

1 International Journal of Software Engineering
and Knowledge Engineering
3 Vol. 15, No. 4 (2005) 1–28
© World Scientific Publishing Company



5 **SUPPORTING SOFTWARE UNDERSTANDING WITH
AUTOMATED REQUIREMENTS TRACEABILITY**

7 ALEXANDER EGYED
Teknowledge Corporation,
9 *4640 Admiralty Way, Suite 1010, Marina Del Rey, CA 90292, USA*

11 PAUL GRÜNBACHER
Systems Engineering and Automation,
Johannes Kepler University, 4040 Linz, Austria

13 Requirements traceability (RT) aims at defining and utilizing relationships between
15 stakeholder requirements and artifacts produced during the software development life-
17 cycle and provides an important means to foster software understanding. Although tech-
19 niques for generating and validating traceability information are available, RT in practice
often suffers from the enormous effort and complexity of creating and maintaining traces.
21 This results in invalid or incomplete trace information which cannot support engineers
in real-world problems. In this paper we present a tool-supported approach that requires
23 the designer to specify some trace dependencies but eases trace acquisition by generat-
ing others automatically. We illustrate the approach using a video-on-demand system
25 and show how the generated traces can be used in various engineering scenarios to im-
prove software understanding. In a case study using an open source software application
we demonstrate that the approach is capable of dealing with large-scale problems and
delivers valid results.

Keywords: Software understanding; software traceability; automation.

27 **1. Introduction**

29 Requirements traceability (RT) is defined as the “ability to describe and follow the
life of a requirement, in both a forward and backward direction” [1] by defining and
31 maintaining relationships to related development artifacts [2] such as stakeholder
needs, architectural components, design model elements, or source code. RT is con-
sidered crucial for establishing and maintaining consistency between heterogeneous
33 models used throughout the development life-cycle [3]. Frequently reported bene-
fits of RT include the facilitation of communication, support for the integration of
35 changes, the preservation of design knowledge, quality assurance, and the preven-
tion of misunderstandings. Trace information fosters software understanding and
37 assists engineers in dealing with critical issues in software development and main-
tenance. For example, engineers might be interested in the origins of a requirement

2 *A. Egyed & P. Grünbacher*

1 (e.g., the stakeholder needs) or the rationale for a particular design choice. They
3 might also need to know how exactly functional or non-functional requirements are
5 realized in the system, or if an implementation completely realizes a given set of
7 requirements. During system evolution and maintenance, RT is also important for
9 analyzing the impact of new requirements or changes to existing ones. Many critical
11 risks in software engineering are architectural [4] and have to do with persistent soft-
13 ware attributes such as performance, reliability, or security [5]. Architectural risks
15 have to be considered, in particular, when new requirements are assessed or exist-
ing requirements are changed. Risk assessment, however, is challenging and relies
on understanding the complex relationship between requirements, desired system
properties, and architectures during development and maintenance [6, 7], so RT
becomes particularly important.

13 The benefits of RT are widely accepted nowadays and sophisticated tool support
is available to record, manage, and retrieve trace information [8]. However, several
15 issues still hamper wide-scale adoption of RT in software engineering practice:

- 17 • Acquiring traces is still mostly a manual process with only little automation
available. This results in enormous effort and complexity [9].
- 19 • The full potential of RT can only be exploited if complete trace information is
available. However, incomplete trace information is a reality due to complex trace
acquisition and maintenance.
- 21 • It is often hard to anticipate the kind of engineering issues that might arise
later and the trace information recorded for one particular purpose might be
23 insufficient for future tasks.
- 25 • Traces have to be identified and recorded among numerous, heterogeneous engi-
neering artifacts (document, models, code, ...). It is often very challenging to
create meaningful relationships in such a complex context.
- 27 • Traces are in a constant state of flux since they may change whenever require-
ments or other development artifacts changes.

29 Automating RT can deal with many of these issues if it goes beyond mere
recording and replaying of trace information [10]. We have thus been developing an
31 automated traceability approach [11, 12] that relies on providing a small number
of easy-to-find trace dependencies as input. The result of the approach are trace
33 dependencies among various artifacts, such as traces among functional artifacts (re-
quirements, architecture), traces between functional and quality (non-functional)
35 requirements, as well as traces among quality requirements. Although the automatic
creation of these dependencies is a significant advancement over traditional trace-
37 ability techniques there is still the challenge to interpret and use the created links.
We have to understand the meaning and implications of these trace dependencies
39 (e.g., conflicts and cooperations between requirements) [13]. The manual investiga-
tion of all trace dependencies, however, is tedious and error prone as there are likely
41 thousands of links one would have to investigate in a large-scale system. We thus

1 also propose some heuristics helping to define the meaning of trace dependencies
based on the types of the bridged requirements.

3 This work is a continuation of our earlier work on identifying trace dependen-
cies using scenarios [11]–[15]. The contribution of this paper is on showing (1) how
5 RT results derived by our approach can be interpreted, (2) how functional and
non-functional requirements are treated, (3) how evolutionary/incremental require-
7 ments engineering is supported, and (4) how well the approach supports large-scale
complex software systems.

9 The remainder of this paper is organized as follows: Section 2 explains our
Trace Analyzer technique using the Video-On-Demand (VOD) example, a simple
11 software application we use for the purpose of illustration. Section 3 discusses how
to interpret and understand the created trace links. In Sec. 4, we demonstrate the
13 capability of the approach in a case study in which we applied the technique to
a large-scale open source software package. Section 5 discusses how the generated
15 traces support various software engineering scenarios. In Sec. 6 we discuss related
work. Conclusions and an outlook on further work round up the paper.

17 2. Automating Software Traceability with Trace Analyzer: The Video-On-Demand Example

19 Trace dependencies describe relationships between different artifacts such as re-
quirements, designs, assumptions, rationale, system components, source code,
21 etc. [2]. The value of recording and maintaining these dependencies is to support
software understanding and to help engineers in answering questions such as “Why
23 is this requirement here?”, or “What happens if I change this design element?”
Trace dependencies describe the origin, rationale, or realization of software devel-
25 opment artifacts.

Trace dependencies are not static but highly dynamic because software evolves.
27 For instance, if a requirement R leads to the implementation of some source code C
then a trace dependency exists between the two. If the requirement changes then the
source code is potentially affected. Conversely, a change to the source could make
29 an update of the requirement necessary. This bi-directionality is very important for
trace analysis and implies that if R depends on C then C depends on, at least, R .

31 The Trace Analyzer [11, 12] defines trace dependencies through (a) shared use
of source code and (b) transitive reasoning:

33 (a) The source code of a software system is a useful medium to identify trace
dependencies. If, say, requirement A depends on some source code C_A and require-
35 ment B depends on some source code C_B then one can infer that A and B depend
on each another if C_A and C_B overlap (i.e., because they are implemented together).
37

(b) Transitivity is an intrinsic property of trace dependencies. It defines A to
39 depend on C if A depends on B and B depends on C .

41 As input, the Trace Analyzer technique takes a small number of known or hy-
pothesized dependencies between software artifacts (e.g., requirements and source

4 *A. Egyed & P. Grünbacher*

1 code). It then builds a graph containing nodes that capture source code and all
2 their overlaps. For example, there are separate nodes for the source code of *A* and
3 *B* but if they overlap then this overlap is explicitly captured in yet another node.
4 The graph is manipulated to move known artifacts among the nodes. The goal is
5 to constrain for all nodes what artifacts they relate to or not. Trace analysis is
6 complicated by imprecise input where single dependencies may include multiple
7 artifacts (*A* or *B* depends on *C*). It is also complicated by open-ended input where
8 only partial knowledge is available (*A* depends on *C* and possibly others). Trace
9 analysis is an iterative process using a large number of rules to manipulate the
10 graph structure. At the end, the graph is traversed to identify all nodes related to
11 individual artifacts. A trace dependency is established if two different artifacts re-
12 late to at least one common node. The graph also helps in determining the strength
13 or reliability of a dependency based on the number of nodes two artifacts have in
14 common.

15 The trace analyzer technique is fully automated and tool supported. The only
16 deficiency, as it may appear, is that some trace dependencies have to be pro-
17 vided as input manually; that is: traces between artifacts and code. Fortunately,
18 we found that this input can be partially generated by executing the source code
19 and observing the lines of code being executed. We use test scenarios to define
20 how to test individual artifacts or groups of artifacts. When executing a sce-
21 nario, we then observe which classes, methods, or lines of code are used. For
22 instance, in a first case study we employed the commercial tool IBM Rational
23 Pure Coverage® to observe the test scenarios of an executing system. In our 2nd
24 case study we used a simple open source tool *org.jmonde.debug.Trace* available
25 from <http://www.geocities.com/mcphailmj/Trace/> to record the necessary trace
26 information.

27 With the help of such tools, trace dependencies between test scenarios and
28 source code can be automatically generated during testing. That is, the tool lists
29 the methods, classes, and packages that are used during the execution of any given
30 test scenario. If a designer now specifies how these test scenarios relate to artifacts
31 (the premise) then one can automatically infer trace dependencies between these
32 artifacts and the code they are using. These trace dependencies are then used as
33 input to the trace analyzer to generate yet other trace dependencies.

