length=96 crc=db83cf07 (12 - 108)
  Type='Floppy Disk' Encoder='SPS (rev 1)' RelID=3211 (rev 1) Origin=f956bf3c
  Cyl=0-81 Head=0-1 Stamp='2010/6/12 1:14:7:765'DiskNum=0 UserId=0
  Platform=Atari ST
```

Following the INFO record we have many IMGE records: one for each track and each side.

Let’s look at the IMGE record for track 00.0:

```
Record IMGE length=80 crc=3a68eb7a (108 - 188)
  Track=00.0 Size=12556 Bits=100448 (D=93056 + G=7392) StartPos=54 StartBits=439
  DataID=1 BlockCnt=12 DensType='Auto' SigType=2us-cell EncProc=0 Flag=0
```

The Track 00 Head/Side 0 contains 12556 MFM (clock + data) bytes equivalent to 100448 bits that are decomposed as 93056 bits of Data and 7392 bits of Gap.

The track starts at position 54 bytes or more precisely 439 bits. This DataID of this first track is set to 1 and contains 12 blocks. The density (cell size) is adjusted automatically with the track size.

Now let’s look at the matching DATA record (DATA Block with the same DataID=1 value).

```
Record DATA length=28 crc=5adc8a4c (13228 - 13256)
  DATA Block BlockSize=6652 BitSize=53216 CRC=12C81870 DataID=1 (13256 - 19908)
```

The DATA Block indicates that there is an **Extra data** record of 6652 bytes that contains: all the block descriptors as well as the Gap and Data stream information. This extra block is “protected” by the provided CRC.

Following the data block we have the block descriptors (9 in this case).

Let’s look at the first one (block 0 at offset 0 in the extra data record):

```
BDesc=0 BlkBits=8992 GapBits=512 GapOfs=384 CType=1 EType=1 Flg=3 GapVal=0 DataOfs=515 (13256 - 13288)
```

This block descriptor 0 indicates that the data area on the track is 8992 bits long in MFM equivalent to 562 bytes (8992/16) of data and has gap areas 512 bit long in MFM equivalent to 32 bytes (512/16).

IPF - Interchangeable Preservation Format Documentation

Let's first look at the Data information. It is interesting to note that the data field contains the FD sector ID field, followed by the intra sector ID-Data gap, followed by the FD sector Data field:

```
-- DATA 8992 bits (562 bytes) @13771
Sync (Type=1) size=6 Bytes (@13771)
13773 44 89 44 89 44 89 D.D.D.
Data (Type=2) size=7 Bytes (@13779)
13781 fe 00 00 01 02 ca 6f .....o
Gap (Type=3) size=34 Bytes (@13788)
13790 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e NNNNNNNNNNNNNNNNN
13806 4e 4e 4e 4e 4e 4e 00 00 00 00 00 00 00 00 00 NNNNNNN.....
13822 00 00 ..
Sync (Type=1) size=6 Bytes (@13824)
13826 44 89 44 89 44 89 D.D.D.
Data (Type=2) size=515 Bytes (@13832)
13835 fb 60 1e 00 00 00 00 00 00 58 00 00 00 02 02 01 .`.....X.....
13851 00 02 70 00 20 03 00 05 00 0a 00 01 00 00 00 00 ..p. ....
...
...
14331 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ad .....
14347 0a b7 2c
```

We can see at the beginning of the sector the 3 sync bytes (\$A1) displayed as MFM encoded bytes \$4489 (\$A1 with missing clock bit):

```
Sync (Type=1) size=6 Bytes (@13771)
13773 44 89 44 89 44 89 D.D.D.
```

This is followed by the sector ID field. We have the IAM \$FE followed by the Track, Side, Sector, Size and the two CRC bytes. We can see that this is sector 1

```
Data (Type=2) size=7 Bytes (@13779)
13781 fe 00 00 01 02 ca 6f .....o
```

This is followed by the intra-sector ID/Data gap: 22 x \$4E bytes + 12 x \$00 bytes

```
Gap (Type=3) size=34 Bytes (@13788)
13790 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e 4e NNNNNNNNNNNNNNNNN
13806 4e 4e 4e 4e 4e 4e 00 00 00 00 00 00 00 00 00 NNNNNNN.....
13822 00 00 ..
```

This is followed by 3 sync bytes (\$A1):

