

A Graphical Interval Logic for Specifying Concurrent Systems

L. K. DILLON, G. KUTTY, L. E. MOSER, P. M. MELLIAR-SMITH, and
Y. S. RAMAKRISHNA
University of California

This article describes a graphical interval logic that is the foundation of a tool set supporting formal specification and verification of concurrent software systems. Experience has shown that most software engineers find standard temporal logics difficult to understand and use. The objective of this article is to enable software engineers to specify and reason about temporal properties of concurrent systems more easily by providing them with a logic that has an intuitive graphical representation and with tools that support its use. To illustrate the use of the graphical logic, the article provides some specifications for an elevator system and proves several properties of the specifications. The article also describes the tool set and the implementation.

Categories and Subject Descriptors: D.2.1 [**Software Engineering**]: Requirements/Specifications—*languages; tools*; D.2.2 [**Software Engineering**]: Tools and Techniques; F.3.1 [**Logics and Meanings of Programs**]: Specifying and Verifying and Reasoning about Programs—*mechanical verification; specification techniques*

General Terms: Human Factors, Languages, Verification

Additional Key Words and Phrases: Automated proof-checking, concurrent systems, formal specifications, graphical interval logic, temporal logic, timing diagrams, visual languages

1. INTRODUCTION

One of the great challenges facing today's software engineers is the development of correct programs for real applications. Recent advances in hardware reliability and fault tolerance technology can assure extremely low hardware failure rates for devices. Unfortunately, technologies for digital hardware design and software engineering have not matched this advance. The use of computers in many critical applications is now primarily limited by the reliability of system designs and implementations.

This research was partially supported by the NSF grant CCR-9014382 with cooperation from ARPA. An early version of the article was presented at the 14th International Conference on Software Engineering, May 1992 (Institution of Engineers Australia, IEEE Computer Society, Association of Computing Machinery, Institution of Radio and Electronic Engineers Australia, and Australian Computer Society).

Authors' address: Department of Computer Science, University of California, Santa Barbara, CA 93106.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1994 1049-331X/94/0400-0131 \$03.50

Often, the most critical real applications involve concurrency, which increases the difficulty of system development and validation. Modern methods of structured programming, which are quite effective for sequential programs, are notoriously inadequate for concurrent ones. Moreover, the nondeterminism inherent in applications that involve concurrency and the reactive character of those applications make them hard to test. Aggravating these problems is the need to explore large spaces of possible executions, which grow exponentially with the number of independent threads of control.

Formal methods for specifying and verifying systems can, in principle, offer a greater assurance of correctness than informal design/code checks or testing. Formal verification methods can demonstrate that a high-level design meets formally specified correctness requirements, thereby reducing the risk that faulty designs will be used as the basis for system development. Formal specifications are valuable for defining interfaces between independently developed software modules and for establishing software and interface standards. Because they provide a succinct and unambiguous statement of system requirements, formal specifications can potentially be analyzed for consistency, a particularly difficult and important problem for concurrent systems. Formal specifications can also be used during the selection of test data to suggest behaviors that should be tested and, later, to determine whether the execution of a test case is correct or erroneous. Thus, system developers can use formal specifications throughout the system life cycle to guide development, maintenance, and enhancement.

In practice, however, system developers seldom make significant use of formal specification and verification methods. We believe that this is due, in large part, to the reliance of those methods on mathematical formalisms that are difficult to understand and use. Formal specification and analysis methods must be made accessible to system designers and software engineers if they are to be used in the development of real-world systems. Users must be able to express the properties of the systems about which they wish to reason as naturally as possible and to confirm mechanically that the specifications, designs, testing criteria, and sample executions have the required properties.

Temporal logics [Barringer et al. 1984; Lamport 1990; Manna and Pnueli 1981; Wolper 1987] are well suited for specifying temporal properties of concurrent systems. Experience has shown, however, that specifications of even moderate-sized systems are too complex to be readily understood. This complexity stems chiefly from the need to establish the temporal context within which properties, such as bounded liveness and invariance, must hold. Interval logics [Halpern and Shoham 1991; Schwartz et al. 1983] address this problem by defining temporal intervals to represent such contexts. For example, to express the requirement that a process that releases a lock on a database must signal that it intends to enter the database before obtaining a new lock, an interval might be used to represent the activity of the system from the time that the process releases the lock until it acquires a new one; the process would be required then to signal its intension within the restricted context represented by the interval (bounded liveness).

Stylized pictures show complex timing relationships and dependencies often more clearly than linear textual representations of the same information. Such diagrams correspond more closely to common conceptualizations than does linear text. Often, software engineers draw timing diagrams, like those used to denote signal levels in hardware designs, when describing and reasoning about properties of systems. Even logicians fluent in temporal logic find timing diagrams helpful to explain the meanings of temporal logic formulas and to motivate lines of reasoning (e.g., Gabbay [1987]). However, in the absence of a formal semantics, timing diagrams cannot be used for rigorous analysis of system properties. Pictorial documentation is typically ad hoc and liable to ambiguous interpretation.

This article describes a visual temporal logic in which formulas resemble the informal timing diagrams familiar to designers of hardware systems and to software engineers. Graphical Interval Logic (GIL) has a formal model-theoretic semantics and is as expressive as propositional temporal logic with Until and without Next [Ramakrishna 1993]. It, thus, provides an intuitive and natural visual notation in which to express system specifications without sacrificing the benefits of a formal notation. A visual editor allows GIL specifications to be easily constructed and to be stored in and retrieved from files. The editor also provides a visual interface to a proof checker and model generator, which permit verification of temporal inferences.

The article provides an overview of GIL in Section 2. Then it presents sample specifications for an elevator system in Section 3 and shows, in Section 4, how a designer uses the specifications to reason about properties of the system. Section 5 describes the GIL tool set, and Section 6 provides an overview of the implementation. Related work is discussed in Section 7, with conclusions and future work presented in Section 8. The Appendix provides a model-theoretic semantics for the logic.

2. GRAPHICAL INTERVAL LOGIC

When reasoning about temporal properties exhibited by a concurrent system during a computation, it is convenient to regard the system as passing through a sequence of states. To model a nonterminating computation, the state sequence must be infinite. A terminating computation can, likewise, be modeled with an infinite state sequence by repeating, or *stuttering*, the final state. This permits a concurrent system to be identified with the set of infinite state sequences that represent its potential computations. GIL specifications for a system describe properties of legal state sequences. That is, the specifications must hold at the first state of every infinite state sequence that represents a computation of the system. We adopt a total-order model of computation, rather than a partial-order model, which has some advantages for representing causality in concurrent systems [Pratt 1986], because total orders are more readily abstracted into meaningful “intervals” that can be represented pictorially at an appropriately high level.

A GIL formula is evaluated at a state in an infinite sequence of states. Infinite state sequences, therefore, provide the *contexts* within which formulas

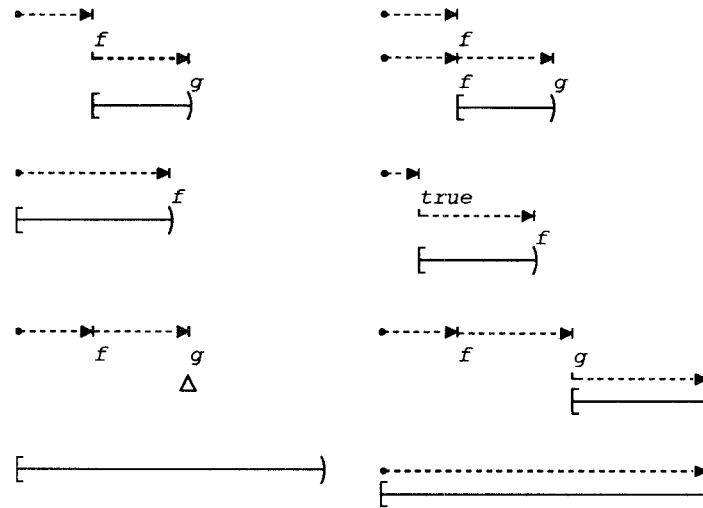
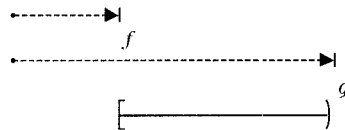


Fig. 1. Examples of some derived intervals (left column) and their definitions (right column).

The extent of an interval is specified by means of a pair of search patterns, which designate searches for locating the left and right ends of the interval. Both searches begin from the same point. We, therefore, draw them one beneath the other, with their start points aligned horizontally. The interval determined by the searches is drawn directly beneath the searches, its left end aligned horizontally with the point located by the first search pattern and its right end aligned horizontally with the point located by the second search pattern. For example:



The interval starts at the point located by the search for its left end and extends up to, but does not include, the point located by the search for its right end. The above diagram, thus, represents the interval that starts with the first point at which f holds and ends just prior to the first point at which g holds. The interval cannot be constructed if either search fails or if the interval is empty (i.e., the point specified by the first search pattern does not precede that specified by the second search pattern).

Figure 1 illustrates conventions that simplify the representation of several common types of intervals. The first abbreviation, in which a single search pattern specifies the extent of an interval, is permitted when the search for the interval's left end is a prefix of the search for its right end. Thus, the interval in the first example begins with the first point at which f holds and extends up to, but does not include, the next point at which g holds. The interval cannot be constructed if f does not hold at any future point, if g does

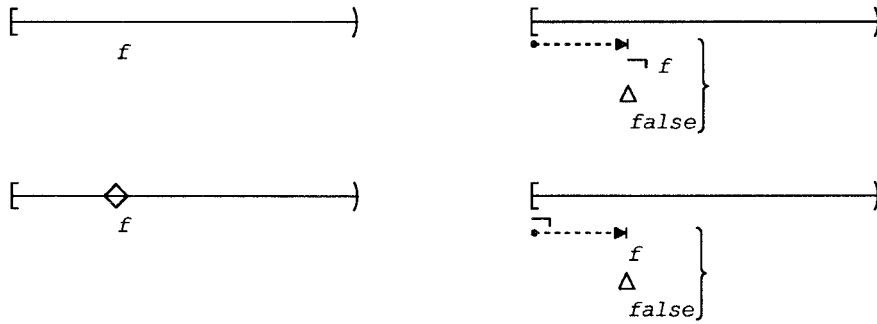


Fig. 2. Representation of invariants (top left) and eventualities (bottom left), and their definitions (right column).

Both layout and precedence rules determine the grouping of operations. GIL formulas obey a variation of Landin’s offside rule [Landin 1966], which requires that every token of a formula lie in the lower right quadrant determined by the upper left corner of the smallest rectangle that contains its first token. The first token that does not obey this rule, called an offside token, terminates the parse of a formula. The precedence of operators (from high to low) is: negation, conjunction, disjunction, implication, and equivalence. Binary operators associate from left to right. Right braces delimit interval formulas and permit explicit grouping of operations.

The weak Until operator U of propositional temporal logic (PTL) is expressed in GIL as follows.

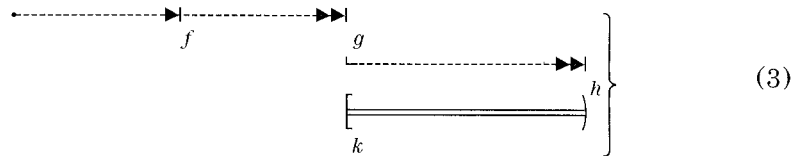
$$\begin{array}{c}
 \text{-----} \rightarrow \\
 \left. \begin{array}{l} \neg f \\ \vee \\ g \\ \Delta \\ g \end{array} \right\} \quad (2)
 \end{array}$$

The formula asserts that *g* holds at the first point where either *f* does not hold or *g* does hold, unless no such point is located. In the latter case, *f* (as well as $\neg g$) holds at all future points. In other words, *f* holds at least until *g* holds.

GIL provides a special syntax for invariants and eventualities. To assert that a formula holds at every point in an interval, the formula is drawn indented directly below the interval. To assert that a formula holds at some point within an interval, the formula is drawn left justified directly below a diamond \diamond drawn on the interval. Figure 2 shows these conventions and their definitions. The definition (top right) of the invariant notation (top left) can be understood as follows. Since *false* does not hold at any point of a context, the point formula holds precisely if the search to $\neg f$ fails, i.e., if *f* holds at all future points. Similarly, the formulas in the bottom row assert that *f* holds at some future point.

