

# C++ Dynamic Cast in Autonomous Space Systems

Damian Dechev<sup>1</sup>, Rabi Mahapatra<sup>1</sup>, Bjarne Stroustrup<sup>1</sup>, David Wagner<sup>2</sup>  
dechev@tamu.edu, rabi@cs.tamu.edu, bs@cs.tamu.edu, david.a.wagner@jpl.nasa.gov

Texas A&M University<sup>1</sup>  
College Station, TX 77843-3112

Jet Propulsion Laboratory, California Institute of Technology<sup>2</sup>  
4800 Oak Grove Drive, M/S 301-270, Pasadena, CA

## Abstract

*The dynamic cast operation allows flexibility in the design and use of data management facilities in object-oriented programs. Dynamic cast has an important role in the implementation of the Data Management Services (DMS) of the Mission Data System Project (MDS), the Jet Propulsion Laboratory's experimental work for providing a state-based and goal-oriented unified architecture for testing and development of mission software. DMS is responsible for the storage and transport of control and scientific data in a remote autonomous spacecraft. Like similar operators in other languages, the C++ dynamic cast operator does not provide the timing guarantees needed for hard real-time embedded systems. In a recent study, Gibbs and Stroustrup (G&S) devised a dynamic cast implementation strategy that guarantees fast constant-time performance. This paper presents the definition and application of a co-simulation framework to formally verify and evaluate the G&S fast dynamic casting scheme and its applicability in the Mission Data System DMS application. We describe the systematic process of model-based simulation and analysis that has led to performance improvement of the G&S algorithm's heuristics by about a factor of 2.*

## 1 Introduction

<sup>1</sup> ISO Standard C++ [8] has become a common choice for hard real-time embedded systems such as the Jet Propulsion Laboratory's Mission Data System [7]. This is so because ISO C++ offers efficient abstraction model, good hardware use, and predictability. C++'s model of computation has helped engineers deliver more correct, maintainable, and comprehensible software compared to code rely-

ing on lower-level programming concepts [19]. However, several C++ features are usually considered unsuitable for programming real-time systems because they do not guarantee predictable constant-time performance [5]. ISO C++ does not provide the necessary timing guarantees for free store (heap) allocation, exception handling, and dynamic casting. In particular, the most common compiler implementations of the dynamic cast operator traverse the representation of the inheritance tree (at run time) searching for a match. Such implementations of dynamic cast are not predictable and are unsuitable for real-time programming. Gibbs and Stroustrup(G&S) [3] describe a technique for implementing dynamic cast that delivers significantly improved and constant-time performance. The key idea is to replace the runtime search through the class hierarchy with a simple (constant-time) calculation, much as the common implementations of the C++ virtual function calls search the class hierarchy at compile time to reduce the runtime action to a simple array subscripting operation. In the G&S scheme, a heuristic algorithm assigns an integer type ID at link time to each class. The type ID assignment rules guarantee that at run time a simple modulo operation can reveal the type information and check the validity of the cast. The requirements for the heuristics assigning the type IDs are that:

- (1) They must keep the size of the type ID to a small number of bits. A 64-bit type ID should be sufficient for very large class hierarchies
- (2) Avoid conflicts and type ID assignments that create ambiguous or erroneous type resolution at run time
- (3) Handle virtual inheritance

There are four heuristic schemes and a few possible optimizations suggested in [3]. However, none of those heuristics guarantee the best solution for every possible class hierarchy. The quality of the type ID assignment heuristics has a critical importance for the performance of the G&S scheme. With better heuristics, a smaller type ID size would be sufficient to facilitate complex and large class hierarchies that

<sup>1</sup>This is the authors' version of the work. It is posted here by permission of the publisher. Not for redistribution. The definitive version is published in Proceedings of 11th IEEE International Symposium on Object/component/service-oriented Real-time Distributed Computing (IEEE ISORC 2008), Orlando, Florida, May 2008.

would certainly need a significantly bigger type ID size with a poor assignment scheme. The main contribution of this work is to present how the algorithm optimization problem encountered has been successfully automated and moreover that its automation has led us to quick but significant improvements of the initial scheme.

