

Assertion Definition Language Project

Research Results and Future Plans [†]

James deRaeve¹
j.deraeve@xopen.co.uk

Shane P. McCarron^{1,‡}
s.mccarron@xopen.co.uk

Bill Gogesch^{2,‡}
bill.gogesch@eng.sun.com

Roger Hayes^{2,‡}
roger.hayes@eng.sun.com

¹X/Open Company Ltd.
²Sun Microsystems Laboratories

This paper presents the findings of a four year project conducting research in the area of a formal interface behavior description grammar and its translation into tests and documentation for the interfaces described. In this paper the research team presents the results of early experimental usage of the research, important findings on the ability to translate formal grammars into natural languages such as English and Japanese, and plans for the future deployment of technology developed in conjunction with this research project.

1. Introduction

Since 1992 the ADL Project team has been working closely with the Information-technology Project Agency to further the goals of Open Systems through the research of Open Systems Testing Tools. The first major phase of this research is now nearly complete. The results of this research, freely available under an X Window System-like copyright and grant of rights, have been very promising. The purpose of this paper is to convey some of the key findings of this project. This paper presents:

- The basic architecture of the ADL system,
- The application of this system to two different examples of testing problems and the results therefrom,
- The applicability of this research to the generation of natural language specifications, and
- The plans of the research team for the future extension and application of this research.

1.1 History

The Assertion Definition Language project arose out of the Open System community's desire to have tools that would speed the development of tests for interfaces defined by open interface specifications¹. The development of these

tests has historically been slow and costly. X/Open, a leader in the development of these "conformance tests", recognized that the rapid expansion of the Open Systems community and resulting standards would soon far outdistance the ability of the industry to test such standards. X/Open felt that it was theoretically possible to describe the behavior of these standards in such a way that they could be automatically tested.

At about the same time IPA announced their Open Fundamental Software Technology Project and invited X/Open and UNIX International to submit a project proposal. These organizations determined that the most important thing they could do for the industry jointly via this project was to sponsor research in the area of automated testing. This proposal was ultimately accepted by IPA, and the ADL Project was officially under way.

The first step of this project was to develop a set of comprehensive requirements for the research. In order to do this, X/Open and UNIX International invited a group of prominent conformance testing experts from across the Open Systems industry to participate. The resulting requirements document formed the basis of the research.

Following this work, X/Open and UNIX International solicited from the industry existing research that was towards the requirements, with the intent that they would partner with an existing research group and speed the project. This solicitation resulted in a number of potential starting points, each of which was carefully evaluated.

[†] This research is conducted under the Open Fundamental Software Technology Project of the Information-technology Promotion Agency, Japan(IPA).

[‡] Researcher, ADL Project, Open Fundamental Software Technology Project

1. Open interfaces are ones controlled by an independent third party and subject to public review, as opposed to proprietary interfaces controlled by a single organization.

The research project selected is original, patented research by Sun Microsystem's Laboratories (SML). This original project, named PrimaVera, was targeted at automating the testing of interfaces described through a syntax much like that of the Object Management Group's (OMG) Interface Definition Language (IDL).

Since 1992, the ADL Project has worked to refine deliver this basic research and deliver it in a form useful to the Open Systems community. The specification language has been extended and ported from IDL to ANSI C, natural language generation has been added, and test generation has been refined and extended.

1.2 Project Status

The project is currently in the middle of the final year of its first major phase. Earlier progress, including delivery of a fully functional and portable translation system, has been reported in previous papers[1],[2]. In this year the focus of the project has shifted to refining the reporting mechanisms and the natural language specification generation and ensuring the research is fully internationalized and localizable. Specifically, this final year the project is conducting research and development in the following areas:

- Extension of Makefile semantics to ease the development of TET-driven test suites
- Better error processing and reporting during ADL interface specification translation
- Improved reporting from provide, relinquish, and auxiliary functions
- Improved coupling of reporting data between grid points and assertions
- Enhanced portability of ADL source and ADL-generated source
- Enhanced translation of ADL interface specifications to natural language and test specifications
- Localization of translation system messages to Japanese
- Development of Prolog engine for generating Japanese natural language specifications and test specifications from ADL input
- Translation of basic ADL Environment documentation to Japanese
- Development of a guide to assist in future localization package development

In addition, the team has continued to extend the basic research by extending the specification languages.

Each of these research areas have already been started. The most recent release (0.9) reflects a nearly final version of

ADL, incorporating all enhancements planned for this first project phase. It is currently being beta tested by a variety of organizations to ensure that it is of sufficient quality to be deemed complete.

The next step is to distribute a final version of ADL. This is scheduled for 15 December. At that time all reported bugs are expected to be addressed, all planned features will be complete, and the Japanese Localization Package will be synchronized. This version will be the first "official" release of the technology resulting from the research, and will be called ADL Release 1.0.

While all this is going on, a parallel phase to evaluate the Japanese Localization Package is being conducted. This phase of the project, being performed by Kanri Kogaku Ltd., is doing the following:

- Carefully reviewing the ADL early releases this year, including the Japanese Localization Package
- Study the POSIX.1 File System I/O interfaces and develop a testing methodology document
- Translate the existing specification into ADL (including a simple Natural Language Dictionary)
- Develop a number of auxiliary functions to support the ADL interface specification, using the testing methodology document as a guide
- Evaluate the generated test suites using a selected product which implements the POSIX.1 File System interfaces

This phase of the project is slightly out of sync with the main research to ensure that there are no dependencies that would cause delays. As a result, the final feedback from this phase will not be available until February 1996. This feedback will be incorporated into the project's final report to IPA.

2. Basic ADL Architecture

The first implementation of the ADL Translation System is designed to generate tests for any open system platform compliant with X/Open's Portability Guide version 4 (XPG4) — most modern UNIX^{®1}-like systems fit this description. The generated test programs are ANSI C. The System itself is written in C++ and ANSI C and will build and execute on an XPG4 compliant platform with a resident C++ compiler. The System builds well using the Free Software Foundation's gcc/g++ compiler. The documentation for the System was developed using Frame Technology Corporation's FrameMaker version 4.0. It is

1. UNIX is a registered trademark in the United States and other countries licensed exclusively by X/Open.

provided in both source and PostScript™ forms in the distribution. In addition to these requirements, a complete ADL Translation System platform may also have the Test Environment Toolkit installed as a test controller (see below).

The current ADL System distribution archive (release version 0.9) contains separate releases for both the ADL system software and the ADL documentation set. The complete system occupies approximately 30 megabytes of disk space. The System builds successfully and with no customization on many platforms, including generic System

V Release 4.0 platforms with gcc version 2.6.0 or above installed.

2.1 Architectural Model

The ADL Translation System architecture is made up of ten basic components. This section provides an overview of how the various components relate to one another from a very high level. The components of the ADL Translation System are best illustrated through the diagram in Figure 1. The elements of the diagram are described below.

