

NAME

parsort - Sort (big files) in parallel

SYNOPSIS

parsort *options for sort*

DESCRIPTION

parsort uses GNU **sort** to sort in parallel. It works just like **sort** but faster on inputs with more than 1 M lines, if you have a multicore machine.

Hopefully these ideas will make it into GNU **sort** in the future.

OPTIONS

Same as **sort**. Except:

--parallel=N

Change the number of sorts run concurrently to *N*. *N* will be increased to number of files if **parsort** is given more than *N* files.

EXAMPLE

Sort files:

```
parsort *.txt > sorted.txt
```

Sort stdin (standard input) numerically:

```
cat numbers | parsort -n > sorted.txt
```

PERFORMANCE

parsort is faster on files than on stdin (standard input), because different parts of a file can be read in parallel.

On a 48 core machine you should see a speedup of 3x over **sort**.

AUTHOR

Copyright (C) 2020-2024 Ole Tange, <http://ole.tange.dk> and Free Software Foundation, Inc.

LICENSE

Copyright (C) 2012 Free Software Foundation, Inc.

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 3 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, see <http://www.gnu.org/licenses/>.

DEPENDENCIES

parsort uses **sort**, **bash**, and **parallel**.

SEE ALSO

sort