

	A	B	C	D	E	F	G	H
1	short description	arbitrary CUDAPM1 item id #	seen in version #	Long description ( <b>changes since last version posted are in bold font</b> )	Status	see also		
2	"Initializing"	1	0.20	Cosmetic change: Initializing	<b>fixed in V0.22</b>	na		
3	UnusedMem ini file entry missing message	2	0.20	Absent ini file UnusedMem directive generates a parse failure message. Literal 100 instead of UNMEM_DFLT used in some code	<b>fixed in V0.22</b>	na		
4	ini file edits	3	0.20	Added UnusedMem section to cudaPm1.ini, changed CUDALucas references to CUDAPm1 throughout, changed default device number from one to zero to fit common single-gpu systems, and zero-based numbering comment added	<b>fixed in V0.22</b>	na		
5	every launch after savefile directory is created, there's a message it could not be created	4	0.20	Reversed case of when savefile directory creation message is printed, modified message accordingly	<b>fixed in V0.22</b>	na		
6	Literal 1,000,000,000 used, MAX_B2?	5	0.20	Modified MAX B2 message, added define of MAX_B2	<b>fixed in V0.22</b>	na		
7	no entry message re threads file	6	0.20	Modified the no entry message for threads file	Source code modified, awaiting compile/ verify/ upload	na		
8	misc code edits	7	0.20	Modified some comments and formatting; version increment	<b>fixed in V0.22</b>	na		
9	error messages splitting between command console and output redirection target	8	0.20	Fprintf(stderr for first part of several messages, printf( for later part was resulting in splitting messages to two destinations if output redirection of stdout occurred. Example: Warning: ignoring line 1: "PFactor=6A7EB8FCE330BCE7F94FF0C29207076A,1,2,185000029,-,77,2" in "worktodo.txt". Reason: appears in cmd.exe window because it's to stderr; the reason "invalid format" appears in stdout which is redirected to a file (1 missing after the -)	<b>fixed in V0.22</b>	51		
10	-r option in CUDAPm1 listed in help message but not functional. This self test mode is not available	9	0.20, <b>0.22</b>	The -r presence in the CUDAPm1 help message output seems to be a holdover from its CUDALucas ancestry. Program's help output indicates "exec residue test." Specifying -r on the command line does not result in any residue check tests running in CUDAPm1; it goes straight to continuation of work present in the worktodo file. If I read the source code correctly, the residue check function did not get implemented for CUDAPm1. Implementation would ideally be driven by an external plain text table, so that adding new cases would not always require code modification and recompilation. A workaround for now is to run one or more P-1 test on one or more exponents that will use the same fft length(s) as untested exponents planned to be run, if suitable exponent and factor combinations can be found. Probably should comment out -r option output in help message output until implemented. Preferred implementation would be -r low high; read the external table, and run residue tests for the subset of entries low<=entry<=high. Specifying -r n would run all entries <=n; specifying -r with no value following would run all residue tests.	identified; no code or volunteer at this time	CUDALucas 28		
11	Some CUDALucas bugs possibly shared	10	expected in all	Since CUDAPm1 was derived from CUDALucas code, CUDAPm1 may share some of CUDALucas' issues. The device renumbering occurrence is generic and more deeply rooted and highly likely	suspected	various		
12	documentation	11	0.20, <b>0.22</b>	Needs a readme, wiki page or something; draft revised ini and readme exist but not posted	pending	na		
13	Are the various executables for linux and windows current? There were various release numbers for V0.20 including r50 and r52	12	0.20	File dates seem to indicate that some executables are not current with the last previous source changes. See <a href="http://www.mersenneforum.org/showpost.php?p=462600&amp;postcount=503">http://www.mersenneforum.org/showpost.php?p=462600&amp;postcount=503</a>	<b>fixed in V0.22</b>	na		

	A	B	C	D	E	F	G	H
14	No date/time stamps; gcd time	13	0.20, 0.22	There are various elapsed time estimates, and total time estimates, but no actual date/time stamps in output. It's unclear if gcd time is included in stage estimated total time, or how long the gcds take	identified; no code or volunteer at this time	74; CUDALucas 33		
15	logging of stdout, stderr to file	14	0.20, 0.22	Planned feature addition, an -o flag to control output to screen and log. Stdout can be duplicated to log, switched to log, suppressed (default is stdout to console, status quo) Stderr can be duplicated to log but not suppressed or switched. Workaround is redirection, or tee (native to linux, part of Windows Powershell)	in development, beginnings of partial code draft exists. Needs ini file and command line extension, global trivial rewrite of printf and fprintf calls, then compile/debug/test	CUDALucas 19		
16	Repeating-zero residue issue	15	0.20	<p>Unknown program error or hardware error initiates repeating 0x00 residues. (Perhaps a failed memory copy.) Program carries this to completion. Of four occurrences known, 3 occurred in stage 1 and one is uncertain. The end result is a normal looking entry in the results file no factor found. Observed to start midstream in a run, following a slightly larger exponent that completed successfully. CUDA reports 1429M of 1536M GPU memory free.</p> <p>Using threads: norm1 128, mult 64, norm2 32.</p> <p>Using up to 1232M GPU memory.</p> <p>Selected B1=2145000, B2=13942500, 4.6% chance of finding a factor</p> <p>Using B1 = 2145000 from savefile.</p> <p>Continuing stage 1 from a partial result of M249500221 fft length = 14336K, iteration = 2050001</p> <p>Iteration 2100000 M249500221, 0xb81566f6b7207983, n = 14336K, CUDAPm1 v0.20 err = 0.14063 (22:30 real, 26.9995 ms/iter, ETA 7:27:50)</p> <p>Iteration 2150000 M249500221, 0x0000000000000000, n = 14336K, CUDAPm1 v0.20 err = 0.13281 (22:22 real, 26.8399 ms/iter, ETA 7:02:49)</p> <p>Iteration 2200000 M249500221, 0x0000000000000000, n = 14336K, CUDAPm1 v0.20 err = 0.11289 (22:21 real, 26.8282 ms/iter, ETA 6:40:17)</p> <p>Iteration 2250000 M249500221, 0x0000000000000000, n = 14336K, CUDAPm1 v0.20 err = 0.09596 (22:23 real, 26.8743 ms/iter, ETA 6:18:34)</p> <p>SIGINT caught, writing checkpoint.</p> <p>... continued into stage 2...</p> <p>Processing 180 - 180 of 480 relative primes.</p> <p>Initializing pass... done. transforms: 251, err = 0.00000, (3.60 real, 14.3510 ms/tran, ETA 16:21:38)</p> <p>SIGINT caught, writing checkpoint.</p> <p>Transforms: 8760 M249500221, 0x0000000000000000, n = 14336K, CUDAPm1 v0.20 err = 0.00000 (2:03 real, 14.0203 ms/tran, ETA 16:19:13)</p> <p>Proposed treatment is to detect at or just prior to checkpoint or save file write, revert to most recent believed-good save file or checkpoint, retry a limited number of times.</p>	bad residues check implemented in v0.22; test it	15; 70; CUDALucas 4		