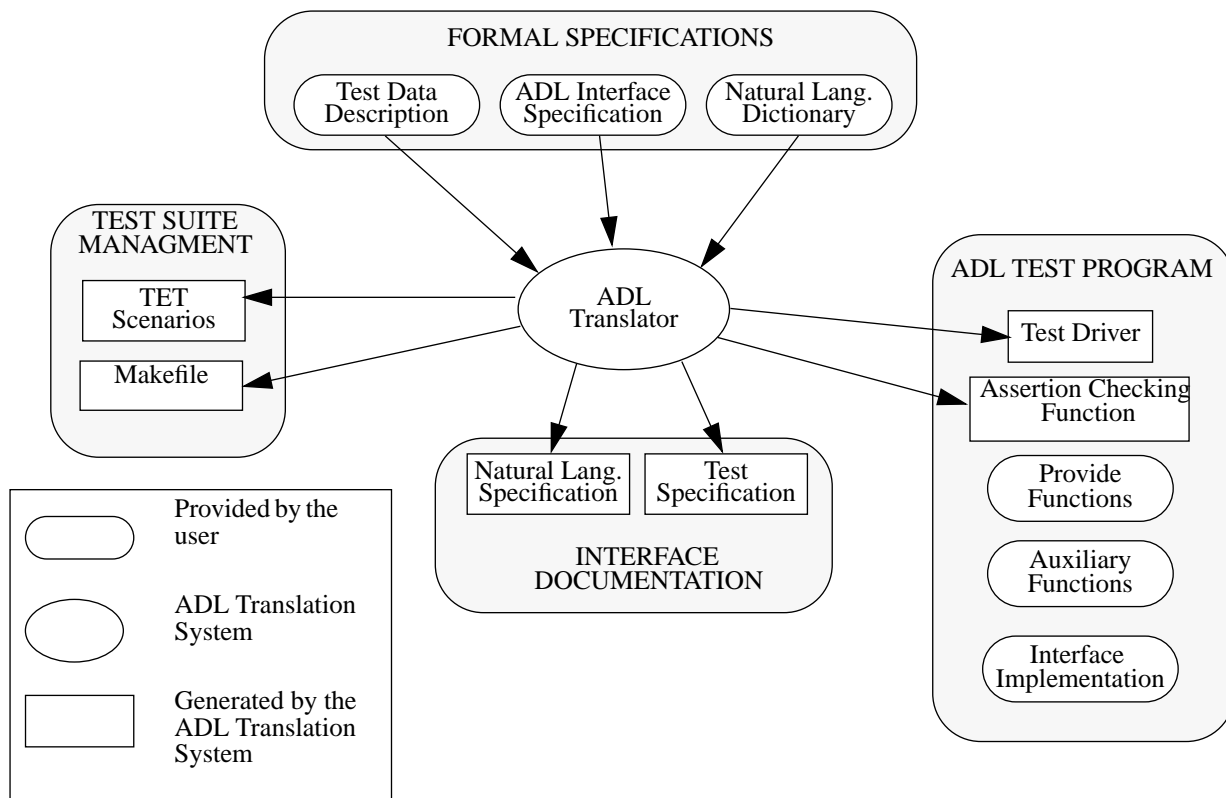


FIGURE 1 - Abstract ADL Architecture

Formal Specifications

Inputs to ADL are in the form of formal specification files. The ADL Translation System defines three specification languages; Assertion Definition Language (ADL), Test Data Description (TDD) language and the Natural Language Dictionary (NLD).

ADL Interface Specification

The ADL Interface specification is the file (or files) in which the signature and behavioral semantics of the interfaces are defined.

Test Data Description Specification

The Test Data Description (TDD) specification is a file (or files) that define the test data and specify the test execution.

Natural Language Dictionary

The Natural Language Dictionary (NLD) is a reference file (or files) that defines the natural language translation for elements of the ADL specification. This file is optional. If it is not provided, the ADL Translator will perform a rudimentary default translation.

Interface Documentation

One of the features of the ADL Translation System is the automated generation of interface documentation. The generated documentation is a natural language translation of the information in the ADL Interface Specification and the TDD specification.

Test Specification

A Test Specification is an ADL Translator-generated file that describes each function under test, the assertions for those functions, and the ways in which

each function will be exercised during a test case run. This file is generated in a natural language based upon information in the ADL Interface Specification, the Test Data Description, and the Natural Language Dictionary.

Natural Language Specification

A Natural Language Specification (NLS) is an ADL Translator-generated file that describes the syntax and semantics of the functions described in the ADL Interface Specification, based upon information in the ADL Interface Specification and the Natural Language Dictionary.

ADL Test Program

The ADL test program comprises the source code for an executable test. The ADL Translator generates appropriate source code based upon the ADL Interface Specification and the data definitions and test directives found in the TDD. This generated source code may be compiled and linked with user-written support code and with the ADLT library to create a test program. Each test program tests one function from the interface; several test programs can be created from one specification.

Test Driver

The ADL Translator generates source code for the Test Driver based on a test specification from the TDD file. The Test Driver is responsible for iterating over the requested list of test cases and managing the test execution.

Assertion Checking Function

The Assertion Checking Function is emitted source code that invokes an implementation of the specified interface and checks the results of invocation against the semantic assertions specified in the ADL specification.

Auxiliary Functions

Auxiliary functions are user-provided functions that assist the ADL-generated Assertion Checking Function in determining the veracity of assertions in the ADL Interface Specification. Auxiliary functions are referenced by using them in assertion statements.

Provide Functions

The provide functions are user-written functions that generate appropriate parameter values given the parameter properties defined in the TDD.

Interface Implementation

The Interface Implementation is the system or function being specified and tested by the ADL Program.

Test Suite Management

The ADL Translator can generate two files to aid in building and executing ADL Test Programs: a TET scenario file and a Makefile.

TET Scenarios

A TET Scenario is a file (or files) that describes the sequence of tests to be built, executed, and/or cleaned up by TET's Test Case Controller. The ADL Translator generates a TET Scenario file that describes a suite of tests comprised of all test directives in the TDD, as well as individual entries for each test directive, in either summary or detailed testing mode (see Section 2.2).

Makefile

The ADL Translator generates a Makefile to build the ADL Test Program. Once the Makefile is generated, the user can use the make facility to create the executable test program from the source (the specification, support code, and implementation). The Makefile will run the ADLT translator, C compiler, and linker as required. New versions of the ADL Test Program are easily compiled and linked with the generated Makefile.

2.2 ADL and the Test Environment Toolkit

The Test Environment Toolkit, or TET, is an Open Systems industry standard testing harness designed in 1989 by the Open Software Foundation, UNIX International, and X/Open. This free-ware harness provides a standard user interface for test users, and APIs in a number of programming languages for test developers. Through TET, the designers hoped to unify the open systems testing industry with the goal of allowing plug-and-play interoperability of test suites and test users. Within that small community, this goal has been largely achieved.

When X/Open first started the ADL Project, the intent was to develop a tool that would work in conjunction with TET to automate the generation of TET-based test suites. After discussion with the team at Sun Microsystems Laboratories, the design goals of the project were changed to allow this interaction with TET, but not to require it. When ADL generated tests are used with TET, the interaction between TET and ADL generated tests suites is seamless. The ADL Translation System has the ability to generate complete test suites that can be built, executed, and cleaned-up by the TET test case controller.

In order to achieve this seamless integration, the System generates a TET scenario file for the generated test cases. A scenario file describes the sequence of tests to be executed, and may have subsets of tests defined within it. ADL generates a scenario file that describes a complete suite of tests for each test directive in the TDD.

Specifically, the ADL Translation System allows the generation of test cases that use the low-level TET primitives for reporting results and interim data. In addition,

the System provides two modes of TET-style result reporting that are aligned with IEEE Std. 1003.3-1991 (the POSIX standard for assertion-based test methods). These modes allow the test case to record the results of assertion evaluation in a manner consistent with the IEEE "POSIX" standards of reporting a single result per assertion. This can be done either in summary mode (where the union of all evaluations of an assertion are considered the result for that assertion), or in detail mode (where each assertion evaluation reports a result).

