

Extrapolation-based Path Invariants for Abstraction Refinement of Fifo Systems

Alexander Heußner¹, Tristan Le Gall², and Grégoire Sutre¹

¹ LaBRI, Université Bordeaux, CNRS {heussner, sutre}@labri.fr

² Université Libre de Bruxelles (ULB) tlegall@ulb.ac.be

Abstract. The technique of counterexample-guided abstraction refinement (CEGAR) has been successfully applied in the areas of software and hardware verification. Automatic abstraction refinement is also desirable for the safety verification of complex infinite-state models. This paper investigates CEGAR in the context of formal models of network protocols, in our case, the verification of fifo systems. Our main contribution is the introduction of *extrapolation-based path invariants* for abstraction refinement. We develop a range of algorithms that are based on this novel theoretical notion, and which are parametrized by different extrapolation operators. These are utilized as subroutines in the refinement step of our CEGAR semi-algorithm that is based on recognizable partition abstractions. We give sufficient conditions for the termination of CEGAR by constraining the extrapolation operator. Our empirical evaluation confirms the benefit of extrapolation-based path invariants.

1 Introduction

Distributed processes that communicate over a network of reliable and unbounded fifo channels are an important model for the automatic verification of client-server architectures and network protocols. We focus on communicating fifo systems that consist of a set of finite automata that model the processes, and a set of reliable, unbounded fifo queues that model the communication channels. This class of infinite-state systems is, unfortunately, Turing-complete even in the case of one fifo queue [BZ83]. In general, two approaches for the automatic verification of Turing-complete infinite-state models have been considered in the literature: (a) exact semi-algorithms that compute forward or backward reachability sets (e.g., [BG99, BH99, FIS03] for fifo systems) but may not terminate, and (b) algorithms that always terminate but only compute an over-approximation of these reachability sets (e.g., [LGJJ06, YBCI08] for fifo systems).

CEGAR. In the last decade, counterexample-guided abstraction refinement [CGJ⁺03] has emerged as a powerful technique that bridges the gap between these two approaches. CEGAR plays a prominent role in the automatic, iterative approximation and refinement of abstractions and has been applied successfully in the areas of software [BR01, HJMS02] and hardware verification [CGJ⁺03]. Briefly, the CEGAR approach to the verification of a safety property utilizes an

abstract-check-refine loop that searches for a counterexample in a conservative over-approximation of the original model, and, in the case of finding a false negative, refines the over-approximation to eliminate this spurious counterexample.

Our Contribution. We present a CEGAR semi-algorithm for safety verification of fifo systems based on finite partition abstractions where equivalence classes are recognizable languages of queue contents, or, equivalently, QDDs [BG99]. The crucial part in CEGAR-based verification is refinement, which must find a new partition that is both (1) precise enough to rule out the spurious counterexample and (2) computationally “simple”. In most techniques, refinement is based on the generation of *path invariants*; these are invariants along the spurious counterexample that prove its unfeasibility (in our case, given by a series of recognizable languages). We follow this approach, and present several generic algorithms to obtain path invariants based on parametrized extrapolation operators for queue contents. Our path invariant generation procedures are fully generic with respect to the extrapolation. Refining the partition consists in splitting abstract states that occur on the counterexample with the generated path invariant.

We formally present the resulting CEGAR semi-algorithm and give partial termination results that, in contrast to the classical CEGAR literature, do not rely on an “a priori finiteness condition” on the set of all possible abstractions. Actually, our results depend mainly on our generic extrapolation-based path invariant generation. In particular we show that our semi-algorithm always terminates if (at least) one of these two conditions is satisfied: (1) the fifo system under verification is unsafe, or (2) it has a finite reachability set and the parametrized extrapolation has a finite image for each value of the parameter.

We have implemented our approach in the tool McscM [Mcs] that performs CEGAR-based safety verification of fifo systems. Experimental results on a suite of (small to medium size) network protocols allow for a first discussion of our approach’s advantages.

Related Work. Exact semi-algorithms for reachability set computations of fifo systems usually apply *acceleration techniques* [BG99, BH99, FIS03] that, intuitively, compute the effect of iterating a given “control flow” loop. The tools LASH [Las] (for counter/fifo systems) and TRex [Tre] (for lossy fifo systems) implement these techniques. However, recognizable languages equipped with Presburger formulas (CQDDs [BH99]) are required to represent (and compute) the effect of *counting* loops [BG99, FIS03]. Moreover such tools may only terminate when the fifo system can be flattened into an equivalent system without nested loops. Our experiments show that our approach can cope with both counting loops and nested loops that cannot be flattened.

The closest approach to ours is *abstract regular model checking* [BHV04], an extension of the generic regular model-checking framework based on the abstract-check-refine paradigm. As in classical regular model-checking, a system is modeled as follows: configurations are words over a finite alphabet and the transition relation is given by a finite-state transducer. The analysis consists in an over-approximated forward exploration (by Kleene iteration), followed, in case of a non-empty intersection with the bad states, by an exact backward

computation along the reached sets. Two parametrized automata abstraction schemes are provided in [BHV04], both based on state merging. These schemes fit in our definition of extrapolation, and therefore can also be used in our framework. Notice that in ARMC, abstraction is performed on the data structures that are used to represent sets of configurations, whereas in our case the system itself is abstracted. After each refinement step, ARMC restarts (from scratch) the approximated forward exploration from the refined reached set, whereas our refinement is *local* to the spurious counterexample path. Moreover, the precision of the abstraction is *global* in ARMC, and may only increase (for the entire system) at each refinement step. In contrast, our path invariant generation procedures only use the precision *required* for each spurious counterexample. Preliminary benchmarks demonstrate the benefit of our local and adaptive approach for the larger examples, where a “highly” precise abstraction is required only for a few control loops. Last, our approach is not tied to words and automata. In this work we only focus on fifo systems, but our framework is fully generic and could be applied to other infinite-state systems (e.g., hybrid systems), provided that suitable parametrized extrapolations are designed (e.g., on polyhedra).

Outline. We recapitulate fifo systems in Section 2 and define their partition abstractions in Section 3. Refinement and extrapolation-based generation of path invariants are developed in Section 4. In Section 5, we present the general CEGAR semi-algorithm, and analyze its correctness and termination. Section 6 provides an overview of the extrapolation used in our implementation. Experimental results are presented in Section 7, along with some perspectives.

Due to space limitations, all proofs were omitted in this paper. A long version with detailed proofs and additional material can be obtained from the authors.

2 Fifo Systems

This section presents basic definitions and notations for fifo systems that will be used throughout the paper. For any set S we write $\wp(S)$ for the set of all subsets of S , and S^n for the set of n -tuples over S (when $n \geq 1$). For any $i \in \{1, \dots, n\}$, we denote by $s(i)$ the i^{th} component of an n -tuple s . Given $s \in S^n$, $i \in \{1, \dots, n\}$ and $u \in S$, we write $s[i \leftarrow u]$ for the n -tuple $s' \in S^n$ defined by $s'(i) = u$ and $s'(j) = s(j)$ for all $j \in \{1, \dots, n\}$ with $j \neq i$.

