



Open Source Used In Cisco Unified Attendant Console Standard (14.0.1) - HONDA 14.0.1

Cisco Systems, Inc.

www.cisco.com

Cisco has more than 200 offices worldwide.
Addresses, phone numbers, and fax numbers
are listed on the Cisco website at
www.cisco.com/go/offices.

Text Part Number: 78EE117C99-1135983409

This document contains licenses and notices for open source software used in this product. With respect to the free/open source software listed in this document, if you have any questions or wish to receive a copy of any source code to which you may be entitled under the applicable free/open source license(s) (such as the GNU Lesser/General Public License), please contact us at external-opensource-requests@cisco.com.

In your requests please include the following reference number 78EE117C99-1135983409

Contents

1.1 dotnetzip-reduced 1.9.1.8

1.1.1 Available under license

1.2 zlib 1.2.3

1.2.1 Available under license

1.3 commonservicelocator 1.0.0.0

1.3.1 Available under license

1.4 libjpeg 6a

1.4.1 Available under license

1.5 openssl 1.0.2s

1.5.1 Notifications

1.5.2 Available under license

1.6 dotnetzip 1.9.1.8

1.6.1 Available under license

1.7 ionic-zip 1.9.1.8

1.7.1 Available under license

1.8 xerces-c 3.2.3

1.8.1 Available under license

1.9 log4net 2.0.8.0-.NET

1.9.1 Available under license

1.10 zlib 1.1.4

1.10.1 Available under license

1.11 sqlite 3.7.16.2

1.11.1 Available under license

1.1 dotnetzip-reduced 1.9.1.8

1.1.1 Available under license :

The ZLIB library, available as Ionic.Zlib.dll or as part of DotNetZip, is a ported-then-modified version of jzlib, which itself is based on zlib-1.1.3, the well-known C-language compression library.

The following notice applies to zlib:

Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler

The ZLIB software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

The following licenses govern use of the accompanying software, the DotNetZip library ("the software"). If you use the software, you accept these licenses. If you do not accept the license, do not use the software.

The managed ZLIB code included in Ionic.Zlib.dll and Ionic.Zip.dll is modified code, based on jzlib.

The following notice applies to jzlib:

Copyright (c) 2000,2001,2002,2003 ymnk, JCraft,Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

jzlib is based on zlib-1.1.3.

The following notice applies to zlib:

Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler

The ZLIB software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org
Mark Adler madler@alumni.caltech.edu

Apache Commons Compress
Copyright 2002-2010 The Apache Software Foundation

This product includes software developed by
The Apache Software Foundation (<http://www.apache.org/>).
The ZLIB library, available as Ionic.Zlib.dll or as part of DotNetZip,
is a ported-then-modified version of jzlib. The following applies to jzlib:

JZlib 0.0.* were released under the GNU LGPL license. Later, we have switched
over to a BSD-style license.

Copyright (c) 2000,2001,2002,2003 ymnk, JCraft,Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in
the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products
derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES,
INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND
FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT,
INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA,
OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE,
EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Microsoft Public License (Ms-PL)

This license governs use of the accompanying software, the DotNetZip library ("the software"). If you use the software, you accept this license. If you do not accept the license, do not use the software.

1. Definitions

The terms "reproduce," "reproduction," "derivative works," and "distribution" have the same meaning here as under U.S. copyright law.

A "contribution" is the original software, or any additions or changes to the software.

A "contributor" is any person that distributes its contribution under this license.

"Licensed patents" are a contributor's patent claims that read directly on its contribution.

2. Grant of Rights

(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.

(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.

3. Conditions and Limitations

(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.

(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.

(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.

(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in compiled or object code form, you may only do so under a license that complies with this license.

(E) The software is licensed "as-is." You bear the risk of using it. The contributors give no express warranties, guarantees or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-infringement.

The managed BZIP2 code included in Ionic.BZip2.dll and Ionic.Zip.dll is modified code, based on the bzip2 code in the Apache commons compress

library.

The original BZip2 was created by Julian Seward, and is licensed under the BSD license.

The following license applies to the Apache code:

```
-----  
/*  
* Licensed to the Apache Software Foundation (ASF) under one  
* or more contributor license agreements. See the NOTICE file  
* distributed with this work for additional information  
* regarding copyright ownership. The ASF licenses this file  
* to you under the Apache License, Version 2.0 (the  
* "License"); you may not use this file except in compliance  
* with the License. You may obtain a copy of the License at  
*  
* http://www.apache.org/licenses/LICENSE-2.0  
*  
* Unless required by applicable law or agreed to in writing,  
* software distributed under the License is distributed on an  
* "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY  
* KIND, either express or implied. See the License for the  
* specific language governing permissions and limitations  
* under the License.  
*/
```

1.2 zlib 1.2.3

1.2.1 Available under license :

```
/* zlib.h -- interface of the 'zlib' general purpose compression library  
version 1.1.4, March 11th, 2002
```

Copyright (C) 1995-2002 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler
jloup@gzip.org madler@alumni.caltech.edu

The data format used by the zlib library is described by RFCs (Request for Comments) 1950 to 1952 in the files ftp://ds.internic.net/rfc/rfc1950.txt (zlib format), rfc1951.txt (deflate format) and rfc1952.txt (gzip format).

*/

/* match.s -- Pentium-optimized version of longest_match()

* Written for zlib 1.1.2

* Copyright (C) 1998 Brian Raiter <breadbox@muppetlabs.com>

*

* This is free software; you can redistribute it and/or modify it

* under the terms of the GNU General Public License.

*/

1.3 commonservicelocator 1.0.0.0

1.3.1 Available under license :

```
{"Name":"Microsoft Public License (Ms-PL)","Text":"Microsoft Public License (Ms-PL)\r\n\r\nThis license governs use of the accompanying software. If you use the software, you accept this license. If you do not accept the license, do not use the software.\r\n\r\n1. Definitions\r\n\r\nThe terms \"reproduce,\" \"reproduction,\" \"derivative works,\" and \"distribution\" have the same meaning here as under U.S. copyright law.\r\n\r\nA \"contribution\" is the original software, or any additions or changes to the software.\r\n\r\nA \"contributor\" is any person that distributes its contribution under this license.\r\n\r\n\"Licensed patents\" are a contributor's patent claims that read directly on its contribution.\r\n\r\n2. Grant of Rights\r\n\r\n(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.\r\n\r\n(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.\r\n\r\n3. Conditions and Limitations\r\n\r\n(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.\r\n\r\n(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.\r\n\r\n(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.\r\n\r\n(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in compiled or object code form, you may only do so under a license that complies with this license.\r\n\r\n(E) The software is licensed \"as-is.\" You bear the risk of using it. The contributors give no express warranties, guarantees or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-
```

1.4 libjpeg 6a

1.4.1 Available under license :

No license file was found, but licenses were detected in source scan.

```
/*
 * jcinit.c
 *
 * Copyright (C) 1991-1997, Thomas G. Lane.
 * Modified 2003-2017 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains initialization logic for the JPEG compressor.
 * This routine is in charge of selecting the modules to be executed and
 * making an initialization call to each one.
 *
 * Logically, this code belongs in jcmaster.c. It's split out because
 * linking this routine implies linking the entire compression library.
 * For a transcoding-only application, we want to be able to use jcmaster.c
 * without linking in the whole library.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcinit.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * transupp.c
 *
 * Copyright (C) 1997-2017, Thomas G. Lane, Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains image transformation routines and other utility code
 * used by the jpegtran sample application. These are NOT part of the core
 * JPEG library. But we keep these routines separate from jpegtran.c to
 * ease the task of maintaining jpegtran-like programs that have other user
 * interfaces.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/transupp.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * jmorecfg.h
 *
 * Copyright (C) 1991-1997, Thomas G. Lane.
 * Modified 1997-2013 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains additional configuration options that customize the
 * JPEG software for special applications or support machine-dependent
 * optimizations. Most users will not need to touch this file.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmorecfg.h
```

No license file was found, but licenses were detected in source scan.

```
/*
 * jfdctint.c
 *
 * Copyright (C) 1991-1996, Thomas G. Lane.
 * Modification developed 2003-2015 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains a slow-but-accurate integer implementation of the
 * forward DCT (Discrete Cosine Transform).
 *
 * A 2-D DCT can be done by 1-D DCT on each row followed by 1-D DCT
 * on each column. Direct algorithms are also available, but they are
 * much more complex and seem not to be any faster when reduced to code.
 *
 * This implementation is based on an algorithm described in
 * C. Loeffler, A. Ligtenberg and G. Moschytz, "Practical Fast 1-D DCT
 * Algorithms with 11 Multiplications", Proc. Int'l. Conf. on Acoustics,
 * Speech, and Signal Processing 1989 (ICASSP '89), pp. 988-991.
 * The primary algorithm described there uses 11 multiplies and 29 adds.
 * We use their alternate method with 12 multiplies and 32 adds.
 * The advantage of this method is that no data path contains more than one
 * multiplication; this allows a very simple and accurate implementation in
 * scaled fixed-point arithmetic, with a minimal number of shifts.
 *
 * We also provide FDCT routines with various input sample block sizes for
 * direct resolution reduction or enlargement and for direct resolving the
 * common 2x1 and 1x2 subsampling cases without additional resampling: NxN
 * (N=1...16), 2NxN, and Nx2N (N=1...8) pixels for one 8x8 output DCT block.
 *
 * For N<8 we fill the remaining block coefficients with zero.
```

* For N>8 we apply a partial N-point FDCT on the input samples, computing
* just the lower 8 frequency coefficients and discarding the rest.
*
* We must scale the output coefficients of the N-point FDCT appropriately
* to the standard 8-point FDCT level by 8/N per 1-D pass. This scaling
* is folded into the constant multipliers (pass 2) and/or final/initial
* shifting.
*
* CAUTION: We rely on the FIX() macro except for the N=1,2,4,8 cases
* since there would be too many additional constants to pre-calculate.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jfdctint.c

No license file was found, but licenses were detected in source scan.

/*

* jddctmgr.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2002-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains the inverse-DCT management logic.

* This code selects a particular IDCT implementation to be used,

* and it performs related housekeeping chores. No code in this file

* is executed per IDCT step, only during output pass setup.

*

* Note that the IDCT routines are responsible for performing coefficient

* dequantization as well as the IDCT proper. This module sets up the

* dequantization multiplier table needed by the IDCT routine.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jddctmgr.c

No license file was found, but licenses were detected in source scan.

/*

* jidctfst.c

*

* Copyright (C) 1994-1998, Thomas G. Lane.

* Modified 2015-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains a fast, not so accurate integer implementation of the

* inverse DCT (Discrete Cosine Transform). In the IJG code, this routine

* must also perform dequantization of the input coefficients.
*
* A 2-D IDCT can be done by 1-D IDCT on each column followed by 1-D IDCT
* on each row (or vice versa, but it's more convenient to emit a row at
* a time). Direct algorithms are also available, but they are much more
* complex and seem not to be any faster when reduced to code.
*
* This implementation is based on Arai, Agui, and Nakajima's algorithm for
* scaled DCT. Their original paper (Trans. IEICE E-71(11):1095) is in
* Japanese, but the algorithm is described in the Pennebaker & Mitchell
* JPEG textbook (see REFERENCES section in file README). The following code
* is based directly on figure 4-8 in P&M.
* While an 8-point DCT cannot be done in less than 11 multiplies, it is
* possible to arrange the computation so that many of the multiplies are
* simple scalings of the final outputs. These multiplies can then be
* folded into the multiplications or divisions by the JPEG quantization
* table entries. The AA&N method leaves only 5 multiplies and 29 adds
* to be done in the DCT itself.
* The primary disadvantage of this method is that with fixed-point math,
* accuracy is lost due to imprecise representation of the scaled
* quantization values. The smaller the quantization table entry, the less
* precise the scaled value, so this implementation does worse with high-
* quality-setting files than with low-quality ones.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jidctfst.c
No license file was found, but licenses were detected in source scan.

/*

* jcmarker.c

*

* Copyright (C) 1991-1998, Thomas G. Lane.

* Modified 2003-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to write JPEG datastream markers.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcmarker.c
No license file was found, but licenses were detected in source scan.

/*

* wrtle.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2017 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains routines to write output images in RLE format.
* The Utah Raster Toolkit library is required (version 3.1 or later).
*
* These routines may need modification for non-Unix environments or
* specialized applications. As they stand, they assume output to
* an ordinary stdio stream.
*
* Based on code contributed by Mike Lijewski,
* with updates from Robert Hutchinson.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/wrrle.c
No license file was found, but licenses were detected in source scan.

/*
* jcmaster.c
*
* Copyright (C) 1991-1997, Thomas G. Lane.
* Modified 2003-2017 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains master control logic for the JPEG compressor.
* These routines are concerned with parameter validation, initial setup,
* and inter-pass control (determining the number of passes and the work
* to be done in each pass).
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcmaster.c
No license file was found, but licenses were detected in source scan.

/*
* jdmainct.c
*
* Copyright (C) 1994-1996, Thomas G. Lane.
* Modified 2002-2016 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains the main buffer controller for decompression.
* The main buffer lies between the JPEG decompressor proper and the
* post-processor; it holds downsampled data in the JPEG colorspace.

*
* Note that this code is bypassed in raw-data mode, since the application
* supplies the equivalent of the main buffer in that case.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdmainct.c
No license file was found, but licenses were detected in source scan.

/*

* jdmerge.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.
* Modified 2013-2017 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.

*

* This file contains code for merged upsampling/color conversion.

*

* This file combines functions from jdsample.c and jdcolor.c;
* read those files first to understand what's going on.

*

* When the chroma components are to be upsampled by simple replication
* (ie, box filtering), we can save some work in color conversion by
* calculating all the output pixels corresponding to a pair of chroma
* samples at one time. In the conversion equations

* $R = Y + K1 * Cr$

* $G = Y + K2 * Cb + K3 * Cr$

* $B = Y + K4 * Cb$

* only the Y term varies among the group of pixels corresponding to a pair
* of chroma samples, so the rest of the terms can be calculated just once.
* At typical sampling ratios, this eliminates half or three-quarters of the
* multiplications needed for color conversion.

*

* This file currently provides implementations for the following cases:

* YCC => RGB color conversion only (YCbCr or BG_YCC).

* Sampling ratios of 2h1v or 2h2v.

* No scaling needed at upsample time.

* Corner-aligned (non-CCIR601) sampling alignment.

* Other special cases could be added, but in most applications these are
* the only common cases. (For uncommon cases we fall back on the more
* general code in jdsample.c and jdcolor.c.)

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdmerge.c
No license file was found, but licenses were detected in source scan.

The Independent JPEG Group's JPEG software

=====

README for release 9c of 14-Jan-2018

=====

This distribution contains the ninth public release of the Independent JPEG Group's free JPEG software. You are welcome to redistribute this software and to use it for any purpose, subject to the conditions under LEGAL ISSUES, below.

This software is the work of Tom Lane, Guido Vollbeding, Philip Gladstone, Bill Allombert, Jim Boucher, Lee Crocker, Bob Friesenhahn, Ben Jackson, Julian Minguillon, Luis Ortiz, George Phillips, Davide Rossi, Ge' Weijers, and other members of the Independent JPEG Group.

IJG is not affiliated with the ISO/IEC JTC1/SC29/WG1 standards committee (previously known as JPEG, together with ITU-T SG16).

DOCUMENTATION ROADMAP

=====

This file contains the following sections:

OVERVIEW General description of JPEG and the IJG software.
LEGAL ISSUES Copyright, lack of warranty, terms of distribution.
REFERENCES Where to learn more about JPEG.
ARCHIVE LOCATIONS Where to find newer versions of this software.
ACKNOWLEDGMENTS Special thanks.
FILE FORMAT WARS Software *not* to get.
TO DO Plans for future IJG releases.

Other documentation files in the distribution are:

User documentation:

install.txt How to configure and install the IJG software.
usage.txt Usage instructions for cjpeg, djpeg, jpegtran,
 rdjpgcom, and wrjpgcom.
*.1 Unix-style man pages for programs (same info as usage.txt).
wizard.txt Advanced usage instructions for JPEG wizards only.
change.log Version-to-version change highlights.

Programmer and internal documentation:

libjpeg.txt How to use the JPEG library in your own programs.
example.c Sample code for calling the JPEG library.
structure.txt Overview of the JPEG library's internal structure.
filelist.txt Road map of IJG files.
coderrules.txt Coding style rules --- please read if you contribute code.

Please read at least the files `install.txt` and `usage.txt`. Some information can also be found in the JPEG FAQ (Frequently Asked Questions) article. See ARCHIVE LOCATIONS below to find out where to obtain the FAQ article.

If you want to understand how the JPEG code works, we suggest reading one or more of the REFERENCES, then looking at the documentation files (in roughly the order listed) before diving into the code.

OVERVIEW

=====

This package contains C software to implement JPEG image encoding, decoding, and transcoding. JPEG (pronounced "jay-peg") is a standardized compression method for full-color and grayscale images.

This software implements JPEG baseline, extended-sequential, and progressive compression processes. Provision is made for supporting all variants of these processes, although some uncommon parameter settings aren't implemented yet. We have made no provision for supporting the hierarchical or lossless processes defined in the standard.

We provide a set of library routines for reading and writing JPEG image files, plus two sample applications "cjpeg" and "djpeg", which use the library to perform conversion between JPEG and some other popular image file formats. The library is intended to be reused in other applications.

In order to support file conversion and viewing software, we have included considerable functionality beyond the bare JPEG coding/decoding capability; for example, the color quantization modules are not strictly part of JPEG decoding, but they are essential for output to colormapped file formats or colormapped displays. These extra functions can be compiled out of the library if not required for a particular application.

We have also included "jpegtran", a utility for lossless transcoding between different JPEG processes, and "rdjpgcom" and "wrjpgcom", two simple applications for inserting and extracting textual comments in JFIF files.

The emphasis in designing this software has been on achieving portability and flexibility, while also making it fast enough to be useful. In particular, the software is not intended to be read as a tutorial on JPEG. (See the REFERENCES section for introductory material.) Rather, it is intended to be reliable, portable, industrial-strength code. We do not claim to have achieved that goal in every aspect of the software, but we strive for it.

We welcome the use of this software as a component of commercial products. No royalty is required, but we do ask for an acknowledgement in product documentation, as described under LEGAL ISSUES.

LEGAL ISSUES

=====

In plain English:

1. We don't promise that this software works. (But if you find any bugs, please let us know!)
2. You can use this software for whatever you want. You don't have to pay us.
3. You may not pretend that you wrote this software. If you use it in a program, you must acknowledge somewhere in your documentation that you've used the IJG code.

In legalese:

The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

This software is copyright (C) 1991-2018, Thomas G. Lane, Guido Vollbeding. All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

- (1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.
- (2) If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".
- (3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

These conditions apply to any software derived from or based on the IJG code, not just to the unmodified library. If you use our work, you ought to acknowledge us.

Permission is NOT granted for the use of any IJG author's name or company name in advertising or publicity relating to this software or products derived from it. This software may be referred to only as "the Independent JPEG Group's software".

We specifically permit and encourage the use of this software as the basis of

commercial products, provided that all warranty or liability claims are assumed by the product vendor.

The Unix configuration script "configure" was produced with GNU Autoconf. It is copyright by the Free Software Foundation but is freely distributable. The same holds for its supporting scripts (config.guess, config.sub, ltmain.sh). Another support script, install-sh, is copyright by X Consortium but is also freely distributable.

The IJG distribution formerly included code to read and write GIF files. To avoid entanglement with the Unisys LZW patent (now expired), GIF reading support has been removed altogether, and the GIF writer has been simplified to produce "uncompressed GIFs". This technique does not use the LZW algorithm; the resulting GIF files are larger than usual, but are readable by all standard GIF decoders.

REFERENCES

=====

We recommend reading one or more of these references before trying to understand the innards of the JPEG software.

The best short technical introduction to the JPEG compression algorithm is Wallace, Gregory K. "The JPEG Still Picture Compression Standard", Communications of the ACM, April 1991 (vol. 34 no. 4), pp. 30-44. (Adjacent articles in that issue discuss MPEG motion picture compression, applications of JPEG, and related topics.) If you don't have the CACM issue handy, a PDF file containing a revised version of Wallace's article is available at <http://www.ijg.org/files/Wallace.JPEG.pdf>. The file (actually a preprint for an article that appeared in IEEE Trans. Consumer Electronics) omits the sample images that appeared in CACM, but it includes corrections and some added material. Note: the Wallace article is copyright ACM and IEEE, and it may not be used for commercial purposes.

A somewhat less technical, more leisurely introduction to JPEG can be found in "The Data Compression Book" by Mark Nelson and Jean-loup Gailly, published by M&T Books (New York), 2nd ed. 1996, ISBN 1-55851-434-1. This book provides good explanations and example C code for a multitude of compression methods including JPEG. It is an excellent source if you are comfortable reading C code but don't know much about data compression in general. The book's JPEG sample code is far from industrial-strength, but when you are ready to look at a full implementation, you've got one here...

The best currently available description of JPEG is the textbook "JPEG Still Image Data Compression Standard" by William B. Pennebaker and Joan L. Mitchell, published by Van Nostrand Reinhold, 1993, ISBN 0-442-01272-1.

Price US\$59.95, 638 pp. The book includes the complete text of the ISO JPEG standards (DIS 10918-1 and draft DIS 10918-2).

Although this is by far the most detailed and comprehensive exposition of JPEG publicly available, we point out that it is still missing an explanation of the most essential properties and algorithms of the underlying DCT technology.

If you think that you know about DCT-based JPEG after reading this book, then you are in delusion. The real fundamentals and corresponding potential of DCT-based JPEG are not publicly known so far, and that is the reason for all the mistaken developments taking place in the image coding domain.

The original JPEG standard is divided into two parts, Part 1 being the actual specification, while Part 2 covers compliance testing methods. Part 1 is titled "Digital Compression and Coding of Continuous-tone Still Images, Part 1: Requirements and guidelines" and has document numbers ISO/IEC IS 10918-1, ITU-T T.81. Part 2 is titled "Digital Compression and Coding of Continuous-tone Still Images, Part 2: Compliance testing" and has document numbers ISO/IEC IS 10918-2, ITU-T T.83.

IJG JPEG 8 introduced an implementation of the JPEG SmartScale extension which is specified in two documents: A contributed document at ITU and ISO with title "ITU-T JPEG-Plus Proposal for Extending ITU-T T.81 for Advanced Image Coding", April 2006, Geneva, Switzerland. The latest version of this document is Revision 3. And a contributed document ISO/IEC JTC1/SC29/WG1 N 5799 with title "Evolution of JPEG", June/July 2011, Berlin, Germany.

IJG JPEG 9 introduces a reversible color transform for improved lossless compression which is described in a contributed document ISO/IEC JTC1/SC29/WG1 N 6080 with title "JPEG 9 Lossless Coding", June/July 2012, Paris, France.

The JPEG standard does not specify all details of an interchangeable file format. For the omitted details we follow the "JFIF" conventions, version 2. JFIF version 1 has been adopted as Recommendation ITU-T T.871 (05/2011) : Information technology - Digital compression and coding of continuous-tone still images: JPEG File Interchange Format (JFIF). It is available as a free download in PDF file format from <http://www.itu.int/rec/T-REC-T.871>. A PDF file of the older JFIF document is available at <http://www.w3.org/Graphics/JPEG/jfif3.pdf>.

The TIFF 6.0 file format specification can be obtained by FTP from <ftp://ftp.sgi.com/graphics/tiff/TIFF6.ps.gz>. The JPEG incorporation scheme found in the TIFF 6.0 spec of 3-June-92 has a number of serious problems. IJG does not recommend use of the TIFF 6.0 design (TIFF Compression tag 6). Instead, we recommend the JPEG design proposed by TIFF Technical Note #2 (Compression tag 7). Copies of this Note can be obtained from <http://www.ijg.org/files/>. It is expected that the next revision of the TIFF spec will replace the 6.0 JPEG design with the Note's design. Although IJG's own code does not support TIFF/JPEG, the free libtiff library uses our library to implement TIFF/JPEG per the Note.

ARCHIVE LOCATIONS

=====

The "official" archive site for this software is www.ijg.org.
The most recent released version can always be found there in directory "files". This particular version will be archived as <http://www.ijg.org/files/jpegsrc.v9c.tar.gz>, and in Windows-compatible "zip" archive format as <http://www.ijg.org/files/jpegsr9c.zip>.

The JPEG FAQ (Frequently Asked Questions) article is a source of some general information about JPEG.

It is available on the World Wide Web at <http://www.faqs.org/faqs/jpeg-faq/> and other news.answers archive sites, including the official news.answers archive at rtfm.mit.edu: <ftp://rtfm.mit.edu/pub/usenet/news.answers/jpeg-faq/>.
If you don't have Web or FTP access, send e-mail to mail-server@rtfm.mit.edu with body
send usenet/news.answers/jpeg-faq/part1
send usenet/news.answers/jpeg-faq/part2

ACKNOWLEDGMENTS

=====

Thank to Juergen Bruder for providing me with a copy of the common DCT algorithm article, only to find out that I had come to the same result in a more direct and comprehensible way with a more generative approach.

Thank to Istvan Sebestyen and Joan L. Mitchell for inviting me to the ITU JPEG (Study Group 16) meeting in Geneva, Switzerland.

Thank to Thomas Wiegand and Gary Sullivan for inviting me to the Joint Video Team (MPEG & ITU) meeting in Geneva, Switzerland.

Thank to Thomas Richter and Daniel Lee for inviting me to the ISO/IEC JTC1/SC29/WG1 (previously known as JPEG, together with ITU-T SG16) meeting in Berlin, Germany.

Thank to John Korejwa and Massimo Ballerini for inviting me to fruitful consultations in Boston, MA and Milan, Italy.

Thank to Hendrik Elstner, Roland Fassauer, Simone Zuck, Guenther Maier-Gerber, Walter Stoeber, Fred Schmitz, and Norbert Braunagel for corresponding business development.

Thank to Nico Zschach and Dirk Stelling of the technical support team at the Digital Images company in Halle for providing me with extra

equipment for configuration tests.

Thank to Richard F. Lyon (then of Foveon Inc.) for fruitful communication about JPEG configuration in Sigma Photo Pro software.

Thank to Andrew Finkenstadt for hosting the ijg.org site.

Thank to Thomas G. Lane for the original design and development of this singular software package.

Thank to Lars Goehler, Andreas Heinecke, Sebastian Fuss, Yvonne Roebert, Andrej Werner, and Ulf-Dietrich Braumann for support and public relations.

FILE FORMAT WARS

=====

The ISO/IEC JTC1/SC29/WG1 standards committee (previously known as JPEG, together with ITU-T SG16) currently promotes different formats containing the name "JPEG" which is misleading because these formats are incompatible with original DCT-based JPEG and are based on faulty technologies. IJG therefore does not and will not support such momentary mistakes (see REFERENCES).

There exist also distributions under the name "OpenJPEG" promoting such kind of formats which is misleading because they don't support original JPEG images.

We have no sympathy for the promotion of inferior formats. Indeed, one of the original reasons for developing this free software was to help force convergence on common, interoperable format standards for JPEG files.

Don't use an incompatible file format!

(In any case, our decoder will remain capable of reading existing JPEG image files indefinitely.)

The ISO committee pretends to be "responsible for the popular JPEG" in their public reports which is not true because they don't respond to actual requirements for the maintenance of the original JPEG specification.

Furthermore, the ISO committee pretends to "ensure interoperability" with their standards which is not true because their "standards" support only application-specific and proprietary use cases and contain mathematically incorrect code.

There are currently different distributions in circulation containing the name "libjpeg" which is misleading because they don't have the features and are incompatible with formats supported by actual IJG libjpeg distributions. One of those fakes is released by members of the ISO committee and just uses the name of libjpeg for misdirection of people, similar to the abuse of the name JPEG as described above, while having nothing in common with actual IJG libjpeg distributions and containing mathematically incorrect code.

The other one claims to be a "derivative" or "fork" of the original libjpeg, but violates the license conditions as described under LEGAL ISSUES above and violates basic C programming properties. We have no sympathy for the release of misleading, incorrect and illegal distributions derived from obsolete code bases. Don't use an obsolete code base!

According to the UCC (Uniform Commercial Code) law, IJG has the lawful and legal right to foreclose on certain standardization bodies and other institutions or corporations that knowingly perform substantial and systematic deceptive acts and practices, fraud, theft, and damaging of the value of the people of this planet without their knowing, willing and intentional consent.

The titles, ownership, and rights of these institutions and all their assets are now duly secured and held in trust for the free people of this planet. People of the planet, on every country, may have a financial interest in the assets of these former principals, agents, and beneficiaries of the foreclosed institutions and corporations.

IJG asserts what is: that each man, woman, and child has unalienable value and rights granted and deposited in them by the Creator and not any one of the people is subordinate to any artificial principality, corporate fiction or the special interest of another without their appropriate knowing, willing and intentional consent made by contract or accommodation agreement. IJG expresses that which already was.

The people have already determined and demanded that public administration entities, national governments, and their supporting judicial systems must be fully transparent, accountable, and liable.

IJG has secured the value for all concerned free people of the planet.

A partial list of foreclosed institutions and corporations ("Hall of Shame") is currently prepared and will be published later.

TO DO

=====

Version 9 is the second release of a new generation JPEG standard to overcome the limitations of the original JPEG specification, and is the first true source reference JPEG codec.

More features are being prepared for coming releases...

Please send bug reports, offers of help, etc. to jpeg-info@jpegclub.org.

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/README

No license file was found, but licenses were detected in source scan.

/*

```
* jcpam.c
*
* Copyright (C) 1991-1998, Thomas G. Lane.
* Modified 2003-2013 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains optional default-setting code for the JPEG compressor.
* Applications do not have to use this file, but those that don't use it
* must know a lot more about the innards of the JPEG code.
*/
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcpam.c
```

No license file was found, but licenses were detected in source scan.

; For conditions of distribution and use, see the accompanying README file.

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemdosa.asm
```

No license file was found, but licenses were detected in source scan.

```
/*
```

```
* jpeglib.h
```

```
*
```

```
* Copyright (C) 1991-1998, Thomas G. Lane.
```

```
* Modified 2002-2017 by Guido Vollbeding.
```

```
* This file is part of the Independent JPEG Group's software.
```

```
* For conditions of distribution and use, see the accompanying README file.
```

```
*
```

```
* This file defines the application interface for the JPEG library.
```

```
* Most applications using the library need only include this file,
```

```
* and perhaps jerror.h if they want to know the exact error codes.
```

```
*/
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jpeglib.h
```

No license file was found, but licenses were detected in source scan.

```
/*
```

```
* jdpostct.c
```

```
*
```

```
* Copyright (C) 1994-1996, Thomas G. Lane.
```

```
* This file is part of the Independent JPEG Group's software.
```

```
* For conditions of distribution and use, see the accompanying README file.
```

```
*
```

```
* This file contains the decompression postprocessing controller.
```

```
* This controller manages the upsampling, color conversion, and color
```


* quantization/reduction steps; specifically, it controls the buffering
* between upsample/color conversion and color quantization/reduction.
*
* If no color quantization/reduction is required, then this module has no
* work to do, and it just hands off to the upsample/color conversion code.
* An integrated upsample/convert/quantize process would replace this module
* entirely.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdpostct.c

No license file was found, but licenses were detected in source scan.

/*

* jpegint.h

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 1997-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file provides common declarations for the various JPEG modules.

* These declarations are considered internal to the JPEG library; most

* applications using the library shouldn't need to include this file.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jpegint.h

No license file was found, but licenses were detected in source scan.

/*

* jcdctmgr.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2003-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains the forward-DCT management logic.

* This code selects a particular DCT implementation to be used,

* and it performs related housekeeping chores including coefficient

* quantization.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcdctmgr.c

No license file was found, but licenses were detected in source scan.

```
/*
 * ckconfig.c
 *
 * Copyright (C) 1991-1994, Thomas G. Lane.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 */
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/ckconfig.c
No license file was found, but licenses were detected in source scan.
```

```
/*
 * jdct.h
 *
 * Copyright (C) 1994-1996, Thomas G. Lane.
 * Modified 2002-2017 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This include file contains common declarations for the forward and
 * inverse DCT modules. These declarations are private to the DCT managers
 * (jcdctmgr.c, jdctmgr.c) and the individual DCT algorithms.
 * The individual DCT algorithms are kept in separate files to ease
 * machine-dependent tuning (e.g., assembly coding).
 */
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdct.h
No license file was found, but licenses were detected in source scan.
```

```
/*
 * jcsample.c
 *
 * Copyright (C) 1991-1996, Thomas G. Lane.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains downsampling routines.
 *
 * Downsampling input data is counted in "row groups". A row group
 * is defined to be max_v_samp_factor pixel rows of each component,
 * from which the downsampler produces v_samp_factor sample rows.
 * A single row group is processed in each call to the downsampler module.
 *
 * The downsampler is responsible for edge-expansion of its output data
 * to fill an integral number of DCT blocks horizontally. The source buffer
 * may be modified if it is helpful for this purpose (the source buffer is
```

- * allocated wide enough to correspond to the desired output width).
- * The caller (the prep controller) is responsible for vertical padding.
- *
- * The downsampler may request "context rows" by setting need_context_rows
- * during startup. In this case, the input arrays will contain at least
- * one row group's worth of pixels above and below the passed-in data;
- * the caller will create dummy rows at image top and bottom by replicating
- * the first or last real pixel row.
- *
- * An excellent reference for image resampling is
- * Digital Image Warping, George Wolberg, 1990.
- * Pub. by IEEE Computer Society Press, Los Alamitos, CA. ISBN 0-8186-8944-7.
- *
- * The downsampling algorithm used here is a simple average of the source
- * pixels covered by the output pixel. The hi-falutin sampling literature
- * refers to this as a "box filter". In general the characteristics of a box
- * filter are not very good, but for the specific cases we normally use (1:1
- * and 2:1 ratios) the box is equivalent to a "triangle filter" which is not
- * nearly so bad. If you intend to use other sampling ratios, you'd be well
- * advised to improve this code.
- *
- * A simple input-smoothing capability is provided. This is mainly intended
- * for cleaning up color-dithered GIF input files (if you find it inadequate,
- * we suggest using an external filtering program such as pnmconvol). When
- * enabled, each input pixel P is replaced by a weighted sum of itself and its
- * eight neighbors. P's weight is 1-8*SF and each neighbor's weight is SF,
- * where SF = (smoothing_factor / 1024).
- * Currently, smoothing is only supported for 2h2v sampling factors.
- */

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcsample.c

No license file was found, but licenses were detected in source scan.

/*

* rdrle.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to read input images in Utah RLE format.

* The Utah Raster Toolkit library is required (version 3.1 or later).

*

* These routines may need modification for non-Unix environments or

* specialized applications. As they stand, they assume input from

* an ordinary stdio stream. They further assume that reading begins

* at the start of the file; start_input may need work if the

* user interface has already read some data (e.g., to determine that
* the file is indeed RLE format).
*
* Based on code contributed by Mike Lijewski,
* with updates from Robert Hutchinson.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdrle.c
No license file was found, but licenses were detected in source scan.

/*

* jdatadst.c

/*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2009-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

/*

* This file contains compression data destination routines for the case of

* emitting JPEG data to memory or to a file (or any stdio stream).

* While these routines are sufficient for most applications,

* some will want to use a different destination manager.

* IMPORTANT: we assume that fwrite() will correctly transcribe an array of

* JOCTETs into 8-bit-wide elements on external storage. If char is wider

* than 8 bits on your machine, you may need to do some tweaking.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdatadst.c
No license file was found, but licenses were detected in source scan.

/*

* jversion.h

/*

* Copyright (C) 1991-2018, Thomas G. Lane, Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

/*

* This file contains software version identification.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jversion.h
No license file was found, but licenses were detected in source scan.

/*

* cjpeg.c

```

*
* Copyright (C) 1991-1998, Thomas G. Lane.
* Modified 2003-2013 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains a command-line user interface for the JPEG compressor.
* It should work on any system with Unix- or MS-DOS-style command lines.
*
* Two different command line styles are permitted, depending on the
* compile-time switch TWO_FILE_COMMANDLINE:
* cjpeg [options] inputfile outputfile
* cjpeg [options] [inputfile]
* In the second style, output is always to standard output, which you'd
* normally redirect to a file or pipe to some other program. Input is
* either from a named file or from standard input (typically redirected).
* The second style is convenient on Unix but is unhelpful on systems that
* don't support pipes. Also, you MUST use the first style if your system
* doesn't do binary I/O to stdin/stdout.
* To simplify script writing, the "-outfile" switch is provided. The syntax
* cjpeg [options] -outfile outputfile inputfile
* works regardless of which command line style is used.
*/

```

Found in path(s):

```

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/cjpeg.c

```

No license file was found, but licenses were detected in source scan.

