THE DISCRETE FOCAL STACK TRANSFORM

F. Pérez Nava [1], J.G. Marichal-Hernández [2], J.M. Rodríguez-Ramos.[2]
[1]Departamento de Estadística, Investigación Operativa y Computación.
[2]Departamento de Física Fundamental y Experimental, Electrónica y Sistemas.
Universidad de La Laguna, 38271, Spain.
email: fdoperez@ull.es, jmariher@ull.es, jmramos@ull.es

ABSTRACT

This paper describes the Discrete Focal Stack Transform (DFST) that computes the focal stack (a collection of photographs focused at different depths) from a lightfield. It proposes a new discretization of the Photography operator based on the interpolation of the lightfield by means of 4D trigonometric polynomials and an approximation to the integration process. Our approach combines the Fourier decomposition of the Photography operator with a generalization of the discrete Radon transform and provides a fast and exact computation of the focal stack. We also use our formulation to generalize the computation of the focal stack to the modified laplacian stack that can be used to solve the depth estimation problem. Finally we provide some computational results that show the validity of the approach.

1. INTRODUCTION

Computational photography is based on capturing and processing discrete samples of all the light rays in the 3D space. These light rays are described by the 7D plenoptic function [1] in terms of position, orientation, spectral content, and time. This can be simplified to a 4D function by considering only the value of each ray as a function of its position and orientation in a static scene, and by constraining each ray to have the same value at every point along its direction of propagation. Compared to conventional photography, which captures 2D images, computational photography captures the entire 4D lightfield. The conventional 2D images are obtained by 2D projections of the 4D lightfield. These images can be used to overcome long-standing limitations of photographic films such as constraints on dynamic range, depth of field, focus, resolution and the extent of scene motion during exposure.

In this paper we will address the refocusing process whose goal is to obtain the focal stack i.e. to produce photographs focused at different depths in the scene. Given the lightfield of a scene it is theoretically possible to obtain a conventional photograph focused at a desired depth by means of the Photography integral operator [2]. Using a plenoptic camera [3] it is feasible to capture samples from the 4D lightfield. A plenoptic camera uses a microlens array to measure the intensity and direction of light rays. There are several variants of plenoptic cameras including hand-held [4] or video cameras [5]. After the capture of the samples from the lightfield and to obtain the focused photograph it is necessary to solve two problems: to interpolate the lightfield and to approximate the integration process. The brute-force approach would interpolate the lightfield by nearest neighbour and approximate the integral by sums. If we assume that the plenoptic camera is composed of $n \times n$ microlenses and that each microlens generates a $n \times n$ image, the brute-force approach would need $O(n^4)$ operations to generate a refocused photography. A significant improvement to this performance was obtained in [2] with the "Fourier Slice Photography" (FSP) technique that reformulates the problem in the Fourier domain. The FSP method is based on the extraction of an appropriate dilated 2D slice in the 4D Fourier transform of the lightfield. In this case interpolation takes place in the 4D Fourier domain using a 4D Kaiser-Bessel filter. This decreases the computational cost of the discretization of the Photography operator to $O(n^4\log(n))$ operations to generate a focal stack with n^2 refocused photographs.

In this paper we propose the Discrete Focal Stack Transform (DFST). Instead of using nearest neighbour interpolation in the 4D spatial domain or Kaiser-Bessel interpolation in the 4D frequency domain we use 4D trigonometric interpolation in the spatial domain and a discretization of the integral operator. The DFST links the trigonometric decomposition that characterizes Fourier approaches with a generalization of the discrete Radon transform (DRT) [6] and solves the Focal Stack problem sharing the main features of the DRT:

- Algebraic exactness: The transform is based in a clear and rigorous definition. It can be computed numerically and analytically. The transform is non-parametric.
- Geometric exactness: The transform uses true geometric subspaces.
- Speed: The transform employs $O(n^4\log(n))$ operations to generate a focal stack composed of $n\log(n)$ refocused photographs.
- Parallels with continuum theory: the transform obeys relations with those of the continuous transform.

There is only one feature of the DRT not shared by the DFST: invertibility. This is not a weakness of our approach. The continuous Focal Stack transform is not invertible in general. This paper is divided in six parts. In Section 2 we present the continuous Focal stack transform. In Section 3 we propose the DFST and discuss its properties. In Section 4 we show an algorithm to compute the DFST and finally in Section 5 we discuss the conclusions and future extensions.

