

Fast Formal Analysis of Requirements via “Topoi Diagrams”

Tim Menzies

Dept. Elec. & Comp. Eng.
Univ. British Columbia,
2356 Main Mall, Vancouver, B.C.
Canada, V6T 1Z4
tim@menzies.com

John Powell

Averstar Inc.
NASA IV&V Facility
100 University Drive,
Fairmont WV, USA, 26554
john.powell@ivv.nasa.gov

Michael E. Houle

IBM Japan[†]
Tokyo Research Laboratory
1623-14 Shimotsuruma, Yamato-shi
Kanagawa-ken 242-8502, Japan
meh@cs.usyd.edu.au

Abstract

Early testing of requirements can decrease the cost of removing errors in software projects. However, unless done carefully, that testing process can significantly add to the cost of requirements analysis. We show here that requirements expressed as topoi diagrams can be built and tested cheaply – using our SP2 algorithm, the formal temporal properties of a large class of topoi can be proven very quickly, in time nearly linear in the number of nodes and edges in the diagram. There are two limitations to our approach. Firstly, topoi diagrams cannot express certain complex concepts such as iteration and sub-routine calls. Hence, our approach is more useful for requirements engineering than for traditional model checking domains. Secondly, our approach is better for exploring the temporal occurrence of properties than the temporal ordering of properties. Within these restrictions, we can express a useful range of concepts currently seen in requirements engineering, and a wide range of interesting temporal properties.

1. Introduction

The case for more formality in requirements engineering is overwhelming. Many errors in software can be traced back to errors in the requirements [21]. Often, the conception of a system is improved as a direct result of the discovery of inadequacies in the current conception. The earlier such inadequacies are found, the better, since the cost of removing errors at the requirements stage can be orders of magnitude cheaper than the cost of removing errors in the final system [22].

The benefit of formally checking a system is that formal proofs can find more errors than standard testing. A single formal first-order query is equivalent to many white-box or black-box test inputs.

The cost of rigorous requirements engineering may be impractically high. These costs include the *modeling cost*, the *execution cost*, the *personnel cost*, and the *development brake*. The *modeling cost* is incurred when analysts create the *systems model*, and the *properties model* needed for formal analysis. Both models are in some machine-readable form. The properties model is often much smaller than the systems model and contains a formal temporal logic¹ description of the invariants that must be proved in the systems model.

A rigorous analysis of formal properties has a high *execution cost*, since it implies a full-scale search through the systems model. For example, if a given systems model has n variables, each of which may take on a finite number of unique values m , then the size of the state space associated with that model is m^n . This space can be too large to explore, even on today’s fast machines. Despite extensive research into speeding up this search (see our *Related Work* section), analysts often must painstakingly rework the systems and properties models into more abstract and succinct forms that are small enough to permit formal analysis.

Analysts skilled in formal methods must be recruited or trained. Since such analysts are generally hard to find and retain, formal methods have a high *personnel cost*.

The above costs can be so high that the requirements must be frozen for some time while one performs the formal analysis. Hence, one of the costs of formal analysis is *development brake* that slows down the requirements process. Slowing down the requirements process is unacceptable for fast moving software companies, such as the start-up dot-coms.

Ideally, a method for reducing the cost of testing require-

¹Temporal logic is classical logic augmented with some temporal operators such as $\Box X$ (always X is true), $\Diamond X$ (eventually X is true), $\bigcirc X$ (X is true at the next time point), $X \cup Y$: (X is true until Y is true).

[†]Dr. Houle is on leave from the Basser Department of Computer Science, The University of Sydney, Sydney, NSW, 2006, Australia.

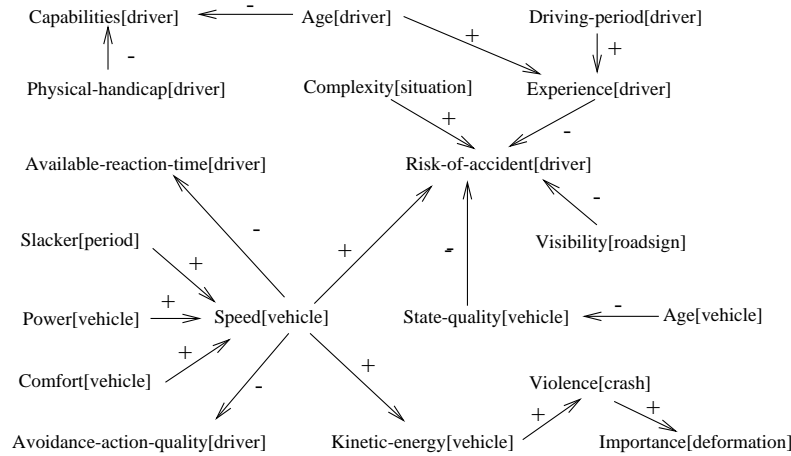


Figure 1. An example topoi from [5]. The formal semantics for topoi is described below. Informally, we say that + indicates “encourages” while – indicates “discourages”.

ments would eliminate the execution cost and reduce the cost and skill involved in building the properties and systems models. If achievable, such a method would also reduce the personnel cost, since it would not require such highly-skilled analysts. Having reduced the personnel, modeling, and execution costs, this hypothetical method would inevitably decrease the development brake.

Some progress has already been made in reducing the cost of properties modeling using *temporal logic patterns*. Dwyer et.al. [7,8] have identified patterns within the temporal logic formula seen in many real-world properties models. For each pattern, they have defined an expansion from the intuitive pseudo-English form of the pattern to a formal temporal logic formula. In this way, analysts are shielded from the complexity of formal logics. For example, the simple pseudo-English statement

always(brake = on) between(danger = seen) and(car = stop)

can be automatically expanded into the more arcane formal statement:

$$\square((danger = seen \wedge \neg(car = stop)) \wedge \diamond(car = stop)) \\ \rightarrow (brake = on \cup (car = stop)))$$

One drawback with temporal logic patterns is that while complex temporal formula can be automatically generated from intuitive pseudo-English, the execution cost remains. That is, even though we can quickly build the properties model, we may not be able to explore all of that model.