```
Sync (Type=1) size=6 Bytes (@13824)
13826 44 89 44 89 44 89 D.D.D.
```

Followed by the actual data field: First the DAM (\$FB) followed by 512 bytes of data, followed by 2 bytes of CRC

```
Data (Type=2) size=515 Bytes (@13832)
13835 fb 60 1e 00 00 00 00 00 00 58 00 00 00 02 02 01 .`.....X.....
13851 00 02 70 00 20 03 00 05 00 0a 00 01 00 00 00 00 ..p. ....
...
...
14331 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ad .....
14347 0a b7 2c
```

It is interesting to note that in the IPF we have kept the ID and Data fields of the sector into one block. Therefore it also contains the inter-field gap into the block. This gap is where a potential “sector write splice” can occur. But as we are working on master FD created on professional duplicator machines there should not be any write splice in this area (otherwise this would indicate a tempered FD).

IPF – Interchangeable Preservation Format Documentation

Now let's look at the Gap stream data. The flag in the block descriptor is equal to 3. This indicates that we have *forward and backward* gap streams. The forward gap stream is from the end of data of this block toward the data of the next block and the backward gap stream is from the beginning of this data block toward the end of the previous data block. We find the following information:

```
-- GAP 512 bits (32 bytes) @13640
Forward Type=1 StreamLength 192 Bits (@13640)
Forward Type=2 Data 8 Bits value=4E (@13642)
Backward Type=1 StreamLength length=64 Bits (@13646)
Backward Type=2 Data length=8 Bits value=00 (@13648)
```

As we can see the forward gap (post-data gap) contains 192 bits (24 bytes) of 4E bytes and the backward gap (the pre-address gap) contains 64 bits (8 bytes) of 00 bytes. We can see that the gap stream exactly matches the gap length:

$$512 = 192 * 2 + 64 * 2$$

The same kind of information is provided for the sectors 2 to 10.

4.2.3 Copylock Protection

The content of sectors 11 and 12 is quite different because these pseudo-sectors are part of the “Copylock short” protection.

The block descriptor for the **Copylock short sgn** (sector 11) is:

```
Encode: 1, 1
Area start: 1
  Area start: 2
    Mark, byte: 3, $4489 *1 <esc>
    ID, byte: 1, $FE *1 <esc>
    Track, byte: 1 *1
    Side, byte: 1 *1
    Sector, byte f0: 1 *1 <esc>
    Length, byte: 1 *1
  _Area stop: 2
  EDC, CRC16: 2, $FFFF *2
  GAP,byte: 28 *1
  Mark, byte: 3, $4489 *1 <esc>
  ID, byte: 1, $FE *1 <esc>
  _Area start: 3
  _Data, byte 8 *1
  _Area stop: 3
  EDC, MFM Encoding: 3, 0 *3
_Area stop:1
```

In the IPF file we find the following information for block descriptor 11:

```
+ BDesc=10 BlkBits=800 GapBits=0 GapOfs=0 CType=1 EType=1 Flg=0 GapVal=0 DataOfs=6315 (13576 - 13608)
```


IPF - Interchangeable Preservation Format Documentation

In the IPF file we find the following information for block descriptor 12

```
+ BDesc=11 BlkBits=2336 GapBits=2272 GapOfs=504 CType=1 EType=1 Flg=2 GapVal=0 DataOfs=6382 (13608 - 13640)
```

If we first look at the data we have:

```
-- DATA 2336 bits (146 bytes) @19638
Sync (Type=1) size=6 Bytes (@19638)
19640 44 89 44 89 44 89 D.D.D.
Data (Type=2) size=7 Bytes (@19646)
19648 fe 00 00 0c 02 bc 33 .....3
Gap (Type=3) size=8 Bytes (@19655)
19657 4e 4e 4e 4e 4e 4e 4e NNNNNNNN
Data (Type=2) size=8 Bytes (@19665)
19667 d9 23 76 c5 e6 d3 31 b2 .#v...1.
Gap (Type=3) size=8 Bytes (@19675)
19677 4e 4e 4e 4e 4e 4e 4e NNNNNNNN
Data (Type=2) size=6 Bytes (@19685)
19687 ff ff ff ff ff fe .....
Sync (Type=1) size=212 Bytes (@19693)
19695 89 12 89 12 89 12 aa 8a 55 55 55 55 55 55 .....UUUUUUUU
19711 55 55 55 55 55 55 55 55 55 55 55 55 55 55 UUUUUUUUUUUUUUUUU
...
19887 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
19903 00 00 00 00
```

We can see at the beginning of the sector the 3 sync bytes (\$A1) displayed as MFM encoded bytes \$4489 (A1 with missing clock bit):

```
Sync (Type=1) size=6 Bytes (@19638)
19640 44 89 44 89 44 89 D.D.D.
```

This is followed by the sector ID field. Here we have the IAM (\$FE) followed by the Track, Side, Sector, Size and the two CRC bytes. Here we see that this is sector 12 (\$0c)

```
Data (Type=2) size=7 Bytes (@19646)
19648 fe 00 00 0c 02 bc 33 .....3
```

This is followed by a short gap: 8 x \$4E bytes

```
Gap (Type=3) size=8 Bytes (@19655)
19657 4e 4e 4e 4e 4e 4e 4e NNNNNNNN
```

Followed by a very short data field of 8 bytes (the Copylock key?)

```
Data (Type=2) size=8 Bytes (@19665)
19667 d9 23 76 c5 e6 d3 31 b2 .#v...1.
```

This is followed by a short gap: 8 x 4E bytes

```
Gap (Type=3) size=8 Bytes (@19675)
19677 4e 4e 4e 4e 4e 4e 4e NNNNNNNN
```

Followed by a very short data field of 6 bytes: 5 x 255 (\$FF) and 1 x 254 (\$FE) as described in the Block descriptor.

```
Data (Type=2) size=6 Bytes (@19685)
19687 ff ff ff ff ff fe .....
```

Followed by an unusual 212 bytes field of Synchron data!

```
Sync (Type=1) size=212 Bytes (@19693)
19695 89 12 89 12 89 12 aa 8a 55 55 55 55 55 55 .....UUUUUUUU
19711 55 55 55 55 55 55 55 55 55 55 55 55 55 55 UUUUUUUUUUUUUUUUU
...
19887 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
19903 00 00 00 00
```

This field starts with 3 x \$8912 sync marks. These \$8912 sync marks correspond to \$4489 sync marks shifted by a half clock. Following the sync marks we have \$AA8A that correspond to the \$FB DAM shifted by a half clock. Following we have many \$5555 marks which are bytes \$00 when decoded. Following we have \$0000 marks. This indicates that there are no clock transitions and therefore correspond to a **No Flux Transition Area** that violates the MFM rules.

IPF – Interchangeable Preservation Format Documentation

The gap data for this last sector contains the following information:

```
-- GAP 2272 bits (142 bytes) @13760
Backward Type=1 StreamLength length=192 Bits (@13760)
Backward Type=2 Data length=8 Bits value=4E (@13762)
Backward Type=1 StreamLength length=64 Bits (@13765)
Backward Type=2 Data length=8 Bits value=00 (@13767)
*** 512 bits out of 2272 = Looping required
```

There is no forward gap description. This indicates no gap following the last chunk of data.

The backward gap description indicates that we first have 64 bits (8 bytes) of 00 bytes in front of the next sector. As sector 12 is the last sector, the next sector is in fact the first sector of the track. These 8 bytes of \$00 are preceded by 192 bits (24 bytes) of \$4E bytes.

The overall length of the gap is 2272 bit and we have only 512 bits specified. Therefore the decoder needs to loop the samples (4E) at the beginning and at the end of the gap until they meet. The point where they meet is the track **write splice**.

5. CAPS/IPF Programing

This section is dedicated to programmers that develop application around the IPF format and Library. More information can be found in the [IPF Project](#) page of my Web

5.1 IPF Decoder Library

The source of the CAPS/IPF library has now been [publicly released under license](#) by SPS. This license is based on the MAME license and is intended for use in all non-commercial projects and environments. Other licensing options are available. Don't hesitate – to contact SPS at licensing@kryoflux.com.

I have created several modified version of the original SPS library sources.

5.1.1 IPF Library Release 4.2A

The original 4.2 source tree released by SPS was compiled using Microsoft Visual C++ 6.0 and does not compile successfully with recent version of Visual C++. Visual C++ 6.0 is the worst compiler ever released by Microsoft: it has a lot of bugs, it is not fully C++ compliant, and furthermore it is no more available!

Therefore I have made a [release 4.2A of the library](#) that includes minimum modifications to compile properly under Microsoft Visual Studio 2010 / 2012.

What has changed?

- StdAfx.h: In order to avoid many warnings you need to define `_WIN32_WINNT`, but most important you must use the `stdint.h` from MS library and **not** the one provided in the SPS 4.2 tree. For that matter you need to replace the `#include "stdint.h"` statement by `#include <stdint.h>`.
- One area where Visual C++ 6.0 is not C++ compliant is about variables declared in a “for loop”. In standard C++ a variable declared in a “for loop” is only known inside the loop and goes out of scope outside the loop. So we have