2.1. *Video-on-demand system*

35 We illustrate the benefits of the Trace Analyzer technique using a simple video-on-
36 demand (VOD) system, which was developed by a third party (see <http://peace.snu.ac.kr/dhkim/java/MPEG/>). This system provides capabilities for searching, se-
37 lecting, and playing movies. It supports playing a movie concurrently while down-
38 loading its data from a remote site. VOD's complex computational logic is well-
39 hidden underneath a simple VCR-like user interface (play, pause, stop button).
40 Both functional and non-functional aspects are important for the requirements of
41 the VOD.

Table 1. VOD requirements.

r0	Download movie data on-demand while playing a movie (<i>Functionality</i>)
r1	Play movie automatically after selection from list (<i>Functionality</i>)
r2	Users should be able to display textual information about a selected movie (<i>Functionality</i>)
r3	User should be able to pause a movie (<i>Functionality</i>)
r4	Three seconds max to load movie list (<i>Efficiency/Time behavior</i>)
r5	Three seconds max to load textual information about a movie (<i>Efficiency/Time behavior</i>)
r6	One second max to start playing a movie (<i>Efficiency/Time behavior</i>)
r7	Novices should be able to use the major system functions (selecting movie, playing/pausing/stopping movie) without training (<i>Understandability</i>)
r8	User should be able to stop a movie (<i>Functionality</i>)
r9	User should be able to (re-)start a movie (<i>Functionality</i>)

1 The VOD system consists of 21 Java application-specific classes, it also adopts
 2 numerous off-the-shelf library classes. Static and behavioral aspects of the VOD
 3 were also modeled using various UML diagrams (see the UML state diagram in
 4 Fig. 1). Requirements were not available but for the purpose of extending the VOD
 5 system (discussed later) we reverse-engineered them. Table 1 depicts a subset of the
 6 VOD requirements. Figure 1 shows a state diagram of the VOD system showing
 7 that it operates either in a movie selection mode (left) or in a movie playing mode
 8 (right). During movie selection, a user can select servers for downloading movie lists,
 9 inspect textual information about movies, and select individual movies for playing.
 10 During playing mode, a selected movie may be paused, stopped, and played again.
 11 The transitions between these states correspond to buttons a user may press in the
 12 VOD’s user interface. For instance, a user may press the “Movies” Button at any
 13 time during movie playing to select another movie.

14 The drawback of the existing requirements and UML diagrams is that no trace
 15 dependencies are known. In some cases, trace dependencies could be guessed fairly
 16 easy but the informal nature of the requirements and the semi-formal UML model
 17 make it hard to manually identify complete and correct trace dependencies.

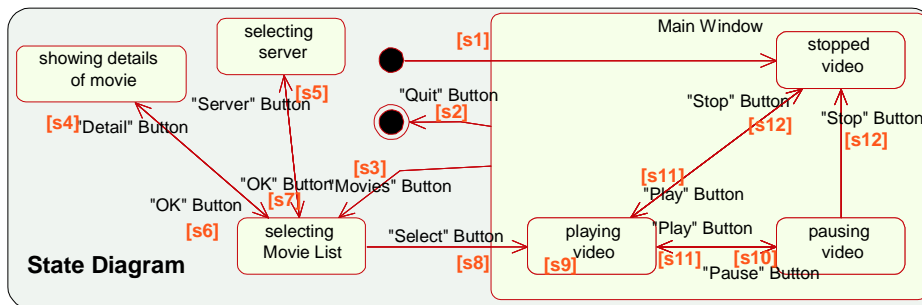


Fig. 1. UML state diagram of VOD system.

1 **2.2. VOD scenarios**

3 Our scenario-based trace analysis approach automatically defines trace dependen-
 4 cies among requirements, between requirements and code, and between require-
 5 ments and model elements (e.g., state transitions). The approach only requires
 6 scenarios that can be tested against the code to identify trace dependencies auto-
 7 matically. Table 2 lists all test scenarios we defined for the case study. For example,
 8 test Scenario 1 uses the VOD to display a list of movies. The details of how to test
 9 this scenario on the system are omitted for brevity but the test scenario describes
 10 how to configure the VOD system and what user interface actions to perform (e.g.,
 11 which buttons to press) in order to achieve the desired results. IBM Rational’s
 12 PureCoverage® was used to monitor the VOD system during the testing of the
 13 scenarios. For example, the Java classes `BorderPanel` (C), `ListFrame` (J), `Server-`
 14 `Req` (R), and `VODClient` (U) were executed while testing Scenario 1. For the sake
 15 of brevity, we only use single letter acronyms for Java classes.

16 Table 2 also shows which artifacts (model element, requirements) the different
 17 test scenarios apply to. For instance, our hypothesis was that test Scenario 1 relates
 18 to the state transition [s3] “Movies” Button in the state diagram (see Fig. 1). While
 19 executing the scenario, it was observed to execute the Java classes (code) [C,J,R,U].
 20 Due to transitivity of trace dependencies, one may conclude that the state transition
 21 [s3] depends on the code [C,J,R,U].

22 Table 2 defines 12 additional scenarios including one test scenario for each re-
 23 quirement (although multiple may exist) and, to make the trace analysis more
 24 realistic, some ambiguous test cases for the state diagram. We call a trace depen-
 25 dency ambiguous if it does not precisely define relationships among artifacts. For
 26 instance, test Scenario 2 defines the state transitions [s4] and [s6] relating to the
 27 code [C,E,J,N,R]. This statement is ambiguous in that it is unclear which subset of
 [C,E,J,N,R] actually belongs to [s4] and which subset belongs to [s6].

Table 2. Scenarios and observed footprints.

Test Scenario	Artifacts	Observed Java Classes
1. view movie list	[s3]	[C,J,R,U0]
2. view textual movie information	[s4,s6][r2]	[C,E,J,N,R]
3. select/play movie	[s8,s9][r6]	[A,C,D,F,G,I,J,K,N,O,R,T,U]
4. press stop button	[s9,s12][r8]	[A,C,D,F,G,I,K,O,T,U]
5. press play button	[s9,s11][r9]	[A,C,D,F,G,I,K,N,O,T,R,U]
6. change server	[s5,s7]	[C,R,J,S]
7. playing	[s9]	[A,C,D,F,G,I,K,O]
8. get textual movie information	[r5]	[N,R]
9. movie list	[r4]	[R]
10. VCR-like UI	[r7]	[A,C,D,F,G,I,K,N,O,R,T,U]
11. select movie	[r0]	[C,J,N,R,T,U1]
12. select/play movie	[r1]	[A,C,D,F,G,I,J,K,N,O,R,T,U]
13. press pause	[s9,s10][r3]	[A,C,D,F,G,I,K,O,U]

1 **2.3. VOD trace analysis**

2 Scenario-based trace analysis is fairly straightforward for unambiguous test sce-
 3 narios. For instance, requirement [r6] defines a maximum delay of one second to
 4 start playing a movie. We know from test Scenario 3 that [r6] executes the Java
 5 classes [A,C,D,F,G,I,J,K,N,O,R,T,U]. Consequently, this Java code needs to be op-
 6 timized to perform as desired. Trace analysis becomes more complicated in case of
 7 ambiguous scenarios. For instance, through Scenario 5 we know that pressing the
 8 play button causes [s11] directly and [s9] (playing the actually movie) indirectly.
 9 Although together [s9,s11] use 13 Java classes [A,C,D,F,G,I,K,N,O,R,T,U], it is left
 10 unspecified which subsets of those classes are used by [s11] or [s9]. Alternatively,
 11 through Scenario 7 we learn that [s9] alone uses the Java classes [A,C,D,F,G,I,K,O],
 12 which is a subset of [s9,s11].

13 For reasons of efficiency and precision, the trace analyzer uses a graph structure
 14 called the footprint graph to infer trace dependencies. Footprints are the observed
 15 lines of code executed while testing scenarios. Figure 2 shows a partial footprint
 16 graph based on Scenarios 1, 4, 5, 7, and 13. The child nodes represent subsets of
 17 parent nodes. This subset relationship applies to both the model elements and Java
 18 classes used. For instance, Scenario 7 is about playing a movie [s9] and it uses a
 19 subset of the lines of code that Scenario 5 uses. Scenario 7 [s9] furthermore refers
 20 to a subset of the model elements referred to by Scenario 5 [s9,s11]. Within the
 21 footprint graph this places the node for Scenario 7 below the node for Scenario 5.

22 Besides having nodes for each scenario, the footprint graph also contains
 23 nodes for all possible overlaps between scenarios. For instance, Scenario 1 over-
 24 laps with Scenario 5 because they both use the Java classes [C,R]. A node, child to
 25 both, is thus introduced to explicitly capture this overlap. It must be noted that

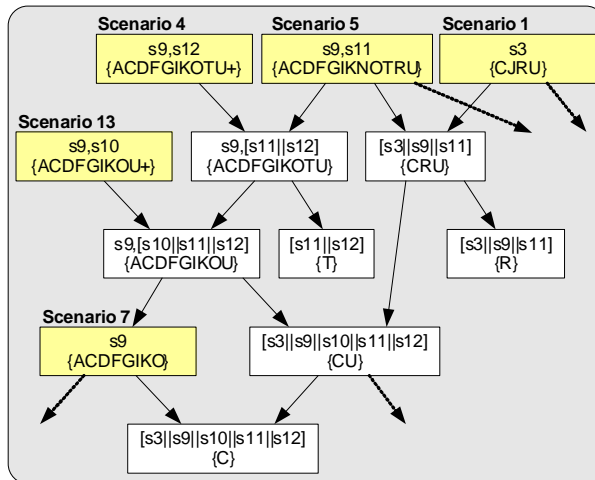


Fig. 2. Partial footprint graph.

