

Performance Evaluation of A Testing Framework Using QuickCheck and Hadoop

YUSUKE WADA^{1,a)} SHIGERU KUSAKABE²

Received: June 14, 2011, Accepted: November 7, 2011

Abstract: Formal methods are mathematically-based techniques for specifying, developing and verifying a component or system for increasing the confidence regarding the reliability and robustness of the target. It can be used at different levels with different techniques, and one approach is to use model-oriented formal languages such as VDM languages in writing specifications. During model development, we can test executable specifications in VDM-SL and VDM++. In a lightweight formal approach, we test formal specifications to increase our confidence as we do in implementing software code with conventional programming languages. For this purpose, millions of tests may be conducted in developing highly reliable mission-critical software in a lightweight formal approach. In this paper, we introduce our approach to supporting a large volume of testing for executable formal specifications using Hadoop, an implementation of the MapReduce programming model. We are able to automatically distribute an interpretation of specifications in VDM languages by using Hadoop. We also apply a property-based data-driven testing tool, QuickCheck, over MapReduce so that specifications can be checked with thousands of tests that would be infeasible to write by hand, often uncovering subtle corner cases that wouldn't be found otherwise. We observed effect to coverage and evaluated scalability in testing large amounts of data for executable specifications in our approaches.

Keywords: cloud computing, formal methods, property-based test, MapReduce

1. Introduction

Formal methods are mathematically-based techniques for the specification, development and verification of the target systems. Performing appropriate mathematical analysis of methods is effective in increasing the confidence regarding to the reliability and robustness of a design of the target system. We can choose a specific technique from various formal methods, such as model checking techniques.

Instead of proving specifications, we test executable specifications to increase our confidence in the specifications as we do in implementing software systems with conventional programming languages. While the specific level of rigor depends on the aim of the project, millions of tests may be conducted in developing highly reliable mission-critical software. For example, in an industrial project using VDM++, a model-oriented formal specification language [1], they developed formal specifications of 100,000 steps including test cases (about 60,000 steps) and comments written in natural language, and they carried out about 7,000 black-box tests and 100 million random tests [2].

Random testing is one of the techniques useful in testing the executable specification in a lightweight way. Random testing is useful to discover corner cases which may not be found by manually generated test cases. In particular, increasing the number of test cases is useful, while it is not realistic to increase the num-

ber of test cases exhaustively. We can execute practical random testing by specifying a property, its input data type, and the range and number of test data.

The coverage information, which represents the degree to which the source code has been tested, can be collected in executing specifications. We can get higher confidence in executable specifications if we have higher test coverage for the specifications. We can expect a higher coverage rate when we can increase the number of test cases in a lightweight formal approach.

In this paper, we discuss our approach to testing executable formal specifications, whose straightforward execution is rather expensive, for a large volume of test data in an elastic way. We try to automatically distribute thousands of test runs of executable specifications in VDM languages over elastic computing resources by using Hadoop, an implementation of the MapReduce programming model. From a practical point of view in large scale testing, we need a testing tool to generate test data, which can provide higher confidence by effectively finding subtle corner cases. We choose QuickCheck [3], an open source testing framework for Haskell. QuickCheck also has a language of testable specifications by which testers can define expected properties of the target under test. QuickCheck employs a simple method to generate test data, random testing, and has a language by which users can control test data generation and generate test data in languages other than Haskell. This type of testing tool encourages a high-level approach to testing in which we can check whether abstract properties are satisfied universally with an arbitrary number of test data generated. We can distribute the generation of large amount of test data, as well as test-runs of executable specification for

¹ Graduate School of Information Science and Electrical Engineering, Kyushu University, Fukuoka 819-0395, Japan

² Faculty of Information Science and Electrical Engineering, Kyushu University, Fukuoka 819-0395, Japan

^{a)} y.wada@ale.csce.kyushu-u.ac.jp

a large amount of test data, and collect coverage information as well as test results from the distributed environment while observing scalable performance.

