

Software Process Validation: Quantitatively Measuring the Correspondence of a Process to a Model

JONATHAN E. COOK

New Mexico State University

and

ALEXANDER L. WOLF

University of Colorado

To a great extent, the usefulness of a formal model of a software process lies in its ability to accurately predict the behavior of the executing process. Similarly, the usefulness of an executing process lies largely in its ability to fulfill the requirements embodied in a formal model of the process. When process models and process executions diverge, something significant is happening. We have developed techniques for uncovering and measuring the discrepancies between models and executions, which we call *process validation*. Process validation takes a process execution and a process model, and measures the level of correspondence between the two. Our metrics are tailorable and give process engineers control over determining the severity of different types of discrepancies. The techniques provide detailed information once a high-level measurement indicates the presence of a problem. We have applied our process validation methods in an industrial case study, of which a portion is described in this article.

Categories and Subject Descriptors: D.2.6 [**Software Engineering**]: Programming Environments; K.6.3 [**Management of Computing and Information Systems**]: Software Management—*software development; software maintenance*

General Terms: Management

Additional Key Words and Phrases: Balboa, process validation, software process, tools

This work was supported in part by the National Science Foundation under grants CCR-93-02739 and CCR-9804067, and the Air Force Materiel Command, Rome Laboratory, and the Defense Advanced Research Projects Agency under Contract Number F30602-94-C-0253. The content of the information does not necessarily reflect the position or the policy of the U.S. Government, and no official endorsement should be inferred.

Authors' addresses: J. E. Cook, Department of Computer Science, New Mexico State University, Las Cruces, NM 88003; email: jcook@cs.nmsu.edu; A. L. Wolf, Software Engineering Research Laboratory, Department of Computer Science, University of Colorado, Boulder, CO 80309-0430; email: alw@cs.colorado.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1999 ACM 1049-331X/99/0400-0147 \$5.00

1. INTRODUCTION

Whenever a model of a system is created, the question arises as to whether that model faithfully captures the system. In software process research, where the model is typically embedded and executed within an automated software engineering environment [Garg and Jazayeri 1996], this question is avoided; the model and process are necessarily in agreement because the model *becomes* the process.

When applied in software process practice, however, this approach suffers from a fundamental flaw. In particular, it assumes that virtually the entire process is executed within the context of the environment. In fact, critical aspects of the process occur off the computer and, therefore, not under the watchful eye of the environment [Sutton 1991; Wolf and Rosenblum 1993a; 1993b]. That being the case, there is no effective way to enforce the process using this approach nor to guarantee the mutual consistency of a process model and a process execution. Moreover, deviations from the process are naturally to be expected [Cugola et al. 1996; Ellis et al. 1995; Grudin 1990], and indeed arise from the basic differences between the machine and the world [Emmerich et al. 1997].

Even if one could completely enforce a process, there still remains the issue of managing change in a process, which might also lead to a discrepancy between the model and the execution. There has, in fact, been considerable work that addresses process evolution [Bandinelli et al. 1993; Jaccheri and Conradi 1993]. Commensurate with the historical approach mentioned above, that work is concerned more with the problem of effecting changes to a process model used for automation, than it is with the problem of uncovering inconsistencies between the model and the execution.

We have developed techniques for detecting and characterizing differences between a formal model of a process and the actual execution of the process. We refer to this activity as *process validation* [Cook and Wolf 1994]. The techniques are neutral with respect to the correctness of the model ("*Does our model reflect what we actually do?*") and the correctness of the execution ("*Do we follow our model?*"). The process engineer has ultimate responsibility for making the appropriate determination of whether a problem lies within the model or within the execution, based on the particular inconsistency uncovered. The validation techniques have been implemented in a prototype tool, which has been used as part of an industrial process data analysis case study [Cook et al. 1998].

Process validation serves several purposes. For one, confidence in a formal process model is raised when it can be shown that the process execution is consistent with the behavior described by the model. This, in turn, raises confidence in the results of any analyses performed on the formal model. For another, process validation can be used as a process enforcement tool, uncovering differences between intended behavior and actual behavior. It is potentially a more flexible enforcement tool than others proposed, since it can accommodate the unavoidable, yet necessary, local perturbations in a process. Finally, process validation can reveal

where a process may need to actually evolve to accommodate new project requirements and activities.

The techniques borrow from various areas of computer science, including distributed debugging, concurrency analysis, and pattern recognition. The techniques go further than simply detecting an inconsistency; they provide a measure of that inconsistency. We believe that developing metrics for process validation is critical because the highly dynamic and exceptional nature of software processes means that simple yes-no answers carry too little information about the significance of any given inconsistency. Managers need to understand where an inconsistency occurs and how severe that inconsistency might be before taking any corrective action.

The next section presents the framework in which this work is cast. Section 3 then states the process validation problem and outlines our approach. Section 4 defines the validation metrics, and Section 5 presents example uses of the metrics. The issue of deriving a characteristic behavior from a process model is discussed in Section 6. Section 7 presents a short review of an industrial case study in which the validation techniques were successfully applied. Section 8 summarizes work related to process validation. Finally, we describe our implementation of a validation tool and discuss some ideas for future work in Section 9.

In general, this article focuses on the theoretical foundations of our validation techniques, and the practical implementation of those techniques in a tool. This focus foregoes the presentation of a broad methodology for applying process validation within a process improvement effort, although the basic approach should be discernable from the examples. We refer the reader elsewhere for a detailed treatment of methodological issues [Cook et al. 1998].

2. AN EVENT-BASED FRAMEWORK FOR PROCESS VALIDATION

The foundation on which our process validation work rests is a view of processes as a sequence of actions performed by agents, either human or automaton, possibly working concurrently. With this, we are taking a decidedly behavioral view of processes, because we are interested in the dynamic activity displayed by the processes, rather than, say, the static roles and responsibilities of the agents or the static relationships among components of the products. This does not mean that other aspects of a process are not worthy of study; it is just that the issues we have chosen to investigate are those having to do with behavior rather than structure.

2.1 Events

Following Wolf and Rosenblum [1993a], we use an event-based model of process actions, where an *event* is used to characterize the dynamic behavior of a process in terms of identifiable, instantaneous actions, such as invoking a development tool or deciding upon the next activity to be performed. The use of events to characterize behavior is already widely accepted in other areas of software engineering, such as program visualiza-

tion [LeBlanc and Robbins 1985], concurrent-system analysis [Avrunin et al. 1991; Cook and Wolf 1998c], and distributed debugging [Bates 1989; Cuny et al. 1993].

The “instant” of an event is relative to the time granularity that is needed or desired; thus, certain activities that are of short duration relative to the time granularity are represented as a single event. An activity spanning some significant period of time is represented by the interval between two or more events. For example, a meeting could be represented by a “begin-meeting” event and “end-meeting” event pair. Similarly, a module compilation submitted to a batch queue could be represented by the three events “enter queue,” “begin compilation,” and “end compilation.”

For purposes of maintaining information about an action, events are typed and can have attributes; one attribute is the time the event occurred. Generally, the other event attributes would be items such as the agents and artifacts associated with an event, the tangible results of the action (e.g., pass/fail from a design review; errors/no-errors from a compilation), and any other information that gives character to the specific occurrence of that type of event. In the work described here, we do not make direct use of attributes other than time, although doing so is an important direction for future work. Nevertheless, all attributes recorded in the data are available to the process engineer through the event definition mechanisms employed by our support tools [Cook and Wolf 1998a]. For example, it is possible to take a “slice” of a data set based on a specific event attribute.

The overlapping activities of a process, then, are represented by a sequence of events, which we refer to as an *event stream*. For simplicity, we assume that a single event stream represents one execution of one process, although depending on the data collection method, this assumption can be relaxed.

The ability to collect events is central to supporting event data analysis. Fortunately, the tools and environments of today provide strong support for logging the events that occur on a system. A version control system, for example, logs the accessing and modification of documents and code. Off-computer events, such as meetings or phone calls, can be logged manually. Email logs (e.g., using a project alias) can be a starting point for collecting communication events. A detailed discussion of the existing support for collecting events and previous studies that have made use of events is beyond the scope of this article, but can be found elsewhere [Bradac et al. 1994; Cook 1996; Cook and Wolf 1998a; Cook et al. 1998; Krishnamurthy and Rosenblum 1995; Selby et al. 1991].

2.2 Relating Models and Events: Event Sites

For our purposes we focus on behavioral process-modeling formalisms. These include models based on state machines (e.g., Statemate [Harel et al. 1988]), Petri nets (e.g., SLANG [Bandinelli et al. 1994] and FUNSOFT Nets [Gruhn and Jegelka 1992]), procedural languages (e.g., APPL/A [Sutton et al. 1995]), and rule-based languages (e.g., Oz [Ben-Shaul and Kaiser 1994]).




Fig. 1. Venn diagram of event types.

We assume that a model described in any such formalism induces one or more event streams, and thus it has places in its behavioral description where events can be recognized. We call these places *event sites*. A state machine, for example, has state transitions as event sites, where a transition is labeled with the event it produces. A Petri net also naturally has its transitions as event sites—a firing sequence is, in effect, an event stream. A rule-based language would have rules as event sites. Note that not all transitions (or rules) must be event sites. There may be internal behavior in a model that does not need to be visible to event stream analyses.

