Towards Diagrammability and Efficiency in Event Sequence Languages*

Kathi Fisler

Department of Computer Science
WPI (Worcester, MA, USA)
kfisler@cs.wpi.edu

Abstract. Many industrial verification teams are developing suitable event sequence languages for hardware verification. Such languages must be expressive, designer friendly, and hardware specific, as well as efficient to verify. While the formal verification community has formal models for assessing the efficiency of an event sequence language, none of these models also accounts for designer friendliness. We propose an intermediate language for event sequences that addresses both concerns. The language achieves usability through a correlation to timing diagrams; its efficiency arises from its mapping into deterministic weak automata. We present the language, relate it to existing event sequence languages, and prove its relationship to deterministic weak automata. These results indicate that timing diagrams can become more expressive while remaining more efficient for symbolic model checking than LTL.

1 Introduction

The increasing adoption of formal verification has led to a flurry of research into property specification languages for hardware verification. Large-scale efforts include Accellera's standardization of Sugar [1], Synopsys' OVA [13], and Intel's FTL [4]. Generally speaking, these are *event sequence languages*: they allow designers to express sequences of events to monitor and check during verification. The proliferation of work from industry on event sequence languages emphasizes that they must be designer friendly, expressive, and specific to the hardware domain in addition to efficient to verify. Although practical experience and theoretical results give insights into how to achieve these goals individually, few formal models attempt to address usability and efficiency simultaneously.

In the space of event sequence languages, timing diagrams provide an appealing combination of usability and efficiency. Designers have established their utility by regularly employing them as an informal design tool. Mappings from formalized timing diagrams to deterministic weak automata [8] provide effectively linear symbolic verification algorithms [5]. That timing diagrams are not more widely used as event sequence languages suggests that they lack the expressiveness needed in industrial verification [3]. Their combination of utility and

* This research is supported through NSF grant CCR-0132659.

D. Geist and E. Tronci (Eds.): CHARME 2003, LNCS 2860, pp. 185–199, 2003.

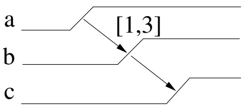

$$\text{LTL: } \neg a \wedge \mathsf{X}(a \wedge ((\neg b \wedge \mathsf{X}(b \wedge \mathsf{F}(\neg c \wedge \mathsf{X}c)))\vee \\ \mathsf{X}(\neg b \wedge \mathsf{X}(b \wedge \mathsf{F}(\neg c \wedge \mathsf{X}c)))\vee \\ \mathsf{XX}(\neg b \wedge \mathsf{X}(b \wedge \mathsf{F}(\neg c \wedge \mathsf{X}c))))$$

Sugar: $\neg a$ & next!(a & next_e![1,3]($\neg b$ & next! (b & eventually! ($\neg c$ & next! c))))

Fig. 1. Expressing an event sequence in three languages.

efficiency, however, raises an interesting question: how expressive can we make an event sequence language while retaining both diagrammability and efficiency?

This paper explores this question by proposing a (textual) intermediate language for capturing event sequence languages. To target diagrammability, we design the core of the language around timing diagrams. To target expressiveness, we extend the core language to capture constructs from other event sequence languages. To target efficiency, we syntactically characterize which expressions in this language map to deterministic weak automata. The results of this work are twofold: first, our language provides a framework in which to assess both usability and efficiency of other event sequence languages; second, our characterization proves that timing diagrams can be extended with several new features—such as partial orders between events, interleaved environmental assumptions, escaping conditions, and event clocks—without losing their mapping to deterministic weak automata. Our long-term goal is to develop formal models that simultaneously characterize both usability and efficiency in event sequence languages. This paper focuses on the efficiency of verifying our proposed language; future papers will treat formal models of diagrammability as a measure of usability.

2 Preliminaries

2.1 Event Sequences and Timing Diagrams

Event sequences, as their name implies, capture sequences of events on signals in a design; they express properties for verification or simulation. Regular expressions and linear temporal logic have similar goals, but also some subtle differences. Event sequences often monitor *transitions* on signals in the design, rather than just boolean values of propositions. In addition, event sequences generally capture timing constraints between events. While both regular expressions and linear temporal logic can capture these features, the resulting expressions can be rather cumbersome, especially in contrast to event sequences and timing diagrams. Figure 1 shows a simple example of the same event sequence expressed as a timing diagram, in linear temporal logic (LTL), and in Sugar.