The rest of this paper is organized as follows. We briefly explain the VDM languages in Section 2. In Section 3, we discuss several issues to reduce the cost of large amounts of testing of VDM formal specification using emerging cloud technology. We outline our framework in Section 4 and evaluate our framework in Section 5. Finally we conclude this paper in Section 6.

2. VDM: Vienna Development Method

VDM (The Vienna Development Method) is one of the model-based formal methods, a collection of techniques for developing computer systems from formally expressed models. While VDM was originally developed in the middle of the 1970s at IBM in Vienna, its support tools, VDMTools, are currently maintained by CSK Corporation in Japan. In order to allow machine-supported analysis, models have to be formulated in a well-defined notation. The formal specification language, VDM-SL, has been used in VDM, which became the ISO standard language (ISO/IEC 13817-1) in 1996, and VDM++ is its object-oriented extension version. VDMTools provide functionality for dynamic checking of formal specifications in the formal specification languages VDM-SL and VDM++ [1], such as the interpretation of executable specifications, in addition to static checks such as syntax checking and type checking of formal specifications. By using the interpreter of VDMTools, we can test executable specifications in VDM-SL and VDM++ to increase our confidence in the specifications as we do in implementing software systems with conventional programming languages.

In this paper, we especially focus on the step in which we validate the specification using systematic testing and rapid prototyping. In light-weight formal methods, we do not rely on highly rigorous means such as theorem proofs, and we use testing of executable specifications in order to increase confidence in our specifications. While the specific level of rigor depends on the aim of the project, thousands of tests may be conducted in developing highly reliable mission-critical software. When we consider performance, it is time-consuming to execute the specification for a large amount of test data, and the performance degradation seems accelerated as the number of tests increases in this case.

3. Large Amount of Testing for Executable Specification

Software testing plays an important role in gaining confidence for quality, robustness, and correctness of software. In this section, first we discuss software testing. Testing executable specifications shares the same issues as software testing, such as high cost and long running time.

3.1 Reducing the Cost of Testing

As the size and complexity of software increase, its test suite becomes larger and its execution time becomes a problem in software development. Several approaches have been used to reduce the cost of time consuming test phases. Selecting a representative subset of the existing test suite reduces the cost of testing [4], [5].

Some other papers describes priority-based approaches [5], [6].

Large software projects may have large test suites. There are industry reports showing that a complete regression test session of thousands of lines of software could take weeks of continuous execution [7]. While each test is independent with each other, the very high level of parallelism provided by a computational grid can be used to speed up test execution [8]. Distributed tests over a set of machines aims at speeding up the test stage by simultaneously executing a test suite [9], [10]. A tool aims at executing software tests on a Grid by distributing the execution of JUnit [11] test suites over the Grid, without requiring modification in the application and hiding the grid complexity from the user [8].

VDMUnit is a framework for unit testing for VDM++ specifications, and it could be possible to distribute the execution of VDMUnit over a Grid like JUnit over GridUnit. Our approach in this paper uses MapReduce [12] rather than the Grid platform to perform a large amount of testing on an elastic cloud computing platform. In our approach, we will be able to automatically distribute the execution of testing specifications by using Hadoop over an elastic cloud computing platform. Using a cloud computing platform may also lower the cost of acquisition and maintenance cost of the test environment.

3.2 Elastic Platform

We consider an approach to leveraging the power of testing by using elastic cloud platforms to perform large scale testing. Increasing the number of tests can be effective in obtaining higher confidence, and increasing the number of machines can be effective in reducing the testing time. We believe the cloud computing paradigm has impact on the field of software engineering and consider an approach to leveraging light-weight formal methods by using cloud computing which has the following aspects [13]:

- (1) The illusion of infinite computing resources available on demand, thereby eliminating the need for cloud computing users to plan far ahead for provisioning;
- (2) The elimination of an up-front commitment by cloud users, thereby allowing organizations to start small and increase hardware resources only when there is an increase in their needs; and
- (3) The ability to pay for use of computing resources on a short-term basis as needed and release them as needed, thereby rewarding conservation by letting machines and storage go when they are no longer useful.

We can prepare a platform of arbitrary number of machines and desired configuration depending on the needs of the project.

