

CONVERGENCE THRESHOLDS OF NEWTON'S METHOD FOR MONOTONE POLYNOMIAL EQUATIONS

JAVIER ESPARZA, STEFAN KIEFER, AND MICHAEL LUTTENBERGER

Institut für Informatik
Technische Universität München, Germany
E-mail address: {esparza,kiefer,luttenbe}@in.tum.de

ABSTRACT. Monotone systems of polynomial equations (MSPEs) are systems of fixed-point equations $X_1 = f_1(X_1, \dots, X_n), \dots, X_n = f_n(X_1, \dots, X_n)$ where each f_i is a polynomial with positive real coefficients. The question of computing the least non-negative solution of a given MSPE $\mathbf{X} = \mathbf{f}(\mathbf{X})$ arises naturally in the analysis of stochastic models such as stochastic context-free grammars, probabilistic pushdown automata, and back-button processes. Etessami and Yannakakis have recently adapted Newton's iterative method to MSPEs. In a previous paper we have proved the existence of a threshold $k_{\mathbf{f}}$ for strongly connected MSPEs, such that after $k_{\mathbf{f}}$ iterations of Newton's method each new iteration computes at least 1 new bit of the solution. However, the proof was purely existential. In this paper we give an upper bound for $k_{\mathbf{f}}$ as a function of the minimal component of the least fixed-point $\mu\mathbf{f}$ of $\mathbf{f}(\mathbf{X})$. Using this result we show that $k_{\mathbf{f}}$ is at most single exponential resp. linear for strongly connected MSPEs derived from probabilistic pushdown automata resp. from back-button processes. Further, we prove the existence of a threshold for arbitrary MSPEs after which each new iteration computes at least $1/w2^h$ new bits of the solution, where w and h are the width and height of the DAG of strongly connected components.

1. Introduction

A *monotone system of polynomial equations* (MSPE for short) has the form

$$\begin{aligned} X_1 &= f_1(X_1, \dots, X_n) \\ &\vdots \\ X_n &= f_n(X_1, \dots, X_n) \end{aligned}$$

where f_1, \dots, f_n are polynomials with *positive* real coefficients. In vector form we denote an MSPE by $\mathbf{X} = \mathbf{f}(\mathbf{X})$. We call MSPEs “monotone” because $\mathbf{x} \leq \mathbf{x}'$ implies $\mathbf{f}(\mathbf{x}) \leq \mathbf{f}(\mathbf{x}')$ for every $\mathbf{x}, \mathbf{x}' \in \mathbb{R}_{\geq 0}^n$. MSPEs appear naturally in the analysis of many stochastic models,

2000 ACM Subject Classification: G.1.5, Mathematics of Computing, Numerical Analysis.

Key words and phrases: Newton's Method, Fixed-Point Equations, Formal Verification of Software, Probabilistic Pushdown Systems.

This work was supported by the project *Algorithms for Software Model Checking* of the *Deutsche Forschungsgemeinschaft (DFG)*. Part of this work was done at the Universität Stuttgart.

such as context-free grammars (with numerous applications to natural language processing [19, 15], and computational biology [21, 4, 3, 17]), probabilistic programs with procedures [6, 2, 10, 8, 7, 9, 11], and web-surfing models with back buttons [13, 14].

By Kleene’s theorem, a feasible MSPE $\mathbf{X} = \mathbf{f}(\mathbf{X})$ (i.e., an MSPE with at least one solution) has a least solution $\mu\mathbf{f}$; this solution can be irrational and non-expressible by radicals. Given an MSPE and a vector \mathbf{v} encoded in binary, the problem whether $\mu\mathbf{f} \leq \mathbf{v}$ holds is in PSPACE and at least as hard as the SQUARE-ROOT-SUM problem, a well-known problem of computational geometry (see [10, 12] for more details).

For the applications mentioned above the most important question is the efficient numerical approximation of the least solution. Finding the least solution of a feasible system $\mathbf{X} = \mathbf{f}(\mathbf{X})$ amounts to finding the least solution of $\mathbf{F}(\mathbf{X}) = \mathbf{0}$ for $\mathbf{F}(\mathbf{X}) = \mathbf{f}(\mathbf{X}) - \mathbf{X}$. For this we can apply (the multivariate version of) *Newton’s method* [20]: starting at some $\mathbf{x}^{(0)} \in \mathbb{R}^n$ (we use uppercase to denote variables and lowercase to denote values), compute the sequence

$$\mathbf{x}^{(k+1)} := \mathbf{x}^{(k)} - (\mathbf{F}'(\mathbf{x}^{(k)}))^{-1} \mathbf{F}(\mathbf{x}^{(k)})$$

where $\mathbf{F}'(\mathbf{X})$ is the Jacobian matrix of partial derivatives.

While in general the method may not even be defined ($\mathbf{F}'(\mathbf{x}^{(k)})$ may be singular for some k), Etesami and Yannakakis proved in [10, 12] that this is not the case for the *Decomposed Newton’s Method (DNM)*, that decomposes the MSPE into *strongly connected components* (SCCs) and applies Newton’s method to them in a bottom-up fashion¹.

The results of [10, 12] provide no information on the number of iterations needed to compute i valid bits of $\mu\mathbf{f}$, i.e., to compute a vector \mathbf{v} such that $|\mu\mathbf{f}_j - v_j| / |\mu\mathbf{f}_j| \leq 2^{-i}$ for every $1 \leq j \leq n$. In a former paper [16] we have obtained a first positive result on this problem. We have proved that for every strongly connected MSPE $\mathbf{X} = \mathbf{f}(\mathbf{X})$ there exists a threshold $k_{\mathbf{f}}$ such that for every $i \geq 0$ the $(k_{\mathbf{f}} + i)$ -th iteration of Newton’s method has at least i valid bits of $\mu\mathbf{f}$. So, loosely speaking, after $k_{\mathbf{f}}$ iterations DNM is guaranteed to compute at least 1 new bit of the solution per iteration; we say that DNM converges *linearly with rate 1*.