Let Σ denote an *alphabet* (i.e., a non-empty set of *letters*). We write Σ^* for the set of all *finite words* (*words* for short) over Σ , and we let ε denote the *empty word*. For any two words $w, w' \in \Sigma^*$, we write $w \cdot w'$ for their *concatenation*. A *language* is any subset of Σ^* . For any language L , we denote by L^* its *Kleene closure* and we write $L^+ = L \cdot L^*$. The *alphabet of L* , written $\text{alph}(L)$, is the least subset A of Σ such that $L \subseteq A^*$. For any word $w \in \Sigma^*$, the singleton language $\{w\}$ will be written simply as word w when no confusion is possible.

Safety Verification of Labeled Transition Systems. We will use labeled transition systems to formally define the behavioral semantics of fifo systems. A *labeled transition system* is any triple $LTS = \langle \mathcal{C}, \Sigma, \rightarrow \rangle$ where \mathcal{C} is a set of *configurations*,




Fig. 1. The Connection/Disconnection Protocol [JR86]

Σ is a finite set of *actions* and $\rightarrow \subseteq \mathcal{C} \times \Sigma \times \mathcal{C}$ is a (labeled) *transition relation*. We say that *LTS* is *finite* when \mathcal{C} is finite. For simplicity, we will often write $c \xrightarrow{l} c'$ in place of $(c, l, c') \in \rightarrow$.

A *finite path* (*path* for short) in *LTS* is any pair $\pi = (c, u)$ where $c \in \mathcal{C}$, and u is either the empty sequence, or a non-empty finite sequence of transitions $(c_0, l_0, c'_0), \dots, (c_{h-1}, l_{h-1}, c'_h)$ such that $c_0 = c$ and $c'_i = c_{i+1}$ for every $0 < i < h$. We simply write π as $c_0 \xrightarrow{l_0} \dots \xrightarrow{l_{h-1}} c_h$. The natural number h is called the *length* of π . We say that π is a *simple path* if $c_i \neq c_j$ for all $0 \leq i < j \leq h$. For any two sets $Init \subseteq \mathcal{C}$ and $Bad \subseteq \mathcal{C}$ of configurations, a *path from $Init$ to Bad* is any path $c_0 \xrightarrow{l_0} \dots \xrightarrow{l_{h-1}} c_h$ such that $c_0 \in Init$ and $c_h \in Bad$. Observe that if $c \in Init \cap Bad$ then c is a path (of zero length) from $Init$ to Bad . The *reachability set* of *LTS* from $Init$ is the set of configurations c such that there is a path from $Init$ to $\{c\}$.

In this paper, we focus on the verification of safety properties on fifo systems. A safety property is in general specified as a set of “bad” configurations that should not be reachable from the initial configurations. Formally, a *safety condition* for a labeled transition system $LTS = \langle \mathcal{C}, \Sigma, \rightarrow \rangle$ is a pair $(Init, Bad)$ of subsets of \mathcal{C} . We say that *LTS* is $(Init, Bad)$ -unsafe if there is a path from $Init$ to Bad in *LTS*, which is called a *counterexample*. We say that *LTS* is $(Init, Bad)$ -safe when it is not $(Init, Bad)$ -unsafe.

Fifo Systems. The asynchronous communication of distributed systems is usually modeled as a set of local processes together with a network topology given by channels between processes. Each process can be modeled by a finite-state machine that sends and receives messages on the channels to which it is connected. Let us consider a classical example, which will be used in the remainder of this paper to illustrate our approach.

Example 2.1. The connection/disconnection protocol [JR86] – abbreviated as *c/d* protocol – between two hosts is depicted in Figure 1. This model is composed of two processes, a client and a server, as well as two unidirectional channels.

To simplify the presentation, we restrict our attention to the case of one finite-state control process. The general case of multiple processes can be reduced to this simpler form by taking the asynchronous product of all processes. For the connection/disconnection protocol, the asynchronous product of the two processes is depicted in Figure 2.

Definition 2.2. A fifo system \mathcal{A} is a 4-tuple $\langle Q, M, n, \Delta \rangle$ where:




Fig. 2. Fifo System Representing the Connection/Disconnection Protocol

- Q is a finite set of control states,
- M is a finite alphabet of messages,
- $n \geq 1$ is the number of fifo queues,
- $\Delta \subseteq Q \times \Sigma \times Q$ is a set of transition rules,
 where $\Sigma = \{1, \dots, n\} \times \{!, ?\} \times M$ is the set of fifo actions over n queues.

Simplifying notation, fifo actions in Σ will be shortly written $i!m$ and $i?m$ instead of $(i, !, m)$ and $(i, ?, m)$. The intended meaning of fifo actions is the following: $i!m$ means “emission of message m on queue i ” and $i?m$ means “reception of message m from queue i ”. The operational semantics of a fifo system \mathcal{A} is formally given by its associated labeled transition system $\llbracket \mathcal{A} \rrbracket$ defined below.

Definition 2.3. The operational semantics of a fifo system $\mathcal{A} = \langle Q, M, n, \Delta \rangle$ is the labeled transition system $\llbracket \mathcal{A} \rrbracket = \langle \mathcal{C}, \Sigma, \rightarrow \rangle$ defined as follows:

- $\mathcal{C} = Q \times (M^*)^n$ is the set of configurations,
- $\Sigma = \{1, \dots, n\} \times \{!, ?\} \times M$ is the set of actions,
- the transition relation $\rightarrow \subseteq \mathcal{C} \times \Sigma \times \mathcal{C}$ is the set of triples $((q, \mathbf{w}), l, (q', \mathbf{w}'))$ such that $(q, l, q') \in \Delta$ and that satisfy the two following conditions:
 - if $l = i!m$ then $\mathbf{w}'(i) = \mathbf{w}(i) \cdot m$ and $\mathbf{w}'(j) = \mathbf{w}(j)$ for all $j \neq i$,
 - if $l = i?m$ then $\mathbf{w}(i) = m \cdot \mathbf{w}'(i)$ and $\mathbf{w}'(j) = \mathbf{w}(j)$ for all $j \neq i$.

Example 2.4. The fifo system $\mathcal{A} = \langle \{00, 01, 10, 11\}, \{o, c, d\}, 2, \Delta \rangle$ that corresponds to the c/d protocol is displayed in Figure 2. The set of initial configurations is $Init = \{(00, \varepsilon, \varepsilon)\}$. A set of bad configurations for this protocol is $Bad = \{00, 10\} \times (c \cdot M^* \times M^*)$. This set contains configurations where the server is in control state 0 but the first message in the first queue is `close`. This is the classical case of an *undefined reception* which results in a (local) *deadlock* for the server. Setting the initial configuration to $c_0 = (00, \varepsilon, \varepsilon)$, a counterexample to the safety condition $(\{c_0\}, Bad)$ is the path $(00, \varepsilon, \varepsilon) \xrightarrow{1!o} (10, o, \varepsilon) \xrightarrow{1?o} (11, \varepsilon, \varepsilon) \xrightarrow{2!d} (10, \varepsilon, d) \xrightarrow{1!c} (00, c, d)$ in $\llbracket \mathcal{A} \rrbracket$.

3 Partition Abstraction for Fifo Systems

In the context of CEGAR-based safety verification, automatic abstraction techniques are usually based on predicates [GS97] or partitions [CGJ⁺03]. In this work, we focus on partition-based abstraction and refinement techniques for fifo systems. A *partition* of a set S is any set P of non-empty pairwise disjoint subsets of S such that $S = \bigcup_{p \in P} p$. Elements p of a partition P are called *classes*. For any element s in S , we denote by $[s]_P$ the class in P containing s .

