More Anti-Chain Based Refinement Checking

Ting Wang¹, Songzheng Song², Jun Sun³, Yang Liu², Jin Song Dong², Xinyu Wang¹
and Shanping Li¹

¹ College of Computer Science and Technology, Zhejiang University
{qdw, wangxinyu, shan}@zju.edu.cn

² National University of Singapore
{songsongzheng@, tslliuya@, dongjs@comp.}nus.edu.sg

³ Singapore University of Technology and Design
sunjun@sutd.edu.sg

Abstract. Refinement checking plays an important role in system verification. It establishes properties of an implementation by showing a refinement relationship between the implementation and a specification. Recently, it has been shown that anti-chain based approaches increase the efficiency of trace refinement checking significantly. In this work, we study the problem of adopting anti-chain for stable failures refinement checking, failures-divergence refinement checking and probabilistic refine checking (i.e., a probabilistic implementation against a non-probabilistic specification). We show that the first two problems can be significantly improved, because the state space of the product model may be reduced dramatically. Though applying anti-chain for probabilistic refinement checking is more complicated, we manage to show improvements in some cases. We have integrated these techniques into the PAT model checking framework. Experiments are conducted to demonstrate the efficiency of our approach.

1 Introduction

Model checking has established itself as an effective technique for system verification. It works by exhaustively searching through the state space in order to show that an implementation model, in certain modeling language, satisfies a property. Properties are often specified using temporal logic formulae such as CTL or LTL, in other words, a language different from the modeling language. An alternative approach is called refinement checking. Different from temporal-logic based model checking, refinement checking shows a refinement relationship between two models in the same language, one modeling an implementation and one modeling a specification. If the specification satisfies certain property and the refinement relationship is strong enough to preserve the property, we imply that the property is satisfied by the implementation. A variety of refinement relationships have been defined, which preserve different classes of properties. For instance, safety can be verified by showing a trace refinement relationship. Combination of safety and liveness is verified by showing a stable failures refinement relationship if the system is divergence-free or otherwise by showing a failures-divergence refinement relationship. The readers are refer to [12] for a discussion on the expressiveness of different refinement.

Refinement checking has been traditionally used to verify CSP [11]. The success of the FDR refinement checker [1], which supports fully automatic checking of the above-mentioned refinement relationships, evidences the usefulness of refinement checking. Recently, *Sun et al.* extended the idea of automated trace refinement checking to probabilistic systems [14], which we refer to as probabilistic refinement checking in this work⁴. The idea is that, given a probabilistic implementation model (which has the semantics of a Markov Decision Process) and a non-probabilistic specification model, probabilistic refinement checking calculates the probability of the implementation exhibiting traces of the specification model. This is useful as, for instance, if the specification model captures desired system behaviors, the result is the probability of the implementation behaving ‘well’.

Due to the non-determinism in the specification, refinement checking often relies on the classic subset construction approach. The subset construction is used to build a deterministic finite-state automaton (DFA) from the specification, which is in general a non-deterministic finite-state automaton (NFA). Next, refinement checking works by computing the synchronous product of the implementation and transforming the problem into a reachability analysis problem in the product. In the worst case, the resultant DFA could have exponentially more states than the original NFA. As a result, refinement checking suffers from state space explosion. Recently, *Wulf et al.* proposed an approach (for solving the language universality problem and trace refinement checking) named anti-chain [16]. It has been shown that anti-chain based trace refinement checking outperforms the previous approaches significantly. The key point of anti-chain based approaches is that the complete subset construction and computing the complete state space of the product are avoided. Given that the existing approaches for checking other refinement relationships are all based on the subset construction, it is only naturally to investigate whether anti-chain can be used for better performance as it did for trace refinement checking.

In this work, we study three kinds of refinement checking, in particular, stable failures refinement, failures-divergence refinement and probabilistic refinement checking. The problem is non-trivial as we need to formally prove that anti-chain works with stable failures semantics and failures-divergence semantics. Furthermore, it is complicated for probabilistic refinement checking as omitting parts of the product would affect the probability. We make the following technical contribution. Firstly, we show that anti-chain can be readily used to improve stable failures refinement and failures-divergence refinement. Secondly, we show that anti-chain can be used to improve probabilistic refinement checking in some particular cases, using an iterative probability calculation method. Lastly, we implement the technique in the PAT model checker [13] and show improvement over existing approaches (significant for stable failures refinement and failures-divergence refinement).

Related works This work is related to research on anti-chain based model checking. *Wulf et al.* proposed the anti-chain based approach for checking the language universality and trace refinement of NFA [16]. It has been shown that the anti-chain based approach may outperform the standard ones by several orders of magnitude. Their fol-

⁴ Probabilistic refinement has been used by different researchers to mean different things.

lowing works show that significant improvements can also be brought to the model checking problem of LTL by using anti-chain based algorithms [8, 17]. Later Abdulla *et al.* improved the approach through exploiting a simulation relation on the states of NFA [2]. Remotely related are anti-chain based methods for solving other problems, e.g., the LTL realizability and synthesis problem [7, 10] and the universality and language inclusion problem of non-deterministic tree automata (NTA) [2, 6]. In our work, we focus on stable failures refinement, failures-divergence refinement and probabilistic refinement checking.

Organization The remainder of the paper is organized as follows. Section 2 reviews trace refinement and anti-chain based trace refinement checking. Section 3 presents algorithms for anti-chain based stable failures refinement checking and failures-divergence refinement checking. Section 4 shows that anti-chain can be used to improve probabilistic refinement checking (with a non-probabilistic specification model). Lastly, Section 5 concludes the paper.

2 Background

In this section, we review previous work on anti-chain based trace refinement checking.

2.1 Trace Refinement

Let Σ be a set of event names; τ denote an invisible event; and Σ_τ denote $\Sigma \cup \{\tau\}$.

Definition 1 (LTS). A labeled transition system (LTS) is a tuple $\mathcal{L} = (S, \text{init}, \text{Act}, T)$ where S is a set of states; $\text{init} \in S$ is an initial state; $\text{Act} \subseteq \Sigma_\tau$ is a set of events and $T : S \times \text{Act} \times S$ is a labeled transition relation.