The problem with this result is that its proof provides no information on $k_{\mathbf{f}}$ other than its existence. In this paper we show that the threshold $k_{\mathbf{f}}$ can be chosen as

$$k_{\mathbf{f}} = 3n^2m + 2n^2 \lceil \log \mu_{\min} \rceil$$

where n is the number of equations of the MSPE, m is such that all coefficients of the MSPE can be given as ratios of m -bit integers, and μ_{\min} is the minimal component of the least solution $\mu\mathbf{f}$.

It can be objected that $k_{\mathbf{f}}$ depends on $\mu\mathbf{f}$, which is precisely what Newton’s method should compute. However, for MSPEs coming from stochastic models, such as the ones listed above, we can do far better. The following observations and results help to deal with μ_{\min} :

- We obtain a syntactic bound on μ_{\min} for probabilistic programs with procedures (having stochastic context-free grammars and back-button stochastic processes as special instances) and prove that in this case $k_{\mathbf{f}} \leq n2^{n+2}m$.

¹A subset of variables and their associated equations form an SCC, if the value of any variable in the subset influences the value of all variables in the subset, see Section 2 for details.

- We show that if every procedure has a non-zero probability of terminating, then $k_{\mathbf{f}} \leq 3nm$. This condition always holds in the special case of back-button processes [13, 14]. Hence, our result shows that i valid bits can be computed in time $\mathcal{O}((nm + i) \cdot n^3)$ in the unit cost model of Blum, Shub and Smale [1], where each single arithmetic operation over the reals can be carried out exactly and in constant time. It was proved in [13, 14] by a reduction to a semidefinite programming problem that i valid bits can be computed in $\text{poly}(i, n, m)$ -time in the classical (Turing-machine based) computation model. We do not improve this result, because we do not have a proof that round-off errors (which are inevitable on Turing-machine based models) do not crucially affect the convergence of Newton's method. But our result sheds light on the convergence of a practical method to compute $\mu\mathbf{f}$.
- Finally, since $\mathbf{x}^{(k)} \leq \mathbf{x}^{(k+1)} \leq \mu\mathbf{f}$ holds for every $k \geq 0$, as Newton's method proceeds it provides better and better lower bounds for μ_{\min} and thus for $k_{\mathbf{f}}$. In the paper we exhibit a MSPE for which, using this fact and our theorem, we can prove that no component of the solution reaches the value 1. This cannot be proved by just computing more iterations, no matter how many.

The paper contains two further results concerning non-strongly-connected MSPEs: Firstly, we show that DNM still converges linearly even if the MSPE has more than one SCC, albeit the convergence rate is poorer. Secondly, we prove that Newton's method is well-defined also for non-strongly-connected MSPEs. Thus, it is not necessary to decompose an MSPE into its SCCs – although decomposing the MSPE may be preferred for efficiency reasons.

The paper is structured as follows. In Section 2 we state preliminaries and give some background on Newton's method applied to MSPEs. Sections 3, 5, and 6 contain the three results of the paper. Section 4 contains applications of our main result. We conclude in Section 7. Missing proofs can be found in a technical report [5].

2. Preliminaries

In this section we introduce our notation and formalize the concepts mentioned in the introduction.

2.1. Notation

\mathbb{R} and \mathbb{N} denote the sets of real, respectively natural numbers. We assume $0 \in \mathbb{N}$. \mathbb{R}^n denotes the set of n -dimensional real valued *column* vectors and $\mathbb{R}_{\geq 0}^n$ the subset of vectors with non-negative components. We use bold letters for vectors, e.g. $\mathbf{x} \in \mathbb{R}^n$, where we assume that \mathbf{x} has the components x_1, \dots, x_n . Similarly, the i^{th} component of a function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is denoted by f_i .

$\mathbb{R}^{m \times n}$ denotes the set of matrices having m rows and n columns. The transpose of a vector or matrix is indicated by the superscript \top . The identity matrix of $\mathbb{R}^{n \times n}$ is denoted by Id .

The *formal Neumann series* of $A \in \mathbb{R}^{n \times n}$ is defined by $A^* = \sum_{k \in \mathbb{N}} A^k$. It is well-known that A^* exists if and only if the spectral radius of A is less than 1, i.e. $\max\{|\lambda| \mid \mathbb{C} \ni \lambda \text{ is an eigenvalue of } A\} < 1$. If A^* exists, we have $A^* = (\text{Id} - A)^{-1}$.

The partial order \leq on \mathbb{R}^n is defined as usual by setting $\mathbf{x} \leq \mathbf{y}$ if $x_i \leq y_i$ for all $1 \leq i \leq n$. By $\mathbf{x} < \mathbf{y}$ we mean $\mathbf{x} \leq \mathbf{y}$ and $\mathbf{x} \neq \mathbf{y}$. Finally, we write $\mathbf{x} \prec \mathbf{y}$ if $x_i < y_i$ in every component.

We use X_1, \dots, X_n as variable identifiers and arrange them into the vector \mathbf{X} . In the following n always denotes the number of variables, i.e. the dimension of \mathbf{X} . While $\mathbf{x}, \mathbf{y}, \dots$ denote arbitrary elements in \mathbb{R}^n , resp. $\mathbb{R}_{\geq 0}^n$, we write \mathbf{X} if we want to emphasize that a function is given w.r.t. these variables. Hence, $\mathbf{f}(\mathbf{X})$ represents the function itself, whereas $\mathbf{f}(\mathbf{x})$ denotes its value for $\mathbf{x} \in \mathbb{R}^n$.

If Y is a set of variables and \mathbf{x} a vector, then by \mathbf{x}_Y we mean the vector obtained by restricting \mathbf{x} to the components in Y .

The *Jacobian* of a differentiable function $\mathbf{f}(\mathbf{X})$ with $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is the matrix $\mathbf{f}'(\mathbf{X})$ given by

$$\mathbf{f}'(\mathbf{X}) = \begin{pmatrix} \frac{\partial f_1}{\partial X_1} & \cdots & \frac{\partial f_1}{\partial X_n} \\ \vdots & & \vdots \\ \frac{\partial f_m}{\partial X_1} & \cdots & \frac{\partial f_m}{\partial X_n} \end{pmatrix}.$$

2.2. Monotone Systems of Polynomials

Definition 2.1. A function $\mathbf{f}(\mathbf{X})$ with $\mathbf{f} : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}^n$ is a *monotone system of polynomials (MSP)*, if every component $f_i(\mathbf{X})$ is a polynomial in the variables X_1, \dots, X_n with coefficients in $\mathbb{R}_{\geq 0}$. We call an MSP $\mathbf{f}(\mathbf{X})$ *feasible* if $\mathbf{y} = \mathbf{f}(\mathbf{y})$ for some $\mathbf{y} \in \mathbb{R}_{\geq 0}^n$.