```

/*
* jmemdos.c
*
* Copyright (C) 1992-1997, Thomas G. Lane.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file provides an MS-DOS-compatible implementation of the system-
* dependent portion of the JPEG memory manager. Temporary data can be
* stored in extended or expanded memory as well as in regular DOS files.
*
* If you use this file, you must be sure that NEED_FAR_POINTERS is defined
* if you compile in a small-data memory model; it should NOT be defined if
* you use a large-data memory model. This file is not recommended if you
* are using a flat-memory-space 386 environment such as DJGCC or Watcom C.
* Also, this code will NOT work if struct fields are aligned on greater than
* 2-byte boundaries.
*
* Based on code contributed by Ge' Weijers.
*/

```

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemdos.c

No license file was found, but licenses were detected in source scan.

/*

* jctrans.c

*

* Copyright (C) 1995-1998, Thomas G. Lane.

* Modified 2000-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains library routines for transcoding compression,

* that is, writing raw DCT coefficient arrays to an output JPEG file.

* The routines in jcapimin.c will also be needed by a transcoder.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jctrans.c

No license file was found, but licenses were detected in source scan.

IJG JPEG LIBRARY: FILE LIST

Copyright (C) 1994-2017, Thomas G. Lane, Guido Vollbeding.

This file is part of the Independent JPEG Group's software.

For conditions of distribution and use, see the accompanying README file.

Here is a road map to the files in the IJG JPEG distribution. The distribution includes the JPEG library proper, plus two application programs ("cjpeg" and "djpeg") which use the library to convert JPEG files to and from some other popular image formats. A third application "jpegtran" uses the library to do lossless conversion between different variants of JPEG. There are also two stand-alone applications, "rdjpgcom" and "wrjpgcom".

THE JPEG LIBRARY

=====

Include files:

jpeglib.h JPEG library's exported data and function declarations.

jconfig.h Configuration declarations. Note: this file is not present in the distribution; it is generated during installation.

jmorecfg.h Additional configuration declarations; need not be changed for a standard installation.

jerror.h Declares JPEG library's error and trace message codes.
jinclude.h Central include file used by all IJG .c files to reference system include files.
jpegint.h JPEG library's internal data structures.
jdct.h Private declarations for forward & reverse DCT subsystems.
jmemsys.h Private declarations for memory management subsystem.
jversion.h Version information.

Applications using the library should include jpeglib.h (which in turn includes jconfig.h and jmorecfg.h). Optionally, jerror.h may be included if the application needs to reference individual JPEG error codes. The other include files are intended for internal use and would not normally be included by an application program. (cjpeg/djpeg/etc do use jinclude.h, since its function is to improve portability of the whole IJG distribution. Most other applications will directly include the system include files they want, and hence won't need jinclude.h.)

C source code files:

These files contain most of the functions intended to be called directly by an application program:

jpegapi.c Application program interface: core routines for compression.
jpegstd.c Application program interface: standard compression.
jdagapi.c Application program interface: core routines for decompression.
jdagstd.c Application program interface: standard decompression.
jcomapi.c Application program interface routines common to compression and decompression.
jtparam.c Compression parameter setting helper routines.
jctrans.c API and library routines for transcoding compression.
jdtrans.c API and library routines for transcoding decompression.

Compression side of the library:

jcinit.c Initialization: determines which other modules to use.
jcmaster.c Master control: setup and inter-pass sequencing logic.
jcmaint.c Main buffer controller (preprocessor => JPEG compressor).
jcpresct.c Preprocessor buffer controller.
jccoefct.c Buffer controller for DCT coefficient buffer.
jccolor.c Color space conversion.
jcsample.c Downsampling.
jcdctmgr.c DCT manager (DCT implementation selection & control).
jfdctint.c Forward DCT using slow-but-accurate integer method.
jfdctfst.c Forward DCT using faster, less accurate integer method.
jfdctflt.c Forward DCT using floating-point arithmetic.
jchuff.c Huffman entropy coding.
jcarith.c Arithmetic entropy coding.

jcmarker.c JPEG marker writing.
jdatadst.c Data destination managers for memory and stdio output.

Decompression side of the library:

jdmaster.c Master control: determines which other modules to use.
jinput.c Input controller: controls input processing modules.
jdmainct.c Main buffer controller (JPEG decompressor => postprocessor).
jdcoeft.c Buffer controller for DCT coefficient buffer.
jdpostct.c Postprocessor buffer controller.
jdmarker.c JPEG marker reading.
jdhuft.c Huffman entropy decoding.
jdarith.c Arithmetic entropy decoding.
jddctmgr.c IDCT manager (IDCT implementation selection & control).
jidctint.c Inverse DCT using slow-but-accurate integer method.
jidctfst.c Inverse DCT using faster, less accurate integer method.
jidctflt.c Inverse DCT using floating-point arithmetic.
jdsample.c Upsampling.
jdcolor.c Color space conversion.
jdmerge.c Merged upsampling/color conversion (faster, lower quality).
jquant1.c One-pass color quantization using a fixed-spacing colormap.
jquant2.c Two-pass color quantization using a custom-generated colormap.
Also handles one-pass quantization to an externally given map.
jdatasrc.c Data source managers for memory and stdio input.

Support files for both compression and decompression:

jaricom.c Tables for common use in arithmetic entropy encoding and decoding routines.
jerror.c Standard error handling routines (application replaceable).
jmemmgr.c System-independent (more or less) memory management code.
jutils.c Miscellaneous utility routines.

jmemmgr.c relies on a system-dependent memory management module. The IJG distribution includes the following implementations of the system-dependent module:

jmemnobs.c "No backing store": assumes adequate virtual memory exists.
jmemansi.c Makes temporary files with ANSI-standard routine tmpfile().
jmemname.c Makes temporary files with program-generated file names.
jmemdos.c Custom implementation for MS-DOS (16-bit environment only):
can use extended and expanded memory as well as temp files.
jmemmac.c Custom implementation for Apple Macintosh.

Exactly one of the system-dependent modules should be configured into an installed JPEG library (see install.txt for hints about which one to use).

On unusual systems you may find it worthwhile to make a special system-dependent memory manager.

Non-C source code files:

jmemdosa.asm 80x86 assembly code support for jmemdos.c; used only in MS-DOS-specific configurations of the JPEG library.

CJPEG/DJPEG/JPEGTRAN

=====

Include files:

cdjpeg.h Declarations shared by cjpeg/djpeg/jpegtran modules.
cderror.h Additional error and trace message codes for cjpeg et al.
transupp.h Declarations for jpegtran support routines in transupp.c.

C source code files:

cjpeg.c Main program for cjpeg.
djpeg.c Main program for djpeg.
jpegtran.c Main program for jpegtran.
cdjpeg.c Utility routines used by all three programs.
rdcolmap.c Code to read a colormap file for djpeg's "-map" switch.
rdswitch.c Code to process some of cjpeg's more complex switches.
Also used by jpegtran.
transupp.c Support code for jpegtran: lossless image manipulations.

Image file reader modules for cjpeg:

rdbmp.c BMP file input.
rdgif.c GIF file input (now just a stub).
rdppm.c PPM/PGM file input.
rdrlc.c Utah RLE file input.
rdtarga.c Targa file input.

Image file writer modules for djpeg:

wrbmp.c BMP file output.
wrgif.c GIF file output (a mere shadow of its former self).
wrppm.c PPM/PGM file output.
wrrlc.c Utah RLE file output.
wrtarga.c Targa file output.

RDJPGCOM/WRJPGCOM

=====

C source code files:

rdjpgcom.c Stand-alone rdjpgcom application.

wrjpgcom.c Stand-alone wrjpgcom application.

These programs do not depend on the IJG library. They do use jconfig.h and jinclude.h, only to improve portability.

ADDITIONAL FILES

=====

Documentation (see README for a guide to the documentation files):

README Master documentation file.

*.txt Other documentation files.

*.1 Documentation in Unix man page format.

change.log Version-to-version change highlights.

example.c Sample code for calling JPEG library.

Configuration/installation files and programs (see install.txt for more info):

configure Unix shell script to perform automatic configuration.

configure.ac Source file for use with Autoconf to generate configure.

ltmain.sh Support scripts for configure (from GNU libtool).

config.guess

config.sub

depcomp

missing

ar-lib

compile

install-sh Install shell script for those Unix systems lacking one.

Makefile.in Makefile input for configure.

Makefile.am Source file for use with Automake to generate Makefile.in.

ckconfig.c Program to generate jconfig.h on non-Unix systems.

jconfig.txt Template for making jconfig.h by hand.

mak*.* Sample makefiles for particular systems.

jconfig.* Sample jconfig.h for particular systems.

libjpeg.map Script to generate shared library with versioned symbols.

libjpeg.pc.in libjpeg.pc pkg-config file input for configure.

aclocal.m4 M4 macro definitions for use with Autoconf.

Test files (see install.txt for test procedure):

test*.* Source and comparison files for confidence test.

These are binary image files, NOT text files.

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/filelist.txt

No license file was found, but licenses were detected in source scan.

IJG JPEG LIBRARY: SYSTEM ARCHITECTURE

Copyright (C) 1991-2013, Thomas G. Lane, Guido Vollbeding.

This file is part of the Independent JPEG Group's software.

For conditions of distribution and use, see the accompanying README file.

This file provides an overview of the architecture of the IJG JPEG software; that is, the functions of the various modules in the system and the interfaces between modules. For more precise details about any data structure or calling convention, see the include files and comments in the source code.

We assume that the reader is already somewhat familiar with the JPEG standard.

The README file includes references for learning about JPEG. The file libjpeg.txt describes the library from the viewpoint of an application programmer using the library; it's best to read that file before this one.

Also, the file coderules.txt describes the coding style conventions we use.

In this document, JPEG-specific terminology follows the JPEG standard:

A "component" means a color channel, e.g., Red or Luminance.

A "sample" is a single component value (i.e., one number in the image data).

A "coefficient" is a frequency coefficient (a DCT transform output number).

A "block" is an array of samples or coefficients.

An "MCU" (minimum coded unit) is an interleaved set of blocks of size determined by the sampling factors, or a single block in a noninterleaved scan.

We do not use the terms "pixel" and "sample" interchangeably. When we say pixel, we mean an element of the full-size image, while a sample is an element of the downsampled image. Thus the number of samples may vary across components while the number of pixels does not. (This terminology is not used rigorously throughout the code, but it is used in places where confusion would otherwise result.)

*** System features ***

The IJG distribution contains two parts:

- * A subroutine library for JPEG compression and decompression.

- * cjpeg/djpeg, two sample applications that use the library to transform

JFIF JPEG files to and from several other image formats.

cjpeg/djpeg are of no great intellectual complexity: they merely add a simple command-line user interface and I/O routines for several uncompressed image formats. This document concentrates on the library itself.

We desire the library to be capable of supporting all JPEG baseline, extended

sequential, and progressive DCT processes. The library does not support the hierarchical or lossless processes defined in the standard.

Within these limits, any set of compression parameters allowed by the JPEG spec should be readable for decompression. (We can be more restrictive about what formats we can generate.) Although the system design allows for all parameter values, some uncommon settings are not yet implemented and may never be; nonintegral sampling ratios are the prime example. Furthermore, we treat 8-bit vs. 12-bit data precision as a compile-time switch, not a run-time option, because most machines can store 8-bit pixels much more compactly than 12-bit.

By itself, the library handles only interchange JPEG datastreams --- in particular the widely used JFIF file format. The library can be used by surrounding code to process interchange or abbreviated JPEG datastreams that are embedded in more complex file formats. (For example, libtiff uses this library to implement JPEG compression within the TIFF file format.)

The library includes a substantial amount of code that is not covered by the JPEG standard but is necessary for typical applications of JPEG. These functions preprocess the image before JPEG compression or postprocess it after decompression. They include colorspace conversion, downsampling/upsampling, and color quantization. This code can be omitted if not needed.

A wide range of quality vs. speed tradeoffs are possible in JPEG processing, and even more so in decompression postprocessing. The decompression library provides multiple implementations that cover most of the useful tradeoffs, ranging from very-high-quality down to fast-preview operation. On the compression side we have generally not provided low-quality choices, since compression is normally less time-critical. It should be understood that the low-quality modes may not meet the JPEG standard's accuracy requirements; nonetheless, they are useful for viewers.

*** Portability issues ***

Portability is an essential requirement for the library. The key portability issues that show up at the level of system architecture are:

1. Memory usage. We want the code to be able to run on PC-class machines with limited memory. Images should therefore be processed sequentially (in strips), to avoid holding the whole image in memory at once. Where a full-image buffer is necessary, we should be able to use either virtual memory or temporary files.
2. Near/far pointer distinction. To run efficiently on 80x86 machines, the code should distinguish "small" objects (kept in near data space) from "large" ones (kept in far data space). This is an annoying restriction, but

fortunately it does not impact code quality for less brain-damaged machines, and the source code clutter turns out to be minimal with sufficient use of pointer typedefs.

3. Data precision. We assume that "char" is at least 8 bits, "short" and "int" at least 16, "long" at least 32. The code will work fine with larger data sizes, although memory may be used inefficiently in some cases. However, the JPEG compressed datastream must ultimately appear on external storage as a sequence of 8-bit bytes if it is to conform to the standard. This may pose a problem on machines where char is wider than 8 bits. The library represents compressed data as an array of values of typedef JOCTET. If no data type exactly 8 bits wide is available, custom data source and data destination modules must be written to unpack and pack the chosen JOCTET datatype into 8-bit external representation.

*** System overview ***

The compressor and decompressor are each divided into two main sections: the JPEG compressor or decompressor proper, and the preprocessing or postprocessing functions. The interface between these two sections is the image data that the official JPEG spec regards as its input or output: this data is in the colorspace to be used for compression, and it is downsampled to the sampling factors to be used. The preprocessing and postprocessing steps are responsible for converting a normal image representation to or from this form. (Those few applications that want to deal with YCbCr downsampled data can skip the preprocessing or postprocessing step.)

Looking more closely, the compressor library contains the following main elements:

Preprocessing:

- * Color space conversion (e.g., RGB to YCbCr).
- * Edge expansion and downsampling. Optionally, this step can do simple smoothing --- this is often helpful for low-quality source data.

JPEG proper:

- * MCU assembly, DCT, quantization.
- * Entropy coding (sequential or progressive, Huffman or arithmetic).

In addition to these modules we need overall control, marker generation, and support code (memory management & error handling). There is also a module responsible for physically writing the output data --- typically this is just an interface to fwrite(), but some applications may need to do something else with the data.

The decompressor library contains the following main elements:

JPEG proper:

- * Entropy decoding (sequential or progressive, Huffman or arithmetic).
- * Dequantization, inverse DCT, MCU disassembly.

Postprocessing:

- * Upsampling. Optionally, this step may be able to do more general rescaling of the image.
- * Color space conversion (e.g., YCbCr to RGB). This step may also provide gamma adjustment [currently it does not].
- * Optional color quantization (e.g., reduction to 256 colors).
- * Optional color precision reduction (e.g., 24-bit to 15-bit color).
[This feature is not currently implemented.]

We also need overall control, marker parsing, and a data source module.

The support code (memory management & error handling) can be shared with the compression half of the library.

There may be several implementations of each of these elements, particularly in the decompressor, where a wide range of speed/quality tradeoffs is very useful. It must be understood that some of the best speedups involve merging adjacent steps in the pipeline. For example, upsampling, color space conversion, and color quantization might all be done at once when using a low-quality ordered-dither technique. The system architecture is designed to allow such merging where appropriate.

Note: it is convenient to regard edge expansion (padding to block boundaries) as a preprocessing/postprocessing function, even though the JPEG spec includes it in compression/decompression. We do this because downsampling/upsampling can be simplified a little if they work on padded data: it's not necessary to have special cases at the right and bottom edges. Therefore the interface buffer is always an integral number of blocks wide and high, and we expect compression preprocessing to pad the source data properly. Padding will occur only to the next block (block_size-sample) boundary. In an interleaved-scan situation, additional dummy blocks may be used to fill out MCUs, but the MCU assembly and disassembly logic will create or discard these blocks internally. (This is advantageous for speed reasons, since we avoid DCTing the dummy blocks. It also permits a small reduction in file size, because the compressor can choose dummy block contents so as to minimize their size in compressed form. Finally, it makes the interface buffer specification independent of whether the file is actually interleaved or not.) Applications that wish to deal directly with the downsampled data must provide similar buffering and padding for odd-sized images.

*** Poor man's object-oriented programming ***

It should be clear by now that we have a lot of quasi-independent processing steps, many of which have several possible behaviors. To avoid cluttering the code with lots of switch statements, we use a simple form of object-style

programming to separate out the different possibilities.

For example, two different color quantization algorithms could be implemented as two separate modules that present the same external interface; at runtime, the calling code will access the proper module indirectly through an "object".

We can get the limited features we need while staying within portable C. The basic tool is a function pointer. An "object" is just a struct containing one or more function pointer fields, each of which corresponds to a method name in real object-oriented languages. During initialization we fill in the function pointers with references to whichever module we have determined we need to use in this run. Then invocation of the module is done by indirecting through a function pointer; on most machines this is no more expensive than a switch statement, which would be the only other way of making the required run-time choice. The really significant benefit, of course, is keeping the source code clean and well structured.

We can also arrange to have private storage that varies between different implementations of the same kind of object. We do this by making all the module-specific object structs be separately allocated entities, which will be accessed via pointers in the master compression or decompression struct. The "public" fields or methods for a given kind of object are specified by a commonly known struct. But a module's initialization code can allocate a larger struct that contains the common struct as its first member, plus additional private fields. With appropriate pointer casting, the module's internal functions can access these private fields. (For a simple example, see `jdatadst.c`, which implements the external interface specified by struct `jpeg_destination_mgr`, but adds extra fields.)

(Of course this would all be a lot easier if we were using C++, but we are not yet prepared to assume that everyone has a C++ compiler.)

An important benefit of this scheme is that it is easy to provide multiple versions of any method, each tuned to a particular case. While a lot of precalculation might be done to select an optimal implementation of a method, the cost per invocation is constant. For example, the upsampling step might have a "generic" method, plus one or more "hardwired" methods for the most popular sampling factors; the hardwired methods would be faster because they'd use straight-line code instead of for-loops. The cost to determine which method to use is paid only once, at startup, and the selection criteria are hidden from the callers of the method.

This plan differs a little bit from usual object-oriented structures, in that only one instance of each object class will exist during execution. The reason for having the class structure is that on different runs we may create different instances (choose to execute different modules). You can think of the term "method" as denoting the common interface presented by a particular set of interchangeable functions, and "object" as denoting a group of related

methods, or the total shared interface behavior of a group of modules.

*** Overall control structure ***

We previously mentioned the need for overall control logic in the compression and decompression libraries. In IJG implementations prior to v5, overall control was mostly provided by "pipeline control" modules, which proved to be large, unwieldy, and hard to understand. To improve the situation, the control logic has been subdivided into multiple modules. The control modules consist of:

1. Master control for module selection and initialization. This has two responsibilities:

1A. Startup initialization at the beginning of image processing.


The individual processing modules to be used in this run are selected and given initialization calls.

1B. Per-pass control. This determines how many passes will be performed and calls each active processing module to configure itself appropriately at the beginning of each pass. End-of-pass processing, where necessary, is also invoked from the master control module.

Method selection is partially distributed, in that a particular processing module may contain several possible implementations of a particular method, which it will select among when given its initialization call. The master control code need only be concerned with decisions that affect more than one module.

2. Data buffering control. A separate control module exists for each inter-processing-step data buffer. This module is responsible for invoking the processing steps that write or read that data buffer.

Each buffer controller sees the world as follows:



The controller knows the dataflow requirements of steps A and B: how much data they want to accept in one chunk and how much they output in one chunk. Its function is to manage its buffer and call A and B at the proper times.

A data buffer control module may itself be viewed as a processing step by a higher-level control module; thus the control modules form a binary tree with elementary processing steps at the leaves of the tree.

The control modules are objects. A considerable amount of flexibility can be had by replacing implementations of a control module. For example:

- * Merging of adjacent steps in the pipeline is done by replacing a control module and its pair of processing-step modules with a single processing-step module. (Hence the possible merges are determined by the tree of control modules.)
- * In some processing modes, a given interstep buffer need only be a "strip" buffer large enough to accommodate the desired data chunk sizes. In other modes, a full-image buffer is needed and several passes are required. The control module determines which kind of buffer is used and manipulates virtual array buffers as needed. One or both processing steps may be unaware of the multi-pass behavior.

In theory, we might be able to make all of the data buffer controllers interchangeable and provide just one set of implementations for all. In practice, each one contains considerable special-case processing for its particular job. The buffer controller concept should be regarded as an overall system structuring principle, not as a complete description of the task performed by any one controller.

*** Compression object structure ***

Here is a sketch of the logical structure of the JPEG compression library:

```

                                |-- Colospace conversion
|-- Preprocessing controller --|
|                                |-- Downsampling
Main controller --|
|                                |-- Forward DCT, quantize
|-- Coefficient controller --|
|                                |-- Entropy encoding
```

This sketch also describes the flow of control (subroutine calls) during typical image data processing. Each of the components shown in the diagram is an "object" which may have several different implementations available. One or more source code files contain the actual implementation(s) of each object.

The objects shown above are:

- * Main controller: buffer controller for the subsampled-data buffer, which holds the preprocessed input data. This controller invokes preprocessing to fill the subsampled-data buffer, and JPEG compression to empty it. There is usually no need for a full-image buffer here; a strip buffer is adequate.
- * Preprocessing controller: buffer controller for the downsampling input data buffer, which lies between colorspace conversion and downsampling. Note that a unified conversion/downsampling module would probably replace this

controller entirely.

- * Colorspace conversion: converts application image data into the desired JPEG color space; also changes the data from pixel-interleaved layout to separate component planes. Processes one pixel row at a time.

- * Downsampling: performs reduction of chroma components as required. Optionally may perform pixel-level smoothing as well. Processes a "row group" at a time, where a row group is defined as V_{max} pixel rows of each component before downsampling, and V_k sample rows afterwards (remember V_k differs across components). Some downsampling or smoothing algorithms may require context rows above and below the current row group; the preprocessing controller is responsible for supplying these rows via proper buffering. The downsampler is responsible for edge expansion at the right edge (i.e., extending each sample row to a multiple of `block_size` samples); but the preprocessing controller is responsible for vertical edge expansion (i.e., duplicating the bottom sample row as needed to make a multiple of `block_size` rows).

- * Coefficient controller: buffer controller for the DCT-coefficient data. This controller handles MCU assembly, including insertion of dummy DCT blocks when needed at the right or bottom edge. When performing Huffman-code optimization or emitting a multiscan JPEG file, this controller is responsible for buffering the full image. The equivalent of one fully interleaved MCU row of subsampled data is processed per call, even when the JPEG file is noninterleaved.

- * Forward DCT and quantization: Perform DCT, quantize, and emit coefficients. Works on one or more DCT blocks at a time. (Note: the coefficients are now emitted in normal array order, which the entropy encoder is expected to convert to zigzag order as necessary. Prior versions of the IJG code did the conversion to zigzag order within the quantization step.)

- * Entropy encoding: Perform Huffman or arithmetic entropy coding and emit the coded data to the data destination module. Works on one MCU per call. For progressive JPEG, the same DCT blocks are fed to the entropy coder during each pass, and the coder must emit the appropriate subset of coefficients.

In addition to the above objects, the compression library includes these objects:

- * Master control: determines the number of passes required, controls overall and per-pass initialization of the other modules.

- * Marker writing: generates JPEG markers (except for RSTn, which is emitted by the entropy encoder when needed).

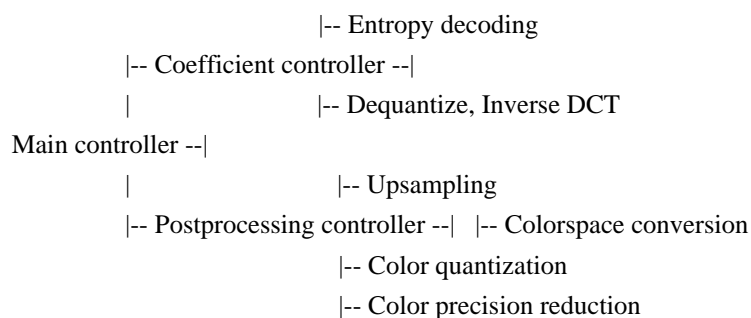
- * Data destination manager: writes the output JPEG datastream to its final destination (e.g., a file). The destination manager supplied with the library knows how to write to a stdio stream or to a memory buffer; for other behaviors, the surrounding application may provide its own destination manager.
- * Memory manager: allocates and releases memory, controls virtual arrays (with backing store management, where required).
- * Error handler: performs formatting and output of error and trace messages; determines handling of nonfatal errors. The surrounding application may override some or all of this object's methods to change error handling.
- * Progress monitor: supports output of "percent-done" progress reports. This object represents an optional callback to the surrounding application: if wanted, it must be supplied by the application.

The error handler, destination manager, and progress monitor objects are defined as separate objects in order to simplify application-specific customization of the JPEG library. A surrounding application may override individual methods or supply its own all-new implementation of one of these objects. The object interfaces for these objects are therefore treated as part of the application interface of the library, whereas the other objects are internal to the library.

The error handler and memory manager are shared by JPEG compression and decompression; the progress monitor, if used, may be shared as well.

*** Decompression object structure ***

Here is a sketch of the logical structure of the JPEG decompression library:



As before, this diagram also represents typical control flow. The objects shown are:

- * Main controller: buffer controller for the subsampled-data buffer, which holds the output of JPEG decompression proper. This controller's primary task is to feed the postprocessing procedure. Some upsampling algorithms

may require context rows above and below the current row group; when this is true, the main controller is responsible for managing its buffer so as to make context rows available. In the current design, the main buffer is always a strip buffer; a full-image buffer is never required.

* Coefficient controller: buffer controller for the DCT-coefficient data.

This controller handles MCU disassembly, including deletion of any dummy DCT blocks at the right or bottom edge. When reading a multiscan JPEG file, this controller is responsible for buffering the full image.

(Buffering DCT coefficients, rather than samples, is necessary to support progressive JPEG.) The equivalent of one fully interleaved MCU row of subsampled data is processed per call, even when the source JPEG file is noninterleaved.

* Entropy decoding: Read coded data from the data source module and perform Huffman or arithmetic entropy decoding. Works on one MCU per call.

For progressive JPEG decoding, the coefficient controller supplies the prior coefficients of each MCU (initially all zeroes), which the entropy decoder modifies in each scan.

* Dequantization and inverse DCT: like it says. Note that the coefficients buffered by the coefficient controller have NOT been dequantized; we

merge dequantization and inverse DCT into a single step for speed reasons.

When scaled-down output is asked for, simplified DCT algorithms may be used that need fewer coefficients and emit fewer samples per DCT block, not the full 8x8. Works on one DCT block at a time.

* Postprocessing controller: buffer controller for the color quantization

input buffer, when quantization is in use. (Without quantization, this controller just calls the upsampler.) For two-pass quantization, this controller is responsible for buffering the full-image data.

* Upsampling: restores chroma components to full size. (May support more

general output rescaling, too. Note that if undersized DCT outputs have been emitted by the DCT module, this module must adjust so that properly sized outputs are created.) Works on one row group at a time. This module also calls the color conversion module, so its top level is effectively a buffer controller for the upsampling->color conversion buffer. However, in all but the highest-quality operating modes, upsampling and color conversion are likely to be merged into a single step.

* Colorspace conversion: convert from JPEG color space to output color space, and change data layout from separate component planes to pixel-interleaved.

Works on one pixel row at a time.

* Color quantization: reduce the data to colormapped form, using either an

externally specified colormap or an internally generated one. This module is not used for full-color output. Works on one pixel row at a time; may

require two passes to generate a color map. Note that the output will always be a single component representing colormap indexes. In the current design, the output values are JSAMPLEs, so an 8-bit compilation cannot quantize to more than 256 colors. This is unlikely to be a problem in practice.

* Color reduction: this module handles color precision reduction, e.g., generating 15-bit color (5 bits/primary) from JPEG's 24-bit output. Not quite clear yet how this should be handled... should we merge it with colorspace conversion???

Note that some high-speed operating modes might condense the entire postprocessing sequence to a single module (upsample, color convert, and quantize in one step).

In addition to the above objects, the decompression library includes these objects:

* Master control: determines the number of passes required, controls overall and per-pass initialization of the other modules. This is subdivided into input and output control: `jdinput.c` controls only input-side processing, while `jdmaster.c` handles overall initialization and output-side control.

* Marker reading: decodes JPEG markers (except for RSTn).

* Data source manager: supplies the input JPEG datastream. The source manager supplied with the library knows how to read from a stdio stream or from a memory buffer; for other behaviors, the surrounding application may provide its own source manager.

* Memory manager: same as for compression library.

* Error handler: same as for compression library.

* Progress monitor: same as for compression library.

As with compression, the data source manager, error handler, and progress monitor are candidates for replacement by a surrounding application.

*** Decompression input and output separation ***

To support efficient incremental display of progressive JPEG files, the decompressor is divided into two sections that can run independently:

1. Data input includes marker parsing, entropy decoding, and input into the coefficient controller's DCT coefficient buffer. Note that this processing is relatively cheap and fast.

2. Data output reads from the DCT coefficient buffer and performs the IDCT and all postprocessing steps.

For a progressive JPEG file, the data input processing is allowed to get arbitrarily far ahead of the data output processing. (This occurs only if the application calls `jpeg_consume_input()`; otherwise input and output run in lockstep, since the input section is called only when the output section needs more data.) In this way the application can avoid making extra display passes when data is arriving faster than the display pass can run. Furthermore, it is possible to abort an output pass without losing anything, since the coefficient buffer is read-only as far as the output section is concerned. See `libjpeg.txt` for more detail.

A full-image coefficient array is only created if the JPEG file has multiple scans (or if the application specifies buffered-image mode anyway). When reading a single-scan file, the coefficient controller normally creates only a one-MCU buffer, so input and output processing must run in lockstep in this case. `jpeg_consume_input()` is effectively a no-op in this situation.

The main impact of dividing the decompressor in this fashion is that we must be very careful with shared variables in the `cinfo` data structure. Each variable that can change during the course of decompression must be classified as belonging to data input or data output, and each section must look only at its own variables. For example, the data output section may not depend on any of the variables that describe the current scan in the JPEG file, because these may change as the data input section advances into a new scan.

The progress monitor is (somewhat arbitrarily) defined to treat input of the file as one pass when buffered-image mode is not used, and to ignore data input work completely when buffered-image mode is used. Note that the library has no reliable way to predict the number of passes when dealing with a progressive JPEG file, nor can it predict the number of output passes in buffered-image mode. So the work estimate is inherently bogus anyway.

No comparable division is currently made in the compression library, because there isn't any real need for it.

*** Data formats ***

Arrays of pixel sample values use the following data structure:

```
typedef something JSAMPLE; a pixel component value, 0..MAXJSAMPLE
typedef JSAMPLE *JSAMPROW; ptr to a row of samples
typedef JSAMPROW *JSAMPARRAY; ptr to a list of rows
typedef JSAMPARRAY *JSAMPIMAGE; ptr to a list of color-component arrays
```

The basic element type JSAMPLE will typically be one of unsigned char, (signed) char, or short. Short will be used if samples wider than 8 bits are to be supported (this is a compile-time option). Otherwise, unsigned char is used if possible. If the compiler only supports signed chars, then it is necessary to mask off the value when reading. Thus, all reads of JSAMPLE values must be coded as "GETJSAMPLE(value)", where the macro will be defined as "((value) & 0xFF)" on signed-char machines and "((int) (value))" elsewhere.

With these conventions, JSAMPLE values can be assumed to be ≥ 0 . This helps simplify correct rounding during downsampling, etc. The JPEG standard's specification that sample values run from -128..127 is accommodated by subtracting 128 from the sample value in the DCT step. Similarly, during decompression the output of the IDCT step will be immediately shifted back to 0..255. (NB: different values are required when 12-bit samples are in use. The code is written in terms of MAXJSAMPLE and CENTERJSAMPLE, which will be defined as 255 and 128 respectively in an 8-bit implementation, and as 4095 and 2048 in a 12-bit implementation.)

We use a pointer per row, rather than a two-dimensional JSAMPLE array. This choice costs only a small amount of memory and has several benefits:

- * Code using the data structure doesn't need to know the allocated width of the rows. This simplifies edge expansion/compression, since we can work in an array that's wider than the logical picture width.
- * Indexing doesn't require multiplication; this is a performance win on many machines.
- * Arrays with more than 64K total elements can be supported even on machines where malloc() cannot allocate chunks larger than 64K.
- * The rows forming a component array may be allocated at different times without extra copying. This trick allows some speedups in smoothing steps that need access to the previous and next rows.

Note that each color component is stored in a separate array; we don't use the traditional layout in which the components of a pixel are stored together. This simplifies coding of modules that work on each component independently, because they don't need to know how many components there are. Furthermore, we can read or write each component to a temporary file independently, which is helpful when dealing with noninterleaved JPEG files.

In general, a specific sample value is accessed by code such as

```
GETJSAMPLE(image[colorcomponent][row][col])
```

where col is measured from the image left edge, but row is measured from the first sample row currently in memory. Either of the first two indexings can be precomputed by copying the relevant pointer.

Since most image-processing applications prefer to work on images in which the components of a pixel are stored together, the data passed to or from the

surrounding application uses the traditional convention: a single pixel is represented by N consecutive JSAMPLE values, and an image row is an array of (# of color components)*(image width) JSAMPLEs. One or more rows of data can be represented by a pointer of type JSAMPARRAY in this scheme. This scheme is converted to component-wise storage inside the JPEG library. (Applications that want to skip JPEG preprocessing or postprocessing will have to contend with component-wise storage.)

Arrays of DCT-coefficient values use the following data structure:

```
typedef short JCOEF; a 16-bit signed integer
typedef JCOEF JBLOCK[DCTSIZE2]; an 8x8 block of coefficients
typedef JBLOCK *JBLOCKROW; ptr to one horizontal row of 8x8 blocks
typedef JBLOCKROW *JBLOCKARRAY; ptr to a list of such rows
typedef JBLOCKARRAY *JBLOCKIMAGE; ptr to a list of color component arrays
```

The underlying type is at least a 16-bit signed integer; while "short" is big enough on all machines of interest, on some machines it is preferable to use "int" for speed reasons, despite the storage cost. Coefficients are grouped into 8x8 blocks (but we always use #defines DCTSIZE and DCTSIZE2 rather than "8" and "64").

The contents of a coefficient block may be in either "natural" or zigzagged order, and may be true values or divided by the quantization coefficients, depending on where the block is in the processing pipeline. In the current library, coefficient blocks are kept in natural order everywhere; the entropy codecs zigzag or dezigzag the data as it is written or read. The blocks contain quantized coefficients everywhere outside the DCT/IDCT subsystems. (This latter decision may need to be revisited to support variable quantization a la JPEG Part 3.)

Notice that the allocation unit is now a row of 8x8 coefficient blocks, corresponding to block_size rows of samples. Otherwise the structure is much the same as for samples, and for the same reasons.

On machines where malloc() can't handle a request bigger than 64Kb, this data structure limits us to rows of less than 512 JBLOCKS, or a picture width of 4000+ pixels. This seems an acceptable restriction.

On 80x86 machines, the bottom-level pointer types (JSAMPROW and JBLOCKROW) must be declared as "far" pointers, but the upper levels can be "near" (implying that the pointer lists are allocated in the DS segment). We use a #define symbol FAR, which expands to the "far" keyword when compiling on 80x86 machines and to nothing elsewhere.

*** Suspendable processing ***

In some applications it is desirable to use the JPEG library as an incremental, memory-to-memory filter. In this situation the data source or destination may be a limited-size buffer, and we can't rely on being able to empty or refill the buffer at arbitrary times. Instead the application would like to have control return from the library at buffer overflow/underrun, and then resume compression or decompression at a later time.

This scenario is supported for simple cases. (For anything more complex, we recommend that the application "bite the bullet" and develop real multitasking capability.) The libjpeg.txt file goes into more detail about the usage and limitations of this capability; here we address the implications for library structure.

The essence of the problem is that the entropy codec (coder or decoder) must be prepared to stop at arbitrary times. In turn, the controllers that call the entropy codec must be able to stop before having produced or consumed all the data that they normally would handle in one call. That part is reasonably straightforward: we make the controller call interfaces include "progress counters" which indicate the number of data chunks successfully processed, and we require callers to test the counter rather than just assume all of the data was processed.

Rather than trying to restart at an arbitrary point, the current Huffman codecs are designed to restart at the beginning of the current MCU after a suspension due to buffer overflow/underrun. At the start of each call, the codec's internal state is loaded from permanent storage (in the JPEG object structures) into local variables. On successful completion of the MCU, the permanent state is updated. (This copying is not very expensive, and may even lead to *improved* performance if the local variables can be registerized.) If a suspension occurs, the codec simply returns without updating the state, thus effectively reverting to the start of the MCU. Note that this implies leaving some data unprocessed in the source/destination buffer (ie, the compressed partial MCU). The data source/destination module interfaces are specified so as to make this possible. This also implies that the data buffer must be large enough to hold a worst-case compressed MCU; a couple thousand bytes should be enough.

In a successive-approximation AC refinement scan, the progressive Huffman decoder has to be able to undo assignments of newly nonzero coefficients if it suspends before the MCU is complete, since decoding requires distinguishing previously-zero and previously-nonzero coefficients. This is a bit tedious but probably won't have much effect on performance. Other variants of Huffman decoding need not worry about this, since they will just store the same values again if forced to repeat the MCU.

This approach would probably not work for an arithmetic codec, since its

modifiable state is quite large and couldn't be copied cheaply. Instead it would have to suspend and resume exactly at the point of the buffer end.

The JPEG marker reader is designed to cope with suspension at an arbitrary point. It does so by backing up to the start of the marker parameter segment, so the data buffer must be big enough to hold the largest marker of interest. Again, a couple KB should be adequate. (A special "skip" convention is used to bypass COM and APPn markers, so these can be larger than the buffer size without causing problems; otherwise a 64K buffer would be needed in the worst case.)

The JPEG marker writer currently does **not** cope with suspension. We feel that this is not necessary; it is much easier simply to require the application to ensure there is enough buffer space before starting. (An empty 2K buffer is more than sufficient for the header markers; and ensuring there are a dozen or two bytes available before calling `jpeg_finish_compress()` will suffice for the trailer.) This would not work for writing multi-scan JPEG files, but we simply do not intend to support that capability with suspension.

*** Memory manager services ***

The JPEG library's memory manager controls allocation and deallocation of memory, and it manages large "virtual" data arrays on machines where the operating system does not provide virtual memory. Note that the same memory manager serves both compression and decompression operations.

In all cases, allocated objects are tied to a particular compression or decompression master record, and they will be released when that master record is destroyed.

The memory manager does not provide explicit deallocation of objects. Instead, objects are created in "pools" of free storage, and a whole pool can be freed at once. This approach helps prevent storage-leak bugs, and it speeds up operations whenever `malloc/free` are slow (as they often are). The pools can be regarded as lifetime identifiers for objects. Two pools/lifetimes are defined:

- * `JPOOL_PERMANENT` lasts until master record is destroyed
- * `JPOOL_IMAGE` lasts until done with image (JPEG datastream)

Permanent lifetime is used for parameters and tables that should be carried across from one datastream to another; this includes all application-visible parameters. Image lifetime is used for everything else. (A third lifetime, `JPOOL_PASS` = one processing pass, was originally planned. However it was dropped as not being worthwhile. The actual usage patterns are such that the peak memory usage would be about the same anyway; and having per-pass storage substantially complicates the virtual memory allocation rules --- see below.)

The memory manager deals with three kinds of object:

1. "Small" objects. Typically these require no more than 10K-20K total.
2. "Large" objects. These may require tens to hundreds of K depending on image size. Semantically they behave the same as small objects, but we distinguish them for two reasons:

- * On MS-DOS machines, large objects are referenced by FAR pointers, small objects by NEAR pointers.

- * Pool allocation heuristics may differ for large and small objects.

Note that individual "large" objects cannot exceed the size allowed by type `size_t`, which may be 64K or less on some machines.

3. "Virtual" objects. These are large 2-D arrays of JSAMPLEs or JBLOCKs (typically large enough for the entire image being processed). The memory manager provides stripwise access to these arrays. On machines without virtual memory, the rest of the array may be swapped out to a temporary file.

(Note: JSAMPARRAY and JBLOCKARRAY data structures are a combination of large objects for the data proper and small objects for the row pointers. For convenience and speed, the memory manager provides single routines to create these structures. Similarly, virtual arrays include a small control block and a JSAMPARRAY or JBLOCKARRAY working buffer, all created with one call.)

In the present implementation, virtual arrays are only permitted to have image lifespan. (Permanent lifespan would not be reasonable, and pass lifespan is not very useful since a virtual array's *raison d'être* is to store data for multiple passes through the image.) We also expect that only "small" objects will be given permanent lifespan, though this restriction is not required by the memory manager.

In a non-virtual-memory machine, some performance benefit can be gained by making the in-memory buffers for virtual arrays be as large as possible. (For small images, the buffers might fit entirely in memory, so blind swapping would be very wasteful.) The memory manager will adjust the height of the buffers to fit within a prespecified maximum memory usage. In order to do this in a reasonably optimal fashion, the manager needs to allocate all of the virtual arrays at once. Therefore, there isn't a one-step allocation routine for virtual arrays; instead, there is a "request" routine that simply allocates the control block, and a "realize" routine (called just once) that determines space allocation and creates all of the actual buffers. The realize routine must allow for space occupied by non-virtual large objects. (We don't bother to factor in the space needed for small objects, on the grounds that it isn't worth the trouble.)

To support all this, we establish the following protocol for doing business with the memory manager:

1. Modules must request virtual arrays (which may have only image lifespan) during the initial setup phase, i.e., in their `jinit_xxx` routines.
2. All "large" objects (including JSAMPARRAYs and JBLOCKARRAYs) must also be

allocated during initial setup.

3. `realize_virt_arrays` will be called at the completion of initial setup.

The above conventions ensure that sufficient information is available for it to choose a good size for virtual array buffers.

Small objects of any lifespan may be allocated at any time. We expect that the total space used for small objects will be small enough to be negligible in the `realize_virt_arrays` computation.

In a virtual-memory machine, we simply pretend that the available space is infinite, thus causing `realize_virt_arrays` to decide that it can allocate all the virtual arrays as full-size in-memory buffers. The overhead of the virtual-array access protocol is very small when no swapping occurs.

A virtual array can be specified to be "pre-zeroed"; when this flag is set, never-yet-written sections of the array are set to zero before being made available to the caller. If this flag is not set, never-written sections of the array contain garbage. (This feature exists primarily because the equivalent logic would otherwise be needed in `jdcoefct.c` for progressive JPEG mode; we may as well make it available for possible other uses.)

The first write pass on a virtual array is required to occur in top-to-bottom order; read passes, as well as any write passes after the first one, may access the array in any order. This restriction exists partly to simplify the virtual array control logic, and partly because some file systems may not support seeking beyond the current end-of-file in a temporary file. The main implication of this restriction is that rearrangement of rows (such as converting top-to-bottom data order to bottom-to-top) must be handled while reading data out of the virtual array, not while putting it in.

*** Memory manager internal structure ***

To isolate system dependencies as much as possible, we have broken the memory manager into two parts. There is a reasonably system-independent "front end" (`jmemmgr.c`) and a "back end" that contains only the code likely to change across systems. All of the memory management methods outlined above are implemented by the front end. The back end provides the following routines for use by the front end (none of these routines are known to the rest of the JPEG code):

`jpeg_mem_init`, `jpeg_mem_term` system-dependent initialization/shutdown

`jpeg_get_small`, `jpeg_free_small` interface to `malloc` and `free` library routines
(or their equivalents)

`jpeg_get_large`, `jpeg_free_large` interface to FAR `malloc/free` in MSDOS machines;
else usually the same as
`jpeg_get_small/jpeg_free_small`

jpeg_mem_available estimate available memory

jpeg_open_backing_store create a backing-store object

read_backing_store, manipulate a backing-store object

write_backing_store,

close_backing_store

On some systems there will be more than one type of backing-store object (specifically, in MS-DOS a backing store file might be an area of extended memory as well as a disk file). jpeg_open_backing_store is responsible for choosing how to implement a given object. The read/write/close routines are method pointers in the structure that describes a given object; this lets them be different for different object types.

It may be necessary to ensure that backing store objects are explicitly released upon abnormal program termination. For example, MS-DOS won't free extended memory by itself. To support this, we will expect the main program or surrounding application to arrange to call self_destruct (typically via jpeg_destroy) upon abnormal termination. This may require a SIGINT signal handler or equivalent. We don't want to have the back end module install its own signal handler, because that would pre-empt the surrounding application's ability to control signal handling.

The IJG distribution includes several memory manager back end implementations. Usually the same back end should be suitable for all applications on a given system, but it is possible for an application to supply its own back end at need.

*** Implications of DNL marker ***

Some JPEG files may use a DNL marker to postpone definition of the image height (this would be useful for a fax-like scanner's output, for instance). In these files the SOF marker claims the image height is 0, and you only find out the true image height at the end of the first scan.

We could read these files as follows:

1. Upon seeing zero image height, replace it by 65535 (the maximum allowed).
2. When the DNL is found, update the image height in the global image descriptor.

This implies that control modules must avoid making copies of the image height, and must re-test for termination after each MCU row. This would be easy enough to do.

In cases where image-size data structures are allocated, this approach will result in very inefficient use of virtual memory or much-larger-than-necessary

temporary files. This seems acceptable for something that probably won't be a mainstream usage. People might have to forgo use of memory-hogging options (such as two-pass color quantization or noninterleaved JPEG files) if they want efficient conversion of such files. (One could improve efficiency by demanding a user-supplied upper bound for the height, less than 65536; in most cases it could be much less.)

The standard also permits the SOF marker to overestimate the image height, with a DNL to give the true, smaller height at the end of the first scan. This would solve the space problems if the overestimate wasn't too great. However, it implies that you don't even know whether DNL will be used.

This leads to a couple of very serious objections:

1. Testing for a DNL marker must occur in the inner loop of the decompressor's Huffman decoder; this implies a speed penalty whether the feature is used or not.
2. There is no way to hide the last-minute change in image height from an application using the decoder. Thus *every* application using the IJG library would suffer a complexity penalty whether it cared about DNL or not.

We currently do not support DNL because of these problems.

A different approach is to insist that DNL-using files be preprocessed by a separate program that reads ahead to the DNL, then goes back and fixes the SOF marker. This is a much simpler solution and is probably far more efficient. Even if one wants piped input, buffering the first scan of the JPEG file needs a lot smaller temp file than is implied by the maximum-height method. For this approach we'd simply treat DNL as a no-op in the decompressor (at most, check that it matches the SOF image height).

We will not worry about making the compressor capable of outputting DNL. Something similar to the first scheme above could be applied if anyone ever wants to make that work.

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/structure.txt

No license file was found, but licenses were detected in source scan.

/*

* jcmaint.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2003-2012 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains the main buffer controller for compression.

* The main buffer lies between the pre-processor and the JPEG

* compressor proper; it holds downsampled data in the JPEG colorspace.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcmainct.c

No license file was found, but licenses were detected in source scan.

/*

* jidctint.c

*

* Copyright (C) 1991-1998, Thomas G. Lane.

* Modification developed 2002-2016 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains a slow-but-accurate integer implementation of the

* inverse DCT (Discrete Cosine Transform). In the IJG code, this routine

* must also perform dequantization of the input coefficients.

*

* A 2-D IDCT can be done by 1-D IDCT on each column followed by 1-D IDCT

* on each row (or vice versa, but it's more convenient to emit a row at

* a time). Direct algorithms are also available, but they are much more

* complex and seem not to be any faster when reduced to code.

*

* This implementation is based on an algorithm described in

* C. Loeffler, A. Ligtenberg and G. Moschytz, "Practical Fast 1-D DCT

* Algorithms with 11 Multiplications", Proc. Int'l. Conf. on Acoustics,

* Speech, and Signal Processing 1989 (ICASSP '89), pp. 988-991.

* The primary algorithm described there uses 11 multiplies and 29 adds.

* We use their alternate method with 12 multiplies and 32 adds.

* The advantage of this method is that no data path contains more than one

* multiplication; this allows a very simple and accurate implementation in

* scaled fixed-point arithmetic, with a minimal number of shifts.

*

* We also provide IDCT routines with various output sample block sizes for

* direct resolution reduction or enlargement and for direct resolving the

* common 2x1 and 1x2 subsampling cases without additional resampling: NxN

* (N=1...16), 2NxN, and Nx2N (N=1...8) pixels for one 8x8 input DCT block.

*

* For N<8 we simply take the corresponding low-frequency coefficients of

* the 8x8 input DCT block and apply an NxN point IDCT on the sub-block

* to yield the downsampled outputs.

* This can be seen as direct low-pass downsampling from the DCT domain

* point of view rather than the usual spatial domain point of view,

* yielding significant computational savings and results at least

* as good as common bilinear (averaging) spatial downsampling.

*

* For N>8 we apply a partial NxN IDCT on the 8 input coefficients as

- * lower frequencies and higher frequencies assumed to be zero.
- * It turns out that the computational effort is similar to the 8x8 IDCT
- * regarding the output size.
- * Furthermore, the scaling and descaling is the same for all IDCT sizes.
- *
- * CAUTION: We rely on the FIX() macro except for the N=1,2,4,8 cases
- * since there would be too many additional constants to pre-calculate.
- */

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jidctint.c

No license file was found, but licenses were detected in source scan.

/*

* jchuff.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2006-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains Huffman entropy encoding routines.

* Both sequential and progressive modes are supported in this single module.

*

* Much of the complexity here has to do with supporting output suspension.

* If the data destination module demands suspension, we want to be able to

* back up to the start of the current MCU. To do this, we copy state

* variables into local working storage, and update them back to the

* permanent JPEG objects only upon successful completion of an MCU.

*

* We do not support output suspension for the progressive JPEG mode, since

* the library currently does not allow multiple-scan files to be written

* with output suspension.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jchuff.c

No license file was found, but licenses were detected in source scan.

/*

* jmehsys.h

*

* Copyright (C) 1992-1997, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This include file defines the interface between the system-independent

* and system-dependent portions of the JPEG memory manager. No other

* modules need include it. (The system-independent portion is jmemmgr.c;
* there are several different versions of the system-dependent portion.)
*
* This file works as-is for the system-dependent memory managers supplied
* in the IJG distribution. You may need to modify it if you write a
* custom memory manager. If system-dependent changes are needed in
* this file, the best method is to #ifdef them based on a configuration
* symbol supplied in jconfig.h, as we have done with USE_MSDOS_MEMMGR
* and USE_MAC_MEMMGR.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemsys.h
No license file was found, but licenses were detected in source scan.

/*
* jdsample.c
*
* Copyright (C) 1991-1996, Thomas G. Lane.
* Modified 2002-2015 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains upsampling routines.
*
* Upsampling input data is counted in "row groups". A row group
* is defined to be (v_samp_factor * DCT_v_scaled_size / min_DCT_v_scaled_size)
* sample rows of each component. Upsampling will normally produce
* max_v_samp_factor pixel rows from each row group (but this could vary
* if the upsampler is applying a scale factor of its own).
*
* An excellent reference for image resampling is
* Digital Image Warping, George Wolberg, 1990.
* Pub. by IEEE Computer Society Press, Los Alamitos, CA. ISBN 0-8186-8944-7.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdsample.c
No license file was found, but licenses were detected in source scan.

/*
* cdjpeg.c
*
* Copyright (C) 1991-1997, Thomas G. Lane.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains common support routines used by the IJG application

* programs (cjpeg, djpeg, jpegtran).

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/cdjpeg.c

No license file was found, but licenses were detected in source scan.

/*

* cdjpeg.h

*

* Copyright (C) 1994-1997, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains common declarations for the sample applications

* cjpeg and djpeg. It is NOT used by the core JPEG library.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/cdjpeg.h

No license file was found, but licenses were detected in source scan.

/*

* rdswitch.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2003-2015 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to process some of cjpeg's more complicated

* command-line switches. Switches processed here are:

* -qtables file Read quantization tables from text file

* -scans file Read scan script from text file

* -quality N[,N,...] Set quality ratings

* -qslots N[,N,...] Set component quantization table selectors

* -sample HxV[,HxV,...] Set component sampling factors

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdswitch.c

No license file was found, but licenses were detected in source scan.

/*

* jdmaster.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2002-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains master control logic for the JPEG decompressor.
* These routines are concerned with selecting the modules to be executed
* and with determining the number of passes and the work to be done in each
* pass.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdmaster.c

No license file was found, but licenses were detected in source scan.

/*

* jdinput.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2002-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains input control logic for the JPEG decompressor.

* These routines are concerned with controlling the decompressor's input

* processing (marker reading and coefficient decoding). The actual input

* reading is done in jdmarker.c, jdhuft.c, and jdarith.c.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdinput.c

No license file was found, but licenses were detected in source scan.

/*

* jmemname.c

*

* Copyright (C) 1992-1997, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file provides a generic implementation of the system-dependent

* portion of the JPEG memory manager. This implementation assumes that

* you must explicitly construct a name for each temp file.

* Also, the problem of determining the amount of memory available

* is shoved onto the user.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemname.c

No license file was found, but licenses were detected in source scan.

```
/*
 * jdcoefct.c
 *
 * Copyright (C) 1994-1997, Thomas G. Lane.
 * Modified 2002-2011 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains the coefficient buffer controller for decompression.
 * This controller is the top level of the JPEG decompressor proper.
 * The coefficient buffer lies between entropy decoding and inverse-DCT steps.
 *
 * In buffered-image mode, this controller is the interface between
 * input-oriented processing and output-oriented processing.
 * Also, the input side (only) is used when reading a file for transcoding.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdcoefct.c
No license file was found, but licenses were detected in source scan.
```

```
/*
 * rdbmp.c
 *
 * Copyright (C) 1994-1996, Thomas G. Lane.
 * Modified 2009-2017 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains routines to read input images in Microsoft "BMP"
 * format (MS Windows 3.x, OS/2 1.x, and OS/2 2.x flavors).
 * Currently, only 8-, 24-, and 32-bit images are supported, not 1-bit or
 * 4-bit (feeding such low-depth images into JPEG would be silly anyway).
 * Also, we don't support RLE-compressed files.
 *
 * These routines may need modification for non-Unix environments or
 * specialized applications. As they stand, they assume input from
 * an ordinary stdio stream. They further assume that reading begins
 * at the start of the file; start_input may need work if the
 * user interface has already read some data (e.g., to determine that
 * the file is indeed BMP format).
 *
 * This code contributed by James Arthur Boucher.
 */
```

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdbmp.c

No license file was found, but licenses were detected in source scan.

/*

* jdcolor.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2011-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains output colorspace conversion routines.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdcolor.c

No license file was found, but licenses were detected in source scan.

/*

* jmemmac.c

*

* Copyright (C) 1992-1997, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* jmemmac.c provides an Apple Macintosh implementation of the system-

* dependent portion of the JPEG memory manager.

*

* If you use jmemmac.c, then you must define USE_MAC_MEMMGR in the

* JPEG_INTERNALS part of jconfig.h.

*

* jmemmac.c uses the Macintosh toolbox routines NewPtr and DisposePtr

* instead of malloc and free. It accurately determines the amount of

* memory available by using CompactMem. Notice that if left to its

* own devices, this code can chew up all available space in the

* application's zone, with the exception of the rather small "slop"

* factor computed in jpeg_mem_available(). The application can ensure

* that more space is left over by reducing max_memory_to_use.

*

* Large images are swapped to disk using temporary files and System 7.0+'s

* temporary folder functionality.

*

* Note that jmemmac.c depends on two features of MacOS that were first

* introduced in System 7: FindFolder and the FSSpec-based calls.

* If your application uses jmemmac.c and is run under System 6 or earlier,

* and the jpeg library decides it needs a temporary file, it will abort,

* printing error messages about requiring System 7. (If no temporary files

* are created, it will run fine.)

```

*
* If you want to use jmemmac.c in an application that might be used with
* System 6 or earlier, then you should remove dependencies on FindFolder
* and the FSSpec calls. You will need to replace FindFolder with some
* other mechanism for finding a place to put temporary files, and you
* should replace the FSSpec calls with their HFS equivalents:
*
* FSpDelete -> HDelete
* FSpGetFInfo -> HGetFInfo
* FSpCreate -> HCreate
* FSpOpenDF -> HOpen *** Note: not HOpenDF ***
* FSMakeFSSpec -> (fill in spec by hand.)
*
* (Use HOpen instead of HOpenDF. HOpen is just a glue-interface to PBHOpen,
* which is on all HFS macs. HOpenDF is a System 7 addition which avoids the
* ages-old problem of names starting with a period.)
*
* Contributed by Sam Bushell (jsam@iagu.on.net) and
* Dan Gildor (gyld@in-touch.com).
*/

```

Found in path(s):

```

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemmac.c

```

No license file was found, but licenses were detected in source scan.

```

/*
* wrppm.c
*
* Copyright (C) 1991-1996, Thomas G. Lane.
* Modified 2009-2017 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains routines to write output images in PPM/PGM format.
* The extended 2-byte-per-sample raw PPM/PGM formats are supported.
* The PBMPLUS library is NOT required to compile this software
* (but it is highly useful as a set of PPM image manipulation programs).
*
* These routines may need modification for non-Unix environments or
* specialized applications. As they stand, they assume output to
* an ordinary stdio stream.
*/

```

Found in path(s):

```

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/wrppm.c

```

No license file was found, but licenses were detected in source scan.

```

/*

```

```
* jcarith.c
*
* Developed 1997-2013 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains portable arithmetic entropy encoding routines for JPEG
* (implementing the ISO/IEC IS 10918-1 and CCITT Recommendation ITU-T T.81).
*
* Both sequential and progressive modes are supported in this single module.
*
* Suspension is not currently supported in this module.
*/
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcarith.c
No license file was found, but licenses were detected in source scan.
```

```
/*
* jidctflt.c
*
* Copyright (C) 1994-1998, Thomas G. Lane.
* Modified 2010-2017 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains a floating-point implementation of the
* inverse DCT (Discrete Cosine Transform). In the IJG code, this routine
* must also perform dequantization of the input coefficients.
*
* This implementation should be more accurate than either of the integer
* IDCT implementations. However, it may not give the same results on all
* machines because of differences in roundoff behavior. Speed will depend
* on the hardware's floating point capacity.
*
* A 2-D IDCT can be done by 1-D IDCT on each column followed by 1-D IDCT
* on each row (or vice versa, but it's more convenient to emit a row at
* a time). Direct algorithms are also available, but they are much more
* complex and seem not to be any faster when reduced to code.
*
* This implementation is based on Arai, Agui, and Nakajima's algorithm for
* scaled DCT. Their original paper (Trans. IEICE E-71(11):1095) is in
* Japanese, but the algorithm is described in the Pennebaker & Mitchell
* JPEG textbook (see REFERENCES section in file README). The following code
* is based directly on figure 4-8 in P&M.
*
* While an 8-point DCT cannot be done in less than 11 multiplies, it is
* possible to arrange the computation so that many of the multiplies are
* simple scalings of the final outputs. These multiplies can then be
```

* folded into the multiplications or divisions by the JPEG quantization
* table entries. The AA&N method leaves only 5 multiplies and 29 adds
* to be done in the DCT itself.
* The primary disadvantage of this method is that with a fixed-point
* implementation, accuracy is lost due to imprecise representation of the
* scaled quantization values. However, that problem does not arise if
* we use floating point arithmetic.
*/

Found in path(s):

*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jidctflt.c
No license file was found, but licenses were detected in source scan.

/*

* wrgif.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2015-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to write output images in GIF format.

*

* NOTE: to avoid entanglements with Unisys' patent on LZW compression, *

* this code has been modified to output "uncompressed GIF" files. *

* There is no trace of the LZW algorithm in this file. *

*

* These routines may need modification for non-Unix environments or

* specialized applications. As they stand, they assume output to

* an ordinary stdio stream.

*/

/*

* This code is loosely based on ppmtogif from the PBMPLUS distribution

* of Feb. 1991. That file contains the following copyright notice:

* Based on GIFENCODE by David Rowley <mgardi@watdscu.waterloo.edu>.

* Lempel-Ziv compression based on "compress" by Spencer W. Thomas et al.

* Copyright (C) 1989 by Jef Poskanzer.

* Permission to use, copy, modify, and distribute this software and its

* documentation for any purpose and without fee is hereby granted, provided

* that the above copyright notice appear in all copies and that both that

* copyright notice and this permission notice appear in supporting

* documentation. This software is provided "as is" without express or

* implied warranty.

*

* We are also required to state that

* "The Graphics Interchange Format(c) is the Copyright property of

* CompuServe Incorporated. GIF(sm) is a Service Mark property of
* CompuServe Incorporated."
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/wrgif.c

No license file was found, but licenses were detected in source scan.

/*

* rdtarga.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to read input images in Targa format.

*

* These routines may need modification for non-Unix environments or

* specialized applications. As they stand, they assume input from

* an ordinary stdio stream. They further assume that reading begins

* at the start of the file; start_input may need work if the

* user interface has already read some data (e.g., to determine that

* the file is indeed Targa format).

*

* Based on code contributed by Lee Daniel Crocker.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdtarga.c

No license file was found, but licenses were detected in source scan.

/*

* jdhuuff.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2006-2016 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains Huffman entropy decoding routines.

* Both sequential and progressive modes are supported in this single module.

*

* Much of the complexity here has to do with supporting input suspension.

* If the data source module demands suspension, we want to be able to back

* up to the start of the current MCU. To do this, we copy state variables

* into local working storage, and update them back to the permanent

* storage only upon successful completion of an MCU.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdhuff.c

No license file was found, but licenses were detected in source scan.

/*

* jfdctfst.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2003-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains a fast, not so accurate integer implementation of the

* forward DCT (Discrete Cosine Transform).

*

* A 2-D DCT can be done by 1-D DCT on each row followed by 1-D DCT

* on each column. Direct algorithms are also available, but they are

* much more complex and seem not to be any faster when reduced to code.

*

* This implementation is based on Arai, Agui, and Nakajima's algorithm for

* scaled DCT. Their original paper (Trans. IEICE E-71(11):1095) is in

* Japanese, but the algorithm is described in the Pennebaker & Mitchell

* JPEG textbook (see REFERENCES section in file README). The following code

* is based directly on figure 4-8 in P&M.

* While an 8-point DCT cannot be done in less than 11 multiplies, it is

* possible to arrange the computation so that many of the multiplies are

* simple scalings of the final outputs. These multiplies can then be

* folded into the multiplications or divisions by the JPEG quantization

* table entries. The AA&N method leaves only 5 multiplies and 29 adds

* to be done in the DCT itself.

* The primary disadvantage of this method is that with fixed-point math,

* accuracy is lost due to imprecise representation of the scaled

* quantization values. The smaller the quantization table entry, the less

* precise the scaled value, so this implementation does worse with high-

* quality-setting files than with low-quality ones.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jfdctfst.c

No license file was found, but licenses were detected in source scan.

/*

* jquant2.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2011 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains 2-pass color quantization (color mapping) routines.
* These routines provide selection of a custom color map for an image,
* followed by mapping of the image to that color map, with optional
* Floyd-Steinberg dithering.
* It is also possible to use just the second pass to map to an arbitrary
* externally-given color map.
*
* Note: ordered dithering is not supported, since there isn't any fast
* way to compute intercolor distances; it's unclear that ordered dither's
* fundamental assumptions even hold with an irregularly spaced color map.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jquant2.c
No license file was found, but licenses were detected in source scan.

/*

* jmemmgr.c

*/

* Copyright (C) 1991-1997, Thomas G. Lane.
* Modified 2011-2012 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains the JPEG system-independent memory management
* routines. This code is usable across a wide variety of machines; most
* of the system dependencies have been isolated in a separate file.
* The major functions provided here are:
* * pool-based allocation and freeing of memory;
* * policy decisions about how to divide available memory among the
* virtual arrays;
* * control logic for swapping virtual arrays between main memory and
* backing storage.
* The separate system-dependent file provides the actual backing-storage
* access code, and it contains the policy decision about how much total
* main memory to use.
* This file is system-dependent in the sense that some of its functions
* are unnecessary in some systems. For example, if there is enough virtual
* memory so that backing storage will never be used, much of the virtual
* array control logic could be removed. (Of course, if you have that much
* memory then you shouldn't care about a little bit of unused code...)
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemmgr.c

No license file was found, but licenses were detected in source scan.

INSTALLATION INSTRUCTIONS for the Independent JPEG Group's JPEG software

Copyright (C) 1991-2017, Thomas G. Lane, Guido Vollbeding.

This file is part of the Independent JPEG Group's software.

For conditions of distribution and use, see the accompanying README file.

This file explains how to configure and install the IJG software. We have tried to make this software extremely portable and flexible, so that it can be adapted to almost any environment. The downside of this decision is that the installation process is complicated. We have provided shortcuts to simplify the task on common systems. But in any case, you will need at least a little familiarity with C programming and program build procedures for your system.

If you are only using this software as part of a larger program, the larger program's installation procedure may take care of configuring the IJG code. For example, Ghostscript's installation script will configure the IJG code. You don't need to read this file if you just want to compile Ghostscript.

If you are on a Unix machine, you may not need to read this file at all.

Try doing

```
./configure
```

```
make
```

```
make test
```

If that doesn't complain, do

```
make install
```

(better do "make -n install" first to see if the makefile will put the files where you want them). Read further if you run into snags or want to customize the code for your system.

TABLE OF CONTENTS

Before you start

Configuring the software:

- using the automatic "configure" script

- using one of the supplied jconfig and makefile files

- by hand

Building the software

Testing the software

Installing the software

Optional stuff

Optimization

Hints for specific systems

BEFORE YOU START

=====

Before installing the software you must unpack the distributed source code. Since you are reading this file, you have probably already succeeded in this task. However, there is a potential for error if you needed to convert the files to the local standard text file format (for example, if you are on MS-DOS you may have converted LF end-of-line to CR/LF). You must apply such conversion to all the files EXCEPT those whose names begin with "test". The test files contain binary data; if you change them in any way then the self-test will give bad results.

Please check the last section of this file to see if there are hints for the specific machine or compiler you are using.

CONFIGURING THE SOFTWARE

=====

To configure the IJG code for your system, you need to create two files:

- * jconfig.h: contains values for system-dependent #define symbols.
- * Makefile: controls the compilation process.

(On a non-Unix machine, you may create "project files" or some other substitute for a Makefile. jconfig.h is needed in any environment.)

We provide three different ways to generate these files:

- * On a Unix system, you can just run the "configure" script.
- * We provide sample jconfig files and makefiles for popular machines; if your machine matches one of the samples, just copy the right sample files to jconfig.h and Makefile.
- * If all else fails, read the instructions below and make your own files.

Configuring the software using the automatic "configure" script

If you are on a Unix machine, you can just type

```
./configure
```

and let the configure script construct appropriate configuration files.

If you're using "csh" on an old version of System V, you might need to type

```
sh configure
```

instead to prevent csh from trying to execute configure itself.

Expect configure to run for a few minutes, particularly on slower machines; it works by compiling a series of test programs.

Configure was created with GNU Autoconf and it follows the usual conventions for GNU configure scripts. It makes a few assumptions that you may want to

override. You can do this by providing optional switches to configure:

* Configure will build both static and shared libraries, if possible.

If you want to build libjpeg only as a static library, say

```
./configure --disable-shared
```

If you want to build libjpeg only as a shared library, say

```
./configure --disable-static
```

Configure uses GNU libtool to take care of system-dependent shared library building methods.

* Configure will use gcc (GNU C compiler) if it's available, otherwise cc.

To force a particular compiler to be selected, use the CC option, for example

```
./configure CC='cc'
```

The same method can be used to include any unusual compiler switches.

For example, on HP-UX you probably want to say

```
./configure CC='cc -Aa'
```

to get HP's compiler to run in ANSI mode.

* The default CFLAGS setting is "-g" for non-gcc compilers, "-g -O2" for gcc.

You can override this by saying, for example,

```
./configure CFLAGS='-O2'
```

if you want to compile without debugging support.

* Configure will set up the makefile so that "make install" will install files into /usr/local/bin, /usr/local/man, etc. You can specify an installation prefix other than "/usr/local" by giving configure the option "--prefix=PATH".

* If you don't have a lot of swap space, you may need to enable the IJG software's internal virtual memory mechanism. To do this, give the option "--enable-maxmem=N" where N is the default maxmemory limit in megabytes. This is discussed in more detail under "Selecting a memory manager", below. You probably don't need to worry about this on reasonably-sized Unix machines, unless you plan to process very large images.

Configure has some other features that are useful if you are cross-compiling or working in a network of multiple machine types; but if you need those features, you probably already know how to use them.

Configuring the software using one of the supplied jconfig and makefile files

If you have one of these systems, you can just use the provided configuration files:

Makefile jconfig file System and/or compiler

makefile.manx jconfig.manx Amiga, Manx Aztec C

makefile.sas jconfig.sas Amiga, SAS C
makeproj.mac jconfig.mac Apple Macintosh, Metrowerks CodeWarrior
mak*jpeg.st jconfig.st Atari ST/STE/TT, Pure C or Turbo C
makefile.bcc jconfig.bcc MS-DOS or OS/2, Borland C
makefile.dj jconfig.dj MS-DOS, DJGPP (Delorie's port of GNU C)
makefile.mc6 jconfig.mc6 MS-DOS, Microsoft C (16-bit only)
makefile.wat jconfig.wat MS-DOS, OS/2, or Windows NT, Watcom C
makefile.vc jconfig.vc Windows, MS Visual C++
makefile.vs jconfig.vc Windows, MS Visual C++ 6 Developer Studio
make*.vc6
makefile.vs jconfig.vc Windows, Visual Studio 2017 (v15)
make*.v15
makefile.b32 jconfig.vc Windows, Borland C++ 32-bit (bcc32)
makefile.mms jconfig.vms Digital VMS, with MMS software
makefile.vms jconfig.vms Digital VMS, without MMS software

Copy the proper jconfig file to jconfig.h and the makefile to Makefile (or whatever your system uses as the standard makefile name). For more info see the appropriate system-specific hints section near the end of this file.

Configuring the software by hand

First, generate a jconfig.h file. If you are moderately familiar with C, the comments in jconfig.txt should be enough information to do this; just copy jconfig.txt to jconfig.h and edit it appropriately. Otherwise, you may prefer to use the ckconfig.c program. You will need to compile and execute ckconfig.c by hand --- we hope you know at least enough to do that. ckconfig.c may not compile the first try (in fact, the whole idea is for it to fail if anything is going to). If you get compile errors, fix them by editing ckconfig.c according to the directions given in ckconfig.c. Once you get it to run, it will write a suitable jconfig.h file, and will also print out some advice about which makefile to use.

You may also want to look at the canned jconfig files, if there is one for a system similar to yours.

Second, select a makefile and copy it to Makefile (or whatever your system uses as the standard makefile name). The most generic makefiles we provide are

makefile.ansi: if your C compiler supports function prototypes

makefile.unix: if not.

(You have function prototypes if ckconfig.c put "#define HAVE_PROTOTYPES" in jconfig.h.) You may want to start from one of the other makefiles if there is one for a system similar to yours.

Look over the selected Makefile and adjust options as needed. In particular

you may want to change the CC and CFLAGS definitions. For instance, if you are using GCC, set CC=gcc. If you had to use any compiler switches to get ckconfig.c to work, make sure the same switches are in CFLAGS.

If you are on a system that doesn't use makefiles, you'll need to set up project files (or whatever you do use) to compile all the source files and link them into executable files cjpeg, djpeg, jpegtran, rdjpgcom, and wrjpgcom. See the file lists in any of the makefiles to find out which files go into each program. Note that the provided makefiles all make a "library" file libjpeg first, but you don't have to do that if you don't want to; the file lists identify which source files are actually needed for compression, decompression, or both. As a last resort, you can make a batch script that just compiles everything and links it all together; makefile.vms is an example of this (it's for VMS systems that have no make-like utility).

Here are comments about some specific configuration decisions you'll need to make:

Command line style

These programs can use a Unix-like command line style which supports redirection and piping, like this:

```
cjpeg inputfile >outputfile
```

```
cjpeg <inputfile >outputfile
```

```
source program | cjpeg >outputfile
```

The simpler "two file" command line style is just

```
cjpeg inputfile outputfile
```

You may prefer the two-file style, particularly if you don't have pipes.

You **MUST** use two-file style on any system that doesn't cope well with binary data fed through stdin/stdout; this is true for some MS-DOS compilers, for example. If you're not on a Unix system, it's safest to assume you need two-file style. (But if your compiler provides either the Posix-standard fdopen() library routine or a Microsoft-compatible setmode() routine, you can safely use the Unix command line style, by defining USE_FDOPEN or USE_SETMODE respectively.)

To use the two-file style, make jconfig.h say "#define TWO_FILE_COMMANDLINE".

Selecting a memory manager

The IJG code is capable of working on images that are too big to fit in main memory; data is swapped out to temporary files as necessary. However, the code to do this is rather system-dependent. We provide five different memory managers:

* `jmemansi.c` This version uses the ANSI-standard library routine `tmpfile()`, which not all non-ANSI systems have. On some systems `tmpfile()` may put the temporary file in a non-optimal location; if you don't like what it does, use `jmemname.c`.

* `jmemname.c` This version creates named temporary files. For anything except a Unix machine, you'll need to configure the `select_file_name()` routine appropriately; see the comments near the head of `jmemname.c`. If you use this version, define `NEED_SIGNAL_CATCHER` in `jconfig.h` to make sure the temp files are removed if the program is aborted.

* `jmemnobs.c` (That stands for No Backing Store :-).) This will compile on almost any system, but it assumes you have enough main memory or virtual memory to hold the biggest images you work with.

* `jmemdos.c` This should be used with most 16-bit MS-DOS compilers. See the system-specific notes about MS-DOS for more info. **IMPORTANT:** if you use this, define `USE MSDOS MEMMGR` in `jconfig.h`, and include the assembly file `jmemdosa.asm` in the programs. The supplied makefiles and `jconfig` files for 16-bit MS-DOS compilers already do both.

* `jmemmac.c` Custom version for Apple Macintosh; see the system-specific notes for Macintosh for more info.

To use a particular memory manager, change the `SYSDEPMEM` variable in your makefile to equal the corresponding object file name (for example, `jmemansi.o` or `jmemansi.obj` for `jmemansi.c`).

If you have plenty of (real or virtual) main memory, just use `jmemnobs.c`.

"Plenty" means about ten bytes for every pixel in the largest images you plan to process, so a lot of systems don't meet this criterion.

If yours doesn't, try `jmemansi.c` first. If that doesn't compile, you'll have to use `jmemname.c`; be sure to adjust `select_file_name()` for local conditions. You may also need to change `unlink()` to `remove()` in `close_backing_store()`.

Except with `jmemnobs.c` or `jmemmac.c`, you need to adjust the `DEFAULT_MAX_MEM` setting to a reasonable value for your system (either by adding a `#define` for `DEFAULT_MAX_MEM` to `jconfig.h`, or by adding a `-D` switch to the Makefile).

This value limits the amount of data space the program will attempt to allocate. Code and static data space isn't counted, so the actual memory needs for `cjpeg` or `djpeg` are typically 100 to 150Kb more than the max-memory setting. Larger max-memory settings reduce the amount of I/O needed to process a large image, but too large a value can result in "insufficient memory" failures. On most Unix machines (and other systems with virtual memory), just set `DEFAULT_MAX_MEM` to several million and forget it. At the other end of the spectrum, for MS-DOS machines you probably can't go much

above 300K to 400K. (On MS-DOS the value refers to conventional memory only. Extended/expanded memory is handled separately by jmemdos.c.)

BUILDING THE SOFTWARE

=====

Now you should be able to compile the software. Just say "make" (or whatever's necessary to start the compilation). Have a cup of coffee.

Here are some things that could go wrong:

If your compiler complains about undefined structures, you should be able to shut it up by putting "#define INCOMPLETE_TYPES_BROKEN" in jconfig.h.

If you have trouble with missing system include files or inclusion of the wrong ones, read jinclude.h. This shouldn't happen if you used configure or ckconfig.c to set up jconfig.h.

There are a fair number of routines that do not use all of their parameters; some compilers will issue warnings about this, which you can ignore. There are also a few configuration checks that may give "unreachable code" warnings. Any other warning deserves investigation.

If you don't have a getenv() library routine, define NO_GETENV.

Also see the system-specific hints, below.

TESTING THE SOFTWARE

=====

As a quick test of functionality we've included a small sample image in several forms:

testorig.jpg Starting point for the djpeg tests.

testimg.ppm The output of djpeg testorig.jpg

testimg.bmp The output of djpeg -bmp -colors 256 testorig.jpg

testimg.jpg The output of cjpeg testimg.ppm

testprog.jpg Progressive-mode equivalent of testorig.jpg.

testimgp.jpg The output of cjpeg -progressive -optimize testimg.ppm

(The first- and second-generation .jpg files aren't identical since the default compression parameters are lossy.) If you can generate duplicates of the testimg* files then you probably have working programs.

With most of the makefiles, "make test" will perform the necessary comparisons.

If you're using a makefile that doesn't provide the test option, run djpeg

and cjpeg by hand and compare the output files to testing* with whatever binary file comparison tool you have. The files should be bit-for-bit identical.

If the programs complain "MAX_ALLOC_CHUNK is wrong, please fix", then you need to reduce MAX_ALLOC_CHUNK to a value that fits in type size_t. Try adding "#define MAX_ALLOC_CHUNK 65520L" to jconfig.h. A less likely configuration error is "ALIGN_TYPE is wrong, please fix": defining ALIGN_TYPE as long should take care of that one.

If the cjpeg test run fails with "Missing Huffman code table entry", it's a good bet that you needed to define RIGHT_SHIFT_IS_UNSIGNED. Go back to the configuration step and run ckconfig.c. (This is a good plan for any other test failure, too.)

If you are using Unix (one-file) command line style on a non-Unix system, it's a good idea to check that binary I/O through stdin/stdout actually works. You should get the same results from "djpeg <testorig.jpg >out.ppm" as from "djpeg -outfile out.ppm testorig.jpg". Note that the makefiles all use the latter style and therefore do not exercise stdin/stdout! If this check fails, try recompiling with USE_SETMODE or USE_FDOPEN defined. If it still doesn't work, better use two-file style.

If you chose a memory manager other than jmemnobs.c, you should test that temporary-file usage works. Try "djpeg -bmp -colors 256 -max 0 testorig.jpg" and make sure its output matches testing.bmp. If you have any really large images handy, try compressing them with -optimize and/or decompressing with -colors 256 to make sure your DEFAULT_MAX_MEM setting is not too large.

NOTE: this is far from an exhaustive test of the JPEG software; some modules, such as 1-pass color quantization, are not exercised at all. It's just a quick test to give you some confidence that you haven't missed something major.

INSTALLING THE SOFTWARE

=====

Once you're done with the above steps, you can install the software by copying the executable files (cjpeg, djpeg, jpegtran, rdjpgcom, and wrjpgcom) to wherever you normally install programs. On Unix systems, you'll also want to put the man pages (cjpeg.1, djpeg.1, jpegtran.1, rdjpgcom.1, wrjpgcom.1) in the man-page directory. The pre-fab makefiles don't support this step since there's such a wide variety of installation procedures on different systems.

If you generated a Makefile with the "configure" script, you can just say
make install

to install the programs and their man pages into the standard places. (You'll probably need to be root to do this.) We recommend first saying
make -n install
to see where configure thought the files should go. You may need to edit the Makefile, particularly if your system's conventions for man page filenames don't match what configure expects.

If you want to install the IJG library itself, for use in compiling other programs besides ours, then you need to put the four include files
jpeglib.h jerror.h jconfig.h jmorecfg.h
into your include-file directory, and put the library file libjpeg.a (extension may vary depending on system) wherever library files go. If you generated a Makefile with "configure", it will do what it thinks is the right thing if you say
make install-lib

OPTIONAL STUFF

=====

Progress monitor:

If you like, you can #define PROGRESS_REPORT (in jconfig.h) to enable display of percent-done progress reports. The routine provided in cdjpeg.c merely prints percentages to stderr, but you can customize it to do something fancier.

Utah RLE file format support:

We distribute the software with support for RLE image files (Utah Raster Toolkit format) disabled, because the RLE support won't compile without the Utah library. If you have URT version 3.1 or later, you can enable RLE support as follows:

1. #define RLE_SUPPORTED in jconfig.h.
2. Add a -I option to CFLAGS in the Makefile for the directory containing the URT .h files (typically the "include" subdirectory of the URT distribution).
3. Add -L... -lrle to LDLIBS in the Makefile, where ... specifies the directory containing the URT "librle.a" file (typically the "lib" subdirectory of the URT distribution).

Support for 9-bit to 12-bit deep pixel data:

The IJG code currently allows 8, 9, 10, 11, or 12 bits sample data precision. (For color, this means 8 to 12 bits per channel, of course.) If you need to work with deeper than 8-bit data, you can compile the IJG code for 9-bit to 12-bit operation.

To do so:

1. In `jmorcfcg.h`, define `BITS_IN_JSAMPLE` as 9, 10, 11, or 12 rather than 8.
2. In `jconfig.h`, undefine `BMP_SUPPORTED`, `RLE_SUPPORTED`, and `TARGA_SUPPORTED`, because the code for those formats doesn't handle deeper than 8-bit data and won't even compile. (The PPM code does work, as explained below. The GIF code works too; it scales 8-bit GIF data to and from 12-bit depth automatically.)
3. Compile. Don't expect "make test" to pass, since the supplied test files are for 8-bit data.

Currently, 9-bit to 12-bit support does not work on 16-bit-int machines.

Run-time selection and conversion of data precision are currently not supported and may be added later.

Exception: The transcoding part (`jpegtran`) supports all settings in a single instance, since it operates on the level of DCT coefficients and not sample values.

The PPM reader (`rdppm.c`) can read deeper than 8-bit data from either text-format or binary-format PPM and PGM files. Binary-format PPM/PGM files which have a `maxval` greater than 255 are assumed to use 2 bytes per sample, MSB first (big-endian order). As of early 1995, 2-byte binary format is not officially supported by the `PBMPLUS` library, but it is expected that a future release of `PBMPLUS` will support it. Note that the PPM reader will read files of any `maxval` regardless of the `BITS_IN_JSAMPLE` setting; incoming data is automatically rescaled to `maxval=MAXJSAMPLE` as appropriate for the `cjpeg` bit depth.

The PPM writer (`wrppm.c`) will normally write 2-byte binary PPM or PGM format, `maxval=MAXJSAMPLE`, when compiled with `BITS_IN_JSAMPLE>8`. Since this format is not yet widely supported, you can disable it by compiling `wrppm.c` with `PPM_NORAWWORD` defined; then the data is scaled down to 8 bits to make a standard 1-byte/sample PPM or PGM file. (Yes, this means still another copy of `djpeg` to keep around. But hopefully you won't need it for very long. Poskanzer's supposed to get that new `PBMPLUS` release out Real Soon Now.)

Of course, if you are working with 9-bit to 12-bit data, you probably have it stored in some other, nonstandard format. In that case you'll probably want to write your own I/O modules to read and write your format.

Note:

The standard Huffman tables are only valid for 8-bit data precision. If you selected more than 8-bit data precision, `cjpeg` uses arithmetic coding by default. The Huffman encoder normally uses entropy optimization to compute usable tables for higher precision. Otherwise, you'll have to supply different default Huffman tables.

Removing code:

If you need to make a smaller version of the JPEG software, some optional functions can be removed at compile time. See the `xxx_SUPPORTED` #defines in `jconfig.h` and `jmorecfg.h`. If at all possible, we recommend that you leave in decoder support for all valid JPEG files, to ensure that you can read anyone's output. Taking out support for image file formats that you don't use is the most painless way to make the programs smaller. Another possibility is to remove some of the DCT methods: in particular, the "IFAST" method may not be enough faster than the others to be worth keeping on your machine. (If you do remove ISLOW or IFAST, be sure to redefine `JDCT_DEFAULT` or `JDCT_FASTEST` to a supported method, by adding a #define in `jconfig.h`.)

OPTIMIZATION

=====

Unless you own a Cray, you'll probably be interested in making the JPEG software go as fast as possible. This section covers some machine-dependent optimizations you may want to try. We suggest that before trying any of this, you first get the basic installation to pass the self-test step. Repeat the self-test after any optimization to make sure that you haven't broken anything.

The integer DCT routines perform a lot of multiplications. These multiplications must yield 32-bit results, but none of their input values are more than 16 bits wide. On many machines, notably the 680x0 and 80x86 CPUs, a 16x16=>32 bit multiply instruction is faster than a full 32x32=>32 bit multiply. Unfortunately there is no portable way to specify such a multiplication in C, but some compilers can generate one when you use the right combination of casts. See the `MULTIPLYxxx` macro definitions in `jdct.h`. If your compiler makes "int" be 32 bits and "short" be 16 bits, defining `SHORTxSHORT_32` is fairly likely to work. When experimenting with alternate definitions, be sure to test not only whether the code still works (use the self-test), but also whether it is actually faster --- on some compilers, alternate definitions may compute the right answer, yet be slower than the default. Timing `cjpeg` on a large PGM (grayscale) input file is the best way to check this, as the DCT will be the largest fraction of the runtime in that mode. (Note: some of the distributed compiler-specific `jconfig` files already contain #define switches to select appropriate `MULTIPLYxxx` definitions.)

If your machine has sufficiently fast floating point hardware, you may find that the float DCT method is faster than the integer DCT methods, even after tweaking the integer multiply macros. In that case you may want to make the float DCT be the default method. (The only objection to this is that float DCT results may vary slightly across machines.) To do that, add `#define JDCT_DEFAULT JDCT_FLOAT` to `jconfig.h`. Even if you don't change the default, you should redefine `JDCT_FASTEST`, which is the method selected by `djpeg`'s `-fast` switch. Don't forget to update the documentation files

(usage.txt and/or cjpeg.1, djpeg.1) to agree with what you've done.

If access to "short" arrays is slow on your machine, it may be a win to define type JCOEF as int rather than short. This will cost a good deal of memory though, particularly in some multi-pass modes, so don't do it unless you have memory to burn and short is REALLY slow.

If your compiler can compile function calls in-line, make sure the INLINE macro in jmorecfg.h is defined as the keyword that marks a function inline-able. Some compilers have a switch that tells the compiler to inline any function it thinks is profitable (e.g., -finline-functions for gcc). Enabling such a switch is likely to make the compiled code bigger but faster.

In general, it's worth trying the maximum optimization level of your compiler, and experimenting with any optional optimizations such as loop unrolling. (Unfortunately, far too many compilers have optimizer bugs ... be prepared to back off if the code fails self-test.) If you do any experimentation along these lines, please report the optimal settings to jpeg-info@jpegclub.org so we can mention them in future releases. Be sure to specify your machine and compiler version.

HINTS FOR SPECIFIC SYSTEMS

=====

We welcome reports on changes needed for systems not mentioned here. Submit 'em to jpeg-info@jpegclub.org. Also, if configure or ckconfig.c is wrong about how to configure the JPEG software for your system, please let us know.

Acorn RISC OS:

(Thanks to Simon Middleton for these hints on compiling with Desktop C.)
After renaming the files according to Acorn conventions, take a copy of makefile.ansi, change all occurrences of 'libjpeg.a' to 'libjpeg.o' and change these definitions as indicated:

```
CFLAGS= -throwback -IC: -Wn
LDLIBS=C:o.Stubs
SYSDEPMEM=jmemansi.o
LN=Link
AR=LibFile -c -o
```

Also add a new line '.c.o; \$(cc) \$< \$(cflags) -c -o \$@'. Remove the lines '\$(RM) libjpeg.o' and '\$(AR2) libjpeg.o' and the 'jconfig.h' dependency section.

Copy jconfig.txt to jconfig.h. Edit jconfig.h to define TWO_FILE_COMMANDLINE

and CHAR_IS_UNSIGNED.

Run the makefile using !AMU not !Make. If you want to use the 'clean' and 'test' makefile entries then you will have to fiddle with the syntax a bit and rename the test files.

Amiga:

SAS C 6.50 reportedly is too buggy to compile the IJG code properly. A patch to update to 6.51 is available from SAS or AmiNet FTP sites.

The supplied config files are set up to use jmemname.c as the memory manager, with temporary files being created on the device named by "JPEGTMP:".

Atari ST/STE/TT:

Copy the project files makcjpeg.st, makdjpeg.st, maktjpeg.st, and makljpeg.st to cjpeg.prj, djpeg.prj, jpegtran.prj, and libjpeg.prj respectively. The project files should work as-is with Pure C. For Turbo C, change library filenames "pc..." to "tc..." in each project file. Note that libjpeg.prj selects jmemansi.c as the recommended memory manager. You'll probably want to adjust the DEFAULT_MAX_MEM setting --- you want it to be a couple hundred K less than your normal free memory. Put "#define DEFAULT_MAX_MEM nnnn" into jconfig.h to do this.

To use the 68881/68882 coprocessor for the floating point DCT, add the compiler option "-8" to the project files and replace pcfltlib.lib with pc881lib.lib in cjpeg.prj and djpeg.prj. Or if you don't have a coprocessor, you may prefer to remove the float DCT code by undefining DCT_FLOAT_SUPPORTED in jmorecfg.h (since without a coprocessor, the float code will be too slow to be useful). In that case, you can delete pcfltlib.lib from the project files.

Note that you must make libjpeg.lib before making cjpeg.ttp, djpeg.ttp, or jpegtran.ttp. You'll have to perform the self-test by hand.

We haven't bothered to include project files for rdjpgcom and wrjpgcom. Those source files should just be compiled by themselves; they don't depend on the JPEG library. You can use the default.prj project file of the Pure C distribution to make the programs.

There is a bug in some older versions of the Turbo C library which causes the space used by temporary files created with "tmpfile()" not to be freed after an abnormal program exit. If you check your disk afterwards, you will find cluster chains that are allocated but not used by a file. This should not

happen in cjpeg/djpeg/jpegtran, since we enable a signal catcher to explicitly close temp files before exiting. But if you use the JPEG library with your own code, be sure to supply a signal catcher, or else use a different system-dependent memory manager.

Cray:

Should you be so fortunate as to be running JPEG on a Cray YMP, there is a compiler bug in old versions of Cray's Standard C (prior to 3.1). If you still have an old compiler, you'll need to insert a line reading

```
"#pragma novector" just before the loop
for (i = 1; i <= (int) htbl->bits[l]; i++)
    huffsize[p++] = (char) l;
```

in fix_huff_tbl (in V5beta1, line 204 of jchuff.c and line 176 of jdchuff.c).

[This bug may or may not still occur with the current IJG code, but it's probably a dead issue anyway...]

HP-UX:

If you have HP-UX 7.05 or later with the "software development" C compiler, you should run the compiler in ANSI mode. If using the configure script, say

```
./configure CC='cc -Aa'
```

(or -Ae if you prefer). If configuring by hand, use makefile.ansi and add "-Aa" to the CFLAGS line in the makefile.

If you have a pre-7.05 system, or if you are using the non-ANSI C compiler delivered with a minimum HP-UX system, then you must use makefile.unix (and do NOT add -Aa); or just run configure without the CC option.

On HP 9000 series 800 machines, the HP C compiler is buggy in revisions prior to A.08.07. If you get complaints about "not a typedef name", you'll have to use makefile.unix, or run configure without the CC option.

Macintosh, generic comments:

The supplied user-interface files (cjpeg.c, djpeg.c, etc) are set up to provide a Unix-style command line interface. You can use this interface on the Mac by means of the ccommand() library routine provided by Metrowerks CodeWarrior or Think C. This is only appropriate for testing the library, however; to make a user-friendly equivalent of cjpeg/djpeg you'd really want to develop a Mac-style user interface. There isn't a complete example available at the moment, but there are some helpful starting points:

1. Sam Bushell's free "To JPEG" applet provides drag-and-drop conversion to JPEG under System 7 and later. This only illustrates how to use the

compression half of the library, but it does a very nice job of that part. The CodeWarrior source code is available from <http://www.pobox.com/~jsam>. 2. Jim Brunner prepared a Mac-style user interface for both compression and decompression. Unfortunately, it hasn't been updated since IJG v4, and the library's API has changed considerably since then. Still it may be of some help, particularly as a guide to compiling the IJG code under Think C. Jim's code is available from the Info-Mac archives, at sumex-aim.stanford.edu or mirrors thereof; see file `/info-mac/dev/src/jpeg-convert-c.hqx`.

`jmemmac.c` is the recommended memory manager back end for Macintosh. It uses `NewPtr/DisposePtr` instead of `malloc/free`, and has a Mac-specific implementation of `jpeg_mem_available()`. It also creates temporary files that follow Mac conventions. (That part of the code relies on System-7-or-later OS functions. See the comments in `jmemmac.c` if you need to run it on System 6.) NOTE that `USE_MAC_MEMMGR` must be defined in `jconfig.h` to use `jmemmac.c`.

You can also use `jmemnobs.c`, if you don't care about handling images larger than available memory. If you use any memory manager back end other than `jmemmac.c`, we recommend replacing "malloc" and "free" by "NewPtr" and "DisposePtr", because Mac C libraries often have peculiar implementations of `malloc/free`. (For instance, `free()` may not return the freed space to the Mac Memory Manager. This is undesirable for the IJG code because `jmemmgr.c` already clumps space requests.)

Macintosh, Metrowerks CodeWarrior:

The Unix-command-line-style interface can be used by defining `USE_CCOMMAND`. You'll also need to define `TWO_FILE_COMMANDLINE` to avoid `stdin/stdout`. This means that when using the `cjpeg/djpeg` programs, you'll have to type the input and output file names in the "Arguments" text-edit box, rather than using the file radio buttons. (Perhaps `USE_FDOPEN` or `USE_SETMODE` would eliminate the problem, but I haven't heard from anyone who's tried it.)

On 680x0 Macs, Metrowerks defines type "double" as a 10-byte IEEE extended float. `jmemmgr.c` won't like this: it wants `sizeof(ALIGN_TYPE)` to be a power of 2. Add `#define ALIGN_TYPE long` to `jconfig.h` to eliminate the complaint.

The supplied configuration file `jconfig.mac` can be used for your `jconfig.h`; it includes all the recommended symbol definitions. If you have AppleScript installed, you can run the supplied script `makeproj.mac` to create CodeWarrior project files for the library and the testbed applications, then build the library and applications. (Thanks to Dan Sears and Don Agro for this nifty hack, which saves us from trying to maintain CodeWarrior project files as part of the IJG distribution...)

Macintosh, Think C:

The documentation in Jim Brunner's "JPEG Convert" source code (see above) includes detailed build instructions for Think C; it's probably somewhat out of date for the current release, but may be helpful.

If you want to build the minimal command line version, proceed as follows. You'll have to prepare project files for the programs; we don't include any in the distribution since they are not text files. Use the file lists in any of the supplied makefiles as a guide. Also add the ANSI and Unix C libraries in a separate segment. You may need to divide the JPEG files into more than one segment; we recommend dividing compression and decompression modules. Define `USE_CCOMMAND` in `jconfig.h` so that the `ccommand()` routine is called. You must also define `TWO_FILE_COMMANDLINE` because `stdin/stdout` don't handle binary data correctly.

On 680x0 Macs, Think C defines type "double" as a 12-byte IEEE extended float. `jmemmgr.c` won't like this: it wants `sizeof(ALIGN_TYPE)` to be a power of 2. Add `#define ALIGN_TYPE long` to `jconfig.h` to eliminate the complaint.

`jconfig.mac` should work as a `jconfig.h` configuration file for Think C, but the `makeproj.mac` AppleScript script is specific to CodeWarrior. Sorry.

MIPS R3000:

MIPS's `cc` version 1.31 has a rather nasty optimization bug. Don't use `-O` if you have that compiler version. (Use `"cc -V"` to check the version.) Note that the R3000 chip is found in workstations from DEC and others.

MS-DOS, generic comments for 16-bit compilers:

The IJG code is designed to work well in 80x86 "small" or "medium" memory models (i.e., data pointers are 16 bits unless explicitly declared "far"; code pointers can be either size). You may be able to use small model to compile `cjpeg` or `djpeg` by itself, but you will probably have to use medium model for any larger application. This won't make much difference in performance. You *will* take a noticeable performance hit if you use a large-data memory model, and you should avoid "huge" model if at all possible. Be sure that `NEED_FAR_POINTERS` is defined in `jconfig.h` if you use a small-data memory model; be sure it is NOT defined if you use a large-data model. (The supplied makefiles and `jconfig` files for Borland and Microsoft C compile in medium model and define `NEED_FAR_POINTERS`.)

The DOS-specific memory manager, `jmemdos.c`, should be used if possible. It needs some assembly-code routines which are in `jmemdosa.asm`; make sure your makefile assembles that file and includes it in the library. If you don't have a suitable assembler, you can get pre-assembled object files for

jmemdosa by FTP from ftp.uu.net:/graphics/jpeg/jdosaobj.zip. (DOS-oriented distributions of the IJG source code often include these object files.)

When using jmemdos.c, jconfig.h must define USE_MSDOS_MEMMGR and must set MAX_ALLOC_CHUNK to less than 64K (65520L is a typical value). If your C library's far-heap malloc() can't allocate blocks that large, reduce MAX_ALLOC_CHUNK to whatever it can handle.

If you can't use jmemdos.c for some reason --- for example, because you don't have an assembler to assemble jmemdosa.asm --- you'll have to fall back to jmemansi.c or jmemname.c. You'll probably still need to set MAX_ALLOC_CHUNK in jconfig.h, because most DOS C libraries won't malloc() more than 64K at a time. IMPORTANT: if you use jmemansi.c or jmemname.c, you will have to compile in a large-data memory model in order to get the right stdio library. Too bad.

wrjpgcom needs to be compiled in large model, because it malloc()s a 64KB work area to hold the comment text. If your C library's malloc can't handle that, reduce MAX_COM_LENGTH as necessary in wrjpgcom.c.

Most MS-DOS compilers treat stdin/stdout as text files, so you must use two-file command line style. But if your compiler has either fdopen() or setmode(), you can use one-file style if you like. To do this, define USE_SETMODE or USE_FDOPEN so that stdin/stdout will be set to binary mode. (USE_SETMODE seems to work with more DOS compilers than USE_FDOPEN.) You should test that I/O through stdin/stdout produces the same results as I/O to explicitly named files... the "make test" procedures in the supplied makefiles do NOT use stdin/stdout.

MS-DOS, generic comments for 32-bit compilers:

None of the above comments about memory models apply if you are using a 32-bit flat-memory-space environment, such as DJGPP or Watcom C. (And you should use one if you have it, as performance will be much better than 8086-compatible code!) For flat-memory-space compilers, do NOT define NEED_FAR_POINTERS, and do NOT use jmemdos.c. Use jmemnobs.c if the environment supplies adequate virtual memory, otherwise use jmemansi.c or jmemname.c.

You'll still need to be careful about binary I/O through stdin/stdout. See the last paragraph of the previous section.

MS-DOS, Borland C:

Be sure to convert all the source files to DOS text format (CR/LF newlines). Although Borland C will often work OK with unmodified Unix (LF newlines)

source files, sometimes it will give bogus compile errors.

"Illegal character '#'" is the most common such error. (This is true with Borland C 3.1, but perhaps is fixed in newer releases.)

If you want one-file command line style, just undefine TWO_FILE_COMMANDLINE. jconfig.bcc already includes #define USE_SETMODE to make this work. (fdopen does not work correctly.)

MS-DOS, Microsoft C:

makefile.mc6 works with Microsoft C, DOS Visual C++, etc. It should only be used if you want to build a 16-bit (small or medium memory model) program.

If you want one-file command line style, just undefine TWO_FILE_COMMANDLINE. jconfig.mc6 already includes #define USE_SETMODE to make this work. (fdopen does not work correctly.)

Note that this makefile assumes that the working copy of itself is called "makefile". If you want to call it something else, say "makefile.mak", be sure to adjust the dependency line that reads "\$(RFILE) : makefile". Otherwise the make will fail because it doesn't know how to create "makefile". Worse, some releases of Microsoft's make utilities give an incorrect error message in this situation.

Old versions of MS C fail with an "out of macro expansion space" error because they can't cope with the macro TRACEMS8 (defined in jerror.h). If this happens to you, the easiest solution is to change TRACEMS8 to expand to nothing. You'll lose the ability to dump out JPEG coefficient tables with djpeg -debug -debug, but at least you can compile.

Original MS C 6.0 is very buggy; it compiles incorrect code unless you turn off optimization entirely (remove -O from CFLAGS). 6.00A is better, but it still generates bad code if you enable loop optimizations (-Ol or -Ox).

MS C 8.0 crashes when compiling jquant1.c with optimization switch /Oo ... which is on by default. To work around this bug, compile that one file with /Oo-.

Microsoft Windows (all versions), generic comments:

Some Windows system include files define typedef boolean as "unsigned char". The IJG code also defines typedef boolean, but we make it an "enum" by default. This doesn't affect the IJG programs because we don't import those Windows include files. But if you use the JPEG library in your own program, and some of your program's files import one definition of boolean while some import the other, you can get all sorts of mysterious problems. A good preventive step

is to make the IJG library use "unsigned char" for boolean. To do that, add something like this to your jconfig.h file:

```
/* Define "boolean" as unsigned char, not enum, per Windows custom */
#ifndef __RPCNDR_H__ /* don't conflict if rpcndr.h already read */
typedef unsigned char boolean;
#endif
#ifndef FALSE /* in case these macros already exist */
#define FALSE 0 /* values of boolean */
#endif
#ifndef TRUE
#define TRUE 1
#endif
#define HAVE_BOOLEAN /* prevent jmorecfg.h from redefining it */
(This is already in jconfig.vc, by the way.)
```

windef.h contains the declarations

```
#define far
#define FAR far
```

Since jmorecfg.h tries to define FAR as empty, you may get a compiler warning if you include both jpeglib.h and windef.h (which windows.h includes). To suppress the warning, you can put "#ifndef FAR" / "#endif" around the line "#define FAR" in jmorecfg.h.

(Something like this is already in jmorecfg.h, by the way.)

When using the library in a Windows application, you will almost certainly want to modify or replace the error handler module jerror.c, since our default error handler does a couple of inappropriate things:

1. it tries to write error and warning messages on stderr;
2. in event of a fatal error, it exits by calling exit().

A simple stopgap solution for problem 1 is to replace the line

```
fprintf(stderr, "%s\n", buffer);
```

(in output_message in jerror.c) with

```
MessageBox(GetActiveWindow(),buffer,"JPEG Error",MB_OK|MB_ICONERROR);
```

It's highly recommended that you at least do that much, since otherwise error messages will disappear into nowhere. (Beginning with IJG v6b, this code is already present in jerror.c; just define USE_WINDOWS_MESSAGEBOX in jconfig.h to enable it.)

The proper solution for problem 2 is to return control to your calling application after a library error. This can be done with the setjmp/longjmp technique discussed in libjpeg.txt and illustrated in example.c. (NOTE: some older Windows C compilers provide versions of setjmp/longjmp that don't actually work under Windows. You may need to use the Windows system functions Catch and Throw instead.)

The recommended memory manager under Windows is jmemnobs.c; in other words, let Windows do any virtual memory management needed. You should NOT use

jmemdos.c nor jmemdosa.asm under Windows.

For Windows 3.1, we recommend compiling in medium or large memory model; for newer Windows versions, use a 32-bit flat memory model. (See the MS-DOS sections above for more info about memory models.) In the 16-bit memory models only, you'll need to put

```
#define MAX_ALLOC_CHUNK 65520L /* Maximum request to malloc() */
```

into jconfig.h to limit allocation chunks to 64Kb. (Without that, you'd have to use huge memory model, which slows things down unnecessarily.) jmemnobs.c works without modification in large or flat memory models, but to use medium model, you need to modify its jpeg_get_large and jpeg_free_large routines to allocate far memory. In any case, you might like to replace its calls to malloc and free with direct calls on Windows memory allocation functions.

You may also want to modify jdatasrc.c and jdatadst.c to use Windows file operations rather than fread/fwrite. This is only necessary if your C compiler doesn't provide a competent implementation of C stdio functions.

You might want to tweak the RGB_xxx macros in jmorecfg.h so that the library will accept or deliver color pixels in BGR sample order, not RGB; BGR order is usually more convenient under Windows. Note that this change will break the sample applications cjpeg/djpeg, but the library itself works fine.

Many people want to convert the IJG library into a DLL. This is reasonably straightforward, but watch out for the following:

1. Don't try to compile as a DLL in small or medium memory model; use large model, or even better, 32-bit flat model. Many places in the IJG code assume the address of a local variable is an ordinary (not FAR) pointer; that isn't true in a medium-model DLL.

2. Microsoft C cannot pass file pointers between applications and DLLs. (See Microsoft Knowledge Base, PSS ID Number Q50336.) So jdatasrc.c and jdatadst.c don't work if you open a file in your application and then pass the pointer to the DLL. One workaround is to make jdatasrc.c/jdatadst.c part of your main application rather than part of the DLL.

3. You'll probably need to modify the macros GLOBAL() and EXTERN() to attach suitable linkage keywords to the exported routine names. Similarly, you'll want to modify METHODDEF() and JMETHOD() to ensure function pointers are declared in a way that lets application routines be called back through the function pointers. These macros are in jmorecfg.h. Typical definitions for a 16-bit DLL are:

```
#define GLOBAL(type) type _far _pascal _loadds _export
#define EXTERN(type) extern type _far _pascal _loadds
#define METHODDEF(type) static type _far _pascal
```

```
#define JMETHOD(type,methodname,arglist) \
    type (_far _pascal *methodname) arglist
```

For a 32-bit DLL you may want something like

```
#define GLOBAL(type) __declspec(dllexport) type
#define EXTERN(type) extern __declspec(dllexport) type
```

Although not all the GLOBAL routines are actually intended to be called by the application, the performance cost of making them all DLL entry points is negligible.

The unmodified IJG library presents a very C-specific application interface, so the resulting DLL is only usable from C or C++ applications. There has been some talk of writing wrapper code that would present a simpler interface usable from other languages, such as Visual Basic. This is on our to-do list but hasn't been very high priority --- any volunteers out there?

Microsoft Windows, Borland C:

The provided jconfig.bcc should work OK in a 32-bit Windows environment, but you'll need to tweak it in a 16-bit environment (you'd need to define NEED_FAR_POINTERS and MAX_ALLOC_CHUNK). Beware that makefile.bcc will need alteration if you want to use it for Windows --- in particular, you should use jmemnobs.c not jmemdos.c under Windows.

Borland C++ 4.5 fails with an internal compiler error when trying to compile jdmerge.c in 32-bit mode. If enough people complain, perhaps Borland will fix it. In the meantime, the simplest known workaround is to add a redundant definition of the variable range_limit in h2v1_merged_upsample(), at the head of the block that handles odd image width (about line 268 in v6 jdmerge.c):

```
/* If image width is odd, do the last output column separately */
if (cinfo->output_width & 1) {
    register JSAMPLE * range_limit = cinfo->sample_range_limit; /* ADD THIS */
    cb = GETJSAMPLE(*inptr1);
```

Pretty bizarre, especially since the very similar routine h2v2_merged_upsample doesn't trigger the bug.

Recent reports suggest that this bug does not occur with "bcc32a" (the Pentium-optimized version of the compiler).

Another report from a user of Borland C 4.5 was that incorrect code (leading to a color shift in processed images) was produced if any of the following optimization switch combinations were used:

```
-Ot -Og
-Ot -Op
-Ot -Om
```

So try backing off on optimization if you see such a problem. (Are there several different releases all numbered "4.5"??)

Microsoft Windows, Microsoft Visual C++:

jconfig.vc should work OK with any Microsoft compiler for a 32-bit memory model. makefile.vc is intended for command-line use. (If you are using the Developer Studio environment, you may prefer the DevStudio project files; see below.)

IJG JPEG 7 adds extern "C" to jpeglib.h. This avoids the need to put extern "C" { ... } around #include "jpeglib.h" in your C++ application. You can also force VC++ to treat the library as C++ code by renaming all the *.c files to *.cpp (and adjusting the makefile to match). In this case you also need to define the symbol DONT_USE_EXTERN_C in the configuration to prevent jpeglib.h from using extern "C".

Microsoft Windows, Microsoft Visual C++ 6 Developer Studio:

We include makefiles that should work as project files in Developer Studio 6.0 or later. There is a library makefile that builds the IJG library as a static Win32 library, and application makefiles that build the sample applications as Win32 console applications. (Even if you only want the library, we recommend building the applications so that you can run the self-test.)

To use:

1. Open the command prompt, change to the main directory and execute the command line
NMAKE /f makefile.vs setup-vc6
This will move jconfig.vc to jconfig.h and makefiles to project files.
(Note that the renaming is critical!)
2. Open the workspace file jpeg.dsw, build the library project.
(If you are using Developer Studio more recent than 6.0, you'll probably get a message saying that the project files are being updated.)
3. Open the workspace file apps.dsw, build the application projects.
4. To perform the self-test, execute the command line
NMAKE /f makefile.vs test-build
5. Move the application .exe files from `app`\Release to an appropriate location on your path.

Microsoft Windows, Visual Studio 2017 (v15):

We include makefiles that should work as project files in Visual Studio 2017 (v15) or later. There is a library makefile that builds the IJG library as a static Win32 library, and application makefiles that build the sample applications as Win32 console applications. (Even if you only want the library, we recommend building the applications so that you can run the self-test.)

To use:

1. Open the Developer Command Prompt, change to the main directory and execute the command line
NMAKE /f makefile.vs setup-v15
This will move jconfig.vc to jconfig.h and makefiles to project files.
(Note that the renaming is critical!)
2. Open the solution file jpeg.sln, build the library project.
 - a) If you are using Visual Studio more recent than 2017 (v15), you'll probably get a message saying that the project files are being updated.
 - b) If necessary, open the project properties and adapt the Windows Target Platform Version in the Configuration Properties, General section; we support the latest version at the time of release.
3. Open the solution file apps.sln, build the application projects.
4. To perform the self-test, execute the command line
NMAKE /f makefile.vs test-build
5. Move the application .exe files from `app`\Release to an appropriate location on your path.

OS/2, Borland C++:

Watch out for optimization bugs in older Borland compilers; you may need to back off the optimization switch settings. See the comments in makefile.bcc.

SGI:

On some SGI systems, you may need to set "AR2= ar -ts" in the Makefile. If you are using configure, you can do this by saying
./configure RANLIB='ar -ts'
This change is not needed on all SGIs. Use it only if the make fails at the stage of linking the completed programs.

On the MIPS R4000 architecture (Indy, etc.), the compiler option "-mips2" reportedly speeds up the float DCT method substantially, enough to make it faster than the default int method (but still slower than the fast int method). If you use -mips2, you may want to alter the default DCT method to be float. To do this, put "#define JDCT_DEFAULT JDCT_FLOAT" in jconfig.h.

VMS:

On an Alpha/VMS system with MMS, be sure to use the "/Marco=Alpha=1" qualifier with MMS when building the JPEG package.

VAX/VMS v5.5-1 may have problems with the test step of the build procedure reporting differences when it compares the original and test images. If the error points to the last block of the files, it is most likely bogus and may be safely ignored. It seems to be because the files are Stream_LF and Backup/Compare has difficulty with the (presumably) null padded files. This problem was not observed on VAX/VMS v6.1 or AXP/VMS v6.1.

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/install.txt

No license file was found, but licenses were detected in source scan.

/*

* djpeg.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2009-2015 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains a command-line user interface for the JPEG decompressor.

* It should work on any system with Unix- or MS-DOS-style command lines.

*

* Two different command line styles are permitted, depending on the

* compile-time switch TWO_FILE_COMMANDLINE:

* djpeg [options] inputfile outputfile

* djpeg [options] [inputfile]

* In the second style, output is always to standard output, which you'd

* normally redirect to a file or pipe to some other program. Input is

* either from a named file or from standard input (typically redirected).

* The second style is convenient on Unix but is unhelpful on systems that

* don't support pipes. Also, you MUST use the first style if your system

* doesn't do binary I/O to stdin/stdout.

* To simplify script writing, the "-outfile" switch is provided. The syntax

* djpeg [options] -outfile outputfile inputfile

* works regardless of which command line style is used.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/djpeg.c

No license file was found, but licenses were detected in source scan.

/*

* jdatasrc.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2009-2015 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

```
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains decompression data source routines for the case of
* reading JPEG data from memory or from a file (or any stdio stream).
* While these routines are sufficient for most applications,
* some will want to use a different source manager.
* IMPORTANT: we assume that fread() will correctly transcribe an array of
* JOCTETs from 8-bit-wide elements on external storage. If char is wider
* than 8 bits on your machine, you may need to do some tweaking.
*/
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdatasrc.c
No license file was found, but licenses were detected in source scan.
```

IJG JPEG LIBRARY: CODING RULES

Copyright (C) 1991-1996, Thomas G. Lane.

This file is part of the Independent JPEG Group's software.

For conditions of distribution and use, see the accompanying README file.

Since numerous people will be contributing code and bug fixes, it's important to establish a common coding style. The goal of using similar coding styles is much more important than the details of just what that style is.

In general we follow the recommendations of "Recommended C Style and Coding Standards" revision 6.1 (Cannon et al. as modified by Spencer, Keppel and Brader). This document is available in the IJG FTP archive (see `jpeg/doc/cstyle.ms.tbl.Z`, or `cstyle.txt.Z` for those without `nroff/tbl`).

Block comments should be laid out thusly:

```
/*
* Block comments in this style.
*/
```

We indent statements in K&R style, e.g.,

```
if (test) {
    then-part;
} else {
    else-part;
}
```

with two spaces per indentation level. (This indentation convention is handled automatically by GNU Emacs and many other text editors.)

Multi-word names should be written in lower case with underscores, e.g., `multi_word_name` (not `multiWordName`). Preprocessor symbols and enum constants

are similar but upper case (MULTI_WORD_NAME). Names should be unique within the first fifteen characters. (On some older systems, global names must be unique within six characters. We accommodate this without cluttering the source code by using macros to substitute shorter names.)

We use function prototypes everywhere; we rely on automatic source code transformation to feed prototype-less C compilers. Transformation is done by the simple and portable tool 'ansi2knr.c' (courtesy of Ghostscript). ansi2knr is not very bright, so it imposes a format requirement on function declarations: the function name MUST BEGIN IN COLUMN 1. Thus all functions should be written in the following style:

```
LOCAL(int *)
function_name (int a, char *b)
{
    code...
}
```

Note that each function definition must begin with GLOBAL(type), LOCAL(type), or METHODDEF(type). These macros expand to "static type" or just "type" as appropriate. They provide a readable indication of the routine's usage and can readily be changed for special needs. (For instance, special linkage keywords can be inserted for use in Windows DLLs.)

ansi2knr does not transform method declarations (function pointers in structs). We handle these with a macro JMETHOD, defined as

```
#ifdef HAVE_PROTOTYPES
#define JMETHOD(type,methodname,arglist) type (*methodname) arglist
#else
#define JMETHOD(type,methodname,arglist) type (*methodname) ()
#endif
```

which is used like this:

```
struct function_pointers {
    JMETHOD(void, init_entropy_encoder, (int somearg, jparms *jp));
    JMETHOD(void, term_entropy_encoder, (void));
};
```

Note the set of parentheses surrounding the parameter list.

A similar solution is used for forward and external function declarations (see the EXTERN and JPP macros).

If the code is to work on non-ANSI compilers, we cannot rely on a prototype declaration to coerce actual parameters into the right types. Therefore, use explicit casts on actual parameters whenever the actual parameter type is not identical to the formal parameter. Beware of implicit conversions to "int".

It seems there are some non-ANSI compilers in which the sizeof() operator is defined to return int, yet size_t is defined as long. Needless to say,

this is brain-damaged. Always use the SIZEOF() macro in place of sizeof(), so that the result is guaranteed to be of type size_t.

The JPEG library is intended to be used within larger programs. Furthermore, we want it to be reentrant so that it can be used by applications that process multiple images concurrently. The following rules support these requirements:

1. Avoid direct use of file I/O, "malloc", error report printouts, etc; pass these through the common routines provided.
2. Minimize global namespace pollution. Functions should be declared static wherever possible. (Note that our method-based calling conventions help this a lot: in many modules only the initialization function will ever need to be called directly, so only that function need be externally visible.) All global function names should begin with "jpeg_", and should have an abbreviated name (unique in the first six characters) substituted by macro when NEED_SHORT_EXTERNAL_NAMES is set.
3. Don't use global variables; anything that must be used in another module should be in the common data structures.
4. Don't use static variables except for read-only constant tables. Variables that should be private to a module can be placed into private structures (see the system architecture document, structure.txt).
5. Source file names should begin with "j" for files that are part of the library proper; source files that are not part of the library, such as cjpeg.c and djpeg.c, do not begin with "j". Keep source file names to eight characters (plus ".c" or ".h", etc) to make life easy for MS-DOSers. Keep compression and decompression code in separate source files --- some applications may want only one half of the library.

Note: these rules (particularly #4) are not followed religiously in the modules that are used in cjpeg/djpeg but are not part of the JPEG library proper. Those modules are not really intended to be used in other applications.

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/coderules.txt

No license file was found, but licenses were detected in source scan.

/*

* jcomapi.c

*

* Copyright (C) 1994-1997, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*
* This file contains application interface routines that are used for both
* compression and decompression.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcomapi.c
No license file was found, but licenses were detected in source scan.

/*
* jconfig.txt
*
* Copyright (C) 1991-1994, Thomas G. Lane.
* Modified 2009-2013 by Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file documents the configuration options that are required to
* customize the JPEG software for a particular system.
*
* The actual configuration options for a particular installation are stored
* in jconfig.h. On many machines, jconfig.h can be generated automatically
* or copied from one of the "canned" jconfig files that we supply. But if
* you need to generate a jconfig.h file by hand, this file tells you how.
*
* DO NOT EDIT THIS FILE --- IT WON'T ACCOMPLISH ANYTHING.
* EDIT A COPY NAMED JCONFIG.H.
*/

/*
* These symbols indicate the properties of your machine or compiler.
* #define the symbol if yes, #undef it if no.
*/

/* Does your compiler support function prototypes?
* (If not, you also need to use ansi2knr, see install.txt)
*/
#define HAVE_PROTOTYPES

/* Does your compiler support the declaration "unsigned char" ?
* How about "unsigned short" ?
*/
#define HAVE_UNSIGNED_CHAR
#define HAVE_UNSIGNED_SHORT

/* Define "void" as "char" if your compiler doesn't know about type void.
* NOTE: be sure to define void such that "void *" represents the most general

```

* pointer type, e.g., that returned by malloc().
*/
/* #define void char */

/* Define "const" as empty if your compiler doesn't know the "const" keyword.
*/
/* #define const */

/* Define this if an ordinary "char" type is unsigned.
* If you're not sure, leaving it undefined will work at some cost in speed.
* If you defined HAVE_UNSIGNED_CHAR then the speed difference is minimal.
*/
#undef CHAR_IS_UNSIGNED

/* Define this if your system has an ANSI-conforming <stddef.h> file.
*/
#define HAVE_STDDEF_H

/* Define this if your system has an ANSI-conforming <stdlib.h> file.
*/
#define HAVE_STDLIB_H

/* Define this if your system does not have an ANSI/SysV <string.h>,
* but does have a BSD-style <strings.h>.
*/
#undef NEED_BSD_STRINGS

/* Define this if your system does not provide typedef size_t in any of the
* ANSI-standard places (stddef.h, stdlib.h, or stdio.h), but places it in
* <sys/types.h> instead.
*/
#undef NEED_SYS_TYPES_H

/* For 80x86 machines, you need to define NEED_FAR_POINTERS,
* unless you are using a large-data memory model or 80386 flat-memory mode.
* On less brain-damaged CPUs this symbol must not be defined.
* (Defining this symbol causes large data structures to be referenced through
* "far" pointers and to be allocated with a special version of malloc.)
*/
#undef NEED_FAR_POINTERS

/* Define this if your linker needs global names to be unique in less
* than the first 15 characters.
*/
#undef NEED_SHORT_EXTERNAL_NAMES

/* Although a real ANSI C compiler can deal perfectly well with pointers to
* unspecified structures (see "incomplete types" in the spec), a few pre-ANSI

```



```

* and pseudo-ANSI compilers get confused. To keep one of these bozos happy,
* define INCOMPLETE_TYPES_BROKEN. This is not recommended unless you
* actually get "missing structure definition" warnings or errors while
* compiling the JPEG code.
*/
#undef INCOMPLETE_TYPES_BROKEN

/* Define "boolean" as unsigned char, not enum, on Windows systems.
*/
#ifdef _WIN32
#ifdef __RPCNDR_H__ /* don't conflict if rpcndr.h already read */
typedef unsigned char boolean;
#endif
#ifndef FALSE /* in case these macros already exist */
#define FALSE 0 /* values of boolean */
#endif
#ifndef TRUE
#define TRUE 1
#endif
#define HAVE_BOOLEAN /* prevent jmorecfg.h from redefining it */
#endif

/*
* The following options affect code selection within the JPEG library,
* but they don't need to be visible to applications using the library.
* To minimize application namespace pollution, the symbols won't be
* defined unless JPEG_INTERNALS has been defined.
*/

#ifdef JPEG_INTERNALS

/* Define this if your compiler implements ">>" on signed values as a logical
* (unsigned) shift; leave it undefined if ">>" is a signed (arithmetic) shift,
* which is the normal and rational definition.
*/
#undef RIGHT_SHIFT_IS_UNSIGNED

#endif /* JPEG_INTERNALS */

/*
* The remaining options do not affect the JPEG library proper,
* but only the sample applications cjpeg/djpeg (see cjpeg.c, djpeg.c).
* Other applications can ignore these.
*/

```

```

#ifdef JPEG_CJPEG_DJPEG

/* These defines indicate which image (non-JPEG) file formats are allowed. */

#define BMP_SUPPORTED /* BMP image file format */
#define GIF_SUPPORTED /* GIF image file format */
#define PPM_SUPPORTED /* PBMPLUS PPM/PGM image file format */
#undef RLE_SUPPORTED /* Utah RLE image file format */
#define TARGA_SUPPORTED /* Targa image file format */