Although timing diagrams present event sequences somewhat intuitively, they are not as expressive as some other event sequence languages. For example, textual event sequence languages easily express disjunctions, while diagrams in general capture disjunctive information poorly. The mapping from timing diagrams to weak automata, which does not hold for full LTL, demonstrates benefits to

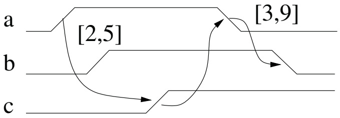

$$\begin{aligned} C &= \langle \{a\uparrow, b\uparrow, c\uparrow, a\downarrow\}; b\downarrow\rangle \\ T &= \{\langle a\uparrow, c\uparrow, 2, 5, \mathsf{true}\rangle, \\ &\qquad \langle c\uparrow, a\downarrow, 1, \infty, \mathsf{true}\rangle, \\ &\qquad \langle a\downarrow, b\downarrow, 3, 9, \mathsf{true}\rangle\} \end{aligned}$$

Fig. 2. A timing diagram with partial orders and its mapping into an event sequence.

this limited expressive power. The question, then, is how far we can push timing diagrams while retaining this mapping. The timing diagram shown in Figure 2, for example, expresses some disjunction as the order of events is left unspecified (a partial order rather than a total one). This extension adds expressive power without sacrificing diagrammability or weakness. We are interested in similar extensions based on constructs from modern event sequence languages.

2.2 Weak Automata

A Büchi automaton $\langle Q, \Sigma, q_0, R, L, \mathcal{F}\rangle$ is *weak* if it has only one fair set and each of its strongly connected components has either all states fair or no states fair [10]. Weak automata are attractive in verification because symbolic cycle detection is effectively linear for weak automata, as opposed to quadratic for full LTL [5]. Deterministic weak automata are particularly interesting for their properties under complementation. Automata-based verification approaches complement automata that capture properties. In the general case, complementing a Büchi automaton can blowup the number of states exponentially. Complementing a deterministic weak automaton, however, requires only complementing the fair set; the structure of an automaton and its complement are otherwise identical. This represents a substantial savings in construction time, and more importantly, in the size of automata used to represent complemented properties.

3 An Intermediate Language for Event Sequences

This section presents a regular-expression-like notation for event sequences. We motivate the development of the language using the example timing diagram shown in Figure 2. We explain the semantics of the diagram informally; the formal details appear elsewhere [7].

To capture the diagram, the language must express transitions on signals and constraints (timing and ordering) between these transitions. Let propositional literals (p, $\neg q$) denote boolean values and propositional variables annotated with arrows ($p\downarrow$, $p\uparrow$) denote falling and rising transitions, respectively. Let semicolons denote concatenation (temporal sequencing) of events. Using these notations and reading off the timing diagram from left to right suggests the expression $\langle a\uparrow; b\uparrow; c\uparrow; a\downarrow; b\downarrow\rangle$. If we interpret semicolons as implying order between events (a common interpretation of concatenation), this expression is inconsistent with the semantics of the timing diagram. The rising transitions on a and b may occur in any order since no constraint orders them (the falling events on a and b, in contrast, must occur in order). The event sequence language must therefore support partial, rather than only total, orders between events.

Timing diagrams consist of totally-ordered regions within which individual events are partially ordered. For sake of generality, our event sequence language supports hierarchical combinations of ordered, unordered and iterated groups of events. In the formal syntax and semantics that follows, we refer to these groups of events as *clusters*. We capture partial orders within unordered clusters using a separate annotation for transition (timing) constraints between events; a timing constraint specifies the events covered, lower and upper bounds on the time between the events, and the clock against which the bounds are measured (true specifies the system clock). This approach treats constraints between events uniformly, whether they occur in ordered or unordered clusters. Figure 2 shows the resulting event sequence for our example timing diagram.

3.1 Syntax

The timing diagram example suggests the following syntax for event sequences:

Definition 1 Clusters are defined hierarchically as follows:

- An *event* is a conjunction of values of and transitions on variables that contains at least one transition. Propositional literals (p, $\neg q$) denote boolean values; propositions with arrows ($p\downarrow$, $p\uparrow$) denote transitions.
- A *cluster* is either:
  - a single event, or
  - an *unordered cluster* $\{C_1, \ldots, C_k\}$ where each C_i is a cluster, or
  - an *ordered cluster* $\langle C_1; \ldots; C_k\rangle$ where each C_i is a cluster, or
  - a *repeating cluster* C^M where C is a non-repeating cluster and M is a positive number, $*$, or $+$ (called a *repetition marker*; markers $*$ and $+$ are called *unbounded*).

An event sequence consists of a (top level) cluster and three kinds of modifiers. Temporal constraints, already motivated, may be relative to a designer-specified *event clock*, as captured by a boolean expression (this is a common feature in event sequence languages). To indicate that certain variables hold value during regions (between events) in a diagram, *holding patterns* constrain variable values within clusters. To allow portions of diagrams to serve as assumptions rather than requirements, *escape conditions* capture circumstances under which the sequence should be immediately rejected or accepted.

Definition 2 An *event sequence* is a tuple $\langle C, H, T, S\rangle$ where C is a cluster, H (the holding patterns) is a partial function from C to propositional formulas, T is a set of temporal constraints and S is a set of escape conditions.

