

Lower Bounds for Possibly Divergent Probabilistic Programs

Shenghua Feng, Mingshuai Chen, Han Su, Benjamin L. Kaminski,
Joost-Pieter Katoen, Naijun Zhan



A FUN FACT

“A drunk man will find his way home, but a drunk bird may get lost forever.”
— Kakutani’ interpretation of [Pólya, Math. Ann. ’21]


Yet, can we show *how likely at least* can a drunk bird return home?

IN A NUTSHELL

Problem: How can we establish **lower bounds** on expected values of probabilistic programs that may exhibit **divergence**?

Example: Determine a non-trivial *lower bound on the termination probability* of the 1-D biased random walk which *diverges with non-zero probability*:


C_{brw} : `while (0 < n) { n := n - 1 [1/3] n := n + 1 }`



Challenge: Existing lower bound induction rules are confined to either (i) *bounded* random variables with a priori knowledge on the termination probability [McIver & Morgan 2005]; or (ii) (universally) *almost-surely terminating* (AST) programs [Hark et al., POPL ’20].

Our solution: The **guard-strengthening technique**. Strengthening the loop guard (and thereby the reachable state space) to forge an AST program, which can then be tackled by existing proof rules.

C_{brw}^M : `while (0 < n < M) { n := n - 1 [1/3] n := n + 1 }`



Our guard-strengthening technique makes a connection between the lower bound for C_{brw} and the lower bound for C_{brw}^M . Moreover, the strengthened program C_{brw}^M features a *stronger* termination property (i.e., becoming AST), and thus is amenable to existing induction rules.

FORMALIZATION IN THE wp-CALCULUS

The expected value of function f after program C terminates is precisely captured by *weakest preexpectations* [Kozen; McIver, Morgan; Kaminski]:

$$wp[C](f)(\sigma) \triangleq \text{Exp}[f(\tau_1) \quad f(\tau_2) \quad \dots \quad f(\tau_m)]$$

Intuitively, $wp[C](f)(\sigma)$ represents the expected value of f evaluated in the final states reached after termination of C on input σ .

The crux of (probabilistic) program verification is to reason about **while-loops**: It amounts to determining the *quantitative least fixed point* (of some monotonic operator Φ_f capturing the loop semantics w.r.t. f) which is often difficult or even impossible to compute [Kaminski et al., Acta Inform. ’19]:

$$wp[\text{while}(\varphi)\{C\}](f) = \text{lfp } \Phi_f = ?$$

LIMITATIONS OF EXISTING LOWER BOUND RULE

As computing the exact least fixed point $\text{lfp } \Phi_f$ is often intractable, researchers seek to bound it from above and/or from below:

- Upper bounds (Park induction [Park, Mach. Intel. ’69]) :

$$\Phi_f(u) \preceq u \quad \text{implies} \quad \text{lfp } \Phi_f \preceq u$$

- Lower bounds [Hark et al., POPL ’20]:

$$l \preceq \Phi_f(l) \quad \wedge \quad l \text{ is } \boxed{\text{uniformly integrable}} \quad \text{implies} \quad l \preceq \text{lfp } \Phi_f$$

almost-sure termination
bounded expectations

...

The above lower bound rule does not apply to *divergent* programs, and even for AST ones, it requires extra proof efforts in, e.g., looking for supermartingales witnessing AST and reasoning about the looping time or establishing bounds on expectations to achieve uni. int. of l .

OUR APPROACH: GUARD-STRENGTHENING

Guard-Strengthening Rule

$$C_{loop}: \text{while}(\varphi)\{C\} \quad \rightsquigarrow \quad C'_{loop}: \text{while}(\varphi')\{C\}$$