In this article, we argue that we can greatly reduce the execution cost for a class of systems models seen in the requirements stage, and for a large class of temporal logic properties. In our approach, we use temporal logic patterns to reduce the cost of properties modeling, and optimization to reduce the execution cost. The key to this reduction

is SP2, a new algorithm for testing temporal properties of *topoi*, which are statements of gradual influences between variables [5]. Topoi can be represented graphically by *topoi diagrams*, an example of which is shown in Figure 1. Topoi are quick to sketch, and so (for requirements that are topoi-compatible) our approach also reduces the systems modeling cost.

These cost-reduction benefits can only be realized if we accept certain restrictions. Firstly, our approach limits the kinds of properties that can be tested. Secondly, the systems model must be expressed as topoi diagrams. Topoi are not very expressive, and exclude statements such as first-order assertions, iterations, sub-routine calls, and assignments. Due to these language limitations, our approach is not suited to domains that need the excluded statements, such as complex protocols seen in concurrent processes.

These restrictions are not fatal to the modeling process, at least at the requirements stage. We will describe how to quickly recognize inadmissible properties statements. Further, we will use the Dwyer et.al. survey to show that within the limits to the properties language, we can represent a wide range of useful temporal logic properties. Also, we will show that topoi diagrams are sufficient to represent diagrams seen in certain approaches to requirements engineering and recording design rationales. Hence, when we say that this approach is practical and useful, we really mean *practical and useful for early life cycle requirements discussions only*.

This work is based on Feldman & Compton’s study of the validation of topoi [9] (which they called qualitative compartmental models). Menzies tried to optimize that validation process and offered an implementation that was orders of magnitude faster than the validation engine built by

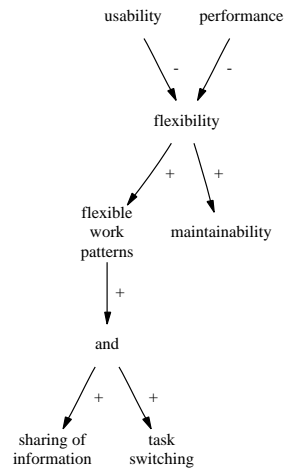


Figure 2.i: A *soft-goal graph*: the *and* node denotes that both sharing of information and task switching are enabled by flexible work patterns.

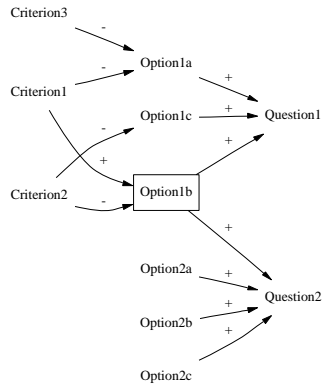


Figure 2.ii: A *questions-options-criteria graph*; from [23].




Figure 2.iii: *Topoi* from Figure 3.




Figure 2.iv: *The Smythe '87 theory*. From [24]. The diagram shows statements of gradual knowledge relating to laboratory experiments on mammals.

Figure 2. Four sample topoi diagrams.

Feldman & Compton. However, he could not reduce the exponential upper-bound on the runtimes [14–16]. Assuming a certain restriction on topoi edge types, Cohen, Menzies, Waugh and Goss showed that the cost of checking temporal properties of topoi-based simulation is a function of the number of time-ticks in the query [17]. This paper improves significantly on the Menzies et.al. result. We assume the same restriction as Menzies et.al. and introduce SP2, a nearly linear-time algorithm for checking a large class of interesting temporal properties (for space reasons, we describe the full details of that algorithm elsewhere [18]). Also, we describe an implementation of SP2 which, in at least one domain, out-performs a state-of-the-art temporal logic model checker (SPIN [11]).

2. About Topoi

Our approach assumes that requirements systems models are expressed in the form of *topoi*; i.e. statements of gradual statements such as (i) the more X, the more Y; (ii) the less X the less Y; (iii) the more X, the less Y; or (iv) the less X the less Y. Dieng et.al. offer many examples of topoi from numerous examples from their records of interviews with experts [5]. For example: *the more there is water infiltration in the roadway body, the worse the foundation risks to be*; and *if there is a punctual undressing and if the roadway is between five and fifteen years old, then the causes 'too old coating' is all the more certain since the roadway is older*.

Our experience has always been that the systems modeling cost with topoi is very low. Topoi graphs can be quickly generated in the requirements stage. Two feuding stakeholders with two marker pens and one whiteboard can generate many, many topoi in just a few hours.

Topoi graphs can be found in many domains. Figure 1 showed a topoi from an insurance domain using the graphical notation of Dieng's 3DKAT tool.

Figure 2.i show some Mylopoulos-style *soft-goal* graphs [19]. Soft-goal graphs represent gradual knowledge about non-functional requirements. In Figure 2.i, an expert describes how to increase business flexibility. Figure 2.ii shows a “questions-options-criteria” (QOC) graph from the design rationale community [23]. In such QOC graphs, questions suggest options and deciding on a certain option can raise other questions. Options shown in a box denote selected options. Options are assessed by criteria and criteria are gradual knowledge; i.e. they *tend to support* or *tend to reject* options. QOCs can succinctly summarize lengthy debates; e.g. 480 sentences uttered in a debate between two analysts on interface options can be displayed in a QOC graph on a single page [13]. Figure 2.iii shows topoi generated from the requirements of a rule-based legal system, shown in Figure 3. This translation assumes that propositions in the rule base are modeled as a belief/strength pair such as infant/X , where X is some continuous number.

Topoi seem quite simple, but often defy manual analysis. For example, Menzies' topoi theorem prover [14] found be-

```

if infant or moron
then not legally_responsible.

if guilty
then jail.

if age < 7
then infant.

if legally_responsible and guilty
then jail.

if motive and means and
  opportunity and witnesses
then guilty.

if guilty and not legally_responsible
then not jail.

