

A Compositional Theory of Refinement for Branching Time

Panagiotis Manolios

Georgia Institute of Technology, CERCs Lab
801 Atlantic Drive
Atlanta, Georgia, 30332, USA
manolios@cc.gatech.edu
<http://www.cc.gatech.edu/~manolios>

Abstract. I develop a compositional theory of refinement for the branching time framework based on stuttering simulation and prove that if one system refines another, then a refinement map always exists. The existence of refinement maps in the linear time framework was studied in an influential paper by Abadi and Lamport. My interest in proving analogous results for the branching time framework arises from the observation that in the context of mechanical verification, branching time has some important advantages. By setting up the refinement problem in a way that differs from the Abadi and Lamport approach, I obtain a proof of the existence of refinement maps (in the branching time framework) that does not depend on any of the conditions found in the work of Abadi and Lamport *e.g.*, machine closure, finite invisible nondeterminism, internal continuity, the use of history and prophecy variables, etc. A direct consequence is that refinement maps always exist in the linear time framework, subject only to the use of prophecy-like variables.

1 Introduction

Computing systems are ubiquitous, controlling everything from cars and airplanes to financial markets and the distribution of information. Such systems tend to be very complicated and often contain costly errors. One approach to dealing with this complexity is to specify a sequence of related systems, starting with an abstract system, the specification, and ending with a concrete system, the implementation. One then proves that every pair of adjacent systems is related, via a suitable, compositional notion of correctness, thereby establishing that the specification is correctly implemented. For example, we can imagine verifying a netlist description of a pipelined microprocessor, the implementation, by relating it via a sequence of refinements to an instruction set level specification—the assembly programmer’s view of the processor.

Two important concepts that notions of correctness must account for are:

- **Stuttering.** Since the specification is defined at a more abstract level than the implementation, notions of correctness should allow for *stuttering*, where the implementation may require several steps before matching a single step of the specification [14].

- **Refinement.** The implementation may contain more state components and may use different data representations than the specification. *Refinement maps* are used to show how to view an implementation state as a specification state [1].

The classic paper on the topic by Abadi and Lamport [1], which has motivated the work appearing in this paper, contains an in-depth discussion of these topics. The main idea is to use refinement maps to prove that systems have related infinite computations, by reasoning *locally*, about states and their successors, instead of *globally*, about infinite paths. Abadi and Lamport prove a theorem about when such refinement maps exist in the linear time framework, where the semantics of systems and properties correspond to sets of infinite sequences.

My approach differs in that I work in the branching time framework, where the semantics of systems are given by sets of infinite trees. Even so, the results can be applied to the linear time framework, as I explain later.

The theorem proved by Abadi and Lamport holds only under certain conditions. Briefly, they allow one to add history and prophecy variables to the implementation, they require that the implementation is machine closed, and they require that the specification has finite invisible nondeterminism and is internally continuous. My theorems do not depend on these conditions, but there are important differences between the two approaches that are explored in depth later.

There are two main reasons why I chose to work in the branching-time framework. The first is that in the simple case where one is dealing with finite-state systems, it makes sense to use algorithms that can check if one finite-state system refines another. For example, in [17] we use algorithms for deciding stuttering bisimulation to complete a proof of correctness for the alternating bit protocol (this is an infinite-state problem that was reduced to a finite-state problem using a theorem prover). The branching time notions of simulation and bisimulation, due to Milner and Park [18, 22], can be decided in polynomial time [21, 7]. The most efficient algorithm for deciding bisimulation equivalence has time complexity $O(m \log n)$, where n is the number of states and m the number of transitions, and is due to Paige and Tarjan [21]. The best known algorithm for checking simulations has time complexity $O(m \cdot n)$ and is due to Henzinger, and Kopke [3, 7]. In contrast, the corresponding linear time notions, trace equivalence and trace containment, are both PSPACE-complete problems [27].

Second, refinement maps allow one to show that one system *simulates* another. This is inherently a branching time notion which has the advantage of being structural and local. However, in order to use refinement maps in a linear time setting other mechanisms are needed to, in essence, hide the branching structure of systems. Thus, we expect the branching time case to be simpler than the linear time case. Obvious questions arise. How much simpler? What conditions in the Abadi and Lamport theorem are there for this purpose? It turns out that by using only prophecy-like variables, which have the effect of destroying

the branching structure of systems, we can get a completeness theorem for the linear time.

Stuttering simulation is based on the notions of simulation and bisimulation, which have had a deep impact on how we think about specifications. I chose to use stuttering simulation as opposed to weak simulation because weak simulation allows infinite stuttering, something I wish to avoid. The literature on this topic is vast and contains many fine surveys [24, 15, 6]. In addition, there have been various extensions of the Abadi and Lamport result [1], including [5, 9, ?, 8].

A very readable discussion of the computational complexity of bisimulation is given by Moller and Smolka [19]. Pnueli gives a balanced comparison between the branching-time and linear time approaches in [24].

In related previous work, Namjoshi [20] gives a sound and complete proof rule for symmetric stuttering bisimulations which has heavily influenced my work; however, Namjoshi does not consider simulations and does not deal with refinement. Stuttering bisimulations and the related notion of WEBs (Well-founded Equivalence Bisimulations) were used to link theorem proving and model checking and to mechanically verify the alternating bit protocol in [17]. In [16], I proposed a notion of correctness for pipelined machines based on WEBs and I showed that the variant of the Burch and Dill notion of correctness [4] in [25, 26] can be satisfied by machines that deadlock. In addition, I used the ACL2 theorem prover [12, 11, 10] to automate much of the verification, *e.g.*, the verification of the simple pipelined machine in [26] is automatic. I also verified variants of the pipelined machine including machines with exceptions, interrupts (which lead to non-determinism), and netlist (gate-level) descriptions and showed that my notion of correctness applies to these extensions. Many of the variant machines were verified in stages, using the WEB compositional proof rule. Unfortunately, stuttering bisimulation and WEBs are often too strong a notion, just as trace equivalence is often too strong a notion in the linear time case. I expect stuttering simulation to be much more applicable, hence my interest in the topic.

The paper is organized as follows. In section 2, I describe my notational conventions and review background material. In section 3, I develop a theory of refinement based on stuttering simulation. In section 4, I discuss refinement in the linear time framework and compare my work with that of Abadi and Lamport; some readers may want to start by skimming this section first. I conclude in section 5.

2 Notation and Mathematical Preliminaries

\mathbb{N} and ω both denote the natural numbers, *i.e.*, $\{0, 1, \dots\}$. The ordered pair whose first component is i and whose second component is j is denoted $\langle i, j \rangle$. $[i..j]$ denotes the closed interval $\{k \in \mathbb{N} : i \leq k \leq j\}$; parentheses are used to denote open and half-open intervals, *e.g.*, $[i..j)$ denotes the set $\{k \in \mathbb{N} : i \leq k < j\}$. The disjoint union operator is denoted by \uplus . Cardinality of a set S is denoted by $|S|$. $\mathcal{P}(S)$ denotes the powerset of S . Function application is sometimes denoted by an infix dot “.” and is left associative. This allows us to use the curried

version of a function when it suits us, *e.g.*, I may write $f.x.y$ instead of $f(x, y)$. For any binary relation R : I abbreviate $\langle s, w \rangle \in R$ by sRw , I write $R(S)$ for the *image* of S under R (*i.e.*, $R(S) = \{y : \langle \exists x : x \in S : xRy \rangle\}$), and $R|_A$ denotes R *left-restricted* to the set A (*i.e.*, $R|_A = \{\langle a, b \rangle : (aRb) \wedge (a \in A)\}$). The *composition* of binary relations R and T is denoted $R;T$ or $T \circ R$, *i.e.*, $R;T = T \circ R = \{\langle r, t \rangle : \langle \exists x :: rRx \wedge xTt \rangle\}$. The *inverse* of binary relation R is denoted R^{-1} and is defined to be $\{\langle a, b \rangle : bRa\}$.