At the labeled transition system level, partition abstraction consists of merging configurations that are equivalent with respect to a given equivalence relation, or a given partition. In practice, it is often desirable to maintain different partitions for different control states, to keep partition sizes relatively small. We follow this approach in our definition of partition abstraction for fifo systems, by associating a partition of $(M^*)^n$ with each control state. To ease notation, we write $\bar{L} = (M^*)^n \setminus L$ for the *complement* of any subset L of $(M^*)^n$.

To effectively compute partition abstractions for fifo systems, we need a family of finitely representable subsets of $(M^*)^n$. A natural candidate is the class of recognizable subsets of $(M^*)^n$, or, equivalently, of QDD-definable subsets of $(M^*)^n$ [BGWW97], since this class is effectively closed under Boolean operations. Recall that a subset L of $(M^*)^n$ is *recognizable* if (and only if) it is a finite union of subsets of the form $L_1 \times \dots \times L_n$ where each L_i is a regular language over M [Ber79]. We extend recognizability in the natural way to subsets of the set $\mathcal{C} = Q \times (M^*)^n$ of configurations. A subset $C \subseteq \mathcal{C}$ is *recognizable* if $\{\mathbf{w} \mid (q, \mathbf{w}) \in C\}$ is recognizable for every $q \in Q$. We denote by $\text{Rec}((M^*)^n)$ the set of recognizable subsets of $(M^*)^n$, and write $\mathbb{P}((M^*)^n)$ for the set of all finite partitions of $(M^*)^n$ where classes are recognizable subsets of $(M^*)^n$.

Definition 3.1. *Consider a fifo system $\mathcal{A} = \langle Q, M, n, \Delta \rangle$ and a partition map $P : Q \rightarrow \mathbb{P}((M^*)^n)$. The partition abstraction of $\llbracket \mathcal{A} \rrbracket$ induced by P is the finite labeled transition system $\llbracket \mathcal{A} \rrbracket_P^\sharp = \langle \mathcal{C}_P^\sharp, \Sigma, \rightarrow_P^\sharp \rangle$ defined as follows:*

- $\mathcal{C}_P^\sharp = \{(q, p) \mid q \in Q \text{ and } p \in P(q)\}$ is the set of abstract configurations,
- $\Sigma = \{1, \dots, n\} \times \{!, ?\} \times M$ is the set of actions,
- the abstract transition relation $\rightarrow_P^\sharp \subseteq \mathcal{C}_P^\sharp \times \Sigma \times \mathcal{C}_P^\sharp$ is the set of triples $((q, p), l, (q', p'))$ such that $(q, \mathbf{w}) \xrightarrow{l} (q', \mathbf{w}')$ for some $\mathbf{w} \in p$ and $\mathbf{w}' \in p'$.

To relate concrete and abstract configurations, we define the *abstraction function* $\alpha_P : \mathcal{C} \rightarrow \mathcal{C}_P^\sharp$, and its extension to $\wp(\mathcal{C}) \rightarrow \wp(\mathcal{C}_P^\sharp)$, as well as the *concretization function* $\gamma_P : \mathcal{C}_P^\sharp \rightarrow \mathcal{C}$, extended to $\wp(\mathcal{C}_P^\sharp) \rightarrow \wp(\mathcal{C})$, as expected:

$$\begin{aligned} \alpha_P((q, \mathbf{w})) &= (q, [\mathbf{w}]_{P(q)}) & \alpha_P(C) &= \{\alpha(c) \mid c \in C\} \\ \gamma_P((q, p)) &= \{q\} \times p & \gamma_P(C^\sharp) &= \bigcup \{\gamma(c^\sharp) \mid c^\sharp \in C^\sharp\} \end{aligned}$$

To simplify notations, we shall drop the P subscript when the partition map can easily be derived from the context. Intuitively, an abstract configuration (q, p) of $\llbracket \mathcal{A} \rrbracket^\sharp$ represents the set $\{q\} \times p$ of (concrete) configurations of $\llbracket \mathcal{A} \rrbracket$. The




Fig. 3. Example Partition Abstraction of the C/D Protocol (Example 3.5)

abstract transition relation \rightarrow^\sharp is the existential lift of the concrete transition relation \rightarrow to abstract configurations.

The following forward and backward language transformers will be used to capture the effect of fifo actions. The functions $post : \Sigma \times \wp((M^*)^n) \rightarrow \wp((M^*)^n)$ and $pre : \Sigma \times \wp((M^*)^n) \rightarrow \wp((M^*)^n)$ are defined by:

$$\begin{aligned} post(i!m, L) &= \{\mathbf{w}[i \leftarrow u] \mid \mathbf{w} \in L, u \in M^* \text{ and } \mathbf{w}(i) \cdot m = u\} \\ post(i?m, L) &= \{\mathbf{w}[i \leftarrow u] \mid \mathbf{w} \in L, u \in M^* \text{ and } \mathbf{w}(i) = m \cdot u\} \\ pre(i!m, L) &= \{\mathbf{w}[i \leftarrow u] \mid \mathbf{w} \in L, u \in M^* \text{ and } \mathbf{w}(i) = u \cdot m\} \\ pre(i?m, L) &= \{\mathbf{w}[i \leftarrow u] \mid \mathbf{w} \in L, u \in M^* \text{ and } m \cdot \mathbf{w}(i) = u\} \end{aligned}$$

Obviously, $post(l, L)$ and $pre(l, L)$ are effectively recognizable subsets of $(M^*)^n$ for any $l \in \Sigma$ and any recognizable subset $L \subseteq (M^*)^n$. Moreover, we may use $post$ and pre to characterize the abstract transition relation of a partition abstraction $\llbracket \mathcal{A} \rrbracket_P^\sharp$, as follows: for any rule $(q, l, q') \in \Delta$ and for any pair $(p, p') \in P(q) \times P(q')$, we have $(q, p) \xrightarrow{l}^\sharp (q', p')$ iff $post(l, p) \cap p' \neq \emptyset$ iff $p \cap pre(l, p') \neq \emptyset$.

Lemma 3.2. *For any fifo system \mathcal{A} and partition map $P : Q \rightarrow \mathbb{P}((M^*)^n)$, $\llbracket \mathcal{A} \rrbracket^\sharp$ is effectively computable. For any recognizable subset $C \subseteq \mathcal{C}$, $\alpha(C)$ is effectively computable.*

We extend α to paths in the obvious way: $\alpha(c_0 \xrightarrow{l_0} \dots \xrightarrow{l_{h-1}} c_h) = \alpha(c_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp \alpha(c_h)$. Observe that $\alpha(\pi)$ is an abstract path in $\llbracket \mathcal{A} \rrbracket^\sharp$ for any concrete path π in $\llbracket \mathcal{A} \rrbracket$. We therefore obtain the following safety preservation property.

Proposition 3.3. *Consider a fifo system \mathcal{A} and a safety condition $(Init, Bad)$ for $\llbracket \mathcal{A} \rrbracket$. For any partition abstraction $\llbracket \mathcal{A} \rrbracket^\sharp$ of $\llbracket \mathcal{A} \rrbracket$, if $\llbracket \mathcal{A} \rrbracket^\sharp$ is $(\alpha(Init), \alpha(Bad))$ -safe then $\llbracket \mathcal{A} \rrbracket$ is $(Init, Bad)$ -safe.*

The converse to this proposition does not hold generally. An abstract counterexample π^\sharp is called *feasible* if there exists a concrete counterexample π such that $\pi^\sharp = \alpha(\pi)$, and π^\sharp is called *spurious* otherwise.