2. THE CONTINOUS FOCAL STACK TRANSFORM

In this section we present our definition for the continuous Focal Stack transform in terms of the Photography transform.

2.1 Definition of the continuous Focal Stack transform

To introduce the Photography transform we parameterize the lightfield defined by all the light rays inside a camera. We will use the two-plane parameterization of this lightfield and write $L_F(\boldsymbol{x},\boldsymbol{u})$ as the radiance travelling from position $\boldsymbol{u}=(u_1, u_2)'$ (apostrophe mean transpose) on the lens plane to position $\boldsymbol{x}=(x_1, x_2)'$ on the sensor plane. F is the distance between the lens and the sensor (see Figure 1 adapted from [2]).

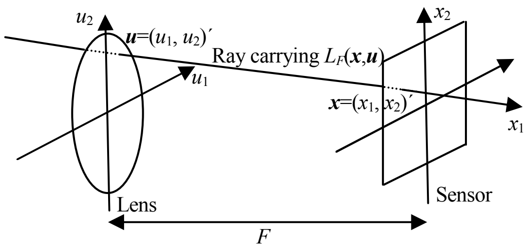

Figure 1 - Two plane parameterization of the lightfield.

The lightfield L_F can be used to compute conventional photographs at any depth αF. Let $\mathcal{P}_\alpha$ be the operator that transforms a lightfield L_F at a sensor depth F into a photograph formed on a film at sensor depth αF, then we have [2]:

$$\mathcal{P}_\alpha[L_F](\boldsymbol{x}) = \frac{1}{\alpha^2 F^2}\int L_F\left(\boldsymbol{u}\left(1-\frac{1}{\alpha}\right)+\frac{\boldsymbol{x}}{\alpha},\boldsymbol{u}\right)d\boldsymbol{u}.$$

This equation explains how to compute photographs at different depths from the lightfield L_F. When we compute the photographs for every sensor depth αF we obtain the focal stack transform $\mathcal{S}$ of the lightfield. We will define $\mathcal{S}$ as:

$$\mathcal{S}[L_F](\boldsymbol{x},\alpha) = \mathcal{P}_\alpha[L_F](\alpha\boldsymbol{x}).$$

2.2 The Fourier Slice Photography (FSP) transform.

Since one of the features of the proposed DFST is to obey the relations with those of the continuous transform we have to examine the Photography operator in the Fourier domain. The study of the operator in the Fourier domain leads to the definition of the Fourier Photograpy operator $\mathfrak{P}$:

$$\mathfrak{P}_\alpha[C](\boldsymbol{a}) = \frac{1}{F^2}C(\alpha\boldsymbol{a},(1-\alpha)\boldsymbol{a}), \qquad \boldsymbol{a}=(a_1,a_2)'.$$

and the following result:

Theorem (Fourier Slice Photography) [2]

"A conventional photograph is the inverse 2D Fourier transform of a dilated 2D slice in the 4D Fourier transform of the lightfield":

$$\mathcal{P}_\alpha \equiv \mathcal{F}_{2\mathrm{D}}^{-1}\circ\mathfrak{P}_\alpha\circ\mathcal{F}_{4\mathrm{D}}.$$

The utility of this theorem lies in the replacement of a complex operator $\mathcal{P}_\alpha$ in the spatial domain by a simpler one $\mathfrak{P}_\alpha$ in the frequency domain.

3. THE DISCRETE FOCAL STACK TRANSFORM (DFST)

In this section we present the DFST transform. In section 3.1 we will study how to define the photography operator for a periodic and continuous signal. Then, in section 3.2 we will define the photography operator for a periodic and discrete signal and we will formulate the discrete counterpart to the Fourier Photography theorem by means of the Fractional Fourier Transform.

