

Sandokan

☒ internal / external quicksort

☒ open-source

☒ 32bit / 64bit

☒ written in C

☒ 3 regimes: fast / mixed / slow

☒ free download at www.sanmayce.com

E:\Sandokan_r3-++\Sandokan_vs_GNU_vs_RP>"sandokan.quickSortExternal_4+GB_r3-++.written by kaze, using Bill Durango's Quicksort source.
Usage: sandokan.quickSortExternal_4+GB_wordlistfile [/fast]/[slow] [/ascend]/[descend]
Purpose: Sandokan sorts all [CRLF] ending lines in a given file. Lines can be 4+ billion i.e. more than 1<<32.
Note1: Sandokan is [still] in look-around stage, for superfast sorter look for FFT,
or Internal-Quicksort applied on biggest physical RAM block followed by chunks merging (all-sequential reads/writes).
Note2: Currently (r.3-++) the idea is to see how quicksort behaves on all-physical, physical-external and all-external RAM.
Note3: The Times in wordlistfile must be up to 2048 chars/bytes and end with CRLF or LF i.e. maximum 2050bytes long.
Note4: The goal is to sort a file unfittable in physical memory, and with 32bit code.
Note5: When /fast is specified wordlistfile is uploaded into internal memory if it is possible.
Note6: When /fast is specified and attempt to upload the entire file is failed then new attempt is made to use 768MB for dynamic chunks uploading to the physical RAM block used for internal quicksort, thus /fast is either all-internal or internal-external.
Note7: When /slow is specified wordlistfile is NOT uploaded into internal memory even if it is possible, thus /slow is all-external.
Note8: Not buffered I/O is just awful, like lbyre Reads/Writes: fread(&fourgram[1], 1, 1, fp.in); fwrite(&fourgram[1], 1, 1, fp.outrg);.
Note9: The most cure thing about pure (r.3-++) quicksort approach is the I/O read-only mode through all the sorting,
of course if all the pointers are in physical (not a part in virtual) RAM (another ugly facet: 8bytes-per-line requirement), the ugly thing is the huge number of random accesses i.e. seeks, to bridge this gap in next revision a new (fourth) mode will be available: quicksorting all chunks (fittable in a given physical block) and then simply merging them - an all-sequential approach except the few seeks needed for reaching the chunk offsets.
Note10: I cannot avoid asking myself:
how come that sorting names of all people walked/walking on Earth (some 20 billion maybe) appeared such a difficult task, if Sandokan (r.3-++) is to sort (in external-internal regime) he (not it) would need 151GB (20,000,000,000x8+~1GB) physical (preferably) RAM as a minimum. An assumption: if names are 50bytes long in average then for all-internal mode the needed memory jumps to 150GB+932GB, ouch.
Note11: Any improvements are welcomed, my email: sanmayce@sanmayce.com, the idea is a free open-source sorter to become available.

E:\Sandokan_r3-++\Sandokan_vs_GNU_vs_RP>