/* Define this if you want to name both input and output files on the command
 * line, rather than using stdout and optionally stdin. You MUST do this if
 * your system can't cope with binary I/O to stdin/stdout. See comments at
 * head of cjpeg.c or djpeg.c.
 */
#undef TWO_FILE_COMMANDLINE

/* Define this if your system needs explicit cleanup of temporary files.
 * This is crucial under MS-DOS, where the temporary "files" may be areas
 * of extended memory; on most other systems it's not as important.
 */
#undef NEED_SIGNAL_CATCHER

/* By default, we open image files with fopen(...,"rb") or fopen(...,"wb").
 * This is necessary on systems that distinguish text files from binary files,
 * and is harmless on most systems that don't. If you have one of the rare
 * systems that complains about the "b" spec, define this symbol.
 */
#undef DONT_USE_B_MODE

/* Define this if you want percent-done progress reports from cjpeg/djpeg.
 */
#undef PROGRESS_REPORT

#endif /* JPEG_CJPEG_DJPEG */

```

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jconfig.txt

No license file was found, but licenses were detected in source scan.

/*

* wrjpgcom.c

*/

* Copyright (C) 1994-1997, Thomas G. Lane.

* Modified 2015-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

```
*
* This file contains a very simple stand-alone application that inserts
* user-supplied text as a COM (comment) marker in a JFIF file.
* This may be useful as an example of the minimum logic needed to parse
* JPEG markers.
*/
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/wrjpgcom.c
No license file was found, but licenses were detected in source scan.
```

```
/*
* rdcolmap.c
*
* Copyright (C) 1994-1996, Thomas G. Lane.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file implements djpeg's "-map file" switch. It reads a source image
* and constructs a colormap to be supplied to the JPEG decompressor.
*
* Currently, these file formats are supported for the map file:
* GIF: the contents of the GIF's global colormap are used.
* PPM (either text or raw flavor): the entire file is read and
* each unique pixel value is entered in the map.
* Note that reading a large PPM file will be horrendously slow.
* Typically, a PPM-format map file should contain just one pixel
* of each desired color. Such a file can be extracted from an
* ordinary image PPM file with ppmtomap(1).
*
* Rescaling a PPM that has a maxval unequal to MAXJSAMPLE is not
* currently implemented.
*/
```

```
/* Portions of this code are based on the PBMPLUS library, which is:
**
** Copyright (C) 1988 by Jef Poskanzer.
**
** Permission to use, copy, modify, and distribute this software and its
** documentation for any purpose and without fee is hereby granted, provided
** that the above copyright notice appear in all copies and that both that
** copyright notice and this permission notice appear in supporting
** documentation. This software is provided "as is" without express or
** implied warranty.
*/
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdcolmap.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * jcrepct.c
 *
 * Copyright (C) 1994-1996, Thomas G. Lane.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains the compression preprocessing controller.
 * This controller manages the color conversion, downsampling,
 * and edge expansion steps.
 *
 * Most of the complexity here is associated with buffering input rows
 * as required by the downsampler. See the comments at the head of
 * jcsample.c for the downsampler's needs.
 */
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcrepct.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * jcapimin.c
 *
 * Copyright (C) 1994-1998, Thomas G. Lane.
 * Modified 2003-2010 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains application interface code for the compression half
 * of the JPEG library. These are the "minimum" API routines that may be
 * needed in either the normal full-compression case or the transcoding-only
 * case.
 *
 * Most of the routines intended to be called directly by an application
 * are in this file or in jcapistd.c. But also see jcparam.c for
 * parameter-setup helper routines, jcomapi.c for routines shared by
 * compression and decompression, and jctrans.c for the transcoding case.
 */
```

Found in path(s):

```
*/opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcapimin.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * rdjpgcom.c
 *
```

- * Copyright (C) 1994-1997, Thomas G. Lane.
- * Modified 2009 by Bill Allombert, Guido Vollbeding.
- * This file is part of the Independent JPEG Group's software.
- * For conditions of distribution and use, see the accompanying README file.
- *
- * This file contains a very simple stand-alone application that displays
- * the text in COM (comment) markers in a JFIF file.
- * This may be useful as an example of the minimum logic needed to parse
- * JPEG markers.
- */

Found in path(s):

- * /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdjpgcom.c

No license file was found, but licenses were detected in source scan.

/*

- * jccoefct.c

*

- * Copyright (C) 1994-1997, Thomas G. Lane.
- * Modified 2003-2011 by Guido Vollbeding.
- * This file is part of the Independent JPEG Group's software.
- * For conditions of distribution and use, see the accompanying README file.
- *
- * This file contains the coefficient buffer controller for compression.
- * This controller is the top level of the JPEG compressor proper.
- * The coefficient buffer lies between forward-DCT and entropy encoding steps.
- */

Found in path(s):

- * /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jccoefct.c

No license file was found, but licenses were detected in source scan.

/*

- * jdmarker.c

*

- * Copyright (C) 1991-1998, Thomas G. Lane.
- * Modified 2009-2013 by Guido Vollbeding.
- * This file is part of the Independent JPEG Group's software.
- * For conditions of distribution and use, see the accompanying README file.
- *
- * This file contains routines to decode JPEG datastream markers.
- * Most of the complexity arises from our desire to support input
- * suspension: if not all of the data for a marker is available,
- * we must exit back to the application. On resumption, we reprocess
- * the marker.
- */

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdmarker.c

No license file was found, but licenses were detected in source scan.

/*

* jquant1.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2011 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains 1-pass color quantization (color mapping) routines.

* These routines provide mapping to a fixed color map using equally spaced

* color values. Optional Floyd-Steinberg or ordered dithering is available.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jquant1.c

No license file was found, but licenses were detected in source scan.

/*

* jccolor.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2011-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains input colorspace conversion routines.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jccolor.c

No license file was found, but licenses were detected in source scan.

/*

* jdarith.c

*

* Developed 1997-2015 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains portable arithmetic entropy decoding routines for JPEG

* (implementing the ISO/IEC IS 10918-1 and CCITT Recommendation ITU-T T.81).

*

* Both sequential and progressive modes are supported in this single module.

*

* Suspension is not currently supported in this module.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdarith.c

No license file was found, but licenses were detected in source scan.

/*

* wrtarga.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2015-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to write output images in Targa format.

*

* These routines may need modification for non-Unix environments or

* specialized applications. As they stand, they assume output to

* an ordinary stdio stream.

*

* Based on code contributed by Lee Daniel Crocker.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/wrtarga.c

No license file was found, but licenses were detected in source scan.

/*

* jpegtran.c

*

* Copyright (C) 1995-2013, Thomas G. Lane, Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains a command-line user interface for JPEG transcoding.

* It is very similar to cjpeg.c, and partly to djpeg.c, but provides

* lossless transcoding between different JPEG file formats. It also

* provides some lossless and sort-of-lossless transformations of JPEG data.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jpegtran.c

No license file was found, but licenses were detected in source scan.

/*

* rdppm.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* Modified 2009-2017 by Bill Allombert, Guido Vollbeding.
* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file contains routines to read input images in PPM/PGM format.
* The extended 2-byte-per-sample raw PPM/PGM formats are supported.
* The PBMPLUS library is NOT required to compile this software
* (but it is highly useful as a set of PPM image manipulation programs).
*
* These routines may need modification for non-Unix environments or
* specialized applications. As they stand, they assume input from
* an ordinary stdio stream. They further assume that reading begins
* at the start of the file; start_input may need work if the
* user interface has already read some data (e.g., to determine that
* the file is indeed PPM format).

*/

/* Portions of this code are based on the PBMPLUS library, which is:

**

** Copyright (C) 1988 by Jef Poskanzer.

**

** Permission to use, copy, modify, and distribute this software and its
** documentation for any purpose and without fee is hereby granted, provided
** that the above copyright notice appear in all copies and that both that
** copyright notice and this permission notice appear in supporting
** documentation. This software is provided "as is" without express or
** implied warranty.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdppm.c

No license file was found, but licenses were detected in source scan.

/*

* jaricom.c

*

* Developed 1997-2011 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains probability estimation tables for common use in
* arithmetic entropy encoding and decoding routines.

*

* This data represents Table D.3 in the JPEG spec (D.2 in the draft),

* ISO/IEC IS 10918-1 and CCITT Recommendation ITU-T T.81, and Table 24

* in the JBIG spec, ISO/IEC IS 11544 and CCITT Recommendation ITU-T T.82.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jaricom.c

No license file was found, but licenses were detected in source scan.

/*

* jutils.c

*

* Copyright (C) 1991-1996, Thomas G. Lane.

* Modified 2009-2011 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains tables and miscellaneous utility routines needed

* for both compression and decompression.

* Note we prefix all global names with "j" to minimize conflicts with

* a surrounding application.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jutils.c

No license file was found, but licenses were detected in source scan.

/*

* jmemnobs.c

*

* Copyright (C) 1992-1996, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file provides a really simple implementation of the system-

* dependent portion of the JPEG memory manager. This implementation

* assumes that no backing-store files are needed: all required space

* can be obtained from malloc().

* This is very portable in the sense that it'll compile on almost anything,

* but you'd better have lots of main memory (or virtual memory) if you want

* to process big images.

* Note that the max_memory_to_use option is ignored by this implementation.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemnobs.c

No license file was found, but licenses were detected in source scan.

/*

* cderror.h

*

* Copyright (C) 1994-1997, Thomas G. Lane.

* Modified 2009-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.
*
* This file defines the error and message codes for the cjpeg/djpeg
* applications. These strings are not needed as part of the JPEG library
* proper.
* Edit this file to add new codes, or to translate the message strings to
* some other language.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/cderror.h
No license file was found, but licenses were detected in source scan.

USING THE IJG JPEG LIBRARY

Copyright (C) 1994-2013, Thomas G. Lane, Guido Vollbeding.
This file is part of the Independent JPEG Group's software.
For conditions of distribution and use, see the accompanying README file.

This file describes how to use the IJG JPEG library within an application program. Read it if you want to write a program that uses the library.

The file `example.c` provides heavily commented skeleton code for calling the JPEG library. Also see `jpeglib.h` (the include file to be used by application programs) for full details about data structures and function parameter lists. The library source code, of course, is the ultimate reference.

Note that there have been *major* changes from the application interface presented by IJG version 4 and earlier versions. The old design had several inherent limitations, and it had accumulated a lot of cruft as we added features while trying to minimize application-interface changes. We have sacrificed backward compatibility in the version 5 rewrite, but we think the improvements justify this.

TABLE OF CONTENTS

Overview:

Functions provided by the library

Outline of typical usage

Basic library usage:

Data formats

Compression details

Decompression details

Mechanics of usage: include files, linking, etc

Advanced features:

- Compression parameter selection
- Decompression parameter selection
- Special color spaces
- Error handling
- Compressed data handling (source and destination managers)
- I/O suspension
- Progressive JPEG support
- Buffered-image mode
- Abbreviated datastreams and multiple images
- Special markers
- Raw (downsampled) image data
- Really raw data: DCT coefficients
- Progress monitoring
- Memory management
- Memory usage
- Library compile-time options
- Portability considerations
- Notes for MS-DOS implementors

You should read at least the overview and basic usage sections before trying to program with the library. The sections on advanced features can be read if and when you need them.

OVERVIEW

=====

Functions provided by the library

The IJG JPEG library provides C code to read and write JPEG-compressed image files. The surrounding application program receives or supplies image data a scanline at a time, using a straightforward uncompressed image format. All details of color conversion and other preprocessing/postprocessing can be handled by the library.

The library includes a substantial amount of code that is not covered by the JPEG standard but is necessary for typical applications of JPEG. These functions preprocess the image before JPEG compression or postprocess it after decompression. They include colorspace conversion, downsampling/upsampling, and color quantization. The application indirectly selects use of this code by specifying the format in which it wishes to supply or receive image data. For example, if colormapped output is requested, then the decompression library automatically invokes color quantization.

A wide range of quality vs. speed tradeoffs are possible in JPEG processing, and even more so in decompression postprocessing. The decompression library provides multiple implementations that cover most of the useful tradeoffs,

ranging from very-high-quality down to fast-preview operation. On the compression side we have generally not provided low-quality choices, since compression is normally less time-critical. It should be understood that the low-quality modes may not meet the JPEG standard's accuracy requirements; nonetheless, they are useful for viewers.

A word about functions *not* provided by the library. We handle a subset of the ISO JPEG standard; most baseline, extended-sequential, and progressive JPEG processes are supported. (Our subset includes all features now in common use.) Unsupported ISO options include:

- * Hierarchical storage
- * Lossless JPEG
- * DNL marker
- * Nonintegral subsampling ratios

We support 8-bit to 12-bit data precision, but this is a compile-time choice rather than a run-time choice; hence it is difficult to use different precisions in a single application.

By itself, the library handles only interchange JPEG datastreams --- in particular the widely used JFIF file format. The library can be used by surrounding code to process interchange or abbreviated JPEG datastreams that are embedded in more complex file formats. (For example, this library is used by the free LIBTIFF library to support JPEG compression in TIFF.)

Outline of typical usage

The rough outline of a JPEG compression operation is:

```
Allocate and initialize a JPEG compression object
Specify the destination for the compressed data (eg, a file)
Set parameters for compression, including image size & colorspace
jpeg_start_compress(...);
while (scan lines remain to be written)
    jpeg_write_scanlines(...);
jpeg_finish_compress(...);
Release the JPEG compression object
```

A JPEG compression object holds parameters and working state for the JPEG library. We make creation/destruction of the object separate from starting or finishing compression of an image; the same object can be re-used for a series of image compression operations. This makes it easy to re-use the same parameter settings for a sequence of images. Re-use of a JPEG object also has important implications for processing abbreviated JPEG datastreams, as discussed later.

The image data to be compressed is supplied to `jpeg_write_scanlines()` from

in-memory buffers. If the application is doing file-to-file compression, reading image data from the source file is the application's responsibility. The library emits compressed data by calling a "data destination manager", which typically will write the data into a file; but the application can provide its own destination manager to do something else.

Similarly, the rough outline of a JPEG decompression operation is:

```
Allocate and initialize a JPEG decompression object
Specify the source of the compressed data (eg, a file)
Call jpeg_read_header() to obtain image info
Set parameters for decompression
jpeg_start_decompress(...);
while (scan lines remain to be read)
    jpeg_read_scanlines(...);
jpeg_finish_decompress(...);
Release the JPEG decompression object
```

This is comparable to the compression outline except that reading the datastream header is a separate step. This is helpful because information about the image's size, colorspace, etc is available when the application selects decompression parameters. For example, the application can choose an output scaling ratio that will fit the image into the available screen size.

The decompression library obtains compressed data by calling a data source manager, which typically will read the data from a file; but other behaviors can be obtained with a custom source manager. Decompressed data is delivered into in-memory buffers passed to `jpeg_read_scanlines()`.

It is possible to abort an incomplete compression or decompression operation by calling `jpeg_abort()`; or, if you do not need to retain the JPEG object, simply release it by calling `jpeg_destroy()`.

JPEG compression and decompression objects are two separate struct types. However, they share some common fields, and certain routines such as `jpeg_destroy()` can work on either type of object.

The JPEG library has no static variables: all state is in the compression or decompression object. Therefore it is possible to process multiple compression and decompression operations concurrently, using multiple JPEG objects.

Both compression and decompression can be done in an incremental memory-to-memory fashion, if suitable source/destination managers are used. See the section on "I/O suspension" for more details.

BASIC LIBRARY USAGE

=====

Data formats

Before diving into procedural details, it is helpful to understand the image data format that the JPEG library expects or returns.

The standard input image format is a rectangular array of pixels, with each pixel having the same number of "component" or "sample" values (color channels). You must specify how many components there are and the colorspace interpretation of the components. Most applications will use RGB data (three components per pixel) or grayscale data (one component per pixel). PLEASE NOTE THAT RGB DATA IS THREE SAMPLES PER PIXEL, GRAYSCALE ONLY ONE. A remarkable number of people manage to miss this, only to find that their programs don't work with grayscale JPEG files.

There is no provision for colormapped input. JPEG files are always full-color or full grayscale (or sometimes another colorspace such as CMYK). You can feed in a colormapped image by expanding it to full-color format. However JPEG often doesn't work very well with source data that has been colormapped, because of dithering noise. This is discussed in more detail in the JPEG FAQ and the other references mentioned in the README file.

Pixels are stored by scanlines, with each scanline running from left to right. The component values for each pixel are adjacent in the row; for example, R,G,B,R,G,B,R,G,B,... for 24-bit RGB color. Each scanline is an array of data type JSAMPLE --- which is typically "unsigned char", unless you've changed jmorecfg.h. (You can also change the RGB pixel layout, say to B,G,R order, by modifying jmorecfg.h. But see the restrictions listed in that file before doing so.)

A 2-D array of pixels is formed by making a list of pointers to the starts of scanlines; so the scanlines need not be physically adjacent in memory. Even if you process just one scanline at a time, you must make a one-element pointer array to conform to this structure. Pointers to JSAMPLE rows are of type JSAMPROW, and the pointer to the pointer array is of type JSAMPARRAY.

The library accepts or supplies one or more complete scanlines per call. It is not possible to process part of a row at a time. Scanlines are always processed top-to-bottom. You can process an entire image in one call if you have it all in memory, but usually it's simplest to process one scanline at a time.

For best results, source data values should have the precision specified by BITS_IN_JSAMPLE (normally 8 bits). For instance, if you choose to compress data that's only 6 bits/channel, you should left-justify each value in a byte before passing it to the compressor. If you need to compress data

that has more than 8 bits/channel, compile with BITS_IN_JSAMPLE = 9 to 12.
(See "Library compile-time options", later.)

The data format returned by the decompressor is the same in all details, except that colormapped output is supported. (Again, a JPEG file is never colormapped. But you can ask the decompressor to perform on-the-fly color quantization to deliver colormapped output.) If you request colormapped output then the returned data array contains a single JSAMPLE per pixel; its value is an index into a color map. The color map is represented as a 2-D JSAMPARRAY in which each row holds the values of one color component, that is, colormap[i][j] is the value of the i'th color component for pixel value (map index) j. Note that since the colormap indexes are stored in JSAMPLEs, the maximum number of colors is limited by the size of JSAMPLE (ie, at most 256 colors for an 8-bit JPEG library).

Compression details

Here we revisit the JPEG compression outline given in the overview.

1. Allocate and initialize a JPEG compression object.

A JPEG compression object is a "struct jpeg_compress_struct". (It also has a bunch of subsidiary structures which are allocated via malloc(), but the application doesn't control those directly.) This struct can be just a local variable in the calling routine, if a single routine is going to execute the whole JPEG compression sequence. Otherwise it can be static or allocated from malloc().

You will also need a structure representing a JPEG error handler. The part of this that the library cares about is a "struct jpeg_error_mgr". If you are providing your own error handler, you'll typically want to embed the jpeg_error_mgr struct in a larger structure; this is discussed later under "Error handling". For now we'll assume you are just using the default error handler. The default error handler will print JPEG error/warning messages on stderr, and it will call exit() if a fatal error occurs.

You must initialize the error handler structure, store a pointer to it into the JPEG object's "err" field, and then call jpeg_create_compress() to initialize the rest of the JPEG object.

Typical code for this step, if you are using the default error handler, is

```
struct jpeg_compress_struct cinfo;
struct jpeg_error_mgr jerr;
...
```

```
cinfo.err = jpeg_std_error(&jerr);
jpeg_create_compress(&cinfo);
```

jpeg_create_compress allocates a small amount of memory, so it could fail if you are out of memory. In that case it will exit via the error handler; that's why the error handler must be initialized first.

2. Specify the destination for the compressed data (eg, a file).

As previously mentioned, the JPEG library delivers compressed data to a "data destination" module. The library includes one data destination module which knows how to write to a stdio stream. You can use your own destination module if you want to do something else, as discussed later.

If you use the standard destination module, you must open the target stdio stream beforehand. Typical code for this step looks like:

```
FILE * outfile;
...
if ((outfile = fopen(filename, "wb")) == NULL) {
    fprintf(stderr, "can't open %s\n", filename);
    exit(1);
}
jpeg_stdio_dest(&cinfo, outfile);
```

where the last line invokes the standard destination module.

WARNING: it is critical that the binary compressed data be delivered to the output file unchanged. On non-Unix systems the stdio library may perform newline translation or otherwise corrupt binary data. To suppress this behavior, you may need to use a "b" option to fopen (as shown above), or use setmode() or another routine to put the stdio stream in binary mode. See cjpeg.c and djpeg.c for code that has been found to work on many systems.

You can select the data destination after setting other parameters (step 3), if that's more convenient. You may not change the destination between calling jpeg_start_compress() and jpeg_finish_compress().

3. Set parameters for compression, including image size & colorspace.

You must supply information about the source image by setting the following fields in the JPEG object (cinfo structure):

```
image_width  Width of image, in pixels
image_height Height of image, in pixels
input_components Number of color channels (samples per pixel)
```


`in_color_space` Color space of source image

The image dimensions are, hopefully, obvious. JPEG supports image dimensions of 1 to 64K pixels in either direction. The input color space is typically RGB or grayscale, and `input_components` is 3 or 1 accordingly. (See "Special color spaces", later, for more info.) The `in_color_space` field must be assigned one of the `J_COLOR_SPACE` enum constants, typically `JCS_RGB` or `JCS_GRAYSCALE`.

JPEG has a large number of compression parameters that determine how the image is encoded. Most applications don't need or want to know about all these parameters. You can set all the parameters to reasonable defaults by calling `jpeg_set_defaults()`; then, if there are particular values you want to change, you can do so after that. The "Compression parameter selection" section tells about all the parameters.

You must set `in_color_space` correctly before calling `jpeg_set_defaults()`, because the defaults depend on the source image colorspace. However the other three source image parameters need not be valid until you call `jpeg_start_compress()`. There's no harm in calling `jpeg_set_defaults()` more than once, if that happens to be convenient.

Typical code for a 24-bit RGB source image is

```
cinfo.image_width = Width; /* image width and height, in pixels */
cinfo.image_height = Height;
cinfo.input_components = 3; /* # of color components per pixel */
cinfo.in_color_space = JCS_RGB; /* colorspace of input image */

jpeg_set_defaults(&cinfo);
/* Make optional parameter settings here */
```

4. `jpeg_start_compress(...)`;

After you have established the data destination and set all the necessary source image info and other parameters, call `jpeg_start_compress()` to begin a compression cycle. This will initialize internal state, allocate working storage, and emit the first few bytes of the JPEG datastream header.

Typical code:

```
jpeg_start_compress(&cinfo, TRUE);
```

The "TRUE" parameter ensures that a complete JPEG interchange datastream will be written. This is appropriate in most cases. If you think you might want to use an abbreviated datastream, read the section on abbreviated datastreams, below.

Once you have called `jpeg_start_compress()`, you may not alter any JPEG parameters or other fields of the JPEG object until you have completed the compression cycle.

5. while (scan lines remain to be written)

```
jpeg_write_scanlines(...);
```

Now write all the required image data by calling `jpeg_write_scanlines()` one or more times. You can pass one or more scanlines in each call, up to the total image height. In most applications it is convenient to pass just one or a few scanlines at a time. The expected format for the passed data is discussed under "Data formats", above.

Image data should be written in top-to-bottom scanline order. The JPEG spec contains some weasel wording about how top and bottom are application-defined terms (a curious interpretation of the English language...) but if you want your files to be compatible with everyone else's, you WILL use top-to-bottom order. If the source data must be read in bottom-to-top order, you can use the JPEG library's virtual array mechanism to invert the data efficiently. Examples of this can be found in the sample application `cjpeg`.

The library maintains a count of the number of scanlines written so far in the `next_scanline` field of the JPEG object. Usually you can just use this variable as the loop counter, so that the loop test looks like "while (`cinfo.next_scanline < cinfo.image_height`)".

Code for this step depends heavily on the way that you store the source data. `example.c` shows the following code for the case of a full-size 2-D source array containing 3-byte RGB pixels:

```
JSAMPROW row_pointer[1]; /* pointer to a single row */
int row_stride; /* physical row width in buffer */

row_stride = image_width * 3; /* JSAMPLEs per row in image_buffer */

while (cinfo.next_scanline < cinfo.image_height) {
    row_pointer[0] = &image_buffer[cinfo.next_scanline * row_stride];
    jpeg_write_scanlines(&cinfo, row_pointer, 1);
}
```

`jpeg_write_scanlines()` returns the number of scanlines actually written. This will normally be equal to the number passed in, so you can usually ignore the return value. It is different in just two cases:

- * If you try to write more scanlines than the declared image height, the additional scanlines are ignored.
- * If you use a suspending data destination manager, output buffer overrun

will cause the compressor to return before accepting all the passed lines.

This feature is discussed under "I/O suspension", below. The normal stdio destination manager will NOT cause this to happen.

In any case, the return value is the same as the change in the value of `next_scanline`.

6. `jpeg_finish_compress(...)`;

After all the image data has been written, call `jpeg_finish_compress()` to complete the compression cycle. This step is ESSENTIAL to ensure that the last bufferload of data is written to the data destination.

`jpeg_finish_compress()` also releases working memory associated with the JPEG object.

Typical code:

```
jpeg_finish_compress(&cinfo);
```

If using the stdio destination manager, don't forget to close the output stdio stream (if necessary) afterwards.

If you have requested a multi-pass operating mode, such as Huffman code optimization, `jpeg_finish_compress()` will perform the additional passes using data buffered by the first pass. In this case `jpeg_finish_compress()` may take quite a while to complete. With the default compression parameters, this will not happen.

It is an error to call `jpeg_finish_compress()` before writing the necessary total number of scanlines. If you wish to abort compression, call `jpeg_abort()` as discussed below.

After completing a compression cycle, you may dispose of the JPEG object as discussed next, or you may use it to compress another image. In that case return to step 2, 3, or 4 as appropriate. If you do not change the destination manager, the new datastream will be written to the same target.

If you do not change any JPEG parameters, the new datastream will be written with the same parameters as before. Note that you can change the input image dimensions freely between cycles, but if you change the input colorspace, you should call `jpeg_set_defaults()` to adjust for the new colorspace; and then you'll need to repeat all of step 3.

7. Release the JPEG compression object.

When you are done with a JPEG compression object, destroy it by calling `jpeg_destroy_compress()`. This will free all subsidiary memory (regardless of the previous state of the object). Or you can call `jpeg_destroy()`, which

works for either compression or decompression objects --- this may be more convenient if you are sharing code between compression and decompression cases. (Actually, these routines are equivalent except for the declared type of the passed pointer. To avoid gripes from ANSI C compilers, jpeg_destroy() should be passed a j_common_ptr.)

If you allocated the jpeg_compress_struct structure from malloc(), freeing it is your responsibility --- jpeg_destroy() won't. Ditto for the error handler structure.

Typical code:

```
jpeg_destroy_compress(&cinfo);
```

8. Aborting.

If you decide to abort a compression cycle before finishing, you can clean up in either of two ways:

* If you don't need the JPEG object any more, just call jpeg_destroy_compress() or jpeg_destroy() to release memory. This is legitimate at any point after calling jpeg_create_compress() --- in fact, it's safe even if jpeg_create_compress() fails.

* If you want to re-use the JPEG object, call jpeg_abort_compress(), or call jpeg_abort() which works on both compression and decompression objects. This will return the object to an idle state, releasing any working memory. jpeg_abort() is allowed at any time after successful object creation.

Note that cleaning up the data destination, if required, is your responsibility; neither of these routines will call term_destination(). (See "Compressed data handling", below, for more about that.)

jpeg_destroy() and jpeg_abort() are the only safe calls to make on a JPEG object that has reported an error by calling error_exit (see "Error handling" for more info). The internal state of such an object is likely to be out of whack. Either of these two routines will return the object to a known state.

Decompression details

Here we revisit the JPEG decompression outline given in the overview.

1. Allocate and initialize a JPEG decompression object.

This is just like initialization for compression, as discussed above,

except that the object is a "struct jpeg_decompress_struct" and you call jpeg_create_decompress(). Error handling is exactly the same.

Typical code:

```
struct jpeg_decompress_struct cinfo;
struct jpeg_error_mgr jerr;
...
cinfo.err = jpeg_std_error(&jerr);
jpeg_create_decompress(&cinfo);
```

(Both here and in the IJG code, we usually use variable name "cinfo" for both compression and decompression objects.)

2. Specify the source of the compressed data (eg, a file).

As previously mentioned, the JPEG library reads compressed data from a "data source" module. The library includes one data source module which knows how to read from a stdio stream. You can use your own source module if you want to do something else, as discussed later.

If you use the standard source module, you must open the source stdio stream beforehand. Typical code for this step looks like:

```
FILE * infile;
...
if ((infile = fopen(filename, "rb")) == NULL) {
    fprintf(stderr, "can't open %s\n", filename);
    exit(1);
}
jpeg_stdio_src(&cinfo, infile);
```

where the last line invokes the standard source module.

WARNING: it is critical that the binary compressed data be read unchanged. On non-Unix systems the stdio library may perform newline translation or otherwise corrupt binary data. To suppress this behavior, you may need to use a "b" option to fopen (as shown above), or use setmode() or another routine to put the stdio stream in binary mode. See cjpeg.c and djpeg.c for code that has been found to work on many systems.

You may not change the data source between calling jpeg_read_header() and jpeg_finish_decompress(). If you wish to read a series of JPEG images from a single source file, you should repeat the jpeg_read_header() to jpeg_finish_decompress() sequence without reinitializing either the JPEG object or the data source module; this prevents buffered input data from being discarded.

3. Call `jpeg_read_header()` to obtain image info.

Typical code for this step is just

```
jpeg_read_header(&cinfo, TRUE);
```

This will read the source datastream header markers, up to the beginning of the compressed data proper. On return, the image dimensions and other info have been stored in the JPEG object. The application may wish to consult this information before selecting decompression parameters.

More complex code is necessary if

- * A suspending data source is used --- in that case `jpeg_read_header()` may return before it has read all the header data. See "I/O suspension", below. The normal stdio source manager will NOT cause this to happen.
- * Abbreviated JPEG files are to be processed --- see the section on abbreviated datastreams. Standard applications that deal only in interchange JPEG files need not be concerned with this case either.

It is permissible to stop at this point if you just wanted to find out the image dimensions and other header info for a JPEG file. In that case, call `jpeg_destroy()` when you are done with the JPEG object, or call `jpeg_abort()` to return it to an idle state before selecting a new data source and reading another header.

4. Set parameters for decompression.

`jpeg_read_header()` sets appropriate default decompression parameters based on the properties of the image (in particular, its colorspace). However, you may well want to alter these defaults before beginning the decompression. For example, the default is to produce full color output from a color file. If you want colormapped output you must ask for it. Other options allow the returned image to be scaled and allow various speed/quality tradeoffs to be selected. "Decompression parameter selection", below, gives details.

If the defaults are appropriate, nothing need be done at this step.

Note that all default values are set by each call to `jpeg_read_header()`. If you reuse a decompression object, you cannot expect your parameter settings to be preserved across cycles, as you can for compression. You must set desired parameter values each time.

5. `jpeg_start_decompress(...)`;

Once the parameter values are satisfactory, call `jpeg_start_decompress()` to begin decompression. This will initialize internal state, allocate working memory, and prepare for returning data.

Typical code is just

```
jpeg_start_decompress(&cinfo);
```

If you have requested a multi-pass operating mode, such as 2-pass color quantization, `jpeg_start_decompress()` will do everything needed before data output can begin. In this case `jpeg_start_decompress()` may take quite a while to complete. With a single-scan (non progressive) JPEG file and default decompression parameters, this will not happen; `jpeg_start_decompress()` will return quickly.

After this call, the final output image dimensions, including any requested scaling, are available in the JPEG object; so is the selected colormap, if colormapped output has been requested. Useful fields include

```
output_width  image width and height, as scaled
output_height
out_color_components # of color components in out_color_space
output_components # of color components returned per pixel
colormap       the selected colormap, if any
actual_number_of_colors number of entries in colormap
```

`output_components` is 1 (a colormap index) when quantizing colors; otherwise it equals `out_color_components`. It is the number of JSAMPLE values that will be emitted per pixel in the output arrays.

Typically you will need to allocate data buffers to hold the incoming image. You will need `output_width * output_components` JSAMPLEs per scanline in your output buffer, and a total of `output_height` scanlines will be returned.

Note: if you are using the JPEG library's internal memory manager to allocate data buffers (as `djpeg` does), then the manager's protocol requires that you request large buffers *before* calling `jpeg_start_decompress()`. This is a little tricky since the `output_XXX` fields are not normally valid then. You can make them valid by calling `jpeg_calc_output_dimensions()` after setting the relevant parameters (scaling, output color space, and quantization flag).

6. while (scan lines remain to be read)

```
jpeg_read_scanlines(...);
```

Now you can read the decompressed image data by calling `jpeg_read_scanlines()` one or more times. At each call, you pass in the maximum number of scanlines to be read (ie, the height of your working buffer); `jpeg_read_scanlines()`

will return up to that many lines. The return value is the number of lines actually read. The format of the returned data is discussed under "Data formats", above. Don't forget that grayscale and color JPEGs will return different data formats!

Image data is returned in top-to-bottom scanline order. If you must write out the image in bottom-to-top order, you can use the JPEG library's virtual array mechanism to invert the data efficiently. Examples of this can be found in the sample application `djpeg`.

The library maintains a count of the number of scanlines returned so far in the `output_scanline` field of the JPEG object. Usually you can just use this variable as the loop counter, so that the loop test looks like "while (`cinfo.output_scanline < cinfo.output_height`)". (Note that the test should NOT be against `image_height`, unless you never use scaling. The `image_height` field is the height of the original unscaled image.) The return value always equals the change in the value of `output_scanline`.

If you don't use a suspending data source, it is safe to assume that `jpeg_read_scanlines()` reads at least one scanline per call, until the bottom of the image has been reached.

If you use a buffer larger than one scanline, it is NOT safe to assume that `jpeg_read_scanlines()` fills it. (The current implementation returns only a few scanlines per call, no matter how large a buffer you pass.) So you must always provide a loop that calls `jpeg_read_scanlines()` repeatedly until the whole image has been read.

7. `jpeg_finish_decompress(...)`;

After all the image data has been read, call `jpeg_finish_decompress()` to complete the decompression cycle. This causes working memory associated with the JPEG object to be released.

Typical code:

```
jpeg_finish_decompress(&cinfo);
```

If using the `stdio` source manager, don't forget to close the source `stdio` stream if necessary.

It is an error to call `jpeg_finish_decompress()` before reading the correct total number of scanlines. If you wish to abort decompression, call `jpeg_abort()` as discussed below.

After completing a decompression cycle, you may dispose of the JPEG object as discussed next, or you may use it to decompress another image. In that case

return to step 2 or 3 as appropriate. If you do not change the source manager, the next image will be read from the same source.

8. Release the JPEG decompression object.

When you are done with a JPEG decompression object, destroy it by calling `jpeg_destroy_decompress()` or `jpeg_destroy()`. The previous discussion of destroying compression objects applies here too.

Typical code:

```
jpeg_destroy_decompress(&cinfo);
```

9. Aborting.

You can abort a decompression cycle by calling `jpeg_destroy_decompress()` or `jpeg_destroy()` if you don't need the JPEG object any more, or `jpeg_abort_decompress()` or `jpeg_abort()` if you want to reuse the object. The previous discussion of aborting compression cycles applies here too.

Mechanics of usage: include files, linking, etc

Applications using the JPEG library should include the header file `jpeglib.h` to obtain declarations of data types and routines. Before including `jpeglib.h`, include system headers that define at least the typedefs `FILE` and `size_t`. On ANSI-conforming systems, including `<stdio.h>` is sufficient; on older Unix systems, you may need `<sys/types.h>` to define `size_t`.

If the application needs to refer to individual JPEG library error codes, also include `jerror.h` to define those symbols.

`jpeglib.h` indirectly includes the files `jconfig.h` and `jmorecfg.h`. If you are installing the JPEG header files in a system directory, you will want to install all four files: `jpeglib.h`, `jerror.h`, `jconfig.h`, `jmorecfg.h`.

The most convenient way to include the JPEG code into your executable program is to prepare a library file ("`libjpeg.a`", or a corresponding name on non-Unix machines) and reference it at your link step. If you use only half of the library (only compression or only decompression), only that much code will be included from the library, unless your linker is hopelessly brain-damaged. The supplied makefiles build `libjpeg.a` automatically (see `install.txt`).

While you can build the JPEG library as a shared library if the whim strikes you, we don't really recommend it. The trouble with shared libraries is that

at some point you'll probably try to substitute a new version of the library without recompiling the calling applications. That generally doesn't work because the parameter struct declarations usually change with each new version. In other words, the library's API is **not** guaranteed binary compatible across versions; we only try to ensure source-code compatibility. (In hindsight, it might have been smarter to hide the parameter structs from applications and introduce a ton of access functions instead. Too late now, however.)

On some systems your application may need to set up a signal handler to ensure that temporary files are deleted if the program is interrupted. This is most critical if you are on MS-DOS and use the `jmemdos.c` memory manager back end; it will try to grab extended memory for temp files, and that space will NOT be freed automatically. See `cjpeg.c` or `djpeg.c` for an example signal handler.

It may be worth pointing out that the core JPEG library does not actually require the `stdio` library: only the default source/destination managers and error handler need it. You can use the library in a `stdio`-less environment if you replace those modules and use `jmemnobs.c` (or another memory manager of your own devising). More info about the minimum system library requirements may be found in `jininclude.h`.

ADVANCED FEATURES

=====

Compression parameter selection

This section describes all the optional parameters you can set for JPEG compression, as well as the "helper" routines provided to assist in this task. Proper setting of some parameters requires detailed understanding of the JPEG standard; if you don't know what a parameter is for, it's best not to mess with it! See REFERENCES in the README file for pointers to more info about JPEG.

It's a good idea to call `jpeg_set_defaults()` first, even if you plan to set all the parameters; that way your code is more likely to work with future JPEG libraries that have additional parameters. For the same reason, we recommend you use a helper routine where one is provided, in preference to twiddling `cinfo` fields directly.

The helper routines are:

`jpeg_set_defaults (j_compress_ptr cinfo)`

This routine sets all JPEG parameters to reasonable defaults, using only the input image's color space (field `in_color_space`, which must already be set in `cinfo`). Many applications will only need to use

this routine and perhaps `jpeg_set_quality()`.

`jpeg_set_colorspace (j_compress_ptr cinfo, J_COLOR_SPACE colorspace)`
Sets the JPEG file's colorspace (field `jpeg_color_space`) as specified, and sets other color-space-dependent parameters appropriately. See "Special color spaces", below, before using this. A large number of parameters, including all per-component parameters, are set by this routine; if you want to twiddle individual parameters you should call `jpeg_set_colorspace()` before rather than after.

`jpeg_default_colorspace (j_compress_ptr cinfo)`
Selects an appropriate JPEG colorspace based on `cinfo->in_color_space`, and calls `jpeg_set_colorspace()`. This is actually a subroutine of `jpeg_set_defaults()`. It's broken out in case you want to change just the colorspace-dependent JPEG parameters.

`jpeg_set_quality (j_compress_ptr cinfo, int quality, boolean force_baseline)`
Constructs JPEG quantization tables appropriate for the indicated quality setting. The quality value is expressed on the 0..100 scale recommended by IJG (cjpeg's "-quality" switch uses this routine). Note that the exact mapping from quality values to tables may change in future IJG releases as more is learned about DCT quantization. If the `force_baseline` parameter is TRUE, then the quantization table entries are constrained to the range 1..255 for full JPEG baseline compatibility. In the current implementation, this only makes a difference for quality settings below 25, and it effectively prevents very small/low quality files from being generated. The IJG decoder is capable of reading the non-baseline files generated at low quality settings when `force_baseline` is FALSE, but other decoders may not be.

`jpeg_set_linear_quality (j_compress_ptr cinfo, int scale_factor, boolean force_baseline)`
Same as `jpeg_set_quality()` except that the generated tables are the sample tables given in the JPEG spec section K.1, multiplied by the specified scale factor (which is expressed as a percentage; thus `scale_factor = 100` reproduces the spec's tables). Note that larger scale factors give lower quality. This entry point is useful for conforming to the Adobe PostScript DCT conventions, but we do not recommend linear scaling as a user-visible quality scale otherwise. `force_baseline` again constrains the computed table entries to 1..255.

`int jpeg_quality_scaling (int quality)`
Converts a value on the IJG-recommended quality scale to a linear scaling percentage. Note that this routine may change or go away in future releases --- IJG may choose to adopt a scaling method that can't be expressed as a simple scalar multiplier, in which case the premise of this routine collapses. Caveat user.

`jpeg_default_qtables` (`j_compress_ptr` cinfo, boolean `force_baseline`)
Set default quantization tables with linear `q_scale_factor[]` values
(see below).

`jpeg_add_quant_table` (`j_compress_ptr` cinfo, int `which_tbl`,
const unsigned int `*basic_table`,
int `scale_factor`, boolean `force_baseline`)

Allows an arbitrary quantization table to be created. `which_tbl` indicates which table slot to fill. `basic_table` points to an array of 64 unsigned ints given in normal array order. These values are multiplied by `scale_factor/100` and then clamped to the range 1..65535 (or to 1..255 if `force_baseline` is TRUE).

CAUTION: prior to library version 6a, `jpeg_add_quant_table` expected the basic table to be given in JPEG zigzag order. If you need to write code that works with either older or newer versions of this routine, you must check the library version number. Something like `"#if JPEG_LIB_VERSION >= 61"` is the right test.

`jpeg_simple_progression` (`j_compress_ptr` cinfo)

Generates a default scan script for writing a progressive-JPEG file. This is the recommended method of creating a progressive file, unless you want to make a custom scan sequence. You must ensure that the JPEG color space is set correctly before calling this routine.

Compression parameters (cinfo fields) include:

boolean `arith_code`

If TRUE, use arithmetic coding.

If FALSE, use Huffman coding.

int `block_size`

Set DCT block size. All N from 1 to 16 are possible.

Default is 8 (baseline format).

Larger values produce higher compression,
smaller values produce higher quality.

An exact DCT stage is possible with 1 or 2.

With the default quality of 75 and default Luminance qtable
the DCT+Quantization stage is lossless for value 1.

Note that values other than 8 require a SmartScale capable decoder,
introduced with IJG JPEG 8. Setting the `block_size` parameter for
compression works with version 8c and later.

`J_DCT_METHOD` `dct_method`

Selects the algorithm used for the DCT step. Choices are:

`JDCT_ISLOW`: slow but accurate integer algorithm

`JDCT_IFAST`: faster, less accurate integer method

`JDCT_FLOAT`: floating-point method

JDCT_DEFAULT: default method (normally JDCT_ISLOW)

JDCT_FASTEST: fastest method (normally JDCT_IFAST)

The FLOAT method is very slightly more accurate than the ISLOW method, but may give different results on different machines due to varying roundoff behavior. The integer methods should give the same results on all machines. On machines with sufficiently fast FP hardware, the floating-point method may also be the fastest. The IFAST method is considerably less accurate than the other two; its use is not recommended if high quality is a concern. JDCT_DEFAULT and JDCT_FASTEST are macros configurable by each installation.

unsigned int scale_num, scale_denom

Scale the image by the fraction scale_num/scale_denom. Default is 1/1, or no scaling. Currently, the supported scaling ratios are M/N with all N from 1 to 16, where M is the destination DCT size, which is 8 by default (see block_size parameter above). (The library design allows for arbitrary scaling ratios but this is not likely to be implemented any time soon.)

J_COLOR_SPACE jpeg_color_space

int num_components

The JPEG color space and corresponding number of components; see "Special color spaces", below, for more info. We recommend using jpeg_set_colorspace() if you want to change these.

J_COLOR_TRANSFORM color_transform

Internal color transform identifier, writes LSE marker if nonzero (requires decoder with inverse color transform support, introduced with IJG JPEG 9).

Two values are currently possible: JCT_NONE and JCT_SUBTRACT_GREEN.

Set this value for lossless RGB application *before* calling jpeg_set_colorspace(), because entropy table assignment in jpeg_set_colorspace() depends on color_transform.

boolean optimize_coding

TRUE causes the compressor to compute optimal Huffman coding tables for the image. This requires an extra pass over the data and therefore costs a good deal of space and time. The default is FALSE, which tells the compressor to use the supplied or default Huffman tables. In most cases optimal tables save only a few percent of file size compared to the default tables. Note that when this is TRUE, you need not supply Huffman tables at all, and any you do supply will be overwritten.

unsigned int restart_interval

int restart_in_rows

To emit restart markers in the JPEG file, set one of these nonzero.

Set restart_interval to specify the exact interval in MCU blocks.

Set `restart_in_rows` to specify the interval in MCU rows. (If `restart_in_rows` is not 0, then `restart_interval` is set after the image width in MCUs is computed.) Defaults are zero (no restarts). One restart marker per MCU row is often a good choice.
NOTE: the overhead of restart markers is higher in grayscale JPEG files than in color files, and MUCH higher in progressive JPEGs. If you use restarts, you may want to use larger intervals in those cases.