3.1 The Photography operator for periodic and continuous signals

To discretize the Focal Stack operator we have to consider the finite extent of the lens and sensor. We will assume then that the lightfield L_F is defined on the hypercube $\boldsymbol{H}$=[$-T_{x1}/2$, $T_{x1}/2$]×[$-T_{x2}/2$, $T_{x2}/2$]×[$-T_{u1}/2$, $T_{u1}/2$]×[$-T_{u2}/2$, $T_{u2}/2$] . However, to computer $\mathcal{P}_\alpha[L_F](\boldsymbol{x})$ we need to extend it outside $\boldsymbol{H}$. We will assume then that the lightfield is periodic with periods $\boldsymbol{T_x}=(T_{x1},T_{x2})'$ $\boldsymbol{T_u}=(T_{u1},T_{u2})'$ and define the Photography transform for periodic functions as the integral over one period:

$$\mathcal{P}_\alpha[L_F](\boldsymbol{x}) = \frac{1}{\alpha^2F^2}\int_{-\boldsymbol{T_u}/2}^{\boldsymbol{T_u}/2} L_F\left(\boldsymbol{u}\left(1-\frac{1}{\alpha}\right)+\frac{\boldsymbol{x}}{\alpha},\boldsymbol{u}\right)d\boldsymbol{u}.$$

Where we use the notation:

$$\int_{-\boldsymbol{T_u}/2}^{\boldsymbol{T_u}/2} f(\boldsymbol{u})\,d\boldsymbol{u} = \int_{-T_{u_1}/2}^{T_{u_1}/2}\int_{-T_{u_2}/2}^{T_{u_2}/2} f(u_1,u_2)du_1du_2.$$

We now present the analogous of the Fourier Slice Photography Theorem:

Theorem (Fourier Slice for Periodic and Continuous Signals)

The Photography Operator for periodic signals can be written as:

$$\mathcal{P}_\alpha[L_F](\boldsymbol{x}) = \sum_{\widehat{\boldsymbol{a}}=-\infty}^{\infty} e^{2\pi i(\boldsymbol{x}'\boldsymbol{a})}\frac{1}{F^2}C(\alpha\boldsymbol{a},(1-\alpha)\boldsymbol{a})|\Delta\boldsymbol{a}|,$$

$$C(\boldsymbol{a},\boldsymbol{b}) = \frac{1}{F^2}\int_{\frac{-\boldsymbol{T_x}}{2}}^{\frac{\boldsymbol{T_x}}{2}}\int_{\frac{-\boldsymbol{T_u}}{2}}^{\frac{\boldsymbol{T_u}}{2}} L_F(\boldsymbol{x},\boldsymbol{u})e^{-2\pi i(\boldsymbol{x}'\boldsymbol{a}+\boldsymbol{u}'\boldsymbol{b})}d\boldsymbol{x}d\boldsymbol{u},$$

where:

$$\boldsymbol{a}=(a_1,a_2)', \widehat{\boldsymbol{a}} = \boldsymbol{a}\cdot(\alpha\boldsymbol{T}_x)\in\mathbb{Z}, |\Delta\boldsymbol{a}| = \frac{1}{\alpha^2|\boldsymbol{T_x}|},$$

and we use the following notation:

$$\sum_{\widehat{\boldsymbol{a}}=-\infty}^{\infty} f(\boldsymbol{a}) = \sum_{a_1=-\infty}^{\infty}\sum_{a_2=-\infty}^{\infty} f(a_1,a_2).$$

$\boldsymbol{x}\cdot\boldsymbol{y}$ stands for the pointwise multiplication of $\boldsymbol{x}$ and $\boldsymbol{y}$ and $|\boldsymbol{x}|$ stands for the product of all components in vector $\boldsymbol{x}$.

3.2 The Photography operator for periodic and discrete signals.

This section introduces our definition for the DFST and the analogous for the Fourier Slice Photography theorem for discrete signals. We suppose that the lightfield L_F is periodic with periods $\boldsymbol{T_x}$ and $\boldsymbol{T_u}$ and that its values denoted by $L_F^d(\boldsymbol{x},\boldsymbol{u})$ are known for the set of points on a grid $\boldsymbol{G}$:

$$\boldsymbol{G} = \{(\hat{\boldsymbol{x}}\cdot\Delta\boldsymbol{x}, \hat{\boldsymbol{u}}\cdot\Delta\boldsymbol{u})\},$$
$$\hat{\boldsymbol{x}} = -\boldsymbol{n_x}\ldots\boldsymbol{n_x}\in\mathbb{Z}^2, \hat{\boldsymbol{u}} = -\boldsymbol{n_u}\ldots\boldsymbol{n_u}\in\mathbb{Z}^2$$

which is a compact notation for a 4D grid where:

$\hat{x}_i = -n_{x_i}\ldots n_{x_i}\in\mathbb{Z}, \hat{u}_i = -n_{u_i}\ldots n_{u_i}\in\mathbb{Z}$, i=1,2.