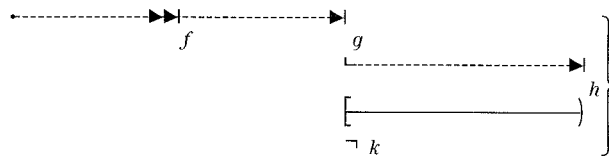
As noted above, an interval formula holds vacuously if any search performed in locating the ends of the interval fails or if the interval is empty. Thus, an interval formula is implicitly predicated on locating all search targets and on locating the left end of the interval before locating the right end. We, therefore, refer to the search operator and interval operator described above as *weak* operators.

GIL also provides *strong* versions of these operators, which are useful in specifications and in expressing the negations of interval formulas. A double arrowhead denotes a strong search and asserts that the search succeeds unless some prior weak search fails. A double line denotes a strong interval and asserts that the point located by the search for the interval's left end strictly precedes that located by the search for its right end unless some search fails. For instance,



holds by default if the search for f fails. However, if this search succeeds, then the formula requires that the subsequent searches for g and h succeed, that the interval is not empty (h does not hold at the point located by the second search), and that k holds at the first state of the interval.

The dual of an interval formula is obtained by changing the senses (strong to weak and weak to strong) of the interval modality and of the searches for the ends of the interval. This dual relationship implies that negation can be moved into an interval formula by changing the senses of the interval and of its searches. For instance, the negation of (3) is equivalent to



The Appendix gives formal definitions for the syntax and semantics of GIL.

3. AN EXAMPLE SPECIFICATION

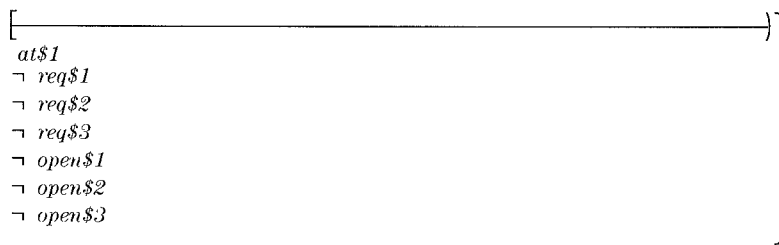
We present a GIL specification for an elevator system to illustrate the ideas in the previous section. The example includes specifications of basic safety and liveness requirements, and also of more complex fairness requirements.

For simplicity, we consider an elevator with three floors. The specification makes use of the following state predicates, for $n = 1, 2, 3$. The predicate *at* n is true when the elevator is at floor n and *false* when it is not. The predicate *goingup* models a physical switch whose setting, when the elevator

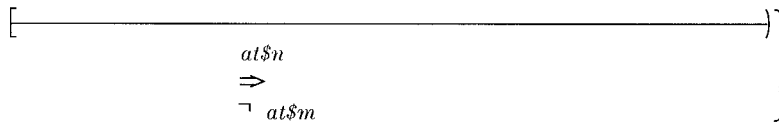
leaves a floor, determines the direction of travel: up, if *goingup* is *true*, and down, if *goingup* is *false*. The predicate *open* n is *true* when the doors to the elevator are open at floor n and *req* n is *true* when there is an outstanding request for service at floor n . Finally, when the elevator is at the second floor, *arriveup* indicates whether it was going up or down when it arrived.

Specifications are read from top to bottom and left to right. By convention, we begin each specification with a context line, which represents a legal execution of the system. The first specification expresses initial requirements, and the remaining specifications describe system invariants. We associate labels (shown in bold) with specifications for reference purposes below.

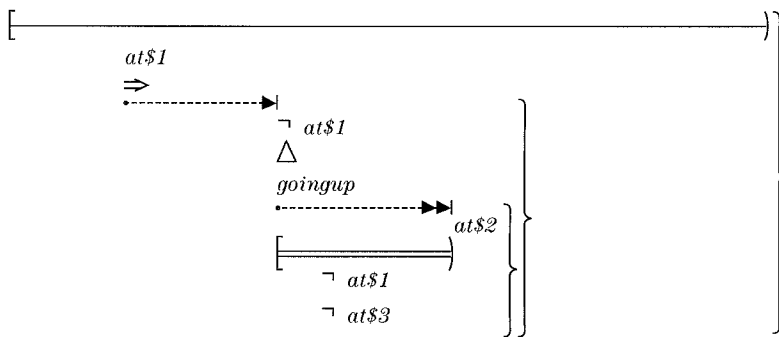
Init. The elevator begins operation at the first floor, all doors are closed, and there are no requests for service.



At n m , $1 \leq n < m \leq 3$. The elevator is never at two different floors simultaneously.



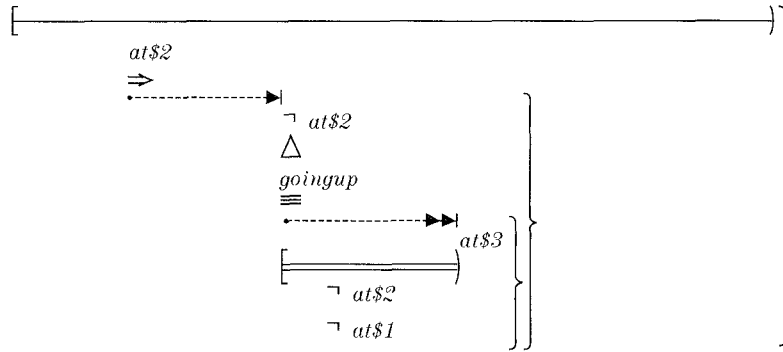
UpFrom\$1. The elevator goes up when it departs the first floor, arriving at the second floor without first visiting any other floors.



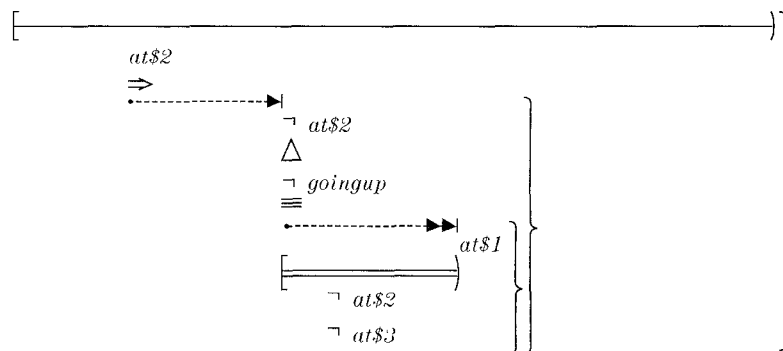
The invariant in this formula is predicated on locating a point at which the elevator has just left the first floor. The specification asserts that the elevator is going up at every such point and that it reaches the second floor before

either of the other floors. The strong search requires that the elevator eventually arrives at the second floor and the strong interval requires that it does not arrive there immediately upon leaving the first floor, but takes some time to do so.

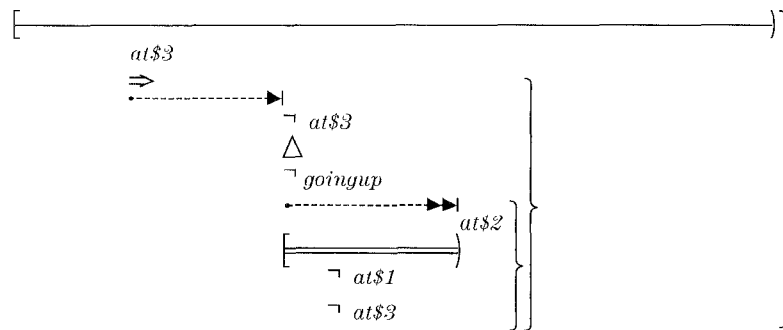
UpFrom\$2 The elevator goes up when it departs the second floor precisely if it goes directly to the third floor.



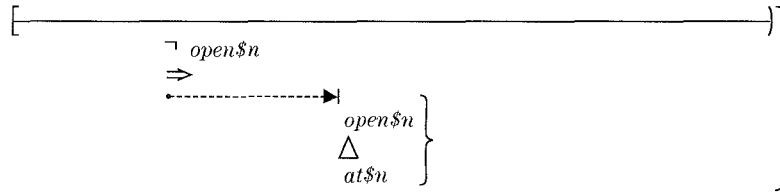
DownFrom\$2 The elevator goes down when it departs the second floor precisely if it goes directly to the first floor.



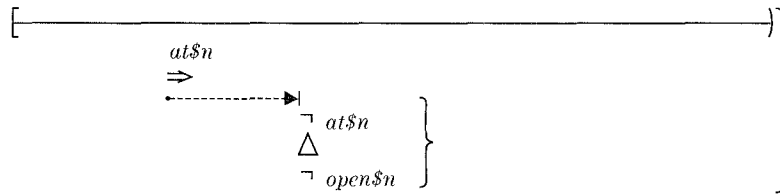
DownFrom\$3. The elevator goes down when it departs the third floor, arriving at the second floor without visiting any other floors first.



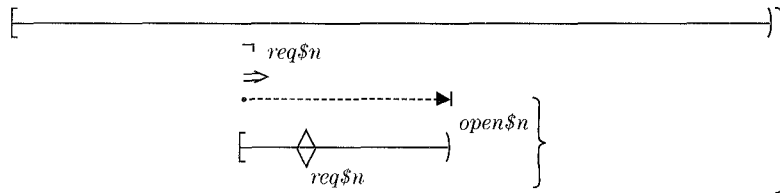
SafeOpen n , $n = 1, 2, 3$. The doors open at a floor only when the elevator is at the floor.



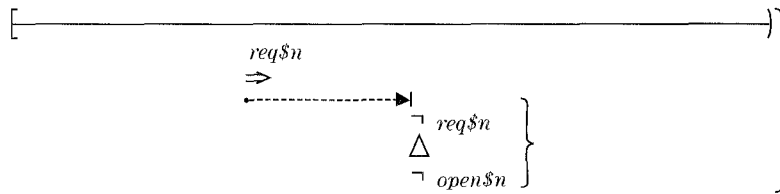
SafeDepart n , $n = 1, 2, 3$. The elevator departs a floor only when the doors at the floor are closed.



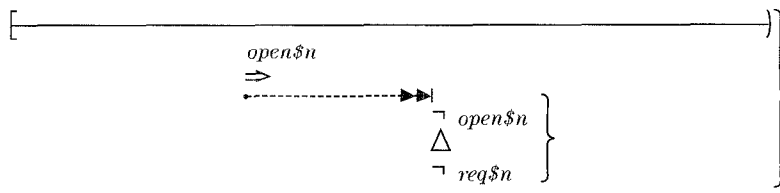
ReqService n , $n = 1, 2, 3$. The doors open at a floor only in response to a request for service at the floor.



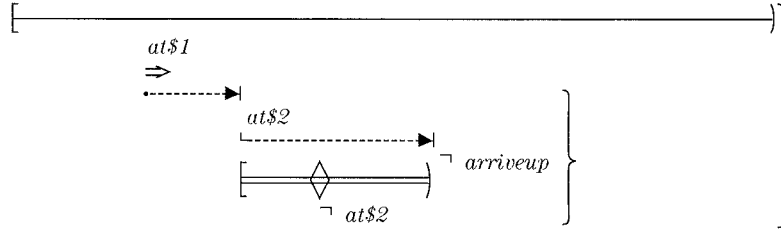
WaitService n , $n = 1, 2, 3$. A request for service at a floor is only canceled if the floor is being serviced (the doors are open).



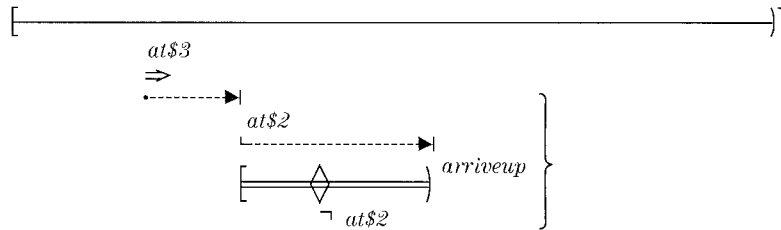
CancelService n , $n = 1, 2, 3$. The doors do not remain open indefinitely, and all requests for service at the current floor are canceled when they close.