17	invalid threadbench timings	16	0.20	Certain threads counts on certain gpu models, or CUDA level & fft length & gpu combinations result in wrong execution and anomalously fast benchmarks that then cause the malfunctioning values to be selected for future use. the only lower bound currently on iteration timings in case of a malfunction is 0.0f. A lower bound computed as a fraction of the average of all timings above a lower fft length cutoff can be used to validate execution of individual thread combinations of a given fft length above the cutoff length.	test for anomalously fast threadbench timings implemented in v0.22; test it	part of CUDALucas 2		
18	invalid fftbench timings	17	0.20	Certain CUDA level & fft length & gpu combinations result in wrong execution and anomalously fast benchmarks that then cause the malfunctioning values to be selected for future use. the only lower bound currently on iteration timings in case of a malfunction is 0.0f. A lower bound computed as a fraction of the average of (each timing above a lower fft length cutoff divided by its fft length) can be used to validate execution of a given fft length above the lower fft length cutoff. Invalidated lengths can be suppressed from the output fft file and prevented from removing other shorter fft lengths from consideration.	test for anomalously fast fftbench timings implemented in v0.22; test it	part of CUDALucas 2		

	A	B	C	D	E	F	G	H
19	autoselect picks b1=0 b2=0 which produces a run with repeating residues 0x01	18	0.20	<p>Some exponents lead to autoselection of B1=0, B2=0. Two cases have been identified; smallest mersenne prime exponents, and composite exponents. The program computes stage one with a repeating residue value of 0x01. It does not warn or error. The end result is a normal looking entry in the results file no factor found. A preferable implementation if some exponents are being excluded intentionally from the full computation is to produce a message about their exclusion and why, and skip the computation. CUDA reports 1434M of 1536M GPU memory free.</p> <p>Index 2 No entry for fft = 8k found. Using default thread sizes. For optimal thread selection, please run ./CUDApm1 -cufftbench 8 8 r for some small r, 0 &lt; r &lt; 6 e.g. Using threads: norm1 512, mult 128, norm2 128. Using up to 1334M GPU memory. Selected B1=0, B2=0, 0% chance of finding a factor Starting stage 1 P-1, M86243, B1 = 0, B2 = 0, fft length = 8K Doing 18 iterations Running careful round off test for 1000 iterations. If average error &gt; 0.25, the test will restart with a longer FFT. Iteration 1000, average error = 0.00000 &lt;= 0.25 (max error = 0.00000), continuing test. M86243, 0x0000000000000001, n = 8K, CUDApm1 v0.20 Stage 1 complete, estimated total time = 0:00 Starting stage 1 gcd. M86243 Stage 1 found no factor (P-1, B1=0, B2=0, e=0, n=8K CUDApm1 v0.20)</p>	bad residues check implemented in v0.22; test it	15; 70		
20	fft length table is from gtx570	19	0.20	<p>Other tables may be suitable for other gpu models. Feature request, ability to load the table from a file. File could be switched out for differing gpu models to optimize for each. int init_ffts() { ... #define COUNT 160 int default_mult[COUNT] = { //this batch from GTX570 timings 2, 8, 10, 14, 16, 18, 20, 32, 36, 42, 48, 50, 56, 60, 64, 70, 80, 84, 96, 112, 120, 126, 128, 144, 160, 162, 168, 180, 192, 224, 256, 288, 320, 324, 336, 360, 384, 392, 400, 448, 512, 576, 640, 648, 672, 720, 768, 784, 800, 864, 896, 900, 1024, 1152, 1176, 1280, 1296, 1344, 1440, 1568, 1600, 1728, 1792, 2048, 2160, 2304, 2352, 2592, 2688, 2880, 3024, 3136, 3200, 3584, 3600, 4096, 4320, 4608, 4704, 5120, 5184, 5600, 5760, 6048, 6144, 6272, 6400, 6480, 7168, 7200, 7776, 8064, 8192, 8640, 9216, 9408, 10240, 10368, 10584, 10800, 11200, 11520, 12096, 12288, 12544, 12960, 13824, 14336, 14400, 16384, 17496, 18144, 19208, 19600, 20000, 20250, 21952, 23328, 23814, 24300, 24500, 25088, 25600, 26244, 27000, 27216, 28000, 28672, 31104, 31250, 32000, 32400, 32768, 33614, 34992, 36000, 36288, 38416, 39200, 39366, 40500, 41472, 42336, 43200, 43904, 47628, 49000, 50000, 50176, 51200, 52488, 54432, 55296, 56000, 57344, 60750, 62500, 64000, 64800, 65536 };</p>		na		
21	B1 extension	20	0.20	<p>There's a branch to "stage 3" which is a placeholder for where B1 extension code would go. Raising B1 would involve checking for and potentially raising powers of primes below old B1, and multiplying by suitable powers of primes between old B1 and extended B1(mod the Mp).</p>		na		
22	*UID:username/systemid-gpuid, "	21	0.20	<p>Support prefacing results lines with UID: &lt;username&gt;/(&lt;compname&gt;), e.g.: UID: JamesHeinrich/GTX1080, &lt;result line here as normal&gt; see <a href="http://www.mersenneforum.org/showpost.php?p=461411&amp;postcount=3">http://www.mersenneforum.org/showpost.php?p=461411&amp;postcount=3</a></p>	conceptual	74; CUDA Lucas 32		

	A	B	C	D	E	F	G	H
23	savefile naming with previous save's residue value in name	22	0.20	<p>First savefile of a run is S81328579.50001.txt, writing checkpoints at 50000 intervals. Second is s81328579.100001.9519ada275083227.txt, third is s81328579.150001.3b78a68dbb318e75.txt, fourth s81328579.200001.eca5c42ceeadae9e.txt etc console output corresponding to these is                      Iteration 50000 M81328579, 0x9519ada275083227, n = 4608K, CUDAPm1 v0.20 err = 0.09863 (8:25 real, 10.1038 ms/iter, ETA 2:48:57)                      Iteration 100000 M81328579, 0x3b78a68dbb318e75, n = 4608K, CUDAPm1 v0.20 err = 0.09814 (8:19 real, 9.9766 ms/iter, ETA 2:38:31)                      Iteration 150000 M81328579, 0xec5c42ceeadae9e, n = 4608K, CUDAPm1 v0.20 err = 0.10352 (8:19 real, 9.9774 ms/iter, ETA 2:30:13)                      Iteration 200000 M81328579, 0x04bb5c6a41d55bcd, n = 4608K, CUDAPm1 v0.20 err = 0.10156 (8:15 real, 9.8952 ms/iter, ETA 2:20:44)                      A comparison shows the residue is missing on the 50000 and its residue is used in the name of the 100000 instead, and following ones follow this off-by-one pattern. this could conceivably be intended behavior, tagging the successive savefile with the residue from which it was begun, however, the first produced of each run is missing the residue tag. It would seem more clear to tag with the residue of the file contents being saved, simplifying the logic of what savefile to attempt a resume from. See <a href="http://www.mersenneforum.org/showpost.php?p=358176&amp;postcount=392">http://www.mersenneforum.org/showpost.php?p=358176&amp;postcount=392</a></p>		na		