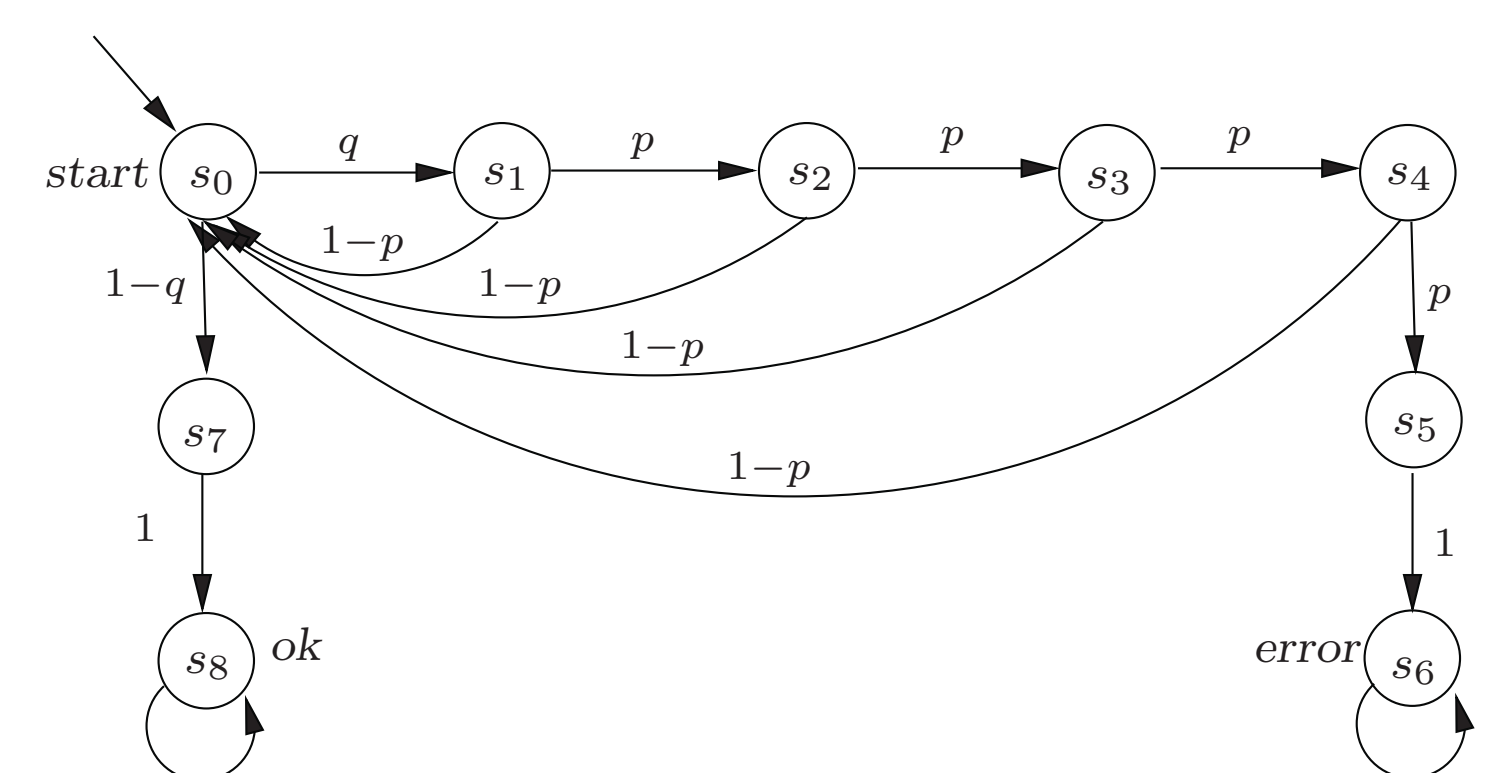
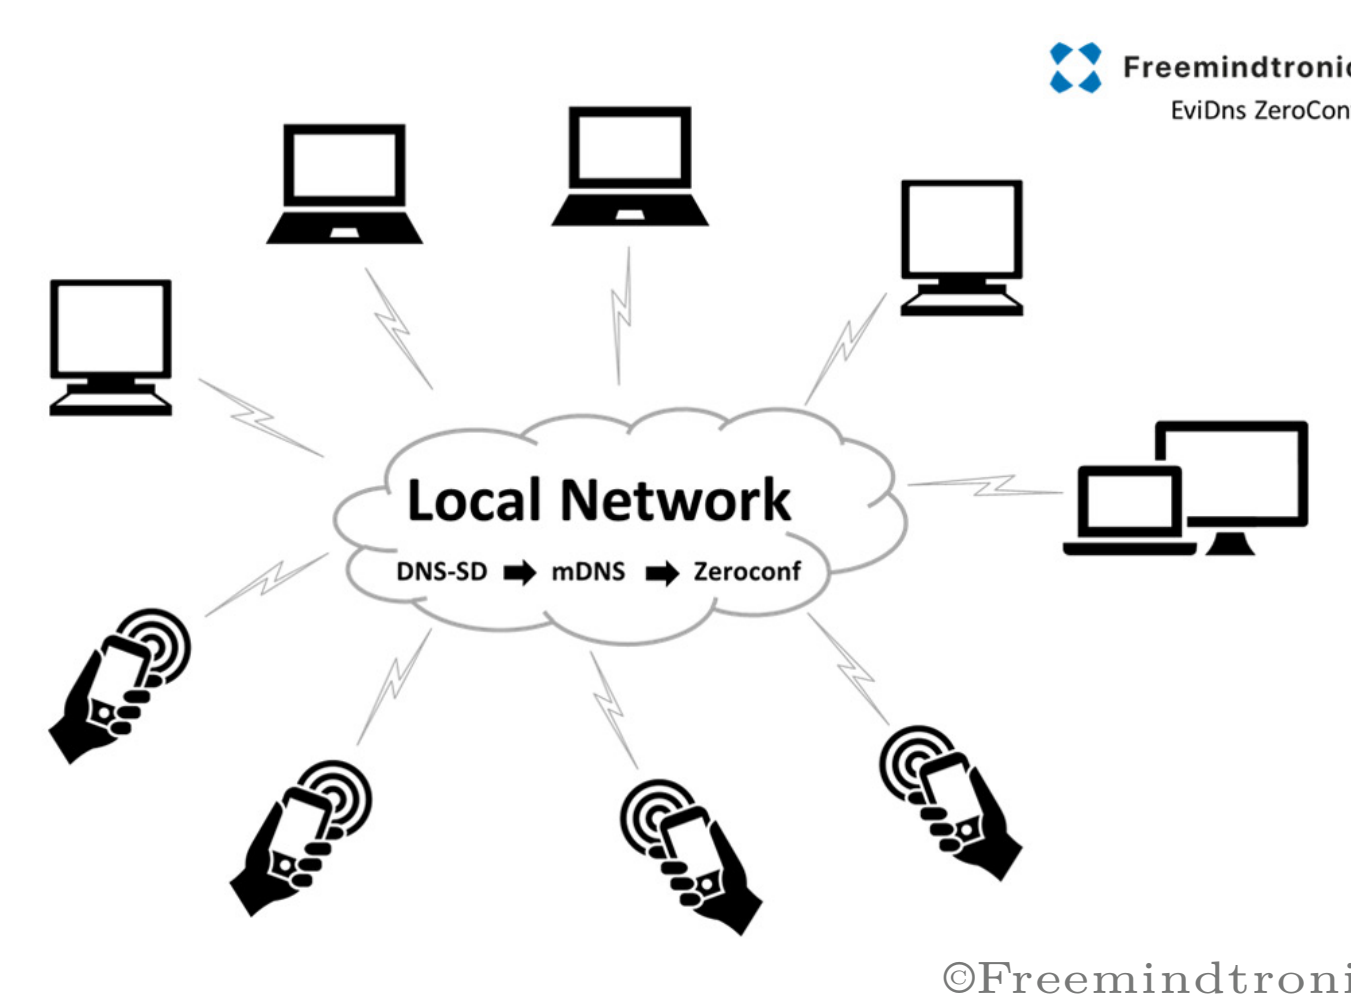
$$\frac{\varphi' \Rightarrow \varphi \quad l \preceq wp[C'_{loop}](\lceil \neg \varphi \rceil \cdot f)}{l \preceq wp[C_{loop}](f)}$$

