In A. Kuijper, K. Bredies, T. Pock, H. Bischof (Eds.): Scale-Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science, Vol. 7893, pp. 368-379, Springer, Berlin, 2013. The final publication is available at link.springer.com.

A Fully Discrete Theory for Linear Osmosis Filtering

Oliver Vogel, Kai Hagenburg, Joachim Weickert, and Simon Setzer

Mathematical Image Analysis Group
Dept. of Mathematics and Computer Science, Campus E1.7
Saarland University, 66041 Saarbrücken, Germany
{vogel,hagenburg,weickert,setzer}@mia.uni-saarland.de
http://www.mia.uni-saarland.de

Abstract. Osmosis filters are based on drift–diffusion processes. They offer nontrivial steady states with a number of interesting applications. In this paper we present a fully discrete theory for linear osmosis filtering that follows the structure of Weickert's discrete framework for diffusion filters. It regards the positive initial image as a vector and expresses its evolution in terms of iterative matrix–vector multiplications. The matrix differs from its diffusion counterpart by the fact that it is unsymmetric. We assume that it satisfies four properties: vanishing column sums, nonnegativity, irreducibility, and positive diagonal elements. Then the resulting filter class preserves the average grey value and the positivity of the solution. Using the Perron–Frobenius theory we prove that the process converges to the unique eigenvector of the iteration matrix that is positive and has the same average grey value as the initial image. We show that our theory is directly applicable to explicit and implicit finite difference discretisations. We establish a stability condition for the explicit scheme, and we prove that the implicit scheme is absolutely stable. Both schemes converge to a steady state that solves the discrete elliptic equation. This steady state can be reached efficiently when the implicit scheme is equipped with a BiCGStab solver.

Keywords: osmosis filtering, drift–diffusion, finite difference methods, BiCGStab

1 Introduction

Osmosis filtering relies on the idea of making diffusion filters unsymmetric. This is achieved by supplementing it with a drift term that allows nontrivial steady states. While specific applications of this idea to the fields of digital halftoning and numerical methods for hyperbolic conservation laws can be found in two earlier publications [1, 2], the first comprehensive description of osmosis models for a variety of visual computing applications is presented in our companion paper [3]. In [3] we demonstrate that osmosis models are powerful tools for compact data representation, for editing an existing image, and for fusing information from different images. Most of these applications go far beyond of what can be achieved with nonlinear diffusion filters, in spite of the fact that the osmosis models in [3] are linear. Osmosis filters have some similarities to gradient domain methods from computer graphics [4, 5], but offer additional advantages such as invariance under multiplicative illumination changes.

Since osmosis can be interpreted as a modification of diffusion filtering and there is a well-established theory for diffusion filters, it is natural to study which results can be generalised from diffusion to osmosis. The goal of the present paper is to provide a fully discrete theory for linear osmosis filtering that has a similar structure as Weickert's discrete framework for diffusion filters [6]. We will see that this theory offers some fundamental differences to diffusion filters, and that it is applicable to the design of osmosis algorithms that are not only reliable, but also efficient.

Our paper is organised as follows. In Section 2 we review the basic structure of continuous osmosis filters, and we consider finite difference discretisations in space and time. This leads us to fully discrete osmosis filters that can be expressed as iterative matrix–vector multiplications. Section 3 provides our theoretical framework for this filter class, in which we establish useful properties such as preservation of positivity and convergence results. In Section 4 we apply this theory to two popular finite difference discretisations: an explicit and an implicit scheme. The performance of these schemes is evaluated in Section 5, and a summary in Section 6 concludes our paper.

2 From Continuous to Discrete Osmosis

Before we can introduce a theory for discrete linear osmosis processes in visual computing, we have to discuss the continuous concept first and show how it can be turned into a discrete filter representation. This is the topic of the present section.