Fact 2.2. Every MSP \mathbf{f} is monotone on $\mathbb{R}_{\geq 0}^n$, i.e. for $\mathbf{0} \leq \mathbf{x} \leq \mathbf{y}$ we have $\mathbf{f}(\mathbf{x}) \leq \mathbf{f}(\mathbf{y})$.

Since every MSP is continuous, Kleene's fixed-point theorem (see e.g. [18]) applies.

Theorem 2.3 (Kleene's fixed-point theorem). *Every feasible MSP $\mathbf{f}(\mathbf{X})$ has a least fixed point $\mu\mathbf{f}$ in $\mathbb{R}_{\geq 0}^n$ i.e., $\mu\mathbf{f} = \mathbf{f}(\mu\mathbf{f})$ and, in addition, $\mathbf{y} = \mathbf{f}(\mathbf{y})$ implies $\mu\mathbf{f} \leq \mathbf{y}$. Moreover, the sequence $(\kappa_{\mathbf{f}}^{(k)})_{k \in \mathbb{N}}$ with $\kappa_{\mathbf{f}}^{(0)} := \mathbf{0}$, and $\kappa_{\mathbf{f}}^{(k+1)} := \mathbf{f}(\kappa_{\mathbf{f}}^{(k)}) = \mathbf{f}^{k+1}(\mathbf{0})$ is monotonically increasing with respect to \leq (i.e. $\kappa_{\mathbf{f}}^{(k)} \leq \kappa_{\mathbf{f}}^{(k+1)}$) and converges to $\mu\mathbf{f}$.*

In the following we call $(\kappa_{\mathbf{f}}^{(k)})_{k \in \mathbb{N}}$ the *Kleene sequence* of $\mathbf{f}(\mathbf{X})$, and drop the subscript whenever \mathbf{f} is clear from the context. Similarly, we sometimes write μ instead of $\mu\mathbf{f}$.

A variable X_i of an MSP $\mathbf{f}(\mathbf{X})$ is *productive* if $\kappa_i^{(k)} > 0$ for some $k \in \mathbb{N}$. An MSP is *clean* if all its variables are productive. It is easy to see that $\kappa_i^{(n)} = 0$ implies $\kappa_i^{(k)} = 0$ for all $k \in \mathbb{N}$. As for context-free grammars we can determine all productive variables in time linear in the size of \mathbf{f} .

Notation 2.4. In the following, we always assume that an MSP \mathbf{f} is clean and feasible. I.e., whenever we write ‘‘MSP’’, we mean ‘‘clean and feasible MSP’’, unless explicitly stated otherwise.

For the formal definition of the *Decomposed Newton's Method (DNM)* (see also Section 1) we need the notion of *dependence* between variables.

Definition 2.5. Let $\mathbf{f}(\mathbf{X})$ be an MSP. X_i *depends directly* on X_k , denoted by $X_i \trianglelefteq X_k$, if $\frac{\partial f_i}{\partial X_k}(\mathbf{X})$ is not the zero-polynomial. X_i *depends* on X_k if $X_i \trianglelefteq^* X_k$, where \trianglelefteq^* is the reflexive transitive closure of \trianglelefteq . An MSP is *strongly connected* (short: an *scMSP*) if all its variables depend on each other.

Any MSP can be decomposed into strongly connected components (SCCs), where an SCC S is a maximal set of variables such that each variable in S depends on each other variable in S . The following result for strongly connected MSPs was proved in [10, 12]:

Theorem 2.6. *Let $\mathbf{f}(\mathbf{X})$ be an scMSP and define the Newton operator $\mathcal{N}_{\mathbf{f}}$ as follows*

$$\mathcal{N}_{\mathbf{f}}(\mathbf{X}) = \mathbf{X} + (\text{Id} - \mathbf{f}'(\mathbf{X}))^{-1}(\mathbf{f}(\mathbf{X}) - \mathbf{X}).$$

We have: (1) $\mathcal{N}_{\mathbf{f}}(\mathbf{x})$ is defined for all $\mathbf{0} \leq \mathbf{x} \prec \mu\mathbf{f}$ (i.e., $(\text{Id} - \mathbf{f}'(\mathbf{x}))^{-1}$ exists). Moreover, $\mathbf{f}'(\mathbf{x})^ = \sum_{k \in \mathbb{N}} \mathbf{f}'(\mathbf{x})^k$ exists for all $\mathbf{0} \leq \mathbf{x} \prec \mu\mathbf{f}$, and so $\mathcal{N}_{\mathbf{f}}(\mathbf{X}) = \mathbf{X} + \mathbf{f}'(\mathbf{X})^*(\mathbf{f}(\mathbf{X}) - \mathbf{X})$. (2) The Newton sequence $(\boldsymbol{\nu}_{\mathbf{f}}^{(k)})_{k \in \mathbb{N}}$ with $\boldsymbol{\nu}_{\mathbf{f}}^{(k)} = \mathcal{N}_{\mathbf{f}}^k(\mathbf{0})$ is monotonically increasing, bounded from above by $\mu\mathbf{f}$ (i.e. $\boldsymbol{\nu}^{(k)} \leq \mathbf{f}(\boldsymbol{\nu}^{(k)}) \leq \boldsymbol{\nu}^{(k+1)} \prec \mu\mathbf{f}$), and converges to $\mu\mathbf{f}$.*

DNM works by substituting the variables of lower SCCs by corresponding Newton approximations that were obtained earlier.

3. A Threshold for scMSPs

In this section we obtain a threshold after which DNM is guaranteed to converge linearly with rate 1.