24	echo the worktodo line or command line inputs to console/log	23	0.20	Echo the worktodo line content or content of command line parameters concerning the exponent to console and if applicable to log, so that after the console session scrolls or the worktodo line is removed, there is still a record of the parameters used for the exponent, such as B1, B2 values, # of LL tests that would be saved by finding a factor, etc.	conceptual	CUDALucas 34		
25	retry rather than halt on round off >0.4	24	0.20	Current behavior is if round off error observed is >0.4, the program prints the error and halts: err = 0.5 >= 0.40, quitting. Frequently a restart produces a lower, acceptable error level and computation can continue. Unattended sessions that halt cause loss of throughput. Batch wrappers can reduce that but it makes screen and log clutter and overhead for restart. A resetting retry counter with a small limit, say 3, perhaps with a brief time delay of several seconds between retries is proposed.	conceptual	CUDALucas 35		
26	simplify resume from save file	25	0.20	A command line option to resume from a specific interim save file: check for seemingly valid residue in filename, rename suspected bad checkpoint files, copy save file to the appropriate filename in working directory, resume	conceptual	CUDALucas 36		
27	device confirmation	26	0.2	<p>In multiple-GPU systems, NVIDIA driver timeout or thermal limits or a combination may cause a device to disappear from the device count, even though Windows Device Manager shows it, and GPU-Z lists it and can find its parameters but not display its sensor readings. Batch wrapper may restart a run on a different device than intended as a result. This may affect execution timing of two sessions sharing one gpu, or may cause a restarted session to fail if it requires more resources than available on a dissimilar card or a device number higher than the reduced Windows count allows. That generates error message                      device_number &gt;= device_count ... exiting                      (This is probably a driver problem)                      Confirming unique device characteristics could allow greater confidence in execution. <b>See also</b> <a href="https://www.mersenneforum.org/showpost.php?p=497512&amp;postcount=2698">https://www.mersenneforum.org/showpost.php?p=497512&amp;postcount=2698</a></p>	conceptual/ flashj has done some development in CUDALucas	CUDALucas 8		
28	feature suggestion: run time estimation for entire worktodo file	27	0.20	It would be useful to have an option in CUDALucas to read the entire worktodo file, and compute and display time span and completion date for each line of the worktodo file, at the outset of a run following the usual header and optional info section. Option could be -w or -work.	conceptual	CUDALucas 9		
29	FFT benchmark and threads benchmark are believed affected by variable clock rate of modern GPUs	28	0.20	Monitoring gpu core clock with GPU-Z 1.18 shows clock rate changes in the early phase of a run, approx 1354 to 1860 Mhz. Nearly all the run occurs at maximum clockrate. A single intermediate clock value 1468 Mhz is logged in the cudalucas output. It is likely to affect the accuracy of the initial fft lengths' benchmarking. Addition of a "warmup" phase that loads the GPU and brings clock up to speed before benchmarking begins is proposed. As a workaround, fft and threads benchmarking could be submitted in a batch file immediately after test of a small exponent, or benchmarking at smaller fft lengths than needed can be performed and the affected lengths not needed for current GIMPS wavefront edited out of the benchmark files.	conceptual	CUDALucas 16		

	A	B	C	D	E	F	G	H
30	Addition of header in fft and threads files, of fairly complete description of conditions in which file was generated	29	0.20	Fft file has no header information subset. Including the following when it is generated would document how it was created and what it's valid for, and conversely, what it may not be valid for. System identifier, Operating system, CUDALucas version, CUDA level, 32 or 64 bit flavor, NVIDIA driver version, GPU model. Currently only the GPU model is documented, and only in the file name for the threads file which contains no header. Program could be enhanced to check that the fft or threads file being used was appropriate for the version running, and warn if not.	conceptual	CUDALucas 20		
31	Automatic setup and verification	30		Add a command line switch to automate initial setup and burn-in (-setup) <b>Possible workaround is a shell script or Windows batch file.</b>	conceptual	CUDALucas 22		
32	PrimeNet interface	31		Automatic PrimeNet interaction. Possible alternative is a separate helper program that reads log file, result file, worktodo file, locks/modifies/unlocks worktodo file	conceptual	CUDALucas 23		
33	FFT file run time cap estimates	32		Add a runtime estimate column for maximum exponent per fft length to <gpu> fft.txt. Format in clock and calendar units; hours, days, weeks, months or years. Adjust file reading to accommodate if necessary.	conceptual	CUDALucas 24		
34	GPUs with vram <=1.5GB (2GB?) have issues with fft benchmark up to <b>65536K; up to 3GB+, to 131072k; under 8GB (6?), to legal program input 262144k</b>	33	0.20; <b>0.22</b>	Fft benchmark on <b>1.5GB</b> Gtx480 runs until it hits a limit and fails to output an fft file, after 32760k, possibly varying somewhat with CUDA level as it did with CUDALucas. Program does not warn of an issue. Successive runs with higher start points were not usable to achieve a complete fft benchmarking; be sure to save first fft file as a separate name to avoid data loss, then assemble the piecemeal results. Modify program to compute whether memory will be adequate for the requested run, suggest or make modifications to calling parameters to allow run to completion. Workaround is to run twice, scaling back upper limit until an FFT file is produced. 1GB Quadro 2000s can go to 32768k.	conceptual	CUDALucas 26		