2.1 Continuous Linear Osmosis Filtering

Let us consider a rectangular image domain $\Omega \subset \mathbb{R}^2$ with boundary $\partial\Omega$, and a positive greyscale image $f : \Omega \to \mathbb{R}_+$. Moreover, assume we are given some *drift vector field* $\boldsymbol{d} : \Omega \mapsto \mathbb{R}^2$. Then a (linear) *osmosis filter* computes a processed version $u(\boldsymbol{x}, t)$ of $f(\boldsymbol{x})$ by solving the drift-diffusion PDE

$$\partial_t u \;=\; \Delta u \;-\; \operatorname{div}(\boldsymbol{d}u) \qquad \text{on } \Omega \times (0, T], \tag{1}$$

with f as initial condition,

$$u(\boldsymbol{x},0) \;=\; f(\boldsymbol{x}) \qquad \text{on } \Omega, \tag{2}$$

and homogeneous Neumann boundary conditions. They specify a vanishing flux in normal direction $\boldsymbol{n}$ to the image boundary $\partial\Omega$:

$$\langle \boldsymbol{\nabla} u - \boldsymbol{d}u, \boldsymbol{n}\rangle \;=\; 0 \qquad \text{on } \partial\Omega \times (0,T]. \tag{3}$$

Let us now sketch three key properties of our osmosis model [3]:

(a) **Preservation of the Average Grey Value:**
Since the osmosis process is in divergence form, its solution preserves the average grey value of the initial image:

$$\frac{1}{|\Omega|}\int_\Omega u(\boldsymbol{x},t)\,d\boldsymbol{x} \;=\; \frac{1}{|\Omega|}\int_\Omega f(\boldsymbol{x})\,d\boldsymbol{x} \qquad \forall\, t>0\,. \tag{4}$$

This property can also be found for diffusion filters.

(b) **Preservation of Positivity:**
One can show that the solution remains positive for all times:

$$u(\boldsymbol{x},t) > 0 \qquad \forall \boldsymbol{x} \in \Omega, \quad \forall t > 0. \tag{5}$$

This is a weaker property than the maximum–minimum principle for diffusion [6]. Osmosis may violate a maximum–minimum principle.

(c) **Convergence to a Nontrivial Steady State:**
The continuous linear osmosis model differs from a homogeneous diffusion filter only by its drift term. However, the drift vector field $\boldsymbol{d}$ is a powerful tool to steer its convergence: If $\boldsymbol{d}$ satisfies

$$\boldsymbol{d} \;=\; \boldsymbol{\nabla}(\ln v) \;=\; \frac{\boldsymbol{\nabla} v}{v} \tag{6}$$

with some positive image v, one can show that the osmosis process converges to v up to a multiplicative constant which ensures preservation of the average grey value of f. Thus, osmosis creates nontrivial steady states. This is a fundamental difference to diffusion that allows only flat steady states [6].

Since $\boldsymbol{d}$ contains the gradient information of $\ln v$, we may regard osmosis as a process for data integration. In that sense it resembles so-called gradient domain methods that are popular in computer graphics [4, 5]. Therefore, it is not surprising that it can also be used for similar applications such as image editing and image fusion. We refer to our companion paper [3] for such applications. Other applications are concerned with alternative numerical schemes for hyperbolic conservation laws [2]. Moreover, also the PDE limit of a lattice Boltzmann model for halftoning [1] is an osmosis equation.

Applying osmosis to colour images is as simple as applying it to greyscale images: One proceeds separately in each RGB channel using the individual drift vector fields of each channel.

2.2 Finite Difference Discretisation

Let us now consider a finite difference space discretisation of the drift–diffusion equation (1). We consider a grid size h in x- and y-direction, and we denote by $u_{i,j}$ an approximation to u in the grid point $((i-\frac{1}{2})h, (j-\frac{1}{2})h))^\top$. Setting $\boldsymbol{d} = (d_1, d_2)^\top$, we approximate (1) by

$$\begin{aligned} u'_{i,j} \; = \; & \frac{u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1} - 4u_{i,j}}{h^2} - \frac{1}{h}\left(d_{1,i+\frac{1}{2},j} \frac{u_{i+1,j}+u_{i,j}}{2} \right. \\ & \left. - \, d_{1,i-\frac{1}{2},j} \frac{u_{i,j}+u_{i-1,j}}{2} \right) - \frac{1}{h}\left(d_{2,i,j+\frac{1}{2}} \frac{u_{i,j+1}+u_{i,j}}{2} - d_{2,i,j-\frac{1}{2}} \frac{u_{i,j}+u_{i,j-1}}{2} \right) \end{aligned} \tag{7}$$