3.3 Property-based Data-driven Testing

Since increasing the number of tests is effective in obtaining higher confidence, huge number of tests may be performed especially in developing mission-critical software. Conceptually, we can increase the number of test cases on elastic cloud computing platforms. However, generating test cases by hands can be a bottleneck in software development.

In this paper, we use a property-based testing tool QuickCheck [3], which supports a high-level approach to testing Haskell programs by automatically generating random




Fig. 1 Concept of map/reduce programming model.

input data. Property-based data-driven testing encourages a high level approach to testing in the form of abstract invariants functions should satisfy universally, with the actual test data. Code can be checked with thousands of tests that would be infeasible to write by hand, often uncovering subtle corner cases that would not be found otherwise. We try to automatically distribute the generation of test data for formal specification in addition to the execution of formal specification.

3.4 MapReduce

While we can prepare a platform of an arbitrary number of computing nodes and generate an arbitrary number of test cases, we need to reduce the cost of managing and administrating of the platform and runtime environment.

MapReduce programming model is proposed in order for processing and generating large data sets on a cluster of machines [12]. Input data-set is split into independent elements, and each mapper task processes the corresponding element in a parallel manner as shown in Fig. 1. Data elements are typically data chunks when processing huge volume of data. The outputs of the mappers are sorted and sent to the reducer tasks as their inputs. The combination of map/reduce phase has flexibility, thus, for example, we can align multiple map phases in front of a reduce phase.

MapReduce programs are automatically parallelized and executed on a large cluster of machines. The runtime system takes care of the details of partitioning the input data, scheduling the program's execution across a set of machines, handling machine failures, and managing the required inter-machine communication. Its implementation allows programmers to easily utilize the resources of a large distributed system without expert skills for parallel and distributed systems.

When using this map/reduce framework, input elements can be test cases, f can be an executable specification in VDM languages or actual code fragment under test, and output elements test results, contains test coverage of an executable specification.

4. Our Approach

Increasing the number of tests can be effective in obtaining higher confidence, and increasing the number of machines can be effective in reducing the testing time. Regarding the number of test cases, preparing an arbitrary large number of test cases by hand is possible but impractical. Among many tools

for testing, a property-based testing tool, QuickCheck, supports a high-level approach toward testing Haskell programs by automatically generating random input data as described later. Users of QuickCheck can customize their test case generation including the number of test cases. We modify QuickCheck to fit to our approach for testing formal specification with Hadoop. As formal specification languages, such as VDM-SL for example, share features with functional programming languages, we can obtain formal description necessary to use QuickCheck to generate test data in VDM-SL. There have been work focusing on their relations [14], [15].

We consider an approach to use QuickCheck on elastic cloud platforms. We can perform testing of arbitrary scale by exploiting such a combination. We can automatically distribute the generation of test data and the execution of tests in a scalable manner with Hadoop. We employ Hadoop framework to easily execute tests in a data-parallel way.

4.1 QuickCheck

QuickCheck is an automatic testing tool for Haskell programs. It defines a formal specification language to state properties. Properties are universally quantified over their arguments implicitly. The function `quickCheck` checks whether the properties hold for randomly generated test cases when they are passed as its arguments. QuickCheck has been widely used and inspired related studies [16], [17], [18].

For the explanation, we use a simple `qsort` example from a book [19].

```
qsort :: Ord a => [a] -> [a]
qsort [] = []
qsort (x:xs) = qsort lhs ++ [x] ++ qsort rhs
  where lhs = filter (< x) xs
        rhs = filter (>= x) xs
```

We use idempotency as an example invariant to check that the function obeys the basic rules a sort program should follow. Applying the function twice has the same result as applying it only once. This invariant can be encoded as a simple property. The QuickCheck convention in writing test properties is prefixing with `prop_` to distinguish them from normal code. This idempotency property is written as a following Haskell function. The function states equality that must hold for any input data that is sorted.

```
prop_idempotent xs = qsort (qsort xs) == qsort xs
```

QuickCheck generates input data for this `prop_idempotent` and passes it to the property via the `quickCheck` function. Following example shows the property holds for the 100 lists generated.

```
> quickCheck (prop_idempotent :: [Integer] -> Bool)
OK, passed 100 tests.
```

While the sort itself is polymorphic, we must specify a fixed type at which the property is to be tested. The type of the property itself determines which data generator is used. The `quickCheck` function checks whether the property is satisfied or not for all the test input data generated. QuickCheck has convenient features such as quantifiers, conditionals, and test data monitors. It provides an embedded language for specifying custom test data generators.