35	GPUs with vram <=1.5GB (2GB?) have issues with threads benchmark up to legal program inputs	34	<b>0.20;</b> <b>0.22</b>	In CUDALucas, Fftlength upper limits relating to limited memory size for threadbench are lower than for fftbench for the same GPU/memory size. Program does not warn of an issue. Threads files do not get overwritten, but added to. Test for similar behavior at probably lower threshold in CUDAPm1, yields a top feasible fft length in threads benchmarking of 55296K for 1.5GB VRAM.	confirmed	CUDALucas 27		
36	Hardening against user error	35	0.20	Detect and warn on illegal user input and substitute sensible default or exit; encode parameters for which fft file and threads file was run into the file itself, then check at use for a match with the parameters of use (CUDA level, software version, etc.)	conceptual	CUDALucas 30		
37	Formatting output	36		Add vertical white space to emphasize stage 1 err >0.4 quitting message; check for other.	preliminary conceptual	na		<b>revisit</b>
38	Halt without warning or error message	37	0.20	May be a program crash, OS or power issue. Not due to a sigint or completion of P-1 stage since it occurs without a sigint message during stage1 or stage2. cause(s) unknown (does it include the NVIDIA driver timeout issue?) Generally restarting these allowed completion of the worktodo item.		63, 64, CUDALucas 37		
39	Exhaustive error checking version	38		For debug purposes, high reliability runs, or making use of lower reliability software, trade performance for error detection and some handling, by checking return values of all function calls (cpu or gpu). optionally dump inputs and outputs of the function to file in case of a nonrecoverable error for later analysis. In some cases gather additional information on machine state specific to the error occurrence	conceptual	CUDALucas 38		
40	Badmem handler	39		Persistent gpu memory error map and lockout. The idea is to map and allocate the bad memory areas, so they can't be allocated to math we want right, which then gets other more reliable portions of memory allocated. Whether this can be made to work on commodity gpus requires some study. Requires reliable device confirmation/device number stability. This technique is used in linux hobbyist systems, and in NVIDIA K20X based supercomputers (see <a href="http://www.mersenneforum.org/showthread.php?t=22471&amp;page=6">http://www.mersenneforum.org/showthread.php?t=22471&amp;page=6</a> <a href="http://rick.vanrein.org/linux/badram/download.html">http://rick.vanrein.org/linux/badram/download.html</a> and <a href="http://on-demand.gputechconf.com/gtc/2015/presentation/S5566-James-Rogers.pdf">http://on-demand.gputechconf.com/gtc/2015/presentation/S5566-James-Rogers.pdf</a> ) If implemented as a separate program, provides a general purpose solution applicable to other gpu computing also (primality testers, factoring, ...)	conceptual	CUDALucas 39		
41	Augment -info output	40		Include OS, OS-version, OS-bitness, app-bitness, app CUDA level compiled for, CUDA level available from driver, device number, identifying properties such as BIOS version string, PCIbus# and PCIdevice#	some development done by flashjh for CUDALucas 2.06beta	26; see also CUDALucas 8		

	A	B	C	D	E	F	G	H
42	error and retry tabulation extension	41		Count recovered errors and retries by type per exponent and summarize in console output, log and result line (see <a href="http://www.mersenneforum.org/showpost.php?p=465207&amp;postcount=58">http://www.mersenneforum.org/showpost.php?p=465207&amp;postcount=58</a> )	conceptual	CUDALucas 41		
43	compute and output to console gpu ram amount required for a stage	42		For a stage 1 or 2, before beginning or resuming, calculate gpu-ram requirements for that exponent, compare to amount available, output figures to console, do not attempt unless sufficient, warn of inadequate memory and <b>switch to next worktodo entry</b> V0.20 warns but does not give indication of max exponent or how much short it might be: Using up to 1584MB of memory. WARNING: There may not be enough GPU memory for stage 2!	conceptual	CUDALucas 42		
44	install script	43		From a source directory created by unzipping the package, running the install script would accomplish some or all the following; log its actions along the way, save info for later uninstall, get info for setup from user, warn about run time of this phase, create shared executable and dll directory, create separate per device working directories, copy executables and dlls into place, copy template ini file into place in each working directory, append customized directives to ini file in each working directory, copy utility batch files into separate working directories, prompt user to customize utility batch files, warn user about run time & give option to return to it later, launch separate setup / benchmark / proveout scripts, clean up after, instruct user to get real work; recommend GPU-Z or similar hardware monitor.	conceptual	CUDALUCAS 22 & 43		
45	uninstall script	44		Warn user to back up first, then wait for ok to proceed, using saved uninstall info, remove exe & dll directory contents & directory, remove each separate worker directory's ini file, and if then empty, the directory	conceptual	CUDALUCAS 44		
46	deactivate current problem worktodo entry as a remark, continue with next	45		If a current work item generates a problem that can not be resolved by retries, deactivate the entry by prefacing it with rem or # and continue on with the next work entry. Implementing this would also add the ability to place remarks in worktodo file without generating error messages for each such. Note, # or \ are comment characters marking the rest of a line as a comment. See <a href="http://www.mersenneforum.org/showpost.php?p=492998&amp;postcount=574">http://www.mersenneforum.org/showpost.php?p=492998&amp;postcount=574</a>	conceptual	CUDALUCAS 45		

	A	B	C	D	E	F	G	H