$\langle Qx : r : b \rangle$ denotes a quantified expression, where Q is the quantifier, x the bound variable, r the range of x (**true** if omitted), and b the body. I sometimes write $\langle Qx \in X : r : b \rangle$ as an abbreviation for $\langle Qx : x \in X \wedge r : b \rangle$, where r is **true** if omitted, as before. From highest to lowest binding power, we have: parentheses, function application, binary relations (*e.g.*, sBw), equality ($=$) and membership (\in), conjunction (\wedge) and disjunction (\vee), implication (\Rightarrow), and finally, binary equivalence (\equiv).

Spacing is used to reinforce binding: more space indicates lower binding.

A binary relation, $B \subseteq X \times X$, is *reflexive* if $\langle \forall x \in X :: xBx \rangle$. B is *symmetric* if $\langle \forall x, y \in X :: xBy \Rightarrow yBx \rangle$. B is *antisymmetric* if $\langle \forall x, y \in X :: xBy \wedge yBx \Rightarrow x = y \rangle$. B is *transitive* if $\langle \forall x, y, z \in X :: xBy \wedge yBz \Rightarrow xBz \rangle$. A binary relation is a *preorder* if it is reflexive and transitive. A preorder that is also symmetric is an *equivalence relation*.

A *finite sequence* is a function from $[0..n)$ for some natural number n . An *infinite sequence* is a function from \mathbb{N} . When I write $x \in \sigma$, for a sequence σ , I mean that x is in the range of σ . A *well-founded structure* is a pair $\langle W, \prec \rangle$ where W is a set and \prec is a binary relation on W such that there are no infinitely decreasing sequences on W , with respect to \prec . I use $<$ to compare natural numbers and \prec to compare ordinal numbers.

A *transition system* (TS) is a structure $\langle S, \rightarrow, L \rangle$, where S is a set of states, $\rightarrow \subseteq S \times S$ is the *transition relation*, L is the *labeling function*: its domain is S and it tells us what is observable at a state. I also require that \rightarrow is *left-total*: for every $s \in S$, there is some $u \in S$ such that $s \rightarrow u$. Notice that a transition system is a labeled graph where the nodes are states and are labeled by L .

A *path* σ is a sequence of states such that for adjacent states s and u , $s \rightarrow u$. A path, σ , is a *fullpath* if it is infinite. $fp.\sigma.s$ denotes that σ is a fullpath starting at state s and σ^i denotes the suffix fullpath $\langle \sigma.i, \sigma(i+1), \dots \rangle$. I use the symbol “;” for concatenation of paths where the left path is finite, *e.g.*, $a;ab = aab$.

Temporal logic was proposed as a formalism for specifying the correctness of computing systems in a landmark paper by Pnueli [23]. I assume that the reader is familiar with temporal logic.

3 Stuttering Simulation Refinement

Stuttering simulation depends on the notion of matching I now define. I start with an informal account. Given a relation B on a set S , we say that an infinite sequence σ (of elements from S) matches an infinite sequence δ (of elements from S) if the sequences can be partitioned into non-empty, finite segments such

that elements in related segments are related by B . For example, if the first segment of σ has three elements and the first segment of δ has seven elements, then each of the three elements is related by B to each of the seven elements. I use matching, where the infinite sequences are fullpaths of a transition system, to define stuttering simulation.

Definition 1. (*match*) Let i range over \mathbb{N} . Let INC be the set of strictly increasing sequences of natural numbers starting at 0; formally, $INC = \{\pi : \pi : \mathbb{N} \rightarrow \mathbb{N} \wedge \pi.0 = 0 \wedge \langle \forall i \in \mathbb{N} :: \pi.i < \pi(i+1) \rangle\}$. The i^{th} segment of an infinite sequence σ with respect to $\pi \in INC$, ${}^\pi\sigma^i$, is given by the sequence $\langle \sigma(\pi.i), \dots, \sigma(\pi(i+1)-1) \rangle$.

For $B \subseteq S \times S$, $\pi, \xi \in INC$, $i, j \in \mathbb{N}$, and infinite sequences σ and δ , I abbreviate $\langle \forall s, w : s \in {}^\pi\sigma^i \wedge w \in {}^\xi\delta^j : sBw \rangle$ by $({}^\pi\sigma^i)B({}^\xi\delta^j)$.

In addition: $corr(B, \sigma, \pi, \delta, \xi) \equiv \langle \forall i \in \mathbb{N} :: ({}^\pi\sigma^i)B({}^\xi\delta^i) \rangle$ and $match(B, \sigma, \delta) \equiv \langle \exists \pi, \xi \in INC :: corr(B, \sigma, \pi, \delta, \xi) \rangle$.

Lemma 1. *Given set S , $B \subseteq S \times S$, and infinite sequences σ and δ ,*
 $\langle \exists \pi, \xi \in INC :: corr(B, \sigma, \pi, \delta, \xi) \rangle$

$$\equiv \langle \exists \pi', \xi' \in INC :: corr(B, \sigma, \pi', \delta, \xi') \wedge \langle \forall i \in \mathbb{N} :: |\pi' \sigma^i| = 1 \vee |\xi' \delta^i| = 1 \rangle \rangle$$

Proof The \Leftarrow direction is clear. For the other case, π', ξ' are just refinements of π, ξ . Suppose $|\pi \sigma^i| = n > 1$ and $|\xi \delta^i| = m > 1$. If $n \geq m$ we subdivide ${}^\pi\sigma^i$ into m segments where the first $m-1$ segments have 1 element and the last segment has $n-m+1$ elements; ${}^\xi\delta^i$ is subdivided into m segments, each consisting of 1 element. As $({}^\pi\sigma^i)B({}^\xi\delta^i)$, each of the refined segments matches its corresponding segment and all of the δ segments are of length 1. If $m > n$ we proceed similarly. π', ξ' have the desired properties. \square

The above lemma allows us to reason about segments using case analysis, where the three cases are: both segments are of length 1, the right segment is of length 1 and the left of length greater than 1, and the left segment is of length 1 and the right of length greater than 1. The definitions of well-founded simulation and well-founded bisimulation are to some extent based on these three cases. I henceforth assume that all segments are of this form, as this assumption simplifies some of the arguments.

3.1 Stuttering Simulation

A relation on $B \subseteq S \times S$ where $\mathcal{M} = \langle S, \dashrightarrow, L \rangle$ is a stuttering simulation, if for every s, w such that sBw , s and w are identically labeled and every fullpath starting at s can be matched by some fullpath starting at w .

Definition 2. (*Stuttering Simulation (STS)*)

$B \subseteq S \times S$ is a stuttering simulation on TS $\mathcal{M} = \langle S, \dashrightarrow, L \rangle$ iff for all s, w such that sBw :

- (Sts1) $L.s = L.w$
(Sts2) $\langle \forall \sigma : fp.\sigma.s : \langle \exists \delta : fp.\delta.w : match(B, \sigma, \delta) \rangle \rangle$

We show that there is a greatest STS. This is accomplished by showing that an arbitrary union of STS's is an STS, hence, there is a greatest STS. We then show that the identity relation is an STS and if A and B are STS's so is $A;B$ —the composition of A and B . Together, the above results imply that the reflexive, transitive closure of an STS is an STS and that the greatest STS is a preorder.

Lemma 2. $(B \subseteq C) \Rightarrow [match(B, \sigma, \delta) \Rightarrow match(C, \sigma, \delta)]$

Proof

$$\begin{aligned}
& match(C, \sigma, \delta) \\
\equiv & \{ \text{Definition of } match, corr \} \\
& \langle \exists \pi, \xi \in INC :: \langle \forall i \in \mathbb{N} :: (\pi \sigma^i)C(\xi \delta^i) \rangle \rangle \\
\Leftarrow & \{ B \subseteq C, \text{monotonicity} \} \\
& \langle \exists \pi, \xi \in INC :: \langle \forall i \in \mathbb{N} :: (\pi \sigma^i)B(\xi \delta^i) \rangle \rangle \\
\equiv & \{ \text{Definition of } match, corr \} \\
& match(B, \sigma, \delta) \quad \square
\end{aligned}$$

Lemma 3. *Let \mathcal{C} be a set of STS's on TS \mathcal{M} , then $G = \langle \cup B : B \in \mathcal{C} : B \rangle$ is an STS on \mathcal{M} .*

Proof If sGw , sBw for some $B \in \mathcal{C}$. $L.s = L.w$ since sBw and Sts1. Our remaining proof obligation follows.

$$\begin{aligned}
& \langle \forall \sigma : fp.\sigma.s : \langle \exists \delta : fp.\delta.w : match(G, \sigma, \delta) \rangle \rangle \\
\Leftarrow & \{ B \subseteq G, \text{lemma 2, monotonicity} \} \\
& \langle \forall \sigma : fp.\sigma.s : \langle \exists \delta : fp.\delta.w : match(B, \sigma, \delta) \rangle \rangle \quad \square
\end{aligned}$$

Corollary 1 *For every TS \mathcal{M} , there is a greatest STS on \mathcal{M} .*

Proof Let $G = \langle \cup B : B \in \mathcal{C} : B \rangle$ where $\mathcal{C} = \{B : B \text{ is an STS on } \mathcal{M}\}$. G is an STS by lemma 3, hence, the greatest. \square

Lemma 4. *If R and S are STS's, so is $T = R;S$.*

Proof Let aTb . By definition of T , there is an x such that aRx and xSb . We show Sts1 and Sts2.