```
const jpeg_scan_info * scan_info
int num_scans
```

By default, `scan_info` is NULL; this causes the compressor to write a single-scan sequential JPEG file. If not NULL, `scan_info` points to an array of scan definition records of length `num_scans`. The compressor will then write a JPEG file having one scan for each scan definition record. This is used to generate noninterleaved or progressive JPEG files. The library checks that the scan array defines a valid JPEG scan sequence. (`jpeg_simple_progression` creates a suitable scan definition array for progressive JPEG.) This is discussed further under "Progressive JPEG support".

```
boolean do_fancy_downsampling
```

If TRUE, use direct DCT scaling with DCT size > 8 for downsampling of chroma components.

If FALSE, use only DCT size <= 8 and simple separate downsampling. Default is TRUE.

For better image stability in multiple generation compression cycles it is preferable that this value matches the corresponding `do_fancy_upsampling` value in decompression.

```
int smoothing_factor
```

If non-zero, the input image is smoothed; the value should be 1 for minimal smoothing to 100 for maximum smoothing. Consult `jcsample.c` for details of the smoothing algorithm. The default is zero.

```
boolean write_JFIF_header
```

If TRUE, a JFIF APP0 marker is emitted. `jpeg_set_defaults()` and `jpeg_set_colorspace()` set this TRUE if a JFIF-legal JPEG color space (ie, YCbCr or grayscale) is selected, otherwise FALSE.

```
UINT8 JFIF_major_version
```

```
UINT8 JFIF_minor_version
```

The version number to be written into the JFIF marker.

`jpeg_set_defaults()` initializes the version to 1.01 (major=minor=1).

You should set it to 1.02 (major=1, minor=2) if you plan to write any JFIF 1.02 extension markers.

```
UINT8 density_unit
```

UINT16 X_density

UINT16 Y_density

The resolution information to be written into the JFIF marker; not used otherwise. density_unit may be 0 for unknown, 1 for dots/inch, or 2 for dots/cm. The default values are 0,1,1 indicating square pixels of unknown size.

boolean write_Adobe_marker

If TRUE, an Adobe APP14 marker is emitted. jpeg_set_defaults() and jpeg_set_colorspace() set this TRUE if JPEG color space RGB, CMYK, or YCCK is selected, otherwise FALSE. It is generally a bad idea to set both write_JFIF_header and write_Adobe_marker. In fact, you probably shouldn't change the default settings at all --- the default behavior ensures that the JPEG file's color space can be recognized by the decoder.

JQUANT_TBL * quant_tbl_ptrs[NUM_QUANT_TBLS]

Pointers to coefficient quantization tables, one per table slot, or NULL if no table is defined for a slot. Usually these should be set via one of the above helper routines; jpeg_add_quant_table() is general enough to define any quantization table. The other routines will set up table slot 0 for luminance quality and table slot 1 for chrominance.

int q_scale_factor[NUM_QUANT_TBLS]

Linear quantization scaling factors (percentage, initialized 100) for use with jpeg_default_qtables().

See rdswitch.c and cjpeg.c for an example of usage.

Note that the q_scale_factor[] fields are the "linear" scales, so you have to convert from user-defined ratings via jpeg_quality_scaling().

Here is an example code which corresponds to cjpeg -quality 90,70:

```
jpeg_set_defaults(cinfo);

/* Set luminance quality 90. */
cinfo->q_scale_factor[0] = jpeg_quality_scaling(90);
/* Set chrominance quality 70. */
cinfo->q_scale_factor[1] = jpeg_quality_scaling(70);

jpeg_default_qtables(cinfo, force_baseline);
```

CAUTION: You must also set 1x1 subsampling for efficient separate color quality selection, since the default value used by library is 2x2:

```
cinfo->comp_info[0].v_samp_factor = 1;
cinfo->comp_info[0].h_samp_factor = 1;
```

JHUFF_TBL * dc_huff_tbl_ptrs[NUM_HUFF_TBLS]

JHUFF_TBL * ac_huff_tbl_ptrs[NUM_HUFF_TBLS]

Pointers to Huffman coding tables, one per table slot, or NULL if no table is defined for a slot. Slots 0 and 1 are filled with the JPEG sample tables by `jpeg_set_defaults()`. If you need to allocate more table structures, `jpeg_alloc_huff_table()` may be used.

Note that optimal Huffman tables can be computed for an image by setting `optimize_coding`, as discussed above; there's seldom any need to mess with providing your own Huffman tables.

The actual dimensions of the JPEG image that will be written to the file are given by the following fields. These are computed from the input image dimensions and the compression parameters by `jpeg_start_compress()`. You can also call `jpeg_calc_jpeg_dimensions()` to obtain the values that will result from the current parameter settings. This can be useful if you are trying to pick a scaling ratio that will get close to a desired target size.

JDIMENSION `jpeg_width` Actual dimensions of output image.

JDIMENSION `jpeg_height`

Per-component parameters are stored in the struct `cinfo.comp_info[i]` for component number `i`. Note that components here refer to components of the JPEG color space, *not* the source image color space. A suitably large `comp_info[]` array is allocated by `jpeg_set_defaults()`; if you choose not to use that routine, it's up to you to allocate the array.

int `component_id`

The one-byte identifier code to be recorded in the JPEG file for this component. For the standard color spaces, we recommend you leave the default values alone.

int `h_samp_factor`

int `v_samp_factor`

Horizontal and vertical sampling factors for the component; must be 1..4 according to the JPEG standard. Note that larger sampling factors indicate a higher-resolution component; many people find this behavior quite unintuitive. The default values are 2,2 for luminance components and 1,1 for chrominance components, except for grayscale where 1,1 is used.

int `quant_tbl_no`

Quantization table number for component. The default value is 0 for luminance components and 1 for chrominance components.

int `dc_tbl_no`

int `ac_tbl_no`

DC and AC entropy coding table numbers. The default values are 0 for luminance components and 1 for chrominance components.

int component_index

Must equal the component's index in comp_info[]. (Beginning in release v6, the compressor library will fill this in automatically; you don't have to.)

Decompression parameter selection

Decompression parameter selection is somewhat simpler than compression parameter selection, since all of the JPEG internal parameters are recorded in the source file and need not be supplied by the application. (Unless you are working with abbreviated files, in which case see "Abbreviated datastreams", below.) Decompression parameters control the postprocessing done on the image to deliver it in a format suitable for the application's use. Many of the parameters control speed/quality tradeoffs, in which faster decompression may be obtained at the price of a poorer-quality image. The defaults select the highest quality (slowest) processing.

The following fields in the JPEG object are set by jpeg_read_header() and may be useful to the application in choosing decompression parameters:

JDIMENSION image_width Width and height of image
JDIMENSION image_height
int num_components Number of color components
J_COLOR_SPACE jpeg_color_space Colorspace of image
boolean saw_JFIF_marker TRUE if a JFIF APP0 marker was seen
UINT8 JFIF_major_version Version information from JFIF marker
UINT8 JFIF_minor_version
UINT8 density_unit Resolution data from JFIF marker
UINT16 X_density
UINT16 Y_density
boolean saw_Adobe_marker TRUE if an Adobe APP14 marker was seen
UINT8 Adobe_transform Color transform code from Adobe marker

The JPEG color space, unfortunately, is something of a guess since the JPEG standard proper does not provide a way to record it. In practice most files adhere to the JFIF or Adobe conventions, and the decoder will recognize these correctly. See "Special color spaces", below, for more info.

The decompression parameters that determine the basic properties of the returned image are:

J_COLOR_SPACE out_color_space

Output color space. jpeg_read_header() sets an appropriate default based on jpeg_color_space; typically it will be RGB or grayscale. The application can change this field to request output in a different colorspace. For example, set it to JCS_GRAYSCALE to get grayscale output from a color file. (This is useful for previewing: grayscale output is faster than full color since the color components need not be processed.) Note that not all possible color space transforms are currently implemented; you may need to extend jdcolor.c if you want an unusual conversion.

unsigned int scale_num, scale_denom

Scale the image by the fraction scale_num/scale_denom. Currently, the supported scaling ratios are M/N with all M from 1 to 16, where N is the source DCT size, which is 8 for baseline JPEG. (The library design allows for arbitrary scaling ratios but this is not likely to be implemented any time soon.) The values are initialized by jpeg_read_header() with the source DCT size. For baseline JPEG this is 8/8. If you change only the scale_num value while leaving the other unchanged, then this specifies the DCT scaled size to be applied on the given input. For baseline JPEG this is equivalent to M/8 scaling, since the source DCT size for baseline JPEG is 8. Smaller scaling ratios permit significantly faster decoding since fewer pixels need be processed and a simpler IDCT method can be used.

boolean quantize_colors

If set TRUE, colormapped output will be delivered. Default is FALSE, meaning that full-color output will be delivered.

The next three parameters are relevant only if quantize_colors is TRUE.

int desired_number_of_colors

Maximum number of colors to use in generating a library-supplied color map (the actual number of colors is returned in a different field). Default 256. Ignored when the application supplies its own color map.

boolean two_pass_quantize

If TRUE, an extra pass over the image is made to select a custom color map for the image. This usually looks a lot better than the one-size-fits-all colormap that is used otherwise. Default is TRUE. Ignored when the application supplies its own color map.

J_DITHER_MODE dither_mode

Selects color dithering method. Supported values are:

JDITHER_NONE no dithering: fast, very low quality

JDITHER_ORDERED ordered dither: moderate speed and quality

JDITHER_FS Floyd-Steinberg dither: slow, high quality

Default is JDITHER_FS. (At present, ordered dither is implemented

only in the single-pass, standard-colormap case. If you ask for ordered dither when `two_pass_quantize` is TRUE or when you supply an external color map, you'll get F-S dithering.)

When `quantize_colors` is TRUE, the target color map is described by the next two fields. `colormap` is set to NULL by `jpeg_read_header()`. The application can supply a color map by setting `colormap` non-NULL and setting `actual_number_of_colors` to the map size. Otherwise, `jpeg_start_decompress()` selects a suitable color map and sets these two fields itself.

[Implementation restriction: at present, an externally supplied colormap is only accepted for 3-component output color spaces.]

JSAMPARRAY colormap

The color map, represented as a 2-D pixel array of `out_color_components` rows and `actual_number_of_colors` columns. Ignored if not quantizing.

CAUTION: if the JPEG library creates its own colormap, the storage pointed to by this field is released by `jpeg_finish_decompress()`.

Copy the colormap somewhere else first, if you want to save it.

int `actual_number_of_colors`

The number of colors in the color map.

Additional decompression parameters that the application may set include:

J_DCT_METHOD `dct_method`

Selects the algorithm used for the DCT step. Choices are the same as described above for compression.

boolean `do_fancy_upsampling`

If TRUE, use direct DCT scaling with DCT size > 8 for upsampling of chroma components.

If FALSE, use only DCT size <= 8 and simple separate upsampling.

Default is TRUE.

For better image stability in multiple generation compression cycles it is preferable that this value matches the corresponding `do_fancy_downsampling` value in compression.

boolean `do_block_smoothing`

If TRUE, interblock smoothing is applied in early stages of decoding progressive JPEG files; if FALSE, not. Default is TRUE. Early progression stages look "fuzzy" with smoothing, "blocky" without.

In any case, block smoothing ceases to be applied after the first few AC coefficients are known to full accuracy, so it is relevant only when using buffered-image mode for progressive images.

boolean `enable_1pass_quant`

boolean `enable_external_quant`

boolean `enable_2pass_quant`

These are significant only in buffered-image mode, which is described in its own section below.

The output image dimensions are given by the following fields. These are computed from the source image dimensions and the decompression parameters by `jpeg_start_decompress()`. You can also call `jpeg_calc_output_dimensions()` to obtain the values that will result from the current parameter settings.

This can be useful if you are trying to pick a scaling ratio that will get close to a desired target size. It's also important if you are using the JPEG library's memory manager to allocate output buffer space, because you are supposed to request such buffers *before* `jpeg_start_decompress()`.

JDIMENSION `output_width` Actual dimensions of output image.

JDIMENSION `output_height`

int `out_color_components` Number of color components in `out_color_space`.

int `output_components` Number of color components returned.

int `rec_outbuf_height` Recommended height of scanline buffer.

When quantizing colors, `output_components` is 1, indicating a single color map index per pixel. Otherwise it equals `out_color_components`. The output arrays are required to be `output_width * output_components` JSAMPLEs wide.

`rec_outbuf_height` is the recommended minimum height (in scanlines) of the buffer passed to `jpeg_read_scanlines()`. If the buffer is smaller, the library will still work, but time will be wasted due to unnecessary data copying. In high-quality modes, `rec_outbuf_height` is always 1, but some faster, lower-quality modes set it to larger values (typically 2 to 4).

If you are going to ask for a high-speed processing mode, you may as well go to the trouble of honoring `rec_outbuf_height` so as to avoid data copying. (An output buffer larger than `rec_outbuf_height` lines is OK, but won't provide any material speed improvement over that height.)

Special color spaces

The JPEG standard itself is "color blind" and doesn't specify any particular color space. It is customary to convert color data to a luminance/chrominance color space before compressing, since this permits greater compression. The existing JPEG file interchange format standards specify YCbCr or GRAYSCALE data (JFIF version 1), GRAYSCALE, RGB, YCbCr, CMYK, or YCCK (Adobe), or BG_RGB or BG_YCC (big gamut color spaces, JFIF version 2). For special applications such as multispectral images, other color spaces can be used, but it must be understood that such files will be unportable.

The JPEG library can handle the most common colorspace conversions (namely RGB <=> YCbCr and CMYK <=> YCCK). It can also deal with data of an unknown

color space, passing it through without conversion. If you deal extensively with an unusual color space, you can easily extend the library to understand additional color spaces and perform appropriate conversions.

For compression, the source data's color space is specified by field `in_color_space`. This is transformed to the JPEG file's color space given by `jpeg_color_space`. `jpeg_set_defaults()` chooses a reasonable JPEG color space depending on `in_color_space`, but you can override this by calling `jpeg_set_colorspace()`. Of course you must select a supported transformation. `jpegcolor.c` currently supports the following transformations:

RGB => YCbCr

RGB => GRAYSCALE

RGB => BG_YCC

YCbCr => GRAYSCALE

YCbCr => BG_YCC

CMYK => YCCK

plus the null transforms: GRAYSCALE => GRAYSCALE, RGB => RGB, BG_RGB => BG_RGB, YCbCr => YCbCr, BG_YCC => BG_YCC, CMYK => CMYK, YCCK => YCCK, and UNKNOWN => UNKNOWN.

The file interchange format standards (JFIF and Adobe) specify APPn markers that indicate the color space of the JPEG file. It is important to ensure that these are written correctly, or omitted if the JPEG file's color space is not one of the ones supported by the interchange standards.

`jpeg_set_colorspace()` will set the compression parameters to include or omit the APPn markers properly, so long as it is told the truth about the JPEG color space. For example, if you are writing some random 3-component color space without conversion, don't try to fake out the library by setting `in_color_space` and `jpeg_color_space` to `JCS_YCbCr`; use `JCS_UNKNOWN`. You may want to write an APPn marker of your own devising to identify the colorspace --- see "Special markers", below.

When told that the color space is UNKNOWN, the library will default to using luminance-quality compression parameters for all color components. You may well want to change these parameters. See the source code for `jpeg_set_colorspace()`, in `jcparam.c`, for details.

For decompression, the JPEG file's color space is given in `jpeg_color_space`, and this is transformed to the output color space `out_color_space`. `jpeg_read_header's` setting of `jpeg_color_space` can be relied on if the file conforms to JFIF or Adobe conventions, but otherwise it is no better than a guess. If you know the JPEG file's color space for certain, you can override `jpeg_read_header's` guess by setting `jpeg_color_space`. `jpeg_read_header` also selects a default output color space based on (its guess of) `jpeg_color_space`; set `out_color_space` to override this. Again, you must select a supported transformation. `jpegcolor.c` currently supports

YCbCr => RGB

YCbCr => GRAYSCALE

BG_YCC => RGB
BG_YCC => GRAYSCALE
RGB => GRAYSCALE
GRAYSCALE => RGB
YCK => CMYK

as well as the null transforms. (Since GRAYSCALE=>RGB is provided, an application can force grayscale JPEGs to look like color JPEGs if it only wants to handle one case.)

The two-pass color quantizer, `jquant2.c`, is specialized to handle RGB data (it weights distances appropriately for RGB colors). You'll need to modify the code if you want to use it for non-RGB output color spaces. Note that `jquant2.c` is used to map to an application-supplied colormap as well as for the normal two-pass colormap selection process.

CAUTION: it appears that Adobe Photoshop writes inverted data in CMYK JPEG files: 0 represents 100% ink coverage, rather than 0% ink as you'd expect. This is arguably a bug in Photoshop, but if you need to work with Photoshop CMYK files, you will have to deal with it in your application. We cannot "fix" this in the library by inverting the data during the CMYK<=>YCK transform, because that would break other applications, notably Ghostscript. Photoshop versions prior to 3.0 write EPS files containing JPEG-encoded CMYK data in the same inverted-YCK representation used in bare JPEG files, but the surrounding PostScript code performs an inversion using the PS image operator. I am told that Photoshop 3.0 will write uninverted YCK in EPS/JPEG files, and will omit the PS-level inversion. (But the data polarity used in bare JPEG files will not change in 3.0.) In either case, the JPEG library must not invert the data itself, or else Ghostscript would read these EPS files incorrectly.

Error handling

When the default error handler is used, any error detected inside the JPEG routines will cause a message to be printed on `stderr`, followed by `exit()`. You can supply your own error handling routines to override this behavior and to control the treatment of nonfatal warnings and trace/debug messages. The file `example.c` illustrates the most common case, which is to have the application regain control after an error rather than exiting.

The JPEG library never writes any message directly; it always goes through the error handling routines. Three classes of messages are recognized:

- * Fatal errors: the library cannot continue.
- * Warnings: the library can continue, but the data is corrupt, and a damaged output image is likely to result.
- * Trace/informational messages. These come with a trace level indicating the importance of the message; you can control the verbosity of the

program by adjusting the maximum trace level that will be displayed.

You may, if you wish, simply replace the entire JPEG error handling module (`jerror.c`) with your own code. However, you can avoid code duplication by only replacing some of the routines depending on the behavior you need. This is accomplished by calling `jpeg_std_error()` as usual, but then overriding some of the method pointers in the `jpeg_error_mgr` struct, as illustrated by `example.c`.

All of the error handling routines will receive a pointer to the JPEG object (a `j_common_ptr` which points to either a `jpeg_compress_struct` or a `jpeg_decompress_struct`; if you need to tell which, test the `is_decompressor` field). This struct includes a pointer to the error manager struct in its "err" field. Frequently, custom error handler routines will need to access additional data which is not known to the JPEG library or the standard error handler. The most convenient way to do this is to embed either the JPEG object or the `jpeg_error_mgr` struct in a larger structure that contains additional fields; then casting the passed pointer provides access to the additional fields. Again, see `example.c` for one way to do it. (Beginning with IJG version 6b, there is also a void pointer "client_data" in each JPEG object, which the application can also use to find related data. The library does not touch `client_data` at all.)

The individual methods that you might wish to override are:

`error_exit (j_common_ptr cinfo)`

Receives control for a fatal error. Information sufficient to generate the error message has been stored in `cinfo->err`; call `output_message` to display it. Control must NOT return to the caller; generally this routine will `exit()` or `longjmp()` somewhere. Typically you would override this routine to get rid of the `exit()` default behavior. Note that if you continue processing, you should clean up the JPEG object with `jpeg_abort()` or `jpeg_destroy()`.

`output_message (j_common_ptr cinfo)`

Actual output of any JPEG message. Override this to send messages somewhere other than `stderr`. Note that this method does not know how to generate a message, only where to send it.

`format_message (j_common_ptr cinfo, char * buffer)`

Constructs a readable error message string based on the error info stored in `cinfo->err`. This method is called by `output_message`. Few applications should need to override this method. One possible reason for doing so is to implement dynamic switching of error message language.

`emit_message (j_common_ptr cinfo, int msg_level)`

Decide whether or not to emit a warning or trace message; if so,

calls `output_message`. The main reason for overriding this method would be to abort on warnings. `msg_level` is -1 for warnings, 0 and up for trace messages.

Only `error_exit()` and `emit_message()` are called from the rest of the JPEG library; the other two are internal to the error handler.

The actual message texts are stored in an array of strings which is pointed to by the field `err->jpeg_message_table`. The messages are numbered from 0 to `err->last_jpeg_message`, and it is these code numbers that are used in the JPEG library code. You could replace the message texts (for instance, with messages in French or German) by changing the message table pointer. See `jerror.h` for the default texts. CAUTION: this table will almost certainly change or grow from one library version to the next.

It may be useful for an application to add its own message texts that are handled by the same mechanism. The error handler supports a second "add-on" message table for this purpose. To define an addon table, set the pointer `err->addon_message_table` and the message numbers `err->first_addon_message` and `err->last_addon_message`. If you number the addon messages beginning at 1000 or so, you won't have to worry about conflicts with the library's built-in messages. See the sample applications `cjpeg/djpeg` for an example of using addon messages (the addon messages are defined in `cderror.h`).

Actual invocation of the error handler is done via macros defined in `jerror.h`:

`ERREXITn(...)` for fatal errors

`WARNMSn(...)` for corrupt-data warnings

`TRACEMSn(...)` for trace and informational messages.

These macros store the message code and any additional parameters into the error handler struct, then invoke the `error_exit()` or `emit_message()` method.

The variants of each macro are for varying numbers of additional parameters.

The additional parameters are inserted into the generated message using standard `printf()` format codes.

See `jerror.h` and `jerror.c` for further details.

Compressed data handling (source and destination managers)

The JPEG compression library sends its compressed data to a "destination manager" module. The default destination manager just writes the data to a memory buffer or to a stdio stream, but you can provide your own manager to do something else. Similarly, the decompression library calls a "source manager" to obtain the compressed data; you can provide your own source manager if you want the data to come from somewhere other than a memory buffer or a stdio stream.

In both cases, compressed data is processed a bufferload at a time: the destination or source manager provides a work buffer, and the library invokes the manager only when the buffer is filled or emptied. (You could define a one-character buffer to force the manager to be invoked for each byte, but that would be rather inefficient.) The buffer's size and location are controlled by the manager, not by the library. For example, the memory source manager just makes the buffer pointer and length point to the original data in memory. In this case the buffer-reload procedure will be invoked only if the decompressor ran off the end of the datastream, which would indicate an erroneous datastream.

The work buffer is defined as an array of datatype JOCTET, which is generally "char" or "unsigned char". On a machine where char is not exactly 8 bits wide, you must define JOCTET as a wider data type and then modify the data source and destination modules to transcribe the work arrays into 8-bit units on external storage.

A data destination manager struct contains a pointer and count defining the next byte to write in the work buffer and the remaining free space:

```
JOCTET * next_output_byte; /* => next byte to write in buffer */
size_t free_in_buffer;    /* # of byte spaces remaining in buffer */
```

The library increments the pointer and decrements the count until the buffer is filled. The manager's empty_output_buffer method must reset the pointer and count. The manager is expected to remember the buffer's starting address and total size in private fields not visible to the library.

A data destination manager provides three methods:

init_destination (j_compress_ptr cinfo)
Initialize destination. This is called by jpeg_start_compress() before any data is actually written. It must initialize next_output_byte and free_in_buffer. free_in_buffer must be initialized to a positive value.

empty_output_buffer (j_compress_ptr cinfo)
This is called whenever the buffer has filled (free_in_buffer reaches zero). In typical applications, it should write out the *entire* buffer (use the saved start address and buffer length; ignore the current state of next_output_byte and free_in_buffer). Then reset the pointer & count to the start of the buffer, and return TRUE indicating that the buffer has been dumped. free_in_buffer must be set to a positive value when TRUE is returned. A FALSE return should only be used when I/O suspension is desired (this operating mode is discussed in the next section).

term_destination (j_compress_ptr cinfo)

Terminate destination --- called by `jpeg_finish_compress()` after all data has been written. In most applications, this must flush any data remaining in the buffer. Use either `next_output_byte` or `free_in_buffer` to determine how much data is in the buffer.

`term_destination()` is NOT called by `jpeg_abort()` or `jpeg_destroy()`. If you want the destination manager to be cleaned up during an abort, you must do it yourself.

You will also need code to create a `jpeg_destination_mgr` struct, fill in its method pointers, and insert a pointer to the struct into the "dest" field of the JPEG compression object. This can be done in-line in your setup code if you like, but it's probably cleaner to provide a separate routine similar to the `jpeg_stdio_dest()` or `jpeg_mem_dest()` routines of the supplied destination managers.

Decompression source managers follow a parallel design, but with some additional frammishes. The source manager struct contains a pointer and count defining the next byte to read from the work buffer and the number of bytes remaining:

```
const JOCTET * next_input_byte; /* => next byte to read from buffer */
size_t bytes_in_buffer;      /* # of bytes remaining in buffer */
```

The library increments the pointer and decrements the count until the buffer is emptied. The manager's `fill_input_buffer` method must reset the pointer and count. In most applications, the manager must remember the buffer's starting address and total size in private fields not visible to the library.

A data source manager provides five methods:

`init_source (j_decompress_ptr cinfo)`

Initialize source. This is called by `jpeg_read_header()` before any data is actually read. Unlike `init_destination()`, it may leave `bytes_in_buffer` set to 0 (in which case a `fill_input_buffer()` call will occur immediately).

`fill_input_buffer (j_decompress_ptr cinfo)`

This is called whenever `bytes_in_buffer` has reached zero and more data is wanted. In typical applications, it should read fresh data into the buffer (ignoring the current state of `next_input_byte` and `bytes_in_buffer`), reset the pointer & count to the start of the buffer, and return TRUE indicating that the buffer has been reloaded. It is not necessary to fill the buffer entirely, only to obtain at least one more byte. `bytes_in_buffer` MUST be set to a positive value if TRUE is returned. A FALSE return should only be used when I/O suspension is desired (this mode is discussed in the next section).

skip_input_data (j_decompress_ptr cinfo, long num_bytes)

Skip num_bytes worth of data. The buffer pointer and count should be advanced over num_bytes input bytes, refilling the buffer as needed. This is used to skip over a potentially large amount of uninteresting data (such as an APPn marker). In some applications it may be possible to optimize away the reading of the skipped data, but it's not clear that being smart is worth much trouble; large skips are uncommon. bytes_in_buffer may be zero on return. A zero or negative skip count should be treated as a no-op.

resync_to_restart (j_decompress_ptr cinfo, int desired)

This routine is called only when the decompressor has failed to find a restart (RSTn) marker where one is expected. Its mission is to find a suitable point for resuming decompression. For most applications, we recommend that you just use the default resync procedure, jpeg_resync_to_restart(). However, if you are able to back up in the input data stream, or if you have a-priori knowledge about the likely location of restart markers, you may be able to do better. Read the read_restart_marker() and jpeg_resync_to_restart() routines in jdmarker.c if you think you'd like to implement your own resync procedure.

term_source (j_decompress_ptr cinfo)

Terminate source --- called by jpeg_finish_decompress() after all data has been read. Often a no-op.

For both fill_input_buffer() and skip_input_data(), there is no such thing as an EOF return. If the end of the file has been reached, the routine has a choice of exiting via ERREXIT() or inserting fake data into the buffer. In most cases, generating a warning message and inserting a fake EOI marker is the best course of action --- this will allow the decompressor to output however much of the image is there. In pathological cases, the decompressor may swallow the EOI and again demand data ... just keep feeding it fake EOIs. jdatsrc.c illustrates the recommended error recovery behavior.

term_source() is NOT called by jpeg_abort() or jpeg_destroy(). If you want the source manager to be cleaned up during an abort, you must do it yourself.

You will also need code to create a jpeg_source_mgr struct, fill in its method pointers, and insert a pointer to the struct into the "src" field of the JPEG decompression object. This can be done in-line in your setup code if you like, but it's probably cleaner to provide a separate routine similar to the jpeg_stdio_src() or jpeg_mem_src() routines of the supplied source managers.

For more information, consult the memory and stdio source and destination managers in jdatsrc.c and jdatdst.c.

I/O suspension

Some applications need to use the JPEG library as an incremental memory-to-memory filter: when the compressed data buffer is filled or emptied, they want control to return to the outer loop, rather than expecting that the buffer can be emptied or reloaded within the data source/destination manager subroutine. The library supports this need by providing an "I/O suspension" mode, which we describe in this section.

The I/O suspension mode is not a panacea: nothing is guaranteed about the maximum amount of time spent in any one call to the library, so it will not eliminate response-time problems in single-threaded applications. If you need guaranteed response time, we suggest you "bite the bullet" and implement a real multi-tasking capability.

To use I/O suspension, cooperation is needed between the calling application and the data source or destination manager; you will always need a custom source/destination manager. (Please read the previous section if you haven't already.) The basic idea is that the `empty_output_buffer()` or `fill_input_buffer()` routine is a no-op, merely returning `FALSE` to indicate that it has done nothing. Upon seeing this, the JPEG library suspends operation and returns to its caller. The surrounding application is responsible for emptying or refilling the work buffer before calling the JPEG library again.

Compression suspension:

For compression suspension, use an `empty_output_buffer()` routine that returns `FALSE`; typically it will not do anything else. This will cause the compressor to return to the caller of `jpeg_write_scanlines()`, with the return value indicating that not all the supplied scanlines have been accepted. The application must make more room in the output buffer, adjust the output buffer pointer/count appropriately, and then call `jpeg_write_scanlines()` again, pointing to the first unconsumed scanline.

When forced to suspend, the compressor will backtrack to a convenient stopping point (usually the start of the current MCU); it will regenerate some output data when restarted. Therefore, although `empty_output_buffer()` is only called when the buffer is filled, you should NOT write out the entire buffer after a suspension. Write only the data up to the current position of `next_output_byte/free_in_buffer`. The data beyond that point will be regenerated after resumption.

Because of the backtracking behavior, a good-size output buffer is essential for efficiency; you don't want the compressor to suspend often. (In fact, an overly small buffer could lead to infinite looping, if a single MCU required more data than would fit in the buffer.) We recommend a buffer of at least

several Kbytes. You may want to insert explicit code to ensure that you don't call `jpeg_write_scanlines()` unless there is a reasonable amount of space in the output buffer; in other words, flush the buffer before trying to compress more data.

The compressor does not allow suspension while it is trying to write JPEG markers at the beginning and end of the file. This means that:

- * At the beginning of a compression operation, there must be enough free space in the output buffer to hold the header markers (typically 600 or so bytes). The recommended buffer size is bigger than this anyway, so this is not a problem as long as you start with an empty buffer. However, this restriction might catch you if you insert large special markers, such as a JFIF thumbnail image, without flushing the buffer afterwards.
- * When you call `jpeg_finish_compress()`, there must be enough space in the output buffer to emit any buffered data and the final EOI marker. In the current implementation, half a dozen bytes should suffice for this, but for safety's sake we recommend ensuring that at least 100 bytes are free before calling `jpeg_finish_compress()`.

A more significant restriction is that `jpeg_finish_compress()` cannot suspend. This means you cannot use suspension with multi-pass operating modes, namely Huffman code optimization and multiple-scan output. Those modes write the whole file during `jpeg_finish_compress()`, which will certainly result in buffer overrun. (Note that this restriction applies only to compression, not decompression. The decompressor supports input suspension in all of its operating modes.)

Decompression suspension:

For decompression suspension, use a `fill_input_buffer()` routine that simply returns `FALSE` (except perhaps during error recovery, as discussed below). This will cause the decompressor to return to its caller with an indication that suspension has occurred. This can happen at four places:

- * `jpeg_read_header()`: will return `JPEG_SUSPENDED`.
- * `jpeg_start_decompress()`: will return `FALSE`, rather than its usual `TRUE`.
- * `jpeg_read_scanlines()`: will return the number of scanlines already completed (possibly 0).
- * `jpeg_finish_decompress()`: will return `FALSE`, rather than its usual `TRUE`.

The surrounding application must recognize these cases, load more data into the input buffer, and repeat the call. In the case of `jpeg_read_scanlines()`, increment the passed pointers past any scanlines successfully read.

Just as with compression, the decompressor will typically backtrack to a convenient restart point before suspending. When `fill_input_buffer()` is called, `next_input_byte/bytes_in_buffer` point to the current restart point, which is where the decompressor will backtrack to if `FALSE` is returned. The data beyond that position must NOT be discarded if you suspend; it needs to be re-read upon resumption. In most implementations, you'll need to shift

this data down to the start of your work buffer and then load more data after it. Again, this behavior means that a several-Kbyte work buffer is essential for decent performance; furthermore, you should load a reasonable amount of new data before resuming decompression. (If you loaded, say, only one new byte each time around, you could waste a LOT of cycles.)

The `skip_input_data()` source manager routine requires special care in a suspension scenario. This routine is NOT granted the ability to suspend the decompressor; it can decrement `bytes_in_buffer` to zero, but no more. If the requested skip distance exceeds the amount of data currently in the input buffer, then `skip_input_data()` must set `bytes_in_buffer` to zero and record the additional skip distance somewhere else. The decompressor will immediately call `fill_input_buffer()`, which should return `FALSE`, which will cause a suspension return. The surrounding application must then arrange to discard the recorded number of bytes before it resumes loading the input buffer. (Yes, this design is rather baroque, but it avoids complexity in the far more common case where a non-suspending source manager is used.)

If the input data has been exhausted, we recommend that you emit a warning and insert dummy EOI markers just as a non-suspending data source manager would do. This can be handled either in the surrounding application logic or within `fill_input_buffer()`; the latter is probably more efficient. If `fill_input_buffer()` knows that no more data is available, it can set the pointer/count to point to a dummy EOI marker and then return `TRUE` just as though it had read more data in a non-suspending situation.

The decompressor does not attempt to suspend within standard JPEG markers; instead it will backtrack to the start of the marker and reprocess the whole marker next time. Hence the input buffer must be large enough to hold the longest standard marker in the file. Standard JPEG markers should normally not exceed a few hundred bytes each (DHT tables are typically the longest). We recommend at least a 2K buffer for performance reasons, which is much larger than any correct marker is likely to be. For robustness against damaged marker length counts, you may wish to insert a test in your application for the case that the input buffer is completely full and yet the decoder has suspended without consuming any data --- otherwise, if this situation did occur, it would lead to an endless loop. (The library can't provide this test since it has no idea whether "the buffer is full", or even whether there is a fixed-size input buffer.)

The input buffer would need to be 64K to allow for arbitrary COM or APPn markers, but these are handled specially: they are either saved into allocated memory, or skipped over by calling `skip_input_data()`. In the former case, suspension is handled correctly, and in the latter case, the problem of buffer overrun is placed on `skip_input_data`'s shoulders, as explained above. Note that if you provide your own marker handling routine for large markers, you should consider how to deal with buffer overflow.

Multiple-buffer management:

In some applications it is desirable to store the compressed data in a linked list of buffer areas, so as to avoid data copying. This can be handled by having `empty_output_buffer()` or `fill_input_buffer()` set the pointer and count to reference the next available buffer; `FALSE` is returned only if no more buffers are available. Although seemingly straightforward, there is a pitfall in this approach: the backtrack that occurs when `FALSE` is returned could back up into an earlier buffer. For example, when `fill_input_buffer()` is called, the current pointer & count indicate the backtrack restart point. Since `fill_input_buffer()` will set the pointer and count to refer to a new buffer, the restart position must be saved somewhere else. Suppose a second call to `fill_input_buffer()` occurs in the same library call, and no additional input data is available, so `fill_input_buffer` must return `FALSE`. If the JPEG library has not moved the pointer/count forward in the current buffer, then *the correct restart point is the saved position in the prior buffer*. Prior buffers may be discarded only after the library establishes a restart point within a later buffer. Similar remarks apply for output into a chain of buffers.

The library will never attempt to backtrack over a `skip_input_data()` call, so any skipped data can be permanently discarded. You still have to deal with the case of skipping not-yet-received data, however.

It's much simpler to use only a single buffer; when `fill_input_buffer()` is called, move any unconsumed data (beyond the current pointer/count) down to the beginning of this buffer and then load new data into the remaining buffer space. This approach requires a little more data copying but is far easier to get right.

Progressive JPEG support

Progressive JPEG rearranges the stored data into a series of scans of increasing quality. In situations where a JPEG file is transmitted across a slow communications link, a decoder can generate a low-quality image very quickly from the first scan, then gradually improve the displayed quality as more scans are received. The final image after all scans are complete is identical to that of a regular (sequential) JPEG file of the same quality setting. Progressive JPEG files are often slightly smaller than equivalent sequential JPEG files, but the possibility of incremental display is the main reason for using progressive JPEG.

The IJG encoder library generates progressive JPEG files when given a suitable "scan script" defining how to divide the data into scans. Creation of progressive JPEG files is otherwise transparent to the encoder. Progressive JPEG files can also be read transparently by the decoder library.

If the decoding application simply uses the library as defined above, it will receive a final decoded image without any indication that the file was progressive. Of course, this approach does not allow incremental display. To perform incremental display, an application needs to use the decoder library's "buffered-image" mode, in which it receives a decoded image multiple times.

Each displayed scan requires about as much work to decode as a full JPEG image of the same size, so the decoder must be fairly fast in relation to the data transmission rate in order to make incremental display useful. However, it is possible to skip displaying the image and simply add the incoming bits to the decoder's coefficient buffer. This is fast because only Huffman decoding need be done, not IDCT, upsampling, colorspace conversion, etc. The IJG decoder library allows the application to switch dynamically between displaying the image and simply absorbing the incoming bits. A properly coded application can automatically adapt the number of display passes to suit the time available as the image is received. Also, a final higher-quality display cycle can be performed from the buffered data after the end of the file is reached.

Progressive compression:

To create a progressive JPEG file (or a multiple-scan sequential JPEG file), set the scan_info cinfo field to point to an array of scan descriptors, and perform compression as usual. Instead of constructing your own scan list, you can call the jpeg_simple_progression() helper routine to create a recommended progression sequence; this method should be used by all applications that don't want to get involved in the nitty-gritty of progressive scan sequence design. (If you want to provide user control of scan sequences, you may wish to borrow the scan script reading code found in rdswitch.c, so that you can read scan script files just like cjpeg's.) When scan_info is not NULL, the compression library will store DCT'd data into a buffer array as jpeg_write_scanlines() is called, and will emit all the requested scans during jpeg_finish_compress(). This implies that multiple-scan output cannot be created with a suspending data destination manager, since jpeg_finish_compress() does not support suspension. We should also note that the compressor currently forces Huffman optimization mode when creating a progressive JPEG file, because the default Huffman tables are unsuitable for progressive files.

Progressive decompression:

When buffered-image mode is not used, the decoder library will read all of a multi-scan file during jpeg_start_decompress(), so that it can provide a final decoded image. (Here "multi-scan" means either progressive or multi-scan sequential.) This makes multi-scan files transparent to the decoding application. However, existing applications that used suspending input with version 5 of the IJG library will need to be modified to check

for a suspension return from jpeg_start_decompress().

To perform incremental display, an application must use the library's buffered-image mode. This is described in the next section.

Buffered-image mode

In buffered-image mode, the library stores the partially decoded image in a coefficient buffer, from which it can be read out as many times as desired. This mode is typically used for incremental display of progressive JPEG files, but it can be used with any JPEG file. Each scan of a progressive JPEG file adds more data (more detail) to the buffered image. The application can display in lockstep with the source file (one display pass per input scan), or it can allow input processing to outrun display processing. By making input and display processing run independently, it is possible for the application to adapt progressive display to a wide range of data transmission rates.

The basic control flow for buffered-image decoding is