$C = \langle a\uparrow^+; \{b\uparrow, c\uparrow\}; d\uparrow\rangle$
$H = \{b\uparrow, c\uparrow\} \to a$
$T = \{\langle c\uparrow, d\uparrow, 2, 5, \mathsf{true}\rangle\}$
$S = \{\text{accept-if-don't-complete}(a\uparrow^+)\}$

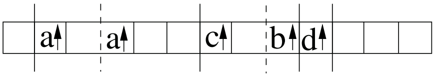

Fig. 3. A sample event sequence and an example of its semantics.

- A *temporal constraint* is a tuple $\langle e_1, e_2, l, u, clk\rangle$ where e_1 and e_2 are (uniquely identified[1]) events in C, l is a positive integer, u is either an integer at least as large as l or the symbol ∞, and clk is a boolean expression (the clock for the constraint; true indicates the system clock). Events e_1 and e_2 may lie in different clusters, but then they must lie in the same repeated clusters.
- An *escape condition* has one of three types, where X is a boolean expression over events (the events need not be in C) and C' is a cluster within C:
  - "accept if don't complete C'"
  - "reject if see X in C'"
  - "accept if see X in C'"

Figure 3 illustrates an event sequence of some number of rising transitions on a, followed by rising transitions on b and c (in either order), followed by a rising transition on d. The transition on d must occur between 2 and 5 ticks (inclusive) after the transition on c (the timing constraint), signal a must remain true until the transition on d occurs (the holding pattern), and the rest of the sequence is only checked if the transition on a occurs (the escape condition).

The language contains some redundancy for sake of clarity: ordered clusters, for example, can be viewed as unordered clusters plus timing constraints. To simplify the semantics and proofs, we assume that all sequences are in *reduced form*, in which all clusters C^+ are replaced with $\langle C; C^*\rangle$, all C^M for a concrete number M are replaced with an ordered cluster of M copies of C, and all ordered clusters $\langle C_1; \ldots; C_k\rangle$ are replaced with unordered clusters and timing constraints that require an event from each C_i to occur before an event from each C_{i+1}.

3.2 Semantics

The semantics of event sequences is defined in terms of languages over infinite words, where each character in a word is an assignment of boolean values to variables. An infinite word models an event sequence if there exists a mapping from the clusters in the sequence to ranges of indices into the word (herein called *windows*) such that the windows assigned to each cluster preserve the cluster's constraints; these mappings are called *index assignments.*

As an example, consider the event sequence and word shown in Figure 3. The word is divided into windows per cluster (demarcated by solid lines), and subwindows as necessary for nested clusters (demarcated by dashed lines). We first formalize the mappings from clusters to windows.

[1] A numbering scheme could distinguish syntactically similar events.

Definition 3 Given a word W, a *window of* W is a subword of W; a pair of indices into W, denoted $[i, j]$ where $i \leq j$, defines a window. Furthermore,

- An individual index i defines a trivial window $[i, i]$.
- Window $[i_1, i_2]$ *contains* window $[i_3, i_4]$ iff $i_1 \leq i_3$ and $i_4 \leq i_2$.
- Window $[i_1, i_2]$ is *earlier* than window $[i_3, i_4]$ iff $i_1 < i_3$ or $i_1 = i_3$ and $i_2 < i_4$.
- Given a window $w = [start, end]$, a sequence $[s_1, e_1], \ldots, [s_k, e_k]$ forms a *non-overlapping covering sequence of windows for* w if $s_1 = start$, $e_k = end$, and for all $1 \leq j < k$, $e_j < s_{j+1}$.

Definition 4 A (partial) *index assignment* for event sequence V and word W is a (partial) function from the clusters in (including nested within) V to non-empty sets of windows of W.

A window must meet certain requirements in order to capture the constraints of a cluster. The following definitions formalize those requirements.

Definition 5 Let $E = v_1 \wedge \ldots \wedge v_k$ where each v_i is a proposition, its negation, or a rising or falling transition on a proposition. Let W be a word and i an index into W. Let $W_i(q)$ denote the value of proposition q at index i into W. Index i *satisfies* E if for every v_i, $W_i(p) = 0$ if $v_i = \neg p$, $W_i(p) = 1$ if $v_i = p$, $W_i(p) = 0$ and $W_{i+1}(p) = 1$ if $v_i = p\uparrow$, and $W_i(p) = 1$ and $W_{i+1}(p) = 0$ if $v_i = p\downarrow$.

Definition 6 Given an unordered cluster $C = \{C_1, \ldots, C_k\}$, a *schedule of* C is a sequence $CO_1, \ldots, CO_j$ of non-empty subsets of C such that

- $CO_1, \ldots, CO_j$ partition C,
- In every CO_i that contains multiple elements of C, all elements of CO_i are single events (rather than other complex clusters), and
- For each timing constraint $\langle e_1, e_2, l, u, clk\rangle$ such that $e_1 \in CO_i$ and $e_2 \in CO_j$, $i < j$.