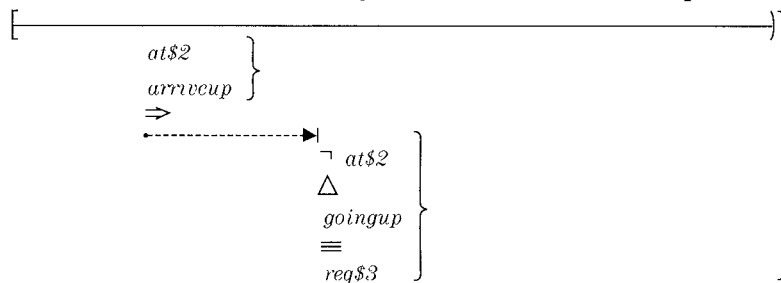
ArriveUp. Whenever the elevator arrives at the second floor from the first floor *arriveup* is *true*, and it remains *true* at least until the elevator departs the second floor.



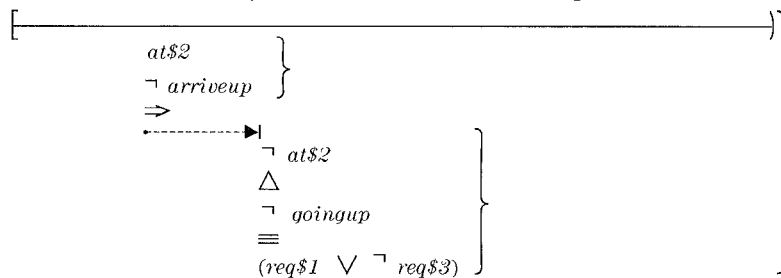
ArriveDown. Whenever the elevator arrives at the second floor from the third floor *arriveup* is *false*, and it remains *false* at least until the elevator departs the second floor.



ContinueUp. If the elevator is going up when it arrives at the second floor, it continues going up when it departs the second floor precisely if someone requires service at the third floor by the time the elevator departs.



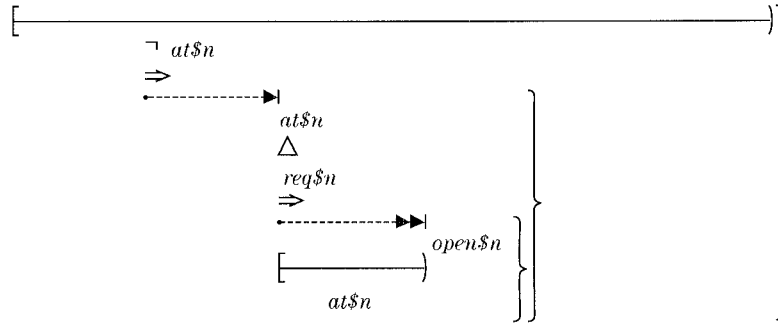
ContinueDown. If the elevator is going down when it arrives at the second floor, it continues going down when it departs the second floor precisely if someone requires service at the first floor or no one requires service at the third floor by the time the elevator departs.



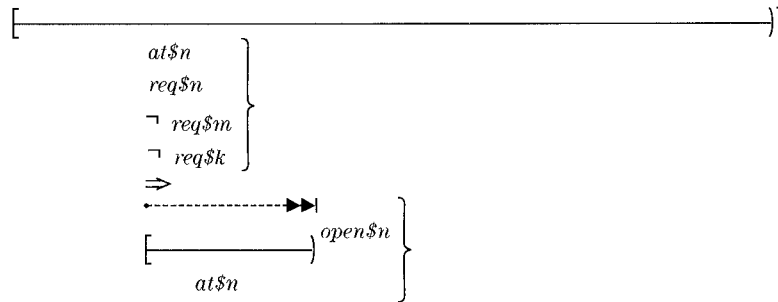
ContinueUp and ContinueDown require that, once the elevator starts traveling in a given direction, it changes directions only if no one requires service at

a floor in that direction. The disjunction in ContinueDown permits (but does not require) the first floor to act as the “home floor,” to which the elevator can return when it is idle.

ServeReqsOnArrival n , $n = 1, 2, 3$. If a passenger requests service at a floor by the time the elevator reaches the floor, the elevator opens its doors before departing the floor.

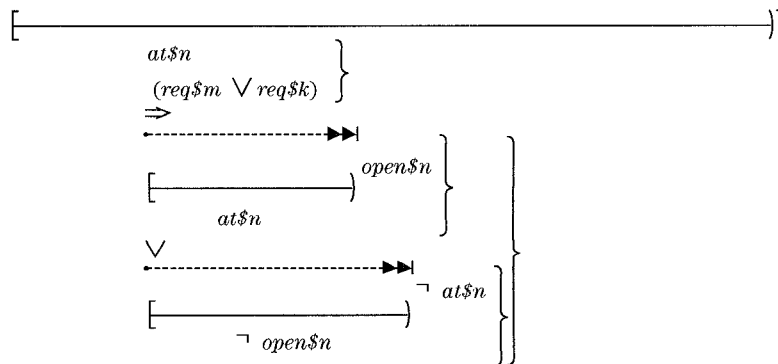


ServoNoConflict n , $n = 1, 2, 3$. If a passenger needs service at a floor while the elevator is at the floor and no one needs service at another floor, then the elevator opens its doors before departing the floor.

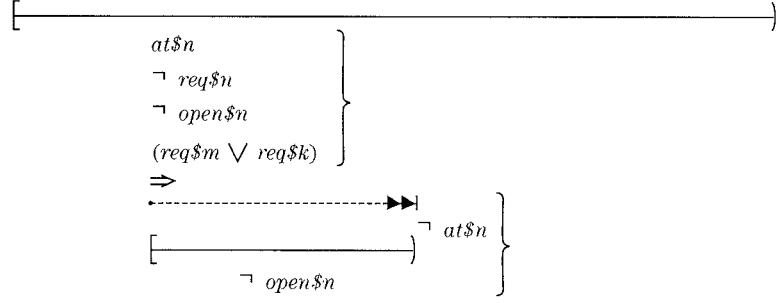


We use m and k in this and the remaining specifications to denote the other floors, that is, $\{m, k\} = \{1, 2, 3\} \setminus \{n\}$, and $m < k$.

ServeOrDepart n , $n = 1, 2, 3$. If the elevator is at a floor and a passenger requires service at a different floor, then either the doors open at the current floor, or the elevator departs the floor (without first opening the doors).



NoServeDepart $n, n = 1, 2, 3$. The elevator departs a floor without first opening its doors whenever no passenger requires service at the floor, the doors are closed, and someone needs service at another floor.



The last four specifications ensure that the elevator makes progress and that it services floors in a timely fashion. If a passenger requests service at a floor by the time the elevator arrives there, the appropriate ServeReqsOnArrival specification guarantees that the elevator stops at the floor before traveling on to other floors. Similarly, the ServeNoConflict specifications ensure that the elevator services a request if both the elevator is at the floor needing service and no one is waiting for service at another floor. However, if a passenger requires the elevator at some other floor, the appropriate NoServeDepart specification ensures that, once the current floor is serviced, the elevator departs the floor without servicing any additional requests for service at the current floor that may be made in the interim. The ServeOrDepart specifications prevent the elevator from sitting idly at a floor if other floors require service.

For purposes of comparison, we show how UpFrom\$2 and ArriveUp are expressed in PTL. Nested until operations are required to limit the scope of subformulas to the appropriate contexts.

UpFrom\$2:

$$\Box(at\$2 \Rightarrow at\$2U[\neg at\$2 \wedge \{goingup \equiv (\Diamond at\$3 \wedge \neg at\$3 \wedge (\neg at\$1 \wedge \neg at\$2)Uat\$3)\}])$$

ArriveUp:

$$\Box(at\$1 \Rightarrow \neg at\$2U[at\$2 \wedge \{\Box arriveup \vee \neg(at\$2U\neg arriveup)\}])$$

4. GRAPHICAL PROOFS OF SYSTEM PROPERTIES

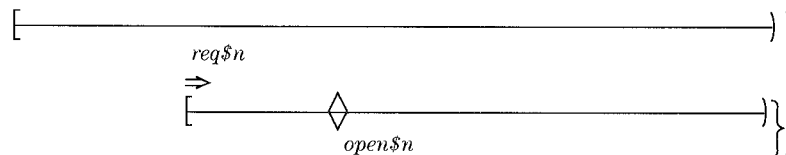
An important benefit of formal specifications is that they can be analyzed for potential consequences. Analysis can demonstrate that specifications express higher-level system requirements correctly and can help the designer learn more about the system under development. If analysis reveals that the specifications admit computations that violate requisite properties, the specifications are incomplete or in error. On the other hand, when desired properties can be proven from the specifications, the designer gains confidence that they provide a complete and accurate description of the system to be built.

The following are examples of properties that are required of the elevator system. The first requirement is one of many safety properties that a designer might wish to establish. The second expresses a minimal fairness requirement.

Safe n , $n = 1, 2, 3$. The elevator must be at a floor for its doors to be open there.



Service n , $n = 1, 2, 3$. The elevator eventually responds to a request for service.

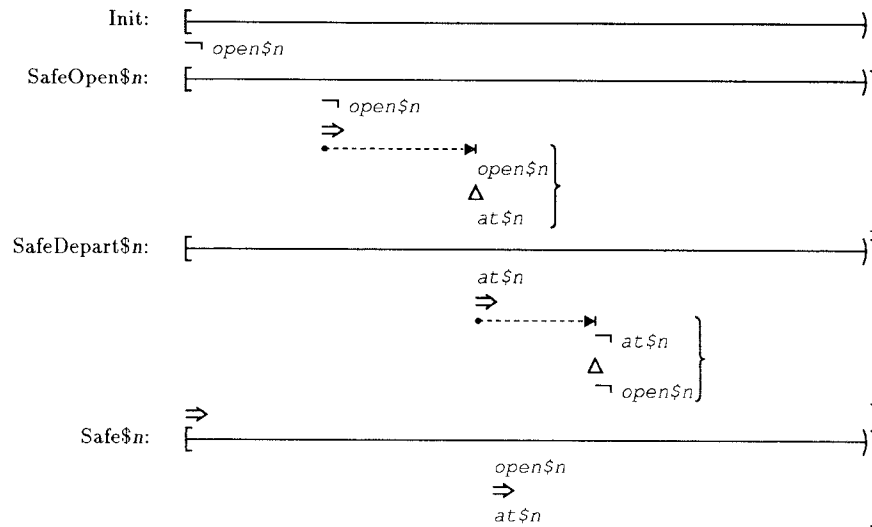


The specifications for a system express temporal constraints on legal computations. Thus, the system satisfies a requirement r if the conjunction s of the specifications implies r , or, equivalently, the implication $s \Rightarrow r$ is valid. In principle, the GIL proof checker can check the validity of this inference. However, in practice, theorem proving requires human assistance in order to be computationally feasible. The designer provides this assistance in our proof method by breaking down a complex proof into inferences that are small enough for the GIL proof checker to validate.

A major advantage of a visual logic, such as GIL, is that a proof can be represented using pictures that show the temporal flow of the argument. The graphical representation of the time line allows one to align appropriate points in the picture. Such alignment helps the designer see the points at which invariants are being instantiated, the intervals and points being aligned to establish bounded liveness and invariance conditions, the relationships between different points and intervals, and so on. These visual cues can be extremely helpful both for constructing proofs and for discovering potential fallacies. This “syntactic sugar” has no semantic content in the proofs below, although we are investigating a technique that will permit the designer to use alignment to specify orderings of points within a specification.

The alignment and ordering of points on a time line has other uses as well. For instance, the GIL tool set provides a model generation facility for producing a counterexample in the case that an inference is invalid. The counterexample can be displayed as a sequence of states or as a timing diagram. Aligning the states in the implication appropriately with this counterexample can help illustrate the fallacy in the inference.

The proof of **Safe** n in Figure 3 uses alignment to highlight the underlying correctness argument. The annotations alongside the picture show the speci-

Fig 3. Proof of $\text{Safe}\$n$, $n = 1, 2, 3$.

fications used in the proof. As shown by the annotations, $\text{Safe}\$n$ is proved from

- Init, which asserts that the doors are not open at floor n when the system starts up,
- $\text{SafeOpen}\$n$, which asserts that the elevator is at floor n when the doors first open at the floor, and
- $\text{SafeDepart}\$n$, which asserts that the doors have closed by the time the elevator departs the floor.