Fig. 2 Outline of our approach to property-based data-driven testing.

In QuickCheck, the function `generate` generates test data. The first argument specifies the range of test data, the second the random seed, and the third the property expression to be checked. An `Int` value is selected randomly from 0 to the first argument, and the `Int` value is used to generate a random value of the input type of the property.

We modified this function to generate test data in a distributed way. We divide the range from which test data is selected. If we use m tasks to generate N data, data range of the i -th task is specified, from $i[\frac{N}{m}]$ to $(i+1)[\frac{N}{m}] - 1$ ($0 \leq i \leq m$).

Conceptually, we can evaluate property expressions in property-based random testing in a data-parallel style by using MapReduce framework. Each mapper evaluates the property for one of the test data and reducer combines the results from mappers. By applying an automatic testing tool such as QuickCheck on MapReduce framework, we expect we can greatly reduce the cost of a large scale testing.

4.2 Implementation

We developed our testing environment by customizing QuickCheck and developing glues to connect components for testing executable specifications on Hadoop framework. In this section, we discuss implementation issues.

Figure 2 outlines our approach. Our approach to implementing property-based testing on Hadoop is to separate the testing process into two phases. We generate test data by using mappers, and store the data into a file in the first phase. Then we can read and split the file, and distribute the data to mappers, where the property function written in Haskell is evaluated.

Hadoop streaming: Hadoop, open source software written in Java, is a software framework implementing MapReduce programming model [20]. We write mapper and reducer functions in Java by default in this Hadoop framework. However, The Apache Hadoop distribution contains a utility, Hadoop streaming, which allows us to create and run jobs with any executable or script as the mapper and/or the reducer. The utility will create a mapper/reducer job, submit the job to an appropriate cluster, and monitor the progress of the job until it completes. When an executable is specified for mappers, each mapper task will launch the executable as a separate process when the mapper is initialized. When an executable is specified for reducers, each reducer task will launch the executable as a separate process when the reducer is initialized. We used this Hadoop streaming for implementing our testing framework.

Figure 3 shows an overview of our property-based data-driven testing. The overview of property-based data-driven testing is as follows: First, we generate test data according to the specified

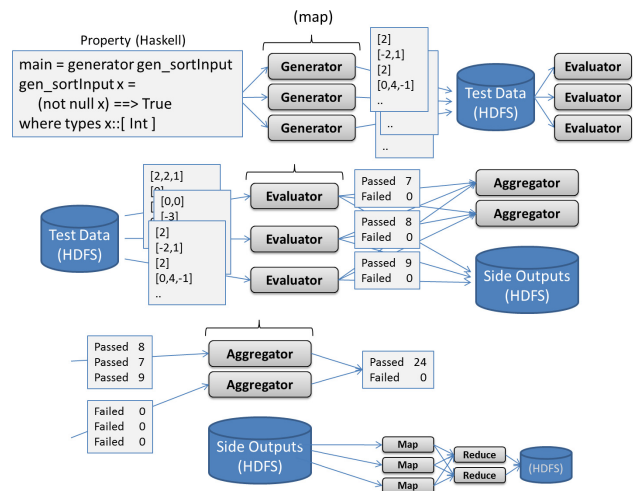


Fig. 3 Progress of property-based data-driven testing.

property. Next, we store the generated data in a file on HDFS. Finally, in the evaluation phase, we pass the test data to mappers through the standard input, and the mappers output the results to the standard output.