$L.a = L.b$ by transitivity of equality, thus Sts1 holds.

We show $\langle \forall \sigma : fp.\sigma.a : \langle \exists \delta : fp.\delta.b : match(T, \sigma, \delta) \rangle \rangle$, namely that Sts2 holds. Since R is an STS and aRx , there exist γ, π, ξ such that $fp.\gamma.x$ and $corr(R, \sigma, \pi, \gamma, \xi)$, i.e., $match(R, \sigma, \gamma)$. Similarly, since S is an STS and xSb , there exist δ, ζ, ρ such that $fp.\delta.b$ and $corr(S, \gamma, \zeta, \delta, \rho)$, i.e., $match(S, \gamma, \delta)$. We show $match(T, \sigma, \delta)$, thereby proving Sts2. To do this, we must relate states in σ to states in δ . Here is an outline of the proof. First, we show that if a state in σ is related by T to a state in δ , then all the states in the corresponding

segments are also related. We then work at the segment level and show how to merge segments in σ and δ to get new segments that match under T . We start by defining a function C from segments in σ to sets of segments in δ that can be thought of as performing the merging in δ . Finally, we use C to define the actual partitions on σ and δ , which are then shown to match under T . \square

Here is an outline of the proof. In step (1), we show that if a state in σ is related by T to a state in δ , then all the states in the corresponding segments are also related. We then work at the segment level and show how to merge segments in σ and δ to get new segments that match under T . We start in step (2) by defining a function C from segments in σ to sets of segments in δ that can be thought of as performing the merging in δ . In step (3), we use C to define the actual partitions on σ and δ , which are then shown to match under T .

1. $L : \mathbb{N} \rightarrow \mathbb{N}$ (Lower) and $U : \mathbb{N} \rightarrow \mathbb{N}$ (Upper) are defined as follows.

- $L.i = j$ such that $\xi.i \in [\zeta.j, \zeta(j+1))$
- $U.i = j$ such that $\xi(i+1) - 1 \in [\zeta.j, \zeta(j+1))$

Note that $\xi.i$ is the index of the first element in ${}^\xi\gamma^i$ and $\xi(i+1) - 1$ is the index of the last element in ${}^\xi\gamma^i$. These elements are each in some segment of γ with respect to ζ . We show that every element in ${}^\pi\sigma^i$ is related to every element in ${}^\rho\delta^k$ for $k \in [L.i, U.i]$, i.e., $\langle \forall i, k : k \in [L.i, U.i] : ({}^\pi\sigma^i)T({}^\rho\delta^k) \rangle$. To see this, let $k \in [L.i, U.i]$, $t \in {}^\rho\delta^k$. By $\text{corr}(S, \gamma, \zeta, \delta, \rho)$, for all $w \in {}^\zeta\gamma^k$, we have wSt . Since ${}^\xi\gamma^i$ intersects ${}^\zeta\gamma^k$ —by definition of U and L —there is $u \in {}^\zeta\gamma^k$ such that $u \in {}^\xi\gamma^i$ and uSt . Since $\text{corr}(R, \sigma, \pi, \gamma, \xi)$, for all $s \in {}^\pi\sigma^i$, sRu , and by definition of T , $\langle \forall s \in {}^\pi\sigma^i :: sTt \rangle$.

2. We define C (corresponds), a function that identifies which segments of δ under ρ correspond to ${}^\pi\sigma^i$. We define $C(-1)$ to be $\{-1\}$.
 - $C.i = \{U.i\}$ if $U.i \in C(i-1)$ or $U.i = \max.C(i-1) + 1$, otherwise
 - if $U.i = U(i+1)$, $C.i = (\max.C(i-1), U.i)$, otherwise
 - $C.i = (\max.C(i-1), U.i]$
3. To prove $\text{match}(T, \sigma, \delta)$ we will exhibit an explicit partition. We do this by defining τ and ϕ as follows.
 - $\tau.0 = 0$
 - $\tau(i+1) = \tau.i + j$ where $j = \min.\{k \in \mathbb{N} : C(\tau.i + k) \neq C(\tau.i)\}$
 - $\phi.i = \min.C(\tau.i)$

The idea is that τ removes duplicate $C.i$'s (duplicate $C.i$'s indicate that neighboring segments of σ under π should be merged into one segment). We now prove $\text{corr}(T, \sigma, \pi \circ \tau, \delta, \rho \circ \phi)$ that expanded is $\langle \forall i, m, n : i \in \mathbb{N} \wedge m \in [\tau.i, \tau(i+1)) \wedge n \in [\phi.i, \phi(i+1)) : ({}^\pi\sigma^m)T({}^\rho\delta^n) \rangle$. Note $m \in [\tau.i, \tau(i+1)) \Rightarrow C.m = C(\tau.i)$ because τ collapses similar $C.i$'s; hence, $({}^\pi\sigma^m)T({}^\rho\delta^j)$ for every $j \in C(\tau.i)$ since $j \in [L(\tau.i), U(\tau.i)]$, but $n \in [\phi.i, \phi(i+1)) \equiv n \in [\min.C(\tau.i), \min(C.\tau(i+1))] \equiv n \in C(\tau.i)$, hence, $({}^\pi\sigma^m)T({}^\rho\delta^n)$. \square

Reasoning about simulations and bisimulations can be tricky. For example, Basten shows that the “simple proof” claimed for the folklore result stating that branching bisimulation (a type of bisimulation closely related to stuttering bisimulation) is an equivalence relation is wrong [2].

Lemma 5. *The reflexive, transitive closure of an STS is an STS.*

Proof The reflexive, transitive closure of B , denoted B^* , can be written as $\langle \cup i \in \mathbb{N} :: B^i \rangle$. The identity relation, B^0 , is an STS and by induction on the natural numbers, using lemma 4, so is B^i for every $i > 0$. By lemma 3 so is $\langle \cup i \in \mathbb{N} :: B^i \rangle$. \square

Theorem 1. *Given TS \mathcal{M} , there is a greatest STS on \mathcal{M} , which is a preorder.*

Proof By corollary 1, there is a greatest STS G on \mathcal{M} . By lemma 5, G^* is also an STS. Now, $G^* \supseteq G$, by definition of G^* and $G^* \subseteq G$ since G is the greatest STS. Thus, $G = G^*$. \square

STSs are not closed under intersection or negation, as the following lemma shows.

Lemma 6. *STSs are not closed under intersection or negation.*

Proof Consider a TS with two states with different labels. The identity relation is an STS, but its negation is not. An example showing that STSs are not closed under intersection appears in figure 1. \square




Fig. 1. An example showing that STSs are not closed under intersection. The transition relation is denoted by a dashed arrow. There are two stuttering simulation relations denoted by a dashed line and a solid line. Every pair of states in the same strongly connected component induced by the dashed line are in the first stuttering simulation relation. Similarly, every pair of states in the same strongly connected component induced by the solid line are in the second stuttering simulation relation. Only the top two states are related in the intersection, but they do not have matching children.