Definition 7 Let V be an event sequence, W a word, and I a partial index assignment for V and W. I is *structurally valid* iff for every cluster C in V:

- If C is an event, then for every $[i, i] \in I(c)$, i satisfies C (Defn 5).
- If C is a repeating cluster C'^*, then for every wp in $I(C'^*)$ there exists a natural number m and some sequence $wp_1, \ldots, wp_m$ of non-overlapping covering windows for wp such that each $wp_i \in I(C')$.
- If C is an unordered cluster $\{C_1, \ldots, C_k\}$, then for every window $w \in I(C)$ there exists a schedule $CO_1, \ldots, CO_j$ for C and a sequence $w_1, \ldots, w_j$ of non-overlapping covering windows for w such that for all $i \leq h \leq j$ and all $e \in CP_h$, $w_h \in I(e)$.

Definition 8 Let $V = \langle C, T, H, S \rangle$ be an event sequence, let W be a word, and let I be an index assignment for V and W. I is *constraint valid* for V and W iff

1. I satisfies the holding patterns, in that for all clusters C', every $x \in H(C')$ and every window $[w_1, w_2] \in I(C')$, every index $w_1 \leq i \leq w_2$ satisfies x, and
2. I satisfies the timing constraints, in that for every $\langle e_1, e_2, l, u, clk \rangle \in T$ and every $t_1 \in I(e_1)$ and $t_2 \in I(t_2)$ such that t_1 and t_2 fall in a common window for the smallest cluster containing both e_1 and e_2, the number of indices satisfying *clk* between t_1 and t_2 (inclusive) is within the range $[l, u]$.

Constraint validity handles timing constraints and holding patterns, but not escape conditions. The next two definitions handle escape conditions. Definition 12 relates words and event sequences based on the existence of index assignments that may or may not invoke escape conditions. Given index assignment I, let $\overline{I}$ be the inverse of I (mapping windows to sets of clusters).

Definition 9 Let V be an event sequence, W a word, and I a structurally valid index assignment for V and W. Let E be an escape condition of type "accept/reject if see X in C". Index i into W *invokes E under I* if $i \in I(C)$, i satisfies X, and I is defined for all clusters in the images of $\overline{I}$ for windows occurring before i. We also say that I *invokes an escape condition of* V.

Definition 10 Let V be an event sequence, W a word, and I a structurally valid index assignment for V and W. *I loops under escape condition E* if E is of the form "accept if don't complete C" and I is defined for all clusters in the images of $\overline{I}$ for windows occurring before i, but not for a window containing i.

For the semantics to yield a deterministic procedure for checking whether a word satisfies an event sequence, index assignments must assign the fewest and earliest possible windows to clusters (in particular, this renders both * and scheduling deterministic). We formally define this notion of minimality as follows:

Definition 11 Let V be an event sequence and let W be a word. Let I and I' be non-equivalent index assignments for V and W. Let Rg denote the range of a function. $I \prec I'$ iff

1. the earliest window in one but not both of $Rg(I)$ and $Rg(I')$ is in $Rg(I)$, or
2. $Rg(I) = Rg(I')$ but for w, the earliest window such that $\overline{I}(w) \neq \overline{I'}(w)$, $\overline{I'}(w) \subset \overline{I}(w)$.

Given a set Σ of index assignments, $I \in \Sigma$ is *minimal in Σ* iff for all $I' \in \Sigma$, $I \prec I'$. ($\prec$ does not order all pairs, but is sufficient for our theorems [9].)

We now define when a word models an event sequence:

Definition 12 Let V be an event sequence and let W be a word. $W \models V$ if there exists a minimal and structurally valid index assignment I for V and W such that I is a total function and constraint valid, or I loops under some escape condition in V, or I invokes some escape condition in V.

The semantics captures one occurrence of an event sequence, rather than the multiple occurrences needed to treat an event sequence as an invariant. The one-occurrence semantics offers two benefits: it provides a foundation for defining different multiple occurrence semantics [7], and it enables the mapping to weak automata. This restriction is not as limiting as it might seem: in prior work [8], we showed that relabeling fair sets and adding a few transitions constructs the automaton for a negated invariant event sequence (the machine most commonly needed for verification) from the machine that accepts one occurrence.

4 Relationship to Existing Event Sequence Languages

To motivate the intersection between our simultaneous goals of diagrammability and efficiency, this section shows how several features of existing event sequence languages do or do not map into the proposed intermediate language.

4.1 Timing Diagrams

Section 3 illustrated the connection between timing diagrams and our proposed event sequence language. The language presented here extends our previous results on the relationship between timing diagrams and weak automata [8] in two ways. The previous result held for timing diagrams with a total order on their transitions and a prefix of the diagram as an environmental assumption (as in, "if the rising transition on a occurs, then match the whole diagram"). As a corollary to the results in this paper, timing diagrams with partial event orders and multiple non-contiguous assumptions on the environment also map to deterministic weak automata. We view environment assumptions as events that are only constrained if they occur [6]; unlike other events, their failure to occur does not violate the diagram's requirements. For the diagram in Figure 2, we could treat the two transitions on a as environment assumptions by rewriting the event chain using nested clusters (as $\langle\{a\uparrow, b\uparrow, c\uparrow, a\downarrow\}; b\downarrow\rangle$ and adding "accept-if-don't-complete" escape conditions on the two clusters for a.