Event sites are typed—i.e., each event site produces a specific type of event. For each event type, then, there is a set of event sites in a process model that produce it.

2.3 Event Domains

Given the distinction between the process model and the process execution, there are really two universes of event types. One universe is the set of event types associated with the model of a process, while the other is the set of event types associated with the execution of the process. Their conceptual relationship is depicted in Figure 1, which indicates that the sets are not necessarily equivalent. For example, consider an organization that executes a particular process. Some members of that organization may informally and unilaterally decide to perform occasional code inspections; a formal model of that process might not account for such an ad hoc activity. Conversely, a model adopted from another organization might include a subprocess for design reviews, but the adopting organization might decide never to perform that activity.

A third set of event types, also shown in Figure 1, includes those that are actually collected as data. Since this set must necessarily be a subset of the execution events (one cannot record something unless it actually occurs), it can be viewed as a window onto the actual process execution. This window might not show the whole execution, because there may be some activities for which no event data are collected. There are several reasons why this might occur; but two obvious ones are that data about a particular event type might be considered inconsequential, or the data might be considered too expensive to collect. For instance, events that occur off the computer,

such as most staff meetings, are likely to be more expensive to collect than events that occur on the computer, simply because off-computer events would require manual, as opposed to automated, collection techniques.

It is important to note that we concentrate here on event types, which are abstractions of the activities in a process. We assume that the model and the execution agree to a significant extent about the set of event types (i.e., activities) involved in the process, although they may not agree on the specific orderings and numbers of events of those types. If they did not largely agree on the basic sets of activities, then it would not be clear what it would mean for them to be relating to the same process. Indeed, modeling and data collection are often closely related, in the sense that models are used to frame the data collection activity, and vice versa.

We do not, however, assume that the particular names for the event types used in the model and found in the collected execution data are equivalent. Fortunately, this is a simple syntactic issue that can be easily dealt with through a name mapping applied to either event stream. In fact, the data analysis framework within which our validation tool is implemented provides a convenient mechanism based on regular expressions for creating and applying such mappings to the execution stream [Cook and Wolf 1998a]. The mechanism additionally allows the analyst to extract events from the data and map them to event types at arbitrary levels of granularity. In particular, the granularities of event types derived from the data can be made to match the granularity of event types found in a model.

3. PROBLEM STATEMENT AND APPROACH

In our framework for process validation, we have an executing process that produces an actual event stream, and, on the other side, we have a model that induces a desired or prescribed event stream. Thus, we can cast the validation problem as quantitatively measuring how close the event stream of the executing process resembles an event stream induced by the model. We call these two event streams the *execution event stream* and the *model event stream*, respectively. Figure 2 depicts the process validation framework.

There are several methods for performing a measurement such as this, but one that seems most applicable is the *string distance metric* [Kruskal 1983]. A string distance metric counts the number of token insertions, deletions, and substitutions needed to transform one string into the other. By applying various mathematical transformations, this method becomes a family of metrics. String distance metrics have been used in applications as varied as DNA/RNA matching [Waterman 1984], substring matching [Kashyap and Oommen 1983; Schneider et al. 1992], spelling-error correction [Du and Chang 1992], syntax error correction [Aho and Peterson 1972; Fischer and Mauney 1992; Röhrich 1980], and text file differencing as in

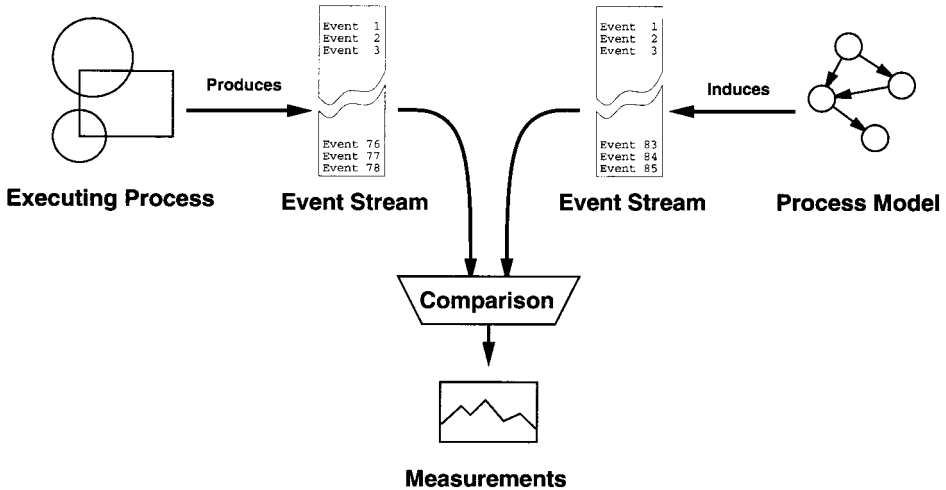


Fig. 2. Process validation framework.

the UNIX tool *diff*. In general, string distance metrics have become the standard method in any domain requiring symbolic sequence comparison.¹

Other methods that could be used to quantify the difference between two event streams do not offer the versatility of string distance metrics. Hamming distance, for example, is the count of the number of tokens that differ, but this method assumes that either the streams are the same length or that they can be suitably matched and padded. In fact, string distance metrics can subsume this method by ignoring the insertion and deletion operations, and just tallying substitutions.

4. VALIDATION METRICS

In this section we introduce two metrics for determining the correspondence between a formal model of a process and an execution of the process. They share the characteristic that they use string distance to compare the event stream produced by a process execution to an event stream representing a possible behavior predicted by the process model. The issue of how the second of these event streams is constructed is an important one and is discussed in Section 6. We defer detailed examples of applying the metrics to Section 5.

4.1 Simple String Distance Metric

The first metric uses a simple, direct approach to measuring string distance, and thus we refer to it as the *simple string distance* (SSD) metric. Under this method, the distance between two strings is measured by counting the minimum number of token insertions, deletions, and substitu-

¹Numeric sequences, which are really a representation of some mathematical function (e.g., a time series of a stock value), is a different topic altogether.

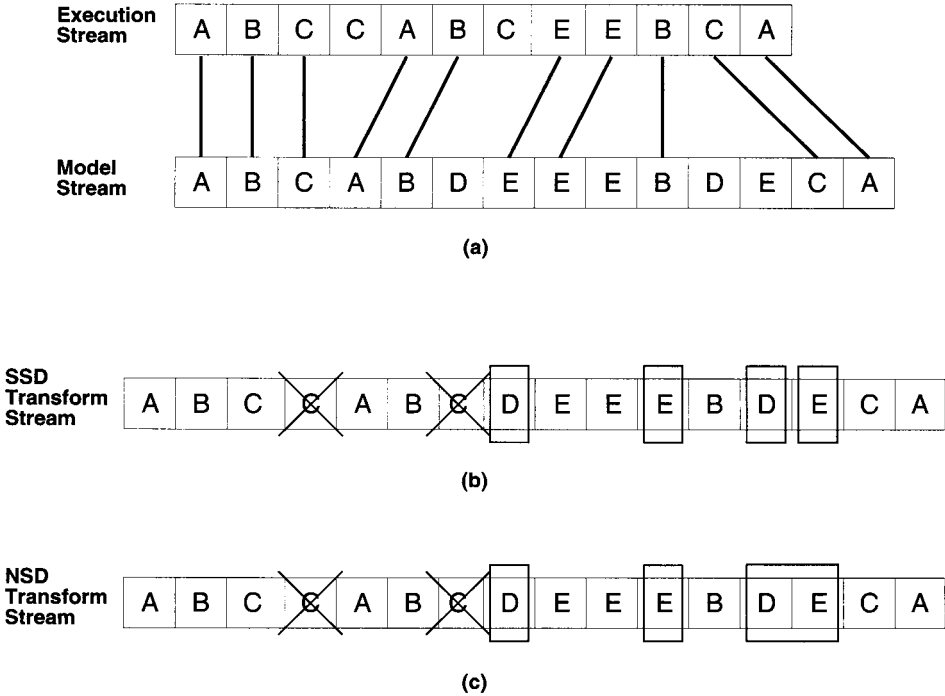


Fig. 3. Example execution and model event streams (a), with execution stream transformed for the SSD metric (b) and the NSD metric (c) calculations.

tions needed to transform one string into the other. For our purposes, we choose the execution event stream as the one to which the operations are applied.² With this choice, insertions represent missed activities (the model predicted them, but the execution did not perform them), and deletions represent extra activities (the model did not predict them, but they were performed, in any case).

As an example, consider the execution and model event streams shown in Figure 3(a), where the lettered boxes represent events. The lines drawn between the two streams indicate one possible correspondence between their respective events. The transformation of the execution stream into the model stream is depicted in Figure 3(b). In particular, we delete a C, substitute a D for a C, and insert an E, a D, and another E. The resulting value for the distance is then 5. This happens to be the minimum transformation required.

To strengthen the metric, weights can be assigned to each of the operation types insertion, deletion, and substitution, giving a relative cost to each. Then, instead of minimizing the number of operations to calculate the distance, the goal would be to minimize the total cost of the operations.

²The operations are isomorphic, so choosing one event stream rather than the other does not change the resulting measurement; it just reverses the senses of insertion and deletion.

Given two strings, one of length M and the other of length N , the minimum total cost of operations can be computed in $O(MN)$ time using a well-known dynamic program [Kruskal 1983].