Theorem 2. *Let B be a STS on \mathcal{M} and let sBw . For every ACTL* $\setminus X$ formula f , if $\mathcal{M}, w \models f$ then $\mathcal{M}, s \models f$.*

Proof The proof is by induction on the structure of ACTL* $\setminus X$ formulas. \square

3.2 Well-Founded Simulation

In order to check that a relation is an STS, we have to show that infinite sequences “match”. This can be problematic when using computer-aided verification techniques. I present the notion of a *well-founded simulation* to remedy this situation. To show that a relation is a well-founded simulation, we need only check local properties; this is analogous to proving program termination by exhibiting a function that maps states into a well-founded relation and showing that the function decreases during every step of the program. As mentioned previously, the intuition is that for every pair of states s, w that are related by an STS and u such that $s \dashrightarrow u$, there are essentially three cases: either there is a v such that $w \dashrightarrow v$ and u is related to v , or u is related to w , or there is a v such that $w \dashrightarrow v$ and s is related to v . In the last two cases, we must also ensure that we do not have an infinite sequence of states, each of which is related to a single state. This is where the well-founded relation comes in: we must show that in these cases there is an appropriate measure function into a well-founded relation that decreases. Formally, we have:

Definition 3. (*Well-Founded Simulation (WFS)*) $B \subseteq S \times S$ is a well-founded simulation on TS $\mathcal{M} = \langle S, \dashrightarrow, L \rangle$ iff:

- (Wfs1) $\langle \forall s, w \in S : sBw : L.s = L.w \rangle$
- (Wfs2) There exists functions, $rankt : S \times S \rightarrow W, rankl : S \times S \times S \rightarrow \mathbb{N}$, such that $\langle W, \prec \rangle$ is well-founded, and
 - $\langle \forall s, u, w \in S : sBw \wedge s \dashrightarrow u :$
 - (a) $\langle \exists v : w \dashrightarrow v : uBv \rangle \vee$
 - (b) $\langle uBw \wedge rankt(u, w) \prec rankt(s, w) \rangle \vee$
 - (c) $\langle \exists v : w \dashrightarrow v : sBv \wedge rankl(v, s, u) < rankl(w, s, u) \rangle \rangle$

3.3 Equivalence

In this section, I show that well-founded simulation completely characterizes stuttering simulation. Thus, we can think of well-founded simulation as a sound and complete proof rule.

Proposition 1 (*Soundness*) *If B is a WFS, then it is an STS.*

Proof Let aBb ; we need to show Sts1 and Sts2. $L.a = L.b$ since B is a WFS (Wfs1), thus Sts1 holds. We show $\langle \forall \sigma : fp.\sigma.a : \langle \exists \delta : fp.\delta.b : match(B, \sigma, \delta) \rangle \rangle$, namely that Sts2 holds. Suppose $fp.\sigma.a$. We define fullpath δ and increasing sequences π, ξ recursively as follows: $\delta.0 = b, \pi.0 = 0, \xi.0 = 0$. The idea is that from $\pi.i, \xi.i, \delta(\xi.i)$ we can define $\pi(i+1), \xi(i+1), \delta(\xi(i+1))$ with $\pi \sigma^i, \xi \delta^i$ matching.

$\delta.0 = b, \pi.0 = 0, \xi.0 = 0$. We write σ^i, δ^i instead of $\pi \sigma^i, \xi \delta^i$, respectively. Also, to make the following more readable, let $s = \sigma(\pi.i), u = \sigma(\pi.i+1)$, and $w = \delta(\xi.i)$. The idea is that from $\pi.i, \xi.i, \delta(\xi.i)$ we can define $\pi(i+1), \xi(i+1), \delta(\xi(i+1))$ with σ^i, δ^i matching. There are three cases:

1. Wfs2a holds, *i.e.*, $\langle \exists v : w \dashrightarrow v : uBv \rangle$. Let $\pi(i+1) = \pi.i + 1, \xi(i+1) = \xi.i + 1, \delta^i = \langle w \rangle, \delta.\xi(i+1) = v$.
2. Wfs2a does not hold, but Wfs2b does, *i.e.*, $\neg \langle \exists v : w \dashrightarrow v : uBv \rangle \wedge uBw \wedge \text{rankt}(u, w) < \text{rankt}(s, w)$. Let n be the maximum number such that $\langle \forall l : 1 \leq l \leq n : \sigma(\pi.i + l)Bw \wedge \neg \langle \exists v : w \dashrightarrow v : \sigma(\pi.i + l)Bv \rangle \rangle$. Note $n \geq 1$ since the above holds for $u = \sigma(\pi.i + 1)$. Also, by the well-foundedness of rankt , n must be finite.
Wfs2 has to hold between $\sigma(\pi.i + n), w, \sigma(\pi.i + n) \dashrightarrow \sigma(\pi.i + n + 1)$. By the maximality of n , we have the following two cases:
 - (a) If $\langle \exists v : w \dashrightarrow v : \sigma(\pi.i + n + 1)Bv \rangle$, then $\sigma(\pi.i + n + 1), v$ mark the beginning of the $(i+1)^{\text{th}}$ segments and $\pi(i+1) = \pi.i + n + 1, \xi(i+1) = \xi.i + 1, \delta^i = \langle w \rangle, \delta.\xi(i+1) = v$.
 - (b) Otherwise, $\neg \langle \sigma(\pi.i + n + 1)Bw \rangle$, but now Wfs2b cannot hold, but neither can Wfs2c, because if it does $\langle \exists v : w \dashrightarrow v : \sigma(\pi.i + n)Bv \rangle$, but by definition of n , this cannot be; hence, Wfs2a must hold, *i.e.*, $\langle \exists v : w \dashrightarrow v : \sigma(\pi.i + n + 1)Bv \rangle$, which is our first case.
3. Only Wfs2c holds. Let $\mathbf{v} = \langle v.0 = w, v.1 = v, \dots, v.n \rangle$ be a maximal (with respect to prefix order) sequence such that $\langle \forall l : 0 \leq l < n : v.l \dashrightarrow v(l+1) \rangle$ and only Wfs2c holds for $s, v.l, s \dashrightarrow u$ for all $l < n$. Such a sequence must contain $v.0 = w$ and $v.1 = v$, so $n \geq 1$; in addition, the sequence is finite by the well-foundedness of rankl . Since $sBv.n$, Wfs2 must hold between $s, v.n, s \dashrightarrow u$, specifically either Wfs2a or Wfs2b (by maximality of \mathbf{v}), but Wfs2b cannot hold because if it did we would have $uBv.n$, but then Wfs2a would hold between $s, v(n-1), s \dashrightarrow u$; hence, Wfs2a must hold between $s, v.n, s \dashrightarrow u$, *i.e.*, there exists an x such that $v.n \dashrightarrow x$ and uBx . Let $\pi(i+1) = \pi.i + 1, \xi(i+1) = \xi.i + n + 1, \delta^i = \langle v.0, \dots, v.n \rangle, \delta.\xi(i+1) = x$; hence, u, x mark the beginning of the $(i+1)^{\text{th}}$ segment. \square

We now prove that every STS is a WFS. For the proof, we have to exhibit the rank functions as per the definition of WFS. Here is a high-level overview.

The value of $\text{rankt}(s, w)$ is important only if sBw , as otherwise there are no restrictions required by the definition of WFS. If sBw , then consider the largest subtree of the computation tree rooted at s such that no node in the subtree matches a successor of w . The “rank” (a kind of height) of this subtree is the value of $\text{rankt}(s, w)$. The “rank” of s is greater than the “rank” of any of its children in the tree, so case Wfs2b is satisfied.

The value of $\text{rankl}(w, s, u)$ is important only if sBw and $s \dashrightarrow u$, as otherwise there are no restrictions required by the definition of WFS. If sBw and $s \dashrightarrow u$, then $\text{rankl}(w, s, u)$ is the length of the shortest path from w that matches s, u . In the case of Wfs2c, we can choose the next successor of w in this path to satisfy the condition.

Given a TS $\mathcal{M} = \langle S, \dashrightarrow, L \rangle$, the notion of the *computation tree* rooted at a state $s \in S$ is standard. It is the tree obtained by unfolding \mathcal{M} starting from s and can be defined as follows. The nodes of the tree are finite sequences over S . The tree is defined to be the smallest tree satisfying the following.