This also holds for boundary points, if we mirror the image at its boundaries and assume a zero drift vector across boundaries. Rearranging (7) gives

$$\begin{aligned} u'_{i,j} = {} & u_{i+1,j} \left(\frac{1}{h^2} - \frac{d_{1,i+\frac{1}{2},j}}{2h} \right) + u_{i-1,j} \left(\frac{1}{h^2} + \frac{d_{1,i-\frac{1}{2},j}}{2h} \right) \\ & + u_{i,j+1} \left(\frac{1}{h^2} - \frac{d_{2,i,j+\frac{1}{2}}}{2h} \right) + u_{i,j-1} \left(\frac{1}{h^2} + \frac{d_{2,i,j-\frac{1}{2}}}{2h} \right) \\ & + u_{i,j} \left(-\frac{4}{h^2} - \frac{d_{1,i+\frac{1}{2},j}}{2h} + \frac{d_{1,i-\frac{1}{2},j}}{2h} - \frac{d_{2,i,j+\frac{1}{2}}}{2h} + \frac{d_{2,i,j-\frac{1}{2}}}{2h} \right). \end{aligned} \tag{8}$$

From now on we restrict ourselves to drift vector fields $(d_1(\boldsymbol{x}), d_2(\boldsymbol{x}))^\top$ with

$$|d_1(\boldsymbol{x})| < \frac{2}{h}, \qquad |d_2(\boldsymbol{x})| < \frac{2}{h} \qquad \forall\, \boldsymbol{x} \in \Omega. \tag{9}$$

This ensures that in (8) the weights of all four neighbours of $u_{i,j}$ are positive. We want to write this discretisation in a more compact notation. To this end, we replace the double indexing in each pixel by a single index and assemble all unknown grey values in a single vector $\boldsymbol{u} \in \mathbb{R}^N$ where N denotes the number of pixels. Then we end up with the following dynamical system:

$$\boldsymbol{u}(0) = \boldsymbol{f}, \tag{10}$$

$$\boldsymbol{u}'(t) = \boldsymbol{A}\, \boldsymbol{u}(t) \tag{11}$$

where the matrix $\boldsymbol{A} \in \mathbb{R}^{N \times N}$ is unsymmetric. This differs from the diffusion scenario that leads to symmetric matrices [6]. Since the weights of the neighbours in (8) are positive, it follows that $\boldsymbol{A}$ has nonnegative off-diagonals. Moreover, one can show that all column sums of $\boldsymbol{A}$ are zero and $\boldsymbol{A}$ is irreducible.

We have different options to discretise this ODE system in time. In the simplest case one can consider the *explicit scheme*:

$$\frac{\boldsymbol{u}^{k+1} - \boldsymbol{u}^k}{\tau} \;=\; \boldsymbol{A}\boldsymbol{u}^k \tag{12}$$

where $\tau > 0$ denotes the time step size, and the upper index k refers to an approximation at time $k\tau$. With $\boldsymbol{P} := \boldsymbol{I} + \tau \boldsymbol{A}$, we can rearrange this scheme to

$$\boldsymbol{u}^{k+1} \;=\; \boldsymbol{P}\boldsymbol{u}^{k} \tag{13}$$

An alternative time discretisation is given by the *implicit scheme*

$$\frac{\boldsymbol{u}^{k+1} - \boldsymbol{u}^{k}}{\tau} \;=\; \boldsymbol{A}\boldsymbol{u}^{k+1}. \tag{14}$$

It requires to solve a linear system in the unknown vector $\boldsymbol{u}^{k+1}$. If the system matrix is invertible, the problem can also be formally written as a matrix–vector multiplication of type (13) with $\boldsymbol{P} := (\boldsymbol{I} - \tau \boldsymbol{A})^{-1}$.

3 A Discrete Osmosis Theory

We have seen that both the explicit and the implicit scheme are examples of numerical methods that can be written in the general form (13). This motivates us to derive a general theory for discrete osmosis processes of this type. Here is our main result.

Proposition 1. [Theory for Discrete Linear Osmosis]
Let $\boldsymbol{f} \in \mathbb{R}_{+}^{N}$ and consider a process

$$\boldsymbol{u}^{0} = \boldsymbol{f}, \tag{15}$$

$$\boldsymbol{u}^{k+1} = \boldsymbol{P}\boldsymbol{u}^{k} \qquad (k = 0, 1, ...) \tag{16}$$

where the (unsymmetric) matrix $\boldsymbol{P} \in \mathbb{R}^{N \times N}$ satisfies the following properties:

(DLO1) All column sums of $\boldsymbol{P}$ are 1.
(DLO2) $\boldsymbol{P}$ is nonnegative.
(DLO3) $\boldsymbol{P}$ is irreducible.
(DLO4) $\boldsymbol{P}$ has only positive diagonal entries.