47	not enough memory for stage 2, executes with unchanging residue	46	0.20	<p>Observed on GTX480 with large exponent, stage 1 executes ok, stage 2 runs with display of repeating residue from stage 1 completion while using approx 1248MB of the 1536MB gpu vram; this can occupy days for large exponents.</p> <p>Using up to 1584M GPU memory.                      WARNING: There may not be enough GPU memory for stage 2!                      Selected B1=2660000, B2=17955000, 5.04% chance of finding a factor                      Starting stage 1 P-1, M299500177, B1 = 2660000, B2 = 17955000, fft length = 18432K Iteration 3800000 M299500177, 0xe97a5cb286fcf801, n = 18432K, CUDAPm1 v0.20 err = 0.03418 (30:13 real, 36.2698 ms/iter, ETA 22:56)                      M299500177, 0x071ac99b54319724, n = 18432K, CUDAPm1 v0.20                      Stage 1 complete, estimated total time = 38:41:27                      Starting stage 1 gcd.                      M299500177 Stage 1 found no factor (P-1, B1=2660000, B2=17955000, e=0, n=18432K CUDAPm1 v0.20)                      Starting stage 2.                      Using b1 = 2660000, b2 = 17955000, d = 2310, e = 2, nrp = 1                      Zeros: 778308, Ones: 811452, Pairs: 143662                      Processing 1 - 1 of 480 relative primes.                      Initializing pass... done. transforms: 170, err = 0.03711, (3.02 real, 17.7483 ms/tran, ETA NA)                      Transforms: 16700 M299500177, 0x071ac99b54319724, n = 18432K, CUDAPm1 v0.20 err = 0.03711 (5:16 real, 18.8891 ms/tran, ETA 42:23:55)</p> <p>Processing 2 - 2 of 480 relative primes.                      Initializing pass... done. transforms: 201, err = 0.03154, (3.99 real, 19.8628 ms/tran, ETA 42:26:28)                      Transforms: 16680 M299500177, 0x071ac99b54319724, n = 18432K, CUDAPm1 v0.20 err = 0.03154 (5:15 real, 18.8892 ms/tran, ETA 42:21:10)                      ...</p> <p>Processing 338 - 338 of 480 relative primes.                      Initializing pass... done. transforms: 257, err = 0.02930, (5.06 real, 19.7013 ms/tran, ETA 12:43:35)                      Transforms: 16660 M299500177, 0x071ac99b54319724, n = 18432K, CUDAPm1 v0.20 err = 0.03027 (5:15 real, 18.8894 ms/tran, ETA 12:38:19)</p>	conceptual	45		
48	some messages say CUDALucas	47	0.20, 0.22	See parse.c for comments and error messages referencing CUDALucas instead of CUDAPm1	conceptual	na		
49	OS-specific messages on other OS	48	0.20, 0.22	<p>Observed in program message output, linux syntax on Windows build. For example:                      For optimal thread selection, please run                      ./CUDAPm1 -cufftbench 8 8 r                      (/ instead of \). Code review of such occurrences shows 3 instances in cudaPm1.cu; appears to be cudaPm1-specific (not found in Mfaktc, CUDALucas); make consistent with target OS with conditional compilation. <b>In v0.22 cudaPm1.cu, printf("./CUDAPm1 -cufftbench 1 8192 r\n");</b></p>	conceptual	na		
50	doesn't begin with Test= or DoubleCheck=.	49	0.20, 0.22	<p>When worktodo file contains a record that is not blank and does not start with Pfactor=, it outputs an error message "doesn't begin with Test= or DoubleCheck=." That's appropriate to an LL test program such as CUDALucas, not a P-1 tester. Revise message to fit P-1. Implement remark handling. Worktodo entry beginning with rem generated the following message at the console: "Warning: ignoring line 1: "rem Pfactor=1,2,152000239,-1,74,2" in "worktodo.txt". Reason:" and the following in redirected output: "doesn't begin with Test= or DoubleCheck=." parse.c contains the following line: case NO_TEST_EQUAL: fprintf(stderr, "doesn't begin with Test= or DoubleCheck=.\n");break;</p>	conceptual	8		
51	no CRC on result record	50	0.20	Implement CRC per <a href="http://www.mersenneforum.org/showpost.php?p=466208&amp;postcount=149">http://www.mersenneforum.org/showpost.php?p=466208&amp;postcount=149</a> , <a href="http://www.mersenneforum.org/showpost.php?p=466222&amp;postcount=154">http://www.mersenneforum.org/showpost.php?p=466222&amp;postcount=154</a>	conceptual	CUDALucas 47, Mfaktc 24		

	A	B	C	D	E	F	G	H
52	Crash near end of stage 2; on resumption, message regarding 481 of 480 relative primes; negative ETAs	51	0.20	<p>Processing 385 - 480 of 480 relative primes.                      Initializing pass... done. transforms: 4493, err = 0.12500, (15.12 real, 3.3654 ms/tran, ETA 22:50)                      Transforms: 21202 M85966123, 0x3d4010076d275b36, n = 4800K, CUDAPm1 v0.20 err = 0.12891 (1:13 real, 3.4487 ms/tran, ETA 21:37)                      Transforms: 20320 M85966123, 0xfb70dd92ab47bb1, n = 4800K, CUDAPm1 v0.20 err = 0.13281 (1:10 real, 3.4472 ms/tran, ETA 20:27)                      ...Transforms: 20790 M85966123, 0xfcfe545618d202, n = 4800K, CUDAPm1 v0.20 err = 0.14063 (1:12 real, 3.4466 ms/tran, ETA 2:25)                      Transforms: 21220 M85966123, 0xb39c618229307025, n = 4800K, CUDAPm1 v0.20 err = 0.12793 (1:13 real, 3.4467 ms/tran, ETA 1:12)                      Transforms: 21048 M85966123, 0x915ddb1ed13b43d8, n = 4800K, CUDAPm1 v0.20 err = 0.12500 (1:13 real, 3.4466 ms/tran, ETA 0:00)                      batch wrapper exiting at Sun 02/18/2018 21:18:23.78                      batch wrapper (re)launching at Sun 02/18/2018 22:15:06.40                      ...Continuing stage 2 from a partial result of M85966123 fft length = 4800K Starting stage 2.                      Using b1 = 740000, b2 = 17575000, d = 2310, e = 6, nrp = 96 Zeros: 795043, Ones: 891197, Pairs: 175430 Processing 481 - 480 of 480 relative primes. Initializing pass... done. transforms: 1773, err = 0.12500, (5.47 real, 3.0833 ms/tran, ETA 0:00) Transforms: 21084 M85966123, 0x915ddb1ed13b43d8, n = 4800K, CUDAPm1 v0.20 err = 0.12500 (1:09 real, 3.2664 ms/tran, ETA -1:-12) Transforms: 21060 M85966123, 0x915ddb1ed13b43d8, n = 4800K, CUDAPm1 v0.20 err = 0.10625 (1:12 real, 3.4104 ms/tran, ETA -2:-24)</p>	new report	na		
53	May report from stage 2 a small incorrect factor not of the proper form for a Mersenne number	52	0.20	<p>After a normal looking run on a GTX1070, cudapm1 reported: M86166293 has a factor: 5 (P-1, B1=740000, B2=17575000, e=6, n=4800K, aid=E279CF2D5D7B0F96B07EF0333C8EE932 umm, no. not of form 2kp+1 and 1 or 7 mod 8. The value 5 would only be valid for exponents up to 2 (f=2kp+1, k=1 is the minimum value) and is not 1 or 7 mod 8, so not a valid potential factor for any higher exponent. A rerun from start gave the same final stage 2 residue but indicated no factor found, so a computation error in the stage 2 gcd is assumed to be the cause. After stage 2 on a GTX1070,                      M86286989 has a factor: 3 (P-1, B1=740000, B2=17575000, e=6, n=4800K, CUDAPm1 v0.20)                      Could add a sanity check for claimed factors in cudapm1 and flag such a result as an error, and perhaps rerun a phase or from start, a limited number of tries.</p>	new report	41, 53, 56, 62		