$\Delta\boldsymbol{x} = (\Delta x_1,\Delta x_2)'$, $\Delta\boldsymbol{u} = (\Delta u_1,\Delta u_2)'$,

$$\Delta x_i = \frac{T_{x_i}}{N_{x_i}}, \Delta u_i = \frac{T_{u_i}}{N_{u_i}},$$

and:

$\boldsymbol{N_x} = (N_{x_1},N_{x_2})', \boldsymbol{N_u} = (N_{u_1},N_{u_2})'$,

$\boldsymbol{N_x} = 2\boldsymbol{n_x}+1, N_u = 2\boldsymbol{n_u}+1$,

$\boldsymbol{n_x} = (n_{x_1},n_{x_2})', \boldsymbol{n_u} = (n_{u_1},n_{u_2})'$.

Then we define L_F^d as the trigonometric polynomial that interpolates those points:

$$L_F^d(\boldsymbol{x},\boldsymbol{u}) = \sum_{\hat{\boldsymbol{a}}=-\boldsymbol{n_x}}^{\boldsymbol{n_x}}\sum_{\hat{\boldsymbol{b}}=-\boldsymbol{n_u}}^{\boldsymbol{n_u}} C^d(\boldsymbol{a}^*,\boldsymbol{b}^*)e^{2\pi i(x'\boldsymbol{a}^*+u'\boldsymbol{b}^*)}|\Delta\boldsymbol{a}^*||\Delta\boldsymbol{b}^*|,$$

$$C^d(\boldsymbol{a}^*,\boldsymbol{b}^*) = \sum_{\hat{\boldsymbol{x}}=-\boldsymbol{n_x}}^{\boldsymbol{n_x}}\sum_{\hat{\boldsymbol{u}}=-\boldsymbol{n_u}}^{\boldsymbol{n_u}} L_F^d(\boldsymbol{x},\boldsymbol{u})\, e^{-2\pi i(x'\boldsymbol{a}^*+u'\boldsymbol{b}^*)}|\Delta\boldsymbol{x}||\Delta\boldsymbol{u}|,$$

where:

$$|\Delta\boldsymbol{a}^*| = \frac{1}{|\boldsymbol{T_x}|}, |\Delta\boldsymbol{b}^*| = \frac{1}{|\boldsymbol{T_u}|},$$

and we use the following notation:

$$\sum_{\boldsymbol{i}=-\boldsymbol{n_i}}^{\boldsymbol{n_i}} f(\boldsymbol{i}) = \sum_{i_1=-n_{i_1}}^{n_{i_1}}\sum_{i_2=-n_{i_2}}^{n_{i_2}} f(i_1,i_2).$$

Then, we approximate the integral by sums and we show a preliminary definition of the discrete Photography operator:

$$\mathcal{P}_\alpha^d[L_F^d](\boldsymbol{x}) = \frac{1}{|\alpha|^2 F}\sum_{\hat{\boldsymbol{u}}=-\boldsymbol{n_u}}^{\boldsymbol{n_u}} L_F^d\left(\boldsymbol{u}\left(1-\frac{1}{\alpha}\right)+\frac{\boldsymbol{x}}{\alpha},\boldsymbol{u}\right)|\Delta\boldsymbol{u}|,$$
$$\boldsymbol{u} = \hat{\boldsymbol{u}}\cdot\Delta\boldsymbol{u}.$$

This definition carries its own slice theorem:

Theorem (Fourier Slice for Periodic and Discrete Signals)
The Photography Operator for periodic and discrete signals can be written as:

$$\mathcal{P}_\alpha^d[L_F^d](\boldsymbol{x}) = \sum_{\hat{\boldsymbol{a}}=-\boldsymbol{n_x}}^{\boldsymbol{n_x}} e^{2\pi i(x'\boldsymbol{a})}\frac{1}{F^2}C^d(\alpha\boldsymbol{a},(1-\alpha)\boldsymbol{a})|\Delta\boldsymbol{a}|,$$

where:

$$\boldsymbol{a} = (a_1,a_2)', \hat{\boldsymbol{a}} = \boldsymbol{a}\cdot(\alpha\boldsymbol{T}_x)\in\mathbb{Z}, |\Delta\boldsymbol{a}| = \frac{1}{\alpha^2|\boldsymbol{T_x}|}.$$

It may seem that the goal has been obtained: a discrete definition of the Photography Operator with its own slice theorem. Unfortunately this is not the case. We still have two problems to solve: duality and periodicity. To simplify the discussion of these problems we will assume from now on that: $\Delta x_1 = \Delta x_2$ and $\Delta u_1 = \Delta u_2$ and denote these common values as Δx and Δu. We will also assume that $N_{x_1} = N_{x_2}$ and $N_{u_1} = N_{u_2}$ and denote those common values by N_x and N_u. These constraints are common in plenoptic cameras.