As pointed out by Lowry [10], the increasing complexity of future space missions, such as the Mars Science Laboratory [20] and Project Constellation [17], raises concerns whether it is possible to establish their reliability in a cost-effective manner. Lowry's analysis indicates that at the present moment the verification and certification cost of mission critical software exceeds its development cost. Perrow [13] studies the risk factors in the modern high technology systems. His work identifies two significant hazard dimensions: *interactions* and *coupling*. Complex interactions represent unexpected and unknown sequences and thus cannot be entirely comprehensible at the time of system development. A tightly-coupled system has a number of time-dependent processes that cannot tolerate delays. Perrow classifies space missions in the riskiest category since both hazard factors are present. The dominant paradigms for software development, assurance, and management at NASA rely on the principle "test-what-you-fly and fly-what-you-test". Born out of experience and hindsight, this methodology had been applied in a large number of robotic space missions at the Jet Propulsion Laboratory. For such missions, it has proven suitable in achieving adherence to some of the most stringent standards of man-rated certification such as the DO-178B [14], the Federal Aviation Administration (FAA) software standard. Its Level A requirements demand 100% coverage of all high and low level assurance policies. However, the present certification methodologies are prohibitively expensive for systems of high complexity [16].

In this paper we present a co-simulation framework based on the SPIN model checker [6] to simulate, evaluate, and formally verify the G&S fast dynamic casting algorithm and its application in mission critical code such as the Data Management Services [21] of the Mission Data System. The aim of the Mission Data System is to provide a unified state-based and goal-oriented architecture for building complete data and control systems for autonomous space missions. The framework's state- and model-based methodology and its associated systems engineering processes and development tools have been successfully applied on a number of test systems including the physical rovers Rocky 7 and Rocky 8 and a simulated Entry, Descent, and Landing (EDL) system for the Mars Science Laboratory mission. We use the feedback from the model checker to perform systematic analysis of the G&S scheme and look for improvements to the heuristics for type ID assignment. SPIN is an on-the-fly, linear-time logic model-

checking tool that was designed for the formal verification of dynamic systems with asynchronously executed processes. The most recent advances in the state space reduction techniques has made it possible to validate large software applications. Model-checking tools have been widely applied for the verification of a large variety of systems, including flight software [4], network protocols [11], and scheduling algorithms [15]. We are unaware of work suggesting its use for the analysis and optimization of compiler heuristics. Compiler verification usually focuses on seeking a proof on the preservation of the program semantics during the various compiler passes [9]. Our work presents the application of a model-checking tool for the analysis and refinement of the combinatorial optimization problem posed by the G&S type ID assignment scheme. Our co-simulation framework consists of the following components:

- (1) An abstract model of the G&S type ID assignment heuristics
- (2) An procedure for exhaustive search of the state space discovering the best type ID assignment

The analysis of the heuristics simulation performed in SPIN provides us with ideas of possible improvements to the G&S type ID assignment. We include and evaluate the proposed improvements in the abstract model in order to seek refinement of the G&S type ID assignment scheme. The experiments we have executed show that the G&S priority assignment is not optimal with respect to the best possible type ID assignment where non-virtual multiple inheritance is used. While potentially dangerous if not constructed carefully, such hierarchies happen to be of significant practical importance [18]. Based on our experiments, we suggest optimizations that lead to significant improvement of the G&S heuristics performance. This paper makes the following contributions:

- (1) Introduces the use of a co-simulation framework based on model-checking for the analysis and improvement of a compiler-heuristics optimization problem
- (2) Verifies and analyzes the G&S C++ fast dynamic casting scheme and its application in mission critical code such as the MDS Data Management Services
- (3) Implements optimizations to the G&S heuristics leading to the discovery of optimal type ID assignment in 85% of the class hierarchies, in contrast to 48% for the original G&S algorithm

The rest of the paper is organized as follows: section 2: a brief description of the G&S fast dynamic cast algorithm, section 3: our approach to co-simulation and improvements to the G&S heuristics, section 4: discussion on the challenges of mission critical code and the applicability of the

G&S dynamic cast section 5: performance results for the G&S algorithm and the proposed improvements, and section 6: conclusion.