Lemma 3.4. *For any fifo system \mathcal{A} , any partition map $P : Q \rightarrow \mathbb{P}((M^*)^n)$, and any safety condition $(Init, Bad)$ for $\llbracket \mathcal{A} \rrbracket$, feasibility of abstract counterexamples is effectively decidable.*

Example 3.5. Continuing the discussion of the c/d protocol, we consider the partition abstraction induced by the following partition map:

$q \in Q$	00	10	01	11
$P(q)$	$\varepsilon \times \varepsilon, \bar{\varepsilon} \times \bar{\varepsilon}$	$\mathfrak{o}^* \times \varepsilon, \bar{\mathfrak{o}}^* \times \bar{\varepsilon}$	$M^* \times M^*$	$M^* \times M^*$

The set of initial abstract configurations is $\alpha(Init) = \{(00, \varepsilon \times \varepsilon)\}$, and the set of bad abstract configurations is $\alpha(Bad) = \{(00, \bar{\varepsilon} \times \bar{\varepsilon}), (10, \bar{\mathfrak{o}}^* \times \bar{\varepsilon})\}$. The resulting partition abstraction is the finite labeled transition system depicted in Figure 3. A simple graph search reveals several abstract counterexamples, for instance $\pi^\sharp = (00, \varepsilon \times \varepsilon) \xrightarrow{1!o}^\sharp (10, \mathfrak{o}^* \times \varepsilon) \xrightarrow{1!c}^\sharp (00, \bar{\varepsilon} \times \bar{\varepsilon})$. This counterexample is spurious since the only concrete path that corresponds to π^\sharp (i.e., whose image under α is π^\sharp) is $\pi = (00, \varepsilon, \varepsilon) \xrightarrow{1!o} (10, \mathfrak{o}, \varepsilon) \xrightarrow{1!c} (00, \mathfrak{o}c, \varepsilon) \notin Bad$.

4 Counterexample-based Partition Refinement

The abstraction-based verification of safety properties relies on refinement techniques that gradually increase the precision of abstractions in order to rule out spurious abstract counterexamples. Refinement for partition abstractions simply consists in splitting some classes into a sub-partition.

Given two partitions P and \tilde{P} of a set S , we say that \tilde{P} *refines* P when each class $\tilde{p} \in \tilde{P}$ is contained in some class $p \in P$. Moreover we then write $[\tilde{p}]_P$ for the class $p \in P$ containing \tilde{p} .

Let us fix, for the remainder of this section, a fifo system $\mathcal{A} = \langle Q, M, n, \Delta \rangle$ and a safety condition $(Init, Bad)$ for $\llbracket \mathcal{A} \rrbracket$. Given two partition maps $P, \tilde{P} : Q \rightarrow \mathbb{P}((M^*)^n)$, we say that \tilde{P} *refines* P if $\tilde{P}(q)$ refines $P(q)$ for every control state $q \in Q$. If \tilde{P} refines P , then for any abstract path $(q_0, \tilde{p}_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, \tilde{p}_h)$ in $\llbracket \mathcal{A} \rrbracket_{\tilde{P}}^\sharp$, it holds that $(q_0, [\tilde{p}_0]_{P(q_0)}) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, [\tilde{p}_h]_{P(q_h)})$ is an abstract path in $\llbracket \mathcal{A} \rrbracket_P^\sharp$. This fact shows that, informally, refining a partition abstraction does not introduce new spurious counterexample.

When a spurious counterexample is found in the abstraction, the partition map must be refined so as to rule out this counterexample. We formalize this concept for an abstract path $\pi_P^\sharp = (q_0, p_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, p_h)$ in $\llbracket \mathcal{A} \rrbracket_P^\sharp$ from $\alpha_P(Init)$ to $\alpha_P(Bad)$ as follows: a refinement \tilde{P} of P is said to *rule out* the abstract counterexample π_P^\sharp if there exists no path $\pi_{\tilde{P}}^\sharp = (q_0, \tilde{p}_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, \tilde{p}_h)$ from $\alpha_{\tilde{P}}(Init)$ to $\alpha_{\tilde{P}}(Bad)$ in $\llbracket \mathcal{A} \rrbracket_{\tilde{P}}^\sharp$ satisfying $\tilde{p}_i \subseteq p_i$ for all $0 \leq i \leq h$.

Note that if π_P^\sharp is a feasible counterexample, then no refinement of P can rule it out. Conversely, if \tilde{P} is a refinement of P that rules out π_P^\sharp then any refinement of \tilde{P} also rules out π_P^\sharp . The main challenge in CEGAR is the discovery of “suitable” refinements, that are computationally “simple” but “precise enough” to rule out spurious counterexamples. In this work, we focus on counterexample-guided refinements based on path invariants.

Definition 4.1. Consider a partition map P and a spurious counterexample $\pi^\sharp = (q_0, p_0) \xrightarrow{l_0, \sharp} \dots \xrightarrow{l_{h-1}, \sharp} (q_h, p_h)$ in $\llbracket \mathcal{A} \rrbracket_P^\sharp$. A path invariant for π^\sharp is any sequence L_0, \dots, L_h of recognizable subsets of $(M^*)^n$ such that:

- (i) we have $(\{q_0\} \times p_0) \cap \text{Init} \subseteq \{q_0\} \times L_0$, and
- (ii) we have $\text{post}(l_i, p_i \cap L_i) \subseteq L_{i+1}$ for every $0 \leq i < h$, and
- (iii) we have $(\{q_h\} \times L_h) \cap \text{Bad} = \emptyset$

Observe that condition (ii) is more general than $\text{post}(l_i, L_i) \subseteq L_{i+1}$ which is classically required for inductive invariants. With this relaxed condition, path invariants are tailored to the given spurious counterexample, and therefore can be simpler (e.g., be coarser or have more empty L_i).

Proposition 4.2. Consider a partition map P and a simple spurious counterexample $\pi^\sharp = (q_0, p_0) \xrightarrow{l_0, \sharp} \dots \xrightarrow{l_{h-1}, \sharp} (q_h, p_h)$. Given a path invariant L_0, \dots, L_h for π^\sharp , the partition map \tilde{P} defined below is a refinement of P that rules out π^\sharp :

$$\tilde{P}(q) = (P(q) \setminus \{p_i \mid i \in I(q)\}) \cup \bigcup_{i \in I(q)} \{p_i \cap L_i, p_i \cap \bar{L}_i\} \setminus \{\emptyset\}$$

where $I(q) = \{i \mid 0 \leq i \leq h, q_i = q\}$ for each control state $q \in Q$.

We propose a generic approach to obtain path invariants by utilizing a parametrized approximation operator for queue contents. The parameter (the k in the definition below) is used to adjust the precision of the approximation.

Definition 4.3. A (parametrized) extrapolation is any function ∇ from \mathbb{N} to $\text{Rec}((M^*)^n) \rightarrow \text{Rec}((M^*)^n)$ that satisfies, for any $L \in \text{Rec}((M^*)^n)$, the two following conditions (with $\nabla(k)$ written as ∇_k):

- (i) we have $L \subseteq \nabla_k(L)$ for every $k \in \mathbb{N}$,
- (ii) there exists $k_L \in \mathbb{N}$ such that $L = \nabla_k(L)$ for every $k \geq k_L$.