Note that the ADL Translation System does not require the use of TET. Instead, it permits test engineers who are designing complete suites of interface tests to use this harness to simplify and automate the building, execution, and cleaning-up of test runs. This is particularly useful in system test environments, where a complete implementation of a set of interfaces must be tested prior to production release. It is also useful for conformance testing, where certifiable passing of a test suite is necessary to ensure conformance to industry standards.

3. Experimental Results

During the last year, the project conducted two experiments, applying an early version of the system to testing two very different styles of interfaces. The resulting tests and documentation were made available to the relevant communities. Members of those communities worked closely with the ADL Project to assist us in evaluating the value of the ADL-generated tests.

These experiments were conducted with an early version of the ADLT Translation system, so that the development team would have time to act on the feedback. It should be noted that the team has attempted to correct the specific faults, and to adapt the design to accommodate the desired mode of use. For example, the documentation has been completely re-written since the experiments, the test reporting has been completely revised to be more suitable for POSIX testing, and test suite creation has been greatly improved.

3.1 Common Object Request Broker Test Suite

The Object Management Group has developed a series of specifications relating to the server and client sides of communication with distributed objects. Central to these is the Common Object Request Broker Architecture specification (at the time of the experiment at version 1.2). This document describes the ways in which object interfaces are specified, how clients find and use objects, and (to a lesser extent) how the objects interface with the object system. The ADL Project Steering Board felt that this specification was a good candidate for testing because it was an emerging standard, require the use of external objects in the tests, and there were implementations available against which the tests could be applied.

The development team for this prototype was Applied Testing and Technology of Los Gatos, California. ApTest is an organization with substantial experience in conformance

test development, and was involved in the original requirements generation for the ADL Project. This section of the paper is a summary of their report at the end of the experiment[4].

ADL as a Test Development Tool

In ApTest's experience with ADL its performance in this area was not in itself an improvement.

Developer productivity in producing the CORBA Test Suite was lower than would be expected from manual development. In part this is due to a very steep learning curve for ADL. However, after this initial learning phase was overcome, productivity remained at below that for manual techniques.

It should be noted that very experienced software developers were involved in this project, more so than has traditionally been considered necessary for test development. ApTest expects an additional dip in productivity in using ADL with less seasoned personnel.

In measuring productivity ApTest considered the time required to complete test development for an interface manually and with ADL. This does not take into account the number of tests produced. This is much higher in the ADL case because of the technique of running a large number of different data points through each test method.

This increase in tests produced is certainly a benefit of ADL in terms of system stress and performance testing. While the large number of times ADL invokes an interface-under-test is also in some sense an increase in test coverage, ApTest found this was not necessarily a meaningful increase. It would only be meaningful if the test data set is carefully constructed.

In the CORBA test suite ApTest did not see an increase in the ability of the ADL generated tests to detect errors versus manually generated ones, despite an orders of magnitude increase in the number of tests. The point here is one made in the ADL documentation: test coverage with ADL is a matter of skilled design, not a necessarily a consequence of the technology. Thus while ADL does generate a lot of test executions, an increase in meaningful test coverage is not an inherent by-product.

In terms of test quality, ApTest again did not see a significant gain. CORBA Test Suite error rates were comparable to what is seen with test suites in general.

In the abstract, higher quality might be achieved through the generation of less code by manual, and thus error prone, means. However, ApTest found the amount of code that had to be written using ADL to be only marginally less than what would be expected without it. Though ADL automates many mundane tasks in sequencing and executing tests, the body of the tests

themselves must still be manually coded. As these aspects of the tests are much more error prone than those ADL automates, and by far the bulk of code necessary in ADL's early releases, ADL's impact on initial test error rates is limited.

In addition, ApTest found added complexity in debugging tests generated with ADL due to the nature of the code generated and the complex flow of control involved. This has a negative impact on quality. It is also an additional consideration in the higher level of experience necessary for a programmer working with ADL as discussed above.

Overall when compared with traditional development methods ApTest found ADL-based test generation to exhibit:

- roughly comparable test error rates
- roughly comparable meaningful test coverage
- lower programmer productivity, and
- higher requirements for the experience level of developers.

It should be noted that while ApTest believes improvements can be made in ADL that bear on these issues, they do not consider expectations for dramatic improvements in terms of programmer productivity, initial quality or test coverage through use of automated test development methods to be realistic.

Given the complexity of the task of software development and its historic resistance to automation (e.g. Artificial Intelligence research in Automatic Programming), a reasonable goal for ADL is that test development can be done within the envelope of performance that can be expected with manual means. As initial development is a relatively small factor in the life-cycle cost of test suites, this limitation of the paradigm can be handsomely compensated for through improving specification quality and test suite maintenance.

In sum, ApTest feels that ADL entails a net increase in initial test development cost as compared to traditional methods but that this increase is sufficiently small as to be mitigated by its benefits in other areas.

ADL as a Specification Tool

The lack of compatibility between implementations based on written specifications is most often traceable to the very poor quality of the specifications themselves. These documents are often vague, imprecise and incomplete, leading to varied and incompatible implementations. From a testing perspective these same issues limit the ability to verify implementation compliance.

Improving the quality of specifications by expressing them in a formal language is the area where ADL can realistically be expected to make to the greatest contribution to the community, and was in fact the most significant positive in ApTest's experience with it.

In capturing the CORBA 1.2 specification in ADL many (over 100) areas of imprecision and ambiguity were found in the original specification. As the specification was considered complete before this project, this demonstrates a clear benefit of ADL to specification writers and test developers, one which it can only be wished was available when the original CORBA specification was drafted!

In terms of developing test specifications, ADL fell short of expectations. The then current ADL language primitives were limited as compared to the semantics of most programming languages. As a result many aspects of tests which should have been exposed in the test specification had to be coded inside ADL provide or auxiliary functions, and are thus were visible within the specification. This issue was noted by many of reviewers of the CORBA Test Suite specification who complained that the test methods used were opaque from reading the specification.

In sum, ApTest feels ADL is very valuable as a specification development tool, though in need of improvement in natural language output and the richness of the ADL language.

ADL as a Maintenance Tool

As on-going maintenance is the largest cost-driver in software costs, and quick turn around for new versions of tests in response to specification changes is a strong requirement, ADL's performance in this area is an important consideration.

Conclusions regarding maintenance are preliminary, as developing the CORBA Test Suite has not provided hands on experience with test suite maintenance using ADL. However the process of developing it can provide insights into how this would work.

ADL has promise in this area through allowing changes to tests solely through modifications of the ADL version of the specification and facilitating easy reuse of existing code.

In practice, realizing these benefits was made difficult by the limitations of the ADL language discussed above. Since ADL specifications must currently be at a very high level, some changes will not automatically update the test suite appropriately and may require changes to or the development of new auxiliary or provide functions.

As well, the requirements for sophisticated programmers in order to utilize ADL discussed above must be taken into consideration. ApTest expects higher

experience levels to be required for ADL test suite support than have been required in traditional situations.

In sum, ADL shows promise in this area, however extensions to the ADL language will be required before meaningful benefits can be expected. Pending the availability of empirical data, ApTest also expects the level of expertise required for ADL maintenance to be a mitigating factor versus gains in productivity.

Product Review

ApTest found the current ADL product to correspond well to its specification.

In addition to those functional issues discussed above, two important deficiencies in the then current functionality were encountered:

- Applicability of the product to building large test suites is limited. Such suites have more complex source tree structures and build procedures than what ADL accommodated.
- ADL's provision for generating test suites under the Test Environment Toolkit (TET) is inadequate. Though some improvements have been made in recent ADL releases it was still not possible to generate the test output or test results expected of TET-based test suites.