## 2 Fast Dynamic Casting Algorithm

The G&S fast constant-time implementation of the dynamic cast operator works as follows: at link time, a static integer type ID number, preferably 32 or 64-bit long, is assigned to each class. The ID numbers are selected so that the operation  $id_a \text{ modulus } id_b$  yields zero if and only if the class with  $id_a$  is derived from the class with  $id_b$ . This is done by exploiting the uniqueness of factorization of integers into prime factors. Each class is assigned a *key* prime number. The *type ID* of a class is calculated by multiplying its *key number* with the key numbers of each of its base classes. In the cases where a class contains more than a single copy of a base class, the type ID is computed by taking the square of the corresponding base class ID. The only constraint of the approach is the desire to limit the ID size to fit the machine's built-in integer types. The key primes are not required to be unique and the same prime key can be used for classes that belong to different groups (i.e. do not share common descendants). Gibbs and Stroustrup suggest four approaches for assigning the type IDs in a space-efficient manner. Each method is based on a preliminary computation of the priority factor of each class. The priority reflects the class impact on the growth of the type ID numbers in the hierarchy. Thus, classes with greater number of descendants should receive higher priority and smaller key prime number values respectively. The four possible schemes suggest that:

- 1 The priority of a class is the *maximum* number of ancestors that any of its descendants has. This scheme was chosen for the initial implementation and testing of the G&S algorithm and also closely followed in the implementation of the abstract model used for our simulation
- 2 ,3, 4. If a range of primes is assigned to every level with wider levels receiving larger initial values, then each node could be assigned an additional value that is proportional to the logarithm of the (2. *minimum*, 3. *mean*, 4. *maximum*) prime in its level. Priorities of hierarchy leaves are computed by taking the sum of these additional values for the leaf itself and all of its ancestor classes

After the priority of each class has been computed, the classes with the highest priority get the smallest prime numbers. According to this scheme, prime numbers can be reused only if there are two classes on the same level of the class hierarchy and only if they do not share common descendants, they are not siblings, and also that none of their parents share a common descendant. According to the ID assignment rules, we know that:




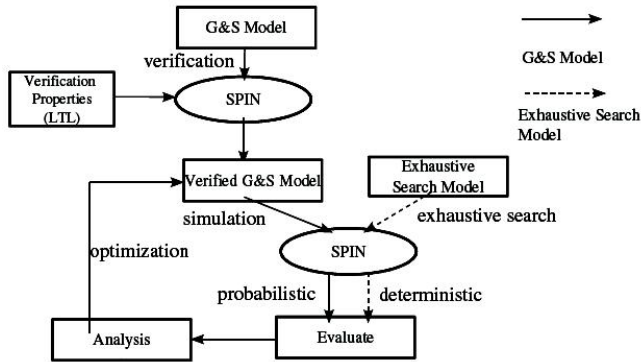
Figure 1. A class hierarchy with 11 classes

- (1)  $id_x = k_x \times (k_a)^2 \times k_{a_1} \times (k_b)^2 \times k_{b_1} \times k_c$
- (2)  $id_y = k_y \times k_c \times k_{c_1} \times (k_d)^2 \times k_{d_1} \times k_b$
- (3)  $id_z = k_z \times k_d \times k_{d_1} \times k_c$

Given a set  $C$  with 11 classes in the hierarchy and the set of the first 11 prime numbers  $P = \{2, 3, 5, 7, 11, 13, 17, 19, 23, 29, 31\}$ , we must assign each class  $V$  a key  $k_v \in P$  such that, the maximum of the set  $id_{leaf} = \{id_x, id_y, id_z\}$ , the set consisting of the ID numbers of all leaf nodes in  $C$ , is minimal. As we already know, prime numbers need not be unique for each class and can be reused in same circumstances.

## 3 A Co-simulation Framework

The goals of the co-simulation framework are to validate the main invariants of the G&S heuristics, improve its performance, and establish its applicability in mission critical systems. The co-simulation process in the framework (Figure 2) consists of three consecutive stages: verification, evaluation, and analysis. The verification phase is a straightforward application of model checking where an abstract description of the system's behavior is checked against a set of invariants. In the evaluation stage the simulation results from the probabilistic approach are contrasted to the outcome of the deterministic approach. The aim of the analysis stage is to closely examine the instances where the solutions yielded by the two implementations differ. We identify patterns among the inconsistent results that reveal the weaknesses of the probabilistic solution. The framework works by executing two independent models, the G&S model and the exhaustive search model. The first input component to the co-simulation framework (Figure 2) is an abstract model of the G&S fast dynamic casting heuristics, implemented in Promela (SPIN's input language) and the embedded C primitives it allows. The G&S abstract model is subsequently used to verify the main invariants of the G&S heuristics and at the same time provide us with a simulation testbed to examine the heuristics performance. The second component of the framework is the exhaustive search model that simply looks into all possible type ID assignments to discover the optimal solution for a given



**Figure 2.** A Co-Simulation Framework for G&S Improvement and Verification

class hierarchy. We employ SPIN’s search engine to perform the exhaustive search. In Algorithm 1 we present the pseudocode of our co-simulation approach. The following sections elaborate in more details on each of the stages of the framework.

