The FITS Test Tape
Version 4

Donald C. Wells
National Radio Astronomy Observatory
Edgemont Road
Charlottesville, VA 22901
(804) 296-0211, FTS 538-1271

Document printed: September 22, 1992

Abstract

The contents of Version 4 of the FITS Test Tape are described. Various implementation details are discussed, as well as the differences of style which are seen in the files on the tape.

Contents

Chapter 1

Introduction

It is difficult to describe a communication protocol like FITS in a manner which is unambiguous for all people who read the written version of the specifications. The same words can mean significantly different things to different people. Indeed, a protocol standard document must be made to read almost like a legal contract in order to try to reduce ambiguities to a minimum. FITS must be implemented in real computer systems, and the vendors of computer systems differ in their interpretation of the meaning of various words used to describe computer hardware and software. It is easy to imagine that differences in interpretation of the FITS documents or of implementation of computer systems by vendors might make various FITS implementations differ in subtle ways, and users of astronomical data systems might be frustrated occasionally by incompatibilities as they carry their data from system to system. The purpose of the Test Tape is to provide a sample set of files which are believed to be in conformance with the FITS specifications and which can be used to 'certify' that FITS reading programs function correctly. These samples of the FITS format are intended to eliminate any residual uncertainties which may exist about the precise meaning of the specifications.

The Test Tape itself must be certified. Its certification is the responsibility of the whole community of FITS users. Errors or problems must be reported if the test tape scheme is to work. The author promises to accept such reports graciously. The author also invites suggestions for improvements in the content, style, etc., of the Test Tape. It is especially important that the rigor of the tests made by the tape should be affirmed by many sites.

The present tape includes files which were written by six different systems at four different observatories. Such diversity tends to insure against error. The author will welcome tapes of samples of the output style of additional systems for inclusion in future versions of the tape.

The astronomical data processing community can be thankful that there have been almost no reports of incompatibility problems with FITS, but it would be foolish to be complacent. The best insurance against incompatibilities is provided by systematic trials of data exchange which are designed to assure that a certain minimum set of the features of the FITS design are implemented consistently at all sites and in all systems. The Test Tape is the most important component of this process. If a site has a program which can read the FITS Test Tape correctly it can use this program to certify the correctness of any FITS writing program which it has, although ultimately this assurance is not as convincing as the knowledge that one or more other systems with independent implementations read the output files correctly. The implementors of FITS reading and writing programs should always be on guard against incompatibilities. Whenever errors are discovered in implementations at other sites they should be privately reported to those sites in a polite manner so that they may be eliminated. If this is not done the errors will propagate farther and farther, and the whole community will suffer.

The FITS specifications must not be regarded as being absolutely complete or frozen. It is particularly important that any ambiguities which may exist in the original specifications of FITS should be discussed publicly. Common errors of interpretation must be discussed, and differences of style which are permissable within the specifications of the design must be examined to determine whether they imply that the specifications should be modified or extended. Extensions of the specifications to cover applications which were not considered by the original designers are certain to be needed in the long run. The IAU Circular provides a forum for all such discussions. They are also an appropriate subject for the AAS Working Group for Astronomical Software, and for the analogous European group.

Chapter 2

Contents of the Tape

A summary of the properties of the 12 files in version 4 of the tape is given below:

File	Name	BITPIX	NAXIS	#pixels	#blocks	#header	#special
1	Text File	(8)	0		22	22	
2	Mandrill	8	3	512*512*3	276	2	
3	M42	16	2	512*512	184	1	
4	NGC4486	32	2	256*256	94	2	
5	Random-Groups	16	6	0*2*4*1*1*1	9	3	6
6	AAT-IPCS	16	2	240*120	21	1	
7	Orion-Xray	16	2	128*128	13	1	
8	M84	16	4	513*513*1*1	354	4	167
9	NGC40,H-beta	16	2	340*340	82	1	
10	NGC40,ratio	16	2	340*340	82	1	
11	Magnetogram	16	2	512*600	215	1	
12	Venus	8	2	600*800	168	1	

The 12 files were chosen to include a wide range of the combinations permitted within the FITS rules. Certain critical special cases of the rules are represented so as to assure that general-purpose FITS software will not misbehave upon encountering special types of files which will be used for FITS extensions. There is no need for all systems to support the random-group format (file 5), but it is important that they should all print reasonable error messages and skip to the next file on the input tape. Also, a FITS reading program should test for the case of NAXIS=0 (e.g., file 1). Special records behind a data matrix (file 8) should be skipped, perhaps with warning messages. If all FITS software performs gracefully with these files we can be sure that a wide variety of data may be recorded using FITS without producing software crashes.

The files may be considered in three groups. The 'Critical Tests' group (files 1-5) is intended to assure the conformance of a FITS reading program to the basic rules of FITS. The 'Styles' group (files 6-8) illustrates different styles and levels of sophistication in the software at three different sites. The 'interesting pictures' group (files 9-12) is intended to give the impression that a wide variety of astronomical imagery can be expressed within the FITS rules. In the following sections of this document we comment upon the details of these groups and of the individual images within them. It will be helpful to have a listing of the FITS header blocks available when reading the discussions of the files.