The proposed language is more expressive than our current timing diagram formalization. Consider the cluster $\langle a\uparrow^{*}; b\uparrow\rangle$. The timing diagram semantics requires all depicted transitions to occur unless an escape condition matches, so this expression (without escape conditions) is currently not expressible as a timing diagram (since $a\uparrow$ might not occur). Similar examples involving repetitions also exist. Enriching the timing diagram notation could resolve some of these issues; this remains an issue for future work.

4.2 LTL, Sugar, and FTL

Sugar and FTL are similar in that each extends conventional LTL. Since there exist LTL formulas that cannot be captured by weak automata, certain FTL and Sugar formulas will not map into our intermediate language. Weakness primarily characterizes the location of fair sets in automata. In LTL, fairness

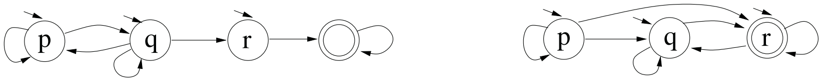

Fig. 4. Automata for two LTL formulas.

constraints arise from combinations of eventualities and cycles (the operators U and G). Figure 4 shows automata that capture two formulas: $(p\ \mathsf{U}\ q)\ \mathsf{U}\ r$ and $p\ \mathsf{U}(\mathsf{G}(q\ \mathsf{U}\ r))$. The first example yields a weak automaton and corresponds to cluster $\langle (p^*; q)^*; r\rangle$. The second corresponds to cluster $\langle p^*; \langle q^*; r^+\rangle^*\rangle$ with escape condition "accept if don't complete r^+"; this expression violates our syntactic restrictions for weakness presented in Theorem 3 (Section 5.3).

One key difference between these two formulas is that the second contains a repetition within its last cluster, while the first does not. This same difference characterizes the automata for the regular expressions $(aa)^*$ and $(aa)^*b$, the first of which cannot be captured by a deterministic weak automaton while the second one can. An automaton can recognize a nonrepeating final pattern without creating a fair set. This motivates our characterization of weakness: the final cluster cannot end with an unbounded repetition marker.

Certain other features of Sugar and FTL do not adversely impact weakness. FTL's *change_on* and *reject_on* constructs indicate when a sequence should be immediately accepted or rejected; escape conditions capture such scenarios in the proposed intermediate language. For example, augmenting $(p\ \mathsf{U}\ q)\ \mathsf{U}\ r$ with escape condition "accept if see *reset* in q" would introduce a new state labeled *reset* with an incoming edge from the state for q; this automaton is also weak.

4.3 OVA

Of the recent event sequence languages discussed in this paper, OVA most closely matches the proposed language. Unlike Sugar and FTL, OVA does not explicitly support LTL or CTL operators. The OVA *istrue* construct maps into holding patterns, and their non-overlapping event clocks map into ours. Unlike the proposed language, however, OVA can express disjunction among sequences and negation of sequences. Our language does not support negation because negated sequences generally cannot be realized diagrammatically. Our language does, however, still support constructing deterministic weak automata for the negations of event sequences, as outlined at the end of Section 3.

5 Relationship to Deterministic Weak Automata

This section characterizes which sequences in our language map to deterministic weak automata; almost all do, with the exception of those with particular interactions between escape conditions and repeated clusters. We construct an automaton corresponding to the semantics, prove the construction sound, then characterize when the resulting machine is both weak and deterministic.

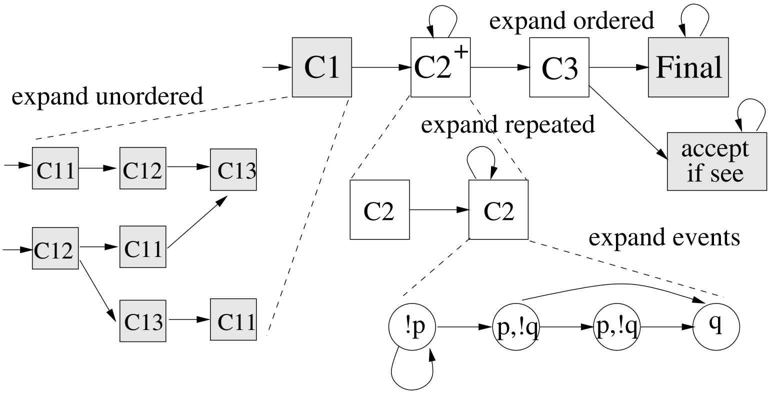

Fig. 5. Overview of the automaton construction algorithm.