1 some overlaps are omitted in the figure for brevity. In [11] we discuss details of
2 building a footprint graph. For example, it is the designer's choice on how to
3 relate the scenarios with the model elements. If the designer says that model
4 element s9 is about Scenario 7 then the designer is confident that s9 is exactly
5 the footprint [A,C,D,F,G,I,K,O]. Thus, our trace analyzer includes s9 in the node
6 [A,C,D,F,G,I,K,O] and it excludes every other node that this is not a subset of
7 [A,C,D,F,G,I,K,O] (e.g., B, E, H, T, U).

8 Once all scenarios are inserted into the footprint graph, the graph contains nodes
9 for every possible overlap between any two scenarios. The graph is then manipu-
10 lated to move artifacts around the nodes to identify for every node the artifacts it
11 could possibly relate to. For instance, parent node [A,C,D,F,G,I,K,O,T,U] has two
12 children. Each child relates to a subset of the Java classes of the parent but both
13 children together relate to the same Java classes as the parent. For consistency and
14 completeness, both children thus relate to the same model elements as their parent.
15 In this case, the parent was defined to relate to [s9] and [s12] (ignore s3 for the
16 moment) and thus each child individually may relate to a subset of [s9,s12]. In case
17 of the left child, we already know that it must relate to [s9]. In fact we know that
18 the other child [T,U] cannot be about [s9] as was discussed previously. Thus, the
19 model element [s12] must be in at least [T,U].

20 The same reasoning applies to the parent node [A,C,D,F,G,I,K,N,O,T,R,U]
21 which was defined to include [s9 and s11]. Since [s9] was defined to be exactly
22 [A,C,D,F,G,I,K] [s9] cannot relate to [N,R,T,U] and it is excluded in that node.
23 Given that the input required that [N,R,T,U] either belong to [s9] or [s11] and
24 given that we now know that it cannot be [s9] we derive that node [N,R,T,U] be-
25 longs to [s11]. However, this reasoning has one flaw. While we excluded [s9] from
26 [N,R,T,U], we did not exclude [s11] from [A,C,D,F,G,I,K]. It is quite possible that
27 a subset of [A,C,D,F,G,I,K] has shared ownership and that subset may also belong
28 to [s11]. This problem was addressed in [11] by explicitly tracking the possibility
29 of shared code. This issue is also discussed in a later section on limitations of the
30 approach.

31 We mentioned earlier that the approach can detect incomplete and incorrect
32 input based on inconsistencies and incompleteness in the footprint graph. In Fig. 2
33 we defined nodes with their included and excluded model elements. For example,
34 node [T,U] includes [s12] but it excludes [s9] and other model elements. We put
35 excluded elements in brackets to separate them from included elements. If a single
36 node includes and excludes the same model element then there is a conflict which is
37 the result of a conflicting input. Similarly, if the union of the included and excluded
38 list does not represent the total set of model elements then the node is incomplete.
39 For example, node [T,U] is incomplete because we do not yet know its relationship
40 to, for example, [s7]. The trace analyzer technique uses a larger set of rules than
41 can be described in this paper and there are many special cases one has to consider
42 to make it reliable. A detailed discussion is published in [11, 12].

Table 3. Artifact to Java class dependencies.

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	R	S	T	U		
r0			F						F					F				F	F	F		
r1	F		F	F		F	F		F	F	F			F	F			F		F	F	
r2			F		F					F				F				F				
r3	F		F	F		F	F		F	F					F						F	
r4																			F			
r5														F				F				
r6	F		F	F		F	F		F	F	F			F	F			F		F	F	
r7	F		F	F		F	F		F	F				F	F			F		F	F	
r8	F		F	F		F	F		F	F				F							F	F
r9	F		F	F		F	F		F	F				F	F			F		F	F	
s3			P						P									P			P	
s9	F		P	F		F	F		F	F				F								
s10			P																		P	
s11			P												P				P		P	
s12			P																		P	

1 All rules have in common that they move model elements within the graph
 2 structure to identify all related model elements for every node. Since in this case
 3 study the leaf nodes refer to individual Java classes, we can infer all model elements
 4 related to a Java class. Table 3 summarizes some dependencies between artifacts
 5 and code that can be interpreted from the graph. Note that leaf nodes may also
 6 refer to packages, methods, or even individual lines of code if a different level of
 7 granularity is desired by the user.

8 From the footprint graph we can interpret that either model element [s11] or
 9 [s12] or both have a dependency to Java class T. Table 3 shows this dependency
 10 using a letter that indicates the confidence of the trace analyzer where column T and
 11 rows [s11] and [s12] intersect: “F” for full confidence; “P” for partial confidence. The
 12 trace analyzer determined that class T either depends on [s11] or [s12]. Consequently
 13 one only has partial confidence that s11 depends on class T. In fact, one may only
 14 then conclude that s11 depends on class T if it becomes known that [s12] does not
 15 depend on class T.

16 In some cases, the trace analyzer technique can reduce ambiguous input. For
 17 instance, Scenario 3 in Table 2 defined [s8] to potentially depend on class F. Yet, the
 18 trace analyzer concluded that class F belongs to [s9]. Although the trace analyzer
 19 technique can reduce ambiguity, it cannot not always avoid it. Ambiguous dependen-
 20 cies are the result of ambiguous and/or incomplete input. The trace analyzer
 21 can also identify some forms of inconsistent input but this discussion is beyond the
 22 scope of this paper.

23 Trace dependencies among requirements are defined based on overlaps among
 24 the lines of code implementing those requirements. Table 3 only captures trace
 25 dependencies between requirements and code, and between model elements and
 26 code. Given the transitive property of trace dependencies, one can also use
 27 Table 3 and the footprint graph in Fig. 2 to infer dependencies among requirements
 and/or model elements. For example, Table 2 shows a trace dependency between

10 A. Egyed & P. Grünbacher

1 Scenario 12 (select movie) and Scenario 5 (press play button) as the latter executes
 2 a subset of the lines of code of the former. Since both scenarios represent test cases
 3 for different requirements ([r1] and [r9]), we can infer a trace dependency between
 4 [r1] and [r9].

5 Trace dependencies can even be inferred for quality requirements. Here, we use
 6 test scenarios executing the part of the system that is relevant for the quality re-
 7 quirement (i.e., the functional context of the requirements). For example, [r9] relates
 8 to [A,C,D,F,G,I,K,N,O,R,T,U] and [r6] relates to [A,C,D,F,G,I,J,K,N,O,R,T,U].
 9 Knowing that [r9] traces to a subset of the classes that [r6] traces to we can infer a
 10 dependency between [r9] and [r6]: the requirement for a “play button” also implies
 11 the non-functional constraint of only having less than one second to start playing
 the movie once the button is pressed.

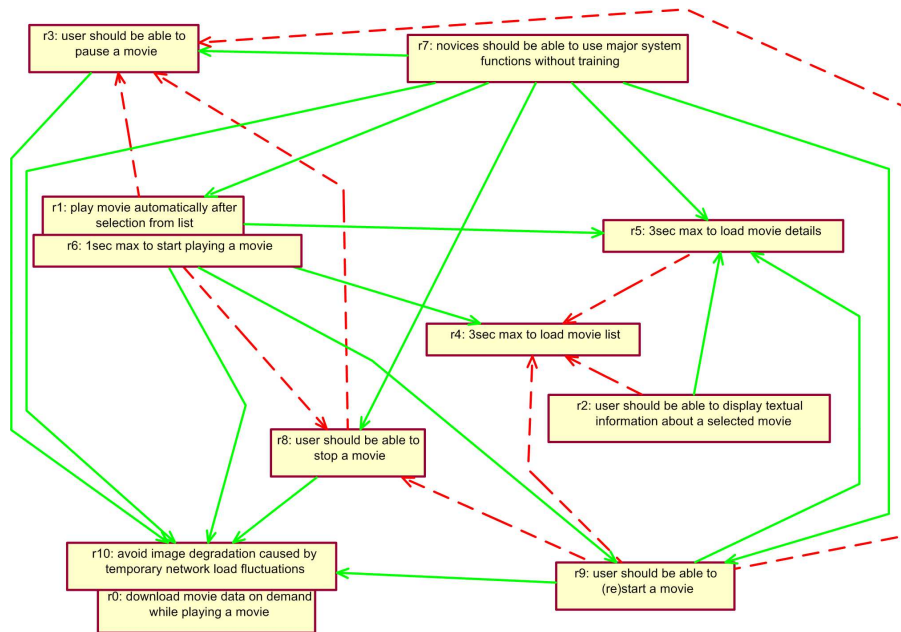


Fig. 3. Automatically created trace links between VOD requirements.