```

Figure 3. Rule-based requirements from a legal system.

behaviors in very small topoi (Figure 2.iv) that had not seen after days of manual analysis. The difficulty in generating possible topoi behavior gets worse as the topoi gets more complicated. When collected from multiple stakeholders, gradual statements can be quite complex, quite large, and contain feedback loops. Smythe extracted a list of gradual influences from a set of articles from different authors relating to human internal physiology. The resulting network contains loops (e.g. Figure 2.iv). The experiments described later in the paper are based on the large topoi of Figure 4.

A pre-experimental concern is that informal topoi are so under-defined that we could use them to infer any properties at all. However, topoi structures can support certain inferences. Recall Figure 2.i and the fragment:

$$usability \rightarrow flexibility \leftarrow performance$$

Note that there is no way to explain the output of increased flexibility (denoted $\{flexibility \uparrow\}$) from the input of increased usability and performance (denoted $\{usability \uparrow, performance \uparrow\}$). That is, while topoi are over-generalized, they may still be restrictive enough to demonstrate what cannot be proved. We describe below experiments which show that large real-world topoi can be restrictive enough to block an interesting number of temporal properties.

2.1. Topoi: Formal Semantics

Formally, we say that a topoi is a directed, possibly cyclic graph G containing vertices and edges $\langle V, E \rangle$. E

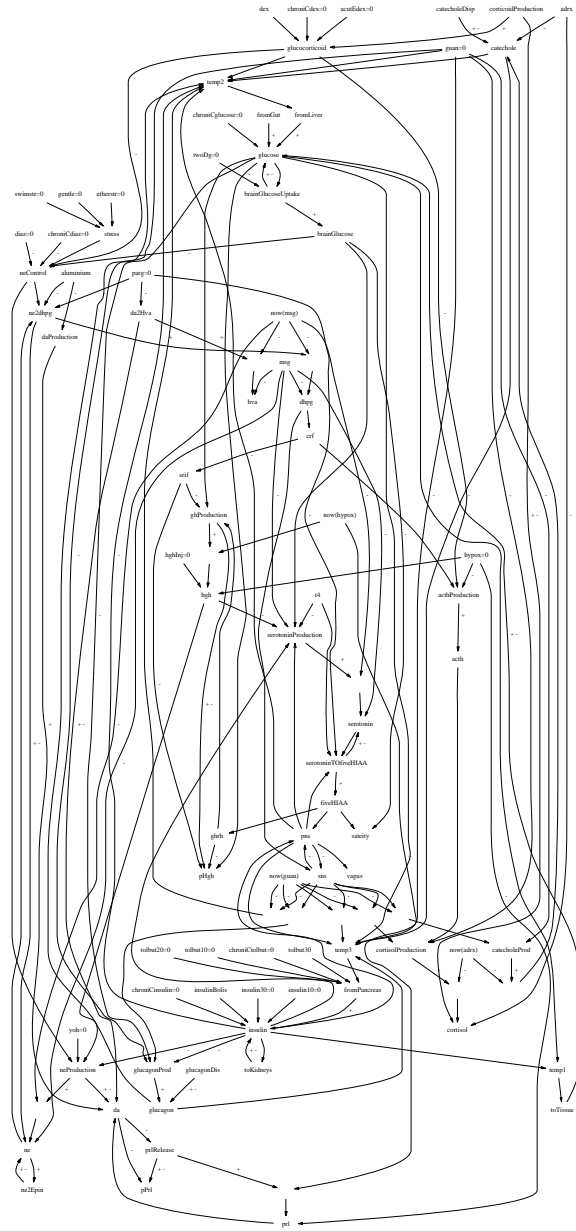


Figure 4. A large topoi with many loops.

are the connectors between variables and are one of a set of pre-defined types; e.g. $\overset{\pm}{\rightarrow}$ or $\overset{-}{\rightarrow}$. That is:

$$E_i = V_i \overset{\pm}{\rightarrow} V_j \text{ or } V_i \overset{-}{\rightarrow} V_j$$

$$G = \langle V, E \rangle$$

The vertices of a topoi can be assigned a finite number of values; e.g. *up*, *down* or *steady*. These values model the sign of the first derivative of these variables (i.e. the rate of change in each value). $X \overset{\pm}{\rightarrow} Y$ denotes that Y being *up*

or *down* could be explained by X being *up* or *down* respectively. That is:

$$V_i \overset{\pm}{\rightarrow} V_j \equiv \begin{cases} V_i \uparrow & \text{implies } V_j \uparrow \\ V_i \downarrow & \text{implies } V_j \downarrow \end{cases} \quad (1)$$

(where \uparrow and \downarrow denote up and down respectively.)

$X \overset{\pm}{\rightarrow} Y$ denotes that Y being *up* or *down* could be explained by X being *down* or *up* respectively. That is:

$$V_i \overset{\pm}{\rightarrow} V_j \equiv \begin{cases} V_i \uparrow & \text{implies } V_j \downarrow \\ V_i \downarrow & \text{implies } V_j \uparrow \end{cases} \quad (2)$$

Tacit in our topoi diagrams are conjunctions of influences. We can view topoi as influences splashing around pipes that connect tubs. Pairs of competing influences can cancel out. That is, we can explain the level of water in a tub remaining steady via conjunction of competing upstream influences; e.g.

$$\left(\begin{array}{c} (V_i \uparrow \text{ implies } V_j \uparrow) \\ \text{and} \\ (V_k \downarrow \text{ implies } V_j \downarrow) \end{array} \right) \text{ implies } \left(\begin{array}{c} (V_i \uparrow \wedge V_k \downarrow) \\ \text{implies} \\ (V_j = \text{steady}) \end{array} \right) \quad (3)$$

This formal semantics is sufficient to guide the translation of topoi for a formal model checker such as SPIN. Figure 5 shows the results of such a translation of Figure 2.iv. In this figure, all the nodes have the values *up*, *down*, *steady* and *unknown* (which is a placeholder for the initial conditions). Also, for convenience, all systems model inputs X are declared to be X_{chg} variables with values *arrived*, *left* denoting the differences between these variables in different experiments. For example, if we increase the injections of *dex*, then we also say that $\text{dex}_{\text{chg}} = \text{arrived}$.