For simplicity, $(s, e, s') \in T$ is sometimes written as $s \xrightarrow{e} s'$. An LTS is deterministic if and only if for all $s \in S$ and $s \in \Sigma_\tau$, if $s \xrightarrow{e} u$ and $s \xrightarrow{e} v$, then $u = v$. We write $\text{enable}(s)$ to denote the set $\{e \mid \exists s'. s \xrightarrow{e} s'\}$. We write $u \rightsquigarrow v$ if there exists a finite sequence of states $\langle s_0, s_1, \dots, s_n \rangle$ such that $s_i \xrightarrow{\tau} s_{i+1}$ for all i and $u = s_0$ and $v = s_n$. We write $u \rightsquigarrow^e v$ if $u \rightsquigarrow u'$ and $u' \xrightarrow{e} v'$ and $v' \rightsquigarrow v$. A finite sequence of events $\langle e_0, e_1, \dots, e_n \rangle$ is a trace of $\mathcal{L}$ if and only if there exists a sequence of state $\langle s_0, s_1, \dots, s_n \rangle$ such that $s_i \rightsquigarrow^{e_i} s_{i+1}$ for all i and $s_0 = \text{init}$. The traces of $\mathcal{L}$ are denoted as $\text{traces}(\mathcal{L})$.

Definition 2 (LTS Synchronous Product). Let $\mathcal{L}_i = (S_i, \text{init}_i, \text{Act}_i, T_i)$ where $i \in \{1, 2\}$ be two LTSs such that $\tau \notin \text{Act}_2$. The synchronous product of $\mathcal{L}_1$ and $\mathcal{L}_2$, written as $\mathcal{L}_1 \times \mathcal{L}_2$, is an LTS $\mathcal{L} = (S, \text{init}, \text{Act}, T)$ such that $S = S_1 \times S_2$; $\text{init} = (\text{init}_1, \text{init}_2)$; $\text{Act} = \text{Act}_1 \cup \text{Act}_2$; and T is the minimum labeled transition relation satisfying the following conditions.

- If $(s_1, \tau, s'_1) \in T_1$, $((s_1, s_2), \tau, (s'_1, s_2)) \in T$ for all $s_2 \in S_2$;
- If $(s_1, e, s'_1) \in T_1$ and $(s_2, e, s'_2) \in T_2$ and $e \notin \tau$, $((s_1, s_2), e, (s'_1, s'_2)) \in T$.

Notice that all events except τ are to be synchronized by the two LTSs.

Definition 3 (Trace Refinement). Let $\mathcal{L}_i$ where $i \in \{1, 2\}$ be two LTSs. $\mathcal{L}_1$ trace-refines $\mathcal{L}_2$ if and only if $\text{traces}(\mathcal{L}_1) \subseteq \text{traces}(\mathcal{L}_2)$.

The standard approach for trace refinement is based on the subset construction. That is, the specification LTS_2 is transformed into trace-equivalent deterministic LTS without τ -transitions through the process of determinization. Let $\mathcal{L} = (S, \text{init}, \text{Act}, T)$ be an LTS. The determinized LTS of $\mathcal{L}$ is $\text{det}(\mathcal{L}) = (S', \text{init}', \text{Act}', T')$ where $S' \subseteq 2^S$ is a set of sets of states, $\text{init}' = \{s \mid \text{init} \rightsquigarrow s\}$; $\text{Act}' = \text{Act} \setminus \{\tau\}$ and T' is a transition relation satisfying the following condition: $(N, e, N') \in T'$ if and only if $N' = \{s' \mid \exists s : N. s \xrightarrow{e} s'\}$. Notice that states which can be reached via the same trace are grouped together in $\text{det}(\mathcal{L})$.

Given an implementation $\mathcal{L}_1$ and a specification $\mathcal{L}_2$, the standard trace refinement checking is to construct (often on-the-fly) the product $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$ and then search for a state of the product (s_1, s_2) (where s_1 is a state of $\mathcal{L}_1$ and s_2 is a set of states in $\mathcal{L}_2$) such that s_2 is an empty set. Such a state is called a TR-witness state. In the worst case, this algorithm has a complexity exponential in the number of states of $\mathcal{L}_2$.

2.2 Trace Refinement Checking with Anti-Chain

It has been shown that trace refinement checking based on anti-chain offers significantly better performance [16]. Given two LTSs $\mathcal{L}_1$ and $\mathcal{L}_2$, the anti-chain method explores a ‘simulation’ relation in $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$. Given any two states (s_1, s_2) and (s'_1, s'_2) of $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$, let $(s'_1, s'_2) \leq (s_1, s_2)$ denote $s_1 = s'_1$ and $s_2 \subseteq s'_2$.

Proposition 1. If $(s'_1, s'_2) \leq (s_1, s_2)$ and $(s_1, s_2) \xrightarrow{e} (u, v)$, then there exists (u', v') such that $(s'_1, s'_2) \xrightarrow{e} (u', v')$ and $u' = u$ and $v \subseteq v'$. $\square$

By the above proposition, it can be readily shown that a TR-witness state is reachable from (s'_1, s'_2) implies that a TR-witness state must be reachable from (s_1, s_2) . As a result, if (s_1, s_2) has been explored, we can skip (s'_1, s'_2) .

Formally, an anti-chain is a set A of sets such that $x \not\subseteq y$ and $y \not\subseteq x$ for all $x \in A$ and $y \in A$, i.e., any pair of sets in A are incomparable. An anti-chain supports two operations. One is to check whether it contains a subset of a given set. let x be the given set, we denote $x \in A$ if and only if there exists $y \in A$ such that $y \subseteq x$. The other is to add a given set x in A . $A \uplus x$ is defined as $\{y \mid y \in A \wedge x \not\subseteq y\} \cup \{x\}$, i.e., $A \uplus x$ contains x and all sets in A which is not a superset of x . Obviously, an empty set is an anti-chain by definition.

Algorithm 1 shows the anti-chain based algorithm. In an abuse of notation, we write $(s, X) \in A$ to denote that the set $(\{s\} \cup X) \in A$; and $A \uplus (s, X)$ to denote $A \uplus (\{s\} \cup X)$. The algorithm works as follows. After initializing the stack *working* and the set *antichain*, the algorithm pops one state $(\text{impl}, \text{spec})$ from *working* and adds it to the set *antichain*, and then generates all successors of the state and adds them to *working* unless $(\text{impl}', \text{spec}') \in \text{antichain}$ is true, till the stack *working* is empty or a TR-witness state is found. We remark that *antichain* keeps to be an anti-chain during this algorithm, because line 5 and line 13 guarantee there are no subsets or supersets of the new added state in the updated *antichain*. Soundness of the algorithm can be referred to in [2] [16].