Aligning the invariant in $\text{SafeDepart}\$n$ with the point located by the search to $open\$n$ in $\text{SafeOpen}\$n$ highlights the fact that, when the invariant is evaluated at this point, it guarantees that $at\$n$ holds continuously from the (arbitrary) point at which $open\$n$ becomes *true* at least until $open\$n$ is *false*.

The proof of $\text{Service}\$n$ is too complex to be accomplished in a single step. Figure 4 shows the last step in the proof. As shown by the annotations alongside the figure, the final deduction makes use of several specifications and of the intermediate result $\text{Arrive}\$n$, which is established independently as another step of the proof. The reasoning illustrated by the picture can be understood as follows. If $req\$n$ holds at a point in a computation, but becomes *false* at a future point, then $\text{WaitServe}\$n$ ensures that the invariant in $\text{Service}\$n$ holds at this point. To highlight this reasoning, we align the invariants in $\text{WaitService}\$n$ and $\text{Service}\$n$ and align the points at which $open\$n$ is asserted to hold. The remaining premises establish $\text{Service}\$n$ in the case that $req\$n$ holds continuously from some point in a computation. We use $\text{Arrive}\$n$ to deduce that there is a future $at\$n$ -point. The $at\$n$ -point is purposely positioned within the span of the search arrow in $\text{WaitService}\$n$ to remind the reader that we are interested in the case where $at\$n$ is *true*

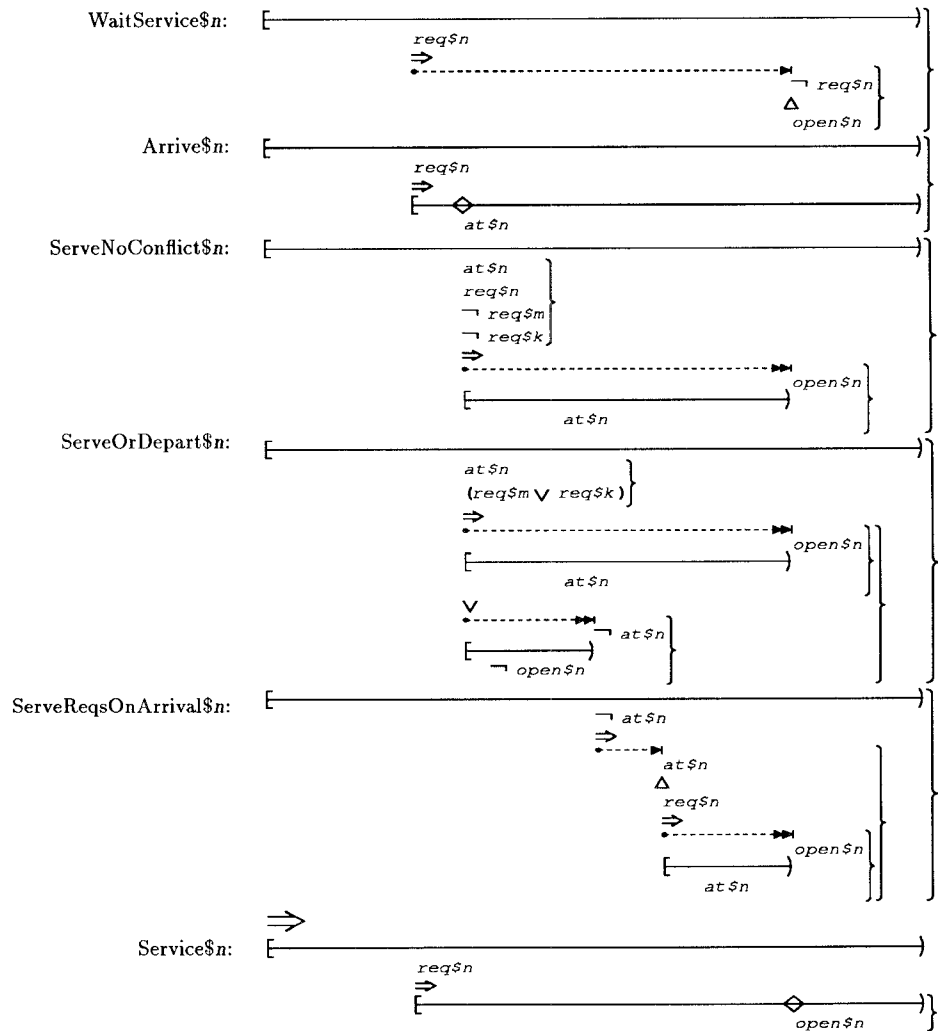


Fig. 4. Final deduction in the proof of Service_n , $n = 1, 2, 3, \{m, k\} = \{1, 2, 3\} \setminus \{n\}$, and $m < k$.

before req_n is *false*. The next three premises represent a case split. The invariant in ServeNoConflict_n establishes the invariant in Service_n in the case that req_m and req_k are both false at the at_n -point. The invariant in ServeOrDepart_n establishes Service_n in the case that req_m or req_k is *true* at the at_n -point and at_n holds throughout the future. Finally, the invariant in $\text{ServeReqsOnArrival}_n$ when instantiated at the next $\neg \text{at}_n$ -point establishes the required invariant in the case that both req_m or req_k is *true* at the at_n -point and there is some future point at which at_n is *false*.

Figure 5 shows how a complex proof is split into more manageable steps by case analysis. It represents the last step in the proof of Arrive_n . In the same

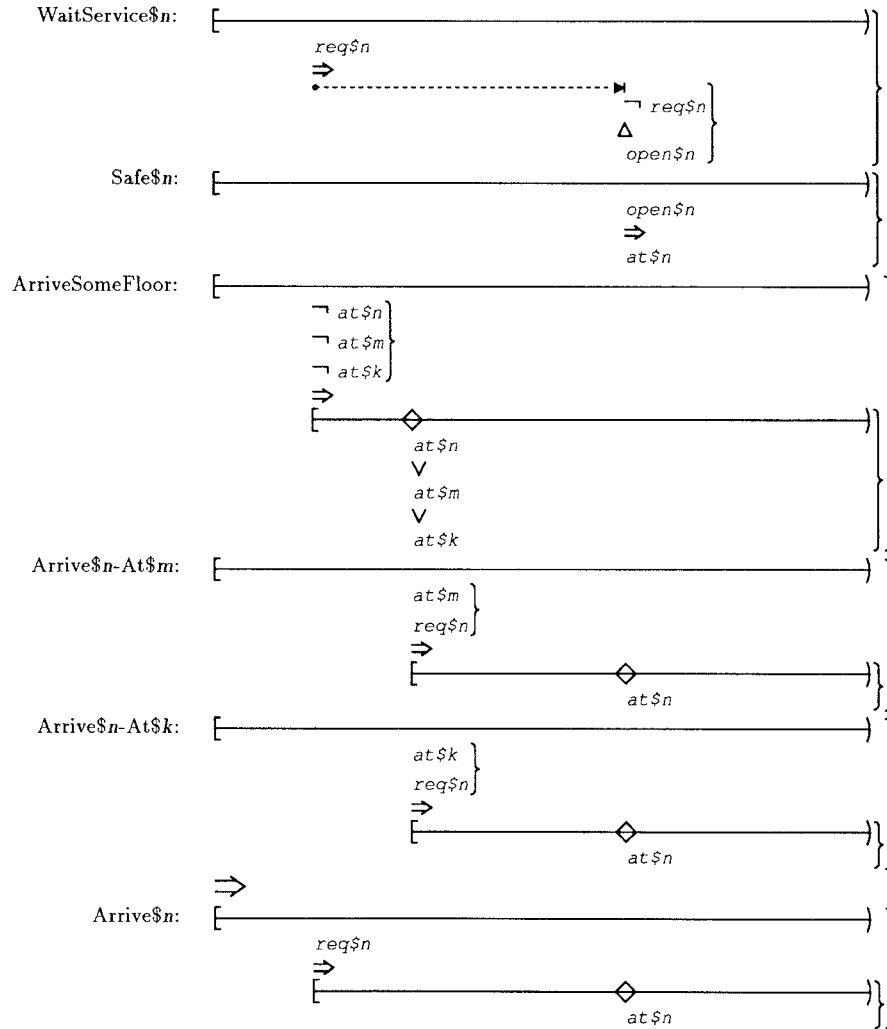


Fig. 5. Final deduction in the proof of Arrive\$n\$, $n = 1, 2, 3, \{m, k\} = \{1, 2, 3\} \setminus \{n\}$, and $m < k$.

style as the previous example, WaitService\$n\$ and Safe\$n\$ establish the required invariant when req\$n\$ is false at some future point. The remaining premises are required when req\$n\$ holds continuously from some point in a computation. ArriveSomeFloor represents a progress requirement needed to ensure that the elevator does not remain in transit indefinitely, but arrives eventually at some floor. This permits the proof to be reduced to the two cases represented by Arrive\$n-At\$m\$ and Arrive\$n-At\$k\$, which assert, respectively, that the elevator arrives eventually at floor \$n\$ from floor \$m\$ and that it arrives eventually at floor \$n\$ from floor \$k\$.

Proofs of ArriveSomeFloor, Arrive\$n-At\$m\$ and Arrive\$n-At\$k\$ are required to complete the proof of Service \$n\$. The requirement ArriveSomeFloor follows

directly from the specifications Init , $\text{UpFrom}\$1$, $\text{UpFrom}\$2$, $\text{DownFrom}\$2$, and $\text{DownFrom}\$3$. The high-level strategy used in the proofs of $\text{Arrive}\$n\text{-At}\m and $\text{Arrive}\$n\text{-At}\k is first to show that the specifications ensure the elevator does not remain at a floor indefinitely if it is needed at a different floor. These “departure results” and $\text{UpFrom}\$1$ ensure $\text{Arrive}\$2\text{-At}\1 . Similarly, the departure results and $\text{DownFrom}\$3$ imply $\text{Arrive}\$2\text{-At}\3 . For the proof of $\text{Arrive}\$1\text{-At}\2 and $\text{Arrive}\$1\text{-At}\3 , the departure results and the specifications are first used to show that, if the elevator is traveling down when it arrives at the second floor, it arrives at the first floor eventually. $\text{Arrive}\$1\text{-At}\3 follows easily from this, the departure results, and the specifications. Finally, we use $\text{Arrive}\$1\text{-At}\3 , the departure results, and the specifications to show that the elevator arrives eventually at the first floor if it is traveling up when it reaches the second one. The proofs of $\text{Arrive}\$3\text{-At}\2 and $\text{Arrive}\$3\text{-At}\1 parallel those of $\text{Arrive}\$1\text{-At}\2 and $\text{Arrive}\$1\text{-At}\3 . The departure results are established by a straightforward (but tedious) case analysis. The full proof is given in Dillon et al. [1993] in the form of intermediate lemmas and annotated proof trees.²

5. THE GRAPHICAL INTERVAL LOGIC TOOL SET

We have built a prototype GIL tool set to demonstrate proof-of-concept and permit experiments with the logic. The prototype includes a visual editor that allows specifications to be easily constructed and to be stored in and retrieved from files, a proof checker that checks mechanically the validity of temporal inferences, and a model generator that exhibits state sequences over which formulas hold. This section provides a brief overview of the GIL tool set.

Figure 6 shows the appearance of the interface of the GIL editor (GILED). Formulas are edited on a canvas, which comprises the main region of the display. The canvas in Figure 6 contains a template for creating a new specification. The template consists of an outer context interval and a box, positioned automatically below the start of the interval, that represents a formula that has yet to be defined. The designer uses the mouse during editing to select formulas in the canvas and editing operations; the box is selected (indicated by shading) in the example. Scroll bars permit the canvas to be scrolled for viewing large formulas.

The buttons in the panel on the lower left side of the display correspond to GIL primitives. The Text button allows a box to be replaced with a state predicate. The remaining buttons in the lower left panel specify GIL operators that apply to appropriate formulas. First are buttons corresponding to the four temporal operators: the interval [—] , eventuality \diamond , invariant \square , and point \triangle operators.³ The last five buttons correspond to the propositional operators; disjunction \vee , conjunction \wedge , negation \neg , implication \Rightarrow , and

²This technical report can be obtained by anonymous ftp from directory/pub at ftp.cs.ucsb.edu.