13 2.4. Validity and complexity of results

14 Figure 3 shows the requirements of the VOD system. The figure contains links
 15 that visualize this subset/superset relationship among executed lines of code for
 16 requirements as discussed above. In particular, a requirement pointed to by the
 17 arrow uses a subset of the lines of code of a pointing requirement. For example,
 18 there is an arrow from [r9] to [r0] because [r9] uses a subset of the lines of code of
 19 [r0]. There are also clusters of requirements (e.g., [r1] and [r6]). Requirements in this

1 cluster execute the exact same lines of code. Graphically this implies bi-directional
2 trace dependencies among the requirements within those clusters which we simply
3 abbreviate by forming clusters.

Here we would like to discuss two challenges:

- 5 • How can we deal with incorrect trace dependencies?
- How can we handle the high number of trace dependencies?

7 *Incorrect trace dependencies.* Figure 3 depicts some trace dependencies. Solid
8 lines imply correct trace dependencies while dashed lines represent incorrect trace
9 dependencies. Incorrect dependencies are identified if the code of two requirements
10 is interleaved such that the execution of one requirement always implies the ex-
11 ecution of the other although they do not interfere with each another; or if the
12 granularity of the trace analysis is not detailed enough. In our case, the latter is at
13 fault. In order to keep the presented information brief in this paper, we chose Java
14 classes as the smallest entities. This can be problematic since different requirements
15 may well use the same Java classes although different methods thereof. In fact, if we
16 would have done the trace analysis by comparing overlapping methods instead of
17 classes we would not have found any erroneous results. Nonetheless, it is important
18 to understand the meaning of trace dependencies to identify false positives. Valid
19 results are crucial to support software understanding and trade-off analysis.

20 Thus far, we used our approach for consistency checking and other forms of rea-
21 soning. In that context, problems generally arise because of the lack of information
22 and not its abundance. The availability of abundant trace information provides
23 better interconnectivity among modeling artifacts and allows deeper manual in-
24 vestigation. It is generally not necessary to comprehend the complete set of trace
25 dependencies but only subsets addressing particular concerns. As pointed out above,
26 our approach also errs in producing incorrect trace dependencies at times. Here,
27 our stance is that it is generally easier to dismiss incorrect trace dependencies,
28 when encountered, instead of discovering missing ones. In the absence of a precise,
29 complete, and automated approach to generating trace dependencies, we have to
30 trade-off completeness and correctness. We believe our approach to be complete in
31 identifying all trace dependencies, however, at the expense of also producing some
32 incorrect ones. We found that it generally produces few incorrect trace dependen-
33 cies compared to the majority of correct ones. Our approach is thus most useful
34 in domains where completeness is desired (i.e., trade-off analysis, automation). In
35 domains where completeness is less important than correctness, it may not fit in as
36 well.

37 *Dealing with complexity.* Figure 3 shows that trace dependencies can get very
38 complex even in this simple example confirming the need for automation. Given
39 complete input (i.e., if the mapping between all model elements and code is known),
40 our approach is exhaustive in generating explicit trace dependencies among all ar-
41 tifact which results in a large number of (potentially incorrect) trace dependencies.
That is, if all traces between model elements and code are known then our approach

12 A. Egyed & P. Grünbacher

1 generates all trace dependencies among model elements. Since there are n^2 traces
among model elements for n traces to the code, it follows that we get n^2 results for
3 n input hypotheses.

3. Some Heuristics for Understanding Trace Dependencies

5 As illustrated so far, the main purpose of the trace analyzer is to identify trace
dependencies. These simple dependencies are already very useful for manual conflict
7 analysis or change management. For example, we know from the example in Sec. 2
that requirement [r1] depends on requirement [r6] and if requirement [r6] changes
9 then requirement [r1] is affected by this change. Even in cases of the incorrect trace
dependencies we identified above (dashed lines), this reasoning is useful since the
11 change of one requirement may unknowingly result in the change of other dependent
requirements. However, while this kind of reasoning is certainly useful, it is still
13 superficial as it says little about how exactly requirements affect one another as
pure dependencies do not convey any meaning or rationale.

15 3.1. Investigating different types of trace dependencies

The trace dependencies derived through our approach are links that merely express
17 existing relationships but human decision makers are required to understand the
true meaning of these relationships. However, upon investigating trace dependencies
19 among requirements, we found that the meaning of these relationships is highly
dependent on the types of requirements they bridge. We will discuss this finding
21 using three examples:

Trace dependency between an efficiency requirement and a functional require-
23 *ment.* An efficiency requirement (time behavior) defines a time constraint on a
(sub)system while a functional requirement defines user/customer requested ca-
25 pability. If there is a trace dependency between an efficiency requirement and a
functional requirement (e.g., the dependency from [r6] to [r1]) one may infer that
27 the execution of this particular function has to satisfy the given efficiency constraint
(e.g., the movies needs to be played within one second after selection from list).

Trace dependency between two efficiency requirements. If our approach identifies
29 a trace dependency between two efficiency requirements (e.g., the dependency from
[r6] to [r5]) one may infer that the [r5] has to be at least as efficient as [r6], i.e.,
31 loading textual information about a movie has to be at least as fast as starting to
33 play the movie. As [r5] is executed as part of executing [r6] it has to execute within
the same or even better performance.

Trace dependency between two functionality requirements. If one functional re-
35 quirement depends on another functional requirement an implication may be that
the implementation of the second functionality requirement is a pre-requisite for
37 the implementation of the first. In other words, eliminating the second requirement
is useless if the first requirement is not eliminated either. Consider, for example,
39

1 requirement [r1] “play movies automatically after selection from list” and, require-
 3 ment [r0] “download movie data on demand from server while playing” where [r1]
 5 depends on [r0]. Clearly, the second requirement needs to be implemented to sup-
 port the first one (i.e., movie data contains information about data location, formal,
 and playback properties that are needed for playing).

Table 4. Implication table (examples).

Dependency Type	Implication
efficiency $e \rightarrow$ function f	Function f has to satisfy efficiency e
efficiency $e1 \rightarrow$ efficiency $e2$	$e2$ has to be at least as efficient as $e1$
efficiency $r \rightarrow$ security s	s needs to be realizable with efficiency r
Understandability $u \rightarrow$ Recoverability r	The recovering action r should not contradict under- standability
Function $f \rightarrow$ Reliability r	f needs to be realizable within reliability r
Security $s \rightarrow$ Function f	Function f must satisfy at least security s
Security $s1 \rightarrow$ Security $s2$	$s2$ has to provide at least the level of $s1$
Function $f1 \rightarrow$ Function $f2$	Implementing function $f2$ is a pre-requisite to implement- ing function $f1$
Function $f \rightarrow$ Efficiency p	A part of the function f has to satisfy performance constraint p

7 Table 4 defines these and additional implications that are generally useful for
 interpreting the meaning of trace dependencies. However, one should be aware that
 9 there are exceptions. For example, two separate functionalities may be implemented
 “close” to one other but in a way that their execution does not affect one another.
 11 As an example, consider the requirements [r1] “play movies automatically after
 selection from list” and [r13] (not listed) “log the playing of movies in a log file”.
 13 Clearly the second requirement is executed as part of the first requirement but the
 first requirement is indifferent towards the second one.

We need to know two things for identifying meaningful trace dependencies this
 15 way (see [13] for a detailed discussion): (1) a classification of requirements (e.g,
 functional, efficiency, security) and (2) trace dependencies among requirements.
 17 Given that we have automated the second part manual effort is only required for
 classifying requirements. As such, a single requirement may be classified into an
 19 arbitrary number of categories. Although a manual activity, we found that it is
 typically easy to categorize requirements this way. Indeed, the classification of re-
 21 quirements is often a byproduct of existing requirements modeling and elicitation
 techniques [16].

23 Table 5 depicts that the implication of trace dependencies can be generalized.
 A trace dependency between a requirement and design element is such that the
 25 requirement provides rationale for the design and the design realizes the require-
 ment. A trace dependency between elements of a state diagram and those of a class
 27 diagram is such that the state elements provide the behavior of their classes while
 the classes provide the structural context of the state elements. This table is rather
 29 generic but can easily be refined.

Table 5. Implication table for other modeling artifacts (examples).

Dependency type	Implication
Requirement \rightarrow Design	Realization of requirement
Design \rightarrow Requirement	Rationale of design
Statechart \rightarrow Class	Behavior of elements
Requirement \rightarrow Code	Implementation
Design \rightarrow Code	Implementation

1 **3.2. Analyzing the degrees of overlaps**

2 Our approach to finding trace dependencies relies on finding overlaps among test
 3 scenarios that belong to requirements (or groups of requirements). If two test sce-
 4 narios for two different requirements overlap in the lines of code they execute (i.e.,
 5 their “footprints” overlap) then we assume a trace dependency. In terms of iden-
 6 tifying the implications of a trace dependency, our approach thus has a unique
 7 advantage. Depending on the degree of overlap in the lines of code executed we can
 8 strengthen and weaken the implications from Table 4.

9 Figure 4 shows how the lines of code of requirements may overlap. If there is no
 10 overlap (case 1) we conclude that there is no trace dependency. In other words, if
 11 a security requirement affects different lines of code than a functional requirement
 12 then it is safe to say that the security requirement does not apply to the functional
 13 requirement. On the other extreme, if there is complete overlap (case 4) we can
 14 deduce that the requirements describe the same part of the system (e.g., clusters in
 15 Fig. 3). In other words, if a security requirement is implemented by the same lines
 16 of code than a functional requirement then it must satisfy the security exactly; and
 17 that the security must be implemented by exactly this functionality (and no other).
 18 Both cases 1 and 4 in Fig. 4 support strong reasoning and we can create reliable
 19 trace dependencies. However, case 4 is the least likely case. Typically, scenarios
 20 do not overlap in the lines of code they execute or they overlap partially (cases 2
 21 and 3).