Given an event sequence V, we construct a Büchi automaton that accepts all words with a prefix that models V. Figure 5 illustrates the intuition behind the expansion. The construction recursively expands states corresponding to clusters until all states correspond to individual events. Holding patterns, escape conditions, and the ordering aspects of timing constraints are incorporated as this expansion proceeds. The durational aspects of timing constraints are handled in a final phase once all states correspond to individual events.

Each intermediate machine during the computation abstracts the final machine, in that if there is no edge from one abstract state to another, then there is no edge from any state in the expansion of the first to the expansion of the second in the final machine. For sake of space, we present the detailed algorithm only up through creating states for each event; this is sufficient for our theorems.

The construction creates edges between abstract states based on which clusters can precede or follow other clusters; it also relies on notions of the first and last subclusters that could be encountered in a cluster. These concepts match intuition. For sake of space, we defer all but the definition of next clusters to the full paper [9]; examples of all four notions follow the definition. The theorem in Section 5.2 also refers to first and last *events*, which are obtained by iterating the first and last computations on clusters until they contain only events.

Definition 13 Let C be a cluster immediately contained in a cluster C^P (if C has no enclosing cluster, next(C) is empty). If $C^P = \langle C_1; \ldots ; C_k\rangle$ and $C = C_i$ for $i < k$, then next(C) is C_{i+1} if C_{i+1} is not an repeating-* cluster and $\{C_{i+1}\}\cup$ next(C_{i+1}) if C_{i+1} is an repeating-* cluster. If $C = C_k$, then next(C) is next(C^P). If C is an repeating-* cluster, next(C) also includes C. The case for unordered clusters unions similar results over all possible schedules, and repeated clusters C have next(C) as $\{C\}\cup$ next(C^P).

Examples: Given sequence $\langle C_1; C_2; C_3^*; C_4^+\rangle$, next($C_2$)=next($C_3$)=$\{C_3, C_4\}$ and prev(C_3)=$\{C_2, C_3\}$. Given sequence $\langle C_1; \{C_{21}, C_{22}^*\}; \{C_{31}, C_{32}\}^*; C_4\rangle$ with a timing constraint from C_{21} to C_{22}, next(C_{21})=$\{C_{22}, C_{31}, C_{32}, C_4\}$. For the first and last sets, first($\{C_{21}, C_{22}^*\}$)=$\{C_{21}\}$, last($\{C_{21}, C_{22}^*\}$)=$\{C_{22}^*\}$, and first($\{C_{31}, C_{32}^*\}$) = last($\{C_{31}, C_{32}^*\}$) = $\{C_{31}, C_{32}^*\}$.

Algorithm 1 To construct an automaton for event sequence $\langle C, T, H, S \rangle$:

```
1. Create a state *Final* with a self loop and mark it fair.
2. Create a state for $C$ and mark it initial, final, and unexpanded.
3. Repeatedly select an unexpanded state $N$ for some non-event cluster $C$ and
     – Add holding patterns and edges for the escape conditions for $C$ to $N$.
     – Expand $N$ according to the type of $C$ and remove $N$.
     – If $N$ was marked initial (resp. final), mark the new states for all first
       (resp. last) clusters of $C$ initial (resp. final). Copy all other propositional
       annotations (including fair) from $N$ to the new states from the expansion.
4. Add an edge from each state marked final to the state *Final*.
```

Expand Repeated Clusters. For a state for repeated cluster C^*, add an edge from the state for each previous cluster of C^* to that for each next cluster of C^*.

Expand Unordered Clusters. For a state N for unordered cluster $C = \{C_1, \ldots, C_k\}$:

- For every schedule $CO_1, \ldots, CO_h$ of C, create a chain of abstract states $CON_1, \ldots, CON_h$. For every non-self-loop edge coming into N, add an edge from the same source to CON_1. For every non-self-loop edge leaving N, add an edge from CON_k to the target of the original edge.[2]
- Eliminate unnecessary nondeterminism by merging states with the same incoming transitions and labels into single states (this shares common prefix states across the various permutations).
- If N had an edge to itself, add an edge from each sink state in the subgraph that expands N to each source state in the subgraph that expands N.

Handle Escape Conditions and Holding Patterns

- For each escape condition E of the form "reject if see X in C", create a new abstract state N_E for E, label N_E with X, add an edge from each abstract state corresponding to C to N_E and add a self-loop at N_E.
- For each escape condition E of the form "accept if see X in C", create a new abstract state N_E for E, label N_E with X, add an edge from each abstract state corresponding to C to N_E, add a self-loop at N_E, and mark N_E as fair (with a new fairness constraint).
- For each escape condition E of the form "accept if don't complete C", mark every abstract for C as fair (with a new fairness constraint).
- For each holding pattern h for cluster C and each abstract state N_C corresponding to or expanded from C, add h as a propositional label to N_C.

[2] To reduce the machine size, we could perform a bisimilarity minimization on the subgraph of all states that expanded N.