1. The root is $\langle s \rangle$.
2. If $\langle s, \dots, w \rangle$ is a node and $w \dashrightarrow v$, then $\langle s, \dots, w, v \rangle$ is a node whose parent is $\langle s, \dots, w \rangle$.

Definition 4. (*tree*) Given an STS B , if $\neg(sBw)$, then $tree(s, w)$ is the empty tree, otherwise $tree(s, w)$ is the largest subtree of the computation tree rooted at s such that for every non-root node of the tree, $\langle s, \dots, x \rangle$, we have that xBw and $\langle \forall v : w \dashrightarrow v : \neg(xBv) \rangle$.

Lemma 7. *Every path of $tree(s, w)$ is finite.*

Proof Suppose not, then there exists a path σ such that $fp.\sigma.s$. Consider σ' , the infinite sequence starting at $\sigma.1$, but otherwise identical to σ . We have $(\sigma'.0)Bw$, so there exists δ such that $match(B, \sigma', \delta)$. By construction, $\langle \forall v : w \dashrightarrow v : \neg((\sigma'.1)Bv) \rangle$, hence, $\delta.1$ marks the beginning of the second segment in δ . Say $\sigma'.j$ marks the beginning of the second segment of σ' , then $(\sigma'.j)B(\delta.1)$, but by construction of $tree(s, w)$ this is impossible. \square

Since the child relation on nodes in $tree.s$ is well-founded, we can recursively define a labeling function, l , that assigns an ordinal to nodes in the tree as follows: $l.n = \langle \cup c : c \text{ is a child of } n : (l.c) + 1 \rangle$. This is the standard “rank” function encountered in set theory [13]. We use the convention that the label of a tree is the label of its root.

Lemma 8. *If $|S| \preceq \kappa$, where κ is an infinite cardinal (i.e., $\omega \preceq \kappa$) then for all $s, w \in S$, $tree(s, w)$ is labeled with an ordinal of cardinality $\preceq \kappa$.*

Proof The i^{th} level of the tree consists of all nodes of distance i from the root. Since paths are finite, i is a natural number. We can show by induction on \mathbb{N} that for all i , the number of nodes at level i is at most κ . At level 0 it is easy to see as there is only 1 node. Suppose level i contains at most κ nodes. Since each node has at most κ children, level $i + 1$ can have at most κ^2 nodes, which is at most κ . We have shown each level has at most κ nodes. Since there is a countable number of levels, the total number of nodes is at most $\omega \times \kappa = max.\{\omega, \kappa\} = \kappa$.

Consider any function $F = \{(n, \alpha) : n \in tree.s \text{ and } \alpha \text{ is an ordinal}\}$ where R , the range of F , is an ordinal (this guarantees that for every α in R , we have $\beta \in R$ for all $\beta \prec \alpha$). Note that the labeling described above is such a function. We have that $F.s$ is an ordinal of cardinality at most κ , otherwise we have a surjection from κ to a larger cardinal, a contradiction. \square

Lemma 9. *If $sBw, s \dashrightarrow u, u \in tree(s, w)$ then $l.tree(u, w) \prec l.tree(s, w)$.*

Proof Since $u \in tree(s, w)$, by construction of $tree(s, w)$, the subtree rooted at u is $tree(u, w)$. By definition of l , $l.tree(u, w) \prec l.tree(s, w)$. \square

Definition 5. (*length*) Given B , an STS, $length(w, s, u) = 0$ if $\neg(sBw)$ or $\neg(s \dashrightarrow u)$, otherwise $length(w, s, u)$ is the length of the shortest initial segment starting at w that matches $\langle s, u \rangle$. Formally:

$$length(w, s, u) = \langle \min \sigma, \delta, \pi, \xi : fp.\sigma.s \wedge \sigma.1 = u \wedge fp.\delta.w \wedge \pi, \xi \in INC \wedge corr(B, \sigma, \pi, \delta, \xi) : |^{\xi} \delta^0 | \rangle$$

As sBw and $s \dashrightarrow u$, the above range is non-empty and $\text{length}(w, s, u) \in \mathbb{N}$.

Lemma 10. *If $sBw, s \dashrightarrow u$ and $\neg(\exists \sigma, \delta, \pi, \xi : fp.\sigma.s \wedge \sigma.1 = u \wedge fp.\delta.w \wedge \pi, \xi \in INC : corr(B, \sigma, \pi, \delta, \xi) \wedge \xi\delta^0 = \langle w \rangle)$, then $\langle \exists v : w \dashrightarrow v : \text{length}(v, s, u) < \text{length}(w, s, u) \wedge sBv \rangle$.*

Proof First, note that by assumption $\text{length}(w, s, u) \geq 2$, because for every σ, δ, π, ξ such that $fp.\sigma.s \wedge \sigma.1 = u \wedge fp.\delta.w \wedge \pi, \xi \in INC \wedge corr(B, \sigma, \pi, \delta, \xi)$, we have $\neg(uB\delta.1)$. Hence, ${}^\pi\sigma^0 = \langle s \rangle, |{}^\xi\delta^0| \geq 2$, and $\text{length}(w, s, u) \geq 2$. Let v be the successor of w in some minimal initial segment, then $\text{length}(v, s, u) = \text{length}(w, s, u) - 1 \geq 1$ and sBv , as required. \square

Proposition 2 (Completeness) *If B is an STS, then B is a WFS.*

Proof Wfs1 follows from Sts1. Let $W = (|S| + \omega)^+$. Note that $+$ denotes cardinal arithmetic; we add ω to $|S|$ to guarantee that we have an infinite cardinal; κ^+ is the successor cardinal to κ .

Clearly, (W, \prec) is well-founded. Let $\text{rankt} = l.\text{tree}$ and let $\text{rankl} = \text{length}$. Let sBw and $s \dashrightarrow u$. There are three cases:

1. $\langle \exists v : w \dashrightarrow v : uBv \rangle$.

By lemma 1, if (1) does not hold, then for every σ, δ, π, ξ such that $fp.\sigma.s \wedge \sigma.1 = u \wedge fp.\delta.w \wedge \pi, \xi \in INC \wedge corr(B, \sigma, \pi, \delta, \xi)$, either s marks the end of ${}^\pi\sigma^0$ or w marks the end of ${}^\xi\delta^0$, but not both.

2. $\langle \exists \sigma, \delta, \pi, \xi : fp.\sigma.s \wedge \sigma.1 = u \wedge fp.\delta.w \wedge \pi, \xi \in INC \wedge \xi\delta^0 = \langle w \rangle : corr(B, \sigma, \pi, \delta, \xi) \rangle$ and (1) does not hold. This implies that $|{}^\pi\sigma^0| > 1$, uBw , and $u \in \text{tree}(s, w)$; hence, $\text{rankt}(u, w) \prec \text{rankt}(s, w)$ by lemma 9.
3. If (1) and (2) do not hold, we must have $\neg(\exists \sigma, \delta, \pi, \xi : fp.\sigma.s \wedge \sigma.1 = u \wedge fp.\delta.w \wedge \pi, \xi \in INC : corr(B, \sigma, \pi, \delta, \xi) \wedge \xi\delta^0 = \langle w \rangle)$. By lemma 10 and the definition of rankl , $\langle \exists v : w \dashrightarrow v : \text{rankl}(v, s, u) < \text{rankl}(w, s, u) \wedge sBv \rangle$. \square

Theorem 3. (Equivalence) *B is an STS iff B is a WFS.*

Proof By propositions 1 and 2. \square

A consequence of the above theorem is that all of the properties proved for STSs carry over to WFSs; I use this fact freely, without reference, in the sequel.

3.4 Refinement

Up to this point, I have developed a theory for relating states. I now show how to apply the theory to transition systems. In this section, I define a notion of refinement and show that STSs can be used in a compositional fashion. For states s and w , I write $s \sqsubseteq w$ to mean that there is an STS B such that sBw .

By theorem 1, $s \sqsubseteq w$ iff sGw , where G is the greatest STS. I now lift this idea to transition systems.

