

ESBMC-GPU

A Context-Bounded Model Checking Tool to Verify CUDA Programs

Felipe R. Monteiro¹, Erickson H. da S. Alves¹, Isabela S. Silva¹,
Hussama I. Ismail¹, Lucas C. Cordeiro^{1,2}, and Eddie B. de Lima Filho^{1,3}

¹*Faculty of Technology, Federal University of Amazonas, Brazil*

²*Department of Computer Science, University of Oxford, United Kingdom*

³*TPV Technology Limited, Brazil*

Abstract

The Compute Unified Device Architecture (CUDA) is a programming model used for exploring the advantages of graphics processing unit (GPU) devices, through parallelization and specialized functions and features. Nonetheless, as in other development platforms, errors may occur, due to traditional software creation processes, which may even compromise the execution of an entire system. In order to address such a problem, ESBMC-GPU was developed, as an extension to the Efficient SMT-Based Context-Bounded Model Checker (ESBMC). In summary, ESBMC processes input code through ESBMC-GPU and an abstract representation of the standard CUDA libraries, with the goal of checking a set of desired properties. Experimental results showed that ESBMC-GPU was able to correctly verify 85% of the chosen benchmarks and it also overcame other existing GPU verifiers regarding the verification of data-race conditions, array out-of-bounds violations, assertive statements, pointer safety, and the use of specific CUDA features.

Keywords: GPU verification, formal verification, model checking, CUDA

1. Introduction

2 The Compute Unified Device Architecture (CUDA) is a development
3 framework that makes use of the architecture and processing power of graph-
4 ics processing units (GPUs) [1]. Indeed, CUDA is also an application pro-
5 gramming interface (API), through which a GPU's parallelization scheme
6 and tools can be accessed, with the goal of executing kernels [1]. Nonethe-
7 less, source code is still written by human programmers, which may result in

8 arithmetic overflow, division by zero, and other violation types. In addition,
9 given that CUDA allows parallelization, problems related to the latter can
10 also occur, due to thread scheduling [2].

11 In order to address the mentioned issues, an extension to the Efficient
12 SMT-Based Context-Bounded Model Checker (ESBMC) [3] was developed,
13 named as ESBMC-GPU [4, 5, 6], with the goal of verifying CUDA-based pro-
14 grams (available online at <http://esbmc.org/gpu/>). ESBMC-GPU consists
15 of an extension for parsing CUDA source code (*i.e.*, a front-end to ESBMC)
16 and a CUDA operational model (COM), which is an abstract representation
17 of the standard CUDA libraries (*i.e.*, the native API) that conservatively
18 approximates their semantics.

19 A distinct feature of ESBMC-GPU, when compared with other approaches
20 [2, 7, 8, 9], is the use of Bounded Model Checking (BMC) [10] allied to
21 Satisfiability Modulo Theories (SMT) [11], with explicit state-space explo-
22 ration [12, 3]. In summary, concurrency problems are tackled, up to a given
23 loop/recursion unwinding and context bound, while each interleaving itself
24 is symbolically handled; however, even with BMC, state-space exploration
25 may become a very time-consuming task, which is alleviated through state
26 hashing and Monotonic Partial Order Reduction (MPOR) [13]. As a conse-
27 quence, redundant interleavings are eliminated, without ignoring a program’s
28 behavior.

29 Finally, existing GPU verifiers often ignore some aspects related to mem-
30 ory leak, data transfer, and overflow, which are normally present in CUDA
31 programs. The proposed approach, in turn, explicitly addresses them, through
32 an accurate checking procedure, which even considers data exchange between
33 main program and kernel. Obviously, it results in higher verification times,
34 but more errors can then be identified and later corrected, in another devel-
35 opment cycle.

36
37 **Existing GPU Verifiers.** In addition to ESBMC-GPU, there are other
38 tools able to verify CUDA programs and each one of them uses its own ap-
39 proach and targets specific property violations. For instance, GPUVerify [2] is
40 based on synchronous, delayed visibility semantics, which focuses on detect-
41 ing data race and barrier divergence, while reducing kernel verification proce-
42 dures for the analysis of sequential programs. GPU+KLEE (GKLEE) [8], in
43 turn, is a concrete plus symbolic execution tool, which considers both *kernels*
44 and *main* functions, while checking deadlocks, memory coalescing, data race,
45 warp divergence, and compilation level issues. In addition, Concurrency In-
46 termediate Verification Language (CIVL) [9], a framework for static analysis
47 and concurrent program verification, uses abstract syntax tree and partial
48 order reduction to detect user-specified assertions, deadlocks, memory leaks,
49 invalid pointer dereference, array out-of-bounds, and division by zero.