Our definition of extrapolation is quite general, in particular, it does not require monotonicity in k or in L , but it is adequate for the design of path invariant generation procedures. The most simple extrapolation is the *identity extrapolation* that maps each $k \in \mathbb{N}$ to the identity on $\text{Rec}((M^*)^n)$. The parametrized automata approximations of [BHV04] and [LGJJ06] also satisfy the requirements of Definition 4.3. The choice of an appropriate extrapolation with respect to the underlying domain of fifo systems is crucial for the implementation of CEGAR’s refinement step, and will be discussed in Section 6.

UPIInv ($\nabla, Init, Bad, \pi_P^\sharp$)

Input: extrapolation ∇ , recognizable subsets $Init, Bad$ of $Q \times (M^*)^n$, spurious counterexample $\pi_P^\sharp = (q_0, p_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, p_h)$

```

1   $k \leftarrow 0$ 
2  do
3     $L_0 \leftarrow \nabla_k(p_0 \cap \{\mathbf{w} \mid (q_0, \mathbf{w}) \in Init\})$ 
4    for  $i$  from 1 upto  $h$ 
5       $F_i \leftarrow post(l_{i-1}, p_{i-1} \cap L_{i-1})$ 
6      if  $p_i \cap F_i = \emptyset$ 
7         $L_i \leftarrow \emptyset$ 
8      else
9         $L_i \leftarrow \nabla_k(F_i)$ 
10      $k \leftarrow k + 1$ 
11  while  $(\{q_h\} \times L_h) \cap Bad \neq \emptyset$ 
12  return  $(L_0, \dots, L_h)$ 

```

Split (∇, L_0, L_1)

Input: extrapolation ∇ , disjoint recognizable subsets L_0, L_1 of $(M^*)^n$

```

1   $k \leftarrow 0$ 
2  while  $\nabla_k(L_0) \cap L_1 \neq \emptyset$ 
3     $k \leftarrow k + 1$ 
4  return  $\nabla_k(L_0)$ 

```

APIInv ($\nabla, Init, Bad, \pi_P^\sharp$)

Input: extrapolation ∇ , recognizable subsets $Init, Bad$ of $Q \times (M^*)^n$, spurious counterexample $\pi_P^\sharp = (q_0, p_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, p_h)$

```

1   $B_h \leftarrow p_h \cap \{\mathbf{w} \mid (q_h, \mathbf{w}) \in Bad\}$ 
2   $i \leftarrow h$ 
3  while  $B_i \neq \emptyset$  and  $i > 0$ 
4     $i \leftarrow i - 1$ 
5     $B_i \leftarrow p_i \cap pre(l_i, B_{i+1})$ 
6  if  $i = 0$ 
7     $I \leftarrow p_0 \cap \{\mathbf{w} \mid (q_0, \mathbf{w}) \in Init\}$ 
8     $L_0 \leftarrow Split(\nabla, I, B_0)$ 
9  else
10    $(L_0, \dots, L_i) \leftarrow ((M^*)^n, \dots, (M^*)^n)$ 
11  for  $j$  from  $i$  upto  $h - 1$ 
12    $L_{j+1} \leftarrow Split(\nabla, post(l_j, p_j \cap L_j), B_{j+1})$ 
13  return  $(L_0, \dots, L_h)$ 

```

Fig. 4. Extrapolation-based Path Invariant Generation Algorithms

Remark 4.4. Extrapolations are closed under various operations, such as functional composition, union and intersection, as well as round-robin combination.

We now present two extrapolation-based path invariant generation procedures (Figure 4). Recall that the parameter k of an extrapolation intuitively indicates the desired precision of the approximation. The first algorithm, **UPIInv**, performs an approximated *post* computation along the spurious counterexample, and iteratively increases the precision k of the approximation until a path invariant is obtained. The applied precision in **UPIInv** is uniform along the counterexample. Due to its simplicity, the termination analysis of **CEGAR** in the following section will refer to **UPIInv**. The second algorithm, **APIInv**, first performs an exact *pre* computation along the spurious counterexample to identify the “bad” coreachable subsets B_i . The path invariant is then computed with a forward traversal that uses the **Split** subroutine to simplify each *post* image while remaining disjoint from the B_i . The precision used in **Split** is therefore tailored to each *post* image, which may lead to simpler path invariants. Naturally, both algorithms may be “reversed” to generate path invariants backwards.

Observe that if the extrapolation ∇ is effectively computable, then all steps in the algorithms **UPIInv**, **Split** and **APIInv** are effectively computable. We now prove correctness and termination of these algorithms. Let us fix, for the remainder of this section, an extrapolation ∇ and a partition map $P : Q \rightarrow \mathbb{P}((M^*)^n)$, and assume that *Init* and *Bad* are recognizable.

Proposition 4.5. *For any spurious abstract counterexample π_P^\sharp , the execution of $\text{UPIInv}(\nabla, \text{Init}, \text{Bad}, \pi_P^\sharp)$ terminates and returns a path invariant for π_P^\sharp .*

Lemma 4.6. *For any two recognizable subsets L_0, L_1 of $(M^*)^n$, if $L_0 \cap L_1 = \emptyset$ then $\text{Split}(\nabla, L_0, L_1)$ terminates and returns a recognizable subset L of $(M^*)^n$ that satisfies $L_0 \subseteq L \subseteq \overline{L_1}$.*

Proposition 4.7. *For any spurious abstract counterexample π_P^\sharp , the execution of $\text{APIInv}(\nabla, \text{Init}, \text{Bad}, \pi_P^\sharp)$ terminates and returns a path invariant for π_P^\sharp .*

Example 4.8. Consider again the c/d protocol, and assume an extrapolation ∇ satisfying $\nabla_0(L \times \varepsilon) = (\text{alph}(L))^* \times \varepsilon$ for all $L \subseteq M^*$, and $\nabla_1(u \times \varepsilon) = u \times \varepsilon$ for each $u \in \{\varepsilon, \text{o}, \text{oc}\}$, e.g., the extrapolation ρ'' presented in Remark 6.1. The **UPIInv** algorithm, applied to the spurious counterexample $(\text{00}, \varepsilon \times \varepsilon) \xrightarrow{1\text{o}, \sharp} (\text{10}, \text{o}^* \times \varepsilon) \xrightarrow{1\text{c}, \sharp} (\text{00}, \overline{\varepsilon \times \varepsilon})$ of Example 3.5, would perform two iterations of the **while**-loop and produce the path invariant $(\varepsilon \times \varepsilon, \text{o} \times \varepsilon, \text{oc} \times \varepsilon)$. These iterations are detailed in the table below. The mark $\frac{1}{2}$ or \checkmark indicates whether the condition at line 11 is satisfied.