54	May report from stage 1 a small incorrect factor not of the proper form for a Mersenne number	53	0.20	<p>After a normal looking run on a GTX1070, cudapm1 reported small factor at end of stage one and skipped stage 2: Iteration 1000000 M86175911, 0xa24012d2d6304811, n = 4800K, CUDAPm1 v0.20 err = 0.14844 (11:00 real, 6.6068 ms/iter, ETA 7:28)                      M86175911, 0x507ca9ce05457d1f, n = 4800K, CUDAPm1 v0.20                      Stage 1 complete, estimated total time = 1:57:48                      Starting stage 1 gcd.                      M86175911 has a factor: 5 (P-1, B1=740000, B2=740000, e=6, n=4800K CUDAPm1 v0.20)</p> <p>Iteration 1000000 M86175937, 0x8cad68456f686913, n = 4800K, CUDAPm1 v0.20 err = 0.14551 (11:00 real, 6.6066 ms/iter, ETA 7:28)                      M86175937, 0x35418143098fd9c8, n = 4800K, CUDAPm1 v0.20                      Stage 1 complete, estimated total time = 1:57:29                      Starting stage 1 gcd.                      M86175937 has a factor: 7 (P-1, B1=740000, B2=740000, e=6, n=4800K CUDAPm1 v0.20)                      umm, no. not of form 2kp+1 and 1 or 7 mod 8. After stage 1 on a GTX1070,                      86175937 reported 7 as a factor. The value 5 would only be valid for exponents up to 2 (2kp+1, k=1 is the minimum value) and is not 1 or 7 mod 8, so not a valid potential factor for any higher exponent. Reruns for each, which began at stage 2, produced "No factor found".. A computation error in the stage 1 gcd is assumed to be the cause of the erroneous low factor. Could add a sanity check for claimed factors in cudapm1 and flag such a result as an error, and perhaps rerun a phase or from start, a limited number of tries, perhaps a single retry, as a program extension.</p>	new report	41, 52, 56, 62		

	A	B	C	D	E	F	G	H
55	gpu goes idle during gcd	54	0.20, 0.22	GPU-Z monitoring shows the gpu going idle apparently for the duration of gcd computations. Workaround is to run another instance of some gpu app, <b>probably tf for its small memory requirement</b> . (May help or may hurt overall throughput depending on apps and gpu model.) If the gpu could effectively be used for the gcd, adding that would be a nice feature and help total system throughput. Or multithread the program, so the gpu can work on a different exponent while the cpu does the gcd. <b>Or speculatively execute stage 2 transforms on the gpu along with stage1 gcd on the cpu, which will pay off in typically ~98% of attempts, those that find no stage 1 factor. Note, although the CUDAPm1 gcd is a single core, the cpu throughput impact is as many cores as are allocated per mprime/prime95 worker, times the gcd durations.</b>	new report	na		
56	Factor found may be composite (multiple)	55	0.20	Factors found may actually be products of multiple prime factors. It would be useful if it did a quick factoring check of factors found.	known behavior / feature request	na		
57	savefile at stage end out before gcd	56	0.20	Extend to save the state before the gcd at stage 1, and at stage 2, and retain until any factor found passes a set of validity tests.	feature request	41, 52, 53, 62		
58	independent screen output and save file periods	57	0.20	Currently savefiles if enabled are tied to the screen output rate. CUDALucas has separate values for these intervals. CUDAPm1 would benefit by having more frequent screen output for status checking, and less frequent checkpoint file saving for efficiency.	feature request	na		check other bug lists for applicability of this item
59	worktodo.add	58	0.20	Add the worktodo.add functionality (See mfakto code or mfaktc 0.21 code)	feature request	CUDALUCAS 49		update xrefs
60	Some phases are single cpu core limited	59	0.20	During load from a save file for a large exponent, the gpu is idle and one cpu core of a 12-core system is maxed out. Multicore processing would improve throughput by reducing the amount of gpu idle time. This probably also applies to gcd computations. GCD is observed to bounce around among all the system's cpu cores (and so also the various chips' caches). GPU computation of the gcd would increase gpu utilization and free cpu cycles for Prime95/mprime.	feature request	na		check other bug lists for applicability of this item
61	threads file ignored	60	0.10	V20130909 Win64 CUDA5.0 and some other available images seem to ignore the threads file and use the ini file setting of threads. If there is no ini file setting for threads, they default to 256 for norm1, ignoring the fft length's entry in the threads file. This is detected if an error message is produced, such as over specifications Grid = n where n> 65535.	new report	na		
62	zu format	61	0.20	The 20130923 Win64 version has zu appearing where parameters should in the info output. The Win64 20130909 (v0.10) and the Win64 20131118 (v0.20) version do not have this issue. CUDAPm1 v0.20 ----- DEVICE 0 ----- name GeForce GTX 1070 Compatibility 6.1 clockRate (MHz) 1708 memClockRate (MHz) 4004 totalGlobalMem zu totalConstMem zu l2CacheSize 2097152 sharedMemPerBlock zu regsPerBlock 65536 warpSize 32 memPitch zu maxThreadsPerBlock 1024 maxThreadsPerMP 2048 multiProcessorCount 15 maxThreadsDim[3] 1024,1024,64 maxGridSize[3] 2147483647,65535,65535 textureAlignment zu DeviceOverlap 1	resolved in 20131118 version	na		
63	savefile at stage end after gcd	62	0.20	Extend to save the state immediately after the gcd at stage 1, and after the gcd at stage 2, and retain until any factor found passes a set of validity tests or the exponent is completed successfully.	feature request	41, 52, 53, 56		

	A	B	C	D	E	F	G	H
64	uncompletable exponents on a gpu	63	0.20	Some exponents were found to crash before completion silently. Restart on the same gpu produced the same result. Relocation to another gpu, even the same model, sometimes permitted completion. This was observed around 84m on Quadro 2000s. Also 86m on a GTX1050Ti	new report	37,64		