50 In fact, ESBMC-GPU differs from the aforementioned approaches due to
 51 its combination of techniques to prune the state-space exploration (*i.e.*, two-
 52 thread analysis, state hashing, and MPOR) with COM, which demonstrated
 53 effectiveness in the verification of data-race conditions, array out-of-bounds
 54 violations, assertive statements, pointer safety, and the use of specific CUDA
 55 features (cf. Section 5).

56 2. Architecture and Implementation

57 ESBMC-GPU is built on top of ESBMC, which is an open-source context-
 58 bounded model checker based on SMT solvers for ANSI-C/C++ programs [12,
 59 3, 14], and adds four essential models, as described below.

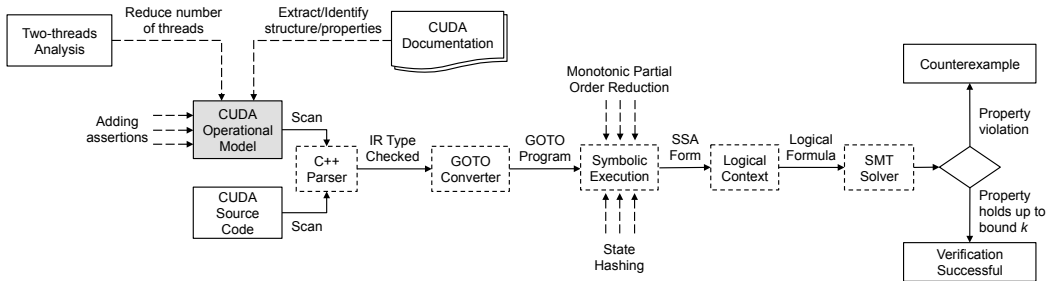


Figure 1: Overview of ESBMC-GPU’s architecture.

60 **CUDA Operational Model.** An operational model for CUDA libraries
 61 that provides support to CUDA functionalities, in conjunction with ESBMC,
 62 is shown in Fig. 1. Such an approach, which was previously attempted in the
 63 verification of C++ programs [14, 15, 16, 17], consists of an abstract repre-
 64 sentation that reliably approximates the CUDA library’s semantics; however,
 65 COM incorporates pre- and post-conditions into verification processes, which
 66 enables ESBMC-GPU to verify specific properties (cf. Section 3). Indeed,
 67 COM allows the necessary control for performing code analysis, where both
 68 CUDA operation and knowledge for model checking its properties are avail-
 69 able. Importantly, COM encloses a representation for the Runtime, Math,
 70 and cuRAND APIs, which are important CUDA libraries widely used in real
 71 applications [1]. In particular, with respect to the number of methods/functions,
 72 COM covers 53%, 31%, and 17% of Runtime, Math, and cuRAND
 73 APIs, respectively.

74 ESBMC was designed to handle multi-threaded software, through the use
 75 of an API called POSIX – ISO/IEC 9945 [18]. In order to support the veri-
 76 fication of CUDA kernels, COM applies code transformations to kernel calls
 77 using ESBMC’s intrinsic functions [6]. In particular, thread/block configu-
 78 rations in a CUDA kernel call are used as parameters in such intrinsic func-
 79 tions, which are responsible for checking for preconditions, configuring block

80 and threads dimension, and translating GPU threads to POSIX ones. Thus,
81 COM is able to support thread interleaving (to create execution paths) and
82 dynamic creation of threads, in order to check data race and specific C/C++
83 programming language failures (*e.g.*, array out-of-bounds and pointer safety).
84 ESBMC models the Pthreads API [12] to support thread synchronization,
85 *i.e.*, mutex locking operations and conditional waiting, and dynamic creation
86 of threads, which makes that representation very similar to the official CUDA
87 scheduler [6], in such a way that our multi-threading model approximates to
88 that of GPU kernels. Therefore, COM is able to use such aspects present in
89 ESBMC, in order to handle variables in different types of memory and also
90 in inter-warp communication.