2.2. Proving Formal Properties in Topoi

We can test topoi using libraries of expected or desired behavior. Such libraries can be quickly built via interviews with users. We have found it useful to structure these interviews in an OO framework. After generating use cases and particular scenarios, we ask our users to clarify exactly what are the expected inputs and required outputs for each scenario. This generates two artifacts. Firstly, it leads to a topoi graph describing how they think influences should propagate around a systems model. Secondly, it leads to the formulation of properties models of the form: *When I do this, I expect to see that* or, in the language of temporal logic used in (e.g.) SPIN:

$$\square (Inputs \rightarrow \diamond Outputs) \quad (4)$$

i.e. always the *inputs* lead, eventually, to the *outputs*.

We encounter problems if we use Equation 4 to check large topoi using standard model checkers. While SPIN

```
#define DOWN 0
#define STEADY 1
#define UP 2
#define UNDEF 3
#define ARRIVED 0
#define LEFT 1

byte chg_cold_swim = UNDEF /* chg_cold_swim = {ARRIVED,LEFT} */
byte chg_dex = UNDEF /* chg_dex = {ARRIVED,SWIM} */
byte cold_swim = UNDEF /* cold_swim = {DOWN,STEADY,UP} */
byte dex = UNDEF /* dex = {DOWN,STEADY,UP} */
byte temp = UNDEF /* temp = {DOWN,STEADY,UP} */
byte nna = UNDEF /* nna = {DOWN,STEADY,UP} */
byte acth = UNDEF /* acth = {DOWN,STEADY,UP} */
byte cortico = UNDEF /* cortico = {DOWN,STEADY,UP} */

active proctype smythe() {
  if
  ::dex == UNDEF -> dex = DOWN
  ::dex == UNDEF -> dex = STEADY
  ::dex == UNDEF -> dex = UP
  fi;
  if
  ::cold_swim == UNDEF -> cold_swim = DOWN
  ::cold_swim == UNDEF -> cold_swim = STEADY
  ::cold_swim == UNDEF -> cold_swim = UP
  fi;
  if
  ::chg_dex == UNDEF -> chg_dex = ARRIVED
  ::chg_dex == UNDEF -> chg_dex = LEFT
  fi;
  if
  ::chg_cold_swim == UNDEF -> chg_cold_swim = ARRIVED
  ::chg_cold_swim == UNDEF -> chg_cold_swim = LEFT
  fi;
  if
  ::chg_dex == ARRIVED -> temp = UP
  ::chg_dex == LEFT -> temp = DOWN
  fi;
  if
  ::chg_cold_swim == ARRIVED -> nna = UP
  ::chg_cold_swim == LEFT -> nna = DOWN
  fi;
  do
  ::(chg_cold_swim == ARRIVED && temp == UP) -> nna = STEADY
  ::(chg_cold_swim == LEFT && temp == DOWN) -> nna = STEADY
  ::temp == DOWN -> nna = UP
  ::temp == UP -> nna = DOWN
  ::temp == DOWN -> acth = UP
  ::temp == UP -> acth = DOWN
  ::nna == UP -> acth = UP
  ::nna == DOWN -> acth = DOWN
  ::acth == UP -> cortico = UP
  ::acth == DOWN -> cortico = DOWN
  ::cortico == UP -> temp = UP
  ::(cortico == UP && chg_dex == LEFT) -> temp = STEADY
  ::cortico == DOWN -> temp = DOWN
  ::(cortico == DOWN && chg_dex == ARRIVED) -> temp = STEADY
  ::(temp == UP && nna == UP) -> acth = STEADY
  ::(temp == DOWN && nna == DOWN) -> acth = STEADY
  od;
}
```

Figure 5. Figure 2.iv expressed in the PROMELA language used in SPIN model checker [11].

checks that Equation 4 holds over Figure 5 in less than a second, it can fail to terminate for larger systems models. In one study, we offered 40 properties of the form of Equation 4 to SPIN along with Figure 4 expressed in the same format as Figure 5. Given 100MB of maximum RAM, SPIN ran out of memory for most of the properties. We suspected that the search space was too big. Figure 4 contains 80 variables, each of which has at least the values *up*, *down*, *steady*, *undef*; i.e. total space of options at least of size ($4^{80} \approx 10^{48}$). In a second study, we reduced the size of the system by removing the *steady* values. This shrank the options to ($3^{80} \approx 10^{38}$). However, even in this reduced system, SPIN ran out of memory and failed to prove anything for 29 of the 40 properties [20].

In summary, while theoretically we can assess topoi us-

ing standard model checkers, in practice, this may not be feasible.

3. SP2: A Model Checker for Topoi

While general topoi defeat general-purpose model checkers, specialized model checkers can quickly check the temporal properties of a restricted class of topoi. Consider a topoi containing two-valued nodes connected by the “+” and “-” edges defined in Equation 1 and Equation 2. Such a topoi has *symmetric edges*; i.e. each edge comments on a connection of every upstream node’s value to every downstream node’s value. Menzies et.al. showed that when every edge of a symmetric topoi comments on all the values of its downstream vertices, then the state space rapidly *saturates* [17]. That is, the granularity of the time axis reduces to the number of values per variable in that theory. For example, in a systems model where every variable has only two values, everything that is reachable can be reached in two time ticks.

Using the result of Cohen et.al we have defined SP2, a specialized model checker for symmetric topoi [18, 20]. SP2 is a variant of Dijkstra’s shortest path algorithm [6]. The algorithm inputs a symmetric topoi with edge set E , node set V , and an initial set $S \subset V$. S contains some value assignments to some nodes and represents the initial conditions of the system. The algorithm outputs a set of edges Z with the following properties. Z is a collection of trees spanning all the nodes reachable from the inputs. For any reachable node z , Z contains the shortest topoi path from the inputs to z . The nodes of V spanned by Z are partitioned into two sets S' and T' , where no edge of Z passes from T' to S' and each set is consistent; i.e., will not contain both $x \uparrow$ and $x \downarrow$. In terms of temporal logic, S' means *now* and T' means *next*.