22 If a requirement uses a subset of the lines of code (case 3) of another one
 23 then the overlap is complete in one direction but not the other. For example, if a
 24 functional requirement uses a subset of the lines of code of a security requirement
 25 then it must fully satisfy the latter; however, the security requirement will only be
 26 implemented partially by the functional requirement resulting in a strong trace link
 27 in one direction and a weak one in the other direction. We found that many trace
 28 dependencies fall under this category. In Fig. 3, all (correct) trace dependencies
 29 (solid links) fall in this category.

30 If a requirement overlaps with another requirement but it has unique source
 31 code (case 2) then the implications we can derive are weakened. For example, a
 32 security requirement that partially overlaps with a functional requirement implies
 33 that a part of the functionality has to implement the security (which part of the




Fig. 4. Types of overlaps among requirements.




Fig. 5. Partially overlapping requirements.

1 functionality remains unknown) and it implies that a part of the security is imple-
 2 mented by the functionality (again the part is unknown).

3 Although this kind of dependency is weaker, it may still produce useful insights.
 4 Depending on how large the overlap is, we can gradually strengthen and weaken
 5 the meaning. If the functional requirement and the security requirement overlap
 6 90% (almost complete overlap; left of Fig. 5) then we can say that most of the
 7 functionality must satisfy the security constraint. Obviously, this is better than if
 8 the overlap is smaller (e.g., 20%; right of Fig. 5). We can use this degree of overlap
 9 to distinguish between more and less reliable meaning. For example, the efficiency
 10 requirements [r6] and the functional requirement [r2] partially overlap such that
 11 the efficiency requirement uses most of the same Java classes as the functional one.
 12 Due to the partial overlap, we cannot imply a precise meaning but because of the
 13 strong overlap, it is fair to say that most of [r2] has to have an efficiency of one
 14 second or less.

15 3.3. Considering groups of related requirements

16 An interesting extension of this discussion is to consider groups of requirements
 17 instead of individual requirements. This is not only important for the purpose of
 18 this paper but has relevance in practical settings. For example, in requirements
 19 negotiation [17] we need to understand the dependencies among requirements in
 20 order to allow meaningful trade-off analyses. It typically does not make sense to
 21 look at individual requirements but we usually have to build packages of related
 22 requirements in order to better handle complexity. In the context of adding meaning
 23 to trace dependencies this “grouping” of requirements has further implications.
 24 Figure 6 shows that both Security1 and Security2 partially overlap with Function1.
 25 We cannot infer which part of the function has to satisfy Security1 or Security2.




Fig. 6. Grouping of requirements.

1 But we see that a package of both security requirements together captures all of the
 2 functionality. Thus, we can conclude that all of Function1 must satisfy the weaker of
 3 the two security requirements; some of the functionality (and this part is unknown)
 4 has to satisfy the stronger of the two security requirements.

5 A similar example can be drafted with response time. If response times 1 and 2
 6 overlap with a function then we can deduce that the response time of the function-
 7 ality has to be less than the combined response times 1 and 2. Future work will
 8 investigate the effect of grouping requirements in more detail.

9 4. Case Study: ArgoUML

10 We performed a thorough case study for validating the trace analyzer. The software
 11 system we have chosen for our study is ArgoUML, an open source software design
 12 tool supporting the Unified Modeling Language (UML). It is written entirely in
 13 the Java programming language. The size of this software is significant contain-
 14 ing over 1300 classes distributed in over 70 packages. The code base has over 200
 15 KLOC. The entire source code and documentation for ArgoUML are available from
 16 <http://www.argouml.org>. A small subset of the ArgoUML requirements is summa-
 17 rized in Table 6.

18 The fundamental capability of ArgoUML is to support software modeling with
 19 the UML. Currently, the tool supports 8 of 9 UML diagrams. The user interface
 20 is quite intuitive and similar to other case tools for the UML. It has three major
 21 areas (see Fig. 7): a diagramming area for creating and modifying graphical symbols
 22 of UML modeling elements, a model explorer for navigating in the evolving UML
 23 model, and an editor for viewing and modifying the properties of individual UML
 24 elements. ArgoUML also provides some innovative features for maintaining todo
 25 lists of open design issues, and a critiquing feature for automatically suggesting
 26 improvements to the design.

27 4.1. Case study process

28 The approach we took for the case study was to compare Trace Analyzer's capa-
 29 bilities to those of a human expert and was carried out in the following steps:

- 30 • *Selection and preparation of requirements.* The ArgoUML documentation has
 31 some sections about the tool's requirements. Although only a few requirements

Table 6. Selected ArgoUML requirements.

R1	The system shall allow creating a class in the current active diagram
R3	The system shall support creating a parent of the currently selected class
R5	The system shall support creating an association with a class from the currently selected class to the new class
R7	The system shall support adding an attribute to the currently selected class
R9	The system shall support creating an association between two existing classes
R14	The system shall allow to display an existing class in multiple diagrams
R17	The system shall support to use the clipboard when working with UML elements
R18	The system shall allow to find a class and navigate to it
R19	The system shall provide a zoom capability
R23	The system shall to maintain a todo-list of modelling tasks
R24	The system shall allow to load the model from a file
R26	The system shall allow to export the model to a file in XML document format
R27	The system shall display a system information
R29	The system shall allow to create Java code for the modelled class diagram
R30	The system shall automatically critique the evolving model and provide suggestions for improving upon request
R31	The system shall allow to change the properties of a class
R33	The system shall support UML use case diagrams
R34	The system shall support UML state diagrams

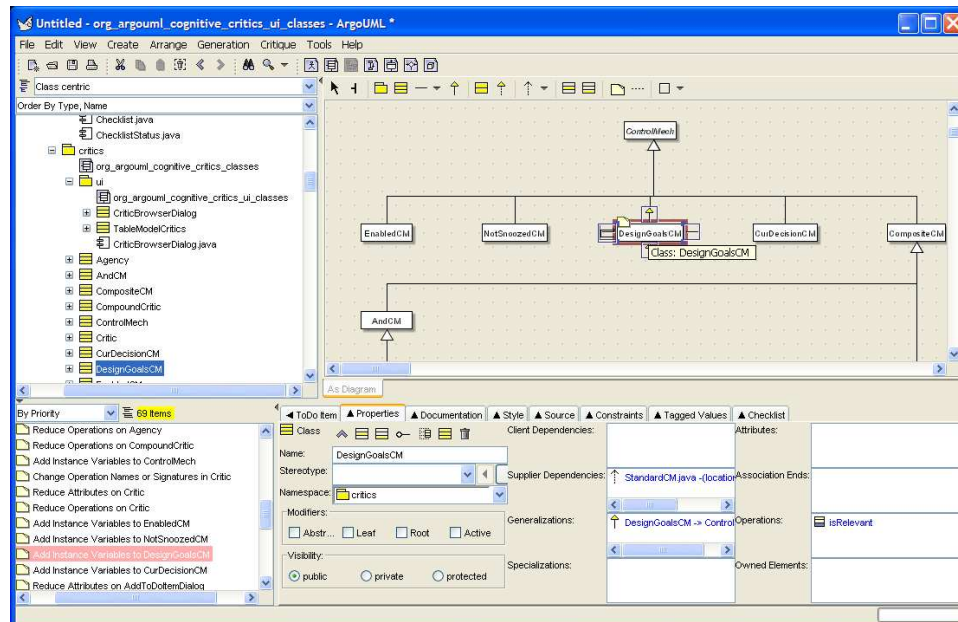


Fig. 7. Snapshot of the ArgoUML user interface showing the model explorer, diagram editor, todo list, and property editor.

1 are documented in the developer documentation, there are further sources for
 3 ArgoUML's requirements including a list of features, and some section in the
 5 user manuals and tutorials. In a first step we analyzed these documents and
 selected and compiled a list of requirements suitable for our purpose. In total 34

- 7 • *Identification of trace links.* The next two steps were done in parallel by an engi-
 9 neer using Trace Analyzer and by a human expert identifying trace links manu-
 11 ally. As described in the preceding sections the engineer using the Trace Analyzer
 defined scenarios and linked them to the identified requirements. He then used
 the tool to identify the traceability links. In parallel the human expert analyzed
 13 the requirements manually and enumerated a list of depending requirements for
 each of the 34 requirements.
- *Comparison and interpretation of results.* In the third step we compared the
 results of the tool supported trace link creation with the manual approach.

15 4.2. Results

17 Figure 8 shows all requirements to code dependencies identified by the trace ana-
 19 lyzer. Row headings are the selected ArgoUML requirements (r01 to r34). Column
 21 headings are different regions of the ArgoUML code (01-61) that were found to be
 23 executed by the chosen test scenarios (those regions represent aggregations of the
 1300 classes of the ArgoUML which would be impossible to visualize here individ-
 ually). A total of 196 trace links were identified by the trace analyzer based on
 an input of only 34 trace links. Figure 8 shows the identified links between the 34
 investigated ArgoUML requirements and 61 source code elements.