³As noted in Section 2, the eventuality, invariant, and point operators are derived from the interval operator. However, they correspond to common conceptualizations that are distinguished by the graphical syntax for visualization purposes.

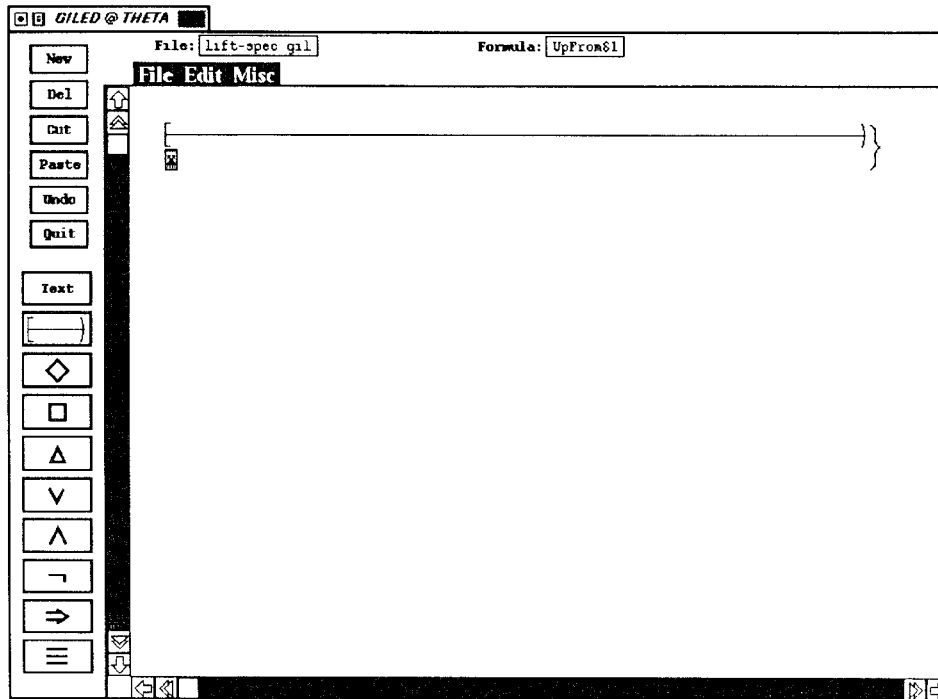
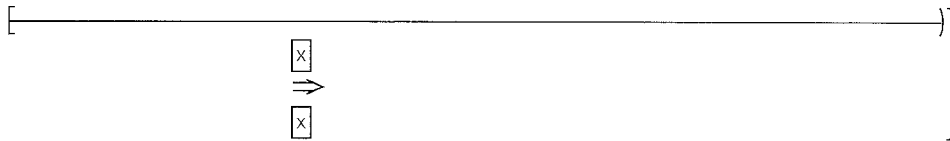


Fig. 6. GILED user interface.

equivalence \equiv . The buttons in the upper left panel provide language-independent editing operations. Commands to override the default layout of formulas and commands for storing and retrieving formulas are found in the Edit and File pull-down menus. Using commands provided in the Misc pull-down menu, the proof checker is invoked and models are displayed. Models are displayed graphically in an accompanying window (not shown in Figure 6).

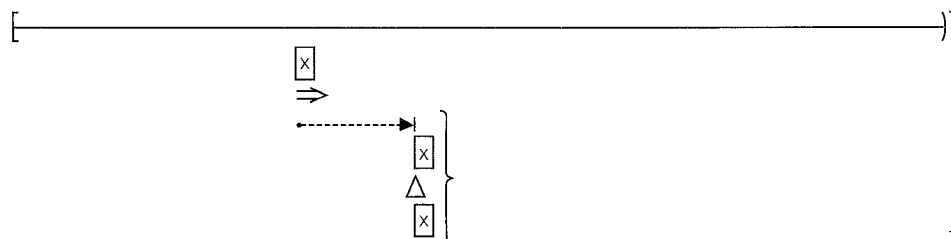
Briefly, to build the formula $UpFrom\$1$, a designer might begin by selecting the New button,⁴ which produces the template shown in Figure 6. The \square and \Rightarrow buttons can then be used to indent (automatically) the box below the context line and expand it into an implication. This produces the following template for an invariant implication.



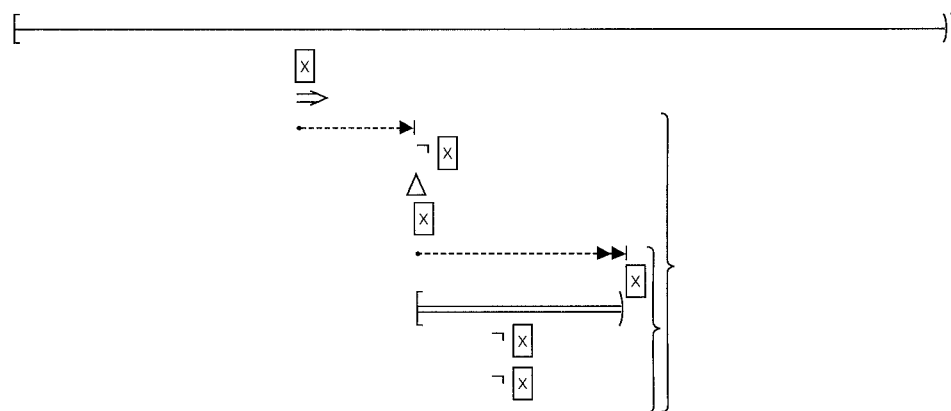
⁴The New button in the current implementation of GILED does not generate automatically the right parenthesis. This will be rectified in the next version of GILED.

GILED selects a box to expand next by default; however, the designer may override the default selection at any time using the mouse.

Selecting the second box and the Δ button converts the consequent into a point formula. For this requirement, the designer uses the mouse to position a single search arrow. GILED then produces a point symbol and a box to represent the search target, as shown below.



The designer can continue in this fashion to produce a template with the required structure.



The interval in this template is created by expanding a box into an appropriate interval template, using the mouse to position the interval and the search arrow. The $[\text{---}]$ button produces weak search arrows and weak intervals by default, so that the designer then clicks the mouse on the appropriate search arrow and interval to obtain their strong counterparts. To convert the above template into $\text{UpFrom}\$1$, the designer selects the pending boxes in turn, clicks on the Text button and types the state predicates.

In addition to the editing operations illustrated above, GILED provides capabilities for cutting and pasting formulas, resizing intervals and search arrows, repositioning invariants and eventualities, and so on. If a formula does not fit in the space allocated, GILED indicates an error and highlights the oversized formula. The designer can correct the error by resizing contexts and searches and repositioning formulas. Automatically, the editor resizes all affected subformulas to scale.

GILED interfaces with the GIL proof checker and model generator, allowing the designer to work entirely with graphical formulas. Functions that access these tool components are provided in the Misc pull-down menu under

the labels Check Proof, which determines if the formula in the canvas follows from premises designated by the designer, Prove, which checks the formula in the canvas for validity, and Construct Models, which determines if the formula in the canvas is satisfiable. Check Proof keeps track of the structure of each proof and checks for circular reasoning. Once verified, a requirement need only be reverified if the designer modifies premises used (either directly or indirectly) in the proof or if the designer wishes to modify the proof structure. In situations where either proof fails or a formula to be proved is not valid, a counterexample can be displayed in a separate window. Alternatively, Construct Models permits a model (infinite state sequence) that satisfies the formula in the canvas to be displayed.

For example, to verify that Safe\$1 follows from the specifications for the elevator system, the designer would create the requirement, or, if Safe\$1 was created and saved in an earlier GILED session, load it from a file. The designer would then invoke Check Proof to begin construction of a proof, or, if Safe\$1 was verified previously, to determine if the proof is up to date and learn the premises used in the proof. If a current proof exists, the designer can opt to see the proof (i.e., the implication constructed automatically by GILED to validate the inference). If a new proof is to be attempted, GILED prompts the designer for the premises to use in the proof. In this case, the designer would designate Init, SafeOpen\$1, and SafeLeave\$1 as premises. Then, GILED would construct a proof similar to that shown in Figure 3 and check that it is valid. The designer can also prove Safe\$1 directly, without invoking Check Proof, by building an implication representing the inference and invoking Prove to determine if the implication is valid.

When an attempt to verify a requirement fails, the designer can request to see a counterexample. Consider, for example, the proof of Arrive\$ n shown in Figure 5. If the designer overlooks the premise Safe\$ n and attempts to prove that Arrive\$ n follows from the other four premises, GILED generates the counterexample shown in Figure 7. The model consists of an infinite sequence of states with the state predicates having the values shown in the rectangles and the shaded state repeated infinitely. (The absence of a state predicate indicates that there is no restriction on the predicate's value in that state.) GILED displays timing diagrams beneath the state sequence to aid visualization. The designer can also invoke Construct Models to generate directly a model that satisfies the formula in the canvas.⁵

6. IMPLEMENTATION OF THE TOOL SET

Figure 8 shows the organization of the GIL tools. Rounded rectangles depict tool components (functions), and square rectangles depict data structures manipulated by the tools. A designer interacts with the tools through the mouse-driven interface provided by GILED. As described above, GILED helps the designer create new graphical formulas and retrieve and modify existing ones. It stores formulas in Unix files as abstract syntax trees with sufficient

⁵The GIL tools can be accessed by anonymous ftp from directory /pub/gil at ftp.cs.ucsb.edu.

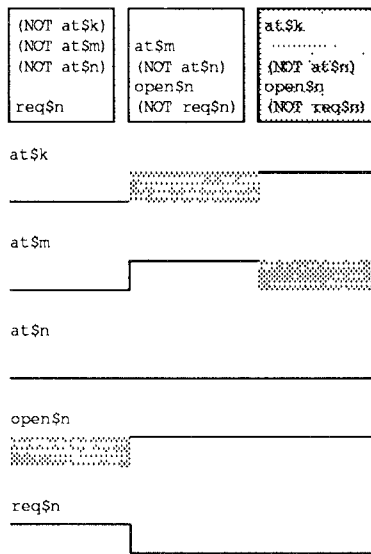


Fig. 7. Countermodel generated if $\text{Safe}n$ is omitted in the proof of $\text{Arrive}n$ shown in Figure 5.

representational information to recreate the layout specified by the designer when creating them. GILED also provides the interface to the proof checker and model generator, both of which make use of an intermediate representation of a formula as a semantic tableau. The procedure that constructs the tableau requires leaner abstract syntax trees in which productions reflect the semantics, and not merely representational variations in formula. Both the proof checker and model generator communicate results back to GILED, which displays them to the designer. The tools run under the X-window system and are written in Common Lisp using the Garnet graphics toolkit [Myers et al. 1990]. The implementations of the GIL tools are discussed in Dillon et al. [1994], Kutty [1993], and Kutty et al. [1993].

7. RELATED WORK

Graphical representations of computer systems have been common in software engineering practice, but have lacked a rigorous formal basis and, thus, have tended to be illustrative and documentary rather than an integral part of the software development process. Some notable exceptions include the statechart visual formalism of Harel [1987], a pictorial version of Milner's CCS, called IDCCS [Giocalone and Smolka 1988], and the \mathcal{V} -automata of Manna and Pnueli [1987]. Environments supporting the specification and verification of concurrent systems have been built around both Statecharts [Harel et al. 1990] and IDCCS. These languages are oriented toward the depiction of states and state transitions, whereas GIL focuses on showing the evolution of properties in time.