91 The ESBMC’s memory model uses static pointer analysis, padding in
92 structures, with the goal of making all fields align to word boundaries, mem-
93 ory access alignment rules enforcement, and byte array allocation, when the
94 type of memory allocation is unclear [19]. Since ESBMC-GPU is built on top
95 of ESBMC, its memory model for the different types of memory is complete.

96
97 **Two-threads Analysis.** Similarly to GPUVerify [2] and PUG [7], ESBMC-
98 GPU also reduces the number of threads (to only two elements), during the
99 verification of CUDA programs, by considering a NVIDIA Fermi GPU archi-
100 tecture [1], in order to improve verification time and avoid the state-space
101 explosion problem. In CUDA programs, whilst threads execute the same
102 parametrized kernel, only two of them are necessary for conflict check. Thus,
103 such an analysis ensures that errors (*e.g.*, data races) detected between two
104 threads, in a given subgroup and due to unsynchronized accesses to shared
105 variables, are enough to justify a property violation [6].

106
107 **State Hashing.** ESBMC-GPU applies state hashing to further eliminate
108 redundant interleavings and also reduce the state space, based on SHA256
109 hashes [20]. In particular, its symbolic state hashing approach computes
110 a summary for a particular state that has already been explored and then
111 indexes the resulting set, in order to reduce the generation of redundant
112 states. Given any state computed during the symbolic execution of a specific
113 CUDA kernel, ESBMC-GPU simply summarizes it and efficiently determines
114 whether it has been explored before or not, along a different computation
115 path. When this behavior is confirmed, which happens during the ESBMC-
116 GPU’s symbolic-execution procedure, then the current computation path
117 does not need to be further explored in the associated reachability tree (RT).
118 This way, if ESBMC-GPU reaches such a state, *i.e.*, where a context switch
119 can be taken (*e.g.*, before a global variable or synchronization primitive) and
120 all shared/local variables and program counters are similar to another ex-
121 plored node, then ESBMC-GPU just considers that an identical node to be

122 further explored, since reachability subtrees associated to them are also sim-
123 ilar [6, 21].

124

125 **Monotonic Partial Order Reduction.** MPOR is used to reduce the
126 number of thread interleavings, by classifying transitions inside a program
127 as dependent or independent. As a consequence, it is possible to determine
128 whether interleaving pairs always lead to the same state and then remove
129 duplicates in a RT, without ignoring any program’s behavior [21].

130 3. Functionalities

131 COM models CUDA libraries and provides a multi-threading model much
132 similar to the CUDA scheduler (cf. Section 2). Moreover, it is able to simu-
133 late CUDA’s program structure and memory, being susceptible of handling
134 CUDA programs. Thus, through the integration of COM into ESBMC (*i.e.*,
135 ESBMC-GPU), one is able to analyze CUDA programs and verify the fol-
136 lowing properties:

137 **Data race.** ESBMC-GPU checks data race conditions, in order to
138 detect if multiple threads perform unsynchronized access to the same
139 memory locations;

140 **Pointer safety.** ESBMC-GPU also ensures that (*i*) a pointer offset
141 does not exceed object bounds and (*ii*) a pointer is neither NULL nor
142 invalid;

143 **Array bounds.** ESBMC-GPU performs array-bound checking, in or-
144 der to ensure that any variable, used as an array index, is within known
145 bounds;

146 **Arithmetic under- and overflow.** ESBMC-GPU checks whether a
147 sum or product exceeds the memory limits that a variable can han-
148 dle, which can cause an error capable of spreading through the entire
149 execution path;

150 **Division by zero.** ESBMC-GPU analyzes whether denominators, in
151 arithmetic expressions, lead to divisions by zero;

152 **User-specified assertions.** ESBMC-GPU considers all assertions
153 specified by users, which is essential to a thorough verification process,
154 as some specific possible violations must be explicitly pointed out.