Chapter 3

Critical Tests

Programmers should assure that all FITS reading programs behave in a graceful manner when reading these five files. It is intended that the files in this group should assure conformance to all of the basic rules of FITS. This does not mean that all programs should fully recognize and use all of the optional keywords and axis types shown in file 5, but rather that they should not crash when they encounter unrecognized keywords, or headers which are very long. Note that although there are no examples of 1-D matrices (e.g., spectra) provided in this group, programmers should assure that they handle such cases gracefully. In particular, it is likely that FITS files of spectra or time series may be created with NAXIS1 very large, maybe even greater than 32767, the INTEGER*2 limit.

3.1 File 1, The Text File

This file illustrates how FITS may be used to transmit a variety of text data. The header simply gives a nominal BITPIX of 8 and then announces that NAXIS is zero (i.e., there is no binary data matrix). It gives the ORIGIN keyword to indicate who wrote the file. And then we see a non-standard keyword: TEXTFILE. The FITS rule about non-standard keywords is: if you don't recognize a keyword just regard it as a comment, and so the use of this keyword should not be regarded as a violation of the rules of the standard. It is best to create new keywords cautiously, and only for good reason. In this case we are trying an idea for documenting text files, and we are using this keyword to separate files of text in the FITS header, and to transmit their VAX/VMS file names. Examples of three kinds of text files are given here: the introductory section of the RUNOFF file of the text of this document (FITS4.RNO), several FORTRAN source files (e.g., PACKFITS.FOR), and a VMS command language procedure file (LISTFITS.COM). The file names are given using FORTRAN-77 list-directed-read notation for strings. The names are longer than 8 characters, and this is a mild violation of the basic standard which wants string values to be able to be used by interpreting only their first 8 characters. But the keyword is non-standard anyway, and so we won't worry very much about this detail.

The programs given in this file were actually used in the preparation of this FITS test tape. PACKFITS and READFITS were built to prepare this file of text data. COPYFITS was used to copy files selectively during the construction of the test tape, and is used to make copies of the test tape. COMPFITS is used during the test tape duplication process to verify the copy by read-back and comparison. And LISTFITS produces the nice listing of the headers and record counts for the files. These programs should run immediately on other VAX/VMS systems, and can probably be adapted to run on other computers with FORTRAN-77 compilers.

3.2 File 2, The Mandrill

This is a famous image in the history of image processing. Its bits have been mangled through many computers and many experimental algorithms for many years. It is also a dramatic image, especially in this full color version. The file tests the 8-bit data type of FITS and it also tests the 3-dimensional capability of FITS.

There is an error in the header which, fortunately, will usually be harmless. This is that the header declares the BLANK keyword, implying that there will be blank pixels in the image, when, in fact, there are none. The original FITS documents are ambiguous on this technical point, but probably it is good practice to specify BLANK only when there really are blanks. In this case the program which generated the file did not have any convenient way to know whether or not there were blanks in the file, and so it was designed to always specify BLANK 'just in case'. A proper FITS reading program will take the BLANK code value of zero seriously, and will set all of the zero pixels of this image to the internal representation for blanks. This is wrong, and eventually the author intends to rewrite this file without the BLANK keyword.

The program used several non-standard keywords to document some relevant facts about the image in the host system, the KPNO IPPS. In retrospect it might be argued that these strings and numeric values should have been encoded into several COMMENT or HISTORY cards rather than appearing as value fields for non-standard keywords, although the rules of FITS permit the invention of any non-reserved keywords which the programmer wants to use. Note that a new axis type code was invented to describe the third axis (RGB=Red-Green-Blue). FITS allows the invention of new axis types, but programmers should be cautious. Existing axis types and units should be used wherever possible so as to minimize interchange difficulties. In this case the axis type covers a common type of data which was not considered by the original designers. Technical points of this general type are mostly matters of taste, and usually will not affect the basic data interchange capability of FITS.

The subject of string notation is raised by some of the details of the header of this file. Consider the ORIGIN card (1/17). The string inside the single quotes is 29 characters long. The basic rule is that the data value fields are to be readable using Fortran-77 list-directed read operations. The FITS specifications also state that the string should be positioned beginning in column 12 and should be interpretable in its first 8 characters. This string conforms to both of these rules (effectively, ORIGIN='KPNO'). The notion of encoding information about the FITS-writing program (named WFITS in this case) beginning in the ninth character of the string is used in other files on the tape (e.g. file 5, line 1/26). Another interesting case is seen in the DATE card (1/18), where three separate strings appear in the value field. Only the first one is really the date, and it obeys FITS rules. In retrospect the additional strings, while strictly within the formal rules of FITS, are a luxury, and it would be better not to run the risk of causing trouble in some FITS reading programs.

Many existing image processing systems are not capable of representing 3-D images in their data structures and formats. Such systems (e.g., the KPNO IPPS, which wrote this file) can produce three separate images from this file.