User Interface

The user interface to ADL is adequate. Limited provisions for making the interface user friendly or facilitating productivity in a development environment have been made. The lack of a GUI-based front end for test developers and the large number of separate files which have to be manipulated to create ADL test suites contribute to the high level of expertise necessary to use it.

Portability

ADL portability was a problem when the CORBA Test Suite project was initiated but was substantially improved in subsequent releases.

By the end of the experiment the only outstanding portability problems are in the Prolog interpreter. This is incorporated into the distributed package. This part of ADL continued to be difficult to port and also exhibited problems in 64-bit environments.

Quality

The development of the CORBA Test Suite uncovered a number of errors in the implementation of ADL. The majority of these errors were corrected during the experiment, and additional regression testing was added to the release process resulting in acceptable quality of the last release reviewed (ADL 0.7).

Performance

Performance of ADL has declined somewhat over the term of the CORBA Test Suite project. Though it remains at an acceptable level for a test suite of this size, performance with a large scale test suite has yet to be evaluated.

Documentation

The documentation of ADL was very poor - the User and Programmer Manuals were effectively unusable. This is partly due to their not having been maintained as ADL evolved and partly to their having been mistargeted in the first place and thus of limited benefit to test developers.

Conclusions and Recommendations

In sum the experience of the CORBA Test Suite project is that ADL offers a valuable improvement in the production of test suites by facilitating much greater accuracy and completeness in interface specifications.

In terms of test development, ADL adds some cost in the initial development cycle. While there is potential for benefits in the maintenance cycle it is not yet clear how well these can be realized.

As a product ADL is reasonably complete and functional and offers acceptable performance.

In terms of the evolution, ApTest offered a number of recommendations for addressing the then current deficiencies in ADL and also for increasing its ability to improve the testing process.

- The ADL language should be extended to increase the degree to which test methods which can be defined in ADL rather than provide or auxiliary functions.
- The ADL interface to TET should be reworked to allow results to be returned by user written code, journal output to be created by user written code, report results in TET (assertion-by-assertion) style rather than ADL (gridpoint-by-gridpoint style), and document assertions in natural language (rather than ADL).
- The interface specifications and test specifications produced by ADL need to be enhanced to be fully in troff format (this is simplistically the case for interface specifications now), allow the programmer to embed troff macros, format complex expressions in a readable fashion, and handle line and page formatting correctly.
- A substantial number of large test suites should be studied and ADL modified to accommodate existing practice in test suite structure and build procedures for such suites.

- The ADL documentation should be rewritten to focus on the experienced test suite developer who will use ADL for verifying real-world specifications.
- The necessity to create separate `adl`, `nld` and `tdl` files should be minimized and ideally eliminated or masked from the user. This complexity reduces productivity.
- The process of debugging ADL generated tests should be investigated with a focus on minimizing difficulty and providing tools that facilitate productivity.
- The specification generation and test development aspects of ADL should be further decoupled so that the ADL's benefits in specification work can be obtained by those who do not wish to use it for test development.
- The Prolog engine should be made reliably portable across Open Systems environments or replaced by something that is.
- The ADL user interface should be supplemented with friendly front end tools that allow less-sophisticated users to more easily utilize ADL.

3.2 Test Environment Toolkit Test Suite

The Test Environment Toolkit (TET) is a collection of tools and interfaces used by test developers. It was chosen for use in this experiment because it has a number of interesting properties that were expected to stretch ADL to its limits:

- TET is used by ADL itself, so the TET under test must NOT be the one that is in use by ADL.
- TET has both commands and interfaces, and both of these must be tested. Since ADL only supports exercising C language interfaces, the testing of the command must be encapsulated.
- TET's API is defined in a programming language-independent manner, and it has several programming language bindings (including C). All of these bindings must be tested, ideally by reusing the assertion definitions or having a single set of assertion definitions that could be applied to all language bindings.

The TET Test Suite experiment was conducted by Mount Bonnell, Inc. of Austin, Texas. MBI is an organization with substantial conformance testing experience and intimate knowledge of the implementation of TET. Their report on the results of their experiment is summarized in this section[5].

Obtaining and Building the ADL Release

The source code and documentation as well as binaries for some systems are available via ftp. ADL is written

largely in C++. Unfortunately the state of standardization and tools in the C++ is somewhat unstable. Fortunately ADL will build with recent versions of GCC. GCC source is available for most Open System platforms, although building it is time consuming.

Once a C++ environment is available, the next step is to build ADL,. Again this is fairly straight forward with the exception of the Prolog interpreter. The only supported configuration files shipped with Prolog were for SunOS 4.1 and Solaris 2.X. Other platforms might require some debugging.

Experiences Building ADL

During this experiment the ADL releases (0.5 and 0.6) were built on 3 different platforms:

- SunOS 4.1
- Linux 1.1.47
- NetBSD 1.0

The SunOS build went very well. At first there was some confusion about which version of the *make* utility to use. Many people using GCC tend to also use other GNU tools. Some earlier versions of ADL would only build with Sun's *make*. This problem was addressed in later versions.

The Linux build only presented a few easily fixed compile problems. Unfortunately, the ADL translator core dumped when run. It appeared that there was a malloc and free problem in this version of Linux and after a few days of debugging the frees were commented out and the code worked.

Attempting to debug the ADL source code revealed an important issue. Some of the C++ code that makes up ADL is generated using a parser-generation tool called *pcct*. In these early releases the source used by *pcct* was not part of the distribution, so debugging required working on the generated code. Working on generated code is notoriously difficult, and doing it in ADL was no exception. This was addressed in subsequent releases - the sources used by *pcct* to generate the portions of ADL are now included in the distribution.

The NetBSD build also went very well except for building the Prolog interpreter. Fortunately, NetBSD is binary compatible with SunOS 4.1 and it was possible to use the distributed binaries on this platform. This worked mostly correctly on the NetBSD release.

All in all ADL appears to be very portable for a C++ program (assuming you have a robust C++ environment). The Prolog part is designed to be portable and as more system configuration files are added it will become easier to build. The one concern in this area is that some organizations may be nervous using a "free" test generator built with a "free"

compiler. There must be a commitment to further support of the tool if it is to be successful.

Assertion Writing Using ADL

The first step to creating a test with ADL is creating the assertions. The team started by creating a set of natural language assertions from the TET documentation. The ADL interface was then written from these. The team felt that this is how ADL will be used to write tests for a specification that is already in use. One of the features of ADL is that it can create the natural language specification directly from the ADL interface specification. Since there was already a specification the team did not use this feature. They did, however, use the generated test specification as a sanity check for the formal assertions against the human-generated ones. This was extremely useful.

The temptation in assertion writing is to just jump right in and do it. Unfortunately this will probably lead to rewriting many ADL interface specification files. The problem here is that in order to get many of the benefits of ADL you need to reuse your auxiliary functions. In order to do this you need to have an architecture in mind for the specification that can frequently reuse the same auxiliary functions. Without a clear architecture in mind, the user may find it difficult to get the ADL interface specification right the first time through. This becomes an iterative process that eventually results in a good set of reusable auxiliary functions, but it is time consuming.

Another problem encountered while writing assertions was the C-like syntax of ADL. In the evaluated versions of ADL, the language did not yet permit the use of all the C language constructs. Since ADL is so C-like, this was very confusing. As a result of this input, ADL was extended to accept most C language constructs with the same semantics they have in C. This change should help alleviate the problem.

Creating Test Data

In order to create a working test the user creates a TDD file. This file contains the data points and test directives. The user must also write the provide and auxiliary functions. The team chose to create a TDD file with a single data point and use that while developing the auxiliary functions.