155 In order to check the aforementioned properties, ESBMC-GPU explicitly
156 explores the possible interleavings (up to the given context bound) and calls

157 the single-threaded BMC procedure on each one, whenever it reaches a RT
158 leaf node. Then, the mentioned procedure will stop if it finds a bug or when
159 all possible RT interleavings have been systematically explored [6]. Further-
160 more, ESBMC-GPU has the following additional command-line options:

161 `--no-assertions`: to ignore assertions;
162 `--no-bounds-check`: to skip array bounds check;
163 `--no-div-by-zero-check`: to skip division by zero check;
164 `--no-pointer-check`: to skip pointer check;
165 `--memory-leak-check`: to enable memory leak check;
166 `--overflow-check`: to enable arithmetic under- and overflow check;
167 `--deadlock-check`: to enable global/local deadlock check with mutex;
168 `--data-races-check`: to enable data races check;
169 `--lock-order-check`: to enable for lock acquisition ordering check;
170 `--atomicity-check`: to enable atomicity check at visible assignments;
171 `--force-malloc-success`: to consider that there is always enough
172 memory available in the device;

173 Thus, ESBMC-GPU is able to check CUDA programs for: deadlock, assertion,
174 lock acquisition error, division by zero, pointer safety, arithmetic over-
175 flow, and/or out-of-bounds array violation. The precision and performance
176 of ESBMC-GPU will be further discussed in Section 5.

177 4. Illustrative Example

178 In this part, ESBMC-GPU usage is demonstrated, by using the CUDA
179 program shown in Fig. 2. First of all, users must replace the default kernel
180 call (line 16) by an intrinsic function of ESBMC-GPU (line 17). Then,
181 the resulting CUDA program can be passed to the command-line version of
182 ESBMC-GPU, as follows:

```
183     esbmc-gpu <file>.cu --unwind <k> --context-switch <c>  
184             --state-hashing -I <path-to-CUDA-OM>,
```

```

1 #include <...>
2 #define BLOCKS 1
3 #define THREADS 2
4
5 --global-- void kernel(int *A) {
6     A[threadIdx.x + 1] = threadIdx.x;
7 }
8
9 int main(){
10     int *a;
11     int *dev_a;
12     int size = THREADS*sizeof(int);
13     a = (int*)malloc(size);
14     cudaMalloc((void*)&dev_a, size);
15     for (int i = 0; i < THREADS; i++)
16         a[i] = 0;
17     cudaMemcpy(dev_a, a, size, cudaMemcpyHostToDevice);
18     // kernel<<<BLOCKS, THREADS>>>(dev_a);
19     ESBMC_verify_kernel(kernel, BLOCKS, THREADS, dev_a);
20     cudaMemcpy(a, dev_a, size, cudaMemcpyDeviceToHost);
21     for (int i = 0; i < THREADS; i++)
22         assert(a[i]==i);
23     cudaFree(dev_a);
24     free(a);
25     return 0;
26 }

```

Figure 2: Illustrative CUDA code example.

185 where `<file>.cu` is the CUDA program, `<k>` is the maximum loop un-
186 rolling, `<c>` is a context-switch bound, `--state-hashing` reduces redundant
187 interleavings, and `<path-to- CUDA-OM>` is the location of the COM library.

188 In the mentioned example, ESBMC-GPU detects an array out-of-bounds
189 violation. Indeed, this CUDA-based program retrieves a memory region that
190 has not been previously allocated, *i.e.*, when `threadIdx.x = 1`, the program
191 tries to access `a[2]`. Importantly, the `cudaMalloc()` function’s operational
192 model has a precondition that checks if the memory size to be allocated is
193 greater than zero. In addition, an assertion checks if the result matches to the
194 expected postcondition (line 22). The verification of this program through
195 ESBMC-GPU produces 54 successful and 3 failed interleavings. For instance,
196 one possible failed interleaving is represented by the threads executions $t_0 :$
197 $a[1] = 0; t_1 : a[2] = 1$, where $a[2] = 1$ represents an incorrect access to the
198 array index a . It is worth noticing that CIVL, ESBMC-GPU, and GKLEE
199 are also able to detect this array out-of-bounds violation, but GPUVerify
200 fails, as it reports a true incorrect result (*i.e.*, a missed bug).