Elsewhere, we have proved that SP2 generates S' and T' correctly, and runs in $O(|V| + |E| \log |V|)$ time in the worst case [18]. SP2 is efficient due to its exploitation of saturation. While spreading out over the topoi, it maintains two sets of nodes: the *now* set (S') and the *later* set (T'). If the algorithm reaches a node that contradicts something else in the *now*, it moves the new node into the *later* set. The repeated application of this rule on a 2-spaced symmetric topoi results in a fast division of the nodes reachable from the initial conditions into the two sets S' and T' .

Using SP2, we can very quickly explore temporal properties that can be proved in two time ticks. A large range of interesting queries can be executed in two time ticks (but see below for a discussion on the properties that require more than two time ticks). Once S' and T' are generated, we can convert our temporal properties into set membership tests of these sets. Figure 6 and Figure 7 show conversion rules for common temporal properties.

exp	→	\Box exp	//always
		\Diamond exp	//eventually
		\bigcirc exp	//next
		exp W exp	// weak until
		exp U exp	//until
		exp \wedge exp	//conjunction
		exp \vee exp	//disjunction
		exp \rightarrow exp	//implication
		! exp	//negation
		x	//proposition
exp ₁ W exp ₂	→	(exp ₁ U exp ₂) \vee (\Box exp ₁)	
\Box x	→	(x \in now) \wedge (x \in later)	
\Diamond x	→	(x \in now) \vee (x \in later)	
\bigcirc x	→	(x \in later)	
x U y	→	(y \in now) \vee ((x \in now) \wedge (y \in later))	
x \wedge y	→	(x \in now) \wedge (y \in now)	
x \vee y	→	(x \in now) \vee (y \in now)	
x \rightarrow y	→	(y \in now) \vee ((x \notin now) \wedge (y \notin now))	
! x	→	(x \notin time)	
x	→	(x \in time)	
time	→	now later	
now	→	s'	
later	→	t'	

Figure 6. Rewrite rules for converting linear temporal logic expressions into set membership tests of SP2’s S' , T' .

SP2 offers two major other advantages over standard temporal reasoning. Firstly, it runs, terminates, returns Z , and then we perform set membership of Z to prove our properties. That is, we do not test for properties till *after* SP2 terminates. Hence, the inference time is not much affected by the complexity of the properties to be tested. Secondly, SP2 uses a shortest-paths tree to build its proofs. That is, when explaining how properties were reached, SP2 will generate the shortest explanation possible. Hence, a user of SP2 need not wade through mountains of trace files in order to understand how the properties were proved.

3.1. Experiments with SP2

Figure 8 shows a comparison of SPIN vs SP2 using properties of the form of Equation 4 and the systems model of Figure 4. Of the 40 properties which were analyzed by both SPIN and SP2, SPIN was able to return a verification result in only 11 out of 40 cases (27.5%) before running out of memory. In every case where SPIN did return a verification result, SP2’s result was in agreement.

Regarding computer resources, SP2 used less than 1% of the RAM required by SPIN. Also, in the case of the unprovable properties, SP2 terminated in less than a second CPU time while SPIN took much longer.

We mentioned earlier that one pre-experimental concern with informal topoi is that they are so under-defined that we

Figure 7.A: Absence properties: p is false

Property	LTL	SP2
Globally	$\Box(\neg p)$	$p \notin S' \wedge p \notin T'$
Before R	$\diamond R \rightarrow (\neg p \cup R)$	$R \in S' \vee R \notin T' \vee p \notin S'$
After Q	$\Box(Q \rightarrow \Box(\neg p))$	$(Q \in S' \vee (p \notin S' \wedge p \notin T')) \wedge (Q \in T' \vee p \notin T')$
Between Q and R	$\Box((Q \wedge \neg R \wedge \diamond R) \rightarrow (\neg p \cup R))$	$(Q \notin T' \vee p \notin T') \wedge (Q \notin S' \vee R \in S' \vee R \notin T' \vee p \notin S')$
After Q until R	$\Box(Q \wedge \neg R \rightarrow (\neg p \cup WR))$	$(Q \notin T' \vee R \in T') \wedge (Q \notin S' \vee R \in S' \vee (p \notin S' \wedge R \in T'))$

Figure 7.B: Existence properties: p becomes true

Globally	$\diamond(p)$	$p \in S' \vee p \in T'$
Before R	$\neg R \wedge (p \wedge \neg R)$	$(p \in S' \wedge (R \notin S' \vee p \in T')) \vee (R \notin S' \wedge p \in T' \wedge R \in T')$
After Q	$\Box(\neg Q \vee \diamond(Q \wedge \diamond p))$	$p \in T' \vee ((Q \notin S' \vee p \in S') \wedge Q \notin T')$
Between Q and R	$\Box(Q \wedge \neg R \rightarrow (\neg R \cup (p \wedge \neg R)))$	$Q \in S' \wedge Q \in T' \wedge p \in T' \wedge (R \in S' \vee p \in S' \vee R \notin T')$
After Q until R	$\Box(Q \wedge \neg R \rightarrow (\neg R \cup (p \wedge \neg R)))$	$Q \in S' \wedge Q \in T' \wedge (R \in T' \vee p \in T') \wedge (R \in S' \vee p \in S' \vee (p \in T' \wedge R \notin T'))$

Figure 7.C: Universality: p always true

Globally	$\Box(p)$	$p \in S' \wedge p \in T'$
Before R	$\diamond R \rightarrow (p \cup R)$	$R \in S' \vee (R \in T' \wedge p \in S')$
After Q	$\Box(Q \rightarrow \Box(p))$	$(Q \notin S' \vee (p \in S' \wedge p \in T')) \wedge Q \notin T' \vee p \in T'$
Between Q and R	$\Box((Q \wedge \neg R \wedge \diamond R) \rightarrow (p \cup R))$	$Q \in S' \vee R \in S' \vee R \notin T' \vee p \in S'$
After Q until R	$\Box(Q \wedge \neg R \rightarrow (p \cup WR))$	$Q \in S' \wedge Q \in T' \wedge (R \in S' \vee p \in S') \wedge (R \in T' \vee p \in T')$