Algorithm 1 Trace Refinement Checking Algorithm with Anti-chain

```

1: let working be a stack containing a pair  $(init_1, \{s \mid init_2 \rightsquigarrow s\})$ ;
2: let antichain  $:= \emptyset$ ;
3: while working  $\neq \emptyset$  do
4:   pop (impl, spec) from working;
5:   antichain  $:= antichain \uplus (impl, spec)$ ;
6:   for all  $(impl, e, impl') \in T_1$  do
7:     if  $e = \tau$  then
8:       spec'  $:= spec$ ;
9:     else
10:      spec'  $:= \{s' \mid \exists s \in spec. s \xrightarrow{e} s'\}$ ;
11:      if spec'  $= \emptyset$  then
12:        return false;
13:      if  $(impl', spec') \in antichain$  is not true then
14:        push (impl', spec') into working;
15: return true;
  
```

Fig. 1. Trace Refinement Checking Algorithm with Anti-Chain

Theorem 1. [16] Algorithm 1 returns true if and only if $traces(\mathcal{L}_1) \subseteq traces(\mathcal{L}_2)$. $\square$

Example 1. Figure 1 shows a simple example of dining philosopher [11] to demonstrate how Algorithm 1 works. The problem is summarized as N philosophers sitting around a round table with a single fork between each pair, and each philosopher requiring both neighboring forks to eat. The LTS $\mathcal{L}_2$ in Figure 1 (b) shows the complete state graph of the system with two philosophers. The invisible events is denoted as $[event]$, i.e., $[get.0.1]$ means that the hidden event is philosopher 0 getting the right fork which is represented as 1. Note that checking $traces(\mathcal{L}_1) \subseteq traces(\mathcal{L}_2)$ is to confirm whether every philosopher can eat. From Figure 1 (c) we can see that the search does not continue from the state $(s_1, \{s'_1, s'_2, s'_3, s'_4, s'_5, s'_6, s'_7, s'_8\})$ because $\{s'_1, s'_2, s'_3, s'_4, s'_5, s'_6\} \subseteq \{s'_1, s'_2, s'_3, s'_4, s'_5, s'_6, s'_7, s'_8\}$. In this case, Algorithm 1 generates 3 states which are labeled with *Anti-Chain*, while the classical algorithm based on subset construction generates 7 states. $\square$

Algorithm 2 Stable Failures Refinement Checking Algorithm with Anti-chain

```
1: let working be a stack containing a pair (init1, {s | init2  $\rightsquigarrow$  s});
2: let antichain :=  $\emptyset$ ;
3: while working  $\neq \emptyset$  do
4:   pop (impl, spec) from working;
5:   antichain := antichain  $\uplus$  (impl, spec);
6:   if refusals(impl)  $\not\subseteq$  refusals(spec) then
7:     return false;
8:   for all (impl, e, impl')  $\in T_1$  do
9:     if e =  $\tau$  then
10:       spec' := spec;
11:     else
12:       spec' := {s' |  $\exists s \in spec. s \xrightarrow{e} s'$ };
13:     if spec' =  $\emptyset$  then
14:       return false;
15:     if (impl', spec')  $\in antichain$  is not true then
16:       push (impl', spec') into working;
17: return true;
```

3 Failures/Divergence Refinement Checking with Anti-Chain

In this section, we demonstrate that anti-chain can be used to improve stable failures refinement checking and failures-divergence refinement checking.

3.1 Stable Failures Refinement Checking

Let $\mathcal{L} = (S, init, Act, T)$ be an LTS. Given a state $s \in S$, s is stable if $\tau \notin enable(s)$. Given a stable state s , the refusals of s , written as *refusals*(s), is defined as $\{X \mid \exists s'. s \rightsquigarrow s' \wedge \tau \notin enable(s') \wedge X \subseteq \Sigma \setminus enable(s')\}$. The failures of $\mathcal{L}$, written as *failures*($\mathcal{L}$), is defined as $\{(tr, X) : \Sigma^* \times 2^\Sigma \mid \exists s. init \xrightarrow{tr} s \wedge X \in refusals(s)\}$ where $init \xrightarrow{tr} s$ denotes that there exists a run $\langle s_0, e_0, s_1, e_1, \dots, e_n, s_{n+1} \rangle$ such that $s_0 = init$ and $s_{n+1} = s$ and $tr = \langle e_0, e_1, \dots, e_n \rangle$.

Definition 4 (Stable Failures Refinement). Let $\mathcal{L}_i$ where $i \in \{1, 2\}$ be two LTSs. $\mathcal{L}_1$ refines $\mathcal{L}_2$ in stable failures semantics if and only if *failures*($\mathcal{L}_1$) $\subseteq$ *failures*($\mathcal{L}_2$).

The existing stable failures refinement checking algorithm [1] works by searching for a state (x, y) of $\mathcal{L}_1 \times det(\mathcal{L}_2)$ such that $y = \emptyset$ or *refusals*(x) $\not\subseteq$ *refusals*(y). Such a state is called a SFR-witness state. In the following, we extend Algorithm 1 for stable failures refinement checking. Given a set states x , we write *refusals*(x) to denote $\{r \mid \exists s \in x. r \in refusals(s)\}$. The algorithm is shown in Algorithm 2.

Lemma 1. For every state (s_1, s_2) of $\mathcal{L}_1 \times det(\mathcal{L}_2)$, for all (s_1, s'_2) in the product, if $s'_2 \subseteq s_2$, then a SFR-witness state is reachable from (s_1, s_2) implies a SFR-witness state is reachable from (s_1, s'_2) .

Fig. 2. Stable Failures Refinement Checking Algorithm with Anti-Chain

Proof By induction. The base case is that (s_1, s_2) is a SFR-witness state, then $s_2 = \emptyset$ or $\text{refusals}(s_1) \not\subseteq \text{refusals}(s_2)$. Because $s_2' \subseteq s_2$ by assumption, $\text{refusals}(s_2') \subseteq \text{refusals}(s_2)$. Then we have $s_2' = \emptyset$, or $\text{refusals}(s_1) \not\subseteq \text{refusals}(s_2')$. Thus, (s_1, s_2') is a SFR-witness state. Next, we prove the induction step. Assume that (s_1, s_2) satisfies the condition, i.e., for any state (s_1, s_2') such that $s_2' \subseteq s_2$, a SFR-witness state is reachable from (s_1, s_2) implies a SFR-witness state is reachable from (s_1, s_2') . Let (x, y) be a state of the product such that $(x, y) \xrightarrow{e} (s_1, s_2)$. For all (x, y') such that $y' \subseteq y$, we can get $(x, y') \xrightarrow{e} (s_1, s_2')$ such that $s_2' \subseteq s_2$. Therefore, the induction step holds by induction hypothesis. Thus, the lemma is true. $\square$

Theorem 2. Algorithm 2 returns true if and only if $\text{failures}(\mathcal{L}_1) \subseteq \text{failures}(\mathcal{L}_2)$.

Proof For a state S of $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$, define $\text{Dist}(S) \in \mathbb{N} \cup \{\infty\}$ as the length of the shortest SFR-witness trace from S (if a SFR-witness state is not reachable from S , $\text{Dist}(S) = \infty$). For a set of states States , if $\text{States} = \emptyset$, $\text{Dist}(\text{States}) = \infty$, otherwise, $\text{Dist}(\text{States}) = \min_{S \in \text{States}} \text{Dist}(S)$. The predicate $\text{SFR}(\text{States})$ is true if and only if all the states in States are not SFR-witness states. Then the correctness of Algorithm 2 can be proved using the two invariants below. The invariants can be proved in a very similar way to [2].

