NAME

cdparanoia − an audio CD reading utility which includes extra data verification features

SYNOPSIS

cdparanoia [**options**] **span** [**outfile**] |**−B**

DESCRIPTION

cdparanoia retrieves audio tracks from CDDA-capable CDROM drives. The data can be saved to a file or directed to standard output in WAV, AIFF, AIFF-C or raw format. Most ATAPI and SCSI and several proprietary CDROM drive makes are supported; **cdparanoia** can determine if the target drive is CDDA capable.

In addition to simple reading, **cdparanoia** adds extra-robust data verification, synchronization, error handling and scratch reconstruction capability.

OPTIONS

−A --analyze-drive
Run and log a complete analysis of drive caching, timing and reading behavior; verifies that cdparanoia is correctly modelling a specific drive's cache and read behavior. Implies -vQL.

−v --verbose
Be absurdly verbose about the auto-sensing and reading process. Good for setup and debugging.

−q --quiet
Do not print any progress or error information during the reading process.

−e --stderr-progress
Force output of progress information to stderr (for wrapper scripts).

−l --log-summary [file]
Save result summary to file, default filename cdparanoia.log.

−L --log-debug [file]
Save detailed device auto-sense and debugging output to a file, default filename cdparanoia.log.

−V --version
Print the program version and quit.

−Q --query
Perform CDROM drive auto-sense, query and print the CDROM table of contents, then quit.

−s --search-for-drive
Forces a complete search for a CDROM drive, even if the /dev/cdrom link exists.

−h --help
Print a brief synopsis of **cdparanoia** usage and options.

−p --output-raw
Output headerless data as raw 16-bit PCM data with interleaved samples in host byte order. To force little or big endian byte order, use **−r** or **−R** as described below.

−r --output-raw-little-endian
Output headerless data as raw 16-bit PCM data with interleaved samples in LSB first byte order.

−R --output-raw-big-endian
Output headerless data as raw 16-bit PCM data with interleaved samples in MSB first byte order.

−w --output-wav
Output data in Micro$oft RIFF WAV format (note that WAV data is always LSB-first byte order).

−f --output-aiff
Output data in Apple AIFF format (note that AIFC data is always in MSB-first byte order).

−a --output-aifc
Output data in uncompressed Apple AIFF-C format (note that AIFF-C data is always in MSB-first byte order).

−B --batch

Cdda2wav-style batch output flag; **cdparanoia** will split the output into multiple files at track boundaries. Output file names are prepended with 'track#.'

−c --force-cdrom-little-endian
Some CDROM drives misreport their endianness (or do not report it at all); it's possible that **cdparanoia** will guess wrong. Use **−c** to force **cdparanoia** to treat the drive as a little endian device.

−C --force-cdrom-big-endian
As above but force cdparanoia to treat the drive as a big endian device.

−n --force-default-sectors *n*
Force the interface backend to do atomic reads of **n** sectors per read. This number can be misleading; the kernel will often split read requests into multiple atomic reads (the automated Paranoia code is aware of this) or allow reads only within a restricted size range. **This option should generally not be used.**

−d --force-cdrom-device *device*
Force the interface backend to read from **device** rather than the first readable CDROM drive it finds. This can be used to specify devices of any valid interface type (ATAPI, SCSI, or proprietary).

−k --force-cooked-device *device*
This option forces use of the old 'cooked ioctl' kernel interface with the specified CDROM device. The cooked ioctl interface is obsolete in Linux 2.6 if it is present at all. **−k** cannot be used with **−d** or **−g**.

−g --force-generic-device *device*
This option forces use of the old 'generic SCSI' (sg) kernel interface with the specified generic SCSI device. **−g** cannot be used with **−k**. **−g** may be used with **−d** to explicitly set both the SCSI carom and generic (sg) devices separately. This option is only useful on obsolete SCSI setups and when using the generic SCSI (sg) driver.

−S --force-read-speed *number*

Use this option explicitly to set the read rate of the CD drive (where supported). This can reduce underruns on machines that have slow disks, or which are low on memory.

−t --toc-offset *number*

Use this option to force the entire disc LBA addressing to shift by the given amount; the value is added to the beginning offsets in the TOC. This can be used to shift track boundaries for the whole disc manually on sector granularity. The next option does something similar...

−T --toc-bias

Some drives (usually random Toshibas) report the actual track beginning offset values in the TOC, but then treat the beginning of track 1 index 1 as sector 0 for all read operations. This results in every track seeming to start too late (losing a bit of the beginning and catching a bit of the next track). **−T** accounts for this behavior. Note that this option will cause**cdparanoia** to attempt to read sectors before or past the known user data area of the disc, resulting in read errors at disc edges on most drives and possibly even hard lockups on some buggy hardware.