3.2.1 The duality problem

The duality problem arises as a consequence of the definition of the Photography operator. The role of the sensor and the lens can be interchanged and there are four equivalent ways to evaluate the Photography operator depending on which are the independent variables. For example we may write:

$$\mathcal{P}_\alpha^{u_1,u_2}[L_F](\boldsymbol{x}^*) = \frac{1}{\alpha^2F^2}\int L_F\left(\boldsymbol{u}\left(1-\frac{1}{\alpha}\right)+\frac{\boldsymbol{x}^*}{\alpha},\boldsymbol{u}\right)d\boldsymbol{u} =$$
$$= \frac{1}{(1-\alpha)^2F^2}\int L_F\left(\boldsymbol{x},\boldsymbol{u}\left(1-\frac{1}{1-\alpha}\right)+\frac{\boldsymbol{x}^*}{1-\alpha}\right)d\boldsymbol{x} =$$
$$= \mathcal{P}_\alpha^{x_1,x_2}[L_F](\boldsymbol{x}^*).$$

Notice that the formula above says that we can evaluate also $\mathcal{P}_\alpha[L_F](\boldsymbol{x})$ by means of $\mathcal{P}_\alpha^{x_1,x_2}[L_F](\boldsymbol{x})$ interchanging the roles of the sensor and the lens and using 1-α instead of α. In the continuous case this is not a problem because both values are the same, but in the discrete case this is a problem since we have four alternative values for the same true value. It can be shown that we have to use our original definition:

$$\mathcal{P}_\alpha^{d:u_1,u_2}[L_F^d](\boldsymbol{x}) = \frac{1}{\alpha^2F^2}\sum_{\hat{\boldsymbol{u}}=-\boldsymbol{n_u}}^{\boldsymbol{n_u}} L_F^d\left(\boldsymbol{u}\left(1-\frac{1}{\alpha}\right)+\frac{\boldsymbol{x}}{\alpha},\boldsymbol{u}\right)|\Delta\boldsymbol{u}| =$$
$$= \sum_{\hat{\boldsymbol{a}}=-\boldsymbol{n_x}}^{\boldsymbol{n_x}} e^{2\pi i(x'\boldsymbol{a})}\frac{1}{F^2}C^d(\alpha\boldsymbol{a},(1-\alpha)\boldsymbol{a})|\Delta\boldsymbol{a}|,$$
$$\boldsymbol{a} = (a_1,a_2)', \hat{\boldsymbol{a}} = \boldsymbol{a}\cdot(\alpha\boldsymbol{T}_x)\in\mathbb{Z}, |\Delta\boldsymbol{a}| = \frac{1}{\alpha^2|\boldsymbol{T_x}|}.$$

when $|\alpha|\Delta x \geq |1-\alpha|\Delta u$, and an alternative definition:

$$\mathcal{P}_\alpha^{d:x_1,x_2}[L_F^d](\boldsymbol{x}^*) = \frac{1}{(1-\alpha)^2F^2}\sum_{\hat{\boldsymbol{x}}=-\boldsymbol{n_x}}^{\boldsymbol{n_x}} L_F^d\left(\boldsymbol{x},\boldsymbol{x}\left(1-\frac{1}{1-\alpha}\right)+\frac{\boldsymbol{x}^*}{1-\alpha}\right)|\Delta\boldsymbol{x}|$$
$$= \sum_{\hat{\boldsymbol{b}}=-\boldsymbol{n_u}}^{\boldsymbol{n_u}} e^{2\pi i(x^{*\prime}\boldsymbol{b})}\frac{1}{F^2}C^d((1-\alpha)\boldsymbol{b},\alpha\boldsymbol{b})|\Delta\boldsymbol{b}|,$$
$$\boldsymbol{b} = (b_1,b_2)', \hat{\boldsymbol{b}} = \boldsymbol{b}\cdot((1-\alpha)\boldsymbol{T}_u)\in\mathbb{Z}, |\Delta\boldsymbol{b}| = \frac{1}{(1-\alpha)^2|\boldsymbol{T_u}|}.$$

when $|\alpha|\Delta x \leq |1-\alpha|\Delta u$.