In some applications of this method, such as DNA/RNA sequencing or text recognition, token substitution in the string distance metric makes sense. For process validation, however, it is not clear that a substituted event should contribute in any way to the measure of the correspondence. To account for this, we can set the weight of substitution to be greater than the sum of the insertion and deletion weights, so that substitution is never applied, since it would then be less costly to apply a deletion and insertion pair at the potential substitution point. We will not consider substitution further in this article.

The SSD metric is formulated as the following equation:

$$SSD = \frac{W_I N_I + W_D N_D}{W_{max} L_E}$$

where W_I and W_D are the weights for the insertion and deletion operations; N_I and N_D are the number of insertion and deletion operations performed on the execution event stream; W_{max} is the maximum of W_I and W_D ; and L_E is the length of the execution event stream. The divisor in the equation normalizes the value to the size of the input and the maximum weight used.

The weights W_I and W_D act as tuning parameters for the metric and can be used to highlight different properties of the process. For example, one could argue that insertions into the execution event stream are more costly than deletions, since they inherently represent missed activities in the process execution. Conversely, deletions from the execution event stream in some sense represent extra work that was performed (from the perspective of what is predicted by the formal model), and extra work probably does not affect the correctness of the process execution. Thus, we can set $W_I \gg W_D$ to reflect this property.

The values of the metric are, for all intents and purposes, bounded between 0 and 1.0. Although technically a value greater than 1.0 could appear (e.g., if all events are deleted and some others are inserted), this is unlikely. Thus, one might pick the standard statistical correlation rules of thumb [DeVore 1991] and say that any measurement less than 0.2 is a strong correspondence, less than 0.5 a moderate correspondence, and greater than 0.5 a weak correspondence.³

4.2 Nonlinear String Distance Metric

A characteristic of the SSD metric is that it is focused narrowly on the costs of individual string transformation operations, since each operation is

³Actually, these are inversions of the standard statistical rules of thumb, but their effect is the same.

weighted separately. In terms of process behavior, though, a sequence of missed activities, for example, can be viewed as a single deviation from the expected (model) behavior and might potentially be a more serious breach of desired execution than can be represented by a simple count of those missed activities. Our next metric accounts for this.

The *nonlinear string distance* (NSD) metric is an enhancement to the SSD metric based on the notion of a sequence of insertions or a sequence of deletions. A sequence of insertions or a sequence of deletions is called a *block*. By sequence we mean an unbroken series of like transformation operations. In Figure 3(c), for example, NSD recognizes the consecutive D and E insertions required at the end of the streams as an insertion block of length 2. All other blocks in the figure are of length 1.

The NSD metric uses block lengths to calculate values. The distance equation then becomes

$$NSD = \frac{\sum_{j=1}^{N_I^B} W_I f(b_j) + \sum_{k=1}^{N_D^B} W_D f(b_k)}{W_{max} L_E}$$

where N_I^B and N_D^B are the numbers of insertion and deletion blocks; b is a particular block length; $f(b)$ is a cost function applied to a block length b ; and all other terms are the same as in the SSD metric. Note that the weights W_I and W_D could be pulled into the cost function f , but we have left them outside to more easily compare the NSD and SSD metric equations.

The definition of the cost function f is an additional tuning parameter in the NSD metric. A rather natural function to use would be an exponential one, such as

$$f(b) = e^{k(b-1)}$$

where k is a constant and is the actual tuning parameter. This equation yields 1.0 for a block length of 1, so if all blocks are kept to a length of 1, then the NSD equation reduces to the SSD equation, as expected. The cost function yields exponentially increasing values for blocks greater than 1. Notice that for $k < 0.7$ and a block length of 2, the function would cause the distance value to be less than the corresponding value given by the SSD metric, which is not what we want. For this reason, we only consider $k > 0.7$ so that the value produced by the NSD metric is always greater than the value produced by the SSD metric for blocks of length greater than 1. Practical k values range from about 1 to 3, with larger values being used when blocks are small but important. Beyond $k = 3$, the exponential factor becomes very extreme for larger block sizes.

Unlike the SSD metric, the NSD metric is unbounded on the high end, although bounded by 0 at the low end. Thus, it is harder for us to say what value might represent a good correspondence between model and execution

and what might represent a bad correspondence. We can, however, derive some values from the rules of thumb we used for the SSD metric (i.e., the 0.2 cutoff for good correspondence and 0.5 for moderate correspondence). What is needed for the NSD rules of thumb is a notion of the average block length that could be expected in an event stream with good correspondence to the model. With this defined as B_{avg} , our derived cutoff for good correspondence for the NSD metric is as follows:

$$C = \frac{0.2e^{k(B_{avg}-1)}}{B_{avg}}$$

This takes the SSD good correspondence cutoff of 0.2 and weights it according to the exponential weight of the average expected block length, taking into account the tuning parameter k . For example, if one sets $B_{avg} = 2.5$ and $k = 1.5$, then the cutoff value for good correspondence would be $C = 0.76$. This value nicely reduces to the SSD cutoff value for $B_{avg} = 1$. For the moderate cutoff value, we would use 0.5 in place of 0.2. Note that as a history of applying the NSD metric to a process is accumulated, the actual value of B_{avg} for that process will be known.

4.3 Event Type Weighting

As discussed above, our validation metrics include a means to differentially weight the insertion and deletion transformation operations. We did this because they represent conceptually different kinds of deviations, namely missed and extra activities, respectively.

But, at a finer level of granularity, it would also be useful to differentiate the relative importance of specific types of events. For example, an event representing the conclusion of a regression-testing activity might be considered more important than an event representing the conclusion of a weekly team training meeting. Thus, one would like the possibility of weighting each event type differently, to distinguish the significance of not matching each kind of event during process validation.

The SSD and NSD metrics easily incorporate such a weighting scheme, and our tools implement this functionality. The process engineer can override the default insertion and deletion weights, and vary the operation weights according to the type of event involved in the operation.⁴ Having the ability to tailor the weights by event type allows the engineer to distinguish between large-grained events versus small-grained events, and important events versus unimportant events. For example, a high insertion weight for a code inspection event would signify that it is an important event to have actually occur during the process execution.

⁴As a practical matter, the default insertion and deletion weights are used in the normalization part of each metric's equation.

4.4 Auxiliary Measures

As presented, the metrics provide a single value as the measure of correspondence. This measure by itself is useful to indicate the presence of a potential problem. But a deeper understanding of the deviation is needed to uncover the source of the problem. The SSD and NSD metrics naturally provide auxiliary measures to the analyst, including

- the number of events that match in the compared streams;
- the number of insertion and deletion operations used to calculate the metric;
- the number, size, and average size of blocks of operations;
- the number of matches, insertions, and deletions per event type;
- the locations in the execution event stream where deviations occur; and
- the locations in the model where deviations occur.

These auxiliary measures enhance the usefulness of the basic metrics. In Section 7 we show how they contribute detailed and important information to an understanding of deviations in an industrial case study.

5. EXAMPLE USE OF THE METRICS

To illustrate the validation metrics introduced in the previous section, we use the Test Unit task from the ISPW 6/7 process problem [Kellner et al. 1990]. This is a simple and small process fragment, but it should give the reader a feel for how the metrics are applied to a process and how they can be used to inform the engineer of process deviations at varying levels of detail.

In this task, a developer and a tester are involved in testing a module that has undergone some change. They are to retrieve the test suite from configuration control, build the test executable, run all the specified tests, and make sure that at least a 95% code coverage has been achieved by the tests. If a failure occurs, either because the new module has an error or the test suite needs updating, then they are to notify the module developers or test developers, as appropriate. Upon a successful completion of the tests, they are to store the test results under configuration control and alert the manager to the new status of the module.

Figure 4 shows a colored Petri net model [Jensen 1986] of this process. Circles denote places, and rectangles denote transitions. Tokens have attributes (i.e., are colored), and those attributes are used by transition predicates to deterministically control the transition firing. Thick rectangles correspond to transitions that are event sites in the model and are labeled with the event that is produced at that site. Thin rectangles correspond to (internal) transitions used to control the model but that are not themselves event sites. To keep the figure simple, we collapse the begin/end event pairs of an activity into one (pseudo) event type; each event