Timing Diagrams [Schlör and Damm 1993] is a graphical notation for expressing precedence and causality relationships between events in a computation. Like GIL, Timing Diagrams can be created using a graphical editor

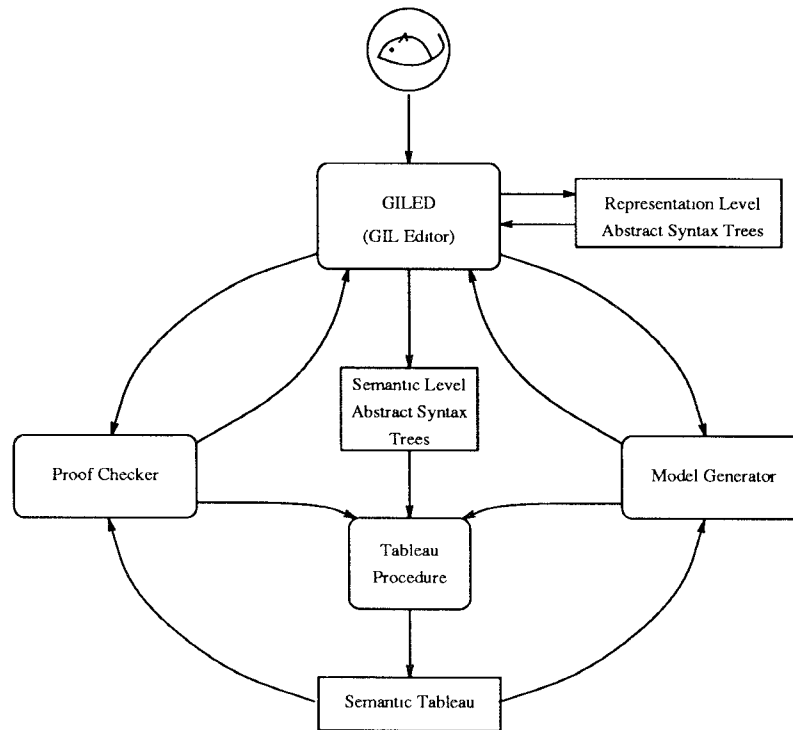


Fig. 8. The GIL tools.

and checked for validity. The semantics of Timing Diagrams is defined by translation to a subset of temporal logic that can be decided very efficiently.

Allen's logic for expressing temporal relationships between intervals of time is the foundation for the TIMELOGIC temporal reasoning system [Koomen 1989]. The logic is textual, but graphical representations are used to show relationships among intervals more clearly.

Moszkowski's Interval Temporal Logic [Halpern et al. 1983] provides an interval-like "chop" operator \mathcal{C} . Informally, $f\mathcal{C}g$ is true of a context if there exists a point that partitions the context into a prefix (subcontext) satisfying f and a suffix (subcontext) satisfying g . While the intuitive semantics of chop is appealing, the decision problem for formulas with chop is nonelementary in the depth of nested alternations of chop and negation. In contrast, intervals in GIL have a more operational semantics, but do not increase the complexity of the decision problem as severely. GIL can express a stronger version of chop, which suffices for expressing the properties of interest for the systems we have considered.

GIL is closest to the IntervalLogic (IL) of Schwartz et al. [1983], from which it is largely inspired. However, there are several presentational and semantic differences between the two logics, which we discuss briefly below.

Both IL and GIL provide explicit construction of intervals using search operations. However, they differ in the way that they construct intervals by the composition of searches. In IL, every search restricts a context, and intervals are obtained by nesting searches, yielding increasingly narrower contexts. In GIL, the start and end of an interval are located independently by means of a sequential composition of searches. Searches in IL are to intervals, rather than to states at which formulas hold, as in GIL. There is no loss or gain in expressiveness in either approach, but we feel that the state-based semantics of GIL are easier to define and understand. Moreover, searching to intervals requires the introduction of event intervals, representing positive transitions of formulas, and of “begin” and “end” operators, which are used to indicate how intervals located by searching further restrict a context. IL permits searches into the past as well as into the future. Allowing unrestricted searches into the past makes the decision procedure for GIL nonelementary [Ramakrishna 1993]. This is a major difference from IL, where the presence of both future and past searches does not appear to affect the complexity of the decision procedure.

Plaisted [1983] demonstrated a decision procedure for IL, obtained through translation to an ω -regular expression-like language with a nonelementary decision problem. PSPACE-completeness of IL was later established by Aaby and Narayana [1988], where they give a translation of IL to an elementary, but nonlogical, fragment of a nonelementary logic. The reduction is tedious and unnatural, and it points out the need for a simpler semantics that retains the advantages of being able to reason within intervals. The proof checker for GIL is based on a direct automata-theoretic decision procedure for the textual interval logic described in the Appendix. The complexity of the automaton construction is 2^m where m is $O(n^k)$ for a formula of size n and depth k of interval nesting [Ramakrishna et al. 1992].

An experiment with a graphical representation of an IL specification for the alternating-bit protocol [Melliari-Smith 1988] demonstrated that a visual representation results in more intuitive and natural specifications. The leaner semantics of our logic makes it more amenable than IL, to a clean graphical representation.

8. CONCLUSION

This article describes a visual logic for specifying concurrent software systems that aids formal reasoning about temporal properties of system. Experiments with the logic have produced graphical specifications for the sliding-window protocol [Kutty et al. 1994], a readers/writers database system [Dillon et al. 1992], a protocol to commit transactions on a shared database [Kutty et al. 1993] and a fair mutual-exclusion algorithm [Dillon et al. 1994], in addition to the elevator system. A prototype tool set supporting the analysis of GIL specifications has been developed.

Current research is addressing issues relating to the display of GIL formulas and the specification of temporal properties. In particular, we are experimenting with vertical spacing and with scaling the size of operator symbols to

improve the visual appearance of complex formulas and make their structure more visually evident. We are also investigating issues relating to the alignment of formulas to reflect known constraints on the partial ordering of points. A recent extension to GILED allows the designer to specify constraints on the ordering of points within a specification. Heuristics for recognizing search patterns that occur commonly in specifications and that impose an ordering on points in the specifications are being investigated. In such cases, GILED could aid the designer by aligning points accordingly. Methods for using a counterexample to realign the points of graphical formulas that constitute an invalid proof as also being explored. This would assist the designer in revising and correcting the proof and specifications.

GIL is very general and certainly admits formulas that lack the immediate visualization of the sample specifications presented in Section 3. For example, the semantics of searching to a formula with nested intervals is subtle and difficult to visualize, even when the search is represented graphically. However, we have not found a need for such searches in the specifications of the concurrent systems that we have considered. Our experiments indicate that most temporal properties of interest for concurrent systems can be specified in a natural and visually appealing manner using the derived operators introduced in this article. On-going research is attempting to identify syntactic restrictions that permit inferences to be checked more efficiently and still allow natural specifications of concurrent systems.

A real-time extension of GIL [Ramakrishna et al. 1993a; 1993b] provides primitives for bounding the duration of intervals. Recently we have modified the GIL proof checker to validate deductions in the extended logic and are currently experimenting with its use. We are also investigating the integration of the GIL decision procedure with an automated reasoning system that provides decision procedures for other useful theories, such as linear inequalities and Presburger arithmetic, and that provides better support for the management of proofs.

The GIL tool set is a prototype. It was developed to demonstrate proof-of-concept and to facilitate experiments with the logic and its graphical representation. Both the logic and the display of formulas have evolved based on our experience with the tools. We expect this process of experimentation and revision to continue as we refine the current tool set into a working environment for specification, validation, and design of concurrent software systems. A robust, user-friendly environment will permit empirical studies needed to determine whether software designers find a visual logic, such as GIL, easier to use than a textual logic.

APPENDIX A

A.1 Semantics

A model-theoretic semantics for our interval logic is presented below. For convenience, the semantics makes use of a textual version of the logic, called

Future Interval Logic (FIL).⁶ The semantics of GIL is then obtained by translation from GIL to FIL.

A.1.1 *Syntax of FIL.* The language of FIL, like that of GIL, is defined relative to a set \mathcal{P} of state predicates. The definition below makes use of the following generic symbols.

State Predicates: p, p_1, p_2, \dots
 Primitive search patterns: q, q_1, q_2, \dots
 General search patterns: Q, Q_1, Q_2, \dots
 Intervals: I, I_1, I_2, \dots
 Formulas: F, F_1, F_2, \dots, G .

The textual syntax of FIL is defined as follows.

$$\begin{aligned} F &::= p \parallel F_1 \vee F_2 \parallel \neg F_1 \parallel (F_1) \parallel IF_1 \\ I &::= [Q_1|Q_2) \\ Q &::= q \parallel q, Q_1 \\ q &::= \rightarrow F \parallel \rightarrow \end{aligned}$$

For convenience, we extend FIL with several abbreviations. As usual, $F_1 \wedge F_2 = \neg(\neg F_1 \vee \neg F_2)$. The temporal Henceforth and Eventually operations are also defined: $\Box F = [\rightarrow \neg F | \rightarrow)false$ and $\Diamond F = \neg[\rightarrow F | \rightarrow)false$. The shorthand $-$ denotes the trivial search $\rightarrow true$. The brackets $[[$ and $]]$ signify a strong interval, which is defined

$$[[Q_1|Q_2))F = [Q_1|Q_2)F \wedge ([Q_1| \rightarrow)false \vee [Q_2| \rightarrow)false \vee \neg[Q_1|Q_2)false).$$

Finally, we define a strong search to a target F , denoted $\rightarrow * F$, as follows:

$$\begin{aligned} \{Q_1, \rightarrow * F, Q_2|Q_3\}G &= \{Q_1, \rightarrow F, Q_2|Q_3\}G \wedge [Q_1| \rightarrow)\Diamond F \\ \{Q_1|Q_2, \rightarrow * F, Q_3\}G &= \{Q_1|Q_2, \rightarrow F, Q_3\}G \wedge ([Q_1| \rightarrow)false \vee [Q_2| \rightarrow)\Diamond F \end{aligned}$$

where the brackets $\{$ and $\}$ may be replaced by either pair of interval bracket delimiters: $[$ and $)$ or $[[$ and $))$. The rules that define strong searches and strong intervals can be shown to be semantically confluent, so that strong searches and strong intervals can be expanded in any order.

A.1.2 *Semantics of FIL.* A *model* (s, i) for an FIL formula consists of a *context* s and an *index* i . A context is either an infinite sequence of states $s = s(0), s(1), \dots$ or the *null* context \perp . A *state* $s(i)$ in a nonnull context s assigns valuations to the state predicates in \mathcal{P} . We identify a state $s(i) \subseteq \mathcal{P}$ with the collection of state predicates true in that state. The index i in a model is a (finite) nonnegative integer. We denote by $F|_{s,i}$ the value of a formula F in the model (s, i) .

For the definitions below, we extend the set of nonnegative integers with an infinite element ω , satisfying $i < \omega$ for all finite i , and define addition

⁶“Future” because searches used in the interval constructions are always into the future.

and subtraction on nonnegative integers (including ω) in the usual manner. For a nonnull state sequence s and indices l and r satisfying $l \leq r$ and $r < \omega$, we define further a sequence $s_{\langle l, r \rangle}$ as follows.

$$s_{\langle l, r \rangle}(k) = \begin{cases} s(k+l) & \text{for } 0 \leq k < r-l+1 \\ s(r) & \text{for } r-l+1 \leq k < \omega \end{cases}$$

Thus, $s_{\langle l, r \rangle}$ denotes the context obtained from s by extracting the subsequence from state $s(l)$ through state $s(r)$ and stuttering the final state.

The following rules provide an inductive definition for the truth value of an FIL formula. They make use of the function Construct for constructing the subcontext specified by an interval. Assuming $s \neq \perp$, we define

- $F|_{\perp, i} = \text{true}$.
- $p|_{s, i} = \text{true}$ if and only if $p \in s(i)$.
- $(F_1 \vee F_2)|_{s, i} = \text{true}$ if and only if $F_1|_{s, i} = \text{true}$ or $F_2|_{s, i} = \text{true}$.
- $(\neg F_1)|_{s, i} = \text{true}$ if and only if $F_1|_{s, i} = \text{false}$.
- $(F_1)|_{s, i} = \text{true}$ if and only if $F_1|_{s, i} = \text{true}$.
- $IF_1|_{s, i} = \text{true}$ if and only if $F_1|_{s', 0} = \text{true}$, where $s' = \text{Construct}(s, i, I)$.

The definition of Construct makes use of the function Locate for locating the state specified by a search pattern. Locate is defined as follows.