65	Repeatable crash on large exponent	64	0.20	<p>After an otherwise unremarkable stage 1 computation, taking nearly 5 days, on a GTX1070, Using up to 6048M GPU memory.                      Selected B1=5745000, B2=157987500, 5.5% chance of finding a factor                      Starting stage 1 P-1, M628318517, B1 = 5745000, B2 = 157987500, fft length = 36864K ...                      Iteration 8280000 M628318517, 0x7b0d0aea8b5d52cf, n = 36864K, CUDAPm1 v0.20 err = 0.14453 (16:28 real, 49.3952 ms/iter, ETA 6:39)                      M628318517, 0x5cdc7c87c9b71cad, n = 36864K, CUDAPm1 v0.20                      Stage 1 complete, estimated total time = 115:20:44batch wrapper exiting at Mon 03/05/2018 13:41:24.14 Program terminates with an appcrash dialog box in Windows, indicating access violation: Problem signature:                      Problem Event Name: APPCRASH                      Application Name: CUDAPm1_win64_20131118_CUDA_50.exe                      Application Version: 0.0.0.0                      Application Timestamp: 5285815f                      Fault Module Name: CUDAPm1_win64_20131118_CUDA_50.exe                      Fault Module Version: 0.0.0.0                      Fault Module Timestamp: 5285815f                      Exception Code: c0000005                      Exception Offset:000000000000dcb4                      OS Version: 6.1.7601.2.1.0.256.48                      Locale ID: 1033                      Additional Information 1: 82af                      Additional Information 2: 82af96a4aa786ad0d7e18faf9f1c0171                      Additional Information 3: 55sec                      Additional Information 4: 55ecfab615f65e3435d809622e0ed32c</p> <p>The program crash occurs with the 20130923 cuda 55 version and other versions tried also. C0000005 is an Access Violation (such as reading past the end of a buffer). Attempts to continue the exponent, including simple restart, restart from a previous save file, run on a different gpu, run on a different system and gpu, run on the GTX1070 with threadbench completed for the fft length, increase fft length, alter the threads values in the threads file for the fft lengths used, run different versions CUDAPm1 executable, have all failed. While there is no output in the log about the gcd, the crash appears to occur early in stage 2 after the gcd, judging by GPU-Z sensor displays and Windows Task Manager cpu utilization display for the process. Reference "excessive stage 2 round-off errors simply halt the program without error messages." <a href="http://www.mersenneforum.org/showpost.php?p=353933&amp;postcount=373">http://www.mersenneforum.org/showpost.php?p=353933&amp;postcount=373</a></p>	new report	37,63		
66	input sanity checking / program output sanity checking	65	0.20 or 0.10?	It appears that the "over specifications Grid =" message output in v0.10 does not appear in v0.20, for the same conditions that produce it in v0.10. Threadbench in v0.20 20131118 executable produced ... 36864 512 32 64. 36864*1024/512 = 73728. v0.10 would warn about such user input and halt. If it's not acceptable input or parameters, the program ought avoid producing and storing it for future use, such as in threadbench operation.	new report	64, 73, 75		check other bug lists for applicability of this item

	A	B	C	D	E	F	G	H
67				<p>On a GTX1050Ti, an exponent repeatedly produced too high a max error and the program terminated with an estimated time spent 0:00 after 58500 iterations. On this gpu, running 4704k fft length, stage 1 iterations are about 15.9msec each so the duration would have been about 15.5 minutes (15:30). Possibly this is due to the absence of an fft file and threads file? Using up to 3858M GPU memory.</p> <p>Selected B1=740000, B2=17390000, 3.29% chance of finding a factor                      Starting stage 1 P-1, M86224591, B1 = 740000, B2 = 17390000, fft length = 4704K                      Doing 1067836 iterations                      Running careful round off test for 1000 iterations. If average error &gt; 0.25, the test will restart with a longer FFT.                      Iteration 100, average error = 0.18522, max error = 0.30078                      Iteration 200, average error = 0.21389, max error = 0.28906                      Iteration 300, average error = 0.22332, max error = 0.28906                      Iteration 400, average error = 0.22775, max error = 0.28125                      Iteration 500, average error = 0.23054, max error = 0.27734                      Iteration 600, average error = 0.23290, max error = 0.30078                      Iteration 700, average error = 0.23410, max error = 0.26758                      Iteration 800, average error = 0.23472, max error = 0.27344                      Iteration 900, average error = 0.23549, max error = 0.29297                      Iteration 1000, average error = 0.23591 &lt;= 0.25 (max error = 0.30078), continuing test.                      Iteration = 58500, err = 0.41809 &gt;= 0.40, quitting.                      Estimated time spent so far: 0:00 it apparently goes from the last saved checkpoint's computed run time (obtaining zero if there has not been one), and ignores the time spent on additional iterations since that point; CUDA reports 1428M of 1536M GPU memory free.                      Using threads: norm1 512, mult 256, norm2 512.                      Using up to 1323M GPU memory.                      Selected B1=725000, B2=16131250, 3.55% chance of finding a factor                      Starting stage 1 P-1, M81258607, B1 = 725000, B2 = 16131250, fft length = 4704K                      Doing 1046122 iterations                      Running careful round off test for 1000 iterations. If average error &gt; 0.25, the test will restart with a longer FFT.                      Iteration 100, average error = 0.03888, max error = 0.05957                      Iteration 200, average error = 0.04437, max error = 0.06348                      Iteration 300, average error = 0.04636, max error = 0.05664                      Iteration 400, average error = 0.04727, max error = 0.05859                      Iteration 500, average error = 0.04798, max error = 0.06250                      Iteration 600, average error = 0.04832, max error = 0.05731                      Iteration 700, average error = 0.04860, max error = 0.06445                      Iteration 800, average error = 0.04881, max error = 0.05933                      Iteration 900, average error = 0.04902, max error = 0.06592                      Iteration 1000, average error = 0.04914 &lt;= 0.25 (max error = 0.06592), continuing test.                      Iteration 50000 M81258607, 0xbb209b8b474a15b5, n = 4704K, CUDAPm1 v0.20 err = 0.07129 (7:22 real, 8.8367 ms/iter, ETA 2:26:42)                      Iteration = 75600, err = 0.49609 &gt;= 0.40, quitting.                      Estimated time spent so far: 7:22</p>				