---

**Algorithm 1** Pseudocode of the co-simulation approach.

---

```

1: const int MAX_NUMBER_TESTS
2: VERIFY :
3: repeat
4:   Formal Verification (G&S Model) → error report
5:   if (no errors) then
6:     goto EVALUATE
7:   else
8:     study counter example
9:     correct G&S
10: until TRUE
11: EVALUATE :
12: count = 0
13: for (count < MAX_NUMBER_TESTS) do
14:   Simulation(G&S Model) → G&S solution
15:   Enumeration(Exhaustive Search Model) → best solution
16:   if (G&S solution ≠ best solution) then
17:     add instance to SIS
18:   count ++
19: ANALYZE :
20: for all i ∈ SIS do
21:   look for a pattern
22:   modify G&S
23:   goto EVALUATE

```

---

### 3.1 Formal Verification

Every G&S implementation operates under the assumption that when a prime number is reused, it is assigned to non-conflicting classes. In addition, the maximum type ID must fit within the boundaries of a memory word. We check

these invariants during the program verification phase. Establishing the validity of the G&S invariants is done by straightforward application of model-checking with SPIN. In SPIN the critical system properties are expressed in the syntax of linear time logic. Based on the G&S abstract specification, the model-checker performs a systematic exploration of all possible states. In case of failure, SPIN provides a counterexample that demonstrates a behavior that has lead to an illegal state. In our model, the invariants are expressed as a *never claim* [6], and are checked just before and after the execution of every statement.

### 3.2 Evaluation


SPIN has been previously employed to implement solutions of scheduling [1] and discrete optimization [15] problems. The problem we face in the G&S heuristics is a combinatorial optimization problem [12]. Given a finite set  $I$ , a collection  $F$  of subsets of  $I$ , and a real-valued function  $w$  defined on  $I$ , a discrete optimization problem could be defined as the task of finding a member  $S$  of  $F$ , such that:  $\sum_{e \in S} w(e)$  is as small (or as large) as possible.

Except for the simplest cases, a discrete optimization problem is difficult because its design space is typically disjoint and nonconvex. Therefore, the optimization methods applied to continuous optimization problems cannot be utilized in this case. In a small discrete problem, it would be possible to exhaustively list all possible combinations. As the number of parameters increase, the state explosion makes optimizations difficult. The two general strategies for approaching a discrete optimization problem can be classified as *deterministic* and *probabilistic*. What we do for the G&S exploration in SPIN could be described as applying a deterministic approach for the evaluation of a set of proposed probabilistic methods. The Branch and Bound method [12] guarantees the discovery of the global optimum in the cases when the problem is linear or convex and is the most frequently used discrete optimization method. It is based on the sequential analysis of the discrete tree of each parameter. The branches that can be estimated to reach invalid or unfeasible solutions are consequently eliminated. This simple optimization could also be applied in some limited cases in the SPIN’s Fast Dynamic Casting exhaustive search. Let us explore a class hierarchy with three classes  $A$ ,  $B$ , and  $C$ , where  $B$  is derived from  $A$ , and  $C$  is derived from both  $A$  and  $B$ . In this case, we have  $C = \{A, B, C\}$ ,  $P = \{2, 3, 5\}$ , and  $id_{leaf} = \{id_c\}$ . The enumeration is given in Table 1. We assume that the computation starts at a state  $S_0$  where all three keys  $k_a$ ,  $k_b$ , and  $k_c$  are uninitialized. Then we assign possible values from the set  $P$  to the key variables of the classes  $A$ ,  $B$ , and  $C$ . The enumeration shown above can be expressed as the computation shown on Figure 3.

$id_c = k_c \times k_b \times (k_a)^2$	$k_a$	$k_b$	$k_c$
60	2	3	5
60	2	5	3
90	3	2	5
90	3	5	2
150	5	2	3
150	5	3	2

**Table 1.** Enumeration of all solutions

The graph shows only the valid states of the computation.



**Figure 3.** Exhaustive search computation

There are a number of invalid states that are not shown on the graph. For example, according to the rules defined in G&S, it is possible to reuse some of the prime numbers in  $P$ . Thus, we can try and add an edge  $k_b = 2$  in state  $S_1$ , however the reuse of 2 in this case is invalid since  $A$  and  $B$  are conflicting classes.