201 5. Experimental Evaluation

202 In order to evaluate ESBMC-GPU’s precision and performance, a bench-
203 mark suite was created, which comprises 20 CUDA kernels from NVIDIA
204 GPU Computing SDK *v2.0* [22], 20 CUDA kernels from Microsoft C++
205 AMP Sample Projects [23], and 114 CUDA-based programs that explore a
206 wide range of CUDA functionalities. In summary, the chosen suite contains
207 47.4% bug-free and 52.6% buggy benchmarks, which were organized into 5

208 sets, in order to simplify our discussion, according to the properties it tackles:
 209 array bounds (5), assertive statements (7), data-race conditions (17), pointer
 210 safety (7), and other specific CUDA functionalities (118). The latter includes
 211 `__device__` function calls, general CUDA functions (*e.g.*, `cudaMemcpy`), gen-
 212 eral libraries in CUDA (*e.g.*, `curand.h`), type modifiers (*e.g.*, `unsigned`), type
 213 definitions, and intrinsic CUDA variables (*e.g.*, `uint4`).

214 The present experiments answer two research questions: (i) How accu-
 215 rate is ESBMC-GPU when verifying the chosen benchmarks? (ii) How does
 216 ESBMC-GPU’s performance compare to other existing verifiers? In order to
 217 answer both research questions, all benchmarks were verified with 4 GPU
 218 verifiers (ESBMC-GPU *v2.0*, GKLEE *v2012*, GPUVerify *v1811*, and CIVL
 219 *v1.7.1*), on an otherwise idle Intel Core i7-4790 CPU 3.60 GHz, with 16 GB
 220 of RAM, running Ubuntu 14.04 OS. Importantly, all presented execution
 221 times are actually CPU times, *i.e.*, only the elapsed time periods spent in
 222 the allocated CPUs, which was measured with the `times` system call (POSIX
 223 system). An overview of the experimental results is shown in Fig. 3, where
 224 *True* represents bug-free benchmarks, *False* represents buggy benchmarks,
 225 *Not supported* represents benchmarks that could not be verified, *Correct* rep-
 226 represents the percentage of benchmarks correctly verified, and *Incorrect* rep-
 227 represents the percentage of benchmarks incorrectly verified (*i.e.*, a verification
 228 tool reports an unexpected result).




Figure 3: Experimental evaluation of ESBMC-GPU against other verifiers.

229 As one may notice, the present experimental results show that ESBMC-
 230 GPU reached a successful verification rate of approximately 85%, while GK-
 231 LEE, GPUVerify, and CIVL reported 72%, 50%, and 35%, respectively. More
 232 precisely, ESBMC-GPU correctly detected all data-race conditions present
 233 in the benchmarks, which is due to the COM’s multi-threading model that
 234 under-approximates the GPU kernels. It also outperformed GKLEE, GPU-

235 Verify, and CIVL, in the verification of array out-of-bounds violations (100%),
236 assertive statements (86%), and pointer safety (72%), which is related to
237 ESBMC’s capacity to handle arrays and pointers [3]. Furthermore, ESBMC-
238 GPU presented the highest coverage rate for specific CUDA functionalities
239 (82%) that is once again due to COM, which incorporates specific pre- and
240 post-conditions into its verification processes.

241

242 **Limitations.** ESBMC-GPU was unable to correctly verify 24 benchmarks,
243 which are related to constant memory access (2%), CUDA’s specific libraries
244 (*e.g.*, `curand.h`) (4.5%), and the use of pointers to functions, structures, and
245 `char` type variables, when passed as kernel call arguments (4.5%). In addi-
246 tion, it only reported 3% of incorrect true, which are due to NULL pointer
247 accesses, and 1% of incorrect false results, due to partial coverage of the *cu-*
248 *daMalloc* function for copies over `float` variables. The remaining verifiers
249 (*i.e.*, GKLEE, GPUVerify, and CIVL) were unable to detect mostly data-
250 race conditions, assertive statements, and array out-of-bounds violations.
251 In addition, they lack support of CUDA specific features, *e.g.*, GPUVer-
252 ify does not support the use of the `memset` function nor function point-
253 ers and CIVL does not support several CUDA features, such as `atomic`
254 functions, `cudaThreadSynchronize`, `threadIdx`, `curand` functions, `dim3`,
255 `math_functions`, `uint4`, `__constant__` variables, among others.