1. $\neg \text{SFR}(\text{antichain} \cup \text{working}) \Rightarrow \neg \text{SFR}(\{(i, \{s \mid \text{init}_2 \rightsquigarrow s\}) \mid i \in \text{init}_1\})$.
2. $\neg \text{SFR}(\{(i, \{s \mid \text{init}_2 \rightsquigarrow s\}) \mid i \in \text{init}_1\}) \Rightarrow \text{Dist}(\text{antichain}) > \text{Dist}(\text{working})$.

Because the number of state is finite and all states are only visited once, Algorithm 2 eventually terminates. Algorithm 2 returns false only if the state spec' is an empty set on line 13, or $(\text{impl}, \text{spec})$ satisfies the condition $\text{refusals}(\text{impl}) \not\subseteq \text{refusals}(\text{spec})$ on line 6. The former case has been proved in Algorithm 1. In the latter case, $(\text{impl}, \text{spec})$ is a SFR-witness state, and hence $\text{SFR}(\text{antichain} \cup \text{working})$ is false. By invariant 1, $\mathcal{L}_1$ cannot refine $\mathcal{L}_2$ in stable failures semantics. Algorithm 2 returns true only when working is empty, which implies that $\text{Dist}(\text{antichain}) > \text{Dist}(\text{working})$ is not true. By invariant 2, $\mathcal{L}_1$ refines $\mathcal{L}_2$ in stable failures semantics. $\square$

Example 2. If Algorithm 2 is applied to the example presented in Figure 1, the reduction remains the same as for trace refinement checking (from 7 states to 3 states). We show another example with some counterexamples, as shown in Figure 2. Notice that

the refusal set of s_2 is $\{b\}$, and s_1, s'_1, s'_2 do not refuse any event. Then the two states with red circles are SFR-witness states. Since $\{s'_1\} \subseteq \{s'_1, s'_2\}$, the search does not continue from the state $(s_1, \{s'_1, s'_2\})$. We can see that a SFR-witness state which is reachable from $(s_1, \{s'_1, s'_2\})$ is also reachable from $(s_1, \{s'_1\})$ after the pruning of states. In this case, Algorithm 2 generates 3 states which are labelled with *Anti-Chain*, while the classical algorithm may generate 4 states before finding a SFR-witness state. Moreover, Algorithm 2 may find a shorter witness trace than the classical algorithm. $\square$

3.2 Failures-Divergence Refinement Checking

In the following, we show how to adopt anti-chain for failures-divergence refinement checking. Let $\mathcal{L} = (S, \text{init}, \text{Act}, T)$ be an LTS. Given a state s , s diverges if and only if s can perform an infinite number of τ -transitions. A trace tr is divergent, written as $\text{div}(tr)$, if and only if there exists a prefix pre of tr or tr itself such that $\text{init} \xrightarrow{pre} s$ and s diverges. We write $\text{divergences}(\mathcal{L})$ to be $\{tr \mid \text{div}(tr)\}$.

Definition 5 (Failures-Divergence Refinement). Let $\mathcal{L}_i$ where $i \in \{1, 2\}$ be two LTSs. $\mathcal{L}_1$ refines $\mathcal{L}_2$ in failures-divergence semantics if and only if $\text{divergences}(\mathcal{L}_1) \subseteq \text{divergences}(\mathcal{L}_2)$ and $\text{failures}(\mathcal{L}_1) \subseteq \text{failures}(\mathcal{L}_2)$.

In the following, we extend Algorithm 2 for failures-divergence refinement checking. The algorithm is shown in Algorithm 3. Given a set of state x , we say that x diverges if there exists $s \in x$ such that s diverges. Like in the existing failures-divergence refinement checking algorithm [1], the idea is to search for a FDR-witness state (x, y) of $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$ such that $y = \emptyset$ or $\text{refusals}(x) \not\subseteq \text{refusals}(y)$ or x diverges but not y .

Lemma 2. For every state (s_1, s_2) of $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$, for all (s_1, s'_2) in the product, if $s'_2 \subseteq s_2$, then a FDR-witness state is reachable from (s_1, s_2) implies a FDR-witness state is reachable from (s_1, s'_2) .

Proof By induction. The base case is that (s_1, s_2) is a FDR-witness state, then $s_2 = \emptyset$ or $\text{refusals}(s_1) \not\subseteq \text{refusals}(s_2)$ or s_1 diverges and s_2 does not. Because $s'_2 \subseteq s_2$ by assumption, $\text{refusals}(s'_2) \subseteq \text{refusals}(s_2)$ and if s'_2 diverges, so does s_2 . Thus, (s_1, s'_2) is a SFR-witness state since $s_2 = \emptyset$ implies $s'_2 = \emptyset$; $\text{refusals}(s_1) \not\subseteq \text{refusals}(s_2)$ implies $\text{refusals}(s_1) \not\subseteq \text{refusals}(s'_2)$; and s_2 not divergent implies that s'_2 does not diverge either. Next, we prove the induction step. Assume that (s_1, s_2) satisfies the condition. Let (x, y) be a state of the product such that $(x, y) \xrightarrow{e} (s_1, s_2)$. For all (x, y') such that $y' \subseteq y$, we can get $(x, y') \xrightarrow{e} (s_1, s'_2)$ such that $s'_2 \subseteq s_2$. Therefore, the induction step holds by induction hypothesis. Thus, the lemma is true. $\square$

Theorem 3. Algorithm 3 returns true if and only if $\text{divergences}(\mathcal{L}_1) \subseteq \text{divergences}(\mathcal{L}_2)$ and $\text{failures}(\mathcal{L}_1) \subseteq \text{failures}(\mathcal{L}_2)$.

Proof Define $\text{Dist}(S) \in N \cup \{\infty\}$ as the length of the shortest FDR-witness trace from a state S of $\mathcal{L}_1 \times \text{det}(\mathcal{L}_2)$ (if a FDR-witness state is not reachable from S , $\text{Dist}(S) = \infty$). Given a set of states States , if $\text{States} = \emptyset$, $\text{Dist}(\text{States}) = \infty$, otherwise, $\text{Dist}(\text{States}) = \min_{S \in \text{States}} \text{Dist}(S)$. The predicate $\text{FDR}(\text{States})$ is true if and only if all the states in States are not FDR-witness states. The correctness of Algorithm 3 can be proved similarly as for Algorithm 2, using the following two invariants.