Creating the data files and the provide functions was the most convenient aspect of ADL. The ability to import (now called "use") other TDD files enables the test developer to define widely used test objects and incorporate them into a number of tests very easily. This feature saved a considerable amount of time in the TET experiment. Further, the ability to change the data set without modifying the test case is a significant advantage of ADL. It allows the user to create a test

suite and then easily come back and make the tests more thorough by adding data points.

The one problem area here is there was no way for a provide function to indicate an error. This ability was added for the 0.9 release, and should help make both provide and auxiliary functions more robust.

Developing Auxiliary Functions

The writing of the auxiliary functions did not go nearly as well as the TDD/provide function writing in this experiment. Ideally auxiliary functions should be written once and used in multiple assertions testing multiple interfaces. Unfortunately, the nature of TET make it very difficult to craft these "generic" auxiliary functions. This is partly due to the fact that the TET tests are run in a subprocess and the auxiliary functions must use inter-process communication facilities in order to find out the state of the subprocess. This is a much more complicated model than most tests would use (and one of the reasons TET was chosen as a good experiment). It did, however, exaggerate a weakness in ADL – that of writing generic auxiliary functions. Without a comprehensive test suite architecture, it is nearly impossible to craft complicated reusable auxiliary functions.

Testing and Maintenance

The challenge with any tool that automates a process is in gaining the trust of the users of that tool. In the case of test suites, the users frequently feel the need to inspect the test code to ensure that it is working right, since they firmly believe that their implementation could not possibly be at fault. Inspecting the code generated by ADL is nearly impossible. It is convoluted, and the generated identifier names are such that they are very hard to follow. While it was never the intent of the ADL Project team that the generated code be inspected, this does raise an interesting point. How can ADL gain the confidence of the test users?

Clearly, the answer is that it cannot do this over night. Test users must be able to inspect the code. However, just as a gcc user trusts it to generate correct assembly language, a test user also need to trust ADL to assemble the pieces of an interface test correctly. Those pieces must be available to the test user for inspection, to be sure. And users can inspect the generated source code if they like. But in the end, ADL, like any compiler must be trusted by its users.

Lessons Learned

From the experiment results reported above, it became clear that a variety of enhancements were required to the basic research. Most of those enhancements are now complete. Some of the most significant changes made as a result of the experiments are described in this list.

- The ADL language was extended to accommodate most C Language constructs. These extensions have greatly improved the ability of ADL users to fully express the behavior of interfaces.
- The syntax and semantics of the Natural Language Dictionary have been enhanced to make resulting translations more expressive.
- It is now possible to include NLD information directly in the ADL Interface Specification instead of in a separate file.
- The NLD translation engine has been changed so that the resulting output is structured more like a traditional manual page and is more readable.
- ADL has been extended overall to accommodate the construction and maintenance of large test suites. This includes the ability to organize a test suite hierarchically, to use libraries of auxiliary functions more readily, and to modify the behavior of the generated Makefiles without actually editing the makefiles themselves.
- The reporting mechanisms have been significantly enhanced. As a result, they are now able to generate reports that more closely resemble the way in which traditional assertion oriented testing is done. They also interact better with the Test Environment Toolkit. Finally, they can report a breakdown of assertion results so that it is possible to easily discern what parts of an assertion caused a failure.
- The Programmer and User documentation has been completely rewritten — targeted at experienced programmers rather than novices.

4. Natural Language Translation

Documentation is an important aspect of software development. Accurate documentation is the product of good communication between development staff and documentation staff. ADLT natural language services help bridge the gap between developer and technical writer by rendering an accurate prose interpretation of the formal ADL specification at every stage of product development. In consequence, the technical writing staff will have a sound basis for documentation at the time of specification. This break from tradition is not only advantageous to the writing staff, but to the development staff, who can base their work on the ADLT-generated documentation.

Using ADLT moves two often neglected aspects of software development to the fore: testing and documentation.

4.1 Objectives

The challenge presented to the document generation component of ADLT is to render the formal assertions of an ADL specification into prose documents. Specifications written in ADL are first order logic assertions on ADL data objects - they are comprehensible mainly to a technical audience. The documents that ADLT generates should be useful to a broader audience.

While translation between natural languages is a task made difficult by context-dependent semantics and the consequent reliance on a knowledge-base and dictionaries for semantic disambiguation, ADL semantics are well-defined and context-free. Translation of ADL assertions into natural language is thus a much simpler task.

ADL is a specification language. It is used to describe the behavior of software systems implemented in some target programming language. The three elements of the ADL assertions that undergo translation are ADL-defined operators, user-defined functions, and data objects.

ADL semantic specifications are made up of expressions, which in turn are made up of operators, data objects, and functions. The semantics of ADL operators are fixed: ADL users cannot rename or redefine the semantics of operators. ADL can be used to define data objects of various types. These data objects may be represented by user-defined identifiers, or by constant expressions. ADL also allows the user to define functions, whose semantics are defined by the user.

Converting ADL assertions to prose requires translating the semantics of each element into a natural language phrase. Our strategy is to combine those phrases in a context-independent manner; this is the principle simplification allowed by the simple structure of the ADL source of the translation. The three element classes described above require slightly different translation strategies.

The unambiguous semantics of ADL-defined operators makes them readily expressible in natural language. Descriptive prose can be substituted directly for each operator. That prose will generally be a verb phrase.

There is generally meaning associated with ADL data objects and functions that is not available from an ADL specification. The meaning intended by an identifier in ADL, the semantics of that data object, may or may not be reflected by the type or the name of the identifier. In translation, the meaning of an identifier is more likely the interesting aspect, that is to say, that the fact that an identifier refers to the balance of a bank account is the idea that must be captured in translation, not the fact that it is an integer. Like most programming languages, ADL relies upon the choice of meaningful names for data objects and functions in order to convey the intent of the specification writer, but there is always a trade-off between descriptiveness and convenience in the choice of a name. Therefore, the default translation of an identifier should be

its name, but it should be possible to substitute a more meaningful phrase.

In the case of data objects, noun phrases are used to convey meaning in translation. Functions are more complex. They are combinations of verb phrases for the function, describing what it does, and noun phrases describing how the function parameters participate in the actions of the function.

ADL assertions comprise combinations of data objects and operators, a mix of elements that are easily translated to prose and elements whose meanings cannot be found in an ADL specification. Useful translations of assertions almost always require some sort of exposition on an object or function in translation to prose.

There is a clear need for the expression of the intended meaning behind ADL objects and user-defined operators in order to produce good natural-language translation.

Summarizing, the problem facing the ADLT translation system is that of transforming the formal specifications of ADL into prose. The translation process requires information not necessarily available from the specification in order to produce useful prose. Furthermore, the system must be sufficiently flexible to produce prose for any locale.

4.2 Design Decisions

This section describes the design decisions made to address the natural-language translation challenges outlined in the previous section.

ADL partitions the task of natural-language translation into three categories:

- ADL-defined operators
- function identifiers
- data object identifiers

Because the semantics of ADL operators are well-defined, translation to a target natural language becomes the task of mapping an operator onto a prose expression of the operator's semantics.

The ADL semantics defined for objects are restricted to type information, which is of limited value in a natural-language translation. In fact, type information is a programming language detail that natural-language documentation seeks to abstract.

We designed ADLT to consult the Natural Language Dictionary (NLD) for the description of identifiers in an ADL specification. The NLD is the place to record the mapping of object identifiers and function identifiers into natural-language descriptions of those identifiers. It is a mechanism for explicating the full meaning of data objects and functions.