Following Algorithm 1, all states correspond to single events but the durations of timing constraints have not been enforced. We handle this using a similar algorithm to that in our prior work [8]. For sake of space, and since the expansion into events does not affect weakness or determinism by construction, we do not reproduce the details here. To handle the event clock *clk* in a timing constraint over events e_1 and e_2, the construction adds a unique label for *clk* to each state between e_1 and e_2, and creates an automaton that outputs this label whenever *clk* is true. A final step cross-products the core machine with the clock machines; this does not affect weakness.

The results on determinism and weakness that follow apply to those event sequences that end with a concrete event rather than a repetition (for reasons motivated in Section 4.2). We call such sequences *event chains*.

Definition 14 An event sequence $\langle C, T, H, S\rangle$ is an *event chain* if the iterative expansion of $\mathrm{last}(C)$ contains no repeated clusters.

5.1 Soundness

Theorem 1. *Let V be an event sequence and let M be the automaton obtained for V from Algorithm 1. Let W be an infinite word. M accepts W iff $W \models V$.*

Proof Sketch: Intuitively, the proof develops a correspondence between states in the abstract machines and the windows in the range of an index assignment for W and V. The theorem follows from an argument that the windows occurring in accepted (resp. rejected) words correspond to accepting (resp. rejecting) paths through the automaton.

5.2 Characterization of Determinism

Theorem 2. *Given an event chain, Algorithm 1 produces a deterministic automaton if all of the following conditions are satisfied:*

- *For every unordered cluster $\{C_1, \ldots, C_k\}$, the first events of each C_i are pairwise logically inconsistent with those of each $C_j \neq C_i$ unless a timing constraint orders C_i and C_j.*
- *For each repeated cluster C^*, the first events of C are pairwise logically inconsistent with the first events of each next cluster of C^* (other than C).*
- *For each "accept/reject when see X in C" escape condition, X is logically inconsistent with all holding patterns for C.*

Proof Sketch: The machine is deterministic if the choice among multiple next states is deterministic. The construction yields multiple next states in four cases: possible transitions to the *Final* state, when choosing between schedules for an unordered cluster, possible skips of repeated clusters, and when invoking escape conditions. The restriction to event chains guarantees that states with transitions to *Final* have no other outgoing transitions. By construction, transitions into the states that expand clusters occur when a first event is recognized for that cluster. If these events are logically inconsistent, then the corresponding transitions must be deterministic. This covers the remaining cases.

5.3 Characterization of Weakness

We call a cluster C *fair* if there exists an escape condition of the form "accept if don't complete C". A cluster is *all-fair* if it is either fair or all of its sub-clusters are all-fair. A cluster is *non-fair* if neither it nor any of its sub-clusters is fair.

Lemma 1. *If an event sequence contains no all-fair repeated clusters, then the automaton from Algorithm 1 requires only one fair set.*

Proof Sketch: If no cycle contains states from more than one fair set, then a single fair set suffices. Cycles can contain states from multiple fair sets under two conditions. First, two "accept don't complete" conditions could exist for clusters C_1 and C_2 where C_1 contains C_2. In this case, a cycle that satisfies C_2 satisfies C_1, so only one fairness constraint is required. Second, a repeated cluster could have all sub-clusters fair, thus creating a cycle that visits each sub-cluster then self-loops for the repeated cluster. The theorem statement rules out this case.

Theorem 3. *Given an event chain, Algorithm 1 produces a weak automaton iff every repeated cluster in the chain is non-fair.*

Proof Sketch: Non-trivial strongly-connected components (SCCs) arise from abstract states with self-loops, which in turn arise from expanding states for repeated clusters. With the exception of the *Final* state and the states for "accept/reject if see" escape conditions (which form their own SCCs), states are marked fair only if they correspond to or expand from clusters that have "accept if don't complete" conditions. If a repeated cluster is non-fair, then it has no fair SCCs embedded within self-loops (other, larger SCCs). If a repeated cluster is all-fair, it requires multiple fair sets and is not weak by definition. All other repeated clusters contain cycles with both fair and non-fair states.

Our mapping to deterministic weak automata is not complete; in other words, our language does not logically characterize deterministic weak automata. Consider the regular expression $ab^* + bc^*$: a deterministic weak automaton accepts it, but it is not expressible in our language due to the use of disjunction.

6 Related Work

We are unaware of logical characterizations of weak automata, much less ones that account for diagrammability or other forms of usability. The original work on the efficiency of verifying weak automata is due to Bloem, Ravi and Somenzi [5]. Other timing diagram formalizations have supported some of the language extensions discussed here [2,6,12], but none related the diagrammatic features of these languages to efficiency in verification.

Amla *et al.'s* work on modular timing diagrams has much in common with this work [3]. Their work makes timing diagrams more expressive by combining them through non-diagrammatic operators for conjunction, iteration, and deterministic choice. Expressions in their language encompass several timing diagrams, while our work pushes the limits of a single timing diagram. Accordingly, they target efficiency through a different model of automata. The core differences between our works appear to be philosophical; ours focuses on understanding the interplay between diagrammability and efficiency, while theirs focuses on building a practical verification framework around timing diagrams. The full paper provides a more detailed comparison [9].