−O --sample-offset *number*

Use this option to force the entire disc to shift sample position output by the given amount; this can be used to shift track boundaries for the whole disc manually on sample granularity. Note that this will cause **cdparanoia** to attempt to read partial sectors before or past the known user data area of the disc, probably causing read errors on most drives and possibly even hard lockups on some buggy hardware.

−Z --disable-paranoia

Disable **all** data verification and correction features. When using -Z, **cdparanoia** reads data exactly as would **cdda2wav**(1) with an overlap setting of zero. This option implies that **−Y** is active.

−z --never-skip[=max_retries]

Do not accept any skips; retry forever if needed. An optional maximum number of retries can be specified; for comparison, default without **-z** is currently 20.

−Y --disable-extra-paranoia

Disables intra-read data verification; only overlap checking at read boundaries is performed. It can wedge if errors occur in the attempted overlap area. Not recommended.

−X --abort-on-skip

If the read skips due to imperfect data, a scratch, or whatever, abort reading this track. If output is to a file, delete the partially completed file.

OUTPUT SMILIES

:-) Normal operation, low/no jitter

:-| Normal operation, considerable jitter

:-/ Read drift

:-P Unreported loss of streaming in atomic read operation

8-| Finding read problems at same point during reread; hard to correct

:-0 SCSI/ATAPI transport error

:-(Scratch detected

;-(Gave up trying to perform a correction

8-X Aborted read due to known, uncorrectable error

:^D Finished extracting

PROGRESS BAR SYMBOLS

<space>

No corrections needed

- Jitter correction required

+ Unreported loss of streaming/other error in read

! Errors found after stage 1 correction; the drive is making the same error through multiple rereads, and **cdparanoia** is having trouble detecting them.

e SCSI/ATAPI transport error (corrected)

V Uncorrected error/skip

SPAN ARGUMENT

The span argument specifies which track, tracks, or subsections of tracks to read. This argument is required, unless batch-mode is used (in batch-mode, cdparanoia will rip all tracks if no span is given). **NOTE:** Unless the span is a simple number, it's generally a good idea to quote the span argument to protect it from the shell.

The span argument may be a simple track number or an offset/span specification. The syntax of an offset/span takes the rough form:

1[ww:xx:yy.zz]-2[aa:bb:cc.dd]

Here, 1 and 2 are track numbers; the numbers in brackets provide a finer-grained offset within a particular track. [aa:bb:cc.dd] is in hours/minutes/seconds/sectors format. Zero fields need not be specified: [::20], [:20], [20], [20.], etc, would be interpreted as twenty seconds, [10:] would be ten minutes, [.30] would be thirty sectors (75 sectors per second).

When only a single offset is supplied, it is interpreted as a starting offset and ripping will continue to the end of the track. If a single offset is preceded or followed by a hyphen, the implicit missing offset is taken to be the start or end of the disc, respectively. Thus:

1:[20.35]

Specifies ripping from track 1, second 20, sector 35 to the end of track 1.

1:[20.35]-

Specifies ripping from 1[20.35] to the end of the disc

−2 Specifies ripping from the beginning of the disc up to (and including) track 2

−2:[30.35]

Specifies ripping from the beginning of the disc up to 2:[30.35]

2-4 Specifies ripping from the beginning of track 2 to the end of track 4.

Again, don't forget to protect square brackets from the shell.

EXAMPLES

A few examples, protected from the shell:

Query only with exhaustive search for a drive and full reporting of auto-sense:

cdparanoia -vsQ

Extract an entire disc, putting each track in a separate file:

cdparanoia -B

Extract from track 1, time 0:30.12 to 1:10.00:

cdparanoia "1[:30.12]-1[1:10]"

Extract from the beginning of the disc up through track 3:

cdparanoia -- -3

The "--" above is to distinguish "-3" from an option flag.

OUTPUT

The output file argument is optional; if it is not specified, **cdparanoia** will output samples to one of **cdda.wav**, **cdda.aifc**, or **cdda.raw** depending on whether **−w**, **−a**, **−r** or, **−R** is used (**−w** is the implicit default). The output file argument of − specifies standard output; all data formats may be piped.

ACKNOWLEDGEMENTS

cdparanoia sprang from and once drew heavily from the interface of Heiko Eissfeldt's (heiko@colossus.escape.de) 'cdda2wav' package. **cdparanoia** would not have happened without it.

Joerg Schilling has also contributed SCSI expertise through his generic SCSI transport library.

AUTHOR

Monty <monty@xiph.org>

cdparanoia's homepage may be found at: http://www.xiph.org/paranoia/