25 We compared the trace links identified by the human expert and the tool-
 27 generated results and derived the following metrics for each requirement.

TATL: The number of trace links found by the trace analyzer.

27 *HETL:* The number of trace links found by the human expert.

29 *OVLV:* Overlapping results, i.e. the number of requirements found by both Trace
 Analyzer and the human expert.




Fig. 8. Requirements to source dependencies created by Trace Analyzer.

Table 7. Requirements to requirements dependencies.

	r01	r02	r03	r04	r05	r06	r07	r08	r09	r10	r11	r12	r13	r14	r15	r16	r17
TATL	16	2	12	12	9	4	12	15	9	9	10	10	12	18	16	11	13
HETL	9	6	7	7	6	3	2	2	6	6	6	2	2	2	2	1	0
OVLP	4	0	7	7	6	1	1	1	6	6	6	2	2	2	1	1	0
TAON	9	2	5	5	3	2	11	14	3	3	3	8	10	16	15	10	13
HEON	5	6	0	0	0	2	1	1	0	0	0	0	0	0	1	0	0
CONS	44%	0%	100%	100%	100%	33%	50%	50%	100%	100%	100%	100%	100%	100%	50%	100%	N/A
POWR	1,8	0,3	1,7	1,7	1,5	1,3	6,0	7,5	1,5	1,5	1,7	5,0	6,0	9,0	8,0	11,0	N/A
	r18	r19	r20	r23	r24	r25	r26	r27	r28	r29	r30	r31	r32	r33	r34		Total
TATL	15	0	5	3	23	20	4	0	3	12	6	18	26	17	12		354
HETL	2	0	0	0	2	2	2	0	0	3	0	2	6	2	3		93
OVLP	0	0	0	0	2	2	2	0	0	2	0	1	6	1	1		70
TAON	15	0	5	3	23	20	2	0	3	10	6	17	20	16	11		283
HEON	2	0	0	0	0	0	0	0	0	1	0	1	0	1	2		23
CONS	0%	N/A	N/A	N/A	100%	100%	100%	N/A	N/A	67%	N/A	50%	100%	50%	33%		75%
POWR	7,5	N/A	N/A	N/A	11,5	10,0	2,0	N/A	N/A	4,0	N/A	9,0	4,3	8,5	4,0		3,8

1 *TAON*: Number of trace links found by the trace analyzer only.

HEON: Number of trace links found by the human expert only.

3 These metrics allowed us to compute two further indicators for comparison.

5 *CONS*: The consistency between the human expert and Trace Analyzer was
 5 derived by computing what percentage of trace links found by the human expert
 was also found by Trace Analyzer.

7 *POWR*: The power of Trace Analyzer indicates by which factors it outperforms
 the human experts in terms of number of trace links identified.

9 Table 7 shows the identified trace dependencies among the selected ArgoUML
 requirements. A total of 354 links were identified by Trace Analyzer, the human
 11 expert came up with 93 links. 70 of these links were also identified by the Trace
 Analyzer, so the consistency of results is about 75%. The trace analyzer approach
 13 identified four times more traces than the human expert and it did so with signif-
 icantly less effort. In addition, it must also be noted that using the trace analyzer
 15 approach not only produced requirements to requirements traces but it also resulted
 in requirements to code traces (i.e., the human expert did not produce these).

17 Of particular interest are the traces identified by the human expert but not
 the tool (HEON). These could indicate incompleteness or incorrectness on part of
 19 Trace Analyzer which is only possible if the test scenarios defined for the require-
 ments were incomplete or incorrect. Or these could be erroneous trace dependencies
 21 generated by the human expert.

5. Benefits and Limitations

23 Automated Requirements Traceability is an important means to facilitate commu-
 nication among the success-critical stakeholders, to ease determining the impact
 25 of changes and support their integration, to preserve knowledge and dependencies
 created during the design process, to assure quality, and to prevent misunderstand-
 27 ing. This section will discuss benefits of our automated approach. We also present
 limitations and potential problems.

1 5.1. Benefits

3 *Understanding requirements origins and rationale.* Traceability between stakeholder
5 needs and requirements can be detected manually but our automated technique pro-
7 vides a more complete traceability. Trace analysis can derive missing relationships
9 between informal user needs and existing design elements. For example, the auto-
mated technique can help to create traceability links from new stakeholder needs
to existing design: In one of our experiments a link from the new user need “Users
should be able to capture movie screen snapshot at any time” to the “Pause button”
GUI element was automatically derived. This gives rationale and explains why an
element is here by providing backward traceability.

11 *Traceability to non-functional requirements.* Using this approach even non-
13 functional requirements can be linked to model elements or code sections. Non-
functional stakeholder needs ultimately always result in some code although this
15 relationship is typically almost impossible to identify. For example, we know that
implementation class [U] exists because of [r7] and [r8]. The requirements “Three
17 seconds max to load textual information about a movie” can be linked to imple-
mentation classes [N,R]. By generating missing trace information the trace analyzer
19 technique can link a new non-functional user need “Novices should be able to use
the most important functions without training” to requirements [r1], [r3], [r8], [r9]
and thereby also show all affected implementation classes.

21 *Identification of conflicting requirements.* It is typically hard to derive all de-
23 pendent requirements because of scalability issues. An automated approach towards
generating dependencies between requirements is thus critical to determine whether
dependent requirements are consistent. For example, the requirement “novices
25 should be able to use system without training” may be in conflict with the require-
ment “3 seconds response time” because such a long delay would not be intuitive.
27 Our approach cannot automatically derive conflicts, but by finding all possible de-
pendencies it is easier to identify potential inconsistencies and conflicts.

29 *Identifying conflicts among requirements.* The following example from VOD il-
31 lustrates the use of a trace dependency during requirements conflict analysis. One
can observe through Table 3 that [r6] depends on [r5] given that [r5] traces to the
Java classes [N,R] (a subset of [A,C,D,F,G,I,J,K,N,O,R,T,U]). This dependency im-
33 plies that in order to start playing a movie one needs to load the textual information
about a movie. The problem is that loading this information is allowed to take up
35 to three seconds which is longer than the allowed 1 second max to start playing
that movie. The finding of this trace dependency implies a conflict between two
37 requirements (see Fig. 9). To a casual observer, this conflict would have been hard
to identify without the help of trace dependencies because it is not obvious that
39 both requirements are related in the lines of code they execute. Once identified, a
potential solution to this conflict is to change requirement [r6] to complete in less
41 than three seconds also.




Fig. 9. Trace dependency leads to a requirements conflict.

Table 8. New/changed requirements.

r3	3 seconds max to start playing a movie
r10	Avoid image degradation caused by temporary network-load fluctuations

Table 9. Artifact to Java class dependencies (update).

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
r10	F			F			F		F		F				F						

1 *Determine impact of new or changed requirements.* The Trace Analyzer technique can also help in analyzing the impact of new or changed requirements, which are common in iterative software processes [18] or in software evolution and maintenance. Normally, it is desirable to validate a new requirement or a modification of an existing requirement prior to implementation. In such a case one can still use the trace analyzer by hypothesizing about the impact of a new requirement or a requirement change. The following discusses one case of adding a new requirement that has not yet been implemented in the VOD system.

9 Figure 9 shows the changed requirement [r6] due to the conflict with [r5] and it also shows a new requirement that deals with image degradation because of network fluctuations. Recall that the VOD system is a video-on-demand system that starts playing a movie as soon as data arrives via the network. If temporary network congestions cause delays then this may negatively affect image quality. A possible approach to satisfy requirement [r10] would be to do some initial caching to overcome this limitation. To find out whether this new requirement clashes with other existing requirements, we can do a preliminary trace analysis. To do this, we hypothesize about the impact of the new requirement and presume that caching can be done solely by modifying the Java classes [A,D,G,I,K,O] plus adding some new ones. We thus define a new, hypothetical trace dependency between [r10] and [A,D,G,I,K,O,+] and repeat our trace analysis with this additional data.

21 If the new hypothesized trace dependency is compared with the other known trace dependencies in Table 9 then we can again determine trace dependencies between [r10] and other artifacts based on their overlapping use of common code. For instance, one can tell that the new requirement [r10] uses a subset of the code that the state [s9] uses. As a result, [s9] fully depends on [r10]. In the following, we are more interested in the trace dependencies among the new requirement [r10] and the other requirements (e.g., [r3], [r6]) as depicted in Table 9.

22 A. Egyed & P. Grünbacher

1 Of particular interest is the previously modified requirement [r6] which depends
 3 fully on [r10]. Recall that [r6] was changed because one second response time was
 5 considered insufficient given that at least three seconds are needed to load textual
 7 information about a movie. We thus relaxed the one second constraint to three
 9 seconds. However, now we see that if requirement [r10] gets implemented there will
 11 be additional delays. A pre-caching of a movie can only happen once movie details
 13 are known. A caching period of 1–2 seconds thus adds to the three seconds already
 15 needed to load and play a movie.

17 This is a conflict that can be identified with ease once one is aware of the trace
 19 dependency. This example again shows that our trace analyzer approach can help
 in pinpointing non-obvious dependencies between artifacts. If those dependencies
 lead to the identification of conflicts then the trace analyzer can further help in
 evaluating potential solutions. In this case, a potential solution is not to pre-cache
 movie data before playing but to incrementally build up a cache while playing. We
 presume that movies download faster than they are played and, consequently, it
 is possible to build up and increase the cache while playing. This solution might
 not be as effective as pre-caching but will no longer conflict with the performance
 requirement [r6] (i.e., note that this changed requirement would execute a different
 part of the system). Figure 10 visualizes the conflict and its potential solution.