Definition 6. (*Simulation Refinement*) Let $\mathcal{M} = \langle S, \rightarrow, L \rangle$, $\mathcal{M}' = \langle S', \rightarrow', L' \rangle$, and $r : S \rightarrow S'$. We say that \mathcal{M} is a simulation refinement of \mathcal{M}' with respect to refinement map r , written $\mathcal{M} \sqsubseteq_r \mathcal{M}'$, if there exists a relation, B , such that $\langle \forall s \in S :: sB(r.s) \rangle$ and B is an STS on the TS $\langle S \uplus S', \rightarrow \uplus \rightarrow', \mathcal{L} \rangle$, where $\mathcal{L}.s = L'(s)$ for s an S' state and $\mathcal{L}.s = L'(r.s)$ otherwise.

In the above definition, it helps to think of \mathcal{M}' as the specification and \mathcal{M} as the implementation. That \mathcal{M} is a simulation refinement of \mathcal{M}' with respect to r implies that every visible behavior of \mathcal{M} (where what is visible depends on r) is a behavior of \mathcal{M}' . There are often other considerations, *e.g.*, it might be that \mathcal{M} and \mathcal{M}' have certain states that are “initial”. In this case one might wish to show that initial states in \mathcal{M} are mapped to initial states in \mathcal{M}' .

One has a great deal of flexibility in choosing refinement maps. The danger is that by choosing a complicated refinement map, one can bypass the verification problem all together. To make this point clear, let PRIME be the system whose single behavior is the sequence of primes and let NAT be the system whose single behavior is the sequence of natural numbers. We do not consider NAT to be an implementation of PRIME, but using the refinement map from NAT to PRIME that maps i to the i^{th} prime, we can indeed prove the peculiar theorem that NAT is a refinement of PRIME. The moral is that we must be careful to not bypass the verification problem with the use of such refinement maps. Simple refinement maps with a clear relationship between implementation states and their image under the map are best. To check that a refinement proof is meaningful, one has to look at the specification system and the refinement map. It is important to look at the specification system to check that in fact it corresponds to what you expected. It is important to look at the refinement map for the reasons stated above, namely, one has to check that refinement maps map implementation states to “related” specification states. It may occur to the reader that we should require that refinement maps preserve (part of) the labeling function of the implementation states. In fact, we consider such a restriction shortly, but the reason we do not place restrictions on refinement maps is that it is not a priori apparent what the “reasonable” relationships between implementation states and specification states might be, *e.g.*, suppose that the specification system represents numbers in decimal but the implementation system represents numbers in binary, or that numbers in the specification are spread across several registers in the implementation, and so on.

Often refinement maps are especially clear, which makes it easy to check that they are in fact appropriate. Suppose that associated with states is a set of variables, each of a particular type. Furthermore, suppose that the variables in the implementation are a superset of the variables in the specification and that the refinement map just hides the implementation variables that do not appear in the specification. Then, it is clear that the refinement map is a reasonable one. More precisely, given TS $\mathcal{M} = \langle S, \rightarrow, L \rangle$, if L has the following structure, we say that \mathcal{M} is *typed*.

Let VAR_S be a set and let $TYPE$ be a function whose domain is VAR_S . Think of VAR_S as the variables of TS \mathcal{M} , where $TYPE$ gives the type of the

variables. For all $s \in S$, let $L.s$ be a function from $VARS$ such that $L.s.v \in TYPE.v$. The lemma below shows why the appropriateness of refinement maps that hide some of the implementation variables is easy to ascertain.

Lemma 11. *If $\mathcal{M} = \langle S, \twoheadrightarrow, L \rangle \sqsubseteq_r \mathcal{M}' = \langle S', \twoheadrightarrow', L' \rangle$, both \mathcal{M} and \mathcal{M}' are typed TSs, and $L'(r.s) = L.s|_V$, then for every pair of states $s, r.s$ such that $s \in S$, and every ACTL* $\setminus X$ formula, f , built out of expressions that only depend on variables in V , we have $\mathcal{M}', r.s \models f \Rightarrow \mathcal{M}, s \models f$.*

Proof Since $\mathcal{M} \sqsubseteq_r \mathcal{M}'$ we have that $s \sqsubseteq r.s$ in the disjoint union of \mathcal{M} and \mathcal{M}' . Recall that the refinement map r does not alter the value of variables in V . The result follows from theorem 2. \square

Lemma 12. *If B is an STS on TS $\mathcal{M} = \langle S \supseteq S_1 \cup S_2, \twoheadrightarrow, L \rangle$, $S_1 \cap S_2 = \emptyset$, states in S_1 can only reach states in S_1 , and states in S_2 can only reach states in S_2 , then $\hat{B} = \{\langle s_1, s_2 \rangle : s_1 \in S_1 \wedge s_2 \in S_2 \wedge s_1 B s_2\}$ is an STS on \mathcal{M} .*

Proof Suppose $s \hat{B} w$, then $s \in S_1$, $w \in S_2$, and $s B w$. Let σ satisfy $fp.\sigma.s$. Since $s B w$, there is a δ such that $fp.\delta.w$ and $match(B, \sigma, \delta)$. Since all states reachable from s are in S_1 , σ contains only S_1 states and similarly δ contains only S_2 states, therefore, $match(\hat{B}, \sigma, \delta)$ by the definitions of \hat{B} and $match$. \square

Theorem 4. (Composition) *If $\mathcal{M} \sqsubseteq_r \mathcal{M}'$ and $\mathcal{M}' \sqsubseteq_q \mathcal{M}''$ then $\mathcal{M} \sqsubseteq_{r;q} \mathcal{M}''$.*

Proof Let $\mathcal{M} = \langle S, \twoheadrightarrow, L \rangle$, $\mathcal{M}' = \langle S', \twoheadrightarrow', L' \rangle$, and $\mathcal{M}'' = \langle S'', \twoheadrightarrow'', L'' \rangle$. Since $\mathcal{M} \sqsubseteq_r \mathcal{M}'$, we have an STS, A , such that $\langle \forall s \in S :: s A(r.s) \rangle$; by lemma 12 we can assume that $A \subseteq S \times S'$. Since $\mathcal{M}' \sqsubseteq_q \mathcal{M}''$, we have an STS, B , such that $\langle \forall s' \in S' :: s' B(q.s') \rangle$; by lemma 12 we can assume that $B \subseteq S' \times S''$. Letting C be $A; B$, we have, by properties of A and B , that $C \subseteq S \times S''$ and $\langle \forall s \in S :: s C(q(r.s)) \rangle$. By lemma 4, we have that C is an STS on the TS $\langle S \uplus S'', \twoheadrightarrow \uplus \twoheadrightarrow'', \mathcal{L} \rangle$, where $\mathcal{L}.s = L''(q(r.s))$ for s an S state and $\mathcal{L}.s = L''(s)$ otherwise. \square

4 The Linear Time Case

The theorem on the existence of refinement maps in the previous section does not apply to the linear time framework because simulation is a stronger property than trace containment. However, note that if we destroy the branching structure of transition system \mathcal{M} to obtain transition system \mathcal{M}' , then $\mathcal{M}' \sqsubseteq_r \mathcal{N}$ iff the set of infinite sequences of \mathcal{M} , labeled by r , is a subset of the set of sequences of \mathcal{N} . We can destroy the branching structure of \mathcal{M} by using an *oracle variable* to record values for every non-deterministic choice made along an infinite path in the computation tree of \mathcal{M} . We have thus sketched a proof of the existence of refinement maps in the linear time framework.

Theorem 5. *If the set of traces of \mathcal{M} is a subset of the traces of \mathcal{N} , then there exists \mathcal{M}' , a transition system obtained from \mathcal{M} by adding an oracle variable, and a refinement map r such that $\mathcal{M}' \sqsubseteq_r \mathcal{N}$.*

I now review the work of Abadi and Lamport on the existence of refinement maps. The review addresses the essential points, but is necessarily concise and readers are urged to read the full paper. I then present several examples, taken from Abadi and Lamport, that are used to justify the conditions appearing in their theorem. At the end of this section, I compare the two approaches.

4.1 Review of Abadi and Lamport Results

I begin by reviewing some initial definitions. A *behavior* is an infinite sequence and a *property* is a set of behaviors closed under finite stuttering. A *specification* is a (possibly infinite) state machine, consisting of externally visible components and internal components, and a *supplementary* property to represent fairness constraints. The *complete property* of a state machine is obtained by closing the set of behaviors allowed by the machine under (possibly infinite) stuttering. The *externally visible property* of a state machine is obtained by projecting the externally visible components of the complete property of the state machine. The *property* defined by a specification is obtained by intersecting the complete property of its state machine with the supplementary property. The *externally visible property* of a specification is obtained by projecting the externally visible components of the property of the specification.