We note that the inequality $|\alpha|\Delta x \geq |1-\alpha|\Delta u$ is the range of focal depths where it can be achieved a perfect reconstruction of the Photography operator from the samples of a band limited lightfield in the continuous case [2].

3.2.2 The periodicity problem

The periodicity assumption that allows us to extend the lightfield beyond the grid $\boldsymbol{G}$ is also a problem because we usually want the lightfield to be zero beyond its domain. We will use the solution of the DRT [6] to this problem. Suppose we have a grid $\boldsymbol{G}$ as in section 3.2 where for simplicity we assume that N_x and N_u are even. We define the extended grid $\boldsymbol{G}^{ext}$ as a grid of $N_x^{ext} = N_x + 2\lceil n_u/2 \rceil + 1, n_u = N_u/2$, elements for x_1 and x_2 and $N_u^{ext} = N_x + N_u + 1$ elements for u_1 and u_2 with the grid $\boldsymbol{G}$ centred on it, where we denote by $\lceil x \rceil$ the ceil function. Then, we set to zero the lightfield values for those extended points obtaining the extended lightfield $L_F^{d;ext}$.

3.2.3 The discrete photography operator

Finally, since the proposed transform is still valid for every $\boldsymbol{x}$ we discretize $\boldsymbol{x}$ and write $x_i = \Delta x \hat{x}_i$ and complete the definition of the transform $\mathcal{P}_\alpha^d[L_F^d]$ as:

$\mathcal{P}_\alpha^{d:u_1,u_2}[L_F^{d;ext}](\alpha\boldsymbol{x})$ for:

$|\alpha|\Delta x \geq |1-\alpha|\Delta u,\ \hat{x}_i = -n_x \dots n_x - 1, n_x = N_x/2$, and

$\mathcal{P}_\alpha^{d:x_1,x_2}[L_F^{d;ext}]((1-\alpha)\boldsymbol{x})$ for

$|\alpha|\Delta x \leq |1-\alpha|\Delta u,\ \hat{x}_i = -n_u \dots n_u - 1, n_u = N_u/2.$

3.2.4 The discrete focal stack transform

We define the discrete focal stack as:

$\mathcal{S}^d[L_F^d](\boldsymbol{x},\alpha) = \mathcal{P}_\alpha^d[L_F^d](\alpha\boldsymbol{x}),\ \hat{x}_i = -n_x \dots n_x - 1$

For the discrete set of α-values:

$$\left(1-\frac{1}{\alpha}\right) = \hat{p}\Delta p, \Delta p = \frac{\Delta x}{n_u \Delta u}, -n_u \leq \hat{p} \leq n_u,\ \hat{p} \in \mathbb{Z}.$$

$$\left(1-\frac{1}{\alpha}\right) = \frac{\Delta p}{\hat{p}}, \Delta p = \frac{n_x \Delta x}{\Delta u}\ -n_x \leq \hat{p} \leq n_x,\ \hat{p} \in \mathbb{Z}.$$

This is a generalization of the polar grid [6]. To compute the focal stack we have to note that the two different definitions of $\mathcal{P}_\alpha^d$ leads to two different ways to compute $\mathcal{S}^d$. The algorithm for the discrete focal stack is based on the (unaliased) 2D discrete Fractional Fourier transform (FrFT)[7]:

$$(\boldsymbol{F}_{\boldsymbol{N}}^{\boldsymbol{\alpha}}[f])(\boldsymbol{a}) = \sum_{\boldsymbol{x}=-\boldsymbol{n}}^{\boldsymbol{n}} f(\boldsymbol{x}) e^{-\frac{2\pi i \alpha_1 x_1 a_1}{N_1}} e^{-\frac{2\pi i \alpha_2 x_2 a_2}{N_2}},$$

$\boldsymbol{\alpha} = (\alpha_1, \alpha_2)', \boldsymbol{N} = 2\boldsymbol{n} + 1, \boldsymbol{n} = (n_1, n_2)',$

$\boldsymbol{a} = (a_1, a_2)', a_i = -n_i \dots n_i, \boldsymbol{x} = (x_1, x_2)'.$

The FrFT transform can be computed with complexity proportional to the Fast Fourier transform (FFT) and has a simple FFT-based algorithm to evaluate it.