Distribution of test data generation: We generate the specified number of random test data with mappers in Hadoop framework in a distributed way. However, naively splitting the number and assigning the sub-numbers to mappers lead to useless computing due to overlap of test data generated among different mappers. We need to avoid increasing the number of redundant test data from the view point of efficiency and coverage. We modified the generator in QuickCheck to avoid this problem. We add one extra parameter to check function in QuickCheck module. The parameter represents the start index of test data and is passed to test function. After the total number of tests is determined, each mapper is given different start index according to the number of tests assigned to mappers.

5. Performance Evaluation


In order to examine the effectiveness of our approach, we measured performance in testing specification of the Enigma cipher machine (Enigma) written in VDM++ in the book [1]. Enigma is basically a typewriter composed of three parts: the keyboard to enter the plain text, the encryption device and a display to show the cipher text. Both the keyboard and the display consist of 26 elements, one for each letter in the alphabet.

5.1 Distribution of Test Data Generation

To investigate the effectiveness of our distribution of test generation, we compare the number of unique (non-redundant) test data in generating 80,000 data for `Int` and `Float` type. **Table 1** is the result. “Original” means using the QuickCheck original data generator on non-Hadoop environment. “Naive” means each data generator has common start index in generating random test data on mappers in Hadoop environment. “Ours” means each data generator knows its own starting index, each of which is different from each other. According to the result, we can see effectiveness of our modification as we see no outstanding difference between “Original” and “Ours” while we have some degradation in

Table 1 The ratio of unique data and standard deviation in generating random data of Int and Float.

		Original	Naive	Ours
Int	Unique(%)	34.5	8.3	34.4
	Std.Dev.	7,754.1	966.3	7,656.2
Float	Unique(%)	99.6	95.7	99.6
	Std.Dev.	7,684.5	962.0	7,729.8

**Fig. 4** Test generation time.

“Naive.”

We can expect the similar effect for other data types, including compound data types. For example, in generating 80,000 tuple data of two Ints, we can generate 79,730 (99.66%) unique test data in “Ours,” 79,742 (99.68%) unique test data in “Original”, and 75,950 (94.93%) unique test data in “Naive”. “Ours” generates almost equal amount of unique test data with QuickCheck original one.

We examined the effectiveness of our distributed test data generation in increasing the number of Mapper tasks to generate 100,000 test data on our eight-slave Hadoop system. Test data is an list of Ints. We show the relation between the number of tasks on the slaves and the generation time shown in **Fig. 4**. As we can see from the figure, the time to generate test data was shortened by increasing the number of tasks until we reach the saturation point. We conclude our framework can generate almost the same unique test data as the original of QuickCheck, and also our framework is effective in gaining speedup on Hadoop systems.


5.2 Coverage

Generally speaking, higher test coverage rate leads to higher confidence in developing software. This also applies to our approach. We rely on testing, not formal proof, in gaining confidence in a formal specification. We expect higher coverage rate when we increase the number of test cases in our property-based approach.

We examine the effectiveness of our approach from the view point of coverage in this section. VDMTools can report coverage of executable specifications, and we can get coverage data in our approach as described later. We can observe coverage rate while we change the number of test cases specified as the parameter value of QuickCheck.

We measured coverage rate in our approach through the following steps:

- (1) We used VDM++ executable specification files for Enigma in the book [1].

**Fig. 5** Coverages of four configuration of patterns in ten trial.**Table 2** Configuration of the platform in performance evaluation.

	NameNode	JobTracker	Slave
CPU	Xeon E5420 *1	Xeon E5420 *1	Xeon X3320 *2
Memory	3.2 GB	8.0 GB	3.2 GB
Disk	2 TB	1 TB	140 GB

Table 3 Elapsed time in tests of Enigma specification.