We say that I , a “concrete” specification (the *Implementation*), *implements* S , an “abstract” specification (the *Specification*) if every externally visible behavior of I is also a behavior of S . Proving that I implements S can require reasoning about arbitrary sequences because one has to show that if I admits the behavior $\langle\langle e_0, z_0 \rangle, \langle e_1, z_1 \rangle, \dots, \langle e_n, z_n \rangle, \dots\rangle$, where the e_i correspond to the externally visible components and the z_i to the internal components, then S admits the behavior $\langle\langle e_0, y_0 \rangle, \langle e_1, y_1 \rangle, \dots, \langle e_n, y_n \rangle, \dots\rangle$. Notice that y_n can depend upon the entire sequence $\langle\langle e_0, z_0 \rangle, \langle e_1, z_1 \rangle, \langle e_2, z_2 \rangle, \dots\rangle$, which can make the proof difficult. We prefer to avoid such global reasoning and would rather reason locally *e.g.*, if there is a function f such that $\langle e_i, y_i \rangle = f(e_i, z_i)$, it can be used to prove that I preserves the safety property of S by reasoning about pairs of states instead of arbitrary sequences of states. If such a function also preserves liveness, it is called a *refinement mapping* and Abadi and Lamport prove the following completeness theorem, showing under what conditions refinement mappings exist.

Theorem 6. *If the machine-closed specification I implements S , a specification that has finite invisible nondeterminism and is internally continuous, then there is a specification I^h obtained from I by adding a history variable and a specification I^{hp} obtained from I^h by adding a prophecy variable such that there exists a refinement mapping from I^{hp} to S .*

The above theorem depends on various conditions, which I now explain. We say that a specification I is *machine-closed* if the supplementary property of I does not specify any safety property not already specified by the state machine of I . A specification S has *finite invisible nondeterminism* if for every finite

prefix of every behavior allowed by S , whenever an infinite nondeterministic choice is made, all but a finite part of that choice is immediately revealed in the externally visible components of the resulting state. A specification S is *internally continuous* if for every behavior, if the behavior is not allowed by the specification, this can be determined by looking at the externally visible part of the behavior and some *finite* portion of the complete behavior.

A *history variable* is used to extend the state space of a specification with a component that records past information, but in a way that does not affect the externally visible behaviors of the specification. Abadi and Lamport give five conditions that must be satisfied in order to show that if a specification S^h is obtained from a specification S by adding a history variable, then the two specifications define the same externally visible property. A *prophecy variable* is the dual of a history variable. Instead of recording past information, it guesses future information. Abadi and Lamport give six conditions that must be satisfied in order to show that if a specification S^p is obtained from a specification S by adding a prophecy variable, then the two specifications define the same externally visible property.

4.2 Examples due to Abadi and Lamport

This section contains several examples that Abadi and Lamport use to explain the conditions found in their completeness theorem. After the examples are introduced, I show how they can be handled using in my framework.

In the first example, system S is a three-bit clock, where only the low-order bit is externally visible and system \mathcal{I} is a one-bit clock. \mathcal{I} implements S since they have the same traces (up to stuttering). However, no refinement mapping can be used to show this because there is no way to define the internal state of S : consider an arbitrary refinement mapping, r , and suppose that $r(\langle 0 \rangle) = \langle 0, y_0 \rangle$ and $r(\langle 1 \rangle) = \langle 1, y_1 \rangle$, then either $\langle 0, y_0 \rangle$ does not transit to $\langle 1, y_1 \rangle$ or $\langle 1, y_1 \rangle$ does not transit to $\langle 0, y_0 \rangle$. This is one reason for introducing history variables and they are used to resolve the dilemma as follows. A history variable is added to \mathcal{I} and the variable “remembers” what \mathcal{I} did in the past. The result is that the state space of \mathcal{I} is expanded so that there are enough states to define an appropriate refinement mapping.

Using the approach outlined in this paper, we find that history variables are not needed as we can define a refinement map that maps the state in \mathcal{I} whose counter is 0 to any state in S whose low-order bit is 0 and similarly with the other state in \mathcal{I} . The equivalence relation that relates states with the same low-order bit in the disjoint union of the two systems is a stuttering simulation.

The second example is used to motivate the need for prophecy variables. System S chooses ten values non-deterministically and displays each in turn, whereas system \mathcal{I} chooses each value as it is displayed. \mathcal{I} implements S since they have the same traces, but there is no refinement mapping that can be used to show this, as should be clear. This example highlights that proofs based on refinement mappings are based on simulation, a branching time notion. Thus, when \mathcal{I} is not a stuttering simulation of S , one cannot directly use refinement

mappings to prove that \mathcal{I} implements \mathcal{S} (in the linear time sense). This is one reason for introducing prophecy variables and they are used to resolve the dilemma as follows. A prophecy variable is added to \mathcal{I} and the variable “guesses” what \mathcal{I} will decide to do in the future. There is now a refinement map, based on this prophecy variable, that can be used to show that \mathcal{I} implements \mathcal{S} . What is happening is that the prophecy variables allow one to push all of the branching in the computation tree of \mathcal{I} up to the root, thereby destroying the branching structure of \mathcal{I} .

This example shows why oracle variables are used in theorem 5. Note that from the branching point of view \mathcal{I} does not implement \mathcal{S} , *e.g.*, from the initial state in \mathcal{I} , there is a successor that has more than one possible future, a branching-time expressible property that does not hold in the initial state of \mathcal{S} . It seems that any refinement-based approach will need a mechanism for dealing with this issue, whether it is by destroying the branching structure of implementations, by adding branching structure to specifications, or by some combination thereof.

The third example shows why a prophecy variable is needed to slow down an implementation that runs faster than a specification, even though the specification is just stuttering. Both \mathcal{I} and \mathcal{S} specify clocks in which the hours and minutes are externally visible, whereas the seconds are internal. Furthermore, \mathcal{I} increments the clock by one second, whereas \mathcal{S} increments the clock by ten seconds. Both \mathcal{I} and \mathcal{S} have the same externally visible behaviors and proving that \mathcal{S} implements \mathcal{I} using refinement mappings is easy. However, there is no way to show that \mathcal{I} implements \mathcal{S} , because there is a behavior of \mathcal{S} such that the minute hand changes every six steps, but any behavior of \mathcal{I} requires at least sixty steps between minute hand changes.

In my formulation, the implementation is allowed to run faster than the specification, as we can both add and remove stuttering steps, thus it is easy to deal with the third example.

Abadi and Lamport present examples showing why the conditions of finite invisible nondeterminism and internal continuity are required. The examples are similar in that the implementation, \mathcal{I} , has the same externally visible behaviors as the specification, \mathcal{S} , but \mathcal{I} has a richer branching structure than \mathcal{S} , *i.e.*, \mathcal{S} is a simulation refinement of \mathcal{I} , but not the other way. As we have seen in the second example, above, prophecy variables can be used to deal with this problem. However, in these examples there are states in \mathcal{I} that are related to an infinite number of states in \mathcal{S} , and Abadi and Lamport’s prophecy variables cannot be used in this case (see their paper for the full details). To summarize, the conditions of internal continuity and finite invisible nondeterminism in the completeness theorem of Abadi and Lamport can be traced to the *branching* structure of the systems involved.

Oracle variables can be used in my approach to deal with these examples. The intuition is that oracle variables allow us to quantify over every possible non-deterministic choice and can be used to transform \mathcal{I} into a linear time equivalent system in which all nondeterministic choices have been made at the onset.

4.3 Comparison with the Approach of Abadi and Lamport

There are various differences between my approach and that of Abadi and Lamport. A major difference is that I deal with branching time notions because in the context of mechanical verification they provide certain advantages, as outlined above. However, in order to simplify the comparison, in this section I consider only the linear time aspects of my results.