```
int a;
for (int foo; foo < 3; foo++) {
    a += foo;
}
foo = a;           // works with VC6 - error with C++ compliant comp.
int foo = 2;      // works with C++ compliant comp. error with VC6
```

- In the CAPS lib some code use variable declared in “for loop” outside of the loop in the following modules: CapsAPI.cpp, CapsEFDC.cpp, CapsmgS.cpp, CapsLdr.cpp, and DiskImg.cpp. The code in these files has been changed so it is C++ compliant and compile under VC2010/2012

- In CapsImgS.cpp there is a wrong cast to the 1st parameter of the Swap function in line 199

As you can see only few lines of the source code need to be modified. I have released the modified source tree as **IPF-42A-120116.rar** file.

5.2 IPF Dump Program

This is a small program that I wrote to dump / check the content of an IPF file. This can be useful to understand the content as described in [IPF file content](#) chapter of this document. This program is based on an original code written in 2011 by Keir Fraser that I have extensively modified to make it more generic and to provide more information.

To run the program type:

```
IPFInfo [-i] [-d] [-n] input_ipf_file [output_file]
```

The `-i` and `-d` options can be used to add the dump of respectively the INFO and DATA records information in the output file. The `-n` option is used to filter the dump of unformatted track (noise track) in the output file.

The `input_ipf_file` is a `.ipf` or `.ipx` file as produced by CAPS Analyzer.

The `output_file` is optional and by default the output goes to the console.

6. Glossary

This section contains few useful definitions.

Bit Cell

This is essentially the distance along the track allocated for the recording of an area of Flux. That is, in each “cell” there are magnetic particles that together indicate a detectable magnetic polarity of set length. Note that bit cells are not read as bits of data directly, they are used to form the flux transitions used to make up the Encoding. Every bit cell contains either a change of polarity (1) or not (0). Writing through the Atari FDC hardware is limited to bytes, and bit writes are clocked by the hardware clock to produce 4us wide bit “cells” for the MFM (Modified Frequency Modulation) Encoding format. For this reason, modifying the width of the bit cells written using Atari hardware is limited. This is one of the reasons why Long Tracks have to be copied using a hardware copier. Most games were mastered more professionally, with programmable duplicator systems - like from Trace. These were capable of writing any cell width desired.

Bit Shifting

Also called peak shift. An effect that happens when reading from magnetic media. Essentially, boundary bits may be interpreted as part of adjacent data values. This means the value of the original element, and the value of the one next to it may change. Another way of putting it, would be to say the timing detected (and hence the data read) for a Flux Transition can change, caused by the influence of a neighboring flux transition. The IPF (Interchangeable Preservation Format) decoder (library) generates bit shifting into the decoded stream automatically at the most likely position it would occur if the data was being written to a real disk for absolutely faithful reproduction of the real media properties. This helps emulation of certain protections, without any kludges or any extra information required to distinguish between data meant for mastering and data meant for actual use (emulation).

Block (aka Sector)

Each Track is made up of one or more blocks. The PC/ Atari ST format usually uses 9 blocks but “generic hardware MFM” formats use whatever number of blocks and sizes as programmed into the controller. Blocks are used to reduce the amount of data stored and encoded in one go, thus modifying just a small amount of data does not require a whole track to be modified, and errors on a block does not render the rest of the data unreadable. Blocks are normally separated by gaps.

Copy Protection

Copy protection was used on media for various systems to protect the investments of the people involved in creating the software content. Various techniques were used to this end. Some of the common copy protection techniques are:

- ◆ Fuzzy Bits: random data is read, which becomes non-random when copied.
- ◆ Long/Short tracks: the data is packed on a track, more and less dense respectively. This needs to be checked by the program code.