	L_0	L_1	L_2	Line 11
$k = 0$	$\varepsilon \times \varepsilon$	$\text{o}^* \times \varepsilon$	$\{\text{o}, \text{c}\}^* \times \varepsilon$	$\frac{1}{2}$
$k = 1$	$\varepsilon \times \varepsilon$	$\text{o} \times \varepsilon$	$\text{oc} \times \varepsilon$	\checkmark

Following Proposition 4.2, the partition map would be refined to:

$q \in Q$	00	10	01, 11
$P(q)$	$\varepsilon \times \varepsilon, \text{oc} \times \varepsilon, (\varepsilon \cup \text{oc}) \times \varepsilon$	$\text{o} \times \varepsilon, (\varepsilon \cup (\text{o} \cdot \text{o}^+)) \times \varepsilon, \overline{\text{o}^*} \times \varepsilon$	$M^* \times M^*$

This refined partition map clearly rules out the spurious counterexample.

5 Safety Cegar Semi-Algorithm for Fifo Systems

We are now equipped with the key ingredients to present our CEGAR semi-algorithm for fifo systems. The semi-algorithm takes as input a fifo system \mathcal{A} , a recognizable safety condition $(Init, Bad)$, an initial partition map P_0 , and a path invariant generation procedure $PathInv$. The initial partition map may be the trivial one, mapping each control state to $(M^*)^n$. We may use any path invariant generation procedure, such as the ones presented in the previous section. The semi-algorithm iteratively refines the partition abstraction until either the abstraction is precise enough to prove that $\llbracket \mathcal{A} \rrbracket$ is $(Init, Bad)$ -safe (line 10), or a feasible counterexample is found (line 4). If the abstract counterexample picked at line 2 is spurious, a path invariant is generated and is used to refine the partition. The new partition map obtained after the **foreach** loop (lines 8–9) is precisely the partition map \tilde{P} from Proposition 4.2, and hence it rules out this abstract counterexample. Recall that Lemmata 3.2 and 3.4 ensure that the steps at lines 1 and 3 are effectively computable. The correctness of the CEGAR semi-algorithm is expressed by the following proposition, which directly follows from Proposition 3.3 and from the definition of feasible abstract counterexamples.

Proposition 5.1. *For any terminating execution of CEGAR($\mathcal{A}, Init, Bad, P_0, PathInv$), if the execution returns \checkmark (resp. ζ) then $\llbracket \mathcal{A} \rrbracket$ is $(Init, Bad)$ -safe (resp. $(Init, Bad)$ -unsafe).*

A detailed example execution of CEGAR on the c/d protocol is provided in the long version. Termination of the CEGAR semi-algorithm cannot be assured as, otherwise, it would solve the general reachability problem, which is known to be undecidable for fifo systems [BZ83]. However, $(Init, Bad)$ -unsafety is semi-decidable for fifo systems by forward or backward symbolic exploration when $Init$ and Bad are recognizable [BG99]. Moreover, this problem becomes decidable for fifo systems having a finite reachability set from $Init$.

We investigate in this section the termination of the CEGAR semi-algorithm when \mathcal{A} is $(Init, Bad)$ -unsafe or has a finite reachability set from $Init$. In contrast to other approaches where abstractions are refined globally (e.g., predicate abstraction [GS97]), partition abstractions [CGJ⁺03] are refined locally by splitting abstract configurations along the abstract counterexample (viz. lines 8–9 of the CEGAR semi-algorithm). The abstract transition relation only needs to be refined locally around the abstract configurations which have been split, and, hence, its refinement can be computed efficiently. However, this local nature of refinement complicates the analysis of the algorithm. We fix an extrapolation

CEGAR $(\mathcal{A}, \text{Init}, \text{Bad}, P_0, \text{PathInv})$

Input: fifo system $\mathcal{A} = \langle Q, M, n, \Delta \rangle$, recognizable subsets Init, Bad of $Q \times (M^*)^n$, partition map $P_0 : Q \rightarrow \mathbb{P}((M^*)^n)$, procedure PathInv

```

1  while  $\llbracket \mathcal{A} \rrbracket_P^\sharp$  is  $(\alpha_P(\text{Init}), \alpha_P(\text{Bad}))$ -unsafe
2      pick a simple abstract counterexample  $\pi^\sharp$  in  $\llbracket \mathcal{A} \rrbracket_P^\sharp$ 
3      if  $\pi^\sharp$  is a feasible abstract counterexample
4          return  $\not\checkmark$ 
5      else
6          write  $\pi^\sharp$  as the abstract path  $(q_0, p_0) \xrightarrow{l_0}^\sharp \dots \xrightarrow{l_{h-1}}^\sharp (q_h, p_h)$ 
7           $(L_0, \dots, L_h) \leftarrow \text{PathInv}(\text{Init}, \text{Bad}, \pi^\sharp)$ 
8          foreach  $i \in \{0, \dots, h\}$ 
9               $P(q_i) \leftarrow (P(q_i) \setminus \{p_i\}) \cup (\{p_i \cap L_i, p_i \cap \overline{L_i}\} \setminus \{\emptyset\})$ 
10     return  $\checkmark$ 
    
```

∇ and we focus on the path invariant generation procedure UPIInv presented in Section 4.

Proposition 5.2. *For any breadth-first execution of $\text{CEGAR}(\mathcal{A}, \text{Init}, \text{Bad}, P_0, \text{UPIInv}(\nabla))$, if the execution does not terminate then the sequence $(h_\theta)_{\theta \in \mathbb{N}}$ of lengths of counterexamples picked at line 2 is nondecreasing and diverges.*

Corollary 5.3. *If $\llbracket \mathcal{A} \rrbracket$ is $(\text{Init}, \text{Bad})$ -unsafe then any breadth-first execution of $\text{CEGAR}(\mathcal{A}, \text{Init}, \text{Bad}, P_0, \text{UPIInv}(\nabla))$ terminates.*

It would also be desirable to obtain termination of the CEGAR semi-algorithm when \mathcal{A} has a finite reachability set from Init . However, as demonstrated in the long version, this condition is not sufficient to guarantee that $\text{CEGAR}(\mathcal{A}, \text{Init}, \text{Bad}, P_0, \text{UPIInv}(\nabla))$ has a terminating execution. It turns out that termination can be guaranteed for fifo systems with a finite reachability set when ∇_k has a finite image for every $k \in \mathbb{N}$. This apparently strong requirement, formally specified in Definition 5.4, is satisfied by the extrapolations presented in [BHV04] and [LGJJ06], which are based on state equivalences up to a certain depth.

Definition 5.4. *An extrapolation ∇ is restricted if for every $k \in \mathbb{N}$, the set $\{\nabla_k(L) \mid L \in \text{Rec}((M^*)^n)\}$ is finite.*

Remark that if ∇ is restricted then for any execution of $\text{CEGAR}(\mathcal{A}, \text{Init}, \text{Bad}, P_0, \text{UPIInv}(\nabla))$, the execution terminates if and only if the number of iterations of the **while**-loop of the algorithm UPIInv is bounded³. As shown by the following proposition, if moreover $\llbracket \mathcal{A} \rrbracket$ has a finite reachability set from Init then the execution necessarily terminates.

Proposition 5.5. *Assume that ∇ is restricted. If $\llbracket \mathcal{A} \rrbracket$ has a finite reachability set from Init , then any execution of $\text{CEGAR}(\mathcal{A}, \text{Init}, \text{Bad}, P_0, \text{UPIInv}(\nabla))$ terminates.*

³ Remark that this bound is not a bound on the length of abstract counterexamples.