Then the following results hold:

(a) The average grey value is preserved:

$$\frac{1}{N}\sum_{i=1}^{N} u_i^{k} \;=\; \frac{1}{N}\sum_{i=1}^{N} f_i \qquad \forall\, k > 0\,. \tag{17}$$

(b) The evolution preserves positivity:

$$u_i^{k} > 0 \qquad \forall i \in \{1, .., N\}, \quad \forall\, k > 0\,. \tag{18}$$

(c) There exists a unique steady state for $k \to \infty$. It is given by the eigenvector $\boldsymbol{v} \in \mathbb{R}_{+}^{N}$ of $\boldsymbol{P}$ to the eigenvalue 1, that has the same average grey value as $\boldsymbol{f}$.

Proof. Average grey value invariance and preservation of positivity are very easily seen, while the convergence result requires some more technicalities.

(a) Average grey value invariance for osmosis has already been shown in [2], where the reasoning is identical to the diffusion case [6, Proposition 4]:

(b) In order to verify preservation of positivity, we observe that applying one osmosis step to the positive initial image $\boldsymbol{f}$ gives

$$u_i^1 \;=\; \underbrace{p_{i,i}}_{>0}\,\underbrace{f_i}_{>0} \;+\; \sum_{\substack{j=1\\ j\neq i}}^{N} \underbrace{p_{i,j}}_{\geq 0}\,\underbrace{f_j}_{>0} \;>\; 0 \qquad \forall\, i \in \{1,...,N\}\,. \tag{19}$$

Applying this reasoning iteratively ensures that $\boldsymbol{u}^k$ is positive for all $k > 0$.

(c) To establish our convergence result, first we show that 1 is an eigenvalue of $\boldsymbol{P}$. As the eigenvalues of $\boldsymbol{P}$ and $\boldsymbol{P}^\top$ are identical, we can exploit the unit row sum of $\boldsymbol{P}^\top$ instead of the unit column sum of $\boldsymbol{P}$. Hence, we can compute

$$\boldsymbol{P}^\top \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} = \begin{pmatrix} \sum_{j=1}^N p_{1,j} \\ \sum_{j=1}^N p_{2,j} \\ \vdots \\ \sum_{j=1}^N p_{N,j} \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}. \tag{20}$$

Thus, 1 is an eigenvalue of $\boldsymbol{P}^\top$ and therefore also of $\boldsymbol{P}$. Note that $(1, 1, \ldots, 1)^\top$ is an eigenvector for $\boldsymbol{P}^\top$, but not for $\boldsymbol{P}$.

Next we prove that all eigenvalues $\lambda \in \mathbb{C}$ with $\lambda \neq 1$ satisfy $|\lambda| < 1$. Since $\boldsymbol{P}$ has unit column sums, its column sum norm satisfies $\|\boldsymbol{P}\|_1 = 1$. Thus, we have $|\lambda| \leq 1$. By Gershgorin's theorem, all eigenvalues of $\boldsymbol{P}$ lie within disks in the complex domain whose centres are given by the diagonal entries, respectively. As the spectrum of eigenvalues of a matrix is the same as the spectrum of eigenvalues of a transposed matrix, we can compute the set of all Gershgorin disks as

$$\Lambda := \bigcup_{j=1}^{N} \underbrace{\left\{ z \in \mathbb{C} \,\middle|\, |z - p_{j,j}| \leq \sum_{i=1, i\neq j}^{N} |p_{i,j}| \right\}}_{=:B_j}. \tag{21}$$

Since $\boldsymbol{P}$ is nonnegative with unit column sums and positive diagonal elements, we conclude that

$$\sum_{i=1, i\neq j}^{N} |p_{i,j}| \;=\; \sum_{i=1, i\neq j}^{N} p_{i,j} \;=\; 1 - p_{j,j} \;<\; 1. \tag{22}$$

As it holds for all j that $B_j \cap \{z \in \mathbb{C} \mid |z| \;= 1\} = \{1\}$, we can describe Λ as

$$\Lambda \subset \{z \in \mathbb{C} \mid |z| < 1\} \cup \{1\}. \tag{23}$$

By the assumptions $\lambda \in \Lambda$ and $\lambda \neq 1$, we have $|\lambda| < 1$.