Algorithm 3 Failures-Divergence Refinement Checking Algorithm with Anti-chain

```
1: let working be a stack containing a pair  $(init_1, \{s \mid init_2 \rightsquigarrow s\})$ ;  
2: let antichain :=  $\emptyset$ ;  
3: while working  $\neq \emptyset$  do  
4:   pop (impl, spec) from working;  
5:   antichain := antichain  $\sqcup$  (impl, spec);  
6:   if impl diverges then  
7:     if spec does not diverge then  
8:       return false;  
9:   else  
10:    if  $refusals(impl) \not\subseteq refusals(spec)$  then  
11:      return false;  
12:    for all (impl, e, impl')  $\in T_1$  do  
13:      if e =  $\tau$  then  
14:        spec' := spec;  
15:      else  
16:        spec' :=  $\{s' \mid \exists s \in spec. s \xrightarrow{e} s'\}$ ;  
17:      if spec' =  $\emptyset$  then  
18:        return false;  
19:      if (impl', spec')  $\in antichain$  is not true then  
20:        push (impl', spec') into working;  
21: return true;
```

1. $\neg FDR(antichain \cup working) \Rightarrow \neg FDR(\{(i, \{s \mid init_2 \rightsquigarrow s\}) \mid i \in init_1\})$.
2. $\neg FDR(\{(i, \{s \mid init_2 \rightsquigarrow s\}) \mid i \in init_1\}) \Rightarrow Dist(antichain) > Dist(working)$.

□

Example 3. We use the example shown in Figure 3 to demonstrate how algorithm 3 works. The state s_2 in LTS $\mathcal{L}_1$ has a self-loop labeled with τ . The two states with red circles are FDR-witness states now. Since $\{s'_1\} \subseteq \{s'_1, s'_2\}$, the search does not continue from the state $(s_1, \{s'_1, s'_2\})$. We can see that a FDR-witness state which is reachable from $(s_1, \{s'_1, s'_2\})$ is also reachable from $(s_1, \{s'_1\})$ after pruning the states. □

3.3 Implementation and Evaluation

The proposed algorithms have been adopted in the Process Analysis Toolkit (PAT) [13]. PAT is designed for systematic validation of distributed/concurrent systems using state-of-the-art model checking techniques. In the following, we evaluate the performance of the algorithms using a range of real-life parameterized systems. All the systems are embedded in the PAT package and available online. The data is obtained with Intel(R) Core(TM) i7-2640M CPU at 2.80GHz and 8GB RAM.

The pairs of LTSs (one as implementation and one as specification) are generated from different systems or same systems with different parameters. The systems include a multi-valued register simulation system with one reader or multiple readers [4], an implementation of concurrent stack with or without linearization point [15], a mailbox system [3], a system of scalable nonzero indicator [9] and the dining philosopher

Fig. 3. Failures-Divergence Refinement Checking Algorithm with Anti-Chain

Fig. 4. Refinement Checking Results of Concurrent Stack Implementation

problem [11]. In total about 300 pairs of LTS were generated to compare the anti-chain algorithms and the classical ones for all three kinds of refinement checking.

Figure 4 shows the statistics of a typical example, i.e., the concurrent stack, from which we can see significant performance improvement. In the figure, the horizontal axis is the sum of the sizes of the two LTSs for refinement checking, and the vertical axis is the execution time (in seconds) of the corresponding algorithm. Each point shows the checking time for a pair of LTSs. We can see that for all three kinds of refinement checking, anti-chain based algorithms offer significantly better performance. The complete experimental results, with the refinement checking assertions always being valid, are summarized in Table 1. Notice that if the pair of LTSs are equivalent in terms of traces or failures or failures-divergence, we can perform the refinement checking in both directions. This is shown in the table using two columns $\subseteq$ and $\supseteq$. A few cases are marked as ‘—’ as the result is false, which are discussed later. Furthermore, ‘unknown’ means either out of memory or running for more than 30 minutes. It can be observed that the speedup differs for different systems. In most cases, the anti-chain approach has a much better performance than the classical one, e.g., in the concurrent stack linearization point implementation, it is 30.28 times faster for stable failures refinement checking and 12.16 times faster for failures-divergence refinement checking. Moreover, the larger the system is, the larger the speedup is. In some cases, anti-chain can not reduce the number of states at all simply because the specifications are deterministic. In some cases (e.g., SNZI), although anti-chain reduces the number of states, the benefit is not significant enough to overcome the computational overhead of the anti-chain operations defined in section 2.2.

System	Size	Trace (Speedup)		Stable Failures (Speedup)		Failures-Divergence (Speedup)	
		$\subseteq$	$\supseteq$	$\subseteq$	$\supseteq$	$\subseteq$	$\supseteq$
Multi-valued Register Simulation with 1 Reader and 1 Writer	0-10000	2.21	1.42	2.32	1.63	3.48	1.46
	10000-100000	4.54	2.14	4.45	2.09	6.61	1.73
	100000-700000	6.74	2.88	6.64	2.92	unknown	2.59
Multi-valued Register Simulation with Multiple Readers	0-10000	2.17	1.49	1.99	1.45	3.33	1.43
	10000-100000	3.32	2.05	3.32	2.11	3.55	1.70
	100000-700000	6.45	2.68	6.14	2.72	unknown	3.16
Concurrent Stack Implementation	0-10000	1.48	1.85	1.63	1.72	2.26	1.60
	10000-100000	1.70	3.72	1.71	3.71	2.71	3.22
	100000-200000	1.62	5.90	1.56	6.44	2.85	4.96
Concurrent Stack Linearization Point Implementation	0-10000	0.75	5.33	—	5.99	—	2.70
	10000-30000	0.84	13.94	—	14.11	—	5.00
	30000-60000	0.91	30.37	—	30.28	—	12.16
Mailbox	0-700000	1.13	1.54	1.11	1.61	1.39	1.01
SNZI	0-50000	0.89	2.45	0.93	2.42	1.03	1.14
Dining Philosopher	0-50000	0.99	1.14	—	1.02	—	1.17

Table 1. Testing on Refinement Checking Assertions which are valid

System	Size	Trace With AC(s)	Trace W/o AC(s)	Stable Failures With AC(s)	Stable Failures W/o AC(s)	Failures-Divergence With AC(s)	Failures-Divergence W/o AC(s)
Multi-valued Register Simulation with Multiple Readers	3175	0.10	0.42	0.09	0.43	0.44	2.88
	24655	1.52	12.12	1.71	12.02	9.78	146.95
	117288	3.92	136.37	3.57	138.50	48.52	985.24
	194455	22.11	294.71	21.51	299.21	180.91	unknown

Table 2. Testing on Refinement Checking Assertions which are invalid

In presence of counterexamples, anti-chain based algorithms may find the counterexample more quickly. For instance, the verification results with the register example (with multiple readers), shown in Table 2 (‘Anti-Chain’ and ‘Without’ are denoted as AC and W/o for short), evidences that anti-chain finds the counterexample more quickly in all three kinds of refinement checking. Nonetheless, we remark that because the algorithms are on-the-fly, whether the counterexample is found earlier depends on the searching order and therefore sometime anti-chain based algorithms may be slower if a ‘wrong’ order is taken.