Nodes/Tests	100	200	300	400	500	600	800	1,000
Single	21.1	67.4	158.1	334.5	492.1	672.4	1,392.3	2,285.5
1	232.9	230.1	272.4	267.0	246.2	262.0	369.9	434.4
2	163.8	160.4	179.3	183.2	168.1	160.1	233.5	285.7
3	126.4	120.2	123.1	127.3	123.9	126.2	157.2	211.1
4	102.9	104.1	104.4	106.1	110.0	114.0	135.1	170.7
5	91.2	90.2	93.9	95.6	99.0	100.0	119.7	145.5
6	86.4	84.7	88.5	88.6	93.5	94.4	110.9	136.0
7	90.3	84.8	85.2	84.4	86.9	91.5	104.4	127.4
8	83.8	78.4	79.4	79.9	83.0	83.9	99.2	117.6

- (2) We conducted property-based data-driven test. We generated test cases by using customized QuickCheck and executed specifications with VDMTools through command line interface.

- (3) We gathered coverage data of VDM++ classes while changing the number of input test data.

We ran our property-based test ten times for each configuration. We changed the number of test cases as one, ten, one hundred and one thousand. **Figure 5** shows the results for a VDM++ file Rotor.vpp, one of the Enigma components, in this experiment. As we can see from the figure, the larger number of test cases we use, the higher coverage we have, and the specification file was covered completely when we use one thousand of test cases.

However, it took long time to execute the specification files for one thousand of test cases. Next, we examine the effectiveness of parallel and distributed execution of formal specifications using Hadoop.

5.3 Speedup in Property-based Testing

The configuration of the platform is shown in **Table 2**. We show the result of elapsed time in **Table 3** and in **Fig. 6**. As we can see from the results, the elapsed time of Hadoop version became shorter when the number of tests was four hundreds and over. Since Hadoop framework is designed for large scale data processing, we have no advantage in elapsed time for small set of test data. The computation node has four processor cores, and we

*1 Intel(R) Xeon(R) CPU E5410 @ 2.50 GHz Quad Core

*2 Intel(R) Xeon(R) CPU X3320 @ 2.50 GHz Quad Core




Fig. 6 Elapsed time in increasing the number of tests of VDM++ Enigma specification on the various number of nodes.




Fig. 7 Speedup ratio in increasing the number of tests of VDM++ Enigma specification on the various number of nodes. Please note each node has a quad-core processor.

can achieve speedup even on a single node as Hadoop can exploit thread-level parallelism on multi-core platforms. Figure 7 shows the scalability with the changing number of nodes. The speedup ratio is calculated against the result of the sequential execution for the tests on a single node. We can acquire many advantages on our platform, so that the ratio is high. As we see in Fig. 7, the increase of the number of slave machines is effective in reducing testing time.

6. Concluding Remarks

In this paper, we explained our approach to testing executable formal specifications in lightweight formal method framework using VDM languages. In order to increase confidence in the specification, we increase the number of test cases with a property-based data-driven approach on a cloud computing oriented programming framework. We apply a property-based data-driven testing tool, QuickCheck, so that specification can be checked with hundreds of tests that would be hard to write by hand. We investigated coverage of executable VDM++ model. However, large amount of test data gains long executing time. Our framework can deal with this problem. We observed scalable performance in conducting large amount of testing for executable specifications. Therefore, we can both of acquiring high coverage because of large amount of test data and tuning our platform so that execute the test in time.

One of our future work is to investigate more detailed performance breakdown to achieve more efficient environment. We will

also try to improve usability of our framework. VDMTools include test coverage printout tool. We will extend our framework to exploit this tool in a parallel and distributed way to inform detail coverage information, such as a VDM++ statement covered but another is not, for more usability.