We showed in [16] that for worst-case results on the convergence of Newton's method it is enough to consider *quadratic* MSPs, i.e., MSPs whose monomials have degree at most 2. The reason is that any MSP (resp. scMSP) \mathbf{f} can be transformed into a quadratic MSP (resp. scMSP) $\tilde{\mathbf{f}}$ by introducing auxiliary variables. This transformation is very similar to the transformation of a context-free grammar into Chomsky normal form. The transformation does not accelerate DNM, i.e., DNM on \mathbf{f} is at least as fast (in a formal sense) as DNM on $\tilde{\mathbf{f}}$, and so for a worst-case analysis, it suffices to consider quadratic systems. We refer the reader to [16] for details.

We start by defining the notion of “valid bits”.

Definition 3.1. Let $\mathbf{f}(\mathbf{X})$ be an MSP. A vector $\boldsymbol{\nu}$ has i *valid bits* of the least fixed point $\mu\mathbf{f}$ if $|\mu\mathbf{f}_j - \nu_j| / |\mu\mathbf{f}_j| \leq 2^{-i}$ for every $1 \leq j \leq n$.

In the rest of the section we prove the following:

Theorem 3.2. *Let $\mathbf{f}(\mathbf{X})$ be a quadratic scMSP. Let c_{\min} be the smallest nonzero coefficient of \mathbf{f} and let μ_{\min} and μ_{\max} be the minimal and maximal component of $\mu\mathbf{f}$, respectively. Let*

$$k_{\mathbf{f}} = n \cdot \log \left(\frac{\mu_{\max}}{c_{\min} \cdot \mu_{\min} \cdot \min\{\mu_{\min}, 1\}} \right).$$

Then $\boldsymbol{\nu}^{(\lceil k_{\mathbf{f}} \rceil + i)}$ has i valid bits of $\mu\mathbf{f}$ for every $i \geq 0$.

Loosely speaking, the theorem states that after $k_{\mathbf{f}}$ iterations of Newton's method, every subsequent iteration guarantees at least one more valid bit. It may be objected that $k_{\mathbf{f}}$ depends on the least fixed point $\mu\mathbf{f}$, which is precisely what Newton's method should compute. However, in the next section we show that there are important classes of MSPs (in fact, those which motivated our investigation), for which bounds on μ_{\min} can be easily obtained.

The following corollary is weaker than Theorem 3.2, but less technical in that it avoids a dependence on μ_{\max} and c_{\min} .

Corollary 3.3. *Let $\mathbf{f}(\mathbf{X})$ be a quadratic scMSP of dimension n whose coefficients are given as ratios of m -bit integers. Let μ_{\min} be the minimal component of $\mu\mathbf{f}$. Let $k_{\mathbf{f}} = 3n^2m + 2n^2 \lceil \log \mu_{\min} \rceil$. Then $\boldsymbol{\nu}^{\lceil k_{\mathbf{f}} \rceil + i}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \geq 0$.*

Corollary 3.3 follows from Theorem 3.2 by a suitable bound on μ_{\max} in terms of c_{\min} and μ_{\min} [5] (notice that, since c_{\min} is the quotient of two m -bit integers, we have $c_{\min} \geq 1/2^m$).

In the rest of the section we sketch the proof of Theorem 3.2. The proof makes crucial use of vectors $\mathbf{d} \succ \mathbf{0}$ such that $\mathbf{d} \geq \mathbf{f}'(\mu\mathbf{f})\mathbf{d}$. We call a vector satisfying these two conditions a *cone vector of \mathbf{f}* or, when \mathbf{f} is clear from the context, just a cone vector.

In a previous paper we have shown that if the matrix $(\text{Id} - \mathbf{f}'(\mu\mathbf{f}))$ is singular, then \mathbf{f} has a cone vector ([16], Lemmata 4 and 8). As a first step towards the proof of Theorem 3.2 we show the following stronger proposition.

Proposition 3.4. *Any scMSP has a cone vector.*


To a cone vector $\mathbf{d} = (d_1, \dots, d_n)$ we associate two parameters, namely the maximum and the minimum of the ratios $\mu\mathbf{f}_1/d_1, \mu\mathbf{f}_2/d_2, \dots, \mu\mathbf{f}_n/d_n$, which we denote by λ_{\max} and λ_{\min} , respectively. The second step consists of showing (Proposition 3.6) that given a cone vector \mathbf{d} , the threshold $k_{\mathbf{f},\mathbf{d}} = \log(\lambda_{\max}/\lambda_{\min})$ satisfies the same property as $k_{\mathbf{f}}$ in Theorem 3.2, i.e., $\boldsymbol{\nu}^{\lceil k_{\mathbf{f},\mathbf{d}} \rceil + i}$ has i valid bits of $\mu\mathbf{f}$ for every $i \geq 0$. This follows rather easily from the following fundamental property of cone vectors: a cone vector leads to an upper bound on the error of Newton's method.

Lemma 3.5. *Let \mathbf{d} be a cone vector of an MSP \mathbf{f} and let $\lambda_{\max} = \max\{\frac{\mu\mathbf{f}_i}{d_i}\}$. Then*

$$\mu\mathbf{f} - \boldsymbol{\nu}^{(k)} \leq 2^{-k} \lambda_{\max} \mathbf{d}.$$

Proof Idea. Consider the ray $\mathbf{g}(t) = \mu\mathbf{f} - t\mathbf{d}$ starting in $\mu\mathbf{f}$ and headed in the direction $-\mathbf{d}$ (the dashed line in the picture below). It is easy to see that $\mathbf{g}(\lambda_{\max})$ is the intersection of \mathbf{g} with an axis which is located farthest from $\mu\mathbf{f}$. One can then prove $\mathbf{g}(\frac{1}{2}\lambda_{\max}) \leq \boldsymbol{\nu}^{(1)}$, where $\mathbf{g}(\frac{1}{2}\lambda_{\max})$ is the point of the ray equidistant from $\mathbf{g}(\lambda_{\max})$ and $\mu\mathbf{f}$. By repeated application of this argument one obtains $\mathbf{g}(2^{-k}\lambda_{\max}) \leq \boldsymbol{\nu}^{(k)}$ for all $k \in \mathbb{N}$.