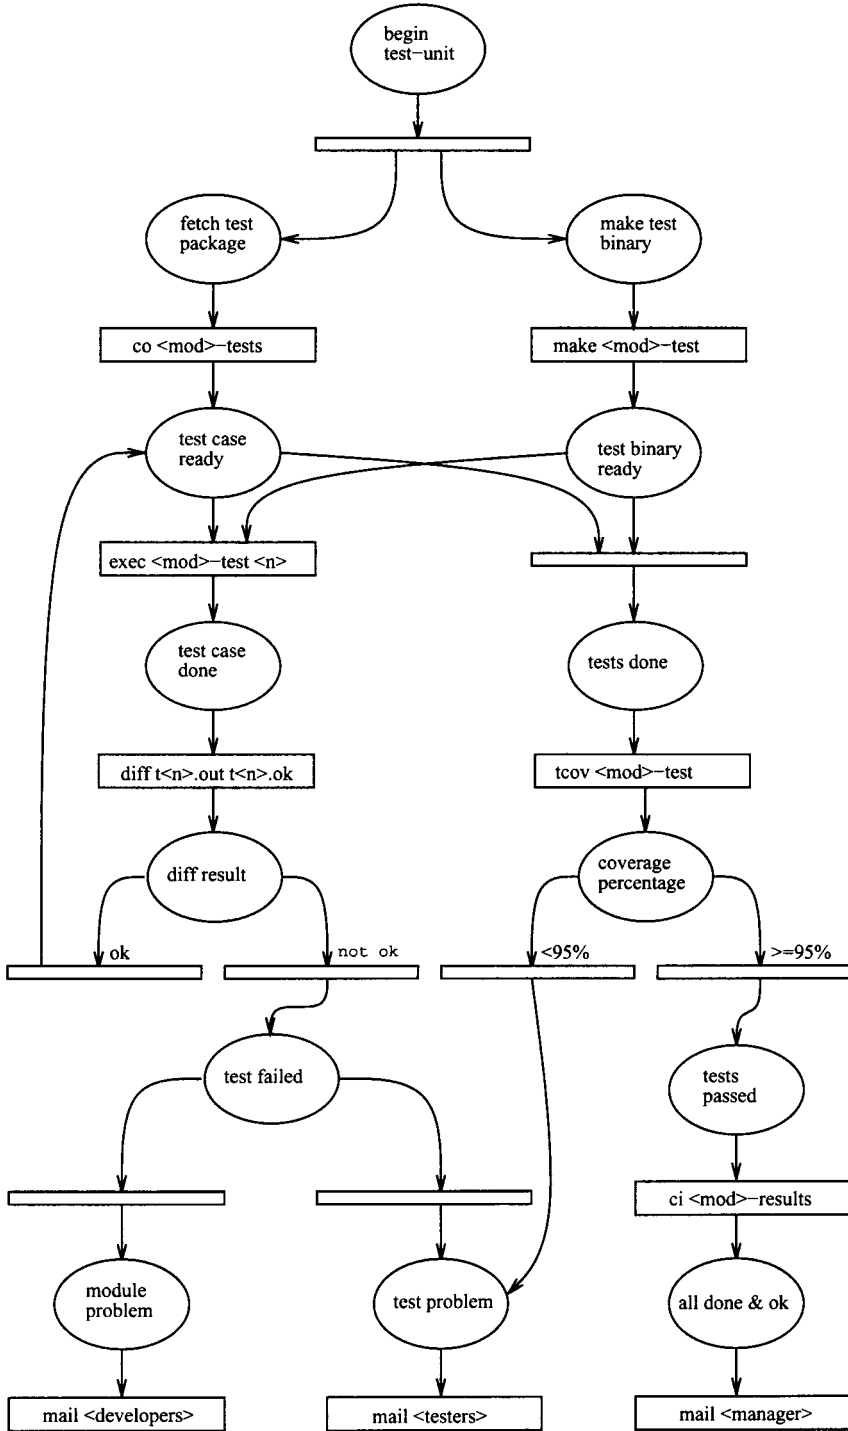


Fig. 4. Petri net model of the ISPW 6/7 test module task.

Table I. Example Pairs of Execution and Model Event Streams

Example 1		Example 2		Example 3		Example 4		Example 5	
Exec.	Model	Exec.	Model	Exec.	Model	Exec.	Model	Exec.	Model
co	co	co	co	co	co	co	co	co	co
make	make	make	make	make	make	make	make	make	make
exec	exec	exec	exec	make		exec	exec	exec	exec
diff	diff	diff	diff	make		diff	diff	diff	diff
exec	exec	exec	exec	exec	exec	exec	exec	exec	exec
diff	diff	diff	diff	diff	diff	diff	diff	diff	diff
tcov	tcov	exec	exec	exec	exec	exec	exec	exec	exec
ci	ci		diff	diff	diff	diff	diff	diff	diff
mail-m	mail-m	tcov	tcov	tcov	tcov	tcov	tcov	exec	exec
		mail-t	mail-t		ci	mail-d		diff	diff
				mail-m	mail-m		ci		tcov
							mail-m		ci
								mail-d	mail-d

site can be thought of as a two-transition sequence with the first producing the begin event and the second producing the end event. We use familiar UNIX command names as the names of event types.⁵

Table I shows five example pairs of execution and model event streams. A blank space in an execution event stream is a point at which the model has predicted that a particular event should have occurred, but in fact that event did not occur. Similarly, a blank space in a model event stream is a point at which an event occurred in the execution that was not predicted by the model. Intuitively, such blanks correspond to either missed or extra activities in the process execution, or to an error in the model. For example, in stream 3, the execution involves three consecutive invocations of the “make” tool, perhaps as a result of some problem performing the build, while the model predicts that only one should have occurred.

The SSD and NSD metric calculations require transformation of an execution event stream into the corresponding model event stream by means of insertion and deletion operations.⁶ We apply an insertion operation at a blank in the execution event stream, and we apply a deletion operation at a blank in the model event stream.

Table II shows validation measurements for the event streams in Table I. Each row contains measurements for the correspondingly numbered example event stream. The first two columns give the raw number of insertions and the raw number of deletions needed to transform the execution event stream into the model event stream.

⁵For those unfamiliar with the UNIX command names appearing in the figure, “co” and “ci” are the check-out and check-in commands for a configuration management tool; “make” is a build tool; “exec” stands for the running of an executable (i.e., a test run, in this case); “tcov” is a test coverage tool; “diff” is a text-differencing tool; and “mail” is an electronic mail tool.

⁶Recall that we chose to use the execution event stream as the one to which the transformation operations are applied (see Section 4).

Table II. Example Event Stream Measurements for the Streams of Table I

Example	# Ins	# Del	SSD	NSD	NSD	SSD	NSD	NSD
			$W_I = 1$ $W_D = 1$	$W_I = 1$ $W_D = 1$ $k = 1.5$	$W_I = 1$ $W_D = 1$ $k = 3$	$W_I = 4$ $W_D = 1$	$W_I = 4$ $W_D = 1$ $k = 1.5$	$W_I = 4$ $W_D = 1$ $k = 3$
1	0	0	—	—	—	—	—	—
2	1	0	0.11	0.11	0.11	0.11	0.11	0.11
3	1	2	0.30	0.55	2.11	0.15	0.21	0.60
4	2	1	0.30	0.55	2.11	0.23	0.47	2.03
5	3	0	0.30	0.55	2.11	0.30	0.55	2.11
Good Cutoff Values			0.20	0.45	2.01	0.20	0.45	2.01

The last six columns of Table II give the results of the parameterized string distance calculations. We vary the relative weights of W_I and W_D for both the SSD and NSD metrics, and vary the exponential constant k for the NSD metric. We present cases where the weights are equal ($W_I = W_D = 1$) and cases where the insertion cost is weighted heavier ($W_I = 4W_D$) to highlight missed events in the execution.⁷ The exponential constant k for the NSD metric is given values 1.5 and 3 to show the magnitude of change and the unboundedness of the metric. The last row of the table shows the cutoff values for the “good” correspondence rules of thumb for each metric; values in a column that are less than the bottom row fall into what we would call a good correspondence between the model and the execution event streams. For the NSD metric cutoffs, the average expected block length, B_{avg} , is taken to be 2.

There are several interesting things to see in the measurements presented in Table II. The first observation is the similarity of values for the three columns with $W_I = W_D = 1$. For streams 3, 4, and 5, which all have one block of length 1 and one of length 2 (though in different operation combinations), these measurements do not differentiate among the discrepancies. On the other hand, if one looks at the measurements with $W_I = 4$ and $W_D = 1$ (the last three columns) for event streams 2 and 3, one can see the effect of weighting insertions heavier than deletions. For the SSD metric, the measurement only changes by 0.04 with the addition of the two deletes (in stream 3), and still remains well within the good correspondence range for the NSD metric. For stream 3, the SSD with $W_I = 4$ also produces a measurement that is in the good correspondence range, whereas the SSD with $W_I = 1$ is in the moderate range (0.2–0.5). Since stream 3 has just one insertion, like stream 2, this weighting better reflects the correspondence of the execution streams than does the $W_I = 1$ weighting.

⁷The cost ratio given by $W_I = 4W_D$ was chosen arbitrarily for this example. In practice, process engineers would explore various ratios that might best reflect the actual situation they are analyzing.

For event stream 4, where the insertions have a block of length 2 rather than the deletions as in stream 3, the last three measurements (with $W_I = 4$ and $W_D = 1$) are significantly greater than for event stream 3. This shows how the metrics can be tuned to place importance on insertions—i.e., on missed events in the process execution. Event stream 5 also shows this in the last two measurements, where it has all insertions, and the difference in the weighting of insertions is evident in comparison to event streams 3 and 4.

6. DERIVING THE MODEL EVENT STREAM

Section 4 presents the process validation metrics assuming the existence of two event streams, an execution stream and a model stream. However, we are given only one stream, the execution stream. Instead of a model stream, we have a model that describes a set of streams. To perform the measurement, we must induce a model stream from the model.

The challenge, of course, is that a formal model of any but the most trivial process likely leads to a large, if not infinite, number of possible event streams. How do we choose one? Because we are measuring correspondence, we need to derive a model event stream that most closely matches the execution event stream, in order to get as useful a measure as possible. By “most closely” we mean the model stream that gives the minimum distance measurement. Any other model stream would imply a greater discrepancy than necessary. It is important to note that closeness is not a fixed property of the relationship between an execution event stream and a model, but also depends on the validation metric being applied and on the weightings used with the metric; different weights on the insertion and deletion operations (and on the event types) will affect which model stream is the closest.