ADL objects and functions have scope defined by their declaration and use in an ADL specification. The translation

system must observe the scope of a particular identifier and apply an NLD translation appropriate to that scope, and the NLD must be able to express the mapping of an identifier in the context of its ADL scope.

The syntax of the NLD is designed to be simple: expressions relating identifiers to their prose descriptions. There are no semantic constraints or grammatical constraints on the descriptions. That choice makes the NLD very easy to write. However, the unconstrained nature of the NLD creates problems for the natural-language translation process. How can the system ensure that NLD prose of arbitrary complexity is grammatically correct, much less sensible?

This introduced complexity is addressed by the current ADLT in a very direct and simple way. We use the three classes of mapping, ADL-defined operator, function and object, in a context-free text substitution process. Comprehension of user-supplied NLD prose by the translation system is thus obviated. This choice means that there is no mechanism to control the introduction of senseless mappings in the NLD. All NLD mappings that pass syntactic analysis will be propagated to the natural-language documents faithfully. NLD authors must craft their mappings with care.

Separation of components by purpose is a guiding principal during this phase of our research. ADL specifications describe the expected behavior of system under consideration, but are distinct from TDD descriptions of the test input data. Both ADL and TDD are separate from system implementation source. Likewise, the NLD is a separate entity that serves to clarify the translation of ADL assertions. As such, the NLD is not a required input. ADLT will produce documents containing default translations in the absence of an NLD. Adding a well-crafted NLD, however, results in better, more descriptive documents.

Another guiding principal of the ADL project is that partial specification must be possible. A system can be minimally specified at first, then, as development progresses, specifications can be extended to render better tests and better documents. In particular, ADLT may be run with no NLD. If the generated documents are found to be inadequate, an NLD can be introduced and ADLT can be rerun to produce better documents.

We chose Prolog as the translation engine because Prolog offers a concise expression of translation rules. Prolog rules represent a convenient localization quantity: each locale has its own set of Prolog rules describing natural-language translation for that locale. Moreover, Prolog offers rapid development and the power and flexibility to perform more complex translations with little perturbation to other components of ADLT as the document generation component evolves.

The C++ component of ADLT manages the user input, parses the ADL, TDD and NLD inputs and generates the test program. The C++ component acts as the coordinator of services that ADLT provides. It runs the Prolog engine as a

child process and opens a UNIX pipe to the Prolog engine. Parent and child communicate bidirectionally over the pipe.

Having chosen this structure, we were faced with a choice of format for the ADL assertions presented to the Prolog engine. Because such a choice minimizes the amount of work done in the C++ code and because Prolog is able to parse infix operator expressions, we chose to present ADL expressions directly as Prolog terms. However, ADL and Prolog have incompatible rules regarding the characters in an identifier. This was addressed by placing all ADL identifiers (and operators) in single quotes for the Prolog engine. In this form, ADL expressions comprise Prolog atoms and operators and are thus compatible with the Prolog parser.

The entire system, including the document generation facility, is to be internationalized early in its development. The first Japanese localization was completed in June of 1995.

Localization of the ADLT C++ code is performed in a manner compatible with X/Open specifications [3]. It requires separate messaging files for each supported locale. Similarly, the translation subsystem requires that a separate set of Prolog rules be written for each locale to describe the translation of ADL assertions into prose for that locale.

The ADLT document generator uses run-time LOCALE [3] environment information to select the Prolog rule files used to produce documents in a particular target language. Changing the target language of the documents is as simple as setting new LOCALE environment information and re-running ADLT.

4.3 Analysis

This section examines in greater detail the process of translating ADL specifications into prose.

4.3.1 The translation process

The principal task of the translation system is to take an ADL assertion, which might look like:

```
(a == b) --> (c == d)
```

and translate it into prose like:

```
if a is equal to b then c is equal to d
```

Each locale has associated with it a static set of Prolog rules that describe translation for that locale. These are the core rules. For example, core rules for English translation define the translation of ADL operators into English prose. Core rules are sufficient to translate assertions in the absence of NLD input. ADLT is designed to look for additional translation information in an NLD file. That information is used to augment the default rules for that locale.

An example of an NLD file which might be used with the ADL file containing the above implication expression follows:

```
a = "the first leg of a triangle",
b = "the second leg of a triangle",
c = "the angle between the first leg and the
    third leg",
d = "the angle between the second leg and the
    third leg"
```

The use of this NLD would result in the following translation:

```
if the first leg of a triangle is equal to
the second leg of a triangle then the angle
between the first leg and the third leg is
equal to the angle between the second leg
and the third leg.
```

That translation occurs as a result of the two cooperating components of the translation process. The C++ component of ADLT initiates a request for translations by sending a question to the Prolog engine. The Prolog engine answers the translation request with the translated text. The C++ component then packages the resulting translation in documents it is called upon to produce.

The following example is offered to illustrate translation of a simple but syntactically complete ADL module. Figure 2 represents an ADL module that makes assertions on some simple C functions that operate on triangles. One type of document generated by ADLT is the Natural Language Specification (NLS). Figure 3 shows the NLS output generated from the ADL input in Figure 2 without an NLD. Figure 4 represents a possible NLD for the ADL in Figure 2. Figure 5 shows the NLS that ADLT produces from the ADL of Figure 2 with the NLD of Figure 4 as input. The generated NLS' in these examples have been abbreviated because of space limitations.

```
module triangle {
boolean
isosceles(int a, int b)
    semantics{
        (a == b) <-> return == TRUE
    };

boolean
equilateral(int a, int b, int c)
    semantics{
        isosceles(a,b)&& isosceles(b,c) <->
return == TRUE
    };
}
```

FIGURE 2 — Triangle ADL

SYNOPSIS:

```
boolean
isosceles(int a, int b)
```

GENERAL BEHAVIOR:

a is equal to b if and only if return is equal to TRUE.

SYNOPSIS:

```
equilateral(int a, int b, int c)
```

GENERAL BEHAVIOR:

The value returned by `isosceles(a,b)` and the value returned by `isosceles(b,c)` if and only if return is equal to TRUE.

FIGURE 3 — NLS without NLD

```
equilateral(p1,p2,p3){
  p1 = "the first leg of the triangle",
  p2 = "the second leg of the triangle",
  p3 = "the third leg of the triangle",
  $DESCRIPTION = "Verify that the triangle
specified by its legs is an equilateral
triangle."
}
```

```
Isosceles(p1,p2) = "test whether" p1 "is
equal to" p2{
  p1 = "leg one",
  p2 = "leg two",
  $DESCRIPTION = "Verify that the two legs
given are equal."
}
```

FIGURE 4 — Triangle NLD

SYNOPSIS:

```
boolean
isosceles(int a, int b)
```

DESCRIPTION:

Verify that the two legs given are equal.

GENERAL BEHAVIOR:

Leg one is equal to leg two if and only if return is equal to TRUE.

SYNOPSIS:

```
equilateral(int a, int b, int c)
```

DESCRIPTION:

Verify that the triangle specified by its legs is an equilateral triangle.

GENERAL BEHAVIOR:

Test whether the first leg of the triangle is equal to the second leg of the triangle and test whether the second leg of the triangle is equal to the third leg of the triangle if and only if return is equal to TRUE.

FIGURE 5 — NLS with NLD

4.3.2 Some Prolog details

When ADLT gets a request for documents it runs the Prolog engine. The Prolog engine is initialized with the core rules that produce the default translations of ADL assertions, unaided by NLD input. ADLT can produce documents for each locale for which such a core set of rules is available. Core rules are stored as files in a hierarchy structured by locale. ADLT selects among the core rules based on its run-time locale information. The initialized Prolog engine runs in query mode connected to ADLT via a UNIX pipe.