4 Probabilistic Refinement Checking with Anti-Chain

In this section, we show that anti-chain can be used to improve a particular kind of probabilistic refinement checking, i.e., the implementation is given as an MDP and the specification is given as an NFA.

4.1 MDP and Probabilistic Model Checking

Given a set of states S , a distribution is a function $\mu : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \mu(s) = 1$. Let $Distr(S)$ be the set of all distributions over S . A Markov Chain is a tuple $\mathcal{M} = (S, init, Act, Pr)$ where S is a countable set of states; $init \in S$ is an initial state; Act is a set of events; and $Pr : S \times Act \times S \rightarrow [0, 1]$ is a labeled transition probability function such that for all state $s \in S$, $\sum_{s' \in S} Pr(s, e, s') = 1$. Notice that Markov Chains are deterministic as there is only one event (and one distribution) at each state.

A sequence of alternating states and events $\pi = \langle s_0, e_0, s_1, e_1, \dots, e_n, s_{n+1} \rangle$ is a path of $\mathcal{M}$ if $Pr(s_i, e_i, s_{i+1}) > 0$ for all i . The probability of executing π from s_0 , written as $Pr(\mathcal{M}, \pi)$, is $Pr(s_0, e_0, s_1) \times Pr(s_1, e_1, s_2) \times \dots \times Pr(s_n, e_n, s_{n+1})$. It is often also interesting to find out the probability of reaching a certain set of states (e.g., what is the probability of reaching the state of system failure?). Given a set of target states G , the probability of reaching any state in G from a starting state s_0 , written as $Pr(\mathcal{M}, s_0, G)$, is the accumulated probability of all paths from s_0 to any state in G , which can be calculated systematically [5]. Given a path π , we define $trace(\pi)$ to be the sequence of visible events in π . We write $Pr(\mathcal{M}, s_0, tr)$ to denote the probability of exhibiting a trace tr from state s_0 , which is the accumulated probability of all paths π from s_0 such that $trace(\pi) = tr$. Given a set of traces Tr , the probability of $\mathcal{M}$ exhibiting any trace in Tr from state s_0 is the accumulated probability $\sum_{tr \in Tr} Pr(\mathcal{M}, s_0, tr)$.

Different from Markov Chains, an MDP can express both probabilistic choices and non-determinism. An MDP is a tuple $\mathcal{D} = (S, init, Act, Pr)$ where S is a set of system states; $init \in S$ is the initial system configuration⁵; Act is a set of actions; $Pr : S \times Act \times Distr(S)$ is a transition probability function such that for all states $s \in S$ and $a \in Act$: $\sum_{s' \in S} Pr(s, a, s') \in \{0, 1\}$. Notice that there could be multiple events at any state. A transition of the system is written as $s \xrightarrow{e} \mu$ where μ is a distribution. A path of $\mathcal{M}$ is a sequence of alternating states, events and distributions $\pi = \langle s_0, e_0, \mu_0, s_1, e_1, \mu_1, \dots \rangle$ such that $s_0 = init$ and $s_i \xrightarrow{e_i} \mu_i$ and $\mu_i(s_{i+1}) > 0$ for all i . Given a path π , we define $trace(\pi)$ to be the sequence of visible events in π .

Intuitively speaking, given a system configuration, firstly an event and a distribution is selected non-deterministically by the *scheduler*, and then one of successor states is reached according to the probability distribution. A scheduler is a function deciding which event and distribution to choose based on the execution history. With a scheduler δ , we effectively obtain a Markov Chain from $\mathcal{D}$, written as $\mathcal{D}_\delta$. Note that with different scheduling, the probability of reaching a state or exhibiting a trace may be different. The measurement of interest is thus the maximum and minimum probability. Given a set of target states G and an MDP $\mathcal{D}$, the maximum probability of reaching any state in G from state s_0 is defined as $\mathcal{P}_{max}(\mathcal{D}, s_0, G) = \sup_\delta Pr(\mathcal{D}_\delta, s_0, G)$. Note that the supremum ranges over all, potentially infinitely many, schedulers. Accordingly, the minimum is written as $\mathcal{P}_{min}(\mathcal{D}, s_0, G)$. Similarly, we define the maximum probability of exhibiting a trace in a set Tr by $\mathcal{D}$ from s_0 .

$$\mathcal{P}_{max}(\mathcal{D}, s_0, Tr) = \sup_\delta (\sum_{tr \in Tr} Pr(\mathcal{D}_\delta, s_0, tr))$$

Accordingly, the minimum is written as $\mathcal{P}_{min}(\mathcal{D}_\delta, s_0, Tr)$.

Example 4. The following shows a simple example MDP.

⁵ This is a simplified definition. In general, there can be an initial distribution.

where s_0 is the initial state and s_2 is a target state. For simplicity, we omit the self-loop of s_2 and s_3 . s_1 has two distributions, following two actions α and β . If s_1 non-deterministically chooses α , then it has equal probability to transfer to s_2 or stay in s_1 ; and if β is chosen, it will transfer to s_3 with probability 1. $\square$

Definition 6 (Refinement Probability). Let $\mathcal{D} = (S, \text{init}, \text{Act}, \text{Pr})$ be an MDP; $\mathcal{L}$ be an LTS. The maximum probability of $\mathcal{D}$ trace-refining $\mathcal{L}$ is $\mathcal{P}_{\max}(\mathcal{D}, \text{init}, \text{traces}(\mathcal{L}))$. The minimum is $\mathcal{P}_{\min}(\mathcal{D}, \text{init}, \text{traces}(\mathcal{L}))$.

Intuitively, the probability of $\mathcal{D}$ refines $\mathcal{L}$ is the probability of $\mathcal{D}$ exhibiting a trace of $\mathcal{L}$. As we mention earlier, the probability may vary due to different scheduling.

Definition 7 (MDP and LTS Synchronous Product). Let $\mathcal{D} = (S_d, \text{init}_d, \text{Act}_d, \text{Pr}_d)$ be an MDP; $\mathcal{L} = (S_l, \text{init}_l, \text{Act}_l, T)$ be an LTS such that $\tau \notin \text{Act}_l$. The synchronous product of $\mathcal{D}$ and $\mathcal{L}$, written as $\mathcal{D} \times \mathcal{L}$, is an MDP $(S, \text{init}, \text{Act}, \text{Pr})$ such that $S = S_d \times S_l$; $\text{init} = (\text{init}_d, \text{init}_l)$; $\text{Act} = \text{Act}_d \cup \text{Act}_l$; and Pr is defined as follows.