Deriving a model stream does not have to be a blind generate-and-test effort. For example, if the execution stream exactly matches the model, then one wants to choose the model stream that *is* the execution stream. If there are small differences, one still wants to choose a model stream that is almost like the execution stream. Thus, the model stream derivation problem can be restated as finding the smallest changes to the execution stream that make it a valid model stream. These changes are exactly the insertions and deletions that contribute to the distance metrics. Clearly, both the execution stream and the model itself can and should be used to derive the model stream used in the metric calculation.

6.1 Background

There are three main areas of related work that have addressed a problem similar to that of our model stream derivation problem: error-correcting parsing, behavior searching, and regular-expression matching. All of the approaches use both a given “stream” and a model. We review the work in these areas to place our approach in context.

6.1.1 *Parsing*. When a syntax error is encountered, a modern compiler for a programming language is expected to report the error and then recover in some way so that it can continue to parse the rest of the program. In essence, the compiler must find a correction between a given string of tokens (the program) and a model (the syntax of the language) so that the model can continue to match the rest of the tokens. A minimal correction is desired to allow the compiler to process as much of the program as possible, and thus this problem is similar to the one we face in process validation.

Compiler research has produced several methods of interest here.

- Aho and Peterson [1972] show a cubic algorithm for globally performing minimum-cost error correction in terms of token insertion and deletion. They do not expect their algorithm to be used, however, because of its high cost. Rather, they propose it as a baseline against which to compare other methods.
- Röhrich [1980] describes an error correction method biased toward insertion of symbols, arguing that as little of the program text should be skipped (deleted) as possible. This method is based on the idea of minimum distance correction, but makes the assumption that one never needs to back up in the input stream to find a good correction.
- Fischer and Mauney [1992] describe a method for local least-cost error correction. They are also biased toward insertions, but include deletions. Their method uses a local search with a priority queue to find a local minimum-cost fix. They show that their method is fast enough to reasonably implement in a compiler.

These techniques have been specifically developed for programming language parsing, and for the grammars that are used in that domain. Some of the critical assumptions, such as not being able to back up in the input stream, are not necessarily valid in the software process domain. Thus, in general, these techniques are not applicable to our model event stream derivation problem. Nevertheless, they demonstrate two important points: optimal solutions are cost prohibitive, and heuristics can be effectively employed in practice.

6.1.2 *Behavior Searching*. Model checking is a technique to efficiently explore a finite-state space for inherent behavioral properties, including whether a particular behavior is allowed by a model. In one example, Burch et al. [1992] describe a model checker based on binary decision diagrams that is able to check models with 10^{20} states, where previous work had only handled 10^8 states. While this is impressive, the models being analyzed were of a pipelined arithmetic logic unit, which has many self-similar states resulting from the width of bits that make up a value. Their model checker directly represents this regularity in the state space, so that they avoid, to a large extent, the state explosion. If models do not have

regularity in the state space, they admit that their techniques will not provide much leverage.

Another example of a technique to search for a behavior is found in the Constrained Expressions framework [Avrunin et al. 1991; Dillon et al. 1988]. This is a method that, given a model, a current simulation state of that model, and a desired event, can answer the question “*Can this event be produced in the future?*”. The model is stated as a system of event sequences, specified by extended regular expressions. This representation, along with the desired event or event sequence to find in the behavior, is reduced to a system of inequalities that are fed into an integer linear solver. The solver produces a binary answer indicating success or failure of the search and, if it finds a solution, several parameters.

If the solver answer is “yes,” heuristics are used along with the parameters from the solver to produce a plausible behavior that leads to the event. This behavior constitutes the next sequence of model events. Unfortunately, if the answer is “no,” the Constrained Expressions framework cannot help in determining a correction to the event stream or model state to continue the analysis of the rest of the event stream.

Our problem of process validation needs techniques that analyze the model in the continual presence of deviations from the model. Thus, it appears that techniques like those of the Constrained Expressions framework are not applicable. The issue is that while they leverage system transformations to gain speed and scalability, the transformations make the system inherently uninspectable.

6.1.3 Regular-Expression Matching. Knight and Myers [1995] and Myers and Miller [1989] describe algorithms for approximately matching a string to a regular expression, using insertion, deletion, and substitution operations. These methods build on the dynamic programming techniques of string-to-string comparison algorithms, and extend this to regular expressions. For simple operation and symbol weightings, equivalent to our SSD metric, their algorithms operate in $O(MN)$ time.

However, dealing with multisymbol blocks (or gaps), as our NSD metric requires, complicates matters significantly. In general, for both string-to-string comparisons [Epstein 1990] and string-to-regular-expression comparisons [Myers and Miller 1989], arbitrary cost functions for blocks require at least $O(MN \max(M, N))$, or cubic time.⁸

⁸Better results can be obtained if one assumes concave block costs, where $F(B_i) - F(B_{i-1}) \geq F(B_{i+1}) - F(B_i)$, i.e., where the difference between the cost of a block of length $i + 1$ and i is nonincreasing as i increases. With this assumption, regular-expression matching takes $O(MN(\log M + \log^2 N))$ time, but also takes $O(MN + N \log^2 N)$ space. Our NSD metric, in general, does not have concave block costs. In fact, in our formulation, we use convex costs because a longer block represents a more serious deviation from the process model. In other domains a block cost function is naturally concave. For example, in DNA matching, the high cost of physically breaking the sequence means that as a block gets longer then the cost of breaking can be amortized over the length of the block. This results in a concave cost function.

The regular-expression algorithm takes advantage of the simplicity of its modeling paradigm. In general, constructs used in process modeling languages are not reducible to regular expressions. More powerful, yet still restricted, constructs have been studied. For example, context-free languages are thought to have high-order polynomial-time algorithms for solving approximate matching [Knight and Myers 1995].

In general, these superquadratic to cubic techniques, while providing optimal answers, are impractical. They also generally have large constants in the actual running times.

6.2 Incremental, Data-Driven Matching

The survey of related approaches presented above leads one to the conclusion that the general model event stream derivation problem has no known efficient, optimal solutions. In fact, some formulations of the problem are known to not have optimal solutions that are efficient.

Fortunately, the problem is not quite so bad as it seems. First, we observe that the execution event stream likely corresponds, at least in places, to any reasonable model of the process. If this were not true, then the model would be so contrary as to be immediately and obviously useless. Second, we can use the execution event stream to help guide our search of the model, thus significantly cutting down on the required search space. In particular, we traverse the execution event stream and incrementally derive events for the model event stream by consulting the model. Where the model and the execution stream match, the model stream will simply mirror the execution stream. Where they do not match, some method that searches the model will need to be employed to find the minimum-cost set of inserted and deleted events so that the execution stream can be changed to continue to match the model. The matched, inserted, and deleted events describe exactly the model stream that is induced from the model and that the metric then uses.

The resulting approach implies a state-space search, one that uses heuristics to control the state explosion. The states in the space are not just the states in the process model state space, but also include the position of the event in the execution stream that is currently being examined, and the operation (match, insertion, or deletion) that led to the state. Having these attributes, along with the method of calculating state costs, means that the search space is a tree, even if the process model's state space is not.

Figure 5 shows a partial view of the search space matching a string to an FSM. The search states are labeled with the FSM state, the position in the given string, and the operation (match, insertion, or deletion) with token type that created this search state. The bold search states represent the lowest-cost path, assuming that deletion is weighted less than insertion. The lowest-cost path deletes one token, and the resulting model string *ABA* is the closest one to the given string, *ABBA*. If insertion is weighted less than deletion, the model string *ABABA* would be the closest one, since inserting the single *A* would be cheaper than deleting a single *B*.




Fig. 5. A two-state FSM model (a) and a partial search tree (b) attempting to match the string *ABBA*.

An obvious candidate for this approach is best-first search. While the standard depth-first and breadth-first searches of a tree of states are exhaustive in a single dimension, best-first search is a heuristic-driven search that determines its search path by following the lowest-cost paths in the state space. For each state S with a parent S_p , a cost is estimated by

$$EstimatedCost(S) = Cost(S_s, S_p) + Cost(S_p, S) + Estimate(S, S_g)$$

where S_s is the start state and S_g is the goal state. In other words, the total estimated cost is the known cost of getting from the start state S_s to S_p , plus the new known cost of getting from the parent S_p to the new state S , plus an estimate of the cost of getting from S to a goal state S_g . The heuristic is in estimating the cost of going from S to a goal state S_g .

The best-first search uses a priority queue of states to be evaluated, and always evaluates the lowest-cost state on the priority queue. When it reaches a goal state, one can either stop or go through one more iteration of states to make sure that the goal found is not likely to be usurped by a lower-cost goal. This method is not guaranteed to find a minimum cost solution unless the heuristic estimator can be proven to always underestimate the true cost [Rich 1983]. If this is true, then the first goal found will always be a lowest-cost goal. Without adding estimations of the cost to a goal state, this method is referred to as *uniform cost*, since it also always finds a lowest-cost goal. Unfortunately, always underestimating the true cost (which includes uniform cost) also guarantees that every state that is lower cost than the goal itself will be inspected.