256

257 **Performance.** MPOR resulted in a performance improvement of approx-
258 imately 80%, by decreasing the verification time from 16 to 3 hours, while
259 the two-threads analysis further reduced that to 789.6 sec. Although such
260 techniques have considerably improved the ESBMC-GPU’s performance, it
261 still takes longer than the other evaluated tools: GPUVerify (98.36 sec), GK-
262 LEE (105.18 sec), and CIVL (708.52 sec). On the one hand, this is due to
263 thread interleavings, which combine symbolic model checking with explicit
264 state-space exploration [6]. On the other hand, ESBMC-GPU still presents
265 the highest accuracy, with less than 6 seconds per benchmark.

266

267 **Availability of Data and Tools.** The performed experiments are based on
268 a set of publicly available benchmarks. All benchmarks, tools, and results,
269 associated with the current evaluation, are available at www.esbmc.org/gpu/.

270 6. Conclusions and Future Work

271 ESBMC-GPU marks the first application of an SMT-based context-BMC
272 tool that recognizes CUDA directives [6]. Besides, it also applies MPOR,
273 two-thread analysis, and state hashing, in order to further simplify verifi-
274 cation models and provides fewer incorrect results, compared with GKLEE,

275 GPUVerify, and CIVL. Indeed, it presents improved ability to detect array
276 out-of-bounds and data race violations.

277 Future work aims to extend ESBMC-GPU, in order to fully support the
278 verification of CUDA (parallel) streams and events [1]. In addition, more
279 models of libraries will be integrated into COM, with the goal of increas-
280 ing the coverage of CUDA’s API such as CUDA Driver API, NPP, and cu-
281 SOLVER. Finally, we also aim to implement further techniques (*e.g.*, invari-
282 ant inference via abstract interpretation [24]), in order to prune the state-
283 space exploration, by taking into account GPU symmetry.

284 Acknowledgements

285 This paper is based on research sponsored by the Institute of Development
286 and Technology (INdT) and by the National Council for Scientific and Tech-
287 nological Development (CNPq) under agreement number 475647/2013 – 0.

288 References

- 289 [1] J. Cheng, M. Grossman, T. McKercher, *Professional CUDA C Program-*
290 *ming*, John Wiley and Sons, Inc., Indianapolis, Indiana, USA, 2014.
- 291 [2] A. Betts, N. Chong, A. Donaldson, S. Qadeer, P. Thomson, *GPUVer-*
292 *ify: A Verifier for GPU Kernels*, in: Proceedings of the ACM Interna-
293 tional Conference on Object Oriented Programming Systems Languages
294 and Applications, ACM, New York, NY, USA, 2012, pp. 113–132, URL
295 <http://doi.acm.org/10.1145/2384616.2384625>.
- 296 [3] L. Cordeiro, B. Fischer, J. Marques-Silva, *SMT-Based Bounded Model*
297 *Checking for Embedded ANSI-C Software*, IEEE Trans. Software Eng.
298 38 (2012) 957–974, URL <https://doi.org/10.1109/TSE.2011.59>.
- 299 [4] P. Pereira, H. Albuquerque, H. Marques, I. Silva, C. Carvalho, V. San-
300 tos, R. Ferreira, L. Cordeiro, *Verificação de Kernels em Programas*
301 *CUDA usando Bounded Model Checking*, in: XV Brazilian Symposium
302 on High-Performance Computing (WSCAD-SSC), Florianópolis, Brazil,
303 2015, pp. 24–35.
- 304 [5] Pereira, P., Albuquerque, H., Marques, H., Silva, I., Carvalho, C.,
305 Santos, V., Ferreira, R., and Cordeiro, L. *Verifying CUDA Pro-*
306 *grams using SMT-Based Context-Bounded Model Checking*. in: Pro-
307 ceedings of the 31st Annual ACM Symposium on Applied Com-
308 puting, ACM, New York, NY, USA, 2016, pp. 1648–1653, URL
309 <http://doi.acm.org/10.1145/2851613.2851830>.