The following picture shows the Newton iterates $\boldsymbol{\nu}^{(k)}$ for $0 \leq k \leq 2$ (shape: \times) and the corresponding points $\mathbf{g}(2^{-k}\lambda_{\max})$ (shape: $+$) located on the ray \mathbf{g} . Notice that $\boldsymbol{\nu}^{(k)} \geq \mathbf{g}(2^{-k}\lambda_{\max})$. ■



Now we easily obtain:

Proposition 3.6. *Let $\mathbf{f}(X)$ be an scMSP and let \mathbf{d} be a cone vector of \mathbf{f} . Let $k_{\mathbf{f},\mathbf{d}} = \log \frac{\lambda_{\max}}{\lambda_{\min}}$, where $\lambda_{\max} = \max_j \frac{\mu \mathbf{f}_j}{d_j}$ and $\lambda_{\min} = \min_j \frac{\mu \mathbf{f}_j}{d_j}$. Then $\nu^{(\lceil k_{\mathbf{f},\mathbf{d}} \rceil + i)}$ has at least i valid bits of $\mu \mathbf{f}$ for every $i \geq 0$.*

We now proceed to the third and final step. We have the problem that $k_{\mathbf{f},\mathbf{d}}$ depends on the cone vector \mathbf{d} , about which we only know that it exists (Proposition 3.4). We now sketch how to obtain the threshold $k_{\mathbf{f}}$ claimed in Theorem 3.2, which is independent of any cone vectors.

Consider Proposition 3.6 and let $\lambda_{\max} = \frac{\mu \mathbf{f}_i}{d_i}$ and $\lambda_{\min} = \frac{\mu \mathbf{f}_j}{d_j}$. We have $k_{\mathbf{f},\mathbf{d}} = \log \left(\frac{d_j}{d_i} \cdot \frac{\mu \mathbf{f}_i}{\mu \mathbf{f}_j} \right)$. The idea is to bound $k_{\mathbf{f},\mathbf{d}}$ in terms of c_{\min} . We show that if $k_{\mathbf{f},\mathbf{d}}$ is very large, then there must be variables X, Y such that X depends on Y only via a monomial that has a very small coefficient, which implies that c_{\min} is very small.

4. Stochastic Models

As mentioned in the introduction, several problems concerning stochastic models can be reduced to problems about the least solution $\mu \mathbf{f}$ of an MSPE \mathbf{f} . In these cases, $\mu \mathbf{f}$ is a vector of probabilities, and so $\mu_{\max} \leq 1$. Moreover, we can obtain information on μ_{\min} , which leads to bounds on the threshold $k_{\mathbf{f}}$.

4.1. Probabilistic Pushdown Automata

Our study of MSPs was initially motivated by the verification of probabilistic pushdown automata. A *probabilistic pushdown automaton* (pPDA) is a tuple $\mathcal{P} = (Q, \Gamma, \delta, Prob)$ where Q is a finite set of *control states*, Γ is a finite *stack alphabet*, $\delta \subseteq Q \times \Gamma \times Q \times \Gamma^*$ is a finite *transition relation* (we write $pX \hookrightarrow q\alpha$ instead of $(p, X, q, \alpha) \in \delta$), and $Prob$ is a function which to each transition $pX \hookrightarrow q\alpha$ assigns its probability $Prob(pX \hookrightarrow q\alpha) \in (0, 1]$ so that for all $p \in Q$ and $X \in \Gamma$ we have $\sum_{pX \hookrightarrow q\alpha} Prob(pX \hookrightarrow q\alpha) = 1$. We write $pX \xrightarrow{x} q\alpha$ instead of $Prob(pX \hookrightarrow q\alpha) = x$. A *configuration* of \mathcal{P} is a pair qw , where q is a control state and $w \in \Gamma^*$ is a *stack content*. A probabilistic pushdown automaton \mathcal{P} naturally induces a possibly infinite Markov chain with the configurations as states and transitions given by: $pX\beta \xrightarrow{x} q\alpha\beta$ for every $\beta \in \Gamma^*$ iff $pX \xrightarrow{x} q\alpha$. We assume w.l.o.g. that if $pX \xrightarrow{x} q\alpha$ is a transition then $|\alpha| \leq 2$.

pPDAs and the equivalent model of recursive Markov chains have been very thoroughly studied [6, 2, 10, 8, 7, 9, 11]. These papers have shown that the key to the analysis of pPDAs are the *termination probabilities* $[pXq]$, where p and q are states, and X is a stack letter, defined as follows (see e.g. [6] for a more formal definition): $[pXq]$ is the probability that, starting at the configuration pX , the pPDA eventually reaches the configuration $q\varepsilon$ (empty stack). It is not difficult to show that the vector of termination probabilities is the least fixed point of the MSPE containing the equation

$$[pXq] = \sum_{pX \xrightarrow{x} rYZ} x \cdot \sum_{t \in Q} [rYt] \cdot [tZq] + \sum_{pX \xrightarrow{x} rY} x \cdot [rYq] + \sum_{pX \xrightarrow{x} q\varepsilon} x$$

for each triple (p, X, q) . Call this quadratic MSPE the *termination MSPE* of the pPDA (we assume that termination MSPEs are clean, and it is easy to see that they are always

feasible). We immediately have that if $\mathbf{X} = \mathbf{f}(\mathbf{X})$ is a termination MSP, then $\mu_{\max} \leq 1$. We also obtain a lower bound on μ_{\min} :

Lemma 4.1. *Let $\mathbf{X} = \mathbf{f}(\mathbf{X})$ be a termination MSPE with n variables. Then $\mu_{\min} \geq c_{\min}^{(2^{n+1}-1)}$.*

Together with Theorem 3.2 we get the following exponential bound for $k_{\mathbf{f}}$.

Proposition 4.2. *Let \mathbf{f} be a strongly connected termination MSP with n variables and whose coefficients are expressed as ratios of m -bit numbers. Then $k_{\mathbf{f}} \leq n2^{n+2}m$.*

We conjecture that there is a lower bound on $k_{\mathbf{f}}$ which is exponential in n for the following reason. We know a family $(\mathbf{f}^{(n)})_{n=1,3,5,\dots}$ of strongly connected MSPs with n variables and irrational coefficients such that $c_{\min}^{(n)} = \frac{1}{4}$ for all n and $\mu_{\min}^{(n)}$ is double-exponentially small in n . Experiments suggest that $\Theta(2^n)$ iterations are needed for the first bit of $\mu_{\mathbf{f}^{(n)}}$, but we do not have a proof.