- $\text{Locate}(\rightarrow, s, i) = \omega$
- $\text{Locate}(\rightarrow F, s, i) = \begin{cases} \min(K) & \text{if } K \neq 0 \\ \epsilon & \text{otherwise} \end{cases}$
where $K = \{k | \omega > k \geq i \text{ and } F|_{s, k} = \text{true}\}$ and the special error value ϵ signifies a failed search
- $\text{Locate}((q, Q), s, i) = \begin{cases} \text{Locate}(Q, s, \text{Locate}(q, s, i)) & \text{if } \text{Locate}(q, s, i) \neq \epsilon \\ \epsilon & \text{otherwise.} \end{cases}$

Given a context, an index, and an interval, Construct produces the subcontext represented by the interval:

$$\text{Construct}(s, i, [Q_1|Q_2]) = \begin{cases} \perp & \text{if } r < l, l = \epsilon \text{ or } r = \epsilon \\ s_{\langle l, r \rangle} & \text{otherwise} \end{cases}$$

where $l = \text{Locate}(Q_1, s, i)$, $r = \text{Locate}(Q_2, s, i) - 1$, and $\epsilon - 1 = \epsilon$.

A.1.3 Formal Syntax for GIL and Translation Rules. The syntax of GIL is specified using a generalization of the attributed multiset grammar model described in Golin and Reiss [1989]. A multiset grammar differs from a context-free grammar in that productions do not impose any order on the symbols in their right-hand sides. A multiset grammar defines a set of multisets (unordered collections of terminal symbols, possibly containing repeated elements) rather than a set of strings. An attributed multiset grammar augments a multiset grammar with

- a set of attributes, which play an integral role in parsing an input
- semantic functions, which define the values of the attributes, and
- constraints, which indicate when a production can be applied.

In the attributed multiset grammar model of Golin and Reiss [1989], parsing attributes are restricted to synthesized attributes. For defining the syntax of GIL, however, we require a more general grammar model, which permits both synthesized and inherited attributes to be used for parsing. The grammar given below can be viewed as belonging to the index set grammar model described in Gillet and Kimura [1986].

In the sequel, therefore, an attributed multiset grammar consists of

- a finite set of terminal symbols
- a finite set of nonterminal symbols
- a finite set of attributes
- a mapping that associates sets of attributes with the terminal and nonterminal symbols and
- a finite set of productions.

Attributes are classified as synthesized or inherited, and attributes associated with terminal symbols are restricted to synthesized attributes. A production specifies a rewrite rule, which can be used to expand the nonterminal on the LHS into the multiset of symbols on the RHS, and associates a finite set of semantic functions and a finite set of constraints with the rule. Each semantic function defines an attribute of one of the rule's nonterminal symbols as a function of the values of other attributes of symbols in the rule. A production provides semantic functions for each synthesized attribute of the nonterminal on a rule's LHS and for each inherited attribute of the nonterminals on the rule's RHS. As is customary, the rewrite rules and semantic functions in a grammar must not admit derivations in which attribute values are defined circularly. The constraints associated with a rule are defined over inherited attributes of the nonterminal on the rule's LHS and synthesized attributes of symbols on the rule's RHS. The constraints specify conditions that must be satisfied in order to apply the rule.

The terminal symbols in our grammar for GIL correspond to the GIL operators, search arrows, interval symbols, and brackets, which are described, in Section 2, and to state formulas, which we denote by the special terminal symbol *state-form*.⁷ The nonterminals are defined as follows.

- F , representing GIL formulas
- Qq , representing a pair of search patterns
- Q , representing a single search pattern
- Qv , representing the continuation of a search pattern⁸
- L , representing a line segment denoting an interval
- Ai , representing a search arrow that begins a search pattern and
- Av , representing a search arrow that is embedded in a search pattern.

⁷State formulas are parsed using a context-free grammar. We omit the details, which are standard.

⁸For simplicity, we do not permit the horizontal shorthand for composing searches used in Section 2.

We regard each instance of a terminal or nonterminal symbol as enclosed by a bounding box. Synthesized attributes associated with the symbols give the position and dimensions of this box:

- t gives the y -coordinate of the top of the box.
- b gives the y -coordinate of the bottom of the box.
- l gives the x -coordinate of the left side.
- r gives the x -coordinate of the right side.

An additional synthesized attribute is associated with interval symbols that denote eventualities:

- p gives the x -coordinate of the center of the diamond.

The remaining attributes are inherited. Two inherited attributes are associated with the nonterminal Qq representing a pair of search patterns:

- lm gives the x -coordinate of the point that the first search pattern locates.
- rm gives the x -coordinate of the point that the second search pattern locates.

Finally, the grammar associates an inherited attribute with each of the non-terminals Q , Qv , and L :

- m gives the x -coordinate of the point that a search pattern Q or Qv locates or gives the x -coordinate of the formula that modifies an interval L .

Figure 9 illustrates the relationships between the attributes for nonterminal symbols. It also illustrates the conventions we use when drawing nonterminals. Lines show order relations between attribute values (vertical lines for y -coordinates and horizontal lines for x -coordinates), with strict ordering denoted by solid lines and nonstrict ordering denoted by broken lines. Thus, for example, the attributes associated with Qq are subject to the following constraints: $b < t$, $l \leq lm$, $lm < rm$, and $rm \leq r$.

Productions are shown in Tables I–VI. The last column of the tables also defines the translation of GIL to FIL. Table I gives sample productions and a translation scheme for the root nonterminal F .⁹ Each row in Table I represents a production for F and the corresponding translation rule. The second column shows the RHS of the rewrite rules, which are named, in the first column, for reference purposes below. Rewrite rules are shown graphically using a two-dimensional format so that relationships among attribute values, which are formally expressed by the semantic functions and constraints, are visually apparent. We use a broken line for the right side of a box when the right side need not coincide with the right side of the box enclosing the LHS nonterminal. The third and fourth columns give the constraints and semantic functions, respectively. Unsubscripted attribute names refer to attributes of the LHS nonterminal, and subscripted attribute names refer to attributes of

⁹For brevity, we have omitted productions for some of the propositional operators. Productions and translations for the missing operators are similar to those given for implication.

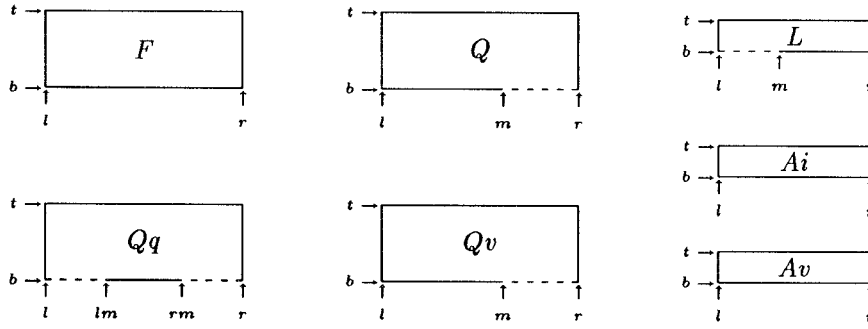


Fig. 9. Nonterminal symbols and their attributes.

Table I. Translation Rules for a GIL Formula F

Rule	F	Constraints	Sem. Functions	$t(F)$
F1	\neg_1 F_2	$l_1 = l_2, t_2 \leq b_1$	$t \leftarrow t_1; b \leftarrow b_2,$ $l \leftarrow l_1; r \leftarrow r_2$	$\neg t(F_2)$
F2	F_1 F_2	$l_1 = l_2, t_2 \leq b_1$	$t \leftarrow t_1; b \leftarrow b_2,$ $l \leftarrow l_1;$ $r \leftarrow \max\{r_1, r_2\}$	$t(F_1) \wedge t(F_2)$
F3	F_1 \Rightarrow_2 F_3	$l_1 = l_2 = l_3,$ $t_2 \leq b_1, t_3 \leq b_2$	$t \leftarrow t_1; b \leftarrow b_3;$ $l \leftarrow l_1;$ $r \leftarrow \max\{r_1, r_2\}$	$t(F_1) \Rightarrow t(F_3)$
F4	F_2 $\}_1$	$t_2 \leq t_1, b_1 \leq b_2,$ $r_2 \leq l_1$	$t \leftarrow t_1; b \leftarrow b_1,$ $l \leftarrow l_2; r \leftarrow r_1$	$t(F_2)$
F5	Qq_2 L_3 F_4 $\}_1$	$t_2 \leq t_1, t_3 \leq b_2,$ $t_4 \leq b_3, b_1 \leq b_4,$ $l_2 \leq l_3, l_3 \leq l_4,$ $r_3 \leq r_2, r_4 \leq r_3,$ $r_2 \leq l_1$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_2; r \leftarrow r_1;$ $lm_2 \leftarrow l_3;$ $m_3 \leftarrow l_4,$ $rm_2 \leftarrow r_3$	$t_{ll}(L_3)t(Qq_2)t_{rr}(L_3)$ $t_{md}(L_3)t(F_4)$
F6	L_2 F_3 $\}_1$	$t_2 \leq t_1, t_3 \leq b_2,$ $b_1 \leq b_3, l_2 \leq l_3,$ $r_2 \leq l_1, r_3 \leq r_2$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_2; r \leftarrow r_1,$ $m_2 \leftarrow l_3$	$t_{md}(L_2)t(F_3)$
F7	Q_2 Δ_3 F_4 $\}_1$	$t_2 \leq t_1, t_3 \leq b_2,$ $t_4 \leq b_3, b_1 \leq b_4,$ $l_2 \leq l_3, r_3 \leq r_2,$ $l_4 = (l_3 + r_3)/2,$ $\max\{r_1, r_2\} \leq l_1,$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_2; r \leftarrow r_1;$ $m \leftarrow l_4$	$t(Q_2) \mapsto t(F_4)$
F8	$state-form_1$		$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	$state-form_1$

Table II. Translation Rules for a Search Pair Qq

Rule	Qq	Constraints	Sem. Functions	$t(Qq)$
Qq1		$l_1 = l_2, t_2 \leq b_1,$ $l_1 < lm \leq r_1,$ $rm \leq r_2$	$t \leftarrow t_1, b \leftarrow b_2, l \leftarrow l_1;$ $r \leftarrow \max\{r_i\}_i;$ $m_1 \leftarrow lm; m_2 \leftarrow rm$	$t(Q_1) \text{ " " } t(Q_2)$
Qq2		$l_1 < l_2, t_2 \leq b_1,$ $l_2 = lm \leq r_1,$ $rm \leq r_2$	$t \leftarrow t_1; b \leftarrow b_2; l \leftarrow l_1,$ $r \leftarrow \max\{r_i\}_i,$ $m_1 \leftarrow lm, m_2 \leftarrow rm$	$t(Q_1) \text{ " " }$ $t(Q_1) \text{ " , " } t(Qv_2)$
Qq3		$l_1 = lm,$ $rm \leq r_1$	$t \leftarrow t_1; b \leftarrow b_1; l \leftarrow l_1,$ $r \leftarrow r_1, m_1 \leftarrow rm$	$\text{ " - " } t(Q_1)$

Table III. Translation Rules for a Search Pattern Q (vertical layout)

Rule	Q	Constraints	Sem. Functions	$t(Q)$
Q1		$r_1 < m,$ $r_1 = l_2 = l_3,$ $t_2 \leq b_1, t_3 \leq b_2,$ $m \leq r_3$	$t \leftarrow t_1, b \leftarrow b_3, l \leftarrow l_1;$ $r \leftarrow \max\{r_i\}_i, m_3 \leftarrow m$	$t(Ai_1)t(F_2) \text{ " , " } t(Qv_3)$
Q2		$r_1 = m, l_2 = r_1,$ $t_2 \leq b_1$	$t \leftarrow t_1; b \leftarrow b_2, l \leftarrow l_1;$ $r \leftarrow r_2$	$t(Ai_1)t(F_2)$
Q3		$r_1 = m$	$t \leftarrow t_1; b \leftarrow b_1; l \leftarrow l_1;$ $r \leftarrow r_1$	 " - "

Table IV. Translation Rules for the Continuation Qv of a Search Pattern (vertical layout)

Rule	Qv	Constraints	Sem Functions	$t(Qv)$
Qv1		$r_1 < m,$ $r_1 = l_2 = l_3,$ $t_2 \leq b_1, t_3 \leq b_2,$ $m \leq r_3$	$t \leftarrow t_1, b \leftarrow b_3, l \leftarrow l_1,$ $r \leftarrow \max\{r_i\}_i; m_3 \leftarrow m$	$t(Av_1)t(F_2) \text{ " , " } t(Qv_3)$
Qv2		$r_1 = m, l_2 = r_1,$ $t_2 \leq b_1$	$t \leftarrow t_1, b \leftarrow b_2, l \leftarrow l_1,$ $r \leftarrow r_2$	$t(Av_1)t(F_2)$
Qv3		$r_1 = m$	$t \leftarrow t_1; b \leftarrow b_1, l \leftarrow l_1;$ $r \leftarrow r_1$	 " - "

the symbols on the RHS. The fifth column of Table I defines the translation of F into a string $t(F)$ representing a FIL formula. For simplicity, we give a fully parenthesized translation.

