

n ro n or op on En n o
n r r or

2. Coordination Problems in IV&V

Independent Verification and Validation (IV&V) is a

built into support tools, it should be possible to detect that Alice uses some parts of a spec to check others, and hence record a dependency relationship. Section 4 discusses our current project to introduce such support tools incrementally into ongoing projects.

3.2. Scenario 2

The second scenario (Figure 2) illustrates a more formal process, in which fewer communication problems can occur. An IV&V analyst (let us call him Carl) is analyzing a section of the requirements document by generating a tabular version of the section, and then running the resulting tables through an automated consistency checking tool. The tool reports an inconsistency, which Carl traces back to a mistake in the original document. He writes a Discrepancy Report (DR). Let us call this DR#101.

Three months later, a new draft of the specification is released. Carl checks the DR database, to see which of his DRs have been addressed in this new draft. DR#101 is marked as having been worked on (by Diane), and is awaiting approval for closure. As originator of the DR, Carl's signature is required before the issue can be closed. He updates his tabular representation to reflect the new draft, runs the new version of the table through the tool again, and confirms that the problem is now fixed. He therefore signs off the DR as closed.

The DR tracking tool removes many potential communication problems, and ensures that closure is

achieved for each reported problem. However, coordination problems can still occur.

For example, Carl could have made mistakes in the translation from the text to the table - it is hard to confirm that the table is a faithful representation of the textual requirements. Similarly, Carl might not be able to trace the inconsistency back to the original requirements. He would then have great difficulty reporting the problem in a DR, unless he includes his tables, a description of the

relationship anywhere else. Either of the two chunks of specification may evolve, but there is no way to trace the “ripples of influence” of any changes. Hence, there is no opportunity for tool support to reason about how changes to one side affect the other.

The scenario illustrates how expensive it can be to develop and maintain an alternative representation of an evolving specification. This may mean that this type of analysis gets delayed until the specification is relatively stable. This is undesirable.

Notice that the relationship is bi-directional. Although the table is generated and updated from the text, Carl needs to be able to trace problems from the table back into the text. It is also highly likely that Carl may want to alter the table to see what possible fixes there are, and then see what effect this has on the text.

4. The Web as an enabling technology

The problems identified in the scenarios show how time-consuming and costly it can be to track changes, especially where there are many dependencies throughout the specifications. A full solution to these problems would require all dependencies between different parts of a specification to be explicitly represented. Such a solution requires significant advances in the capture and representation of dependencies between specification elements. Partial solutions exist for individual methods (e.g. the consistency checking for SCR [3]). A general solution for multiple methods is still a long way off.

To explore such a general solution, we have adopted an incremental, empirical approach. We need to put into place the infrastructure for recording data about each chunk of specification, including annotations and relationships with other chunks. However, we also need to integrate this infrastructure with the existing project support systems on the projects we wish to study, to minimize the disruption caused. We will not be able to proceed with our empirical study unless each step is relatively painless for the project.

The infrastructure we need to put in place must satisfy two major criteria. It must be adaptable enough to fit in with a wide range of existing project support tools on different platforms, using heterogeneous networks, and accessing existing project data in a variety of different formats. Second, it must provide the ability to record and track arbitrary relationships between chunks of specifications [4].

the notation defined in the style slot, following the strategy defined by the work plan, for a particular problem domain. A development history is maintained in the work record. This framework encourages multiple representations, and is a deliberate move away from attempts to develop monolithic specification languages. It is also independent from any particular software development method.

The WHERE project will implement this framework using the Web. The core functionality will be provided by