- 310 [6] P. Pereira, H. Albuquerque, I. Silva, H. Marques, F. R. Monteiro, R.
311 Ferreira, L. Cordeiro, *SMT-Based Context-Bounded Model Checking for*
312 *CUDA Programs*, Concurrency Computat.: Pract. Exper. (2016), URL
313 <http://dx.doi.org/10.1002/cpe.3934>.
- 314 [7] G. Li, G. Gopalakrishnan, *Scalable SMT-based Verification of*
315 *GPU Kernel Functions*, in: Proceedings of the 18th ACM SIG-
316 SOFT International Symposium on Foundations of Software Engi-
317 neering, ACM, New York, NY, USA, 2010, pp. 187–196, URL
318 <http://doi.acm.org/10.1145/1882291.1882320>.
- 319 [8] G. Li, P. Li, G. Sawaya, G. Gopalakrishnan, I. Ghosh, S. Rajan, *GK-*
320 *LEE: Concolic Verification and Test Generation for GPUs*, in: Proceed-
321 ings of the 17th ACM SIGPLAN Symposium on Principles and Practice
322 of Parallel Programming, ACM, New York, NY, USA, 2012, pp. 215–
323 224, URL <http://doi.acm.org/10.1145/2145816.2145844>.
- 324 [9] M. Zheng, M. Rogers, Z. Luo, M. Dwyer, S. Siegel, *CIVL: Formal Ver-*
325 *ification of Parallel Programs*, in: Proceedings of the 30th IEEE/ACM
326 International Conference on Automated Software Engineering, IEEE
327 Computer Society, Washington, DC, USA, 2015, pp. 830–835, URL
328 <http://dx.doi.org/10.1109/ASE.2015.99>.
- 329 [10] A. Biere, *Bounded Model Checking*, in: A. Biere, M. Heule, H. van
330 Maaren, T. Walsh (Eds.), Handbook of Satisfiability, IOS Press, Am-
331 sterdam, The Netherlands, 2009, pp. 457–481.
- 332 [11] C. Barrett, R. Sebastiani, S. Seshia, C. Tinelli, *Satisfiability Modulo*
333 *Theories*, in: A. Biere, M. Heule, H. van Maaren, T. Walsh (Eds.),
334 Handbook of Satisfiability, IOS Press, Amsterdam, The Netherlands,
335 2009, pp. 825–885.
- 336 [12] L. Cordeiro, B. Fischer, *Verifying Multi-threaded Software us-*
337 *ing SMT-based Context-Bounded Model Checking*, in: Proceedings
338 of the 32nd ACM/IEEE International Conference on Software En-
339 gineering, ACM, New York, NY, USA, 2011, pp. 331–340, URL
340 <http://doi.acm.org/10.1145/1810295.1810396>.
- 341 [13] V. Kahlon, C. Wang, A. Gupta, *Monotonic Partial Order Reduction:*
342 *An Optimal Symbolic Partial Order Reduction Technique*, in: Proceed-
343 ings of the 21st International Conference on Computer Aided Verifi-
344 cation, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 398–413, URL
345 http://dx.doi.org/10.1007/978-3-642-02658-4_31.

- 346 [14] M. Ramalho, M. Freitas, F. R. Monteiro, H. Marques, L. Cordeiro,
347 B. Fischer, *SMT-Based Bounded Model Checking of C++ Programs*,
348 in: Proceedings of the 20th Annual IEEE International Conference
349 and Workshops on the Engineering of Computer Based Systems, IEEE
350 Computer Society, Washington, DC, USA, 2013, pp. 147–156, URL
351 <http://dx.doi.org/10.1109/ECBS.2013.15>.
- 352 [15] F. R. Monteiro, L. Cordeiro, E. B. de Lima Filho, *Bounded Model*
353 *Checking of C++ Programs Based on the Qt Framework*, in: IEEE
354 4th Global Conference on Consumer Electronics, IEEE Consumer Elec-
355 tronics Society, Washington, DC, USA, 2015, pp. 179–447, URL
356 <https://doi.org/10.1109/GCCE.2015.7398699>.
- 357 [16] M. Garcia, F. R. Monteiro, L. Cordeiro, E. B. de Lima Filho,
358 *ESBMC^{QtOM}: A Bounded Model Checking Tool to Verify Qt Appli-*
359 *cations*, in: 23rd International Symposium on Model Checking Soft-
360 ware, Springer International Publishing, Cham, 2016, pp. 97–103, URL
361 http://dx.doi.org/10.1007/978-3-319-32582-8_6.
- 362 [17] F. R. Monteiro, M. Garcia, L. Cordeiro, E. B. de Lima Filho,
363 *Bounded Model Checking of C++ Programs based on the Qt Cross-*
364 *Platform Framework*, *Softw Test Verif Reliab.* 27 (2017) e1632, URL
365 <http://dx.doi.org/10.1002/stvr.1632>.
- 366 [18] Institute of Electrical and Electronics Engineers Inc., *IEEE*
367 *Standard for Information Technology-Portable Operating System*
368 *Interface (POSIX) Base Specifications*, IEEE Std. 1003.1-
369 2008 (Revision of IEEE Std. 1003.1-2004). (2008) c1-3826, URL
370 <https://doi.org/10.1109/IEEESTD.2008.4694976>.
- 371 [19] J. Morse, M. Ramalho, L. Cordeiro, D. Nicole, B. Fischer, *ES-*
372 *BMC 1.22 (Competition Contribution)*, in: 20th International Con-
373 ference on Tools and Algorithms for the Construction and Analysis
374 of Systems, Springer, Berlin, Heidelberg, 2014, pp. 405–407, URL
375 http://dx.doi.org/10.1007/978-3-642-54862-8_31.
- 376 [20] National Institute of Standards and Technology, *Secure Hash Standard*
377 *(SHS)*, Federal Information Processing Standards Publication 180-4.
378 (2015), URL <http://dx.doi.org/10.6028/NIST.FIPS.180-4>.
- 379 [21] J. Morse, *Expressive and Efficient Bounded Model Checking of Concur-*
380 *rent Software*, University of Southampton, PhD Thesis, Southampton,
381 UK, 2015.