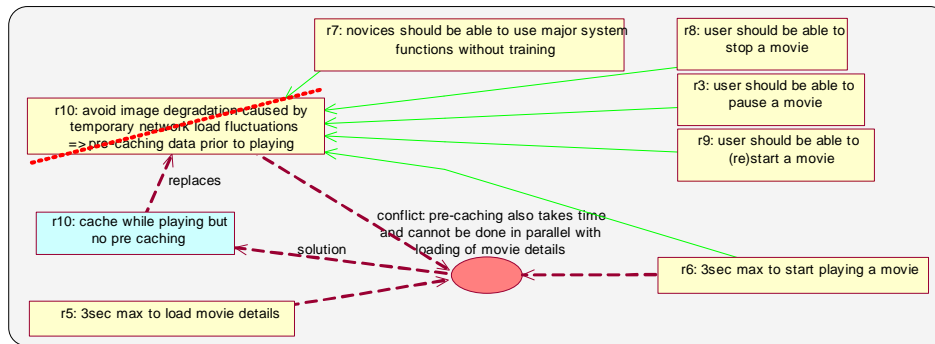


Fig. 10. Trace dependency leads to a conflict with new requirement

21 *Traces between requirements and design.* Besides finding trace dependencies be-
 23 tween different requirements, the trace analyzer technique also finds dependencies
 25 between requirements and design elements. For instance, the requirement [r0] “dis-
 27 play and select movie from list” depends on the state transition [s3] because [s3]
 may at most relate to [C,J,R,U] whereas [r0] is known to relate to [C,J,N,R,U] (a
 superset). Using the same method, one may identify many more trace dependencies.

Verification of requirements. An important task of a software engineer is to
 determine whether the requirements have been realized properly. We specify accep-
 tance test through scenarios for all requirements. We can thereby make sure that

1 scenarios sufficiently cover requirements. The case study shows that we tested all
requirements and know what sections of the code realize them.

3 *Identification of missing requirements.* The approach can also be used to identify
missing requirements. For example, the analysis reveals that we have no require-
5 ments defining scenarios for “S” or “P”. Does this mean that the system implements
something not stated in the requirements? Through the generated trace dependen-
7 cies we know that implementation class “S” is about [s5] and [s7] (selecting server)
and we can now reason that no requirement was defined that allows the user to
9 change servers. Besides detecting missing requirements, we can also reason about
missing or incomplete designs. For instance, we find that design element [s1] was
11 not defined in any requirement or any implementation.

13 *Determination of change impact.* Assume that the response time in requirement
“3 second max to start playing a movie” has to be reduced. Trace analysis reveals
15 the impact of such a change onto other requirements, onto design, and onto code.
But trace analysis also reveals the reverse impacts. For example, we know without
any manual creation of trace information that if code element [T] (Video.java)
17 changes then the design elements [s11,s12] are affected by that change.

19 *Understanding the level of strength of dependencies.* New requirements are an
interesting case for deriving trace dependencies where parts of the system have
not even been built. Section 3.5 discussed that by hypothesizing what model ele-
21 ments/code might be affected by a new requirements the trace analyzer can predict
which requirements and other development artifacts might be affected. It must be
23 noted that our technique can determine strength of dependencies where strength is
defined in terms of how many classes (or methods or lines of code) two artifacts have
25 in common. For instance, it can be observed that that [r3] uses 33% of the classes
of [r0] and [r0] uses 22% of the classes of [r3] (the percentage applies to the number
27 of overlapping classes versus total classes). Although the dependency between [r3]
and [r0] is not very strong it still implies that a change in [r3] has a 33% chance
29 that it will also affect [r0]. This percentage of course presumes that all classes are
of equal size which they are not. For more precise dependency numbers, the trace
31 analysis could be conducted on methods or lines of code. Note that the strength of
a dependency is not to be confused with the confidence in a dependency. Whereas
33 the confidence (full/partial) defines the likelihood of false positives, strength simply
describes the degree of overlaps.

35 Upon inspection of the generated traces between requirements, we find that most
requirements trace to most other requirements at least partially. This is not very
37 surprising since requirements tend to be very generic descriptions. For a more useful
determination of trace dependencies between requirements one should focus more on
39 the extremes, i.e., 0% and 100%. If there is 0% overlap between two requirements
then there is no dependency between them. The requirement [r3] (pause movie)
41 has nothing in common with requirement [r4] (three seconds max to load movie
list). If there is a 100% overlap between two requirements then there is a strong
43 dependency between them. For instance, [r6] uses 100% of [r5] which implies a
strong dependency.

1 *Distinguishing domain specific code vs. generic code.* The approach can also be
used to distinguish project or domain-specific code from generic code. For specific
3 analyzes it might be necessary to ignore generic code since it is likely to be used
for different purposes and may obscure analysis. For instance, two classes may use
5 a common third class to create and modify a file but this does not mean those two
classes are related to one another (note: those two classes may be related if they
7 modify the same file but the trace analyzer approach cannot detect that).

Determining artifacts needing attention. Special care has to be spent on very
9 complex and/or very important artifacts. The importance of a development arti-
fact depends on how many other artifacts it constrains (e.g., design element s9 is
11 important in that it (partially) defines 8 implementation classes). The complexity
of an artifact depends on how many other artifacts constrain it (e.g., design ele-
13 ment s9 is also complex since 6 out of the 10 requirements impose themselves on
it). Trace analysis can simply pinpoint these kinds of metrics, e.g., for complexity
15 versus importance trace offs.

Balancing granularity of requirements. In an early project stage requirements
17 will be typically fairly generic; later on requirements will be more specific unless, of
course, a major change comes along. For instance, [r5] is a lower-level requirement
19 than [r6] because [r5] uses a subset of the code than [r6] does. Trace analyzer can
find requirements that are very generic (e.g., they affect many classes). This can
21 assist the engineer to balance out requirements by increasing precision.

Cost and Effort in Computing the Input Required. Our approach is not free but
23 requires the engineer to define input hypotheses on how model elements (artifacts)
relate to some common representation (e.g., source code). Yet, this activity is man-
25 dated in standards and is often performed by engineers to proof implementation.
We also demonstrated that test scenarios may be used to ease the model element to
27 code mapping if available. Again, engineers are required to test their systems and
the overhead required to observe the test executions is small. Thus, if the source
29 code and test scenarios are available then the cost of using our approach is solely
the hypotheses on how test scenarios relate to model elements. Even this task is
31 supported by our approach in that we allow the engineer to group model elements
(e.g., [a,b] is [1,2] is easier to define than the exact value for [a] and [b] separately)
33 or use uncertainties ([a,b] isAtLeast/isAtMost/isNode/IsExactly [1,2] as defined in
[11]). Yet, in return, our approach computes trace dependencies among all model el-
35 ements. Since there are over n^2 such potential dependencies, our approach computes
a quadratic number of output traces for a linear number of input hypotheses.

37 5.2. Limitations

Dependent on high-quality input. The trace analyzer relies on the capability of a
39 software engineer to relate the test scenarios to some requirements and model el-
ements. Three errors are possible that may impact the trace analysis in different
41 ways: (1) the engineer omits a link between a test scenario and a requirement,

1 (2) the engineer creates a wrong link, or (3) there is a mismatch between a require-
3 ment and the specified tests (for example, the test case only exercise the wrong
or only a partial functionality). Although the technique has some means of detect-
ing inconsistencies among links it can be fooled this way and engineers need to be
5 careful when doing their specifications.

Handling of shared code. A shared code is a part of the source code that is exe-
7 cuted by two or more requirements but should not be considered an overlap during
the trace analysis. For example, during the trace analysis of the ArgoUML case
9 study, we found that a range of user interface classes existed that were triggered
(executed) during testing but did not relate to the test scenario at hand (i.e., simply
11 hovering over a user interface icon causes lines of code being executed). It is impor-
tant to identify shared code to reduce the number of false positives (see also [11]).
13 In fact, during the ArgoUML case study we've updated some Trace Analyzer capa-
bilities to allow better identification of shared code. These recent extensions allow
15 to better eliminate shared from the analysis and helps to significantly reduce the
amount of "noise" generated.

Understanding of granularity. The more granular the trace analysis, the more
17 test scenarios are needed to ensure that all parts of the source code are executed
that are related (i.e., belong to a given requirement). This effort can be reduced
19 by performing the trace analysis on less granular information (e.g., classes instead
of methods). The downside of less granular trace analysis is that more overlaps
21 exist (e.g., two requirements may use different methods of the same class and thus
overlap in the common use of the same class), thus leading to more false positives.
23 In our future work we will investigate this issue in more detail.

25 6. Related Work

Different approaches have been developed to automate the acquisition of trace
27 information. Typically these approaches support the creation or recovery of traces
between different artifacts (e.g., between design and code, code and documentation,
29 requirements and architectures).