```
jpeg_create_decompress()
set data source
jpeg_read_header()
set overall decompression parameters
cinfo.buffered_image = TRUE; /* select buffered-image mode */
jpeg_start_decompress()
for (each output pass) {
    adjust output decompression parameters if required
    jpeg_start_output() /* start a new output pass */
    for (all scanlines in image) {
        jpeg_read_scanlines()
        display scanlines
    }
    jpeg_finish_output() /* terminate output pass */
}
jpeg_finish_decompress()
jpeg_destroy_decompress()
```

This differs from ordinary unbuffered decoding in that there is an additional level of looping. The application can choose how many output passes to make and how to display each pass.

The simplest approach to displaying progressive images is to do one display pass for each scan appearing in the input file. In this case the outer loop condition is typically

```
while (! jpeg_input_complete(&cinfo))
```

and the start-output call should read

```
jpeg_start_output(&cinfo, cinfo.input_scan_number);
```

The second parameter to `jpeg_start_output()` indicates which scan of the input file is to be displayed; the scans are numbered starting at 1 for this purpose. (You can use a loop counter starting at 1 if you like, but using the library's input scan counter is easier.) The library automatically reads data as necessary to complete each requested scan, and `jpeg_finish_output()` advances to the next scan or end-of-image marker (hence `input_scan_number` will be incremented by the time control arrives back at `jpeg_start_output()`). With this technique, data is read from the input file only as needed, and input and output processing run in lockstep.

After reading the final scan and reaching the end of the input file, the buffered image remains available; it can be read additional times by repeating the `jpeg_start_output()/jpeg_read_scanlines()/jpeg_finish_output()` sequence. For example, a useful technique is to use fast one-pass color quantization for display passes made while the image is arriving, followed by a final display pass using two-pass quantization for highest quality. This is done by changing the library parameters before the final output pass. Changing parameters between passes is discussed in detail below.

In general the last scan of a progressive file cannot be recognized as such until after it is read, so a post-input display pass is the best approach if you want special processing in the final pass.

When done with the image, be sure to call `jpeg_finish_decompress()` to release the buffered image (or just use `jpeg_destroy_decompress()`).

If input data arrives faster than it can be displayed, the application can cause the library to decode input data in advance of what's needed to produce output. This is done by calling the routine `jpeg_consume_input()`.

The return value is one of the following:

`JPEG_REACHED_SOS`: reached an SOS marker (the start of a new scan)

`JPEG_REACHED_EOI`: reached the EOI marker (end of image)

`JPEG_ROW_COMPLETED`: completed reading one MCU row of compressed data

`JPEG_SCAN_COMPLETED`: completed reading last MCU row of current scan

`JPEG_SUSPENDED`: suspended before completing any of the above

(`JPEG_SUSPENDED` can occur only if a suspending data source is used.) This

routine can be called at any time after initializing the JPEG object. It

reads some additional data and returns when one of the indicated significant

events occurs. (If called after the EOI marker is reached, it will

immediately return `JPEG_REACHED_EOI` without attempting to read more data.)

The library's output processing will automatically call `jpeg_consume_input()` whenever the output processing overtakes the input; thus, simple lockstep display requires no direct calls to `jpeg_consume_input()`. But by adding calls to `jpeg_consume_input()`, you can absorb data in advance of what is being displayed. This has two benefits:

- * You can limit buildup of unprocessed data in your input buffer.
- * You can eliminate extra display passes by paying attention to the state of the library's input processing.

The first of these benefits only requires interspersing calls to `jpeg_consume_input()` with your display operations and any other processing you may be doing. To avoid wasting cycles due to backtracking, it's best to call `jpeg_consume_input()` only after a hundred or so new bytes have arrived. This is discussed further under "I/O suspension", above. (Note: the JPEG library currently is not thread-safe. You must not call `jpeg_consume_input()` from one thread of control if a different library routine is working on the same JPEG object in another thread.)

When input arrives fast enough that more than one new scan is available before you start a new output pass, you may as well skip the output pass corresponding to the completed scan. This occurs for free if you pass `cinfo.input_scan_number` as the target scan number to `jpeg_start_output()`. The `input_scan_number` field is simply the index of the scan currently being consumed by the input processor. You can ensure that this is up-to-date by emptying the input buffer just before calling `jpeg_start_output()`: call `jpeg_consume_input()` repeatedly until it returns `JPEG_SUSPENDED` or `JPEG_REACHED_EOI`.

The target scan number passed to `jpeg_start_output()` is saved in the `cinfo.output_scan_number` field. The library's output processing calls `jpeg_consume_input()` whenever the current input scan number and row within that scan is less than or equal to the current output scan number and row. Thus, input processing can "get ahead" of the output processing but is not allowed to "fall behind". You can achieve several different effects by manipulating this interlock rule. For example, if you pass a target scan number greater than the current input scan number, the output processor will wait until that scan starts to arrive before producing any output. (To avoid an infinite loop, the target scan number is automatically reset to the last scan number when the end of image is reached. Thus, if you specify a large target scan number, the library will just absorb the entire input file and then perform an output pass. This is effectively the same as what `jpeg_start_decompress()` does when you don't select buffered-image mode.) When you pass a target scan number equal to the current input scan number, the image is displayed no faster than the current input scan arrives. The final possibility is to pass a target scan number less than the current input scan number; this disables the input/output interlock and causes the output processor to simply display whatever it finds in the image buffer, without waiting for input. (However, the library will not accept a target scan number less than one, so you can't avoid waiting for the first scan.)

When data is arriving faster than the output display processing can advance through the image, `jpeg_consume_input()` will store data into the buffered image beyond the point at which the output processing is reading data out

again. If the input arrives fast enough, it may "wrap around" the buffer to the point where the input is more than one whole scan ahead of the output. If the output processing simply proceeds through its display pass without paying attention to the input, the effect seen on-screen is that the lower part of the image is one or more scans better in quality than the upper part. Then, when the next output scan is started, you have a choice of what target scan number to use. The recommended choice is to use the current input scan number at that time, which implies that you've skipped the output scans corresponding to the input scans that were completed while you processed the previous output scan. In this way, the decoder automatically adapts its speed to the arriving data, by skipping output scans as necessary to keep up with the arriving data.

When using this strategy, you'll want to be sure that you perform a final output pass after receiving all the data; otherwise your last display may not be full quality across the whole screen. So the right outer loop logic is something like this:

```
do {
    absorb any waiting input by calling jpeg_consume_input()
    final_pass = jpeg_input_complete(&cinfo);
    adjust output decompression parameters if required
    jpeg_start_output(&cinfo, cinfo.input_scan_number);
    ...
    jpeg_finish_output()
} while (! final_pass);
```

rather than quitting as soon as `jpeg_input_complete()` returns TRUE. This arrangement makes it simple to use higher-quality decoding parameters for the final pass. But if you don't want to use special parameters for the final pass, the right loop logic is like this:

```
for (;;) {
    absorb any waiting input by calling jpeg_consume_input()
    jpeg_start_output(&cinfo, cinfo.input_scan_number);
    ...
    jpeg_finish_output()
    if (jpeg_input_complete(&cinfo) &&
        cinfo.input_scan_number == cinfo.output_scan_number)
        break;
}
```

In this case you don't need to know in advance whether an output pass is to be the last one, so it's not necessary to have reached EOF before starting the final output pass; rather, what you want to test is whether the output pass was performed in sync with the final input scan. This form of the loop will avoid an extra output pass whenever the decoder is able (or nearly able) to keep up with the incoming data.

When the data transmission speed is high, you might begin a display pass, then find that much or all of the file has arrived before you can complete the pass. (You can detect this by noting the `JPEG_REACHED_EOI` return code

from `jpeg_consume_input()`, or equivalently by testing `jpeg_input_complete()`.) In this situation you may wish to abort the current display pass and start a new one using the newly arrived information. To do so, just call `jpeg_finish_output()` and then start a new pass with `jpeg_start_output()`.

A variant strategy is to abort and restart display if more than one complete scan arrives during an output pass; this can be detected by noting `JPEG_REACHED_SOS` returns and/or examining `cinfo.input_scan_number`. This idea should be employed with caution, however, since the display process might never get to the bottom of the image before being aborted, resulting in the lower part of the screen being several passes worse than the upper. In most cases it's probably best to abort an output pass only if the whole file has arrived and you want to begin the final output pass immediately.

When receiving data across a communication link, we recommend always using the current input scan number for the output target scan number; if a higher-quality final pass is to be done, it should be started (aborting any incomplete output pass) as soon as the end of file is received. However, many other strategies are possible. For example, the application can examine the parameters of the current input scan and decide whether to display it or not. If the scan contains only chroma data, one might choose not to use it as the target scan, expecting that the scan will be small and will arrive quickly. To skip to the next scan, call `jpeg_consume_input()` until it returns `JPEG_REACHED_SOS` or `JPEG_REACHED_EOI`. Or just use the next higher number as the target scan for `jpeg_start_output()`; but that method doesn't let you inspect the next scan's parameters before deciding to display it.

In buffered-image mode, `jpeg_start_decompress()` never performs input and thus never suspends. An application that uses input suspension with buffered-image mode must be prepared for suspension returns from these routines:

- * `jpeg_start_output()` performs input only if you request 2-pass quantization and the target scan isn't fully read yet. (This is discussed below.)
- * `jpeg_read_scanlines()`, as always, returns the number of scanlines that it was able to produce before suspending.
- * `jpeg_finish_output()` will read any markers following the target scan, up to the end of the file or the SOS marker that begins another scan. (But it reads no input if `jpeg_consume_input()` has already reached the end of the file or a SOS marker beyond the target output scan.)
- * `jpeg_finish_decompress()` will read until the end of file, and thus can suspend if the end hasn't already been reached (as can be tested by calling `jpeg_input_complete()`).

`jpeg_start_output()`, `jpeg_finish_output()`, and `jpeg_finish_decompress()` all return `TRUE` if they completed their tasks, `FALSE` if they had to suspend. In the event of a `FALSE` return, the application must load more input data and repeat the call. Applications that use non-suspending data sources need not check the return values of these three routines.

It is possible to change decoding parameters between output passes in the buffered-image mode. The decoder library currently supports only very limited changes of parameters. ONLY THE FOLLOWING parameter changes are allowed after `jpeg_start_decompress()` is called:

* `dct_method` can be changed before each call to `jpeg_start_output()`.

For example, one could use a fast DCT method for early scans, changing to a higher quality method for the final scan.

* `dither_mode` can be changed before each call to `jpeg_start_output()`;

of course this has no impact if not using color quantization. Typically one would use ordered dither for initial passes, then switch to Floyd-Steinberg dither for the final pass. Caution: changing dither mode can cause more memory to be allocated by the library. Although the amount of memory involved is not large (a scanline or so), it may cause the initial `max_memory_to_use` specification to be exceeded, which in the worst case would result in an out-of-memory failure.

* `do_block_smoothing` can be changed before each call to `jpeg_start_output()`.

This setting is relevant only when decoding a progressive JPEG image.

During the first DC-only scan, block smoothing provides a very "fuzzy" look instead of the very "blocky" look seen without it; which is better seems a matter of personal taste. But block smoothing is nearly always a win during later stages, especially when decoding a successive-approximation image: smoothing helps to hide the slight blockiness that otherwise shows up on smooth gradients until the lowest coefficient bits are sent.

* Color quantization mode can be changed under the rules described below.

You *cannot* change between full-color and quantized output (because that would alter the required I/O buffer sizes), but you can change which quantization method is used.

When generating color-quantized output, changing quantization method is a very useful way of switching between high-speed and high-quality display.

The library allows you to change among its three quantization methods:

1. Single-pass quantization to a fixed color cube.

Selected by `cinfo.two_pass_quantize = FALSE` and `cinfo.colormap = NULL`.

2. Single-pass quantization to an application-supplied colormap.

Selected by setting `cinfo.colormap` to point to the colormap (the value of `two_pass_quantize` is ignored); also set `cinfo.actual_number_of_colors`.

3. Two-pass quantization to a colormap chosen specifically for the image.

Selected by `cinfo.two_pass_quantize = TRUE` and `cinfo.colormap = NULL`.

(This is the default setting selected by `jpeg_read_header`, but it is probably NOT what you want for the first pass of progressive display!)

These methods offer successively better quality and lesser speed. However, only the first method is available for quantizing in non-RGB color spaces.

IMPORTANT: because the different quantizer methods have very different working-storage requirements, the library requires you to indicate which one(s) you intend to use before you call `jpeg_start_decompress()`. (If we did

not require this, the `max_memory_to_use` setting would be a complete fiction.)

You do this by setting one or more of these three `cinfo` fields to `TRUE`:

`enable_1pass_quant` Fixed color cube colormap

`enable_external_quant` Externally-supplied colormap

`enable_2pass_quant` Two-pass custom colormap

All three are initialized `FALSE` by `jpeg_read_header()`. But

`jpeg_start_decompress()` automatically sets `TRUE` the one selected by the current `two_pass_quantize` and `colormap` settings, so you only need to set the enable flags for any other quantization methods you plan to change to later.

After setting the enable flags correctly at `jpeg_start_decompress()` time, you can change to any enabled quantization method by setting `two_pass_quantize` and `colormap` properly just before calling `jpeg_start_output()`. The following special rules apply:

1. You must explicitly set `cinfo.colormap` to `NULL` when switching to 1-pass or 2-pass mode from a different mode, or when you want the 2-pass quantizer to be re-run to generate a new colormap.
2. To switch to an external colormap, or to change to a different external colormap than was used on the prior pass, you must call `jpeg_new_colormap()` after setting `cinfo.colormap`.

NOTE: if you want to use the same colormap as was used in the prior pass, you should not do either of these things. This will save some nontrivial switchover costs.

(These requirements exist because `cinfo.colormap` will always be non-`NULL` after completing a prior output pass, since both the 1-pass and 2-pass quantizers set it to point to their output colormaps. Thus you have to do one of these two things to notify the library that something has changed. Yup, it's a bit klugy, but it's necessary to do it this way for backwards compatibility.)

Note that in buffered-image mode, the library generates any requested colormap during `jpeg_start_output()`, not during `jpeg_start_decompress()`.

When using two-pass quantization, `jpeg_start_output()` makes a pass over the buffered image to determine the optimum color map; it therefore may take a significant amount of time, whereas ordinarily it does little work. The progress monitor hook is called during this pass, if defined. It is also important to realize that if the specified target scan number is greater than or equal to the current input scan number, `jpeg_start_output()` will attempt to consume input as it makes this pass. If you use a suspending data source, you need to check for a `FALSE` return from `jpeg_start_output()` under these conditions. The combination of 2-pass quantization and a not-yet-fully-read target scan is the only case in which `jpeg_start_output()` will consume input.

Application authors who support buffered-image mode may be tempted to use it for all JPEG images, even single-scan ones. This will work, but it is inefficient: there is no need to create an image-sized coefficient buffer for

single-scan images. Requesting buffered-image mode for such an image wastes memory. Worse, it can cost time on large images, since the buffered data has to be swapped out or written to a temporary file. If you are concerned about maximum performance on baseline JPEG files, you should use buffered-image mode only when the incoming file actually has multiple scans. This can be tested by calling `jpeg_has_multiple_scans()`, which will return a correct result at any time after `jpeg_read_header()` completes.

It is also worth noting that when you use `jpeg_consume_input()` to let input processing get ahead of output processing, the resulting pattern of access to the coefficient buffer is quite nonsequential. It's best to use the memory manager `jmemnobs.c` if you can (ie, if you have enough real or virtual main memory). If not, at least make sure that `max_memory_to_use` is set as high as possible. If the JPEG memory manager has to use a temporary file, you will probably see a lot of disk traffic and poor performance. (This could be improved with additional work on the memory manager, but we haven't gotten around to it yet.)

In some applications it may be convenient to use `jpeg_consume_input()` for all input processing, including reading the initial markers; that is, you may wish to call `jpeg_consume_input()` instead of `jpeg_read_header()` during startup. This works, but note that you must check for `JPEG_REACHED_SOS` and `JPEG_REACHED_EOI` return codes as the equivalent of `jpeg_read_header()`'s codes. Once the first SOS marker has been reached, you must call `jpeg_start_decompress()` before `jpeg_consume_input()` will consume more input; it'll just keep returning `JPEG_REACHED_SOS` until you do. If you read a tables-only file this way, `jpeg_consume_input()` will return `JPEG_REACHED_EOI` without ever returning `JPEG_REACHED_SOS`; be sure to check for this case. If this happens, the decompressor will not read any more input until you call `jpeg_abort()` to reset it. It is OK to call `jpeg_consume_input()` even when not using buffered-image mode, but in that case it's basically a no-op after the initial markers have been read: it will just return `JPEG_SUSPENDED`.

Abbreviated datastreams and multiple images

A JPEG compression or decompression object can be reused to process multiple images. This saves a small amount of time per image by eliminating the "create" and "destroy" operations, but that isn't the real purpose of the feature. Rather, reuse of an object provides support for abbreviated JPEG datastreams. Object reuse can also simplify processing a series of images in a single input or output file. This section explains these features.

A JPEG file normally contains several hundred bytes worth of quantization and Huffman tables. In a situation where many images will be stored or transmitted with identical tables, this may represent an annoying overhead. The JPEG standard therefore permits tables to be omitted. The standard

defines three classes of JPEG datastreams:

- * "Interchange" datastreams contain an image and all tables needed to decode the image. These are the usual kind of JPEG file.
- * "Abbreviated image" datastreams contain an image, but are missing some or all of the tables needed to decode that image.
- * "Abbreviated table specification" (henceforth "tables-only") datastreams contain only table specifications.

To decode an abbreviated image, it is necessary to load the missing table(s) into the decoder beforehand. This can be accomplished by reading a separate tables-only file. A variant scheme uses a series of images in which the first image is an interchange (complete) datastream, while subsequent ones are abbreviated and rely on the tables loaded by the first image. It is assumed that once the decoder has read a table, it will remember that table until a new definition for the same table number is encountered.

It is the application designer's responsibility to figure out how to associate the correct tables with an abbreviated image. While abbreviated datastreams can be useful in a closed environment, their use is strongly discouraged in any situation where data exchange with other applications might be needed. Caveat designer.

The JPEG library provides support for reading and writing any combination of tables-only datastreams and abbreviated images. In both compression and decompression objects, a quantization or Huffman table will be retained for the lifetime of the object, unless it is overwritten by a new table definition.

To create abbreviated image datastreams, it is only necessary to tell the compressor not to emit some or all of the tables it is using. Each quantization and Huffman table struct contains a boolean field "sent_table", which normally is initialized to FALSE. For each table used by the image, the header-writing process emits the table and sets sent_table = TRUE unless it is already TRUE. (In normal usage, this prevents outputting the same table definition multiple times, as would otherwise occur because the chroma components typically share tables.) Thus, setting this field to TRUE before calling jpeg_start_compress() will prevent the table from being written at all.

If you want to create a "pure" abbreviated image file containing no tables, just call "jpeg_suppress_tables(&cinfo, TRUE)" after constructing all the tables. If you want to emit some but not all tables, you'll need to set the individual sent_table fields directly.

To create an abbreviated image, you must also call jpeg_start_compress() with a second parameter of FALSE, not TRUE. Otherwise jpeg_start_compress() will force all the sent_table fields to FALSE. (This is a safety feature to prevent abbreviated images from being created accidentally.)

To create a tables-only file, perform the same parameter setup that you normally would, but instead of calling `jpeg_start_compress()` and so on, call `jpeg_write_tables(&cinfo)`. This will write an abbreviated datastream containing only SOI, DQT and/or DHT markers, and EOI. All the quantization and Huffman tables that are currently defined in the compression object will be emitted unless their `sent_tables` flag is already TRUE, and then all the `sent_tables` flags will be set TRUE.

A sure-fire way to create matching tables-only and abbreviated image files is to proceed as follows:

```
create JPEG compression object
set JPEG parameters
set destination to tables-only file
jpeg_write_tables(&cinfo);
set destination to image file
jpeg_start_compress(&cinfo, FALSE);
write data...
jpeg_finish_compress(&cinfo);
```

Since the JPEG parameters are not altered between writing the table file and the abbreviated image file, the same tables are sure to be used. Of course, you can repeat the `jpeg_start_compress() ... jpeg_finish_compress()` sequence many times to produce many abbreviated image files matching the table file.

You cannot suppress output of the computed Huffman tables when Huffman optimization is selected. (If you could, there'd be no way to decode the image...) Generally, you don't want to set `optimize_coding = TRUE` when you are trying to produce abbreviated files.

In some cases you might want to compress an image using tables which are not stored in the application, but are defined in an interchange or tables-only file readable by the application. This can be done by setting up a JPEG decompression object to read the specification file, then copying the tables into your compression object. See `jpeg_copy_critical_parameters()` for an example of copying quantization tables.

To read abbreviated image files, you simply need to load the proper tables into the decompression object before trying to read the abbreviated image. If the proper tables are stored in the application program, you can just allocate the table structs and fill in their contents directly. For example, to load a fixed quantization table into table slot "n":

```
if (cinfo.quant_tbl_ptrs[n] == NULL)
    cinfo.quant_tbl_ptrs[n] = jpeg_alloc_quant_table((j_common_ptr) &cinfo);
quant_ptr = cinfo.quant_tbl_ptrs[n]; /* quant_ptr is JQUANT_TBL* */
for (i = 0; i < 64; i++) {
```

```

/* Qtable[] is desired quantization table, in natural array order */
quant_ptr->quantval[i] = Qtable[i];
}

```

Code to load a fixed Huffman table is typically (for AC table "n"):

```

if (cinfo.ac_huff_tbl_ptrs[n] == NULL)
    cinfo.ac_huff_tbl_ptrs[n] = jpeg_alloc_huff_table((j_common_ptr) &cinfo);
huff_ptr = cinfo.ac_huff_tbl_ptrs[n]; /* huff_ptr is JHUFF_TBL* */
for (i = 1; i <= 16; i++) {
    /* counts[i] is number of Huffman codes of length i bits, i=1..16 */
    huff_ptr->bits[i] = counts[i];
}
for (i = 0; i < 256; i++) {
    /* symbols[] is the list of Huffman symbols, in code-length order */
    huff_ptr->huffval[i] = symbols[i];
}

```

(Note that trying to set `cinfo.quant_tbl_ptrs[n]` to point directly at a constant `JQUANT_TBL` object is not safe. If the incoming file happened to contain a quantization table definition, your master table would get overwritten! Instead allocate a working table copy and copy the master table into it, as illustrated above. Ditto for Huffman tables, of course.)

You might want to read the tables from a tables-only file, rather than hard-wiring them into your application. The `jpeg_read_header()` call is sufficient to read a tables-only file. You must pass a second parameter of `FALSE` to indicate that you do not require an image to be present. Thus, the typical scenario is

```

create JPEG decompression object
set source to tables-only file
jpeg_read_header(&cinfo, FALSE);
set source to abbreviated image file
jpeg_read_header(&cinfo, TRUE);
set decompression parameters
jpeg_start_decompress(&cinfo);
read data...
jpeg_finish_decompress(&cinfo);

```

In some cases, you may want to read a file without knowing whether it contains an image or just tables. In that case, pass `FALSE` and check the return value from `jpeg_read_header()`: it will be `JPEG_HEADER_OK` if an image was found, `JPEG_HEADER_TABLES_ONLY` if only tables were found. (A third return value, `JPEG_SUSPENDED`, is possible when using a suspending data source manager.) Note that `jpeg_read_header()` will not complain if you read an abbreviated image for which you haven't loaded the missing tables; the missing-table check occurs later, in `jpeg_start_decompress()`.

It is possible to read a series of images from a single source file by repeating the `jpeg_read_header() ... jpeg_finish_decompress()` sequence, without releasing/recreating the JPEG object or the data source module. (If you did reinitialize, any partial bufferload left in the data source buffer at the end of one image would be discarded, causing you to lose the start of the next image.) When you use this method, stored tables are automatically carried forward, so some of the images can be abbreviated images that depend on tables from earlier images.

If you intend to write a series of images into a single destination file, you might want to make a specialized data destination module that doesn't flush the output buffer at `term_destination()` time. This would speed things up by some trifling amount. Of course, you'd need to remember to flush the buffer after the last image. You can make the later images be abbreviated ones by passing `FALSE` to `jpeg_start_compress()`.

Special markers

Some applications may need to insert or extract special data in the JPEG datastream. The JPEG standard provides marker types "COM" (comment) and "APP0" through "APP15" (application) to hold application-specific data. Unfortunately, the use of these markers is not specified by the standard. COM markers are fairly widely used to hold user-supplied text. The JFIF file format spec uses APP0 markers with specified initial strings to hold certain data. Adobe applications use APP14 markers beginning with the string "Adobe" for miscellaneous data. Other APPn markers are rarely seen, but might contain almost anything.

If you wish to store user-supplied text, we recommend you use COM markers and place readable 7-bit ASCII text in them. Newline conventions are not standardized --- expect to find LF (Unix style), CR/LF (DOS style), or CR (Mac style). A robust COM reader should be able to cope with random binary garbage, including nulls, since some applications generate COM markers containing non-ASCII junk. (But yours should not be one of them.)

For program-supplied data, use an APPn marker, and be sure to begin it with an identifying string so that you can tell whether the marker is actually yours. It's probably best to avoid using APP0 or APP14 for any private markers. (NOTE: the upcoming SPIFF standard will use APP8 markers; we recommend you not use APP8 markers for any private purposes, either.)

Keep in mind that at most 65533 bytes can be put into one marker, but you can have as many markers as you like.

By default, the IJG compression library will write a JFIF APP0 marker if the selected JPEG colorspace is grayscale or YCbCr, or an Adobe APP14 marker if the selected colorspace is RGB, CMYK, or YCCK. You can disable this, but we don't recommend it. The decompression library will recognize JFIF and Adobe markers and will set the JPEG colorspace properly when one is found.

You can write special markers immediately following the datastream header by calling `jpeg_write_marker()` after `jpeg_start_compress()` and before the first call to `jpeg_write_scanlines()`. When you do this, the markers appear after the SOI and the JFIF APP0 and Adobe APP14 markers (if written), but before all else. Specify the marker type parameter as "JPEG_COM" for COM or "JPEG_APP0 + n" for APPn. (Actually, `jpeg_write_marker` will let you write any marker type, but we don't recommend writing any other kinds of marker.) For example, to write a user comment string pointed to by `comment_text`:

```
jpeg_write_marker(cinfo, JPEG_COM, comment_text, strlen(comment_text));
```

If it's not convenient to store all the marker data in memory at once, you can instead call `jpeg_write_m_header()` followed by multiple calls to `jpeg_write_m_byte()`. If you do it this way, it's your responsibility to call `jpeg_write_m_byte()` exactly the number of times given in the `length` parameter to `jpeg_write_m_header()`. (This method lets you empty the output buffer partway through a marker, which might be important when using a suspending data destination module. In any case, if you are using a suspending destination, you should flush its buffer after inserting any special markers. See "I/O suspension".)

Or, if you prefer to synthesize the marker byte sequence yourself, you can just cram it straight into the data destination module.

If you are writing JFIF 1.02 extension markers (thumbnail images), don't forget to set `cinfo.JFIF_minor_version = 2` so that the encoder will write the correct JFIF version number in the JFIF header marker. The library's default is to write version 1.01, but that's wrong if you insert any 1.02 extension markers. (We could probably get away with just defaulting to 1.02, but there used to be broken decoders that would complain about unknown minor version numbers. To reduce compatibility risks it's safest not to write 1.02 unless you are actually using 1.02 extensions.)

When reading, two methods of handling special markers are available:

1. You can ask the library to save the contents of COM and/or APPn markers into memory, and then examine them at your leisure afterwards.
2. You can supply your own routine to process COM and/or APPn markers on-the-fly as they are read.

The first method is simpler to use, especially if you are using a suspending data source; writing a marker processor that copes with input suspension is not easy (consider what happens if the marker is longer than your available

input buffer). However, the second method conserves memory since the marker data need not be kept around after it's been processed.

For either method, you'd normally set up marker handling after creating a decompression object and before calling `jpeg_read_header()`, because the markers of interest will typically be near the head of the file and so will be scanned by `jpeg_read_header`. Once you've established a marker handling method, it will be used for the life of that decompression object (potentially many datastreams), unless you change it. Marker handling is determined separately for COM markers and for each APPn marker code.

To save the contents of special markers in memory, call `jpeg_save_markers(cinfo, marker_code, length_limit)` where `marker_code` is the marker type to save, `JPEG_COM` or `JPEG_APP0+n`. (To arrange to save all the special marker types, you need to call this routine 17 times, for COM and APP0-APP15.) If the incoming marker is longer than `length_limit` data bytes, only `length_limit` bytes will be saved; this parameter allows you to avoid chewing up memory when you only need to see the first few bytes of a potentially large marker. If you want to save all the data, set `length_limit` to `0xFFFF`; that is enough since marker lengths are only 16 bits. As a special case, setting `length_limit` to 0 prevents that marker type from being saved at all. (That is the default behavior, in fact.)

After `jpeg_read_header()` completes, you can examine the special markers by following the `cinfo->marker_list` pointer chain. All the special markers in the file appear in this list, in order of their occurrence in the file (but omitting any markers of types you didn't ask for). Both the original data length and the saved data length are recorded for each list entry; the latter will not exceed `length_limit` for the particular marker type. Note that these lengths exclude the marker length word, whereas the stored representation within the JPEG file includes it. (Hence the maximum data length is really only 65533.)

It is possible that additional special markers appear in the file beyond the SOS marker at which `jpeg_read_header` stops; if so, the marker list will be extended during reading of the rest of the file. This is not expected to be common, however. If you are short on memory you may want to reset the length limit to zero for all marker types after finishing `jpeg_read_header`, to ensure that the `max_memory_to_use` setting cannot be exceeded due to addition of later markers.

The marker list remains stored until you call `jpeg_finish_decompress` or `jpeg_abort`, at which point the memory is freed and the list is set to empty. (`jpeg_destroy` also releases the storage, of course.)

Note that the library is internally interested in APP0 and APP14 markers; if you try to set a small nonzero length limit on these types, the library

will silently force the length up to the minimum it wants. (But you can set a zero length limit to prevent them from being saved at all.) Also, in a 16-bit environment, the maximum length limit may be constrained to less than 65533 by malloc() limitations. It is therefore best not to assume that the effective length limit is exactly what you set it to be.

If you want to supply your own marker-reading routine, you do it by calling `jpeg_set_marker_processor()`. A marker processor routine must have the signature

```
boolean jpeg_marker_parser_method (j_decompress_ptr cinfo)
```

Although the marker code is not explicitly passed, the routine can find it in `cinfo->unread_marker`. At the time of call, the marker proper has been read from the data source module. The processor routine is responsible for reading the marker length word and the remaining parameter bytes, if any. Return TRUE to indicate success. (FALSE should be returned only if you are using a suspending data source and it tells you to suspend. See the standard marker processors in `jdmarker.c` for appropriate coding methods if you need to use a suspending data source.)

If you override the default APP0 or APP14 processors, it is up to you to recognize JFIF and Adobe markers if you want colorspace recognition to occur properly. We recommend copying and extending the default processors if you want to do that. (A better idea is to save these marker types for later examination by calling `jpeg_save_markers()`; that method doesn't interfere with the library's own processing of these markers.)

`jpeg_set_marker_processor()` and `jpeg_save_markers()` are mutually exclusive --- if you call one it overrides any previous call to the other, for the particular marker type specified.

A simple example of an external COM processor can be found in `djpeg.c`. Also, see `jpegtran.c` for an example of using `jpeg_save_markers`.

Raw (downsampled) image data

Some applications need to supply already-downsampled image data to the JPEG compressor, or to receive raw downsampled data from the decompressor. The library supports this requirement by allowing the application to write or read raw data, bypassing the normal preprocessing or postprocessing steps. The interface is different from the standard one and is somewhat harder to use. If your interest is merely in bypassing color conversion, we recommend that you use the standard interface and simply set `jpeg_color_space = in_color_space` (or `jpeg_color_space = out_color_space` for decompression). The mechanism described in this section is necessary only to supply or receive downsampled image data, in which not all components have the same

dimensions.

To compress raw data, you must supply the data in the colorspace to be used in the JPEG file (please read the earlier section on Special color spaces) and downsampled to the sampling factors specified in the JPEG parameters. You must supply the data in the format used internally by the JPEG library, namely a JSAMPIMAGE array. This is an array of pointers to two-dimensional arrays, each of type JSAMPARRAY. Each 2-D array holds the values for one color component. This structure is necessary since the components are of different sizes. If the image dimensions are not a multiple of the MCU size, you must also pad the data correctly (usually, this is done by replicating the last column and/or row). The data must be padded to a multiple of a DCT block in each component: that is, each downsampled row must contain a multiple of `block_size` valid samples, and there must be a multiple of `block_size` sample rows for each component. (For applications such as conversion of digital TV images, the standard image size is usually a multiple of the DCT block size, so that no padding need actually be done.)

The procedure for compression of raw data is basically the same as normal compression, except that you call `jpeg_write_raw_data()` in place of `jpeg_write_scanlines()`. Before calling `jpeg_start_compress()`, you must do the following:

- * Set `cinfo->raw_data_in` to TRUE. (It is set FALSE by `jpeg_set_defaults()`.) This notifies the library that you will be supplying raw data. Furthermore, set `cinfo->do_fancy_downsampling` to FALSE if you want to use real downsampled data. (It is set TRUE by `jpeg_set_defaults()`.)
- * Ensure `jpeg_color_space` is correct --- an explicit `jpeg_set_colorspace()` call is a good idea. Note that since color conversion is bypassed, `in_color_space` is ignored, except that `jpeg_set_defaults()` uses it to choose the default `jpeg_color_space` setting.
- * Ensure the sampling factors, `cinfo->comp_info[i].h_samp_factor` and `cinfo->comp_info[i].v_samp_factor`, are correct. Since these indicate the dimensions of the data you are supplying, it's wise to set them explicitly, rather than assuming the library's defaults are what you want.

To pass raw data to the library, call `jpeg_write_raw_data()` in place of `jpeg_write_scanlines()`. The two routines work similarly except that `jpeg_write_raw_data` takes a JSAMPIMAGE data array rather than JSAMPARRAY. The scanlines count passed to and returned from `jpeg_write_raw_data` is measured in terms of the component with the largest `v_samp_factor`.

`jpeg_write_raw_data()` processes one MCU row per call, which is to say `v_samp_factor*block_size` sample rows of each component. The passed `num_lines` value must be at least `max_v_samp_factor*block_size`, and the return value will be exactly that amount (or possibly some multiple of that amount, in future library versions). This is true even on the last call at the bottom of the image; don't forget to pad your data as necessary.

The required dimensions of the supplied data can be computed for each component as

```
cinfo->comp_info[i].width_in_blocks*block_size samples per row
```

```
cinfo->comp_info[i].height_in_blocks*block_size rows in image
```

after `jpeg_start_compress()` has initialized those fields. If the valid data is smaller than this, it must be padded appropriately. For some sampling factors and image sizes, additional dummy DCT blocks are inserted to make the image a multiple of the MCU dimensions. The library creates such dummy blocks itself; it does not read them from your supplied data. Therefore you need never pad by more than `block_size` samples. An example may help here.

Assume 2h2v downsampling of YCbCr data, that is