7 Conclusions and Future Work

The relationships between timing diagrams and deterministic weak automata suggest that there exist formal models of event sequences that simultaneously address both usability and efficiency. A traditional theoretical approach to designing languages towards efficiency would be to find a syntactic (logical) characterization of weak automata. This approach, however, fails to account for the usability of that logical characterization. This is perhaps justifiable, as ”usability” is an inherently informal notion. If we refine our notion of usability to mean diagrammability, however, formal models become possible. Formal characterizations of diagrammability usually rely on topological or spatial arguments [11]; appropriate characterizations for discrete linear events remain an open problem.

The event sequence language proposed in this paper targets diagrammability by allowing only a restricted form of disjunction; in particular, disjunction governs the *ordering* of events, but not their *occurrence*. This is consistent with diagrams' tendency to imply that all depicted items actually exist (maps, for example, indicate that all depicted features are actually there). Such nuances in the different uses of logical operations appear fundamental to formal models of diagrammability. This limited nature of disjunction also targets efficiency by supporting our criteria for deterministic automata. Restricted forms of iteration enable the mapping to weak automata. Single timing diagrams support limited forms of iteration, and hence satisfy the criteria for weakness. Overall, the generality of our language substantially enriches the set of features our timing diagrams can support while retaining efficiency for verification.

Several avenues remain open for future work. Given that the proposed language is more expressive than our current timing diagrams, characterizing diagrammability is an important next problem in this project. We expect restrictions on cluster nesting similar to those in timing diagrams to be key to such a characterization. We also plan to explore formal relationships between other event sequence languages and ours; this would help identify subsets of other languages that could be visualized and verified efficiently through a mapping to weak automata. Finally, many general questions remain regarding the nature of diagrammatic representations and their relationship to computational concerns such as efficiency and decidability that are so important in verification. We hope that our work will contribute to better understanding of these issues.

Acknowledgements. The author thanks the anonymous reviewer whose thoughtful and extensive comments led to significant improvements in this paper.

References

1. Accellera Working Group. Property specification language reference manual (version 1.0). Available at `http://www.eda.org/vfv/docs/psl_lrm-1.0.pdf`, 2003.
2. N. Amla, E. A. Emerson, and K. S. Namjoshi. Efficient decompositional model checking for regular timing diagrams. In *IFIP Conference on Correct Hardware Design and Verification Methods*, 1999.
3. N. Amla, E. A. Emerson, K. S. Namjoshi, and R. J. Trefler. Visual specifications for modular reasoning about asynchronous systems. In *International Conference on Formal Techniques for Networked and Distributed Systems*, pages 226–242, 2002.
4. R. Armoni et al. The ForSpec temporal logic: A new temporal property-specification language. In *Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 296–211, 2002.
5. R. Bloem, K. Ravi, and F. Somenzi. Efficient decision procedures for model checking of linear time logic properties. In *International Conference on Computer-Aided Verification*, number 1633 in Lecture Notes in Computer Science, pages 222–235. Springer-Verlag, 1999.
6. K. Feyerabend and B. Josko. A visual formalism for real-time requirement specifications. In M. Bertran and T. Rus, editors, *Transformation-Based Reactive Systems Development, Proc. 4th International AMAST Workshop on Real-Time Systems and Concurrentand Distributed Software, ARTS'97*, volume 1231, pages 156–168. Springer-Verlag, 1997.
7. K. Fisler. Timing diagrams: Formalization and algorithmic verification. *Journal of Logic, Language, and Information*, 8:323–361, 1999.
8. K. Fisler. On tableau constructions for timing diagrams. In *NASA Langley Formal Methods Workshop*, 2000.
9. K. Fisler. An event sequence language and its relationship to weak automata. Technical Report WPI-CS-TR-03-24, WPI, Department of Computer Science, July 2003.
10. O. Kupferman and M. Y. Vardi. Freedom, weakness, and determinism: From linear-time to branching-time. In *IEEE Symposium on Logic in Computer Science*, 1998.
11. O. Lemon. Comparing the efficacy of visual languages. In D. Barker-Plummer, D. I. Beaver, J. van Benthem, and P. S. di Luzio, editors, *Words, Proofs, and Diagrams*, pages 47–70. CSLI Publications, 2002.
12. Y. Ramakrishna, L. Dillon, L. Moser, P. Melliar-Smith, and G. Kutty. A real-time interval logic and its decision procedure. In *Proc. Thirteenth Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 761 of *Lecture Notes in Computer Science*, pages 173–192. Springer-Verlag, December 1993.
13. Synopsys, Inc. Openvera assertions. Available online for download at `http://www.open-vera.com/technical/technical.html`, 2002.