The illustrated automation in Figure 3 provides a foundation for the construction of a Promela model for the deterministic solution. Each possible prime number assignment to a given class key is represented by a separate state transition in the exhaustive search model. SPIN initiates the optimum search at state  $S_0$  and visits all possible states. At each end state the value of the minimum of the set of leaves, in this case represented only by  $id_c$ , is computed and compared to the current minimum. This approach is similar to the algorithm described by Ruys in [15] and shown in Algorithm 2. For such an application, we use the model checker in a somewhat unusual fashion. In this scenario, the validation property checks whether the value of  $id_c$  is greater than the current minimum. Each time this condition is violated, the current minimum is updated and the process is automatically repeated until SPIN confirms that there are no routes violating the specification. Since the solution is deterministic, it is guaranteed to discover the global optimum for type ID assignment. The performance of the G&S heuristics is measured by running a simulation of the G&S model that has been used earlier for verification. Now we are left with only one important task (not automated at this stage), the comparison of the results from the probabilistic and deterministic solutions. Once we identify a set of incon-

---

**Algorithm 2** Finding the global minimum in the state space.

---

- 1: *input* : Promela model  $M$
  - 2: *output* : the optimal minimum for the problem  $M$
  - 3:  $min =$  (worst case) maximum value for  $id$
  - 4: **repeat**
  - 5:     use SPIN to check  $M$  with condition ( $id_c > min$ )
  - 6:     **if** (error found) **then**
  - 7:          $min = id_c$
  - 8: **until** (error found)
- 

sistent results, we try to find a pattern and refine the G&S heuristics. Then the refined scheme is implemented in the probabilistic model and the evaluation process is reiterated.

### 3.3 Analysis

The simulation and enumeration models are continuously executed until, if possible, a set of instances with inconsistent solutions can be identified. Thus, each instance in the Set of Inconsistent Solutions ( $SIS$ ), represents a given class hierarchy for which the deterministic and probabilistic approach have discovered different solutions. The class hierarchies for each test could be guided or created in a random fashion. For the generation of the test data in our experiments we implemented a pseudo random class hierarchy generation algorithm, in a manner similar to the TGFF (Task-Graphs-For-Free) method as described in [2]. We look for patterns among the collected hierarchies in  $SIS$  and seek clues that can lead us to improvements of the G&S scheme. Potential improvements are tested by adding them to the G&S model and evaluate their effect. Such scheme modifications are carefully selected since it is possible that they might enhance a given G&S feature and at the same time weaken another. Ideally, the improvements lead to a heuristic scheme that provides the best solutions for a larger number of the test hierarchies and at the same time has a time complexity equal to or less than the earlier heuristic scheme.

Despite the numerous advanced state space reduction techniques utilized by the SPIN model checker, little can be done to further optimize the exhaustive search. The main goal of our experiments is to reach quick and effective optimization of the G&S scheme, thus the class hierarchies considered were not the largest and most complex that our models can handle. The models developed for our experiments are capable of handling class hierarchies of double or triple the size of the ones presented in the paper, and even larger number of classes can be facilitated with increased computational power. In the framework, the exhaustive search is used to identify flaws in the G&S type ID assignment scheme, thus, there is no need to create and simulate much larger hierarchies. In this work our goal is to demonstrate that the current size of the class hierarchies is sufficient to discover significant flaws of the original heuristics.

## 4 Application in Mission-Critical Software

Modern space mission systems have evolved from simple embedded devices into complex computing platforms with high autonomy and an exceptionally large demand for human-computer interaction. Consequently, such systems require reliable and flexible data systems managing the collection, storage, and transportation of data. The Mission Data System(MDS) is the Jet Propulsion Laboratory's state-and goal-oriented framework for building embedded control systems with a high degree of autonomy. MDS provides the building blocks for the implementation of embedded platforms based on the concepts of state estimation and control. The Data Management Services(DMS) is the MDS component responsible for the production, storage, processing, and transfer of control and scientific data. In [21] Wagner defines the challenges of data management in MDS as the problems of producing and storing data and converting the data to various formats as needed by its consumers. In addition, DMS needs to ensure the secure and lossless transport of the data with limited resources and through unreliable physical medium. To design and relate the data system entities, DMS employs concepts from high-level ISO C++ including templates, object-oriented class encapsulation, and dynamic casting necessary for the conversion of the data formats.