Figure 7.D: Precedence: S precedes p

Globally	$\neg p \cup WS$	$S \notin S' \vee (p \in S' \wedge (p \notin T' \vee S \in T'))$
Before R	$\diamond R \rightarrow (\neg p \cup (S \vee R))$	$S \in S' \vee R \in S' \vee R \notin T' \vee p \notin S'$
After Q	$\Box(\neg Q \vee \diamond(Q \wedge (\neg p \cup WS)))$	$(Q \notin S' \wedge Q \notin T') \vee (Q \in S' \wedge (S \in S' \vee (p \notin S' \wedge (p \notin T' \vee S \in T')))) \vee (Q \in T' \vee p \notin T' \vee S \in T')$
Between Q and R	$\Box((Q \wedge \neg R \wedge \diamond R) \rightarrow (\neg p \cup (S \vee R)))$	Inexpressible: needs > 2 time ticks
After Q until R	$\Box(Q \wedge \neg R \rightarrow (\neg p \cup (S \vee WR)))$	$(S \in S' \vee R \in S' \vee (p \notin S' \wedge p \notin T')) \vee (p \in S' \wedge (S \in T' \vee R \in T')) \wedge (S \in T' \vee R \in T' \vee p \notin T')$

Figure 7.E: Response: S responds to p

Globally	$\Box(p \rightarrow \diamond S)$	$S \in T' \vee (p \notin T' \wedge (p \notin S' \vee S \in S'))$
Before R	$\diamond R \rightarrow (p \rightarrow (\neg R \cup (S \wedge \neg R))) \cup R$	Inexpressible: needs > 2 time ticks
After Q	$\Box(Q \rightarrow \Box(p \rightarrow \diamond S))$	$S \in T' \vee ((Q \notin S' \vee (p \notin T' \wedge (p \notin S' \vee S \in S'))) \wedge (Q \notin T' \vee p \notin T'))$
Between Q and R	$\Box((Q \wedge \neg R \wedge \diamond R) \rightarrow (p \rightarrow (\neg R \cup (S \wedge \neg R)))) \cup R$	Inexpressible: needs > 2 time ticks
After Q until R	$\Box(Q \wedge \neg R \rightarrow ((p \rightarrow (\neg R \cup (S \wedge \neg R))) \cup WR))$	$Q \in S' \wedge Q \in T' \wedge (R \in T' \vee S \in T' \vee p \notin T')$

Figure 7. Common temporal logic queries converted into set membership tests of SP2's S' , T' . This table was generated by applying the re-write rules of Figure 6 to a survey of common temporal logic queries [7, 8]. From [20].

	SPIN	SP2	Number of Cases
	??	proved	21
	??	unproved	8
	proved	proved	11
	unproved	unproved	0
	proved	unproved	0
	unproved	proved	0
RAM used (max)	100MB	< 1MB	

Figure 8. Proving properties of Figure 4 in SPIN and SP2. “??” denotes that SPIN did not terminate in 100MB of RAM.

could use them to infer any set of properties at all. Figure 8 shows that this is not always true. In the case of 8 of the 40 properties, SP2 could not prove them across the large under-defined topoi of Figure 4.

3.2. Limits to SP2

What are the practical implications of SP2's restrictions? We discuss below two important implications: restrictions of the properties that can be proved and the need for special tools to handle conjunctions.

Inadmissible Properties: Figure 9 shows Dwyer et.al.'s classification of over five hundred linear temporal logic (LTL) properties [7, 8]. Those properties divide into eight groups and each group contains the five temporal scopes seen in Figure 7; i.e. globally; before event R ; after event Q ; between events Q and R ; and after event Q until event R . $\frac{22}{40}$ of these scopes are expressible in terms of two time ticks [20]. The inexpressible scopes all require proving some ordering of > 2 events. By definition, such an ordering cannot be expressed using merely the two time intervals of S' and T' generated by SP2.

$$\begin{array}{l}
\text{Occurrence } \left(\frac{5+5+5+0=15}{20} \right) \\
\text{Order } \left(\frac{4+3+0+0=7}{20} \right)
\end{array}
\left\{ \begin{array}{l}
\text{Absence } \left(\frac{5}{5} \right) \\
\text{Universality } \left(\frac{5}{5} \right) \\
\text{Existence } \left(\frac{5}{5} \right) \\
\text{Bounded Existence } \left(\frac{0}{5} \right) \\
\text{Precedence } \left(\frac{4}{5} \right) \\
\text{Response } \left(\frac{3}{5} \right) \\
\text{Chain Response } \left(\frac{0}{5} \right) \\
\text{Chain Precedence } \left(\frac{0}{5} \right)
\end{array} \right.$$

Figure 9. Coverage of the Dwyer corpus of temporal properties by SP2. Each right-hand-side group of properties contains five scopes. Fractions denote how many of those scopes can be handled by SP2, as seen in Figure 7. Adapted from [20].

Figure 9 shows us that SP2-style inference on symmetric topoi can say more about the *occurrence* of a given event/state during system execution than about the *ordering* in which multiple events/states occur. It is a simple matter to detect the temporal properties that are inadmissible for SP2. All such properties require more than two time ticks; e.g. until operators nested to a depth greater than two such as:

$$\left(\text{day} = \text{sunday} \cup \left(\text{day} = \text{monday} \cup \text{day} = \text{tuesday} \right) \right)$$

Handling Conjunctions: Another problem is that symmetric topoi have no special knowledge of and-nodes. This can lead to some less-than-desirable results. Consider the following topoi:

$$\text{usability} \bar{\rightarrow} \text{flexibility} \bar{\leftarrow} \text{performance}$$

Equation 3 says that the conjunction of competing upstream influence can result in a steady value in a downstream variable; i.e.

$$\begin{array}{l}
\text{usability} \uparrow \rightarrow \text{and001} \\
\text{performance} \downarrow \rightarrow \text{and001} \\
\text{and001} \rightarrow \text{flexibility} = \text{steady} \quad (5)
\end{array}$$

where *and001* is an and-node especially created for this conjunction. A reasonable temporal interpretation of and-nodes is that all pre-conditions must appear before or at the same time as the post-conditions. Suppose we seek to prove $\text{flexibility} = \text{steady} \in S'$, but SP2 computes a node partition in which $\text{usability} \uparrow \in T'$ and $\text{performance} \downarrow \in T'$. We would like to be able to coax these pre-conditions back in time to S' such that they do not occur at a time that is later than $\text{flexibility} = \text{steady} \in S'$.