For the final step of our convergence analysis, we need the following results from the Perron-Frobenius theory (see e.g. Theorem 8.4.4 in [7]):
If $\boldsymbol{A} \in \mathbb{R}^{N\times N}$ is irreducible and nonnegative, then its spectral radius $\rho(\boldsymbol{A})$ is a simple eigenvalue of $\boldsymbol{A}$. Moreover, there exists a positive eigenvector to $\rho(\boldsymbol{A})$.
Since $\rho(\boldsymbol{P}) = 1$, this theorem states that $\lambda = 1$ is a simple eigenvalue and has a positive eigenvector. Hence, the iteration (15)–(16) attenuates all components outside the eigenspace of $\lambda = 1$ to zero. Therefore, the process converges to a vector $\boldsymbol{v}$ in the eigenspace of $\lambda = 1$. Since $\boldsymbol{f} \in \mathbb{R}^N_+$ and the iteration preserves the positive average grey value, it converges to a vector $\boldsymbol{v} \in \mathbb{R}^N$ with the same positive average grey value as $\boldsymbol{f}$. Because of the cited Perron-Frobenius result we know that $\boldsymbol{v}$ is positive. □

Our framework for discrete linear osmosis allows to analyse osmosis algorithms in a very simple way: All one has to do is to check the four properties (DLO1)–(DLO4). If they are satisfied, we can be sure that the filter preserves the average grey value and the positivity of the original image, and we have full control over its steady state.

It should be mentioned that this theory is very general: It does not rely on any specific space discretisation on a regular grid. Without any alterations, it is applicable to osmosis processes acting on graphs, on surface data, or on higher dimensional data sets.

4 Application to Finite Difference Discretisations

Let us now apply our discrete osmosis theory to two important finite difference discretisations that we have already mentioned: the explicit and the implicit scheme. We will see that they are not only useful for computing the parabolic time evolution, but also for the elliptic steady state.

4.1 The Parabolic Time Evolution

Applying Proposition 1 to the explicit and the implicit scheme gives the following result.

Proposition 2. [Finite Difference Discretisations]
Let $\boldsymbol{f} \in \mathbb{R}^N_+$ and consider the semidiscrete linear osmosis evolution

$$\boldsymbol{u}(0) = \boldsymbol{f}, \tag{24}$$

$$\boldsymbol{u}'(t) = \boldsymbol{A}\,\boldsymbol{u}(t) \tag{25}$$

where the (unsymmetric) matrix $\boldsymbol{A} = (a_{i,j}) \in \mathbb{R}^{N\times N}$ fulfils the following properties:

(SLO1) *All column sums of* $\boldsymbol{A}$ *are* 0.
(SLO2) $\boldsymbol{A}$ *has only nonnegative off-diagonal entries.*
(SLO3) $\boldsymbol{A}$ *is irreducible.*

Then the following results hold:

(a) The explicit scheme

$$\boldsymbol{u}^{k+1} \;=\; (\boldsymbol{I} + \tau \boldsymbol{A})\, \boldsymbol{u}^{k} \tag{26}$$

satisfies the requirements (DLO1)–(DLO4) for discrete linear osmosis processes provided that

$$\tau < \frac{1}{|a_{i,i}|} \qquad \forall\, i \in \{1, ..., N\}. \tag{27}$$

(b) The implicit scheme

$$(\boldsymbol{I} - \tau \boldsymbol{A})\, \boldsymbol{u}^{k+1} \;=\; \boldsymbol{u}^{k} \tag{28}$$

satisfies (DLO1)–(DLO4) for all time step sizes $\tau > 0$.

Proof. We check (DLO1)–(DLO4) by applying classical matrix analysis.

(a) It holds that $a_{i,i} \neq 0$ because otherwise (SLO1) implies that the whole column i of $\boldsymbol{A}$ is 0, and thus the digraph associated with $\boldsymbol{A}$ is not strongly connected. This contradicts the irreducibility of $\boldsymbol{A}$. The unit column sum property (DLO1) follows directly from the zero column sums of $\boldsymbol{A}$. Moreover, $\boldsymbol{I}+\tau\boldsymbol{A}$ is nonnegative (DLO2) with positive diagonal elements (DLO4), since (SLO2) holds true and τ fulfils (27). Clearly, (27) guarantees that $\boldsymbol{I} + \tau\boldsymbol{A}$ and $\boldsymbol{A}$ have the same digraph. Thus, $\boldsymbol{I} + \tau\boldsymbol{A}$ is also irreducible (DLO3).