Antoniol *et al.* discuss a technique for automatically recovering traceability links
31 between object-oriented design models and code based on determining the similar-
ity of paired elements from design and code [19]. Basic class attributes are used as
33 traceability anchors. The focus of this work is, however, not to support trace de-
pendencies between requirements and code. Murphy *et al.* [20] aim at automating
35 the identification of links between high-level models and source code. Their ap-
proach uses software reflexion models to find out whether an engineer's high-level
37 model agrees with and where it differs from the source. While our approach to
identifying requirements traceability is similar to design traceability [12], RT has
39 a range of special considerations: requirements are often captured informally but
they are often categorized into functional and non-functional groups (i.e., qualities).
41 While notations exist that express hierarchies and data/control dependencies among

1 design elements, such notations are typically not used for requirements. A particular
2 focus of our work is thus on the semantic implications of requirements, their hier-
3 archies, and dependencies. Antoniol *et al.* describe an approach to automatically
recovering trace information between code and documentation [21].

5 Gruber *et al.* discuss the problems of design rationale capture and demand
the need for automatically inferring rationale information [10] from background
7 knowledge and information captured during design. Their approach emphasizes
design dependency management and rationale by demonstration. One of their key
9 observations is related to our paper stating that rationales are not just statements
of fact, but explanations about dependencies among facts.

11 Many approaches discuss specific traceability issues without particularly focus-
ing on automation: Arlow *et al.* emphasize the need to establish and maintain
13 traceability between requirements and UML design and present Literate Model-
ing as an approach to ease this task [22]. Gotel and Finkelstein extend the view of
15 artifact based RT and focus on understanding the social network of people that con-
tributed to the development of requirements [23]. Pohl *et al.* describe an approach
17 based on scenarios and meta-models to bridge requirements and architectures [7].
Grünbacher *et al.* discuss the CBSP approach that improves traceability between in-
19 formal requirements and architectural models by developing an intermediate model
based on architectural dimensions [6].

21 Other traceability approaches also emphasize the automation aspect. Zisman
and her colleagues have presented a rule-based approach for automatically generat-
23 ing and maintaining traceability relations. The artifacts and rules are described in
XML and supported by a prototype tool [24]. The approach has also been applied
25 to organizational models specified in i^* and software systems models represented
in UML [25].

27 Our trace analyzer approach generates traceability based on source code already
available. A forward engineering approach complementing our approach is taken in
29 the context of program synthesis by Richardson and Green [26]. The approach helps
to automatically derive traceability relations between parts of a specification and
31 parts of the synthesized program. The generality of the technique is demonstrated
by applying it to the synthesis of Kalman Filter programs from specifications using
33 a program synthesis system, and generation of assembly language programs from
C source code using the GCC C compiler.

35 7. Conclusions and Further Work

37 In this paper we presented an approach supporting the automated generation of
trace information. We discussed the approach in the context of a video-on-demand
39 system and showed that it automates the generation of trace dependencies between
the different models and artifacts of the system. We then discussed how the derived
traces can support engineers in understanding software. A major strength of the
41 approach is that it creates many non-obvious dependencies allowing more thorough

1 reasoning and pinpointing of non-standard situations. A key contribution of our
2 approach is that it reduces the enormous effort and complexity of acquiring traces by
3 automatically deriving trace information from a small set of obvious hypothesized
4 traces. This leads to more complete traces and the full potential of RT can be
5 exploited: For example, traces to pre-requirements explaining where requirements
6 come from or traces from/to non-functional requirements are typically difficult to
7 create and maintain using manual approaches. The automated approach also creates
8 traces that engineers typically could not anticipate. This improves the applicability
9 of our approach in different contexts or non-standard engineering problems.

10 Further work will concentrate on developing automated support assisting en-
11 gineers in exploring and using the automatically derived trace dependencies. For
12 example, by highlighting artifacts and situations that require special attention.
13 Another thread of our research will focus on experimenting with different levels
14 of granularity of coverage measurement: The technique allows specifying this level
15 arbitrarily (e.g., class, method, or statement). We aim at developing heuristics al-
16 lowing software engineers to determine the optimum level of granularity in a given
17 situation. We also intend to apply our technique and findings to other large-scale
18 systems.

19 References

- 20 1. O. C. Z. Gotel and A. C. W. Finkelstein, An analysis of the requirements traceability
21 problem, in *1st Int. Conf. on Requirements Engineering*, 1994, pp. 94–101.
- 22 2. B. Ramesh and M. Jarke, Toward reference models for requirements traceability, *IEEE*
23 *Trans. on Software Engineering* **27**(4) (2001) 58–93.
- 24 3. N. Medvidovic, P. Grünbacher, A. Egyed, and B. W. Boehm, Bridging models across
25 the software lifecycle, *Journal of Systems and Software* **68**(3) (2003) 199–215.
- 26 4. B. W. Boehm, Software risk management: Principles and practices, *IEEE Software*
27 **8**(1) (1991) 32–41.
- 28 5. ISO/IEC-9126, Information technology — software product evaluation — quality
29 characteristics and guidelines for their use, Technical report, 1991.
- 30 6. P. Grünbacher, N. Medvicovic, and A. Egyed, Reconciling software requirements and
31 architectures with intermediate models, *Journal on Software and System Modeling*,
2003.
- 32 7. K. Pohl, M. Brandenburg, and A. Glich, Integrating requirement and architecture
33 information: A scenario and meta-model based approach, in *REFSQ Workshop*, 2001.
- 34 8. INCOSE, Requirements management tool survey, online at <http://www.incose.org>.
35 Technical report, INCOSE.
- 36 9. B. Ramesh, C. Stubbs, and M. Edwards, Lessons learned from implementing require-
37 ments traceability, *Crosstalk — Journal of Defense Software Engineering* **8**(4) (1995)
38 11–15.
- 39 10. T. R. Gruber and D. M. Russell, *Generative Design Rationale*, Lawrence Erlbaum
40 Associates, 1994.
- 41 11. A. Egyed, Resolving uncertainties during trace analysis, in *Proc. 12th ACM SIGSOFT*
42 *Symposium on Foundations of Software Engineering (FSE)*, Irvine, CA, November
43 2004, pp. 3–12.
- 44 12. A. Egyed, A scenario-driven approach to trace dependency analysis, *IEEE Trans. on*
45 *Software Engineering* **29**(2) (2003) 116–132.

28 A. Egyed & P. Grünbacher

- 1 13. A. Egyed and P. Grünbacher, Identifying requirements conflicts and cooperation: How
quality attributes and automated traceability can help, *IEEE Software* **21**(6) (2004).
- 3 14. A. Egyed and P. Grünbacher, Automating requirements traceability: Beyond the
record and replay paradigm, in *Proc. 17th IEEE Int. Conf. on Automated Software
5 Engineering (ASE02)*, Edinburgh, 2002, IEEE CS, pp. 163–171.
- 7 15. A. Egyed and P. Grünbacher, Towards understanding implications of trace dependen-
cies among quality requirements, in *2nd Int. Workshop on Traceability in Emerging
Forms of Software Engineering*, eds. G. Spanoudakis and A. Zisman, Montreal, 2003.
- 9 16. S. Robertson and J. Robertson, *Mastering the Requirements Process*, Addison-Wesley,
1999.
- 11 17. B. W. Boehm, P. Grünbacher, and R. O. Briggs, Developing groupware for require-
ments negotiation: Lessons learned, *IEEE Software* **18**(3) (2001) 46–55.
- 13 18. B. W. Boehm, A. Egyed, J. Kwan, D. Port, A. Shah, and R. Madachy, Using the
winwin spiral model: A case study, *IEEE Computer* (7) (1998) 33–44.
- 15 19. G. Antoniol, B. Caprile, A. Potrich, and P. Tonella, Design-code traceability recovery:
Selecting the basic linkage properties, *Science of Computer Programming* **40**(2–3)
17 (2001) 213–234.
- 19 20. G. C. Murphy, D. Notkin, and K. Sullivan, Software reflexion models: Bridging the
gap between source and high-level models, in *Proc. Third ACM SIGSOFT Symposium
on the Foundations of Software Engineering*, ACM, New York, 1995, pp. 18–28.
- 21 21. G. Antoniol, G. Canfora, A. De Lucia, and G. Casazza, Information retrieval models
for recovering traceability links between code and documentation, in *Proceedings of
23 the International Conference on Software Maintenance*, 2000, p. 40.
- 25 22. J. Arlow, W. Emmerich, and J. Quinn, Literate modelling — capturing business
knowledge with the uml, in *UML'98: Beyond the Notation*, 1998.
- 27 23. O. Gotel and A. Finkelstein, Extended requirements traceability: Results of an in-
dustrial case study, in *Proc. 3rd Int. Symposium on Requirements Engineering RE97*,
IEEE CS Press, 1997, pp. 169–178.
- 29 24. G. Spanoudakis, A. Zisman, E. Pérez-Minana, and P. Krause, Rule-based generation
of requirements traceability relations, *J. Systems and Software* **72**(2) (2004) 105–127.
- 31 25. F. G. Cysneiros, A. Zisman, and G. A. Spanoudakis, Traceability approach for i*
and uml models, in *Workshop Report (SELMAS 03): Software Eng. for Large-Scale
33 Multi-Agent Systems*, 2003.
- 35 26. J. Richardson and J. Green, Automating traceability for generated software artifacts,
in *ASE*, 2004, pp. 24–33.