Another case where we want to coax edge weights is the *bad-and* situation. The rules of symmetric topoi require that if we create the edges shown in Equation 5, then we must also create the following complementary rules:

$$\begin{array}{l}
\text{usability} \downarrow \rightarrow \overline{\text{and001}} \\
\text{performance} \uparrow \rightarrow \overline{\text{and001}} \\
\overline{\text{and001}} \rightarrow \overline{\text{flexibility} = \text{steady}} \quad (6)
\end{array}$$

where \bar{X} is an invented node representing “all the other values of X ”. The addition of the nodes $\overline{\text{and001}}$ and $\overline{\text{flexibility} = \text{steady}}$ is required to ensure the symmetry properties upon which SP2 is dependent. However, they are just nonsense symbols that should never appear in any explanation of how certain inputs lead to certain properties. That is, pathways from inputs to properties should never include these nonsense symbols. Hence, if possible, SP2 should be ‘coaxed’ into producing shortest path trees in which these spurious nodes appear at the leaves.

SP2 contains a mechanism to implement such coaxing: each edge in the topoi is augmented with an edge weight, which SP2 uses to compute shortest paths – the length of a path is simply the sum of the weights of the edges along the path. At the core of SP2 is a priority queue. At runtime, the next edge to be explored is one of the edges with lowest weight within the queue. This means that by adjusting weights and re-running the algorithm, we can choose to explore edges at some earlier time or later time. Hence, to coax $\text{usability} \uparrow$ and $\text{performance} \downarrow$ into S' , we can adjust the weights upstream of those nodes. In coaxing, the weights can be adjusted arbitrarily, provided that the any symmetric pair of edges receives the same weight for both edges. Elsewhere [18] we define a set of minimal automatic edge adjustment heuristics which input SP2’s shortest path tree Z , the cut set C containing the edges that connect S' to T' and which outputs changes to the edge weights.

A major pre-experimental concern was that the nearly linear-time processing of SP2 could be followed by an indefinitely long coaxing process. After much experimentation, we can report that we have never seen this worst-case behavior in practice. In those experiments we used SP2 to explore randomly generated properties of the form of Equation 4 over dozens of randomly generated topoi graphs. We varied topoi fanout (2 to 6 edges per node) and the frequency of and-nodes (from 5% to 75%). Initially, each experiment was terminated when the percent of provable properties reached some plateau. In all the experiments, the plateau was reached after < 5 iterations of SP2+coaxing. Next, we compared the plateau reached after 10 coaxes and the plateau reached after 100 coaxes. No significant differences were observed in our experiments. Further, SP2 never used more than 1MB of memory or one minute of runtime. Our conclusion from these experiments is that the need for

heuristic coaxing does not diminish the time and space efficiency of SP2.

4. Related Work

We are hardly the first to explore formal methods for requirements engineering. For example, in the KAOS system [25], analysts generate a properties model by incrementally augmenting object-oriented scenario diagrams with temporal logic statements. Potentially, this research reduces the costs of formal requirements analysis by integrating the generation of the properties model into the rest of the system development. Our reading of the KAOS work is that while the resulting model may be more formal, the level of skill required to write the temporal logic can significantly increase the personnel cost. Further, the extra time required for the augmentation could increase the effect of the development brake. In other work, Schneider et.al. [22] explored reducing the manual modeling costs using *lightweight formal methods*. In the lightweight approach, only partial descriptions of the systems and properties models were constructed using the SPIN formal analysis tool [11]. Despite their incomplete nature, Schneider et.al. found that such partial models could still detect significant systems errors. While exciting research, this approach still incurs the personnel cost since scarce expertise is required to drive tools like SPIN.

Nor are we the first to explore optimizing temporal logic model checking. Elaborate techniques have been developed to tame the state space explosion problem. A sample of these techniques are described below. *Abstraction or partial ordering*: only use the part of the space required for a particular proof. Implementations exploiting this technique can restrain how the space is traversed [10], or constructed in the first place [22]. *Clustering*: divide the systems model into sub-systems which can be reasoned about separately [2]. *Meta-knowledge*: avoid studying the entire space. Instead, study only succinct meta-knowledge of the space. One example used an eigenvector analysis of the long-term properties of the systems model under study [12]. *Exploit symmetry*: Prove properties in some part of the systems model, then reuse those proofs if ever those parts are found elsewhere in the systems model [1]. *Semantic minimization*: Replace the space with some smaller, equivalent space. For example, the BANDERA system [4] reduces both the systems modeling cost and the execution cost via automatically extracting (slicing) the minimum portions of a JAVA program's bytecodes which are relevant to particular properties models.

While the above techniques have all proved useful in their test domains, they may not be universally applicable. Certain optimizations require expensive pre-processing, such as [12]. Also, these methods may rely on certain topo-

logical features of the system being studied. Exploiting symmetry is only useful if the system under study is highly symmetric. Clustering generally fails for tightly connected models. Further, for requirements engineering, systems like BANDERA are not suitable. BANDERA only works on implemented systems; that is, not until long after the requirements phase has ended.

Hence, in the general case, only small models can be tested. Further, these models must be precisely specified. In contrast, this work describes methods for quickly proving properties in large models that have been hastily sketched.

Another class of tools for optimizing searching the space of a program are the tools based on integer programming (for a good introduction and comparison notes on integer programming, see [3]). Preliminary experiments with integer programming in this domain are underway, with no clear results at this time.