ADLT makes translation requests in the form of Prolog questions. Those questions are formed with a term defined in the core rules, `xlate/2`. An ADL expression like:

```
a == b
```

is sent to the Prolog engine on the pipe as the following Prolog question:

```
xlate('a' '==' 'b', ctxt).
```

The `ctxt` term is a representation of the ADL scope of the ADL expression being translated, which is the first term of `xlate/2`.

The following predicates are from the core rules for the English locale:

```
natural(M '==' N,Ctxt,[En,is,equal,
  to,En]) :-
  natural(M,Ctxt,Em),
  natural(N,Ctxt,En).
```

...

```
xlate(A, Ctxt) :-
  natural(A,Ctxt,Et),
  flatten(Et,S),
  unlist_write(S),
```

ADLT, in order to find the translation of `a == b`, writes `xlate('a' '==' 'b', ctxt)` on the pipe for the Prolog engine to read. Prolog reads that question and answers it by writing the translation "a is equal to b" on its end of the pipe. ADLT then reads that translation from the pipe and incorporates it in the current document.

When an NLD input is available, the Prolog engine is started as a child process and initialized just as when no NLD is present. ADLT then transforms NLD dictionary entries into Prolog facts, which are added to the core rules with the `asserta/1` predicate. As with other communications between ADLT and the Prolog engine, those facts are sent over the pipe to the Prolog engine.

Recall the three classes of substitution translations performed by ADLT were operator, object and function. The NLD does not offer the ability to change the definitions of ADL operators. Thus, there are two classes of substitution affected by the NLD: object substitutions and function reference substitutions. Object substitutions are complicated

by namespace rules of ADL for compound data types (structures, unions, enums) but they are nonetheless illustrated by the work ADLT performs on a simple object. An example of simple object substitution follows.

An NLD entry like:

```
a = "the phrase"
```

is transformed by ADLT into:

```
asserta(map('a', ctxt, 'the phrase'))
```

where *ctxt* is a unique atom for the ADL scope in which the substitution should be performed. *asserta/1* is used to add the *map/3* fact to the core rules. Note that the assert occurs before any requests for translation are sent to the Prolog engine. Interpreting the NLD and sending the resultant facts to the Prolog engine are initialization steps the documentation system must take before any translations are rendered.

Object translation is defined within the context of ADL scopes. Each of the translation requests has a reference to the ADL scope over which it is valid. ADLT must assert some facts about the nesting structure of the ADL scopes. They are of the form:

```
asserta(nested(b,a)).
```

Where *b* is a scope that is nested within the scope *a*.

ADL is structured by modules. Modules can import other modules. Importation determines the visibility of identifiers. Facts relating to the module importation structure must also be asserted on the core rules. They are of the form:

```
asserta(imports(b,a)).
```

Where *b* is a module that imports the module *a*.

The following Prolog fragment is from the core rules for the English locale:

```
natural(X,Ctxt,X1) :-
    atomic(X),
    map(X,Ctxt,X1).

...

map(X,Ctxt,X1) :-
    nested(Ctxt,CE),
    map(X,CE,X1),
    not X = X1.

map(X,Ctxt,X1) :-
    imports(Ctxt,Imports),
    maplist(X,Imports,X1).

map(X,_,X).

maplist(X,[H|T],X1):-
    (map(X,H,X2),
     X = X2,
     maplist(X,T,X1));
    map(X,H,X1).
```

During translation, the predicate *map/3* is applied to the atomic nodes of ADL expression trees. Those facts map identifiers to their translations for the appropriate scope and module visibility. Where no applicable *map/3* facts have been asserted, the rule *map(X,_,X)* performs the identity translation.

NLD entries for functions are a bit more complicated. An NLD entry like:

```
a(x,y)="swap" y "for" x {
    x = "param one",
    y = "param two"
}
```

is transformed by ADLT into:

```
asserta(mapfn(a,['swap',2,'for',1]))
asserta(map('x',ctxt,'param one'))
asserta(map('y',ctxt,'param two'))
```

mapfn/2 is used in the core rules to match the function (*a* in this case) with its translation. The translation involves interpolating the translation of the actual parameters of the function reference with the text from the NLD entry for that function, in the order specified by the NLD. The rules represent formal parameters as integers that correspond to the parameter's position in the NLD specification. Note that an ADL function expression is treated as a Prolog term and a function with no parameters in ADL is actually a term with one parameter, a blank in the following fragment from the core rules for English translation:

```

mapfun(X,C,Out) :-
    functor(X,F,A),
    A > 0,
    mapfn(F,M),
    xlatefn(X,C,M,Out).

mapfun(P,_, [the,value,returned,by,P]) :-
    functor(P,_,Ar),
    Ar > 0.

mapfun(X,_,X):- !.

xlatefn(_,C,[],[]).

xlatefn(XX,C, [Mh|Mt], [Oh|Ot]) :-
    ((integer(Mh), arg(Mh,XX,P),
     natural(P,C,Oh));
     (not integer(Mh), map(Mh,C,Oh))),
     xlatefn(XX,C,Mt,Ot).

```

4.4 Conclusions

Though it makes no pretense at linguistic sophistication, the substitution-based translation used in ADLT represents a surprisingly useful level of translation service. It establishes a minimum level of translation service upon which future extensions might yield better prose. Even in its current form, the translation system has received positive reviews from early ADLT users.

Not only does the system deliver a useful level of translation, it does so with relatively small core rule files. There are about 370 lines in the core rules for English. About 300 of those must be replaced in order to create rules for other locales. In the case of Japanese localization, that task has been performed by someone with only a rudimentary knowledge of Prolog. Localization of ADLT is thus relatively easy, even taking into account the automated document generation system.

Typically, ADLT requires little work beyond the creation of an ADL specification to produce accurate documentation. In fact, it requires no work if the default translations are found to be adequate. In the event that the generated prose is found lacking, the work required of the user is the creation of an optional NLD file and possibly some editorial work on the generated documents. System specifications can be structured to reuse NLD files to minimize the number of dictionary entries a user must create.

In common practice, document preparation is a very labor intensive, error-prone and expensive process. Synchronizing documents with specifications and implementations that evolve is a management nightmare. ADLT delivers a document generation system that is capable of producing draft quality documents from the formally specified system. If the system changes, the system documentation changes with it. ADLT-generated draft-quality documents will always be up to date with the system specification.

ADLT includes a Prolog-based document generation system that:

- is easy to localize
- produces draft-quality documents with little user input
- ensures documentation correctness

5. Future Plans

Even though the initial major phase of the research is nearing its end, the ADL Project team feels that enough positive results have been achieved to warrant continued research and development in this area. In fact, the team is planning a variety of interesting projects based upon this research. Several of these are described in this section.

5.1 Deployment Plans

The team is working to encourage the broad deployment of ADL. The first tangible results of this work are the following large testing projects:

5.1.1 CORBA 2 Test Suite

The Object Management Group has recently completed work on revision 2 of their CORBA specification. This revision addresses many of the problems identified through the ADL Project's experimental application to CORBA 1.2. It also contains language bindings for C++ and Smalltalk, a definition of interoperability requirements, and some needed enhancements to the C Language Binding and the Interface Definition Language description.

X/Open and the Object Management Group are working together to develop tests for CORBA 2. These tests will be based upon the CORBA 1.2 testing work done last year. This project has already started with an initial testability analysis of the new CORBA 2 specification. Test Suite development is expected to start before the end of the year.