3.3 File 3, The Orion Nebula (M42)

This file tests the 16-bit format, including blanks. This is an example of the most common pixel format of astronomical imagery, and programmers should assure that their FITS reading programs handle this file properly.

3.4 File 4, The Nucleus and Jet of M87 (NGC4486)

This file tests the 32-bit data format, with blanks. It will not be sufficient to interpret only the low-order 16 bits of the 32-bit pixels. Nor will it be sufficient to interpret only the high-order 16 bits. The pixel values have been deliberately scaled so as to require that the reading program process the 'middle' bits of the numbers in order to get a proper image.

3.5 File 5, The Random-Group Case

This file tests another one of the outer boundaries of the FITS design, namely 'special records'. It has a very complicated specification of the axes for a 6-dimensional matrix, but all of this complexity is nullified for an ordinary FITS reading program by the keyword NAXIS1 (card 1/04) with value zero. This means that there is no ordinary matrix in this file. So the ordinary program will be looking for the tapemark after it sees the END card in the third header record

(3/31). But it will actually see 6 more blocks before the tapemark. One of the fundamental rules of FITS is that the reading program must be prepared to skip any records occurring after the data matrix and before the tapemark. This file is merely a special case of the general rule, special in that there is no data matrix. The data in the special records are the samples of complex visibility measured by the VLA during an observation of a radio source. These can be converted to a conventional image by means of a Fourier Transform operation. This file is an example of the random-group extension of FITS which is described in the second FITS paper (Greisen, E. W., and Harten, R. H., Astron. Astrophys. Suppl. Ser. 44, 371-374, June 1981) .pg

Chapter 4

Styles

4.1 File 6, A Photon-Counting TV Image from the AAT

This image is mainly provided to illustrate alternative styles in the use of the capabilities of FITS. The image was brought from the AAT in Australia to KPNO by Harvey Butcher of KPNO. The FITS reading program on the Cyber computer at KPNO read this image on the first try. Easy data interchange over intercontinental distances between programmers who cannot communicate is what FITS is all about.

The non-standard keywords in the header of this file (RA, DEC, etc.) are presumably taken from the AAT telescope readouts, and are not intended to refer to any particular pixel of the image. Note that cards 1/20 through 1/22 appear to be comment text typed in on a terminal during the observation. This is an interesting use of the comments capability of FITS. In line 1/07 a string value is given ('XM') with the second quote mark as the third character of the string, rather than the ninth, or farther to the right. This is a violation of the recommended practices for FITS in which all strings should be at least eight characters long. In fact, most existing FITS reading programs are more general and are able to read this string correctly. Anyway, this is a non-standard keyword and will be ignored by most FITS reading programs.

4.2 File 7, An X-Ray Image of Orion

This file was produced from Einstein satellite data, and was supplied by Claire Ambler of Columbia Astrophysics Laboratory. Note that there are no slashes to terminate the value fields. The slashes are recommended, but not really required, and so this file is a test of yet another boundary of the FITS standard.

4.3 File 8, A Self-calibrated VLA Image of M84

This file illustrates FITS in all of its glory. It shows the style of the NRAO AIPS system, whose internal data base mimics FITS very closely. Most of the optional keywords of the original FITS specifications appear somewhere in this header. It is a 4-D matrix (in other situations the third and fourth axes would have dimensions greater than one). It has extensive documentation of the processing to which the data were subjected. Another interesting fact is that there are 167 'special records' following the data matrix. These special records are being used to transmit a table of the coordinates and fluxes of 12000 CLEAN components which were subtracted from the original image in the CLEAN process. The format of these special records is still slightly provisional, but anyone who is curious about it should examine the first two special records as 80-character ASCII card images (just like a FITS header). This image, which was provided by Robert Laing of NRAO, is an example of the quality of aperture synthesis imagery produced by recent 'self-calibration' algorithms.

Chapter 5

Interesting Pictures

The files discussed in this section are intended to show the ability of FITS to represent a wide variety of astronomical data. At the present time this section contains only digital images. This is not really a limitation of FITS. Future versions of the Test Tape can and should contain examples of time series or spectroscopic data (1-D matrices).

5.1 Files 9 and 10, Processed Images of NGC 40

These two files, of planetary nebula NGC 40, illustrate the transmission of heavily processed data. The BSCALE values (0.1 and 0.001) show that the original images were in floating point. The second file is a line-ratio map, the logarithm of an H-Alpha image divided by an H-Beta image. These images were provided by Larry Goad of KPNO.

5.2 File 11, A Solar Magnetogram

This image shows the longitudinal magnetic field strength at a certain level in the atmosphere of the Sun with a spatial resolution of about one arcsecond. It was provided by Jack Harvey of KPNO. Note that this image has both positive and negative pixel values.

5.3 File 12, Ultraviolet Image of Venus

The final image of the group is an image of Venus which was generated and processed at JPL. It has been geometrically transformed to a Mercator projection. It was copied from a VICAR-format tape into the KPNO IPPS and was later written back out in FITS format. The data were provided by Mike Belton of KPNO. Note that 72-character VICAR header card images can be copied onto FITS HISTORY cards, although they were not in this case.