- If $(s_1, \tau, \mu) \in \text{Pr}_d$, $((s_1, s_2), \tau, \mu') \in \text{Pr}$ for all $s_2 \in S_l$ such that for all $s'_1 \in S_d$, $\mu'((s'_1, s_2)) = \mu(s'_1)$;
- If $(s_1, e, \mu) \in \text{Pr}_d$ and $(s_2, e, s'_2) \in T$, $((s_1, s_2), e, \mu') \in \text{Pr}$ such that for all $s'_1 \in S_d$, $\mu'((s'_1, s'_2)) = \mu(s'_1)$.

A state (s_1, s_2) of the product $\mathcal{D} \times \mathcal{L}$ is a TR-witness state if and only if $s_2 = \emptyset$. In [14], we show that the refinement probability can be calculated systematically by (1) building the deterministic LTS $\text{det}(\mathcal{L})$; (2) computing the synchronous product of $\mathcal{D}$ and $\text{det}(\mathcal{L})$; (3) calculating the maximum/minimum probability of reaching any TR-witness state.

Theorem 4. Let $\mathcal{D} = (S_d, \text{init}_d, \text{Act}_d, \text{Pr}_d)$ be an MDP; $\mathcal{L} = (S_l, \text{init}_l, \text{Act}_l, T)$. Let G be the set of TR-witness states of $\mathcal{D} \times \text{det}(\mathcal{L})$.

- $\mathcal{P}_{\max}(\mathcal{D}, \text{init}_d, \text{traces}(\mathcal{L})) = \mathcal{P}_{\max}(\mathcal{D} \times \text{det}(\mathcal{L}), (\text{init}_d, \text{init}_l), G)$
- $\mathcal{P}_{\min}(\mathcal{D}, \text{traces}(\mathcal{L})) = \mathcal{P}_{\min}(\mathcal{D} \times \text{det}(\mathcal{L}), (\text{init}_d, \text{init}_l), G)$ $\square$

Based on the above theorem, the probabilistic refinement checking problem is reduced to a probabilistic reachability problem, which can be solved by two standard methods. One is by solving a linear program. That is, we firstly associate a variable x_s to each state s in $\mathcal{P}$ to represent the probability of reaching any target state from s ; then we construct a linear program which constraints the value of every x_s using a set of linear inequalities, based on the probability transition function; and lastly, we solve the linear program to get the maximum/minimum value of each x_s . Notice that the solution of state $(\text{init}_d, \text{init}_l)$ is the refinement probability. The other is to iteratively approximate the probability through graph traversing. Notice that for systems having large state space, it is impractical to store the entire linear program and solve it directly, therefore the iterative calculation approach is more widely used in probabilistic verification. As a result, in this paper we just focus on this approach.

Example 5. In the following, we show how the iterative calculation method works using the simple example shown in Example 4. That is, the maximal probability from initial state s_0 to accepting state s_2 is calculated step by step. Assume p_i^k is the maximal

probability of s_i after the k -th iteration. Starting from the target state s_2 , in k -th iteration we update the probability of states which could reach s_2 in exact k steps. Obviously, $p_0^0 = p_1^0 = 0$. As $p_2^k = 1$ and $p_3^k = 0$ for any k , k is ignored in these two states. In the 1st iteration, only p_1 can be updated, and $p_1^1 = \max\{0.5 \times p_1^0 + 0.5 \times p_2, 1 \times p_3\} = \max\{0.5, 0\} = 0.5$; in the 2nd iteration, both p_0 and p_1 can be updated. It is trivial to show $p_0^2 = p_1^1 = 0.5$, and $p_1^2 = \max\{0.5 \times p_1^1 + 0.5 \times p_2, 1 \times p_3\} = \max\{0.75, 0\} = 0.75$. Iteratively, p_0 and p_1 in the long run can be calculated. A user-defined threshold is usually necessary to terminate the calculation, according to the desired precision. $\square$

4.2 Anti-chain Based Approach

Now we introduce how anti-chain can be used to speed up the iterative calculation approach, by first introducing a lemma.

Lemma 3. Let $\mathcal{D} = (S_d, init_d, Act_d, Pr_d)$ be an MDP; $\mathcal{L} = (S_l, init_l, Act_l, T)$. Let $\mathcal{P}$ be $\mathcal{D} \times det(\mathcal{L})$. Let G be the set of TR-witness states of $\mathcal{P}$. For all state (u_1, v_1) and (u_2, v_2) of $\mathcal{P}$ s.t. $(u_2, v_2) \leq (u_1, v_1)$, $Pr_{max}(\mathcal{P}, (u_1, v_1), G) \geq Pr_{max}(\mathcal{P}, (u_2, v_2), G)$ and $Pr_{min}(\mathcal{P}, (u_1, v_1), G) \geq Pr_{min}(\mathcal{P}, (u_2, v_2), G)$.

Proof The above can be proved with an induction. The base case is that (u_2, v_2) is in G . By definition, (u_1, v_1) must be in G and therefore the lemma holds. Next, we show the induction step. Assume that (u'_2, v'_2) satisfies the lemma above. For every distribution μ_2 from (u_2, v_2) , by Definition 7, there must exist a distribution μ_1 from (u_1, v_1) and for every state (u'_2, v'_2) , there exists (u'_1, v'_1) such that $\mu_2((u'_2, v'_2)) = \mu_1((u'_1, v'_1))$ and $(u'_2, v'_2) \leq (u'_1, v'_1)$. By induction hypothesis, we have $Pr_{max}(\mathcal{P}, (u'_1, v'_1), G) \geq Pr_{max}(\mathcal{P}, (u'_2, v'_2), G)$ and $Pr_{min}(\mathcal{P}, (u'_1, v'_1), G) \geq Pr_{min}(\mathcal{P}, (u'_2, v'_2), G)$. Thus we have $Pr_{max}(\mathcal{P}, (u_1, v_1), G) \geq Pr_{max}(\mathcal{P}, (u_2, v_2), G)$ and $Pr_{min}(\mathcal{P}, (u_1, v_1), G) \geq Pr_{min}(\mathcal{P}, (u_2, v_2), G)$. Therefore, we conclude that the lemma holds. $\square$

Compared to probabilistic reachability calculation for a general MDP, the above lemma gives us additional information, which can be potentially useful in speeding up the calculation. In the following, we discuss how we can make use of the information so as to improve the probabilistic refinement checking using the iterative calculation method.

The first step is building the product MDP meanwhile finding the target states, which is shown in Algorithm 4. The implementation and specification are defined in Definition 7. Different from the non-probabilistic cases, *the state space cannot be reduced in the probabilistic models*; instead, we define a function *sub* of the product state s satisfying $s.sub = \{t \mid t \in S \wedge s \leq t\}$, where S is the state space of the product MDP. Then the refinement checking is reduced to probabilistic reachability of a set of target states, denoted by *Target*. During the iterative calculation, whenever the probability of state s is updated, e.g., to p , according to lemma 3, **all states in $s.sub$ whose probability is less than p could be set to p directly**. This could speed up each iteration and potentially improve probabilistic refinement checking.