Since we are trying to calculate minimum-cost distance metrics, it is natural to use the metrics to assign actual costs to the states during our

state search. In fact, we must use the metrics as actual costs in order to guarantee we are minimizing the correct function. But the question of how to estimate the cost of reaching a goal remains. The goal state in validation can be defined as the state matching (or deleting) the last event in the execution event stream. One might want to define a goal state as a state that contains a termination state of the process model, but this is not as flexible, since it does not allow validation of incomplete process executions.

The SSD and NSD metrics are already normalized with respect to the length of the event stream, so the position in the event stream is factored out of the cost assigned to a state. Thus, we might say that the estimated total cost of reaching a goal state is just the value of the SSD or NSD metric that has been calculated so far—i.e., the cost of the current state. This assumes that the event stream processed so far is representative of the total stream, and that the metric calculation will not greatly change. This heuristic estimator (and cost metric) does not always underestimate the goal state cost, so it is not guaranteed to find the minimum-cost goal. In particular, from a given state there may be continual matches (zero cost) to the end of the stream, thus allowing the distance metric to diminish toward zero as the length increases. Trying to construct a complex estimator that is tailored toward identifying these anomalous cases is counterproductive. However, we have found it to be a reliable estimator, and in practice we have never seen a case where the minimum-cost goal was not reached.

6.3 Pruning

The technique of *pruning* a state space has proven to be useful in reducing the cost of finding a low-cost goal [Rich 1983]. Pruning discards portions of the state space that look unpromising. By pruning, one cannot guarantee a lowest-cost goal. But in some domains, such as game playing, “smart” pruning has negligible effects on the outcome of the search while dramatically reducing search costs. Pruning takes many forms and can use vastly different methods and heuristics.

One heuristic that we employ is to discard any newly generated state that has an estimated cost higher than some threshold relative to the current best-looking state. We refer to this as *cost pruning*. The hypotheses behind cost pruning are that the estimated costs are fairly accurate, or at least predictable, and that a state’s actual cost is not likely to be vastly better than its estimate. Additionally, cost pruning assumes that one can set a single threshold for the whole state space, and that this threshold will work consistently.

Our initial observations show that the variability in costs assigned to states should change over time; in the beginning, especially, there is much larger variability. Thus, for our pruning method, we set a threshold as a fraction of the current standard deviation of state costs. This lets the threshold account for some of the variability during the state space search.

Another pruning method we employ is to discard any state that is some specified distance behind in the execution event stream from the current




Fig. 6. Search tree before and after position pruning.

farthest state. We refer to this as *position pruning*, and its heuristic assumes that the most likely paths to the lowest-cost goal state will be examined in an interleaved fashion—i.e., their costs will not fluctuate too widely from each other, and thus the paths will be expanded close to each other. Then, any unexamined state that is far enough behind (i.e., more than the specified position-pruning parameter) in the event stream is ignored as unlikely to be on a path to a good goal state. For example, if a state that was at the 36th event in the event stream was the furthest along in the stream, and the position-pruning parameter was set at 5, all unexamined states that were at events previous to the 31st event would be discarded. Figure 6 shows the effect of position pruning on a small search tree. All of the open, not-yet-searched states (the leaves) that were behind the furthest state by more than 20 were discarded, along with the parent states whose children have all been discarded. This narrows the search space, resulting in direct, linear state paths to the areas of the state space still being searched.

Our experience with pruning has shown that position pruning consistently performs well for the model event stream derivation problem, while cost pruning is highly variable and often poor in its performance. We have studied our pruning methods extensively, but detailed discussion is beyond

the scope of this article. We refer the reader elsewhere for more information [Cook 1996].

7. USING VALIDATION IN AN INDUSTRIAL CASE STUDY

We recently performed a case study of an industrial software process [Cook et al. 1998]. The goal of the study was to statistically identify process behaviors that correlated with successful and unsuccessful executions of the process. One component of that study involved the use of our validation techniques. In this section we review the study as an example application of process validation in a real-world setting.

The study focused on a change request process for a large telecommunications software system. In this process, a customer reports a problem, the problem is assigned to a developer, the developer completes and tests a fix, and the fix is sent out to the customer. A successful execution of the process results in the fix being accepted by the customer, while an unsuccessful execution results in the fix being rejected. There was no existing formal model of the process, but enough documentation and informal knowledge existed to be able to create one. In addition, data analysis techniques were employed to discover possible model fragments from the data themselves [Cook and Wolf 1998b]. Figure 7 depicts a state machine model of the process.

We were able to obtain data for 159 executions of this process by extracting events from several historical archives and merging them to form complete execution streams. The archives included a source code control system, a modification request tracking system, a customer response database, and loose-leaf binders of code inspection reports. These 159 streams were divided into two populations based on their success: one where the fix was accepted (141 streams) and one where the fix was rejected (18 streams).

Table III shows the validation metrics calculated for the case study, where the statistical test for significant difference between the two populations was the Wilcoxon Rank Sum test, shown in column 3. The other columns are the p-values (column 2) and the means and standard deviations of each metric for the two populations. In general, the process execution was highly variable, with somewhat less than 65% of the behavior matching the model. With equal insertion and deletion weights, the NSD metric, using $k = 1$, showed statistically significant differences between the successful (accepted fix) and unsuccessful (rejected fix) process executions, while the SSD metric did not. Looking at the components of the metrics, we see that only insertions were significant, but then only weakly so.

By using the detailed information provided by the validation techniques—the event type that was matched, inserted, or deleted, and the model state in which each such operation was applied—we were able to localize the statistical differences between the two populations in terms of where in the model the differences appeared and on what specific type of

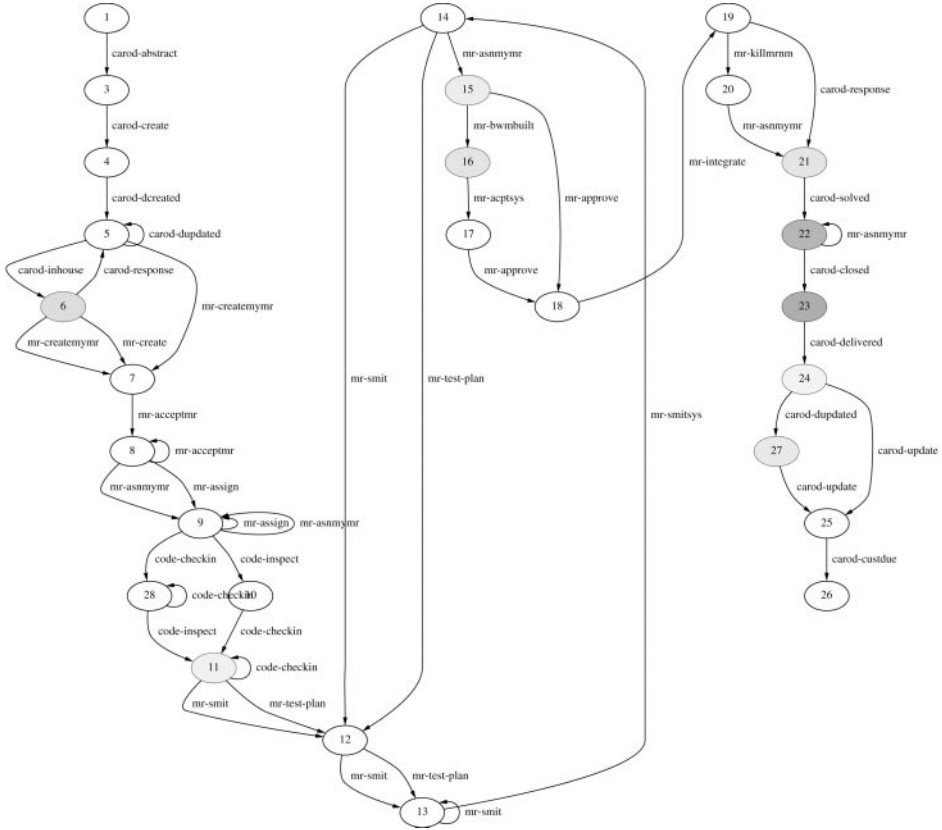


Fig. 7. Process model used in the case study.

Table III. Validation Metrics Calculated in the Case Study

Measure	P-Value (Two-Tailed)	Sig Test (W)	Accept Pop. (N = 141)		Reject Pop. (N = 18)	
			Mean	Std Dev	Mean	Std Dev
SSD	0.41	0.82	0.56	0.22	0.61	0.23
NSD	0.00	3.65	24.49	83.90	59.96	95.82
Matches	0.79	0.27	21.15	16.25	21.39	8.55
Insertions	0.10	1.63	7.18	2.90	8.28	3.16
Deletions	0.27	1.11	7.71	5.74	10.17	7.77
Insertion blocks	0.49	-0.68	3.48	1.51	3.22	1.52
Deletion blocks	0.74	-0.34	4.20	1.91	4.22	2.34

events. The shaded states in Figure 7 are those locations in the model. That is, the difference in the relative amount of events matched from a shaded state (instead of inserted or deleted) is statistically significant between the accepted-fix and rejected-fix populations. In this study, the accepted-fix population had more matches, and thus followed the process model more closely.

The multileveled analysis of this study shows the power of the information that the validation methods provide: once a gross metric (such as NSD)

indicates a problem, the detailed information available lets one perform an in-depth analysis, directly seeing where the process might be breaking down. Indeed, our results led to suggestions about where in the process some adjustments might be useful.