We now show how to efficiently compute the DFST:

Theorem (DFST in terms of the FrFT)

For $\Delta p \in \mathbb{R}, \hat{p} = -n_u \dots n_u \in \mathbb{Z}$,

$\boldsymbol{x} = \Delta x \hat{\boldsymbol{x}}, \hat{\boldsymbol{x}} = -\boldsymbol{n}_x \dots \boldsymbol{n}_x, \boldsymbol{x}^* = \Delta x \hat{\boldsymbol{x}}^*, \hat{\boldsymbol{x}}^* = -\boldsymbol{n}_x \dots \boldsymbol{n}_x.$

If $\left(1-\frac{1}{\alpha}\right) = \hat{p}\Delta p, |\alpha|\Delta x \geq |1-\alpha|\Delta u,$

$$P_\alpha^{d:u_1,u_2}[L_F^d](\alpha\boldsymbol{x}) = \frac{|\Delta\boldsymbol{u}|}{\alpha^2 F^2}\frac{1}{|\boldsymbol{N}_x|} \times$$

$$\times \boldsymbol{F}_{\boldsymbol{N}_x}^{-1,-1}\left[(\boldsymbol{F}_{\boldsymbol{N}_u}^{-\Delta p\left(\frac{T_u}{T_x}\right)\hat{\boldsymbol{a}}}(\boldsymbol{F}_{\boldsymbol{N}_x}^{1,1}[\hat{L}_F^d(\cdot,\hat{\boldsymbol{u}})])(\hat{\boldsymbol{a}})])(\hat{p},\hat{p})\right](\hat{\boldsymbol{x}}).$$

If $\left(1-\frac{1}{\alpha}\right) = \hat{p}\Delta p, \hat{p} = -n_x \dots n_x \in \mathbb{Z}, |\alpha|\Delta x \leq |1-\alpha|\Delta u$

$$P_\alpha^{d:x_1,x_2}[L_F^d](\alpha\boldsymbol{x}^*) = \frac{|\Delta\boldsymbol{x}|}{|1-\alpha|^2 F^2}\frac{1}{|\boldsymbol{N}_u|} \times$$

$$\times \boldsymbol{F}_{\boldsymbol{N}_u}^{\hat{p}\Delta p\left(\frac{\Delta x}{\Delta u},\frac{\Delta x}{\Delta u}\right)}\left[(\boldsymbol{F}_{\boldsymbol{N}_x}^{-\Delta p\left(\frac{T_u}{T_x}\right)\hat{\boldsymbol{b}}}(\boldsymbol{F}_{\boldsymbol{N}_u}^{1,1}[\hat{L}_F^d(\hat{\boldsymbol{x}},.)])(\hat{\boldsymbol{b}})])(\hat{p},\hat{p})\right](\widehat{\boldsymbol{x}^*}).$$

$$\hat{\boldsymbol{a}} = -\boldsymbol{n}_x \dots \boldsymbol{n}_x, \hat{\boldsymbol{b}} = -\boldsymbol{n}_u \dots \boldsymbol{n}_u$$

with $\hat{L}_F^d(\hat{\boldsymbol{x}},\hat{\boldsymbol{u}}) = L_F^d(\hat{\boldsymbol{x}}\Delta x, \hat{\boldsymbol{u}}\Delta u).$

Even though the formulas may seem complicated, the resulting algorithm is quite simple as we will see in Section 4.

3.2.5 The rendering process

To render the focal stack images we must take care that we have to compute $\alpha^2 F^2 |\Delta\boldsymbol{x}| \mathcal{S}^d[L_F^d](\boldsymbol{x},\alpha)$. Also to show the images in the same range as the plenoptic image we have to divide the result by $|\boldsymbol{N}_u|$.

4. THE DFST ALGORITHM

We now show the algorithm for $|\alpha|\Delta x \geq |1-\alpha|\Delta u$ the other case is similar:

DFST rendering algorithm

Input: Sampled lightfield $\hat{L}_F^d(\hat{\boldsymbol{x}},\hat{\boldsymbol{u}})$ and N_x, N_u (both even).