Now we evaluate whether the above method is indeed beneficial. The proposed probabilistic refinement checking algorithm has also been implemented in PAT. We evaluate it using a modified system based on the implementation of a distributed concurrent stack example [15]. Probabilistic choices are used to model a concurrent stack

Algorithm 4 Building MDP in Probabilistic Refinement Checking with Anti-chain

```
1: let working be a stack containing a pair  $(init_d, \{s \mid init_l \rightsquigarrow s\})$ ;
2: let visited :=  $\{(init_d, \{s \mid init_l \rightsquigarrow s\})\}$ ; let Target =  $\emptyset$ ; ;
3: while working  $\neq \emptyset$  do
4:   pop (impl, spec) from working;
5:   for all (impl, e,  $\mu$ )  $\in Pr_d$  do
6:     if e =  $\tau$  then
7:       spec' := spec;
8:     else
9:       spec' :=  $\{s' \mid \exists s \in spec. s \xrightarrow{e} s'\}$ ;
10:    for all impl'  $\in S_d$  do
11:      if  $\mu(impl') > 0 \wedge (impl', spec') \notin visited$  then
12:        push (impl', spec') into working;
13:        visited := visited  $\cup$  (impl', spec');
14:        if spec' =  $\emptyset$  then
15:          Target := Target  $\cup$  (impl', spec');
16:        for all (impl', spec'')  $\in visited$  do
17:          if  $(impl', spec'') \leq (impl', spec')$  then
18:            (impl', spec'').sub.Add(impl', spec');
19:          else if  $(impl', spec') \leq (impl', spec'')$  then
20:            (impl', spec').sub.Add(impl', spec'');
21: return true;
```

System	Size	Verification Time (s)			#States Involved in Iterations		
		W/o AC	With AC	Gain	W/o AC	With AC	Gain
K = 2	20600	2.74	2.21	19.3%	4.2M	3M	28.6%
K = 3	45584	15.98	12.04	24.6%	18.6M	11.7M	37.1%
K = 4	86704	48.72	37.50	22.6%	55.5M	36.2M	34.8%
K = 5	117408	123.9	80.83	34.9%	130.7M	76.3M	41.6%
K = 6	231440	271.2	182.6	32.7%	272.1M	160.7M	40.9%
K = 7	342544	511.1	340.3	33.5%	515.2M	298.8M	42.0%

Table 3. Experiments: Probabilistic Concurrent Stack Implementation

model composed by *two processes*, so as to capture the situation in which the communication between different processes fails from time to time. Failures do exist in real world cases and the experiments results are summarized in Table 3.

We compare the efficiency of the implementation with and without (W/o) Anti-chain (AC) using several cases. *K* means length of the stack; *Size* indicates the number of states in the whole system; *#States Involved in Iterations* represents the total number of states involved in the iterative calculation. For example, a state *s* updates its probability in two iterations, then *#States* should increase two. From the experiments, we can see that the anti-chain approach could reduce the total number of states accumulated during the calculation, through dynamically updating states' probability based on the subset relation *sub*. This speeds up the verification around 29%. We remark that the gains here are not as significant as the non-probabilistic cases, because the state space cannot be reduced; however, in some cases, it does shorten the verification time.

5 Conclusion

In this work, we proposed to adopt anti-chain approach to improve stable failures refinement, failures-divergence refinement and probabilistic refinement checking. These algorithms have been implemented in model checking framework PAT, and some experiments based on benchmark systems demonstrated the dramatic improvement of the verification efficiency of our method. To our best of knowledge, we are the first to investigate anti-chain approaches for these refinement checking.

As for future work, we are trying to extend anti-chain based refinement checking approach in real-time system; meanwhile, we are exploring the refinement relation between probabilistic models, which may also benefit from anti-chain based method.

References

1. A. W. Roscoe. *Model-checking CSP*, chapter 21. Prentice-Hall, 1994.
2. P. A. Abdulla, Y.-F. Chen, L. Holk, R. Mayr, and T. Vojnar. When simulation meets antichains. In *TACAS*, volume 6015 of *Lecture Notes in Computer Science*, pages 158–174. Springer, 2010.
3. M. Aguilera, E. Gafni, and L. Lamport. The mailbox problem. *Distributed Computing*, 2008.
4. H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics, 2nd Edition*. The Oxford University Press, 2004.
5. C. Baier and J. Katoen. *Principles of Model Checking*. The MIT Press, 2008.
6. A. Bouajjani, P. Habermehl, L. Holk, T. Touili, and T. Vojnar. Antichain-based universality and inclusion testing over nondeterministic finite tree automata. In *CIAA*, volume 5148 of *Lecture Notes in Computer Science*, pages 57–67. Springer, 2008.
7. K. Chatterjee, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. In *CSL*, volume 4207 of *Lecture Notes in Computer Science*, pages 287–302. Springer, 2006.
8. L. Doyen and J. F. Raskin. Antichains for the automata-based approach to model checking. *Logical Methods in Computer Science*, 5(1:5):1–20, 2009.
9. F. Ellen, Y. Lev, V. Luchangco, and M. Moir. SNZI: Scalable nonzero indicators. In *PODC*, pages 13–22. ACM, 2007.
10. E. Filiot, N. Jin, and J.-F. Raskin. An antichain algorithm for ltl realizability. In *CAV*, volume 5643 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 2009.
11. C. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
12. A. W. Roscoe. On the expressive power of CSP refinement. *Formal Aspects of Computing*, 17(2):93–112, 2005.
13. J. Sun, Y. Liu, J. S. Dong, and J. Pang. PAT: Towards flexible verification under fairness. In *CAV*, volume 5643 of *Lecture Notes in Computer Science*, pages 709–714. Springer, 2009.
14. J. Sun, S. Song, and Y. Liu. Model checking hierarchical probabilistic systems. In *ICFEM*, volume 6447 of *Lecture Notes in Computer Science*, pages 388–403. Springer, 2010.
15. R. K. Treiber. Systems programming: Coping with parallelism. Technical report, IBM Almaden Research Center, 1986.
16. M. D. Wulf, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Antichains: A new algorithm for checking universality of finite automata. In *CAV*, volume 4144 of *Lecture Notes in Computer Science*, pages 17–30. Springer, 2006.
17. M. D. Wulf, L. Doyen, N. Maquet, and J.-F. Raskin. Antichains: Alternative algorithms for ltl satisfiability and model-checking. In *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 63–77. Springer, 2008.