(b) We start by observing that $\boldsymbol{I} - \tau\boldsymbol{A}$ is strictly column diagonally dominant: From the zero column sum property (SLO1) it follows that

$$-a_{j,j} = \sum_{i=1, i \neq j}^{N} a_{i,j} \qquad \forall\, j \in \{1, \ldots, N\} \tag{29}$$

and thus

$$1 - \tau a_{j,j} > \tau \sum_{i=1, i \neq j}^{N} a_{i,j} \qquad \forall\, \tau > 0. \tag{30}$$

By (SLO2) the off-diagonals of $\boldsymbol{A}$ are nonnegative. Hence, we can apply Gershgorin's theorem to the columns of $\boldsymbol{I}-\tau\boldsymbol{A}$ and conclude that this matrix is nonsingular. Let us consider the row vector $\boldsymbol{e} := (1, 1, \ldots, 1)$ with N components. Clearly, (SLO1) means that $\boldsymbol{I} - \tau\boldsymbol{A}$ has unit column sums. The same holds true for its inverse since

$$\boldsymbol{e}\,(\boldsymbol{I} - \tau\boldsymbol{A}) \;=\; \boldsymbol{e} \quad \Longleftrightarrow \quad \boldsymbol{e} \;=\; \boldsymbol{e}\,(\boldsymbol{I} - \tau\boldsymbol{A})^{-1}. \tag{31}$$

This proves (DLO1). The nonpositivity of the off-diagonals of $\boldsymbol{I} - \tau\boldsymbol{A}$ and its strict column diagonal dominance imply that $\boldsymbol{I} - \tau\boldsymbol{A}$ is a nonsingular M-matrix, cf. [8, Theorem 6.2.3 (C_{10})]. For any nonsingular M-matrix, it holds that its inverse has only strictly positive entries; see [8, Theorem 6.2.7]. This shows (DLO2)–(DLO4). □

Proposition 2 gives stability results with respect to preservation of positivity. Since also the average grey value is preserved, it follows that $0 < u_j^k < \sum_i f_i$ for all $j \in \{1, ..., N\}$ and for all $k > 0$. This ensures that also the ℓ_p norms of the solution remain bounded for $p > 1$. Note that osmosis does not allow to give stability results in terms of decreasing ℓ_p norms for $p > 1$, since comparable properties for the L_p norms do not hold for the continuous equation: An osmosis process that starts with a flat image and converges to a nonflat one with identical average grey value may serve as counterexample. This shows that preservation of positivity is a very natural stability criterion for osmosis.

For a spatial grid size of $h = 1$, the condition (9) becomes

$$|d_1(\boldsymbol{x})| < 2, \qquad |d_2(\boldsymbol{x})| < 2 \qquad \forall\, \boldsymbol{x} \in \Omega, \tag{32}$$

and inspecting the central weight in (8) shows that $|a_{i,i}| < 8$. Thus, the stability condition (27) for the explicit scheme becomes $\tau < \frac{1}{8}$. This stability bound is half as large as the well-known stability limit of an explicit scheme for the homogeneous 2-D diffusion equation $\partial_t u = \Delta u$.

The absolute stability of the implicit scheme is in full accordance with the corresponding diffusion result from [6, Theorem 8]. The implicit scheme yields nonsymmetric pentadiagonal systems of linear equations that are strictly diagonally dominant in their columns. Using the classical theory of regular splittings [9], one can show that the Gauß-Seidel algorithm converges under these circumstances. More efficient alternatives include Krylov subspace methods such as the BiCGStab method [10] and its preconditioned variants [11]. Implementing these iterative methods is fairly straightforward. Also multigrid methods [12] appear promising, but are more cumbersome to implement.

4.2 The Elliptic Steady State

For many applications of osmosis – such as the ones discussed in [3] – one is mainly interested in the osmotic steady state. Thus, it appears tempting to approximate the elliptic PDE

$$\Delta u - \operatorname{div}(\boldsymbol{d}u) \;=\; 0 \tag{33}$$

and its homogeneous Neumann boundary conditions directly with numerical solvers. However, this can become unpleasant since the elliptic problem has infinitely many solutions: For any solution $w(\mathbf{x})$, also $c\, w(\mathbf{x})$ with some arbitrary constant c is a solution.