The actual telemetry data objects in MDS communicate with each other via byte streams produced by the transport protocol (e.g. spacecraft to ground communication). The receiver of the telemetry data needs to recreate the data object from the byte stream and thus invoke type casting in numerous occasions. Constant-time dynamic cast is also needed by the MDS Goal Network in the case when a controller or estimator [21] passes a goal via the Coordinating Goal Network(CGN), typically a large dynamic data structure. In CGN the goal is propagated using only its abstract attributes(start and end time, and the associated state variable). The achiever object who eventually picks up the goal needs to reconstruct the data object via dynamic down-casting to the specific type that conveys the state-specific achievement criteria. The application of the common compiler implementation of dynamic cast has proved to be unacceptable due to poor performance and the lack of the timing guarantees.

The G&S scheme was devised as a solution to a real industrial problem related to C++ use for hard real time code. Inquiries in the C++ community revealed that the problem was fundamental and common, rather than isolated: developers simulate dynamic casting with other language features, leading to type-unsafe special-purpose code or the avoidance of best object-oriented practices. Naturally, such workaround code slows down development, complicates maintenance, and increases the need for testing.

## 5 Results

We applied the co-simulation process described in the previous section to a large number of class hierarchies. The tested hierarchies are not built into our models. Instead, we have followed a pseudo random generation methodology similar to TGFF task graph generation as described in [2] to automatically generate hundreds of possible test cases. For illustration, we show the results from a set of seven pseudo random class hierarchies (Appendix A). The results of the G&S heuristics model and the exhaustive search are shown in Table 2. A brief comparison of the results indicates that the G&S heuristics do not give optimal performance for class hierarchies with non-virtual multiple-inheritance. A closer look at the algorithm reveals that the priority calculation routine takes into account only the number of descendants that each class has. Let us consider the class hierarchy from test case 7. We notice that according to the current scheme, the base classes 0, 1, and 2 all get the same priority rank since they all share the descendant 6. Class 6 is at the lowest level of the hierarchy and has the largest number of ancestors. If we would like to optimize the heuristics, we must find a way to increase the priority of base class 2. Our reasoning is derived from the fact that Class 2 is ambiguous and the leaf Class 6 contains two copies of it. Similarly, let us have a closer look at test case 1. In the optimal solution, Class 5 takes the lower prime number (11) compared to Class 4, despite the fact that its only descendant has less ancestors compared to Class 4. The reason for this result is the fact that the derived Class 3 contains two ambiguous bases while Class 4 contains only one ambiguous base. As a result of our analysis we conclude that higher priority should be given to derived classes and their ancestors who contain more ambiguous base classes. To fix these weaknesses, we extend the G&S heuristics by adding two simple rules:

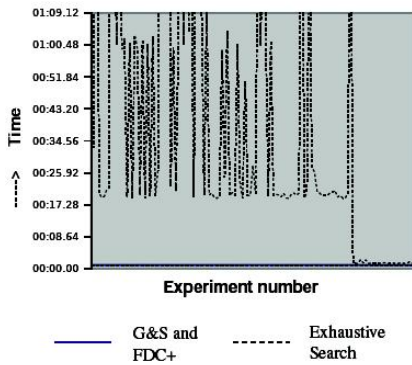
- (1) We count every ambiguous ancestor twice when we determine the number of ancestors to each class
- (2) For each base class, we count the number of derived classes that include more than one copy of it, and add that number directly to its priority

We call this enhanced G&S heuristics Fast Dynamic Casting Plus(FDC+). As Table 2 shows, for the initial set of test cases, FDC+ performance is 100% equivalent to the performance of the deterministic approach. In the performed tests, we have generated 127 pseudo random class hierarchies and applied G&S, FDC+, and the exhaustive search to each one of them. The experimental results showed that FDC+ was able to yield the best type ID assignment in 85% of the class hierarchies compared to 48% for the G&S heuristics. The time performance of the three schemes is shown in Figure