	zero time estimate	66	0.20	That estimated time corresponds to 50,000 iterations at the GTX480's 8.84msec/iter timing.	new report	na		

	A	B	C	D	E	F	G	H
68				<p>If a run is continued on a gpu with more memory than it was started on, new bounds are calculated and then the program indicates it will continue with the bounds in the save file. After the run is completed, the result record contains the B2 found from the selection calculation, not the value from the save file that was used. Using threads: norm1 512, mult 256, norm2 128. Stage 2 checkpoint found.                      Using up to 3780M GPU memory.                      Selected B1=3100000, B2=<u>62000000</u>, 3.18% chance of finding a factor                      Using B1 = 2840000 from savefile.                      Continuing stage 2 from a partial result of M425000083 fft length = 24192K                      Starting stage 2.                      Using b1 = 2840000, b2 = 34080000, d = 2310, e = 2, nrp = 4                      Zeros: 1632727, Ones: 1613513, Pairs: 274647                      Processing 9 - 12 of 480 relative primes.                      Initializing pass... done. transforms: 235, err = 0.20313, (9.42 real, 40.0848 ms/tran, ETA 49:43:53)                      Transforms: 54058 M425000083, 0xed32e096fa463f09, n = 24192K, CUDAPm1 v0.20 err = 0.25439 (38:33 real, 42.7948 ms/tran, ETA 57:59:19)...                      Processing 477 - 480 of 480 relative primes.                      Initializing pass... done. transforms: 299, err = 0.21094, (12.89 real, 43.1271 ms/tran, ETA 39:01)                      Transforms: 53916 M425000083, 0x7efe91810f60cfa3, n = 24192K, CUDAPm1 v0.20 err = 0.25000 (38:28 real, 42.8098 ms/tran, ETA 0:46)</p> <p>Stage 2 complete, 6506485 transforms, estimated total time = 76:55:59                      Starting stage 2 gcd.                      M425000083 Stage 2 found no factor (P-1, B1=2840000, B2=<u>62000000</u>, e=2, n=24192K                      CUDAPm1 v0.20)</p>				
	B2 bounds reported is reselected value not continued value	67	0.20		new report	<a href="http://www.mersenneforum.org/showpost.php?p=490696&amp;postcount=569">http://www.mersenneforum.org/showpost.php?p=490696&amp;postcount=569</a>		

	A	B	C	D	E	F	G	H
69				<p>M300000991 Stage 1 found no factor (P-1, B1=2295000, B2=58555000, e=2, n=16800K CUDAPm1 v0.20)                      Starting stage 2.                      Using b1 = 2295000, b2 = 58555000, d = 4620, e = 2, nrp = 28                      cudaPm1_win64_20130923_CUDA_55.exe exited, attempting disable/enable cycle on gpu device at Sun 07/08/2018 5:46:53.30                      PCI\VEN_10DE&amp;DEV_06C0&amp;SUBSYS_14803842&amp;REV_A3\4&amp;1EF39D8B&amp;0&amp;0018: Disabled 1 device(s) disabled.                      PCI\VEN_10DE&amp;DEV_06C0&amp;SUBSYS_14803842&amp;REV_A3\4&amp;1EF39D8B&amp;0&amp;0018: Enabled 1 device(s) are enabled.                      batch wrapper reports (re)launch of cudaPm1_win64_20130923_CUDA_55.exe on Quadro 5000 at Sun 07/08/2018 5:46:56.17 reset count 1 of max 32                      CUDAPm1 v0.20                      ----- DEVICE 1 -----                      name Quadro 5000                      Compatibility 2.0                      clockRate (MHz) 1026                      memClockRate (MHz) 1500                      totalGlobalMem zu                      totalConstMem zu                      l2CacheSize 655360                      sharedMemPerBlock zu                      regsPerBlock 32768                      warpSize 32                      memPitch zu                      maxThreadsPerBlock 1024                      maxThreadsPerMP 1536                      multiProcessorCount 11                      maxThreadsDim[3] 1024,1024,64                      maxGridSize[3] 65535,65535,65535                      textureAlignment zu                      deviceOverlap 1</p> <p>CUDA reports 2471M of 2560M GPU memory free.                      Using threads: norm1 256, mult 128, norm2 64.                      No stage 2 checkpoint.                      Using up to 2362M GPU memory.                      Selected B1=2295000, B2=43031250, 3.31% chance of finding a factor                      Using B1 = 2295000 from savefile.                      Continuing stage 2 from a partial result of M300000991 fft length = 16800K                      Starting stage 2.                      Using b1 = 2295000, b2 = 43031250, d = 2310, e = 2, nrp = 8                      Zeros: 2053713, Ones: 2073807, Pairs: 363550                      Processing 1 - 8 of 480 relative primes.                      Initializing pass... )</p> <p>Quitting, estimated time spent = 0:04</p>				
	NRP value excessive on first try	68	0.20		new report	na		
70	silent fail very early in stage 1	69	0.20	<p>Some exponents will halt very early in stage 1 with no error message, just a notification quitting. These can be resolved by doubling the norm1 value in the threads file for the fft length being used. See also <a href="http://www.mersenneforum.org/showpost.php?p=491191&amp;postcount=570">http://www.mersenneforum.org/showpost.php?p=491191&amp;postcount=570</a></p>	new report	63		



	A	B	C	D	E	F	G	H
83								
84								
85								
86				Note, behavior analogous to CUDALucas issues has not in all cases been tested for yet				
87								
88	questions			do CUDAPM1, CUDALucas etc check for or enable or disable ECC on gpus? Per flashjh CUDALucas checks enable state but does not change it; see NVIDIA Control Panel for enabling /disabling for GPUs supporting ECC; there's a performance hit.				
89				Jacobi check applies somewhat to CUDAPm1 calculations, stage 1, stage 2, and to stage 1's 2 E p computation. Which are worthwhile?				
90				More reliable error detection at 0.1% overhead; is it applicable to the P-1 calculations? <a href="http://www.mersenneforum.org/showpost.php?p=465431&amp;postcount=88">http://www.mersenneforum.org/showpost.php?p=465431&amp;postcount=88</a>				