8. RELATED WORK

There is related work in the area of process improvement that uses data to characterize processes, but none that uses data in a process validation activity:

- Chmura et al. [1990] and Bhandari et al. [1993] try to deduce problems in the process by looking at defect data in the products. Specifically, they statistically analyze change data and effort data to determine the behavior of the process. For example, they can see ripple effects from interface changes and high percentages of fix-on-fix changes.
- Garg et al. [1993] employ a manual process history analysis in the context of a metaprocess for creating and validating domain-specific process models and software toolkits.
- Pérez et al. [1994] propose methods for evaluating the congruence of process models. Congruence is a measure of how well an environment can accommodate a given process model, based on the tools and activities already in that environment. The effort here is to predict how well a specific process will fit into an environment, rather than whether or not the model is followed once it is deployed.
- Cugola et al. [1996] define a formal framework for reasoning about inconsistencies and deviations in a process. Their approach is directed toward processes that are controlled by a process support system using an enacted model. Their goal is to enable these systems to allow, coordinate, and resolve deviations from the model. In this respect, their work is similar to ours, but they do not use data derived from the process execution to measure the deviations.

We feel that our work effectively complements these other approaches to process improvement by raising confidence in the correspondence between formal models and executions of processes.

In using event-based data to compare an execution with a formal model, our work also relates to that of distributed debugging and history checking.

- Bates [1989] uses “event-based behavioral abstraction” to characterize the behavior of programs. He attempts to match the event data to a model based on regular expressions. However, he only marks the points at which the data and model do not match, rather than attempting to provide aggregate measures of disparity.
- Cuny et al. [1993] build on the work of Bates, attempting to deal with large amounts of event data by providing query mechanisms for event

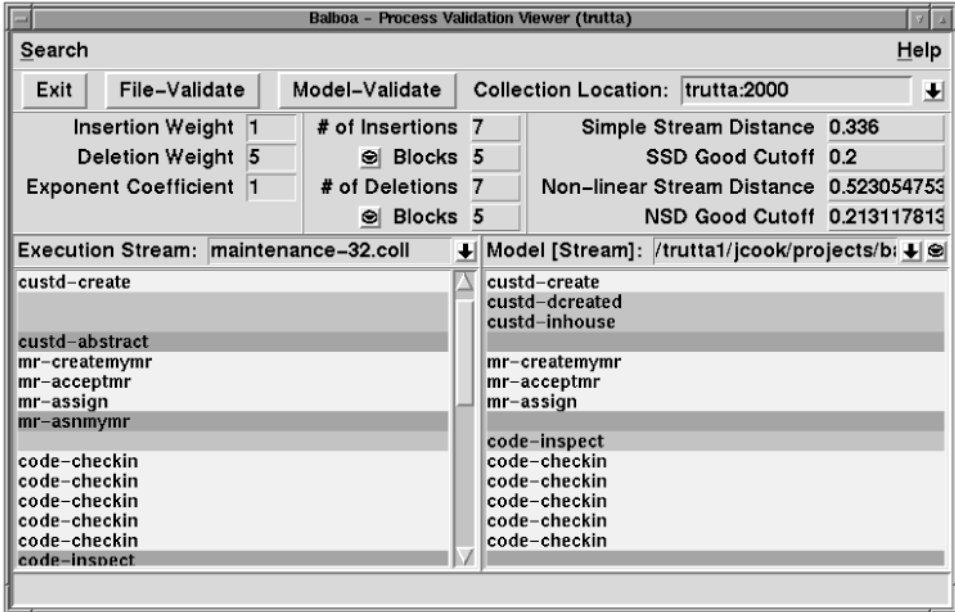


Fig. 8. User interface of the validation tool.

relationships. They assume that there is some problem somewhere in the event stream and that one is trying to locate that problem.

—Felder and Morzenti [1992] and Felder et al. [1994] describe a method and tool by which one can compare an execution history against a temporal logic specification to decide the correctness of that execution with respect to the model. Our goal is to quantify discrepancies, and therefore we take a more pragmatic approach toward “correctness.”

9. CONCLUSION

We have developed two metric-oriented techniques for process validation, from a linear distance measure in terms of event insertions and deletions, to a nonlinear distance measure that takes into account the size of discrepancies. The metrics are independent of any specific behavioral process-modeling paradigm, and thus have wide applicability.

The process validation techniques have been implemented as part of the Balboa process data analysis framework [Cook and Wolf 1998a]. The current implementation works with finite-state machine models of processes. The user interface for selecting execution streams and process models, and for viewing the results of a process validation, is shown in Figure 8.

The upper portion of the window provides a quantitative view of the validation results at three levels of detail: the counts of individual insertions and deletions; the counts of the blocks of insertions and deletions; and

the calculated distance measurements. To the left of these results is a summary of the parameters used to control the metrics calculations.

The lower portion of the window provides a scrollable, visual summary of the detailed differences between the process as executed and the process as predicted. Extra events in the execution stream are highlighted in one color (shown here as dark gray bands), while missing events in the execution stream are highlighted in another color (shown here as light gray bands).

There are many directions that future work in process validation can take, including the following:

- Identifying additional properties of process models that can be exploited in performing validation:* For example, points in a model where one can fix the execution stream and ignore previous behavior could help reduce the search cost in a large model. This is similar to the concept of trace change points [Eckert and Nutt 1996]. Matching attributes to model parameters and among events is also an important direction in this vein.
- Developing improved techniques for visualizing the results of validation:* For example, overlaying the differences onto the process model rather than onto the model event stream may help a process engineer better understand the problems in the process.
- Investigating other analyses for process executions and process models:* For example, time-oriented metrics, perhaps derived from the area of real-time analysis [Felder and Morzenti 1992; Schwartz et al. 1983], would be a useful extension to execution stream analysis. Methods for measuring the efficiency of a process would be another useful analysis method. Both would help in the optimization of a process that has already been behaviorally validated.

We intend to explore these and other directions, to continue improving the practicality and usefulness of process validation techniques, and to further experiment with their application in industrial settings.

ACKNOWLEDGMENTS

We appreciate the many helpful comments on this work provided by Clarence (Skip) Ellis, Dennis Heimbigner, David Rosenblum, Lawrence Votta, and Benjamin Zorn.

REFERENCES

- AHO, A. V. AND PETERSON, T. G. 1972. A minimum distance error-correcting parser for context-free languages. *SIAM J. Comput.* 1, 4 (Dec.), 305–312.
- AVRUNIN, G. S., BUY, U. A., CORBETT, J. C., DILLON, L. K., AND WILEDEN, J. C. 1991. Automated analysis of concurrent systems with the constrained expression toolset. *IEEE Trans. Softw. Eng.* 17, 11 (Nov.), 1204–1222.
- BANDINELLI, S., FUGGETTA, A., AND GHEZZI, C. 1993. Software process model evolution in the SPADE environment. *IEEE Trans. Softw. Eng.* 19, 12 (Dec.), 1128–1144.
- BANDINELLI, S., FUGGETTA, A., GHEZZI, C., AND LAVAZZA, L. 1994. SPADE: An environment for software process analysis, design, and enactment. In *Software Process Modelling and*