Case No	G&S	Exhaustive search	FDC+
Case 1 (keys)	(2, 3, 5, 7, 11, 13, 17)	(3, 2, 5, 7, 13, 11, 17)	(3, 2, 5, 7, 13, 11, 17)
Case 1 ( <i>i</i> ds of all leaves)	(16380, 16830)	(13860, 13260)	(13860, 13260)
Case 2 (keys)	(2, 13, 3, 5, 17, 7, 11)	(2, 13, 3, 5, 17, 7, 11)	(2, 13, 3, 5, 17, 7, 11)
Case 2 ( <i>i</i> ds of all leaves)	(1326, 2310)	(1326, 2310)	(1326, 2310)
Case 3 (keys)	(2, 3, 13, 5, 7, 17, 11)	(2, 3, 13, 5, 7, 17, 11)	(2, 3, 13, 5, 7, 17, 11)
Case 3 ( <i>i</i> ds of all leaves)	(26, 51, 2310)	(26, 51, 2310)	(26, 51, 2310)
Case 4 (keys)	(2, 3, 5, 7, 11, 13, 17)	(2, 3, 5, 7, 11, 13, 17)	(2, 3, 5, 7, 11, 13, 17)
Case 4 ( <i>i</i> ds of all leaves)	(2310, 1547)	(2310, 1547)	(2310, 1547)
Case 5 (keys)	(2, 3, 5, 7, 11, 7, 11)	(2, 3, 5, 7, 11, 7, 11)	(2, 3, 5, 7, 11, 7, 11)
Case 5 ( <i>i</i> ds of all leaves)	(42, 66, 70, 110)	(42, 66, 70, 110)	(42, 66, 70, 110)
Case 6 (keys)	(2, 3, 5, 11, 13, 7, 17)	(2, 3, 5, 11, 13, 7, 17)	(2, 3, 5, 11, 13, 7, 17)
Case 6 ( <i>i</i> ds of all leaves)	(66, 78, 420, 170)	(66, 78, 420, 170)	(66, 78, 420, 170)
Case 7 (keys)	(2, 3, 5, 7, 11, 13, 17)	(3, 5, 2, 7, 11, 13, 17)	(3, 5, 2, 7, 11, 13, 17)
Case 7 ( <i>i</i> ds of all leaves)	(2552550)	(1021020)	(1021020)

**Table 2.** Co-simulation of the seven cases from Appendix A

4. While the time performances of the G&S and FDC+ algorithms are equal and both run in a very low constant-time (the function at 00:01 min on Figure 4), logically the time performance of the exhaustive search increases exponentially with the increase of the number of classes nodes in a given class hierarchy. The analysis of the test results indi-



**Figure 4.** Search time for type ID assignment

cated that FDC+ finds a better type ID compared to the G&S approach in 39% of the test scenarios. For the greater part of the test cases, FDC+ matched the optimal type ID assignment computed by the exhaustive search. This efficiency boost is due to the optimized performance of FDC+ in the cases where multiple non-virtual inheritance is present in the class hierarchy. We have also observed that the implementation of these optimizations does not lead to efficiency loss in other scenarios and the performance of FDC+ is always at least as good as the performance of G&S. Our experimental results have indicated that the introduced optimizations in FDC+ have fixed a weakness of the original G&S approach and have improved the success rate in finding the best type ID assignment. The G&S scheme requires a key of a memory size that is a function of the size and shape of a class hierarchy. Thus, the improved heuristics almost double the size of class hierarchies that can be han-

dled by a given key size. Since the scheme gets significantly slower when a key gets too large for a machine word, the improvements to the heuristics address the main limitation of the G&S scheme.

## 6 Conclusion

In this work we applied co-simulation of the deterministic and probabilistic solutions to the combinatorial optimization problem posed by the G&S type ID assignment scheme. Our framework proved successful in verifying and refining the existing G&S heuristics. We demonstrated how we use the simulation results to devise improvements to the G&S algorithm and evaluate them. The results from our experiments indicate that the improved G&S heuristics (FDC+) provide the optimal type ID assignment in 85% of the class hierarchies, compared to 48% for the regular G&S algorithm. The efficiency of the type ID assignment scheme has significant importance for the performance of the fast dynamic casting by Gibbs and Stroustrup [3]. This paper presented a practical approach of how to discover improvements to the type ID assignment scheme in a simple and effective manner. The main advantage of the presented approach is the ease and simplicity of the discovery and test for potential improvements. The improved heuristics that we have described in this work almost doubles the size of class hierarchies that can be handled by a given key size. A more extensive simulation might suggest further improvements to the type ID assignment scheme. Our main goal in this work has been to demonstrate how an algorithm optimization problem encountered has been successfully automated and moreover that its automation has led us to quick but significant improvements of the initial scheme. In the future we intend to utilize a static analysis tool for automatic class hierarchy analysis and extraction.