Reference

- [1] Larsen, P.G., Mukherjee, P., Plat, N., Verhoef, M. and Fitzgerald, J.: Validated Designs For Object-oriented Systems, Springer Verlag (1998).
- [2] Kurita, T., Chiba, M. and Nakatsugawa, Y.: Application of a formal specification language in the development of the mobile felica IC chip firmware for embedding in mobile phone, *FM 2008: FORMAL METHODS*, pp.425–429 (2008).
- [3] Claessen, K. and Hughes, J.: Quickcheck: A lightweight tool for random testing of haskell programs, *ACM SIGPLAN Notices*, Vol.35, No.9, pp.268–279 (2000).
- [4] Graves, T.L., Harrold, M.J., Kim, J.-M., Porter, A. and Rothermel, G.: An empirical study of regression test selection techniques, *ACM Trans. Softw. Eng. Methodology*, Vol.10, No.2, pp.184–208 (2001).
- [5] Wong, W., Horgan, J.R., London, S. and Agrawal, H.: A study of effective regression testing in practice, *Proc. 8th International Symposium on Software Reliability Engineering* (1997).
- [6] Kim, J.-M. and Porter, A.: A history-based test prioritization technique for regression testing in resource constrained environments, *Proc. 24th International Conference on Software Engineering* (2002).
- [7] Elbaum, S., Malishevsky, A.G. and Rothermel, G.: Prioritizing test cases for regression testing, *Proc. International Symposium on Software Testing and Analysis*, pp.102–112, ACM Press (2000).
- [8] Duade, A., Cirne, W., Brasileiro, F. and Macado, P.: Gridunit: Software testing on the grid, *Proc. 28th ACM/IEEE International Conference on Software Engineering*, Vol.28, p.779, ACM (2006).
- [9] Kapfhammer, G.M.: Automatically and transparently distributing the execution of regression test suites, *Proc. 18th International Conference on Testing Computer Software* (2001).
- [10] Hughes, D., Greenwood, P. and Coulson, G.: A framework for testing distributed systems, *Proc. 4th IEEE International Conference on Peer-to-Peer computing (P2P'04)* (2004).
- [11] Gamma, E. and Beck, K.: Junit: A cook's tour, *Java Report*, Vol.4, No.5, pp.27–38 (May 1999).
- [12] Dean, J. and Ghemawat, S.: MapReduce: Simplified data processing on large clusters, *Comm. ACM*, Vol.51, No.1, pp.107–113 (Jan. 2008).
- [13] Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R.H., Konwinski, A., Lee, G., Patterson, D.A., Rabkin, A., Stoica, I. and Zaharia, M.: Above the clouds: A berkeley view of cloud computing, Technical Report, UCB/EECS-2009-28, Reliable Adaptive Distributed Systems Laboratory (Feb. 2009).
- [14] Borba, P. and Meira, S.: From vdm specifications to functional prototypes, *J. Syst. Softw.*, Vol.21, No.3, pp.267–278 (June 1993).
- [15] Visser, J., Oliveira, J.N., Barbosa, L.S., Ferreira, J.F. and Mendes, A.S.: Camila revival: VDM meets haskell, *1st Overture Workshop* (2005).
- [16] Arts, T., Hughes, J., Johansson, J. and Wiger, U.: Testing telecoms software with quviq quickcheck, *ERLANG'06: Proc. 2006 ACM SIGPLAN Workshop on Erlang*, pp.2–10, New York, NY, USA (2006).
- [17] Boberg, J.: Early fault detection with model-based testing, *Erlang Workshop*, pp.9–20 (2008).
- [18] Claessen, K., Palka, M., Smallbone, N., Hughes, J., Svensson, H., Arts, T. and Wiger, U.: Finding race conditions in erlang with quickcheck and pulse, *ICFP*, Vol.35, No.9, pp.268–279 (2009).
- [19] O'Sullivan, B., Goerzen, J. and Stewart, D.: Real World Haskell, *Oreilly & Associates Inc* (2008).
- [20] Apache: Hadoop, available from (<http://hadoop.apache.org/>) (accessed 2012-01-12).



Yusuke Wada is Graduate student at Kyushu University since April 2010. Research: Efficient software development. Efficiency of software development using disciplined software and formal methods. Speed-up and energy saving in multi-threading and cloud computing. Education: Studies at Graduate School of Information Science and Electrical Engineering, including Social Information Systems Engineering Course.



Shigeru Kusakabe is Associate Professor at Kyushu University since October 1998. Research: Efficiency of software development using disciplined software process, formal methods and applied behavioral analysis. Speed-up and energy saving in multi-threading and cloud computing. Education: Lectures at the Department of Advanced Information Technology in the Graduate School of Information Science and Electrical Engineering, including Social Information Systems Engineering Course. Social: SEA, IPSJ, IEICE, ACM, IEEE-CS, J-ABA.