5.1.2 Common Desktop Environment Test Suite

Another important emerging standard is the Common Desktop Environment. This collection of specifications, published by X/Open, defines an enhanced graphical user interface based upon the X Window System and Motif.

The testing challenges presented by these specifications are significant. The specifications are not as tightly written as other, more mature standards. As a result, testing these specifications via ADL is expected to identify a number of ambiguities even as the implementations of the specifications are being completed. This should help speed the availability of consistent conforming products from a variety of Open Systems vendors.

X/Open is just now starting the process of developing the ADL interface specifications for CDE. Test suite development is expected to start in earnest by year end, with results in the second quarter of 1996.

5.1.3 X/Open Specifications

X/Open is the Open System industry's principle source of specifications and conformance test. As such, it incurs significant costs in the development of these specifications and tests, as well as in the maintenance of both. X/Open views the ADL research project as a potential way to reduce these costs and at the same time improve the quality of the specifications and the tests.

In order to determine the extent of these benefits, X/Open will start using ADL to develop complete specifications and their associated tests early next year. While the exact specification against which ADL will be applied has not yet been determined, X/Open has a general plan for its application:

- Train X/Open work group participants in ADL early in 1996.
- Have that work group develop their specification in ADL, including the production of the natural language dictionary.
- Work toward a mature draft of the specification in ADL by the third quarter.
- Start development of a test suite using the ADL interface specification in the fourth quarter, even as the base specification is being completed.
- Have input from the test developer assist in completing the specification.
- Publish a specification at the beginning of 1997.
- Have an ADL generated test suite available against the specification shortly thereafter.

Based upon the experience gained by this experiment, X/Open plans to start full deployment of ADL within its community at the beginning of 1997.

5.2 Enhancement Plans

Even as ADL Release 1.0 is being deployed, the research into extending and enhancing ADL will continue. Some of the expected research areas are described in this section.

5.2.1 ADL for C++

C++ is an emerging standard in the computing community. It is vitally important that ADL be able to be applied to the specification and testing of C++ interfaces and objects. X/Open and Sun Microsystems Labs are currently exploring the possibility of doing research in this space. While these plans are not finalized, we hope to begin research this calendar year, with results beginning to emerge by the second quarter of 1996.

5.2.2 ADL for IDL

While C++ is emerging as the programming language of choice for application developers, OMG's Interface

Definition Language (IDL) is becoming the *de facto* standard for the deployment of objects and object services. Since the ADL Project team believes that object-oriented application development will be commonplace in the near future, we are investigating the possibility of producing direct support for IDL within ADL. This project is not expected to start until late 1996, with results emerging in the second quarter of 1997.

5.2.3 ADL Localization Packages

One of the most significant advantages of ADL is its ability to produce quality specifications in any natural language. The support of these languages is not, however, automatic. A localization package, or set of files that describe the rules and default translations, for each language must be developed.

The ADL Project team is still considering which languages should be done, but we expect to develop packages for German and French at the very least. This work should be done in the second quarter of 1996.

5.3 Follow-on Research

Our intention is to make ADL the foundation for a family of related test and specification tools. In this way we make maximum use of the user's effort in creating the initial specification. By driving a family of tools from the same specification, we obtain several advantages:

- The investment in creating the specification is put to maximum use.
- Consistency is assured; the implementor and the various tests and documents are guaranteed to use the same specification.
- In many cases, the support software used for one test can be re-used in another test framework.

New testing capabilities can be added by both the ADL team (several efforts are underway) and by outside researchers (efforts are underway at the University of California at Irvine, the University of Texas at Austin, Oxford University, and others).

5.3.1 Unit Testing

ADLT is intended as a tool for unit testing. It is aimed specifically at the problem of conformance testing, checking each procedure in an interface separately to ensure that it meets its documented specification.

Each interface must be tested separately because each procedure may be used separately. The goal of unit testing is to test each element of the interface separately, to some standard of coverage that gives reasonable confidence that the element meets its specification.

5.3.2 System Testing

The purpose of integration, or system, testing is to ensure that the elements of an interface work correctly in various combinations. The number of possible combinations is astronomical, too large to attempt anything approaching complete coverage; the test criterion is to check some typical fragments that might be expected to occur in user code.

System testing provides a less exhaustive but more realistic check of an interface. It is not sufficient by itself, because the number of possible combinations makes it impossible to get good coverage; it is necessary as part of a complete testing program, because it may expose bugs that are not revealed by unit testing.

The existing ADLT tool can be extended to a tool for testing code fragments as well as individual procedure invocations. A design and a prototype for this extension have been completed.

5.3.3 Random Testing

Random testing is an alternate strategy for system testing. Rather than pre-selected code fragments, it executes a random sequence of calls from the interface. In previous experiments, we found that a pure random test sequence was usually not very interesting; it ends up exercising trivial error checking code over and over again. Since then, we have used some of the mechanisms from the present ADLT tests to create more interesting random tests. This work has advanced as far as a successful experimental prototype; we plan to continue this promising avenue of research.

5.3.4 Stress Testing

A stress test is one that subjects the software under test to uncommon stress, to determine that it fails gracefully when resources are exhausted. A stress test determines what the resource limits of an implementation are, and often exposes bugs in the implementation, because the error checking and catching code paths are seldom executed in normal circumstances.

An ADLT test is often, by its mechanical nature, a stress test for some aspects of the software; for example, a fault in resource reclamation will often be exposed by an ADLT test that cycles through thousands of instances of a resource. More specific stress testing is a possible topic for future investigation.

5.3.5 Interference Testing

Another task is interference testing, to check that simultaneous use of an interface by several clients does not lead to problems.

It is important to make a distinction between simultaneous independent use, in which the actions by one user are not supposed to affect the other user, and unrestricted concurrent use, in which the actions of one user may affect the other. In general, one user can affect another only if there are shared elements between the two users. For example, if two users are accessing different directories of a file system, the system should behave as if each one was the sole client; if they are accessing the same directory, they should see the result of each others actions as well as their own.

Of the two testing problems, the one of checking for interference between independent clients is much simpler. It can be testing by the present ADLT test. We are investigating possible solutions for the harder problem of concurrent testing.

5.3.6 External Research

In addition to the ADLT team's research efforts mentioned above, there are collaborative research efforts at several universities and companies. We welcome such efforts and will support them to the best of our abilities. Part of the reason that the ADLT code is in the public domain is to enable experimentation into other test strategies.

6. Conclusions

The ADL Project has been, and continues to be, a success. When the ADL Project team started this work, it was expected to realize a certain set of objectives. These objectives have been largely realized. The research results have been promising enough that X/Open and its members and Sun Microsystems Laboratories expect to continue evolving the basic research into the foreseeable future. While some of the initial hypotheses associated with the project have not proven correct, the overall result has been that a useful tool has been developed and presented to the industry for use. By 15 December 1995 this tool will be complete. At that time the team fully expects it to be put to use by a broad variety of software and system developers.

7. References

- 1) Sriram Sankar and Roger Hayes, *Specifying and Testing Software Components using ADL*, Sun Microsystems, 1993
- 2) Matt Evans and Shane P. McCarron, *Assertion Definition Language Translation System*, SRI Quality Week, 1994
- 3) X/Open Company Limited, *Internationalisation Guide*
- 4) Andrew Silverman, *A Review of ADL*, Applied Testing and Technology, 1995
- 5) Barry Books, *ADL Report*, Mount Bonnell, Inc., 1995