This suggests to use also our parabolic time evolution schemes to obtain the desired solution that is positive and has the same average grey value as the initial image f. For the explicit scheme (26) the steady state $\boldsymbol{w}$ is characterised by

$$\boldsymbol{w} \;=\; (\boldsymbol{I} + \tau \boldsymbol{A})\, \boldsymbol{w} \tag{34}$$

and for the implicit scheme (28), it satisfies

$$(\boldsymbol{I} - \tau \boldsymbol{A})\, \boldsymbol{w} \;=\; \boldsymbol{w}\,. \tag{35}$$

Interestingly both equations (34) and (35) are equivalent to

$$\boldsymbol{A}\boldsymbol{w} \;=\; \boldsymbol{0} \tag{36}$$

which is a space discretisation of the elliptic PDE (33). Thus, we have the remarkable situation that any stable time step size τ gives the correct elliptic steady state $\boldsymbol{w}$. This makes the implicit scheme with large τ attractive for this task, if one has an efficient solver for the resulting linear systems of equations.

5 Experimental Evaluation

The preceding discrete osmosis framework provides general criteria that guarantee the *reliability* of osmosis schemes. However, it tells us nothing about their *speed*. To evaluate the practical performance of the explicit and the implicit scheme, let us now consider a typical image editing problem where one is interested in the osmotic steady state.

For our experiment we want to combine the two images from Figure 1(a) and (b). They depict contemporary paintings of famous US presidents. The task is to replace the face of George Washington with the face of Abraham Lincoln in a seamless way. The image of Washington serves as initialisation of our osmosis process. In order to apply osmosis, we first have to specify its drift vectors. We choose the drift vectors of the Washington image where the binary mask image of Fig. 1(c) is black, and the drift vectors of the Lincoln image where the mask image is white. At the interface we perform arithmetic averaging of both drift vector fields. With this combined drift vector field we compute the osmosis evolution. Its steady state gives the seamlessly cloned image in Fig. 1(d).

Now let us discuss some numerical details. For a positive image f, we use the following discretisation of its canonical drift vector field $(d_1, d_2)^\top = \frac{\boldsymbol{\nabla} f}{f}$ in the sense of (6):

$$d_{1,i+\frac{1}{2},j} = \frac{2\,(f_{i+1,j} - f_{i,j})}{h\,(f_{i+1,j} + f_{i,j})}\,, \qquad d_{2,i,j+\frac{1}{2}} = \frac{2\,(f_{i,j+1} - f_{i,j})}{h\,(f_{i,j+1} + f_{i,j})}\,. \tag{37}$$

These vectors are fed into our space discretisation (7), and as time discretisation we use the explicit and the implicit scheme. In the implicit case, we have tested different solvers for the linear system of equations, including Gauß-Seidel, SOR, BiCGStab, and two preconditioned BiCGStab variants. Because BiCGStab without preconditioning offered the best performance, we only report results for this solver here. Since we approach our steady state solution iteratively, we need a stopping criterion: We compute the average ℓ_1 distance per pixel between our numerical solution and a precomputed ground truth. The iterations are stopped if this error is less than 0.1, where the initial range of each colour channel is $[1, 256]$.

Table 1 shows a comparison of the CPU times for the explicit and the implicit scheme for three different image sizes. The run times are obtained with a

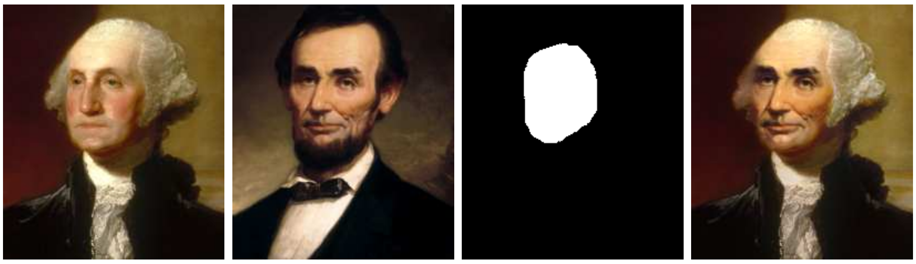

Fig. 1. Seamless image cloning with osmosis. From left to right: **(a)** Painting of George Washington by Gilbert Stuart (Source: Wikimedia Commons, public domain work). **(b)** Painting of Abraham Lincoln by George Story (Source: Wikimedia Commons, public domain work). **(c)** Mask for the seamless image cloning. **(d)** Osmotic steady state using combined drift vector fields.

