

Package ‘fstcore’

May 4, 2020

Type Package

Title R Bindings to the 'Fstlib' Library

Description

The 'fstlib' library provides multithreaded serialization of compressed data frames using the 'fst' format. The 'fst' format allows for random access of stored data and compression with the 'LZ4' and 'ZSTD' compressors.

Version 0.9.4

Date 2020-04-28

LazyData true

Depends R (>= 3.0.0)

Imports Rcpp

LinkingTo Rcpp

SystemRequirements little-endian platform

RoxygenNote 7.1.0

Suggests testthat, lintr

License AGPL-3

Encoding UTF-8

Copyright This package includes sources from the LZ4 library owned by Yann Collet, sources of the ZSTD library owned by Facebook, Inc., sources of the libdivsufsort-lite library owned by Yuta Mori and sources of the fstlib library owned by Mark Klik

URL <https://fstpackage.github.io/fstcore>

BugReports <https://github.com/fstpackage/fst/issues>

NeedsCompilation yes

Author Mark Klik [aut, cre, cph],

Yuta Mori [ctb, cph] (Yuta Mori is author and copyright holder of the bundled 'libdivsufsort-lite' code, part of 'ZSTD'),
Przemysław Skibinski [ctb, cph] (Przemysław Skibinski is author and copyright holder of the bundled price functions, part of 'ZSTD'),

Tino Reichardt [ctb, cph] (Tino Reichardt is author and copyright holder of bundled sources from the 'zstdmt' library, part of 'ZSTD'),

Yann Collet [ctb, cph] (Yann Collet is author of the bundled 'LZ4' and 'ZSTD' code and copyright holder of 'LZ4'),

Facebook, Inc. [cph] (Bundled 'ZSTD' code)

Maintainer Mark Klik <markklik@gmail.com>

Repository CRAN

Date/Publication 2020-05-04 13:40:02 UTC

R topics documented:

fstcore-package	2
threads_fstlib	5

Index	7
--------------	----------

fstcore-package	<i>R bindings to the fstlib library</i>
-----------------	---

Description

R package `fstcore` contains R bindings to the C++ `fstlib` library which allows interfacing with `fst` files. It also contains the LZ4 and ZSTD sources used for compression. `fstcore` exists as a package separate from the `fst` package to facilitate independent updates to the `fstlib`, LZ4 and ZSTD libraries and is used as a backend to `fst`.

Details

The `fstcore` package is based on six C++ libraries or components:

- **fstlib**: library containing code to write, read and compute on files stored in the `fst` format. Written and owned by Mark Klik.
- **LZ4**: library containing code to compress data with the LZ4 compressor. Written and owned by Yann Collet.
- **ZSTD**: library containing code to compress data with the ZSTD compressor. Written by Yann Collet and owned by Facebook, Inc.
- **libdivsufsort-lite**: a lightweight suffix array construction algorithm. This code is bundled with the ZSTD compression library. Written and owned by Yuta Mori.
- **sources from zstdmt**: a multithreading library for Brotli, Lizard, LZ4, LZ5 and Zstandard. This code is bundled with the ZSTD library. Written and owned by Tino Reichardt.
- **source file zstd_opt.h**: Price functions for optimal parser. Written and owned by Yann Collet and Przemyslaw Skibinski. This code is bundled with the ZSTD library.

The following copyright notice, list of conditions and disclaimer apply to the use of the ZSTD library in the fstcore package:

BSD License

For Zstandard software

Copyright (c) 2016-present, Facebook, 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:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 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.
- Neither the name Facebook nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "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 COPYRIGHT HOLDER 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 following copyright notice, list of conditions and disclaimer apply to the use of the LZ4 library in the fstcore package:

LZ4 Library Copyright (c) 2011-2016, Yann Collet All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 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.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "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 COPYRIGHT HOLDER 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 following copyright notice, list of conditions and disclaimer apply to the use of the fstlib library in the fstcore package:

fstlib - A C++ library for ultra fast storage and retrieval of datasets

Copyright (C) 2017-present, Mark AJ Klik

This file is part of fstlib.

fstlib is free software: you can redistribute it and/or modify it under the terms of the GNU Affero General Public License version 3 as published by the Free Software Foundation.

fstlib 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 Affero General Public License for more details.

You should have received a copy of the GNU Affero General Public License along with fstlib. If not, see <http://www.gnu.org/licenses/>.

The following copyright notice, list of conditions and disclaimer apply to the use of the libdivsufsort-lite library in the fstcore package:

Copyright (c) 2003-2008 Yuta Mori All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

The following copyright notice, list of conditions and disclaimer apply to the use of sources from the zstdmt library included in the fstcore package as part the ZSTD library:

Copyright (c) 2016 Tino Reichardt All rights reserved.

This source code is licensed under both the BSD-style license (found in the LICENSE file in the root directory of this source tree) and the GPLv2 (found in the COPYING file in the root directory of this source tree).

You can contact the author at:

- zstdmt source repository: <https://github.com/mcmilk/zstdmt>

The following copyright notice, list of conditions and disclaimer apply to the use of `zstd_opt.h` included in the `fstcore` package as part the `ZSTD` library:

Copyright (c) 2016-present, Przemyslaw Skibinski, Yann Collet, Facebook, Inc. All rights reserved.

This source code is licensed under the BSD-style license found in the LICENSE file in the root directory of <https://github.com/facebook/zstd>.

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License version 2 as published by the Free Software Foundation. This program is dual-licensed; you may select either version 2 of the GNU General Public License ("GPL") or BSD license ("BSD").

threads_fstlib

Get or set the number of threads used in parallel operations

Description

For parallel operations, the performance is determined to a great extent by the number of threads used. More threads will allow the CPU to perform more computational intensive tasks simultaneously, speeding up the operation. Using more threads also introduces some overhead that will scale with the number of threads used. Therefore, using the maximum number of available threads is not always the fastest solution. With `threads_fstlib` the number of threads can be adjusted to the users specific requirements. As a default, `fstcore` uses a number of threads equal to the number of logical cores in the system.

Usage

```
threads_fstlib(nr_of_threads = NULL, reset_after_fork = NULL)
```

Arguments

`nr_of_threads` number of threads to use or NULL to get the current number of threads used in multithreaded operations.

`reset_after_fork`

when `fstcore` is running in a forked process, the usage of OpenMP can create problems. To prevent these, `fstcore` switches back to single core usage when it detects a fork. After the fork, the number of threads is reset to its initial setting. However, on some compilers (e.g. Intel), switching back to multi-threaded mode can lead to issues. When `reset_after_fork` is set to FALSE, `fstcore` is left in single-threaded mode after the fork ends. After the fork, multithreading can be activated again manually by calling `threads_fstlib` with an appropriate value for `nr_of_threads`. The default (`reset_after_fork = NULL`) leaves the fork behavior unchanged.

Details

The number of threads can also be set with `options(fst_threads = N)`. NOTE: This option is only read when the package's namespace is first loaded, with commands like `library`, `require`, or `::`. If you have already used one of these, you must use `threads_fstlib` to set the number of threads.

Value

the number of threads (previously) used

Examples

```
# get current number of threads
threads_fstlib()

# set the number of threads
threads_fstlib(12)

# leave in single threaded mode after a fork
threads_fstlib(12, reset_after_fork = FALSE)

# reset number of threads after a fork
threads_fstlib(12, reset_after_fork = TRUE)
```

Index

`fstcore-package`, [2](#)

`threads_fstlib`, [5](#)