- Technology*, A. Finkelstein, J. Kramer, and B. Nuseibeh, Eds. Research Studies Press Advanced Software Development Series. Research Studies Press Ltd., Taunton, UK, 223–247.
- BATES, P. 1989. Debugging heterogeneous distributed systems using event-based models of behavior. *SIGPLAN Not.* 24, 1 (Jan.), 11–22.
- BEN-SHAUL, I. S. AND KAISER, G. E. 1994. A paradigm for decentralized process modeling and its realization in the Oz environment. In *Proceedings of the 16th International Conference on Software Engineering* (May). IEEE Computer Society Press, Los Alamitos, CA.
- BHANDARI, I., HALLIDAY, M., TARVER, E., BROWN, D., CHAAR, J., AND CHILLAREGE, R. 1993. A case study of software process improvement during development. *IEEE Trans. Softw. Eng.* 19, 12 (Dec.), 1157–1170.
- BRADAC, M. G., PERRY, D. E., AND VOTTA, L. G. 1994. Prototyping a process monitoring experiment. *IEEE Trans. Softw. Eng.* 20, 10 (Oct.), 774–784.
- BURCH, J. R., CLARKE, E. M., McMILLAN, K. L., DILL, D. L., AND HWANG, L. J. 1992. Symbolic model checking: 10^{20} states and beyond. *Inf. Comput.* 98, 2 (June), 142–170.
- CHMURA, L. J., WICINSKI, T. J., AND NORCIO, A. F. 1990. Evaluating software design processes by analyzing change data over time. *IEEE Trans. Softw. Eng.* 16, 7 (July), 729–740.
- COOK, J. E. 1996. Process discovery and validation through event-data analysis. Tech. Rep. CU-CS-817-96. Department of Computer Science, University of Colorado at Boulder, Boulder, CO.
- COOK, J. E. AND WOLF, A. L. 1994. Toward metrics for process validation. In *Proceedings of the 3rd International Conference on the Software Process* (Oct.). IEEE Computer Society Press, Los Alamitos, CA, 33–44.
- COOK, J. E. AND WOLF, A. L. 1998a. Balboa: A framework for event-based process data analysis. In *Proceedings of the 5th International Conference on the Software Process* (June). International Software Process Assoc., 99–110.
- COOK, J. E. AND WOLF, A. L. 1998b. Discovering models of software processes from event-based data. *ACM Trans. Softw. Eng. Methodol.* 7, 3 (July), 215–249.
- COOK, J. E. AND WOLF, A. L. 1998c. Event-based detection of concurrency. In *Proceedings of the 6th ACM SIGSOFT Symposium on the Foundations of Software Engineering* (Nov.). ACM Press, New York, NY, 35–45.
- COOK, J. E., VOTTA, L. G., AND WOLF, A. L. 1998. Cost-effective analysis of in-place software processes. *IEEE Trans. Softw. Eng.* 24, 8 (Aug.).
- CUGOLA, G., DI NITTO, E., FUGGETTA, A., AND GHEZZI, C. 1996. A framework for formalizing inconsistencies and deviations in human-centered systems. *ACM Trans. Softw. Eng. Methodol.* 5, 3 (July), 191–230.
- CUNY, J., FORMAN, G., HOUGH, A., KUNDU, J., LIN, C., SNYDER, L., AND STEMPLER, D. 1993. The Adriane debugger: Scalable application of an event-based abstraction. In *Proceedings of the ACM/ONR Workshop on Parallel and Distributed Debugging* (San Diego, CA, May 17–18), B. P. Miller and C. McDowell, Eds. ACM Press, New York, NY, 85–95.
- DEVORE, J. L. 1991. *Probability and Statistics for Engineering and the Sciences*. 3rd ed. Brooks/Cole Publishing Co., Pacific Grove, CA.
- DILLON, L. K., AVRUNIN, G. S., AND WILEDEN, J. C. 1988. Constrained expressions: Toward broad applicability of analysis methods for distributed software systems. *ACM Trans. Program. Lang. Syst.* 10, 3 (July), 374–402.
- DU, M. W. AND CHANG, S. C. 1992. A model and a fast algorithm for multiple errors spelling correction. *Acta Inf.* 29, 3 (June), 281–302.
- ECKERT, Z. K. F. AND NUTT, G. J. 1996. Trace extrapolation for parallel programs on shared memory multiprocessors. Tech. Rep. TR CU-CS-804-96. Department of Computer Science, University of Colorado at Boulder, Boulder, CO.
- ELLIS, C., KEDDARA, K., AND ROZENBERG, G. 1995. Dynamic change within workflow systems. In *Proceedings of the Conference on Organizational Computing Systems* (Milpitas, CA, Aug. 13–16), N. Comstock and C. Ellis, Eds. ACM Press, New York, NY, 10–21.
- EMMERICH, W., FINKELSTEIN, A., AND MONTANGERO, C. 1997. “The world and the machine”: A critical perspective on process technology. In *Workshop on Research Direction in Process Technology* (Nancy, France).

- EPSTEIN, D. 1990. Sequence comparison with mixed convex and concave costs. *J. Alg.* 11, 1 (Mar.), 85–101.
- FELDER, M. AND MORZENTI, A. 1992. Validating real-time systems by history-checking TRIO specifications. In *Proceedings of the 14th International Conference on Software Engineering (ICSE '92, Melbourne, Australia, May 11–15)*, T. Montgomery, Ed. ACM Press, New York, NY, 199–211.
- FELDER, M., MANDRIOLI, D., AND MORZENTI, A. 1994. Proving properties of real-time systems through logical specifications and Petri net models. *IEEE Trans. Softw. Eng.* 20, 2 (Feb.), 127–141.
- FISCHER, C. N. AND MAUNEY, J. 1992. A simple, fast, and effective LL(1) error repair algorithm. *Acta Inf.* 29, 2 (Apr.), 109–120.
- GARG, P. K. AND M. JAZAYERI, 1996. Process-centered software engineering environments: A grand tour. In *Software Process*, A. Fuggetta and A. Wolf, Eds. Trends in Software, vol. 4. John Wiley and Sons Ltd., Chichester, UK, 25–52.
- GARG, P., JAZAYERI, M., AND CREECH, M. 1993. A meta-process for software reuse, process discovery and evolution. In *Proceedings of the 6th International Workshop on Software Reuse* (Nov.).
- GRUDIN, J. 1990. Groupware and cooperative work: Problems and prospects. In *The Art of Human Computer Interface Design*, B. Laurel, Ed. Addison-Wesley, Reading, MA.
- GRUHN, V. AND JEGELKA, R. 1992. An evaluation of FUNSOFT nets. In *Proceedings of the 2nd European Workshop on Software Process Technology*. Lecture Notes in Computer Science, vol. 635. Springer-Verlag, Berlin, Germany, 196–214.
- HAREL, D., LACHOVER, H., NAAMAD, A., PNUELI, A., POLITI, M., SHERMAN, R., AND SHTUL-TRAURING, A. 1988. Statemate. In *Proceedings of the 10th International Conference on Software Engineering* (Singapore, Apr. 11–15), T. N. Nam, Ed. IEEE Computer Society Press, Los Alamitos, CA, 396–406.
- JACCHERI, M. L. AND CONRADI, R. 1993. Techniques for process model evolution in EPOS. *IEEE Trans. Softw. Eng.* 19, 12 (Dec.), 1145–1156.
- JENSEN, K. 1987. Coloured Petri nets. In *Petri Nets: Central Models and Their Properties* (Bad Honnef, Sept. 8–19), W. Brauer, W. Reisig, and G. Rozenberg, Eds. Lecture Notes in Computer Science, vol. 254. Springer-Verlag, Berlin, Germany, 248–299.
- KASHYAP, R. L. AND OOMMEN, B. J. 1983. The noisy substring matching problem. *IEEE Trans. Softw. Eng.* 9, 3, 365–370.
- KELLNER, M., FEILER, P., FINKELSTEIN, A., KATAYAMA, T., OSTERWEIL, L., PENEDO, M., AND ROMBACH, H. 1990. Software process modeling example problem. In *Proceedings of the 6th International Software Process Workshop* (Oct.). 19–29.
- KNIGHT, J. R. AND MYERS, E. W. 1995. Approximate regular expression pattern matching with concave gap penalties. *Algorithmica* 14, 85–121.
- KRISHNAMURTHY, B. AND ROSENBLUM, D. S. 1005. Yeast: A general purpose event-action system. *IEEE Trans. Softw. Eng.* 21, 10 (Oct.), 845–857.
- KRUSKAL, J. B. 1983. An overview of sequence comparison. In *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison*, D. Sankoff and J. B. Kruskal, Eds. Addison-Wesley, Reading, MA, 1–44.
- LEBLANC, R. AND ROBBINS, A. 1985. Event-driven monitoring of distributed programs. In *Proceedings of the 5th International Conference on Distributed Computing Systems* (May). IEEE Computer Society, Washington, DC, 515–522.
- MYERS, E. W. AND MILLER, W. 1989. Approximate matching of regular expressions. *Bull. Math. Bio.* 51, 1, 5–37.
- PÉREZ, G., EL EMAM, K., AND MADHAVJI, N. H. 1994. A system for evaluating the congruence of software process models. Tech. Rep. SE-94-7. McGill University, Montreal, Canada.
- RICH, E. 1983. *Artificial Intelligence*. McGraw-Hill, Inc., New York, NY.
- RÖHRICH, J. 1980. Methods for the automatic construction of error correcting Parsers. *Acta Inf.* 13, 115–139.
- SCHNEIDER, M., LIM, H., AND SHOAF, W. 1992. The utilization of fuzzy sets in the recognition of imperfect strings. *Fuzzy Sets Syst.* 49, 3 (Aug.), 331–337.

- SCHWARTZ, R. L., MELLIAR-SMITH, P. M., AND VOGT, F. H. 1983. An interval logic for higher-level temporal reasoning. In *Proceedings of the 2nd ACM Symposium on Principles of Distributed Computing* (Aug.). ACM Press, New York, NY, 173–186.
- SELBY, R. W., PORTER, A. A., SCHMIDT, D. C., AND BERNEY, J. 1991. Metric-driven analysis and feedback systems for enabling empirically guided software development. In *Proceedings of the 13th International Conference on Software Engineering (ICSE '91, Austin, TX, May)*. IEEE Computer Society Press, Los Alamitos, CA, 288–298.
- SUTTON, S. M., JR. 1991. Accomodating manual activities in automated process programs. In *Proceedings of the 7th International Software Process Workshop* (Oct.). IEEE Computer Society Press, Los Alamitos, CA.
- SUTTON, S. M., JR., OSTERWEIL, L. J., AND HEIMBIGNER, D. 1995. APPL/A: A language for software process programming. *ACM Trans. Softw. Eng. Methodol.* 4, 3 (July), 221–286.
- WATERMAN, M. S. 1984. General methods of sequence comparison. *Bull. Math. Bio.* 46, 473–501.
- WOLF, A. L. AND ROSENBLUM, D. S. 1993a. A study in software process data capture and analysis. In *Proceedings of the 2nd International Conference on the Software Process* (Feb.–Mar.). IEEE Computer Society Press, Los Alamitos, CA, 115–124.
- WOLF, A. L. AND ROSENBLUM, D. S. 1993b. Process-centered environments (only) support environment-centered processes. In *Proceedings of the 8th International on Software Process Workshop* (Mar.). IEEE Computer Society Press, Los Alamitos, CA, 148–149.

Received: May 1997; revised: December 1997 and October 1998; accepted: December 1998