5. Conclusion

We need better formal testing for our requirements. Applying formal methods can lead to an unacceptable brake on the development process. Cost-effective formal methods have to reduce the cost and skill involved in modeling systems and their properties. The cost of properties modeling can be reduced via temporal logic patterns. However, the execution cost of the resulting properties model may require expensive rework of the properties generated from the patterns.

In the specific case of requirements that can be mapped into symmetric topoi, we have shown that the systems modeling cost is reduced (since the topoi can be sketched quickly). For such symmetric topoi, we can reduce the execution cost for proving formal properties to time that is nearly linear on the number of edges and nodes in the topoi.

The combination of easy specification of properties and systems models implies that the personnel cost of formal modeling is reduced. This cost-reduction can only be achieved in domains where the systems model can be expressed as topoi and the properties model refers more to temporal occurrence properties than temporal ordering properties. We have argued that requirements engineering is one such domain.

Having built the SP2 engine, our next goal is the construction of a shell that exploits this engine. Our current research goal is the construction of the RAPTURE shell. RAPTURE exploits SP2 to enable the fast formal analysis of topoi-compliant descriptions of software systems.

ACKNOWLEDGEMENTS

This work was partially supported by NASA through cooperative agreement #NCC 2-979.

References

- [1] E. Clark and T. Filkorn. Exploiting symmetry in temporal logic model checking. In *Fifth International Conference on Computer Aided Verification*. Springer-Verlag, 1993.
- [2] E. Clark and D. E. Long. Compositional model checking. In *Fourth Annual Symposium on Logic in Computer Science*, 1989.
- [3] J. Corbett. An empirical evaluation of three methods for deadlock analysis of ada tasking programs. In *Proceedings of the 1994 International Symposium on Software Testing and Analysis (ISSTA)*, 1994.
- [4] J. Corbett, M. Dwyer, J. Hatcliff, S. Laubach, C. Pasarenu, Robby, and H. Zheng. Bandera: Extracting finite-state models from java source code. In *Proceedings ICSE2000, Limerick, Ireland*, pages 439–448, 2000.
- [5] R. Dieng, O. Corby, and S. Lapalut. Acquisition and exploitation of gradual knowledge. *International Journal of Human-Computer Studies*, 42:465–499, 1995.
- [6] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
- [7] M. Dwyer, G. Avrunin, and J. Corbett. Patterns in property specifications for finite-state verification. In *ICSE98: Proceedings of the 21st International Conference on Software Engineering*, May 1998.
- [8] M. B. Dwyer, G. S. Avrunin, and J. Corbett. A system specification of patterns. <http://www.cis.ksu.edu/santos/spec-patterns/>, 1997.
- [9] B. Feldman, P. Compton, and G. Smythe. Towards Hypothesis Testing: JUSTIN, Prototype System Using Justification in Context. In *Proceedings of the Joint Australian Conference on Artificial Intelligence, AI '89*, pages 319–331, 1989.
- [10] P. Godefroid. On the costs and benefits of using partial-order methods for the verification of concurrent systems (invited papers). In *The 1996 DIMACS workshop on Partial Order Methods in Verification, July 24-26, 1996*, pages 289–303, 1997.
- [11] G. Holzmann. The model checker SPIN. *IEEE Transactions on Software Engineering*, 23(5):279–295, May 1997.
- [12] Y. Ishida. Using global properties for qualitative reasoning: A qualitative system theory. In *Proceedings of IJCAI '89*, pages 1174–1179., 1989.
- [13] A. MacLean, R. Young, V. Bellotti, and T. Moran. Questions, options and criteria: Elements of design space analysis. In T. Moran and J. Carroll, editors, *Design Rationale: Concepts, Techniques, and Use*, pages 53–106. Lawrence Erlbaum Associates, 1996.
- [14] T. Menzies. *Principles for Generalised Testing of Knowledge Bases*. PhD thesis, University of New South Wales, 1995. Available from <http://tim.menzies.com/pdf/95thesis.pdf>.
- [15] T. Menzies. Applications of abduction: Knowledge level modeling. *International Journal of Human Computer Studies*, 45:305–355, 1996. Available from <http://tim.menzies.com/pdf/96abk1.pdf>.
- [16] T. Menzies. On the practicality of abductive validation. In *ECAI '96*, 1996. Available from <http://tim.menzies.com/pdf/96ok.pdf>.
- [17] T. Menzies, R. Cohen, S. Waugh, and S. Goss. Applications of abduction: Testing very long qualitative simulations. *IEEE Transactions of Data and Knowledge Engineering (to appear)*, 2001. Available from <http://tim.menzies.com/pdf/97iedge.pdf>.
- [18] T. Menzies, M. Houle, and J. Powell. Rapture/sp2: Efficient testing of temporal properties without search space explosion, 1999. NASA IV&V Facility Technical Report.
- [19] J. Mylopoulos, L. Cheng, and E. Yu. From object-oriented to goal-oriented requirements analysis. *Communications of the ACM*, 42(1):31–37, January 1999.
- [20] J. Powell. The rapture/sp2 approach to model checking: An explanation and viability experimentation, 1999.
- [21] D. Reifer. Software failure modes and effects analysis. *IEEE Transactions on Reliability*, pages 247–249, 1979.
- [22] F. Schneider, S. Easterbrook, J. Callahan, G. Holzmann, W. Reinholtz, A. Ko, and M. Shahabuddin. Validating requirements for fault tolerant systems using model checking. In *3rd IEEE International Conference On Requirements Engineering*, 1998.
- [23] S. B. Shum and N. Hammond. Argumentation-based design rationale: What use at what cost? *International Journal of Human-Computer Studies*, 40(4):603–652, 1994.
- [24] G. Smythe. Hypothalamic noradrenergic activation of stress-induced adrenocorticotropin (ACTH) release: Effects of acute and chronic dexamethasone pre-treatment in the rat. *Exp. Clin. Endocrinol. (Life Sci. Adv.)*, pages 141–144, 6 1987.
- [25] A. van Lamsweerde and L. Willemet. Inferring declarative requirements specifications from operational scenarios. *IEEE Transactions on Software Engineering, Special Issue on Scenario Management*, November 1998.