Example: $l_M = [n < 0] + [0 \leq n \leq M] \cdot ((1/2)^n - (1/2)^M)$

$$\forall M \in \mathbb{N}: \quad l_M \stackrel{\text{Hark}}{\preceq} wp[C_{brw}^M](\lceil n \leq 0 \rceil \cdot 1) \stackrel{\text{Stren.}}{\preceq} wp[C_{brw}](1)$$

- Complete:** where there’s l , there’s φ' (albeit not “good” enough).
- General:** applicable to possibly *divergent* C_{loop} and *unbounded* f, l .
- Tight:** the underapproximation error *approaches* 0 as $\varphi' \rightarrow \varphi$.
- Automatable:** reducible to *probabilistic model checking* for finite-state C'_{loop} (which presents a solution to the drunk bird problem).

A “REAL” APPLICATION: ZEROCONF PROTOCOL



©[Baier & Katoen 2008]

Self-configuring IP network

Markov-chain snippet ($N = 4$)

C_{zc} : `start = 1; established = 0; probe = 0;`
`while (start ≤ 1 ∧ established ≤ 0 ∧ probe < N ∧ N ≥ 4) {`
`if (start = 1) {`
`{ start := 0 } [0.5] { start := 0; established := 1 }`
`} else { { probe := probe + 1 } [0.001] { start := 1; probe := 0 } }`

By strengthening the guard of C_{zc} with $N \leq 10$, we are able to establish

$$\text{Pr}(\text{“starting within the loop guard, } C_{zc} \text{ terminates with } established = 1\text{”}) \geq 0.9999999999$$