Tables II, III, and IV provide productions and translation rules for Qq , Q , and Qv , respectively. As shown, Q and Qv differ only in the type of arrow

Table V. Translation Rules for a Line Segment L

Rule	L	Constraints	Sem Functions	$t_{li}(I)$	$t_{ri}(L)$	$t_{md}(L)$
L1	$\boxed{\text{---}} \rightarrow$		$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	"[")"	if $l_1 < m$: "□" if $l_1 = m$: " "
L2	$\boxed{\text{====}} \rightarrow$		$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	"[")"	if $l_1 < m$: "□" if $l_1 = m$: " "
L3	$\boxed{\text{---}} \rightarrow \diamond$	$m = p_1$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	"[")"	"◇"
L4	$\boxed{\text{====}} \rightarrow \diamond$	$m = p_1$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	"[")"	"◇"

 Table VI. Translation Rules for the Arrows A_i and A_v

Rule	A_i	Rule	A_v	Sem. Functions	$t(A_i)/t(A_v)$
Ai1	$\bullet \text{---} \rightarrow \vdash_1$	Av1	$\text{---} \rightarrow \vdash_1$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	"→"
Ai2	$\bullet \text{---} \rightarrow \rightarrow \vdash_1$	Av2	$\text{---} \rightarrow \rightarrow \vdash_1$	$t \leftarrow t_1; b \leftarrow b_1;$ $l \leftarrow l_1; r \leftarrow r_1$	"→*"

with which an instance begins. Table V defines three translation functions for line segments: t_{li} generates a left interval-delimiter; t_{ri} generates a right interval-delimiter; and t_{md} generates a Henceforth or Eventually symbol or, when the formula that modifies the interval is positioned at the first state of the interval, an empty string. Table VI gives the productions and translations of both types of search arrows. The rewrite rules for the two kinds of arrows do not require any constraints, and their semantic functions and translations are identical.

We show some steps in the translation of ArriveDown below. Annotations over a derivation arrow indicate the translation rules used at each step.

$$\begin{array}{l}
 t(F) \xrightarrow{F6} t_{md}(L_1) ("t(F_1)") \\
 \xrightarrow{L1} \square ("t(F_1)") \\
 \xrightarrow{F3, F8} \square ((at \$3) \Rightarrow ("t(F_3)")) \\
 \xrightarrow{F5} \square ((at \$3) \Rightarrow ("t_{li}(L_2)t(Qq_1)t_{ri}(L_2)t_{md}(L_2) ("t(F_4)")))) \\
 \xrightarrow{L4} \square ((at \$3) \Rightarrow (["t(Qq_1)"] \diamond ("t(F_4)")))) \\
 \xrightarrow{Qq2} \square ((at \$3) \Rightarrow (["t(Q_1)"] "t(Q_1)", t(Qv_1)) \diamond ("t(F_4)")))) \\
 \xrightarrow{Q2, Ai1, F8} \square ((at \$3) \Rightarrow (["\rightarrow at \$2"] \rightarrow at \$2, "t(Qv_1)") \diamond ("t(F_4)")))) \\
 \xrightarrow{Qv2, Av1, F8} \square ((at \$3) \Rightarrow (["\rightarrow at \$2"] \rightarrow at \$2, \rightarrow arriveup) \diamond (\neg at \$2)))
 \end{array}$$

The GIL tool set does not include a parser. Formulas are constructed using a syntax-directed editor. Currently, we are investigating the effects of characteristics of attributed multiset grammars on the efficiency of parsing.

ACKNOWLEDGMENT

The authors would like to thank Ron Dolin for implementing several last minute modifications to GILED to improve the appearance of formulas.

REFERENCES

- AABY, A. A., AND NARAYANA, K. T. 1988. Propositional temporal interval logic in PSPACE complete. In *Proceedings of the 9th International Conference on Automated Deduction*. Lecture Notes in Computer Science, vol. 193. Springer-Verlag, Berlin, 218–237.
- BARRINGER, H., KUIPER, R., AND PNUELI, A. 1984. Now you may compose temporal logic specifications. In *Proceedings of the 16th ACM Symposium on Theory of Computing*. ACM, New York, 51–63.
- DILLON, L. K., KUTTY, G., MELLIAR-SMITH, P. M., MOSER, L. E., AND RAMAKRISHNA, Y. S. 1994. Visual specifications for temporal reasoning. *J. Vis. Lang. Comput.* 5, 1, 61–81.
- DILLON, L. K., KUTTY, G., MOSER, L. E., MELLIAR-SMITH, P. M., AND RAMAKRISHNA, Y. S. 1993. A graphical interval logic for specifying concurrent systems. Tech. Rep. TRCS 93-16, Computer Science Dept., Univ. of California, Santa Barbara, Calif.
- DILLON, L. K., KUTTY, G., MOSER, L. E., MELLIAR-SMITH, P. M. AND RAMAKRISHNA, Y. S. 1992. Graphical specifications for concurrent software systems. In *Proceedings of the 14th IEEE International Conference on Software Engineering*. IEEE, New York, 213–224.
- GABBAY, D. M. 1987. The declarative past and imperative future. In *Proceedings of the Conference on Temporal Logic in Specification*. Lecture Notes in Computer Science, vol. 398, Springer-Verlag, 409–448.
- GIACALONE, A., AND SMOLKA, S. A. 1988. Integrated environments for formally well-founded design and simulation of concurrent systems. *IEEE Trans. Softw. Eng.* 14, 6 (June), 787–801.
- GILLET, W. D., AND KIMURA, T. D. 1986. Parsing two-dimensional languages. In *Proceedings of the IEEE 10th International Conference of Computer Software and Applications*. IEEE, New York, 472–477.
- GOLIN, E. AND REISS, S. P. 1989. The specification of visual language syntax. In *Proceedings of the IEEE Workshop on Visual Languages*, IEEE, New York, 105–110.
- HALPERN, J. Y., AND SHOHAM, Y. 1991. A propositional modal logic of time intervals. *J. ACM* 38, 4 (Oct.), 935–962.
- HALPERN, J. Y., MANNA, Z., AND MOSZKOWSKI, B. 1983. A hardware semantics based on temporal intervals. In *Proceedings of the 10th International Conference on Automata, Languages and Programming*. Eur. Assoc. for Theoretical Computer Science, 278–291.
- HAREL, D. 1987. Statecharts: A visual formalism for complex systems. *Sci. Comput. Program.* 8, 3 (June), 231–274.
- HAREL, D., LACHOVER, H., NAAMAD, A., PNUELI, A., POLITI, M., SHERMAN, R., SHTULL-TRAURING, A., AND TRAKHTENBROT, M. 1990. STATEMATE: A working environment for the development of complex reactive systems. *IEEE Trans. Softw. Eng.* 16, 4 (Apr.), 403–414.
- KOOMEN, J. A. G. M. 1987. The TIMELOGIC temporal reasoning system. Tech. Rep., Dept. of Computer Science, Univ. of Rochester, N.Y. (Revised March 1989).
- KUTTY, G. 1993. A tool for the interactive generation of Graphical Interval Logic formulas. Tech. Rep., 9307, Dept. of Electrical and Computer Engineering, Univ. of California, Santa Barbara, Calif.
- KUTTY, G., MOSER, L. E., MELLIAR-SMITH, P. M., DILLON, L. K., AND RAMAKRISHNA, Y. S. 1994. First-order future interval logic. In *Proceedings of the 1st International Conference on Temporal Logic*. Lecture Notes in Artificial Intelligence, vol. 827. Springer-Verlag, Berlin, 195–209.
- KUTTY, G., RAMAKRISHNA, Y. S., MOSER, L. E., DILLON, L. K., AND MELLIAR-SMITH, P. M. 1993. A graphical interval logic toolset for verifying concurrent systems. In *Proceedings of the 4th ACM Transactions on Software Engineering and Methodology*, Vol 3, No 2, April 1994

- Conference on Computer Aided Verification*. Lecture Notes in Computer Science, vol. 697. Springer-Verlag, Berlin, 138–153.
- LAMPORT, L. 1990. A temporal logic of actions. Tech. Rep. 57, DEC Systems Research Center, Palo Alto, Calif.
- LAMPORT, L. 1983. What good is temporal logic? In *Proceedings of the IFIP Congress*. IFIP, Washington, D.C., 657–668.
- LANDIN, P. J. 1966. The next 700 programming languages. *Commun. ACM*, 9, 3 (Mar.) 157–166.
- MANNA, Z., AND PNUELI, A. 1987. Specification and verification of concurrent programs by \forall -automata. In *Proceedings of the Conference on Temporal Logic in Specification*. Lecture Notes in Computer Science, vol. 348, Springer-Verlag, Berlin, 124–187.
- MANNA, Z., AND PNUELI, A. 1981. Verification of concurrent programs: The temporal framework. In *The Correctness Problem in Computer Science*, R. S. Boyer and J. S. Moore, Eds. Academic Press, New York, 215–273.
- MELLIAR-SMITH, P. M. 1988. A graphical representation of interval logic. In *Proceedings of the International Conference on Concurrency*. Lecture Notes in Computer Science, vol. 335, Springer-Verlag, Berlin, 106–120.
- MYERS, B. A., GUISE, D. A., DANNENBERG, R. B., VANDER ZANDEN, B., KOSBIE, D. S., PERVIN, E., MICKISH, A., AND MARCHAL, P. 1990. Garnet: Comprehensive support for graphical highly interactive user interfaces. *IEEE Comput.* 18, 11 (Nov.), 71–85.
- PLAISTED, D. 1983. A low level language for obtaining decision procedures for classes of temporal logics. In *Proceedings of the CMU Workshop on Logics of Programs*. Lecture Notes in Computer Science, Vol. 164. Springer-Verlag, Berlin, 403–420.
- PRATT, V. 1986. Modeling concurrency with partial orders. *Int. J. Parallel Program.* 15, 1, 33–71.
- RAMAKRISHNA, Y. S. 1993. Interval Logics for Temporal Specification and Verification. Ph.D. thesis, Dept. of Computer and Electrical Engineering, Univ. of California, Santa Barbara, Calif.
- RAMAKRISHNA, Y. S., DILLON, L. K., MOSER, L. E., MELLIAR-SMITH, P. M., AND KUTTY, G. 1993a. A real-time interval logic and its decision procedure. In *Proceedings of the 13th Conference on the Foundations of Software Technology and Theoretical Computer Science*. Lecture Notes in Computer Science, vol. 761. Springer-Verlag, Berlin, 173–192.
- RAMAKRISHNA, Y. S., MELLIAR-SMITH, P. M., MOSER, L. E., DILLON, L. K., AND KUTTY, G. 1993b. Really visual temporal reasoning. In *Proceedings of the 14th IEEE Real-Time Systems Symposium*, IEEE, New York, 262–273.
- RAMAKRISHNA, Y. S., DILLON, J. K., MOSER, L. E., MELLIAR-SMITH, P. M. AND KUTTY, G. 1992. An automata-theoretic decision procedure for future interval logic. In *Proceedings of the 12th Conference on the Foundations of Software Technology and Theoretical Computer Science*. Lecture Notes in Computer Science, vol. 652. Springer-Verlag, Berlin, 51–67.
- SCHLÖR, R., AND DAMM, W. 1993. Specification of system-level hardware designs using timing diagrams. In *Proceedings of the European Conference on Design Automation and European Event in ASIC Design*. IEEE Computer Society Press, Los Alamitos, Calif., 518–524.
- 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 the Principles of Distributed Computing*. ACM, New York, 173–186.
- WOLPER, P. 1987. On the relation of programs and computations to models of temporal logic. In *Proceedings of the Conference on Temporal Logic in Specification*. Lecture Notes in Computer Science, vol. 398. Springer-Verlag, Berlin, 75–123.

Received June 1992, revised October 1992; accepted March 1994.