- 382 [22] NVIDIA Corporation, *CUDA Toolkit Archive*. <https://developer.nvidia.com/cuda-toolkit-archive>, 2015 (accessed 01.06.2017).
383
- 384 [23] D. Moth, *C++ AMP Sample Projects for Download*. Microsoft Corpora-
385 tion, [blogs.msdn.com/b/nativeconcurrency/archive/2012/01/30/
386 c-amp-sample-projects-for-download.aspx](blogs.msdn.com/b/nativeconcurrency/archive/2012/01/30/c-amp-sample-projects-for-download.aspx), 2012 (accessed
387 01.06.2017).
- 388 [24] W. Rocha, H. Rocha, H. Ismail, L. Cordeiro, B. Fischer, *DepthK: A k-
389 Induction Verifier Based on Invariant Inference for C Programs*, in: *23rd
390 International Conference on Tools and Algorithms for the Construction
391 and Analysis of Systems*, Springer, Berlin, Heidelberg, 2017, pp. 360–
392 364, URL http://dx.doi.org/10.1007/978-3-662-54580-5_23.

393 **Required Metadata**

394 **Current executable software version**

395 Ancillary data table required for sub version of the executable software.

Nr.	(executable) Software metadata description	Please fill in this column
S1	Current software version	2.0
S2	Permanent link to executables of this version	http://esbmc.org/gpu/
S3	Legal Software License	Apache v2.0
S4	Computing Operating System	Ubuntu Linux OS
S5	Installation requirements & dependencies	GNU Libtool; Automake; Flex & Bison; Boost C++ Libraries; Multi-precision arithmetic library developers tools (libgmp3-dev package); SSL development libraries (libssl-dev package); CLang 3.8; LLDB 3.8; GNU C++ compiler (multilib files); libc6 and libc6-dev packages
S6	Link to user manual	http://esbmc.org/gpu/
S7	Support email for questions	lucas.cordeiro@cs.ox.ac.uk

Table 1: Software metadata (optional)

396 **Current code version**

397 Ancillary data table required for subversion of the codebase.

Nr.	Code metadata description	Please fill in this column
C1	Current code version	<i>v2.0</i>
C2	Permanent link to code/repository used for this code version	https://github.com/ssvlab/esbmc-gpu
C3	Legal Code License	GNU Public License
C4	Code versioning system used	git
C5	Software code languages, tools, and services used	C++
C6	Compilation requirements, operating environments & dependencies	GNU Libtool; Automake; Flex & Bison; Boost C++ Libraries; Multiprecision arithmetic library developers tools (libgmp3-dev package); SSL development libraries (libssl-dev package); CLang 3.8; LLDB 3.8; GNU C++ compiler (multilib files); libc6 and libc6-dev packages
C7	Link to developer documentation	http://esbmc.org/gpu
C8	Support email for questions	lucas.cordeiro@cs.ox.ac.uk

Table 2: Code metadata (mandatory)