Table 1. CPU times [s] and number of iterations for different image sizes and different osmosis schemes. For the explicit scheme we use $\tau = 0.12$, and in the implicit case $\tau = 10^5$.

image size	explicit:	time[s]	iterations	implicit:	time[s]	iterations
100×115		14.689	61184		0.3179	2
200×230		359.49	240115		4.5454	2
400×460		4487.6	948484		61.909	3

double precision C implementation on a standard desktop PC with an Intel Xeon processor, clocked at 3.2 GHz with single threading and without GPU support. We observe that the implicit scheme with BiCGStab allows to reach the desired steady state solution up to 79 times faster than the explicit scheme.

6 Summary and Conclusions

We have introduced a fully discrete theory for osmosis filters that can be expressed in terms of linear drift-diffusion equations. Its prerequisites differ from the ones for discrete diffusion filtering by the fact that the iteration matrix is not symmetric. We have seen that this seemingly small difference has a substantial impact on properties such as maximum-minimum principles and nontrivial steady states. The possibility to design interesting steady states is a key feature of osmosis filtering, and our paper has provided a discrete characterisation of the osmotic steady state. Moreover, we have established stability results in terms of preservation of positivity which is a very natural stability concept for osmosis. We have shown that our theory is applicable to important finite difference approximations such as explicit and implicit schemes. Finally, we have demonstrated that an implicit scheme with a BiCGStab solver also constitutes an efficient method for obtaining the osmotic steady state. This method is not very difficult to implement and can be two orders of magnitude faster than the explicit scheme.

In our ongoing work we are exploring alternative numerical options such as multigrid solvers [12] and additive operator splittings (AOS) [13, 14]. Moreover, we are establishing semidiscrete and continuous theories for osmosis filtering that have a similar structure as their diffusion counterparts.

Acknowledgments. This work has been performed while Kai Hagenburg and Oliver Vogel were members of the Mathematical Image Analysis Group. Financial support by the *Deutsche Forschungsgemeinschaft (DFG)* is gratefully acknowledged.

References

1. Hagenburg, K., Breuß, M., Vogel, O., Weickert, J., Welk, M.: A lattice Boltzmann model for rotationally invariant dithering. In Bebis, G., Boyle, R., Parvin, B., Koracin, D., Kuno, Y., Wang, J., Pajarola, R., Lindstrom, P., Hinkenjann, A., Encarnacao, M.L., Silva, C.T., Coming, D., eds.: Advances in Visual Computing. Volume 5876 of Lecture Notes in Computer Science. Springer, Berlin (2009) 949–959
2. Hagenburg, K., Breuß, M., Weickert, J., Vogel, O.: Novel schemes for hyperbolic PDEs using osmosis filters from visual computing. In Bruckstein, A.M., ter Haar Romeny, B., Bronstein, A.M., Bronstein, M.M., eds.: Scale Space and Variational Methods in Computer Vision. Volume 6667 of Lecture Notes in Computer Science. Springer, Berlin (2012) 532–543
3. Weickert, J., Hagenburg, K., Breuß, M., Vogel, O.: Linear osmosis models for visual computing. (2013) Submitted.
4. Fattal, R., Lischinski, D., Werman, M.: Gradient domain high dynamic range compression. In: Proc. SIGGRAPH 2002, San Antonio, TX (July 2002) 249–256
5. Pérez, P., Gagnet, M., Blake, A.: Poisson image editing. ACM Transactions on Graphics **22**(3) (2003) 313–318
6. Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner, Stuttgart (1998)
7. Horn, R.A., Johnson, C.R.: Matrix Analysis. Cambridge University Press, Cambridge, UK (1990)
8. Berman, A., Plemmons, R.J.: Nonnegative Matrices in the Mathematical Sciences. Academic Press, New York (1979)
9. Varga, R.A.: Matrix Iterative Analysis. Prentice Hall, Englewood Cliffs (1962)
10. van der Vorst, H.A.: Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. SIAM Journal on Scientific and Statistical Computing **13**(2) (1992) 631–644
11. Saad, Y.: Iterative Methods for Sparse Linear Systems. Second edn. SIAM, Philadelphia (2003)
12. Briggs, W.L., Henson, V.E., McCormick, S.F.: A Multigrid Tutorial. Second edn. SIAM, Philadelphia (2000)
13. Lu, T., Neittaanmäki, P., Tai, X.C.: A parallel splitting up method and its application to Navier–Stokes equations. Applied Mathematics Letters **4**(2) (1991) 25–29
14. Weickert, J., ter Haar Romeny, B.M., Viergever, M.A.: Efficient and reliable schemes for nonlinear diffusion filtering. IEEE Transactions on Image Processing **7**(3) (March 1998) 398–410