4.2. Strict pPDAs and Back-Button Processes

A pPDA is *strict* if for every $pX \in Q \times \Gamma$ the transition relation contains a pop-rule $pX \xrightarrow{x} q\epsilon$ for some $q \in Q$ and some $x > 0$. Essentially, strict pPDAs model programs in which every procedure has at least one terminating execution that does not call any other procedure. The termination MSP of a strict pPDA is of the form $\mathbf{b}(\mathbf{X}, \mathbf{X}) + \mathbf{lX} + \mathbf{c}$ for $\mathbf{c} \succ \mathbf{0}$. So we have $\mu_{\mathbf{f}} \geq \mathbf{c}$, which implies $\mu_{\min} \geq c_{\min}$. Together with Theorem 3.2 we get:

Proposition 4.3. *Let \mathbf{f} be a strongly connected termination MSP with n variables and whose coefficients are expressed as ratios of m -bit numbers. If \mathbf{f} is derived from a strict pPDA, then $k_{\mathbf{f}} \leq 3nm$.*

Since in most applications m is small, we obtain an excellent convergence threshold.

In [13, 14] Fagin et al. introduce a special class of strict pPDAs called *back-button processes*: in a back-button process there is only one control state p , and any rule is of the form $pA \xrightarrow{b_A} p\epsilon$ or $pA \xrightarrow{l_{AB}} pBA$. So the stack corresponds to a path through a finite graph with Γ as set of nodes and edges $A \rightarrow B$ for $pA \xrightarrow{l_{AB}} pBA$.

In [13, 14] back-button processes are used to model the behaviour of web-surfers: Γ is the set of web-pages, l_{AB} is the probability that a web-surfer uses a link from page A to page B , and b_A is the probability that the surfer pushes the “back”-button of the web-browser while visiting A . Thus, the termination probability $[pAp]$ is simply the probability that, if A is on top of the stack, A is eventually popped from the stack. The termination probabilities are the least solution of the MSPE consisting of the equations

$$[pAp] = b_A + \sum_{pA \xrightarrow{l_{AB}} pBA} l_{AB}[pBp][pAp] = b_A + [pAp] \sum_{pA \xrightarrow{l_{AB}} pBA} l_{AB}[pBp].$$

4.3. An Example

As an example of application of Theorem 3.2 consider the following scMSPE $\mathbf{X} = \mathbf{f}(\mathbf{X})$.

$$\begin{pmatrix} X_1 \\ X_2 \\ X_3 \end{pmatrix} = \begin{pmatrix} 0.4X_2X_1 + 0.6 \\ 0.3X_1X_2 + 0.4X_3X_2 + 0.3 \\ 0.3X_1X_3 + 0.7 \end{pmatrix}$$

The least solution of the system gives the revocation probabilities of a back-button process with three web-pages. For instance, if the surfer is at page 2 it can choose between following links to pages 1 and 3 with probabilities 0.3 and 0.4, respectively, or pressing the back button with probability 0.3.

We wish to know if any of the revocation probabilities is equal to 1. Performing 14 Newton steps (e.g. with Maple) yields an approximation $\boldsymbol{\nu}^{(14)}$ to the termination probabilities with

$$\begin{pmatrix} 0.98 \\ 0.97 \\ 0.992 \end{pmatrix} \leq \boldsymbol{\nu}^{(14)} \leq \begin{pmatrix} 0.99 \\ 0.98 \\ 0.993 \end{pmatrix}.$$

We have $c_{\min} = 0.3$. In addition, since Newton's method converges to $\mu\mathbf{f}$ from below, we know $\mu_{\min} \geq 0.97$. Moreover, $\mu_{\max} \leq 1$, as $\mathbf{1} = \mathbf{f}(\mathbf{1})$ and so $\mu\mathbf{f} \leq \mathbf{1}$. Hence $k_{\mathbf{f}} \leq 3 \cdot \log \frac{1}{0.97 \cdot 0.3 \cdot 0.97} \leq 6$. Theorem 3.2 then implies that $\boldsymbol{\nu}^{(14)}$ has (at least) 8 valid bits of $\mu\mathbf{f}$. As $\mu\mathbf{f} \leq \mathbf{1}$, the absolute errors are bounded by the relative errors, and since $2^{-8} \leq 0.004$ we know:

$$\mu\mathbf{f} \prec \boldsymbol{\nu}^{(14)} + \begin{pmatrix} 2^{-8} \\ 2^{-8} \\ 2^{-8} \end{pmatrix} \prec \begin{pmatrix} 0.994 \\ 0.984 \\ 0.997 \end{pmatrix} \prec \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

So Theorem 3.2 gives a proof that all 3 revocation probabilities are strictly smaller than 1.

5. Linear Convergence of the Decomposed Newton's Method

Given a strongly connected MSP \mathbf{f} , Theorem 3.2 states that, if we have computed $k_{\mathbf{f}}$ preparatory iterations of Newton's method, then after i additional iterations we can be sure to have computed at least i bits of $\mu\mathbf{f}$. We call this linear convergence with rate 1. Now we show that DNM, which handles non-strongly-connected MSPs, converges linearly as well. We also give an explicit convergence rate.

Let $\mathbf{f}(\mathbf{X})$ be any quadratic MSP (again we assume *quadratic* MSPs throughout this section), and let $h(\mathbf{f})$ denote the height of the DAG of strongly connected components (SCCs). The convergence rate of DNM crucially depends on this height: In the worst case one needs asymptotically $\Theta(2^{h(\mathbf{f})})$ iterations in each component per bit, assuming one performs the same number of iterations in each component.

To get a sharper result, we suggest to perform a different number of iterations in each SCC, depending on its *depth*. The depth of an SCC S is the length of the longest path in the DAG of SCCs from S to a top SCC.

In addition, we use the following notation. For a depth t , we denote by $\text{comp}(t)$ the set of SCCs of depth t . Furthermore we define $C(t) := \bigcup \text{comp}(t)$ and $C_{>}(t) := \bigcup_{t' > t} C(t')$ and, analogously, $C_{<}(t)$. We will sometimes write \mathbf{v}_t for $\mathbf{v}_{C(t)}$ and $\mathbf{v}_{>t}$ for $\mathbf{v}_{C_{>}(t)}$ and $\mathbf{v}_{<t}$ for $\mathbf{v}_{C_{<}(t)}$, where \mathbf{v} is any vector.