- ◆ Density timing. Various patterns of variable density are used across a track and checked to exist. The physical properties of magnetic media can be used in many ways to hinder copying of the data stored on it. For magnetic media, it is difficult to read something without knowing how it is structured, and it is impossible to reliably write something without knowing how to write it. This “how” varies between publishers and game developers, but they used a wide variety of techniques. The whole emphasis in Copy Protection is to write the data in such a way that it can be read, but cannot be written using consumer products, and hence cannot normally be reproduced.

Custom Disk Format

Floppy disk based systems organize data on the media in different ways. Some examples of possible variations are:

- ◆ Physical Encoding
- ◆ Number of blocks
- ◆ Data/value representation

IPF – Interchangeable Preservation Format Documentation

- ◆ Header information (current Track number, head number, data size, etc.)
- ◆ Location and method of integrity information
- ◆ Location and value of the Sync (aka Mark)
- ◆ Location, value and size of Gap
- ◆ Location and size of actual data part

This is simply just layout (aka geometry) information. Computer designers would either create their own format, or they would use an existing one depending on their system's hardware. If they used a fully-fledged FDC (Floppy Disk Controller), the format had to conform to predefined rules of geometry as allowed by that hardware. One reason to use custom format was to prevent the disk from being copied in a normal way.

Density

How closely bit cells are packed together on a section of track. This is effectively determined by the bit cell width. Higher density (using bitcells of shorter width) means more data can be packed into the same length of track. Commercial duplicator machines were able to write bit cells of different lengths. For protections varying density was used across a track (e.g. Copylock, Speedlock. The Atari can read disks of varying density with a limit of roughly +/- 10% of normal density.

Encoding

Data encoding is the process of converting binary information (programs and data) into magnetic impulses that can be stored on the magnetic surface of the disk. There are several different ways that this can be done. Floppy disks generally use the MFM (Modified Frequency Modulation) encoding technique, which is also used on old hard disks.

Due to the nature of magnetic recording, you cannot just store a meaningful sequence of bits in the same number of bit cells, since you would not be able to read it. Instead, a bit is stored by the presence, or lack of, a Flux Transition. That is, the change of polarity between two bitcells. This is further complicated because you cannot (reliably) read this data unless you use some form of clocking so you know when a flux transition has occurred (a binary 1) or could have occurred but did not (a binary 0).

FDC Emulator

An FDC emulator job is to implement a definitive, low level, and most importantly accurate emulation for various brands of generic hardware floppy controllers. Almost of these "off the shelf" controllers are derivatives of the NEC765A FDC. The FDC emulator acts as a "plugin", or "library" for emulators to allow inclusion of true floppy system emulation without needing to re-implement this complicated piece of functionality themselves.

Gap

This is a data area that can be found either side of a Block (aka Sector) or Track, though it does not indicate the separation like a Sync (aka Mark). The size of a gap area may be altered via mastering if allowed by the format. Any errors found after the gap can normally be ignored, though a few protections have been found that store data in this area.

Gaps are normally used to allow re-writing of specific blocks on a track. As this operation is not timing accurate, the gap provides the threshold needed to cover for the inaccuracies that occur in the timing of the write.

There is at least one gap between the last block on one track, and the first block of the next. Its purpose is to compensate for the different properties of a track written on specific locations on the disk as well differences between drive speed alignment. If the drive speed is different within reason, this gap will have some of its content overlapping i.e., overwritten when mastering a whole track.

Index Pulse

Indicates the hardwired beginning of a Track (actually the full rotation cycle of the disk), that can be used as a fixed reference point by software or hardware.

Jitter:

This is the difference between the expected time of arrival of some data, and the real time of arrival of that data. You would be tempted to think of this as tolerance, but it isn't. Tolerance means to

IPF – Interchangeable Preservation Format Documentation

compensate for jitter, or “de-jitter”. Most FDC's incorporate some kind of jitter hardware so that they are tolerant within some boundaries if the Bit Cell widths are not quite the ones expected within a given amount of time, i.e. the clock speed selected for cell width. This is why it is possible for the FDC to read bit cells that are not exactly the same width as expected - and indeed they are not usually the width expected, due to the nature of magnetic recording.

This fact was well used by Copy Protection designers who made the bit cells different widths on purpose, something that could only be reproduced by special mastering hardware, and due to the tolerance built into the FDC's, the data could be read properly, it was just impossible to write by the user. This hardware is built into most FDC's, and what you normally read using the FDC is not actually the same as the data as it is present on the disk surface. This is a very important fact to bear in mind when preserving magnetic media! *For preservation, you need to insure that you get 1:1 what is on the disk surface, NOT what the FDC interprets it has.*

Long Track

One of the simplest and most effective Floppy Disk Copy Protection schemes. This is where the disk mastering hardware reduces the size of the bit cells, which allows more data to be written on any one track. This is something which cannot be done on an Atari platform. An alternate way of producing a long track is to slow down the disk drive which in practice it has the same effect - the bit cells take less space, so there can be more cells written to the same area. Writing short tracks uses exactly the same techniques, but obviously using a wider bit cell width rather than shorter.

MFM (Modified Frequency Modulation)

Modified Frequency Modulation, also known as Delay Modulation or Miller Code is a self-clocking Encoding used on floppy disks of many systems including the Commodore Amiga, Atari ST and PC. It was an enhancement over FM encoding, storing roughly double the amount of “meaningful” data using the same amount of bits due its more efficient clocking, which worked by considering not just the current bit, but also the previous one. This is why disks are marked as DD or 2D - “Double Density” since they are MFM not FM.

Sync (aka Mark)

Special marker used to separate different areas of data. The sync is (in theory) not produced by the Encoding used on a data area. Therefore it can mark the beginning (sometimes not really the beginning) of a new physical/logical data area, known as a Block (aka Sector). The PC and the Atari format both use MFM “0×4489” as the sync pattern. In the analyzer it is referred to as the “Mark” to be more in line with Trace conventions. It is treated as data, but the encoding can be modified, but it is used by the analyzer to spot areas on a track in the first place.

Track:

A “ring” around one side of a disk containing data. Tracks begin and end at the Index Pulse.

7. References

- [CAPS Library and API Documentation Revision 1.02](#)
- SPS WIP: [Time for an even more generic way to describe disk formats...](#)
- SPS WIP: [Analyzer: Generic MFM Support - Gap Analyzer](#)
- SPS WIP: [Analyzer: Generic MFM Support - Gap](#)
- [CAPS Source Code Release 4.2](#)
- [All you always wanted to know about IPF](#)
- [Amiga Disk Utilities & IPF creation and disk writing via Kryoflux](#) by Keir Fraser
- [Structure of descriptors and data for a track](#) by Keir Fraser
- [Atari Protection based on Key Disk](#) by Jean Louis-Guerin
- [KFAnalyze Program](#) and [documentation](#)
- Microsoft Walkthrough: [Creating and Using a Dynamic Link Library](#)
- Microsoft Walkthrough: [Creating and Using a Static Library \(C++\)](#)

8. History

- V1.2 September 2013: Fixed errors, added more references to SPS and Kryoflux forum.
- V1.1 June 2013: Added information on IPX format and CTEI/CTEX blocks. Fixed errors.
- V1.0 June 2013: First official release. Fix information based on feedback. Added the Copylock protection description in [Analysis of Theme Park Mystery](#). TODO: Some record type are not described (DUMP, TRCK, CTEX, CTEI) but they do not seems to be used in standard IPF files. The CAPS/IPF Programming section needs to be completed / reworked.
- V0.0 January 2012: Working review version for feedback