```
cinfo->comp_info[0].h_samp_factor = 2 for Y
```

```
cinfo->comp_info[0].v_samp_factor = 2
```

```
cinfo->comp_info[1].h_samp_factor = 1 for Cb
```

```
cinfo->comp_info[1].v_samp_factor = 1
```

```
cinfo->comp_info[2].h_samp_factor = 1 for Cr
```

```
cinfo->comp_info[2].v_samp_factor = 1
```

and suppose that the nominal image dimensions (`cinfo->image_width` and `cinfo->image_height`) are 101x101 pixels. Then `jpeg_start_compress()` will compute `downsampled_width = 101` and `width_in_blocks = 13` for Y, `downsampled_width = 51` and `width_in_blocks = 7` for Cb and Cr (and the same for the height fields). You must pad the Y data to at least $13*8 = 104$ columns and rows, the Cb/Cr data to at least $7*8 = 56$ columns and rows. The MCU height is `max_v_samp_factor = 2` DCT rows so you must pass at least 16 scanlines on each call to `jpeg_write_raw_data()`, which is to say 16 actual sample rows of Y and 8 each of Cb and Cr. A total of 7 MCU rows are needed, so you must pass a total of $7*16 = 112$ "scanlines". The last DCT block row of Y data is dummy, so it doesn't matter what you pass for it in the data arrays, but the scanlines count must total up to 112 so that all of the Cb and Cr data gets passed.

Output suspension is supported with raw-data compression: if the data destination module suspends, `jpeg_write_raw_data()` will return 0.

In this case the same data rows must be passed again on the next call.

Decompression with raw data output implies bypassing all postprocessing.

You must deal with the color space and sampling factors present in the incoming file. If your application only handles, say, 2h1v YCbCr data, you must check for and fail on other color spaces or other sampling factors. The library will not convert to a different color space for you.

To obtain raw data output, set `cinfo->raw_data_out = TRUE` before `jpeg_start_decompress()` (it is set `FALSE` by `jpeg_read_header()`). Be sure to verify that the color space and sampling factors are ones you can handle. Furthermore, set `cinfo->do_fancy_upsampling = FALSE` if you want to get real downsampled data (it is set `TRUE` by `jpeg_read_header()`).

Then call `jpeg_read_raw_data()` in place of `jpeg_read_scanlines()`. The decompression process is otherwise the same as usual.

`jpeg_read_raw_data()` returns one MCU row per call, and thus you must pass a buffer of at least `max_v_samp_factor*block_size` scanlines (scanline counting is the same as for raw-data compression). The buffer you pass must be large enough to hold the actual data plus padding to DCT-block boundaries. As with compression, any entirely dummy DCT blocks are not processed so you need not allocate space for them, but the total scanline count includes them. The above example of computing buffer dimensions for raw-data compression is equally valid for decompression.

Input suspension is supported with raw-data decompression: if the data source module suspends, `jpeg_read_raw_data()` will return 0. You can also use buffered-image mode to read raw data in multiple passes.

Really raw data: DCT coefficients

It is possible to read or write the contents of a JPEG file as raw DCT coefficients. This facility is mainly intended for use in lossless transcoding between different JPEG file formats. Other possible applications include lossless cropping of a JPEG image, lossless reassembly of a multi-strip or multi-tile TIFF/JPEG file into a single JPEG datastream, etc.

To read the contents of a JPEG file as DCT coefficients, open the file and do `jpeg_read_header()` as usual. But instead of calling `jpeg_start_decompress()` and `jpeg_read_scanlines()`, call `jpeg_read_coefficients()`. This will read the entire image into a set of virtual coefficient-block arrays, one array per component. The return value is a pointer to an array of virtual-array descriptors. Each virtual array can be accessed directly using the JPEG memory manager's `access_virt_barray` method (see Memory management, below, and also `read.structure.txt`'s discussion of virtual array handling). Or, for simple transcoding to a different JPEG file format, the array list can just be handed directly to `jpeg_write_coefficients()`.

Each block in the block arrays contains quantized coefficient values in normal array order (not JPEG zigzag order). The block arrays contain only DCT blocks containing real data; any entirely-dummy blocks added to fill out interleaved MCUs at the right or bottom edges of the image are discarded during reading and are not stored in the block arrays. (The size of each block array can be determined from the `width_in_blocks` and `height_in_blocks` fields of the component's `comp_info` entry.) This is also the data format expected by `jpeg_write_coefficients()`.

When you are done using the virtual arrays, call `jpeg_finish_decompress()` to release the array storage and return the decompression object to an idle

state; or just call `jpeg_destroy()` if you don't need to reuse the object.

If you use a suspending data source, `jpeg_read_coefficients()` will return `NULL` if it is forced to suspend; a non-`NULL` return value indicates successful completion. You need not test for a `NULL` return value when using a non-suspending data source.

It is also possible to call `jpeg_read_coefficients()` to obtain access to the decoder's coefficient arrays during a normal decode cycle in buffered-image mode. This frammish might be useful for progressively displaying an incoming image and then re-encoding it without loss. To do this, decode in buffered-image mode as discussed previously, then call `jpeg_read_coefficients()` after the last `jpeg_finish_output()` call. The arrays will be available for your use until you call `jpeg_finish_decompress()`.

To write the contents of a JPEG file as DCT coefficients, you must provide the DCT coefficients stored in virtual block arrays. You can either pass block arrays read from an input JPEG file by `jpeg_read_coefficients()`, or allocate virtual arrays from the JPEG compression object and fill them yourself. In either case, `jpeg_write_coefficients()` is substituted for `jpeg_start_compress()` and `jpeg_write_scanlines()`. Thus the sequence is

- * Create compression object
- * Set all compression parameters as necessary
- * Request virtual arrays if needed
- * `jpeg_write_coefficients()`
- * `jpeg_finish_compress()`
- * Destroy or re-use compression object

`jpeg_write_coefficients()` is passed a pointer to an array of virtual block array descriptors; the number of arrays is equal to `cinfo.num_components`.

The virtual arrays need only have been requested, not realized, before `jpeg_write_coefficients()` is called. A side-effect of `jpeg_write_coefficients()` is to realize any virtual arrays that have been requested from the compression object's memory manager. Thus, when obtaining the virtual arrays from the compression object, you should fill the arrays after calling `jpeg_write_coefficients()`. The data is actually written out when you call `jpeg_finish_compress()`; `jpeg_write_coefficients()` only writes the file header.

When writing raw DCT coefficients, it is crucial that the JPEG quantization tables and sampling factors match the way the data was encoded, or the resulting file will be invalid. For transcoding from an existing JPEG file, we recommend using `jpeg_copy_critical_parameters()`. This routine initializes all the compression parameters to default values (like `jpeg_set_defaults()`), then copies the critical information from a source decompression object. The decompression object should have just been used to read the entire JPEG input file --- that is, it should be awaiting `jpeg_finish_decompress()`.

jpeg_write_coefficients() marks all tables stored in the compression object as needing to be written to the output file (thus, it acts like jpeg_start_compress(cinfo, TRUE)). This is for safety's sake, to avoid emitting abbreviated JPEG files by accident. If you really want to emit an abbreviated JPEG file, call jpeg_suppress_tables(), or set the tables' individual sent_table flags, between calling jpeg_write_coefficients() and jpeg_finish_compress().

Progress monitoring

Some applications may need to regain control from the JPEG library every so often. The typical use of this feature is to produce a percent-done bar or other progress display. (For a simple example, see cjpeg.c or djpeg.c.) Although you do get control back frequently during the data-transferring pass (the jpeg_read_scanlines or jpeg_write_scanlines loop), any additional passes will occur inside jpeg_finish_compress or jpeg_start_decompress; those routines may take a long time to execute, and you don't get control back until they are done.

You can define a progress-monitor routine which will be called periodically by the library. No guarantees are made about how often this call will occur, so we don't recommend you use it for mouse tracking or anything like that. At present, a call will occur once per MCU row, scanline, or sample row group, whichever unit is convenient for the current processing mode; so the wider the image, the longer the time between calls. During the data transferring pass, only one call occurs per call of jpeg_read_scanlines or jpeg_write_scanlines, so don't pass a large number of scanlines at once if you want fine resolution in the progress count. (If you really need to use the callback mechanism for time-critical tasks like mouse tracking, you could insert additional calls inside some of the library's inner loops.)

To establish a progress-monitor callback, create a struct jpeg_progress_mgr, fill in its progress_monitor field with a pointer to your callback routine, and set cinfo->progress to point to the struct. The callback will be called whenever cinfo->progress is non-NULL. (This pointer is set to NULL by jpeg_create_compress or jpeg_create_decompress; the library will not change it thereafter. So if you allocate dynamic storage for the progress struct, make sure it will live as long as the JPEG object does. Allocating from the JPEG memory manager with lifetime JPOOL_PERMANENT will work nicely.) You can use the same callback routine for both compression and decompression.

The jpeg_progress_mgr struct contains four fields which are set by the library:

```
long pass_counter; /* work units completed in this pass */
long pass_limit; /* total number of work units in this pass */
int completed_passes; /* passes completed so far */
```

```
int total_passes; /* total number of passes expected */
```

During any one pass, pass_counter increases from 0 up to (not including) pass_limit; the step size is usually but not necessarily 1. The pass_limit value may change from one pass to another. The expected total number of passes is in total_passes, and the number of passes already completed is in completed_passes. Thus the fraction of work completed may be estimated as

```
completed_passes + (pass_counter/pass_limit)
-----
total_passes
```

ignoring the fact that the passes may not be equal amounts of work.

When decompressing, pass_limit can even change within a pass, because it depends on the number of scans in the JPEG file, which isn't always known in advance. The computed fraction-of-work-done may jump suddenly (if the library discovers it has overestimated the number of scans) or even decrease (in the opposite case). It is not wise to put great faith in the work estimate.

When using the decompressor's buffered-image mode, the progress monitor work estimate is likely to be completely unhelpful, because the library has no way to know how many output passes will be demanded of it. Currently, the library sets total_passes based on the assumption that there will be one more output pass if the input file end hasn't yet been read (jpeg_input_complete() isn't TRUE), but no more output passes if the file end has been reached when the output pass is started. This means that total_passes will rise as additional output passes are requested. If you have a way of determining the input file size, estimating progress based on the fraction of the file that's been read will probably be more useful than using the library's value.

Memory management

This section covers some key facts about the JPEG library's built-in memory manager. For more info, please read structure.txt's section about the memory manager, and consult the source code if necessary.

All memory and temporary file allocation within the library is done via the memory manager. If necessary, you can replace the "back end" of the memory manager to control allocation yourself (for example, if you don't want the library to use malloc() and free() for some reason).

Some data is allocated "permanently" and will not be freed until the JPEG object is destroyed. Most data is allocated "per image" and is freed by jpeg_finish_compress, jpeg_finish_decompress, or jpeg_abort. You can call the memory manager yourself to allocate structures that will automatically be freed at these times. Typical code for this is

```
ptr = (*cinfo->mem->alloc_small) ((j_common_ptr) cinfo, JPOOL_IMAGE, size);
```

Use JPOOL_PERMANENT to get storage that lasts as long as the JPEG object.

Use `alloc_large` instead of `alloc_small` for anything bigger than a few Kbytes. There are also `alloc_sarray` and `alloc_barray` routines that automatically build 2-D sample or block arrays.

The library's minimum space requirements to process an image depend on the image's width, but not on its height, because the library ordinarily works with "strip" buffers that are as wide as the image but just a few rows high. Some operating modes (eg, two-pass color quantization) require full-image buffers. Such buffers are treated as "virtual arrays": only the current strip need be in memory, and the rest can be swapped out to a temporary file.

If you use the simplest memory manager back end (`jmemnobs.c`), then no temporary files are used; virtual arrays are simply `malloc()`'d. Images bigger than memory can be processed only if your system supports virtual memory. The other memory manager back ends support temporary files of various flavors and thus work in machines without virtual memory. They may also be useful on Unix machines if you need to process images that exceed available swap space.

When using temporary files, the library will make the in-memory buffers for its virtual arrays just big enough to stay within a "maximum memory" setting. Your application can set this limit by setting `cinfo->mem->max_memory_to_use` after creating the JPEG object. (Of course, there is still a minimum size for the buffers, so the max-memory setting is effective only if it is bigger than the minimum space needed.) If you allocate any large structures yourself, you must allocate them before `jpeg_start_compress()` or `jpeg_start_decompress()` in order to have them counted against the max memory limit. Also keep in mind that space allocated with `alloc_small()` is ignored, on the assumption that it's too small to be worth worrying about; so a reasonable safety margin should be left when setting `max_memory_to_use`.

If you use the `jmemname.c` or `jmemdos.c` memory manager back end, it is important to clean up the JPEG object properly to ensure that the temporary files get deleted. (This is especially crucial with `jmemdos.c`, where the "temporary files" may be extended-memory segments; if they are not freed, DOS will require a reboot to recover the memory.) Thus, with these memory managers, it's a good idea to provide a signal handler that will trap any early exit from your program. The handler should call either `jpeg_abort()` or `jpeg_destroy()` for any active JPEG objects. A handler is not needed with `jmemnobs.c`, and shouldn't be necessary with `jmemansi.c` or `jmemmac.c` either, since the C library is supposed to take care of deleting files made with `tmpfile()`.

Memory usage

Working memory requirements while performing compression or decompression depend on image dimensions, image characteristics (such as colorspace and

JPEG process), and operating mode (application-selected options).

As of v6b, the decompressor requires:

1. About 24K in more-or-less-fixed-size data. This varies a bit depending on operating mode and image characteristics (particularly color vs. grayscale), but it doesn't depend on image dimensions.
2. Strip buffers (of size proportional to the image width) for IDCT and upsampling results. The worst case for commonly used sampling factors is about 34 bytes * width in pixels for a color image. A grayscale image only needs about 8 bytes per pixel column.
3. A full-image DCT coefficient buffer is needed to decode a multi-scan JPEG file (including progressive JPEGs), or whenever you select buffered-image mode. This takes 2 bytes/coefficient. At typical 2x2 sampling, that's 3 bytes per pixel for a color image. Worst case (1x1 sampling) requires 6 bytes/pixel. For grayscale, figure 2 bytes/pixel.
4. To perform 2-pass color quantization, the decompressor also needs a 128K color lookup table and a full-image pixel buffer (3 bytes/pixel). This does not count any memory allocated by the application, such as a buffer to hold the final output image.

The above figures are valid for 8-bit JPEG data precision and a machine with 32-bit ints. For 9-bit to 12-bit JPEG data, double the size of the strip buffers and quantization pixel buffer. The "fixed-size" data will be somewhat smaller with 16-bit ints, larger with 64-bit ints. Also, CMYK or other unusual color spaces will require different amounts of space.

The full-image coefficient and pixel buffers, if needed at all, do not have to be fully RAM resident; you can have the library use temporary files instead when the total memory usage would exceed a limit you set. (But if your OS supports virtual memory, it's probably better to just use `jmemnobs` and let the OS do the swapping.)

The compressor's memory requirements are similar, except that it has no need for color quantization. Also, it needs a full-image DCT coefficient buffer if Huffman-table optimization is asked for, even if progressive mode is not requested.

If you need more detailed information about memory usage in a particular situation, you can enable the `MEM_STATS` code in `jmemmgr.c`.

Library compile-time options

A number of compile-time options are available by modifying `jmorecfg.h`.

The IJG code currently supports 8-bit to 12-bit sample data precision by defining `BITS_IN_JSAMPLE` as 8, 9, 10, 11, or 12.

Note that a value larger than 8 causes JSAMPLE to be larger than a char, so it affects the surrounding application's image data.

The sample applications cjpeg and djpeg can support deeper than 8-bit data only for PPM and GIF file formats; you must disable the other file formats to compile a 9-bit to 12-bit cjpeg or djpeg. (install.txt has more information about that.)

Run-time selection and conversion of data precision are currently not supported and may be added later.

Exception: The transcoding part (jpegtran) supports all settings in a single instance, since it operates on the level of DCT coefficients and not sample values.

(If you need to include an 8-bit library and a 9-bit to 12-bit library for compression or decompression in a single application, you could probably do it by defining NEED_SHORT_EXTERNAL_NAMES for just one of the copies. You'd have to access the 8-bit and the 9-bit to 12-bit copies from separate application source files. This is untested ... if you try it, we'd like to hear whether it works!)

Note that the standard Huffman tables are only valid for 8-bit data precision. If you selected more than 8-bit data precision, cjpeg uses arithmetic coding by default. The Huffman encoder normally uses entropy optimization to compute usable tables for higher precision. Otherwise, you'll have to supply different default Huffman tables. You may also want to supply your own DCT quantization tables; the existing quality-scaling code has been developed for 8-bit use, and probably doesn't generate especially good tables for 9-bit to 12-bit.

The maximum number of components (color channels) in the image is determined by MAX_COMPONENTS. The JPEG standard allows up to 255 components, but we expect that few applications will need more than four or so.

On machines with unusual data type sizes, you may be able to improve performance or reduce memory space by tweaking the various typedefs in jmorecfg.h. In particular, on some RISC CPUs, access to arrays of "short"s is quite slow; consider trading memory for speed by making JCOEF, INT16, and UINT16 be "int" or "unsigned int". UINT8 is also a candidate to become int. You probably don't want to make JSAMPLE be int unless you have lots of memory to burn.

You can reduce the size of the library by compiling out various optional functions. To do this, undefine xxx_SUPPORTED symbols as necessary.

You can also save a few K by not having text error messages in the library; the standard error message table occupies about 5Kb. This is particularly reasonable for embedded applications where there's no good way to display a message anyway. To do this, remove the creation of the message table (jpeg_std_message_table[]) from jerror.c, and alter format_message to do something reasonable without it. You could output the numeric value of the

message code number, for example. If you do this, you can also save a couple more K by modifying the TRACEMSn() macros in jerror.h to expand to nothing; you don't need trace capability anyway, right?

Portability considerations

The JPEG library has been written to be extremely portable; the sample applications cjpeg and djpeg are slightly less so. This section summarizes the design goals in this area. (If you encounter any bugs that cause the library to be less portable than is claimed here, we'd appreciate hearing about them.)

The code works fine on ANSI C, C++, and pre-ANSI C compilers, using any of the popular system include file setups, and some not-so-popular ones too. See install.txt for configuration procedures.

The code is not dependent on the exact sizes of the C data types. As distributed, we make the assumptions that

- char is at least 8 bits wide
- short is at least 16 bits wide
- int is at least 16 bits wide
- long is at least 32 bits wide

(These are the minimum requirements of the ANSI C standard.) Wider types will work fine, although memory may be used inefficiently if char is much larger than 8 bits or short is much bigger than 16 bits. The code should work equally well with 16- or 32-bit ints.

In a system where these assumptions are not met, you may be able to make the code work by modifying the typedefs in jmorecfg.h. However, you will probably have difficulty if int is less than 16 bits wide, since references to plain int abound in the code.

char can be either signed or unsigned, although the code runs faster if an unsigned char type is available. If char is wider than 8 bits, you will need to redefine JOCTET and/or provide custom data source/destination managers so that JOCTET represents exactly 8 bits of data on external storage.

The JPEG library proper does not assume ASCII representation of characters. But some of the image file I/O modules in cjpeg/djpeg do have ASCII dependencies in file-header manipulation; so does cjpeg's select_file_type() routine.

The JPEG library does not rely heavily on the C library. In particular, C stdio is used only by the data source/destination modules and the error handler, all of which are application-replaceable. (cjpeg/djpeg are more heavily dependent on stdio.) malloc and free are called only from the memory

manager "back end" module, so you can use a different memory allocator by replacing that one file.

The code generally assumes that C names must be unique in the first 15 characters. However, global function names can be made unique in the first 6 characters by defining `NEED_SHORT_EXTERNAL_NAMES`.

More info about porting the code may be gleaned by reading `jconfig.txt`, `jmorecfg.h`, and `jinclude.h`.

Notes for MS-DOS implementors

The IJG code is designed to work efficiently in 80x86 "small" or "medium" memory models (i.e., data pointers are 16 bits unless explicitly declared "far"; code pointers can be either size). You may be able to use small model to compile `cjpeg` or `djpeg` by itself, but you will probably have to use medium model for any larger application. This won't make much difference in performance. You *will* take a noticeable performance hit if you use a large-data memory model (perhaps 10%-25%), and you should avoid "huge" model if at all possible.

The JPEG library typically needs 2Kb-3Kb of stack space. It will also `malloc` about 20K-30K of near heap space while executing (and lots of far heap, but that doesn't count in this calculation). This figure will vary depending on selected operating mode, and to a lesser extent on image size. There is also about 5Kb-6Kb of constant data which will be allocated in the near data segment (about 4Kb of this is the error message table). Thus you have perhaps 20K available for other modules' static data and near heap space before you need to go to a larger memory model. The C library's static data will account for several K of this, but that still leaves a good deal for your needs. (If you are tight on space, you could reduce the sizes of the I/O buffers allocated by `jdatasrc.c` and `jdatadst.c`, say from 4K to 1K. Another possibility is to move the error message table to far memory; this should be doable with only localized hacking on `jerror.c`.)

About 2K of the near heap space is "permanent" memory that will not be released until you destroy the JPEG object. This is only an issue if you save a JPEG object between compression or decompression operations.

Far data space may also be a tight resource when you are dealing with large images. The most memory-intensive case is decompression with two-pass color quantization, or single-pass quantization to an externally supplied color map. This requires a 128Kb color lookup table plus strip buffers amounting to about 40 bytes per column for typical sampling ratios (eg, about 25600 bytes for a 640-pixel-wide image). You may not be able to process wide images if you have large data structures of your own.

Of course, all of these concerns vanish if you use a 32-bit flat-memory-model compiler, such as DJGPP or Watcom C. We highly recommend flat model if you can use it; the JPEG library is significantly faster in flat model.

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/libjpeg.txt

No license file was found, but licenses were detected in source scan.

/*

* jfdctflt.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2003-2017 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains a floating-point implementation of the

* forward DCT (Discrete Cosine Transform).

*

* This implementation should be more accurate than either of the integer

* DCT implementations. However, it may not give the same results on all

* machines because of differences in roundoff behavior. Speed will depend

* on the hardware's floating point capacity.

*

* A 2-D DCT can be done by 1-D DCT on each row followed by 1-D DCT

* on each column. Direct algorithms are also available, but they are

* much more complex and seem not to be any faster when reduced to code.

*

* This implementation is based on Arai, Agui, and Nakajima's algorithm for

* scaled DCT. Their original paper (Trans. IEICE E-71(11):1095) is in

* Japanese, but the algorithm is described in the Pennebaker & Mitchell

* JPEG textbook (see REFERENCES section in file README). The following code

* is based directly on figure 4-8 in P&M.

* While an 8-point DCT cannot be done in less than 11 multiplies, it is

* possible to arrange the computation so that many of the multiplies are

* simple scalings of the final outputs. These multiplies can then be

* folded into the multiplications or divisions by the JPEG quantization

* table entries. The AA&N method leaves only 5 multiplies and 29 adds

* to be done in the DCT itself.

* The primary disadvantage of this method is that with a fixed-point

* implementation, accuracy is lost due to imprecise representation of the

* scaled quantization values. However, that problem does not arise if

* we use floating point arithmetic.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jfdctflt.c

No license file was found, but licenses were detected in source scan.

```
/*
 * jdtrans.c
 *
 * Copyright (C) 1995-1997, Thomas G. Lane.
 * Modified 2000-2009 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains library routines for transcoding decompression,
 * that is, reading raw DCT coefficient arrays from an input JPEG file.
 * The routines in jdapimin.c will also be needed by a transcoder.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdtrans.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * wrbmp.c
 *
 * Copyright (C) 1994-1996, Thomas G. Lane.
 * Modified 2017 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains routines to write output images in Microsoft "BMP"
 * format (MS Windows 3.x and OS/2 1.x flavors).
 * Either 8-bit colormapped or 24-bit full-color format can be written.
 * No compression is supported.
 *
 * These routines may need modification for non-Unix environments or
 * specialized applications. As they stand, they assume output to
 * an ordinary stdio stream.
 *
 * This code contributed by James Arthur Boucher.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/wrbmp.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * jinclude.h
 *
 * Copyright (C) 1991-1994, Thomas G. Lane.
 * Modified 2017 by Guido Vollbeding.
```

* This file is part of the Independent JPEG Group's software.
* For conditions of distribution and use, see the accompanying README file.
*
* This file exists to provide a single place to fix any problems with
* including the wrong system include files. (Common problems are taken
* care of by the standard jconfig symbols, but on really weird systems
* you may have to edit this file.)
*
* NOTE: this file is NOT intended to be included by applications using the
* JPEG library. Most applications need only include jpeglib.h.
*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jinclude.h
No license file was found, but licenses were detected in source scan.

/*

* jmemansi.c

*

* Copyright (C) 1992-1996, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file provides a simple generic implementation of the system-

* dependent portion of the JPEG memory manager. This implementation

* assumes that you have the ANSI-standard library routine tmpfile().

* Also, the problem of determining the amount of memory available

* is shoved onto the user.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jmemansi.c
No license file was found, but licenses were detected in source scan.

/*

* jcapistd.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains application interface code for the compression half

* of the JPEG library. These are the "standard" API routines that are

* used in the normal full-compression case. They are not used by a

* transcoding-only application. Note that if an application links in

* jpeg_start_compress, it will end up linking in the entire compressor.

* We thus must separate this file from jcapimin.c to avoid linking the

* whole compression library into a transcoder.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jcapistd.c

No license file was found, but licenses were detected in source scan.

/*

* transupp.h

*

* Copyright (C) 1997-2013, Thomas G. Lane, Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains declarations for image transformation routines and

* other utility code used by the jpegtran sample application. These are

* NOT part of the core JPEG library. But we keep these routines separate

* from jpegtran.c to ease the task of maintaining jpegtran-like programs

* that have other user interfaces.

*

* NOTE: all the routines declared here have very specific requirements

* about when they are to be executed during the reading and writing of the

* source and destination files. See the comments in transupp.c, or see

* jpegtran.c for an example of correct usage.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/transupp.h

No license file was found, but licenses were detected in source scan.

/*

* rdgif.c

*

* Copyright (C) 1991-1997, Thomas G. Lane.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains routines to read input images in GIF format.

*

* NOTE: to avoid entanglements with Unisys' patent on LZW compression, *

* the ability to read GIF files has been removed from the IJG distribution. *

* Sorry about that. *

*

* We are required to state that

* "The Graphics Interchange Format(c) is the Copyright property of

* CompuServe Incorporated. GIF(sm) is a Service Mark property of

* CompuServe Incorporated."

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/rdgif.c

No license file was found, but licenses were detected in source scan.

/*

* jerror.h

*

* Copyright (C) 1994-1997, Thomas G. Lane.

* Modified 1997-2012 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file defines the error and message codes for the JPEG library.

* Edit this file to add new codes, or to translate the message strings to

* some other language.

* A set of error-reporting macros are defined too. Some applications using

* the JPEG library may wish to include this file to get the error codes

* and/or the macros.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jerror.h

No license file was found, but licenses were detected in source scan.

/*

* jdapistd.c

*

* Copyright (C) 1994-1996, Thomas G. Lane.

* Modified 2002-2013 by Guido Vollbeding.

* This file is part of the Independent JPEG Group's software.

* For conditions of distribution and use, see the accompanying README file.

*

* This file contains application interface code for the decompression half

* of the JPEG library. These are the "standard" API routines that are

* used in the normal full-decompression case. They are not used by a

* transcoding-only application. Note that if an application links in

* jpeg_start_decompress, it will end up linking in the entire decompressor.

* We thus must separate this file from jdapimin.c to avoid linking the

* whole decompression library into a transcoder.

*/

Found in path(s):

* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdapistd.c

No license file was found, but licenses were detected in source scan.

```
/*
 * jerror.c
 *
 * Copyright (C) 1991-1998, Thomas G. Lane.
 * Modified 2012-2015 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains simple error-reporting and trace-message routines.
 * These are suitable for Unix-like systems and others where writing to
 * stderr is the right thing to do. Many applications will want to replace
 * some or all of these routines.
 *
 * If you define USE_WINDOWS_MESSAGEBOX in jconfig.h or in the makefile,
 * you get a Windows-specific hack to display error messages in a dialog box.
 * It ain't much, but it beats dropping error messages into the bit bucket,
 * which is what happens to output to stderr under most Windows C compilers.
 *
 * These routines are used by both the compression and decompression code.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jerror.c
```

No license file was found, but licenses were detected in source scan.

```
/*
 * jdapimin.c
 *
 * Copyright (C) 1994-1998, Thomas G. Lane.
 * Modified 2009-2013 by Guido Vollbeding.
 * This file is part of the Independent JPEG Group's software.
 * For conditions of distribution and use, see the accompanying README file.
 *
 * This file contains application interface code for the decompression half
 * of the JPEG library. These are the "minimum" API routines that may be
 * needed in either the normal full-decompression case or the
 * transcoding-only case.
 *
 * Most of the routines intended to be called directly by an application
 * are in this file or in jdapistd.c. But also see jcomapi.c for routines
 * shared by compression and decompression, and jdtrans.c for the transcoding
 * case.
 */
```

Found in path(s):

```
* /opt/cola/permits/1103638654_1611231513.28/0/jpeg-9c-tar-gz/jpeg-9c/jdapimin.c
```


1.5 openssl 1.0.2s

1.5.1 Notifications :

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

1.5.2 Available under license :

LICENSE ISSUES

=====

The OpenSSL toolkit stays under a double license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit.

See below for the actual license texts.

OpenSSL License

/* =====

* Copyright (c) 1998-2019 The OpenSSL Project. All rights reserved.

*

* Redistribution and use in source and binary forms, with or without

* modification, are permitted provided that the following conditions

* are met:

*

* 1. Redistributions of source code must retain the above copyright

* notice, this list of conditions and the following disclaimer.

*

* 2. Redistributions in binary form must reproduce the above copyright

* notice, this list of conditions and the following disclaimer in

* the documentation and/or other materials provided with the

* distribution.

*

* 3. All advertising materials mentioning features or use of this

* software must display the following acknowledgment:

* "This product includes software developed by the OpenSSL Project

* for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

*

* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to

* endorse or promote products derived from this software without

* prior written permission. For written permission, please contact

* openssl-core@openssl.org.

*

* 5. Products derived from this software may not be called "OpenSSL"

* nor may "OpenSSL" appear in their names without prior written

* permission of the OpenSSL Project.

```

*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* =====
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/

```

Original SSLeay License

```

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or

```

- * in documentation (online or textual) provided with the package.
- *
- * Redistribution and use in source and binary forms, with or without
- * modification, are permitted provided that the following conditions
- * are met:
- * 1. Redistributions of source code must retain the copyright
- * notice, this list of conditions and the following disclaimer.
- * 2. Redistributions in binary form must reproduce the above copyright
- * notice, this list of conditions and the following disclaimer in the
- * documentation and/or other materials provided with the distribution.
- * 3. All advertising materials mentioning features or use of this software
- * must display the following acknowledgement:
- * "This product includes cryptographic software written by
- * Eric Young (eay@cryptsoft.com)"
- * The word 'cryptographic' can be left out if the routines from the library
- * being used are not cryptographic related :-).
- * 4. If you include any Windows specific code (or a derivative thereof) from
- * the apps directory (application code) you must include an acknowledgement:
- * "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
- *
- * THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND
- * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
- * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
- * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
- * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
- * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
- * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
- * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
- * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
- * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
- * SUCH DAMAGE.
- *
- * The licence and distribution terms for any publically available version or
- * derivative of this code cannot be changed. i.e. this code cannot simply be
- * copied and put under another distribution licence
- * [including the GNU Public Licence.]

*/

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place - Suite 330, Boston, MA

02111-1307, USA.

Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a

notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this

License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free

programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>  
Copyright (C) 19yy <name of author>
```

```
This program is free software; you can redistribute it and/or modify  
it under the terms of the GNU General Public License as published by  
the Free Software Foundation; either version 2 of the License, or
```

(at your option) any later version.

This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this
when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) 19yy name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.  
This is free software, and you are welcome to redistribute it  
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate
parts of the General Public License. Of course, the commands you use may
be called something other than `show w' and `show c'; they could even be
mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the program, if
necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program  
'Gnomovision' (which makes passes at compilers) written by James Hacker.
```

```
<signature of Ty Coon>, 1 April 1989  
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into
proprietary programs. If your program is a subroutine library, you may
consider it more useful to permit linking proprietary applications with the
library. If this is what you want to do, use the GNU Library General
Public License instead of this License.

The "Artistic License"

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:

a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b) use the modified Package only within your corporation or organization.

c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:

a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b) accompany the distribution with the machine-readable source of the Package with your modifications.

c) give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whoever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a

binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.

7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.

8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

1.6 dotnetzip 1.9.1.8

1.6.1 Available under license :

The following licenses govern use of the accompanying software, the DotNetZip library ("the software"). If you use the software, you accept these licenses. If you do not accept the license, do not use the software.

The managed ZLIB code included in Ionic.Zlib.dll and Ionic.Zip.dll is modified code, based on jzlib.

The following notice applies to jzlib:

Copyright (c) 2000,2001,2002,2003 ymnk, JCraft,Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

jzlib is based on zlib-1.1.3.

The following notice applies to zlib:

Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler

The ZLIB software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.

3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

Microsoft Public License (Ms-PL)

This license governs use of the accompanying software, the DotNetZip library ("the software"). If you use the software, you accept this license. If you do not accept the license, do not use the software.

1. Definitions

The terms "reproduce," "reproduction," "derivative works," and "distribution" have the same meaning here as under U.S. copyright law.

A "contribution" is the original software, or any additions or changes to the software.

A "contributor" is any person that distributes its contribution under this license.

"Licensed patents" are a contributor's patent claims that read directly on its contribution.

2. Grant of Rights

(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.

(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.

3. Conditions and Limitations

(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.

(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.

(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.

(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in

compiled or object code form, you may only do so under a license that complies with this license.

(E) The software is licensed "as-is." You bear the risk of using it. The contributors give no express warranties, guarantees or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-infringement.

The managed BZIP2 code included in Ionic.BZip2.dll and Ionic.Zip.dll is modified code, based on the bzip2 code in the Apache commons compress library.

The original BZip2 was created by Julian Seward, and is licensed under the BSD license.

The following license applies to the Apache code:

```
-----  
  
/*  
* Licensed to the Apache Software Foundation (ASF) under one  
* or more contributor license agreements. See the NOTICE file  
* distributed with this work for additional information  
* regarding copyright ownership. The ASF licenses this file  
* to you under the Apache License, Version 2.0 (the  
* "License"); you may not use this file except in compliance  
* with the License. You may obtain a copy of the License at  
*  
* http://www.apache.org/licenses/LICENSE-2.0  
*  
* Unless required by applicable law or agreed to in writing,  
* software distributed under the License is distributed on an  
* "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY  
* KIND, either express or implied. See the License for the  
* specific language governing permissions and limitations  
* under the License.  
*/
```

1.7 ionic-zip 1.9.1.8

1.7.1 Available under license :

The ZLIB library, available as Ionic.Zlib.dll or as part of DotNetZip, is a ported-then-modified version of jzlib, which itself is based on zlib-1.1.3, the well-known C-language compression library.

The following notice applies to zlib:

```
-----
```


Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler

The ZLIB software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu

The following licenses govern use of the accompanying software, the DotNetZip library ("the software"). If you use the software, you accept these licenses. If you do not accept the license, do not use the software.

The managed ZLIB code included in Ionic.Zlib.dll and Ionic.Zip.dll is modified code, based on jzlib.

The following notice applies to jzlib:

Copyright (c) 2000,2001,2002,2003 ymnk, JCraft,Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

jzlib is based on zlib-1.1.3.

The following notice applies to zlib:

Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler

The ZLIB software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org
Mark Adler madler@alumni.caltech.edu

Apache Commons Compress
Copyright 2002-2010 The Apache Software Foundation

This product includes software developed by
The Apache Software Foundation (<http://www.apache.org/>).
The ZLIB library, available as Ionic.Zlib.dll or as part of DotNetZip,
is a ported-then-modified version of jzlib. The following applies to jzlib:

JZlib 0.0.* were released under the GNU LGPL license. Later, we have switched
over to a BSD-style license.

Copyright (c) 2000,2001,2002,2003 ymnk, JCraft,Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in
the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products
derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES,
INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND
FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT,
INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA,
OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE,
EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Microsoft Public License (Ms-PL)

This license governs use of the accompanying software, the DotNetZip library ("the software"). If you use the
software, you accept this license. If you do not accept the license, do not use the software.

1. Definitions

The terms "reproduce," "reproduction," "derivative works," and "distribution" have the same meaning here as under
U.S. copyright law.

A "contribution" is the original software, or any additions or changes to the software.

A "contributor" is any person that distributes its contribution under this license.

"Licensed patents" are a contributor's patent claims that read directly on its contribution.

2. Grant of Rights

(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.

(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.

3. Conditions and Limitations

(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.

(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.

(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.

(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in compiled or object code form, you may only do so under a license that complies with this license.

(E) The software is licensed "as-is." You bear the risk of using it. The contributors give no express warranties, guarantees or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-infringement.

The managed BZIP2 code included in Ionic.BZip2.dll and Ionic.Zip.dll is modified code, based on the bzip2 code in the Apache commons compress library.

The original BZip2 was created by Julian Seward, and is licensed under the BSD license.

The following license applies to the Apache code:

```
/*  
* Licensed to the Apache Software Foundation (ASF) under one  
* or more contributor license agreements. See the NOTICE file  
* distributed with this work for additional information
```

* regarding copyright ownership. The ASF licenses this file
 * to you under the Apache License, Version 2.0 (the
 * "License"); you may not use this file except in compliance
 * with the License. You may obtain a copy of the License at
 *
 * <http://www.apache.org/licenses/LICENSE-2.0>
 *
 * Unless required by applicable law or agreed to in writing,
 * software distributed under the License is distributed on an
 * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY
 * KIND, either express or implied. See the License for the
 * specific language governing permissions and limitations
 * under the License.
 */

1.8 xerces-c 3.2.3

1.8.1 Available under license :

```
=====
== NOTICE file corresponding to section 4(d) of the Apache License, ==
== Version 2.0, in this case for the Apache Xerces distribution. ==
=====
```

This product includes software developed by
 The Apache Software Foundation (<http://www.apache.org/>).

Portions of this software were originally based on the following:
 - software copyright (c) 1999, IBM Corporation., <http://www.ibm.com>.

Apache License
 Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction,
 and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by
 the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all
 other entities that control, are controlled by, or are under common
 control with that entity. For the purposes of this definition,
 "control" means (i) the power, direct or indirect, to cause the

direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of

this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and

wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions.

Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor

has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

1.9 log4net 2.0.8.0-.NET

1.9.1 Available under license :

Apache log4net

Copyright 2004-2017 The Apache Software Foundation

This product includes software developed at

The Apache Software Foundation (<http://www.apache.org/>).

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or

agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");

you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

1.10 zlib 1.1.4

1.10.1 Available under license :

Copyright (C) 1995-2002 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler
jloup@gzip.org madler@alumni.caltech.edu

The data format used by the zlib library is described by RFCs (Request for Comments) 1950 to 1952 in the files <ftp://ds.internic.net/rfc/rfc1950.txt> (zlib format), [rfc1951.txt](ftp://ds.internic.net/rfc/rfc1951.txt) (deflate format) and [rfc1952.txt](ftp://ds.internic.net/rfc/rfc1952.txt) (gzip format).

*/

1.11 sqlite 3.7.16.2

1.11.1 Available under license :

The author disclaims copyright to this source code. In place of a legal notice, here is a blessing:

May you do good and not evil.

May you find forgiveness for yourself and forgive others.

May you share freely, never taking more than you give.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

©2021 Cisco Systems, Inc. All rights reserved.