Figure 1 shows the Decomposed Newton's Method (DNM) for computing an approximation $\boldsymbol{\nu}$ for $\mu\mathbf{f}$, where $\mathbf{f}(\mathbf{X})$ is any quadratic MSP. The authors of [10] recommend to

run Newton’s Method in each SCC S until “approximate solutions for S are considered ‘good enough’”. Here we suggest to run Newton’s Method in each SCC S for a number of steps that depends (exponentially) on the depth of S and (linearly) on a parameter j that controls the number of iterations (see Figure 1).

```

function DNM ( $\mathbf{f}, j$ )
/* The parameter  $j$  controls the number of iterations. */
for  $t$  from  $h(\mathbf{f})$  downto 0
  forall  $S \in \text{comp}(t)$  /* all SCCs  $S$  of depth  $t$  */
     $\nu_S := \mathcal{N}_{\mathbf{f}_S}^{j \cdot 2^t}(\mathbf{0})$  /*  $j \cdot 2^t$  iterations */
    /* apply  $\nu_S$  in the depending SCCs */
     $\mathbf{f}_{<t}(\mathbf{X}) := \mathbf{f}_{<t}(\mathbf{X})[\mathbf{X}_S/\nu_S]$ 
return  $\nu$ 

```

Figure 1: Decomposed Newton’s Method (DNM) for computing an approximation ν of $\mu\mathbf{f}$

Recall that $h(\mathbf{f})$ was defined as the height of the DAG of SCCs. Similarly we define the width $w(\mathbf{f})$ to be $\max_t |\text{comp}(t)|$. Notice that \mathbf{f} has at most $(h(\mathbf{f}) + 1) \cdot w(\mathbf{f})$ SCCs. We have the following bound on the number of iterations run by DNM.

Proposition 5.1. *The function $\text{DNM}(\mathbf{f}, j)$ of Fig. 1 runs at most $j \cdot w(\mathbf{f}) \cdot 2^{h(\mathbf{f})+1}$ iterations of Newton’s method.*

We will now analyze the convergence behavior of DNM asymptotically (for large j). Let $\Delta_S^{(j)}$ denote the error in S when running DNM with parameter j , i.e., $\Delta_S^{(j)} := \mu_S - \nu_S^{(j)}$. Observe that the error $\Delta_t^{(j)}$ can be understood as the sum of two errors:

$$\Delta_t^{(j)} = \mu_t - \nu_t^{(j)} = (\mu_t - \widetilde{\mu}_t^{(j)}) + (\widetilde{\mu}_t^{(j)} - \nu_t^{(j)}),$$

where $\widetilde{\mu}_t^{(j)} := \mu(\mathbf{f}_t(\mathbf{X})[\mathbf{X}_{>t}/\nu_{>t}^{(j)}])$, i.e., $\widetilde{\mu}_t^{(j)}$ is the least fixed point of \mathbf{f}_t after the approximations from the lower SCCs have been applied. So, $\Delta_t^{(j)}$ consists of the *propagation error* $(\mu_t - \widetilde{\mu}_t^{(j)})$ and the newly inflicted *approximation error* $(\widetilde{\mu}_t^{(j)} - \nu_t^{(j)})$.

The following lemma, technically non-trivial to prove, gives a bound on the propagation error.

Lemma 5.2 (Propagation error). *Let $\nu_{>t}$ be some approximation of $\mu_{>t}$, i.e., $\mathbf{0} \leq \nu_{>t} \leq \mu_{>t}$. Let $\widetilde{\mu}_t = \mu(\mathbf{f}_t(\mathbf{X})[\mathbf{X}_{>t}/\nu_{>t}])$. Then there is a constant $c > 0$ such that*

$$\|\mu_t - \widetilde{\mu}_t\| \leq c \cdot \sqrt{\|\mu_{>t} - \nu_{>t}\|}.$$

Intuitively, Lemma 5.2 states that if $\nu_{>t}$ has k valid bits of $\mu_{>t}$, then $\widetilde{\mu}_t$ has roughly $k/2$ valid bits of μ_t . In other words, (at most) one half of the valid bits are lost on each level of the DAG due to the propagation error.

The following theorem assures that after combining the propagation error and the approximation error, DNM still converges linearly.

Theorem 5.3. *Let \mathbf{f} be a quadratic MSP. Let $\nu^{(j)}$ denote the result of calling $\text{DNM}(\mathbf{f}, j)$ (see Figure 1). Then there is a $k_{\mathbf{f}} \in \mathbb{N}$ such that $\nu^{(k_{\mathbf{f}}+i)}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \geq 0$.*

We conclude that increasing i by one gives us asymptotically at least one additional bit in each component and, by Proposition 5.1, costs $w(\mathbf{f}) \cdot 2^{h(\mathbf{f})+1}$ additional Newton iterations.

In the technical report [5] we give an example that shows that the bound above is essentially optimal in the sense that an exponential (in $h(\mathbf{f})$) number of iterations is in general needed to obtain an additional bit.

6. Newton's Method for General MSPs

Etessami and Yannakakis [10] introduced DNM because they could show that the matrix inverses used by Newton's method exist if Newton's method is run on each SCC separately (see Theorem 2.6).

It may be surprising that the matrix inverses used by Newton's method exist even if the MSP is *not* decomposed. More precisely one can show the following theorem, see [5].

Theorem 6.1. *Let $\mathbf{f}(\mathbf{X})$ be any MSP, not necessarily strongly connected. Let the Newton operator $\mathcal{N}_{\mathbf{f}}$ be defined as before:*

$$\mathcal{N}_{\mathbf{f}}(\mathbf{X}) = \mathbf{X} + (\text{Id} - \mathbf{f}'(\mathbf{X}))^{-1}(\mathbf{f}(\mathbf{X}) - \mathbf{X})$$

Then the Newton sequence $(\boldsymbol{\nu}_{\mathbf{f}}^{(k)})_{k \in \mathbb{N}}$ with $\boldsymbol{\nu}^{(k)} = \mathcal{N}_{\mathbf{f}}^k(\mathbf{0})$ is well-defined (i.e., the matrix inverses exist), monotonically increasing, bounded from above by $\mu\mathbf{f}$ (i.e. $\boldsymbol{\nu}^{(k)} \leq \boldsymbol{\nu}^{(k+1)} \prec \mu\mathbf{f}$), and converges to $\mu\mathbf{f}$.