There are differences in how stuttering is dealt with; namely, Abadi and Lamport allow infinite stuttering, whereas I do not. Consider the example of pipelined machine verification. Using the Abadi and Lamport approach, we would define the instruction set architecture using a state machine, say where every component is externally visible. By definition, the property generated by the state machine includes infinite stuttering, *e.g.*, it includes the behavior where nothing happens. Thus, a supplementary property would be used to rule out such behaviors by requiring that non-stuttering steps are eventually taken, a liveness property. In contrast, in my approach, every step of the transition system modeling the instruction set architecture corresponds to the execution of an instruction, with the stuttering being handled by the definition of stuttering simulation. Notice that no supplementary property is required. In addition, the condition that a pipelined machine makes progress is now a safety property, because the number of steps required is bounded by the number of stages in the pipeline [16].

Lamport and Abadi require that systems have the same externally visible states. They make the point that one cannot say whether the value 11111100 corresponds to -3 without knowing how to interpret a sequence of bits as an integer. They go on to say that given such an interpretation, they can translate the externally visible states to the appropriate representation. In my case, instead of having a separate interpretation phase, I allow refinement maps to alter the labels of states directly. I have found that in practice this extra power is necessary. For example, when proving that a pipelined machine implements the instruction set architecture, I have used refinement maps that either modify the value of the program counter (when using my “commit” approach to correctness) or modify the register file and memory (using the Burch and Dill “flushing” approach to correctness) [16]. The point is that when using my commit approach to correctness, if we consider the program counter to be externally visible then we cannot use the Abadi and Lamport approach to prove that a pipelined machine implements the instruction set architecture. Similarly, when using the Burch and Dill approach, if we consider the register file or memory to be externally visible, then we cannot use the Abadi and Lamport approach to prove that a pipelined machine implements the instruction set architecture.

The refinement mappings of Abadi and Lamport are required to preserve the supplementary property of the specification. As they point out, this is not a local condition, but one can apply local methods such as well-founded induction for the proof. Unfortunately, they do not provide any guidance on constructing such arguments. In my case, the proof of proposition 2 (if B is an STS, then B is a WFS) shows how to construct the appropriate well-founded relations

and measure functions, *rankt* and *rankl*. The proof also shows that two measure functions, one from pairs of states and one from triples of states to the naturals, are enough regardless of the transition systems involved.

Finally, my theorems are stronger than the ones given by Abadi and Lamport. For example, they show that even when \mathcal{S} is not internally continuous a refinement map exists to show that \mathcal{I} satisfies the safety property specified by \mathcal{S} . They continue “We do not know if anything can be said about proving arbitrary liveness properties.” Since my refinement theorems apply to any systems, a simple corollary is that, with my approach, refinement maps can always be used to prove both safety and liveness properties. This is something that we used in [17] where we used theorem proving to reduce an infinite-state system to a finite-state system in such a way that stuttering-insensitive properties, including liveness, were preserved. We then model checked the reduced system and were able to lift the results to the original system.

5 Conclusions

I have introduced compositional notions of refinement for stuttering simulation. I have shown that if one system refines another in the branching time framework, then a refinement map always exists, without relying on any of the conditions present in the approach taken by Abadi and Lamport, *e.g.*, machine closure, finite invisible nondeterminism, internally continuity, the use of history and prophecy variables, etc. I also showed that refinement maps always exist in the linear time framework, subject only to the use of oracle variables.

My main motivation is the mechanical verification of systems. Notions of refinement based on stuttering bisimulation have proved useful for this purpose [17, 16]. However, stuttering bisimulation is applicable only in limited contexts, as usually specifications contain more nondeterminism than implementations. Thus, I expect that stuttering simulation will turn out to be more useful than stuttering bisimulation.

References

- [1] M. Abadi and L. Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, 1991.
- [2] T. Basten. Branching bisimilarity is an equivalence indeed. *Information Processing Letters*, 58(3):141–147, 1996.
- [3] B. Bloom and R. Paige. Transformational design and implementation of a new efficient solution to the ready simulation problem. *Science of Computer Programming*, 24(3):189–220, 1995.
- [4] J. R. Burch and D. L. Dill. Automatic verification of pipelined microprocessor control. In *Computer-Aided Verification (CAV '94)*, volume 818 of *LNCS*, pages 68–80. Springer-Verlag, 1994.
- [5] K. Engelhardt and W.-P. de Roever. Generalizing Abadi & Lamport’s method to solve a problem posed by A. Pnueli. In *FME '93: Industrial-Strength Formal Methods*, volume 670 of *LNCS*, pages 294–313. Springer-Verlag, 1993.

- [6] R. J. v. Glabbeek. The linear time – branching time spectrum I; the semantics of concrete, sequential processes. In J. A. Bergstra, A. Ponse, and S. A. Smolka, editors, *Handbook of Process Algebra*, chapter 1, pages 3–99. Elsevier, 2001.
- [7] M. R. Henzinger, T. A. Henzinger, and P. W. Kopke. Computing simulations on finite and infinite graphs. In *IEEE Symposium on Foundations of Computer Science*, pages 453–462. IEEE Computer Society Press, 1995.
- [8] W. Hesselink. Eternity variables to simulate specifications. In E. Boiten and B. Möller, editors, *Proceedings Mathematics of Program Construction, Dagstuhl, 2002*, volume 2386 of *LNCS*, pages 117–130. Springer-Verlag, 2002.
- [9] B. Jonsson, A. Pnueli, and C. Rump. Proving refinement using transduction. *Distributed Computing*, 12(2–3):129–149, 1999.
- [10] M. Kaufmann, P. Manolios, and J. S. Moore, editors. *Computer-Aided Reasoning: ACL2 Case Studies*. Kluwer Academic Publishers, June 2000.
- [11] M. Kaufmann, P. Manolios, and J. S. Moore. *Computer-Aided Reasoning: An Approach*. Kluwer Academic Publishers, July 2000.
- [12] M. Kaufmann and J. S. Moore. ACL2 homepage. See URL <http://www.cs.utexas.edu/users/moore/acl2>.
- [13] K. Kunen. *Set Theory - An Introduction to Independence Proofs*, volume 102 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, 1980.
- [14] L. Lamport. What good is temporal logic? *Information Processing*, 83:657–688, 1983.
- [15] N. A. Lynch and F. W. Vaandrager. Forward and backward simulations – part I: untimed systems. *Information and Computation*, 121(2):214–233, 1995.
- [16] P. Manolios. Correctness of pipelined machines. In W. A. Hunt, Jr. and S. D. Johnson, editors, *Formal Methods in Computer-Aided Design-FMCAD 2000*, volume 1954 of *LNCS*, pages 161–178. Springer-Verlag, 2000.
- [17] P. Manolios, K. Namjoshi, and R. Sumners. Linking theorem proving and model-checking with well-founded bisimulation. In N. Halbwegs and D. Peled, editors, *Computer-Aided Verification-CAV '99*, volume 1633 of *LNCS*, pages 369–379. Springer-Verlag, 1999.
- [18] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1990.
- [19] F. Moller and S. A. Smolka. On the complexity of bisimulation. *ACM Computing Surveys*, 27(2):287–289, June 1995.
- [20] K. S. Namjoshi. A simple characterization of stuttering bisimulation. In *17th Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 1346 of *LNCS*, pages 284–296, 1997.
- [21] R. Paige and R. E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
- [22] D. Park. Concurrency and automata on infinite sequences. In *Theoretical Computer Science*, volume 104 of *LNCS*, pages 167–183. Springer-Verlag, 1981.
- [23] A. Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science*, pages 46–57, Providence, Rhode Island, 1977. IEEE.
- [24] A. Pnueli. Linear and branching structures in the semantics and logics of reactive systems. In W. Brauer, editor, *Proceedings of 12th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 194 of *LNCS*, pages 15–32. Springer-Verlag, 1985.
- [25] J. Sawada. *Formal Verification of an Advanced Pipelined Machine*. PhD thesis, University of Texas at Austin, Dec. 1999. See URL <http://www.cs.utexas.edu/users/sawada/dissertation/>.

- [26] J. Sawada. Verification of a simple pipelined machine model. In Kaufmann et al. [10], pages 137–150.
- [27] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time. In *STOC: ACM Symposium on Theory of Computing (STOC)*, pages 1–9, 1973.