6 Overview of the (Colored) Bisimulation Extrapolation

This section briefly introduces the bisimulation-based extrapolation underlying the widening operator introduced in [LGJJ06]. This extrapolation assumes an automata representation of recognizable subsets of $(M^*)^n$, and relies on bounded-depth bisimulation over the states of the automata. For simplicity, we focus in this section on fifo systems with a single queue, i.e., $n = 1$. In this simpler case, recognizable subsets of $(M^*)^n$ are regular languages contained in M^* , which can directly be represented by finite automata over M . The general case of $n \geq 2$, which is discussed in detail in the long version, requires the use of QDDs, that are finite automata accepting recognizable subsets of $(M^*)^n$ via an encoding of n -tuples in $(M^*)^n$ by words over an extended alphabet. Still, the main ingredients rest the same.

In the remainder of this section, we lift our discussion from regular languages in M^* to finite automata over M . Consider a finite automaton over M with a set Q of states. As in abstract regular model checking [BHV04], we use quotienting under equivalence relations on Q to obtain over-approximations of the automaton. However, we follow the approach of [LGJJ06], and focus on bounded-depth bisimulation equivalence (other equivalence relations were used in [BHV04]).

Given a priori an equivalence relation col on Q , also called “coloring”, and a bound $k \in \mathbb{N}$, the (colored) *bisimulation equivalence of depth k* is the equivalence relation \sim_k^{col} on Q defined as expected: $\sim_0^{col} = col$ and two states are equivalent for \sim_{k+1}^{col} if (1) they are \sim_k^{col} -equivalent and (2) they have \sim_k^{col} -equivalent m -successors for each letter $m \in M$. The ultimately stationary sequence $\sim_0^{col} \supseteq \sim_1^{col} \supseteq \dots \supseteq \sim_k^{col} \supseteq \sim_{k+1}^{col} \supseteq \dots$ of equivalence relations on Q leads to the colored bisimulation-based extrapolation.

We define a coloring std , called *standard coloring*, by $(q_1, q_2) \in std$ if either q_1 and q_2 are both final states or q_1 and q_2 are both non-final states. The *bisimulation extrapolation* is the function ρ from \mathbb{N} to $\mathcal{R}ec(M^*) \rightarrow \mathcal{R}ec(M^*)$ defined by $\rho_k(L) = L / \sim_k^{std}$, where L is identified to a finite automaton accepting it. Notice that ρ is a restricted extrapolation.

Remark 6.1. We could also choose other colorings or define the sequence of equivalences in a different way. For instance, better results are sometimes obtained in practice with the extrapolation ρ' that first (for $k = 0$) applies a quotienting with respect to the equivalence relation $Q \times Q$ (i.e., all states are merged), and then behaves as ρ_{k-1} (for $k \geq 1$). Analogously, the extrapolation ρ'' defined by $\rho''_0 = \rho'_0$ and $\rho''_k = \rho_k$ for $k \geq 1$ was used in Example 4.8.

Example 6.2. Consider the regular language $L = \{aac, baaa\}$ over the alphabet $M = \{a, b, c, d, e\}$, represented by the automaton FA_L of Figure 5a. The previously defined extrapolation ρ applies to L as follows: ρ_0 splits the states of FA_L according to std , hence, $\rho_0(L) = \{a, b\}^* \cdot \{a, c\}$ (viz. Figure 5c). Then ρ_1 merges the states that are bisimulation equivalent up to depth 1, i.e., the states \bigcirc of FA_L (Figure 5d). As all states of FA_L are non equivalent for \sim_k^{std} with $k \geq 2$, we




Fig. 5. Finite Automata Representations for Extrapolating L (Example 6.2)

have $\rho_k(L) = L$ (again Figure 5a). The variants ρ' and ρ'' mentioned previously would lead to $\rho'_0(L) = \rho''_0(L) = (\text{alph}(L))^* = \{a, b, c\}^*$ (viz. Figure 5b).

The benefits of the bisimulation extrapolation for the abstraction of fifo systems were already discussed in [LGJJ06]. The following example shows that this extrapolation can, in some common cases, discover *exact* repetitions of message sequences in queues, without the need for acceleration techniques.

Example 6.3. Let us continue the running example of the c/d protocol, considered here as having a single-queue by restricting it to operations on the first queue. The application of acceleration techniques on the path $(00, \varepsilon) \xrightarrow{1!o} \xrightarrow{1!c}$ $(00, oc) \xrightarrow{1!o} \xrightarrow{1!c} (00, ococ) \dots$ would produce the set of queue contents $(oc)^+$. The bisimulation extrapolation ρ applied to the singleton language $\{ococ\}$, represented by the obvious automaton, produces the following results for the first two parameters: $\rho_0(\{ococ\}) = \{o, c\}^* \cdot c$ and $\rho_1(\{ococ\}) = (oc)^+$.

7 Experimental Evaluation

Our prototypical tool McScM that implements the previous algorithms is written in OCAML and relies on a library by Le Gall and Jeannet [Scm] for the classical finite automata and QDD operations, the fifo *post/pre* symbolic computations, as well as the colored bisimulation-based extrapolation. The standard coloring with final and non-final states is used by default in our tool, but several other variants are also available.

Our implementation includes the two path invariant generation algorithms UPIInv and APIInv of Section 4. We actually implemented a “single split” backward variant of APIInv, reminiscent of the classical CEGAR implementation [CGJ⁺03] (analogous to APIInv but applying the split solely to the “failure” abstract configuration). Therefore our implemented variant APIInv' leads to more CEGAR loops than would be obtained with APIInv, and this explains in part why UPIInv globally outperforms APIInv' for larger examples. Several pluggable subroutines can be used to search for counterexamples (depth-first or breadth-first exploration).

We tested the prototype on a suite of protocols that includes the classical alternating bit protocol ABP [AJ96], a simplified version of TCP – also in the setting of one server with two clients that interfere on their shared channels, a sliding window protocol, as well as protocols for leader election due to Peterson

protocol	states/trans.	refmnt.	time [s]	mem [MiB]	loops	states [#] /trans [#]
ABP	16/64	APIInv'	0.30	1.09	72	87/505
		UPIInv	2.13	1.58	208	274/1443
c/d protocol	5/17	APIInv'	0.02	0.61	8	12/51
		UPIInv	0.01	0.61	6	11/32
nested c/d protocol	6/17	APIInv'	0.68	1.09	80	85/314
		UPIInv	1.15	1.09	93	100/339
non-regular protocol	9/18	APIInv'	0.02	0.61	13	21/47
		UPIInv	0.06	0.61	14	25/39
Peterson	10648/56628	APIInv'	7.05	32.58	299	10946/58536
		UPIInv	2.14	32.09	51	10709/56939
(simplified) TCP	196/588	APIInv'	2.19	3.03	526	721/3013
		UPIInv	1.38	2.06	183	431/1439
server with 2 clients	255/2160	APIInv'	(> 1h)	—	—	—
		UPIInv	9.61	4.97	442	731/7383
token ring	625/4500	APIInv'	85.15	19.50	1720	2344/19596
		UPIInv	4.57	6.42	319	1004/6956
sliding window	225/2010	APIInv'	16.43	9.54	1577	1801/15274
		UPIInv	0.93	2.55	148	388/2367

Table 1. Benchmark results of McScM on a suite of protocols.

and token passing in a ring topology. Further, we provide certain touchstones for our approach: an enhancement of the c/d protocol with nested loops for the exchange of data, and a protocol with a non-recognizable reachability set. A detailed presentation of the protocols is provided in the long version. Except for the c/d protocol, which is unsafe, all other examples are safe.