By exploiting Theorem 5.3 and Theorem 6.1 one can show the following theorem which addresses the convergence *speed* of Newton's Method in general.

Theorem 6.2. *Let \mathbf{f} be any quadratic MSP. Then the Newton sequence $(\boldsymbol{\nu}^{(k)})_{k \in \mathbb{N}}$ is well-defined and converges linearly to $\mu\mathbf{f}$. More precisely, there is a $k_{\mathbf{f}} \in \mathbb{N}$ such that $\boldsymbol{\nu}^{(k_{\mathbf{f}}+i \cdot (h(\mathbf{f})+1) \cdot 2^{h(\mathbf{f})})}$ has at least i valid bits of $\mu\mathbf{f}$ for every $i \geq 0$.*

Again, the $2^{h(\mathbf{f})}$ factor cannot be avoided in general as shown by an example in [5].

7. Conclusions

We have proved a threshold $k_{\mathbf{f}}$ for strongly connected MSPEs. After $k_{\mathbf{f}}+i$ Newton iterations we have i bits of accuracy. The threshold $k_{\mathbf{f}}$ depends on the representation size of \mathbf{f} and on the least solution $\mu\mathbf{f}$. Although this latter dependence might seem to be a problem, lower and upper bounds on $\mu\mathbf{f}$ can be easily derived for stochastic models (probabilistic programs with procedures, stochastic context-free grammars and back-button processes). In particular, this allows us to show that $k_{\mathbf{f}}$ depends linearly on the representation size for back-button processes. We have also shown by means of an example that the threshold $k_{\mathbf{f}}$ improves when the number of iterations increases.

In [16] we left the problem whether DNM converges linearly for non-strongly-connected MSPEs open. We have proven that this is the case, although the convergence rate is poorer: if h and w are the height and width of the graph of SCCs of \mathbf{f} , then there is a threshold $\tilde{k}_{\mathbf{f}}$ such that $\tilde{k}_{\mathbf{f}} + i \cdot w \cdot 2^{h+1}$ iterations of DNM compute at least i valid bits of $\mu\mathbf{f}$, where the exponential factor cannot be avoided in general.

Finally, we have shown that the Jacobian of the whole MSPE is guaranteed to exist, whether the MSPE is strongly connected or not.

Acknowledgments.

The authors wish to thank Kousha Etessami and anonymous referees for very valuable comments.

References

- [1] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the Amer. Math. Society*, 21(1):1–46, 1989.
- [2] T. Brázdil, A. Kučera, and O. Stražovský. On the decidability of temporal properties of probabilistic pushdown automata. In *Proceedings of STACS'2005*, volume 3404 of *LNCS*, pages 145–157. Springer, 2005.
- [3] R.D. Dowell and S.R. Eddy. Evaluation of several lightweight stochastic context-free grammars for RNA secondary structure prediction. *BMC Bioinformatics*, 5(71), 2004.
- [4] R. Durbin, S.R. Eddy, A. Krogh, and G.J. Michison. *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, 1998.
- [5] J. Esparza, S. Kiefer, and M. Luttenberger. Convergence thresholds of Newton’s method for monotone polynomial equations. Technical report, Technische Universität München, 2007.
- [6] J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proceedings of LICS 2004*, pages 12–21, 2004.
- [7] J. Esparza, A. Kučera, and R. Mayr. Quantitative analysis of probabilistic pushdown automata: Expectations and variances. In *Proceedings of LICS 2005*, pages 117–126. IEEE Computer Society Press, 2005.
- [8] K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic systems. In *Proceedings of TACAS 2005*, LNCS 3440, pages 253–270. Springer, 2005.
- [9] K. Etessami and M. Yannakakis. Checking LTL properties of recursive Markov chains. In *Proceedings of 2nd Int. Conf. on Quantitative Evaluation of Systems (QEST'05)*, 2005.
- [10] K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. In *STACS*, pages 340–352, 2005.
- [11] K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proceedings of ICALP 2005*, volume 3580 of *LNCS*, pages 891–903. Springer, 2005.
- [12] K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations, 2006. Draft journal submission, http://homepages.inf.ed.ac.uk/kousha/bib_index.html.
- [13] R. Fagin, A.R. Karlin, J. Kleinberg, P. Raghavan, S. Rajagopalan, R. Rubinfeld, M. Sudan, and A. Tomkins. Random walks with “back buttons” (extended abstract). In *STOC*, pages 484–493, 2000.
- [14] R. Fagin, A.R. Karlin, J. Kleinberg, P. Raghavan, S. Rajagopalan, R. Rubinfeld, M. Sudan, and A. Tomkins. Random walks with “back buttons”. *Annals of Applied Probability*, 11(3):810–862, 2001.
- [15] S. Geman and M. Johnson. Probabilistic grammars and their applications, 2002.
- [16] S. Kiefer, M. Luttenberger, and J. Esparza. On the convergence of Newton’s method for monotone systems of polynomial equations. In *Proceedings of STOC*, pages 217–226. ACM, 2007.
- [17] B. Knudsen and J. Hein. Pfold: RNA secondary structure prediction using stochastic context-free grammars. *Nucleic Acids Research*, 31(13):3423–3428, 2003.
- [18] W. Kuich. *Handbook of Formal Languages*, volume 1, chapter 9: Semirings and Formal Power Series: Their Relevance to Formal Languages and Automata, pages 609 – 677. Springer, 1997.
- [19] C. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, 1999.
- [20] J.M. Ortega and W.C. Rheinboldt. *Iterative solution of nonlinear equations in several variables*. Academic Press, 1970.
- [21] Y. Sakabikara, M. Brown, R. Hughey, I.S. Mian, K. Sjolander, R.C. Underwood, and D. Haussler. Stochastic context-free grammars for tRNA. *Nucleic Acids Research*, 22:5112–5120, 1994.