## 7 Acknowledgements

We thank Peter Pirkelbauer, Kirk Reinholtz, Gerard Holzmann, and the anonymous referees for their helpful suggestions.

## A Appendix




Figure 5. Test Cases 1 and 2




Figure 6. Test Cases 3 and 4




Figure 7. Test Cases 5 and 6




Figure 8. Test Case 7

## References

- [1] E. Brinksma and A. Mader. Verification and Optimization of a PLC Control Schedule. In *Proceedings of the 7th International SPIN Workshop on SPIN Model Checking and Software Verification*, pages 73–92, London, UK, 2000. Springer-Verlag.
- [2] R. P. Dick, D. L. Rhodes, and W. Wolf. TGFF: task graphs for free. In *CODES/CASHE '98: Proceedings of the 6th international workshop on Hardware/software codesign*, pages 97–101, Washington, DC, USA, 1998. IEEE Computer Society.
- [3] M. Gibbs and B. Stroustrup. Fast dynamic casting. *Softw. Pract. Exper.*, 36(2):139–156, 2006.
- [4] R. Gluck and G. Holzmann. Using spin model checker for flight software verification. In *Proceedings of the 2002 IEEE Aerospace Conference*, 2002.
- [5] L. Goldthwaite. Technical Report on C++ Performance. In *ISO/IEC PDTR 18015*, February 2006.
- [6] G. Holzmann. *The Spin Model Checker, Primer and Reference Manual*. Addison-Wesley, Reading, Massachusetts, 2003.
- [7] M. Ingham, R. Rasmussen, M. Bennett, and A. Moncada. Engineering Complex Embedded Systems with State Analysis and the Mission Data System. In *Proceedings of First AIAA Intelligent Systems Technical Conference 2004*, 2004.
- [8] ISO/IEC 14882 International Standard. *Programming languages C++*. American National Standards Institute, September 1998.
- [9] S. Lerner, T. Millstein, and C. Chambers. Automatically proving the correctness of compiler optimizations. In *PLDI '03: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation*, pages 220–231, New York, NY, USA, 2003. ACM Press.
- [10] M. R. Lowry. Software Construction and Analysis Tools for Future Space Missions. In J.-P. Katoen and P. Stevens, editors, *TACAS*, volume 2280 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2002.
- [11] M. Musuvathi and D. R. Engler. Model checking large network protocol implementations. In *NSDI'04: Proceedings of the 1st conference on Symposium on Networked Systems Design and Implementation*, pages 12–12, Berkeley, CA, USA, 2004. USENIX Association.



- [12] G. L. Nemhauser and L. A. Wolsey. *Integer and combinatorial optimization*. Wiley-Interscience, New York, NY, USA, 1988.
- [13] C. Perrow. *Normal Accidents*. Princeton University Press, September 1999.
- [14] RTCA. *Software Considerations in Airborne Systems and Equipment Certification (DO-178B)*, 1992.
- [15] T. C. Ruys. Optimal scheduling using branch and bound with spin 4.0. In T. Ball and S. K. Rajamani, editors, *Model Checking Software, Proceedings of the 10th International SPIN Workshop*, volume 2648 of *Lecture notes in Computer Science*, pages 1–17, Berlin, 2003. Springer Verlag.
- [16] J. Schumann and W. Visser. *Autonomy Software: V&V Challenges and Characteristics*. In *In Proceedings of the 2006 IEEE Aerospace Conference*, 2006.
- [17] A. Stoica, D. Keymeulen, A. Csaszar, Q. Gan, T. Hidalgo, J. Moore, J. Newton, S. Sandoval, and J. Xu. Humanoids for lunar and planetary surface operations. In *In Proceedings of the 2005 IEEE International Conference on Systems, Man and Cybernetics*, October 2005.
- [18] B. Stroustrup. *The C++ Programming Language*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2000.
- [19] B. Stroustrup. Abstraction and the c++ machine model. In Z. Wu, C. Chen, M. Guo, and J. Bu, editors, *ICISS*, volume 3605 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 2004.
- [20] R. Volpe. Rover Technology Development and Mission Infusion Beyond Mars Exploration Rover. In *IEEE Aerospace Conference*, March 2005.
- [21] D. Wagner. Data Management in the Mission Data System. In *In Proceedings of the IEEE System, Man, and Cybernetics Conference*, 2005.