Table 1 gives a summary of the results obtained by McScM on an off-the-shelf computer (2.4 GHz Intel Core 2 Duo). Breadth-first exploration was applied in all examples to search for abstract counterexamples. The bisimulation extrapolation ρ presented in Section 6 was used except for the server with 2 clients, where we applied the variant ρ' of ρ presented in Remark 6.1, as the analysis did not terminate after one hour with ρ . All examples are analyzed with UPIInv in a few seconds, and memory is not a limitation.

We compared McScM with TRex [Tre], which is, to the best of our knowledge, the sole publicly available and directly usable model-checker for the verification of *unbounded* fifo systems. Note, however, that the comparison is biased as TRex focuses on *lossy* channels. We applied TRex to the first six examples of Table 1. TRex has an efficient implementation based on *simple regular expressions* (and not general QDDs as we do), and needs in most cases less than 1 second to build the reachability set (the latter allows to decide the reachability of bad configurations by a simple additional look-up). Further, TRex implements communicating timed and counter automata that are – at this stage – beyond the focus of our tool. Nonetheless, TRex assumes a *lossy* fifo semantics, and, therefore, is not able to verify all *reliable* fifo protocols correctly (e.g., when omitting the `disconnect` messages in the c/d protocol, TRex is still able to reach *Bad* due to the possible

loss of messages, albeit the protocol is safe). Moreover, TRex suffers (as would also a symbolic model-checker based on the LASH library [Las]) from the main drawback of acceleration techniques, which in general cannot cope with nested loops, whereas they seem to have no adverse effect on our tool (viz. nested c/d protocol on which TRex did not finish after one hour). McScM can also handle a simple non-regular protocol (with a counting loop) that is beyond the QDD-based approaches [BG99], as the representation of the reachability set would require recognizable languages equipped with Presburger formulas (CQDDs [BH99]).

To obtain a finer evaluation of our approach, we prototypically implemented the global abstraction refinement scheme of [BHV04] in our tool. While this ARMC implementation seems to be advantageous for some small protocols, larger examples confirm that the local and adaptive approach refinement approach developed in this paper outperforms a global refinement one in protocols that demand a “highly” precise abstraction only for a few control loops (e.g., Peterson’s leader election and token ring). Further, our ARMC implementation was not able to handle the non-regular protocol nor the server with 2 clients.

8 Conclusion and Perspectives

Our prototypical implementation confirms our expectations that the proposed CEGAR framework with extrapolation-based path invariants is a promising alternative approach to the automatic verification of fifo systems.

Our approach relies on partition abstractions where equivalence classes are recognizable languages of queue contents. Our main contribution is the design of generic path invariant generation algorithms based on parametrized extrapolation operators for queue contents. Because of the latter, our CEGAR semi-algorithm enjoys additional partial termination properties.

The framework developed in this paper is not specific to fifo systems, and we intend to investigate its practical relevance to other families of infinite-state models. Future work also includes the safety verification of more complex fifo systems that would allow the exchange of unbounded numerical data over the queues, or include parametrization (e.g., over the number of clients). Several decidable classes of fifo systems have emerged in the literature (in particular lossy fifo systems) and we intend to investigate termination of our CEGAR semi-algorithm (when equipped with the path invariant generation algorithms developed in this paper) for these classes.

Acknowledgments. We thank the anonymous reviewers for supporting and guiding the genesis of this publication, and we are especially grateful for the fruitful and challenging discussions with Jérôme Leroux and Anca Muscholl.

References

- [AJ96] P. A. Abdulla and B. Jonsson. Verifying Programs with Unreliable Channels. *Information and Computation*, 127(2):91–101, 1996.

- [Ber79] J. Berstel. *Transductions and Context-Free Languages*. Teubner, 1979.
- [BG99] B. Boigelot and P. Godefroid. Symbolic Verification of Communication Protocols with Infinite State Spaces using QDDs. *Formal Methods in System Design*, 14(3):237–255, 1999.
- [BGWW97] B. Boigelot, P. Godefroid, B. Willems, and P. Wolper. The Power of QDDs. In *Proc. Static Analysis Symposium 1997*, volume 1302 of *LNCS*, pages 172–186. Springer, 1997.
- [BH99] A. Bouajjani and P. Habermehl. Symbolic Reachability Analysis of FIFO-Channel Systems with Nonregular Sets of Configurations. *Theoretical Computer Science*, 221(1-2):211–250, 1999.
- [BHV04] A. Bouajjani, P. Habermehl, and T. Vojnar. Abstract Regular Model Checking. In *Proc. Computer Aided Verification 2004*, volume 3114 of *LNCS*, pages 372–386. Springer, 2004.
- [BR01] T. Ball and S. K. Rajamani. Automatically Validating Temporal Safety Properties of Interfaces. In *Proc. Model Checking Software, SPIN Workshop 2001*, volume 2057 of *LNCS*, pages 103–122. Springer, 2001.
- [BZ83] D. Brand and P. Zafropulo. On Communicating Finite-State Machines. *Journal of the ACM*, 30(2):323–342, 1983.
- [CGJ⁺03] E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided Abstraction Refinement for Symbolic Model Checking. *Journal of the ACM*, 50(5):752–794, 2003.
- [FIS03] A. Finkel, S. P. Iyer, and G. Sutre. Well-Abstracted Transition Systems: Application to FIFO Automata. *Information and Computation*, 181(1):1–31, 2003.
- [GS97] S. Graf and H. Säïdi. Construction of Abstract State Graphs with PVS. In *Proc. Computer Aided Verification 1997*, volume 1245 of *LNCS*, pages 72–83. Springer, 1997.
- [HJMS02] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy Abstraction. In *Proc. Symposium on Principles of Programming Languages 2002*, pages 58–70. ACM, 2002.
- [JR86] C. Jard and M. Raynal. De la Nécessité de Spécifier des Propriétés pour la Vérification des Algorithmes Distribués. Rapports de Recherche 590, IRISA Rennes, December 1986.
- [Las] Liège Automata-based Symbolic Handler (LASH). Tool Homepage. <http://www.montefiore.ulg.ac.be/~boigelot/research/lash/>.
- [LGJJ06] T. Le Gall, B. Jeannet, and T. Jérón. Verification of Communication Protocols using Abstract Interpretation of FIFO queues. In *Proc. Algebraic Methodology and Software Technology 2006*, volume 4019 of *LNCS*, pages 204–219. Springer, 2006.
- [Mcs] Model Checker for Systems of Communicating Fifo Machines (McsCM). Tool Homepage. <http://www.labri.fr/~heussner/mcscm/>.
- [Scm] Tools and Libraries for Static Analysis and Verification. Tool Homepage. <http://gforge.inria.fr/projects/bjeannet/>.
- [Tre] Tool for Reachability Analysis of Complex Systems (Trex). Tool Homepage. <http://www.liafa.jussieu.fr/~sighirea/trex/>.
- [YBCI08] F. Yu, T. Bultan, M. Cova, and O. H. Ibarra. Symbolic String Verification: An Automata-Based Approach. In *Proc. Model Checking Software, SPIN Workshop 2008*, volume 5156 of *LNCS*, pages 306–324. Springer, 2008.