$\hat{x}_i = -n_x \dots n_x - 1, n_x = N_x/2, \hat{u}_i = -n_u \dots n_u - 1, n_u = N_u/2.$

Output: $\hat{S}(\hat{\boldsymbol{x}},\hat{p}) = \alpha^2 F^2 |\Delta\boldsymbol{x}| \mathcal{S}^d[L_F^d](\boldsymbol{x},\alpha), \boldsymbol{x} = \hat{\boldsymbol{x}}\Delta x$

$\left(1-\frac{1}{\alpha}\right) = \hat{p}\frac{\Delta x}{n_u \Delta u}\ -n_u \leq \hat{p} \leq n_u,$

Step 0

For every $\hat{\boldsymbol{u}}$ add $\lceil n_u/2 \rceil$ zero rows under the image $\hat{L}(\cdot,\hat{\boldsymbol{u}})$ and $\lceil n_u/2 \rceil + 1$ zero rows over $\hat{L}_F^d(\cdot,\hat{\boldsymbol{u}})$. Also add $\lceil n_u/2 \rceil$ zero columns on the left of $\hat{L}_F^d(\cdot,\hat{\boldsymbol{u}})$ and $\lceil n_u/2 \rceil + 1$ zero columns on the right. Extend $\hat{\boldsymbol{u}}$ with zeros so that $\hat{L}_F^d(\cdot,\hat{u}_1,\hat{u}_2)$=0 when $\hat{u}_1$=$N_u + 1$ or $\hat{u}_2$=$N_u + 1$. We obtain then the extended lightfield $L_F^{d;ext}(\hat{\boldsymbol{x}},\hat{\boldsymbol{u}})$ with size: $N_x^{ext} = N_x + 2\lceil n_u/2 \rceil + 1$, $N_u^{ext} = N_u + 1$, $n_u = N_u/2$,and $\boldsymbol{N}_{\boldsymbol{x}}^{ext} = (N_x^{ext}, N_x^{ext})'$, $\boldsymbol{N}_{\boldsymbol{u}}^{ext} = (N_u^{ext}, N_u^{ext})'$.

Step 1

For every $\hat{\boldsymbol{u}}$ compute S1$(\hat{\boldsymbol{a}},\hat{\boldsymbol{u}})$= $(\boldsymbol{F}_{\boldsymbol{N}_x^{ext}}^{1,1}[L_F^{d;ext}(.,\hat{\boldsymbol{u}})])(\hat{\boldsymbol{a}})$

Step 2

For every $\hat{\boldsymbol{a}}$ compute S2$(\hat{\boldsymbol{a}},\hat{\boldsymbol{p}}) = (\boldsymbol{F}_{\boldsymbol{N}_u^{ext}}^{-\frac{N_u^{ext}}{n_u N_x^{ext}}\hat{\boldsymbol{a}}}[\mathrm{S1}(\hat{\boldsymbol{a}},\cdot)])(\hat{\boldsymbol{p}})$

Step 3

For every $\hat{p}$ compute S3$(\hat{\boldsymbol{x}},\hat{p})$= $\boldsymbol{F}_{\boldsymbol{N}_x^{ext}}^{-1,-1}[\mathrm{S2}(\cdot,\hat{p},\hat{p})]\ (\hat{\boldsymbol{x}})$

Step 4

Build $\hat{S}(\hat{\boldsymbol{x}},\hat{p})$ extracting from S3$(\hat{\boldsymbol{x}},\hat{p})$ the $\hat{\boldsymbol{x}}$ values $-\boldsymbol{n}_x \leq \hat{\boldsymbol{x}} \leq \boldsymbol{n}_x - 1$ and finally divide $\hat{S}(\hat{\boldsymbol{x}},\hat{p})$ by $(N_x^{ext} N_u)^2$.

The computational complexity of the algorithm above is O(n^4log(n)) (dominated by Step 1) and generates n refocused photographs. It is presented here because of its simplicity though it is not optimal. It can be shown that the diagonal values S2$(\cdot,\hat{p},\hat{p})$ can be generated in O(n^2) and that each different $\hat{p}$ set of size n costs O(n^2)[7]. Therefore we can generate log(n) of such sets without changing the complexity for a total of $n$log(n) $\hat{p}$-values. We note that [2] claims n^2 $\hat{p}$-values but the independence of the Kaiser-Bessel window size with respect to n is not proved.

5. EXPERIMENTAL RESULTS

In this section we present some experimental results of the DFST. In figure 2 we see a plenoptic image [8] and a detail from it. The image has 4096×4096 pixels and is composed of an array of 256×256 microlenses. Each microlens generates a 16×16 image.

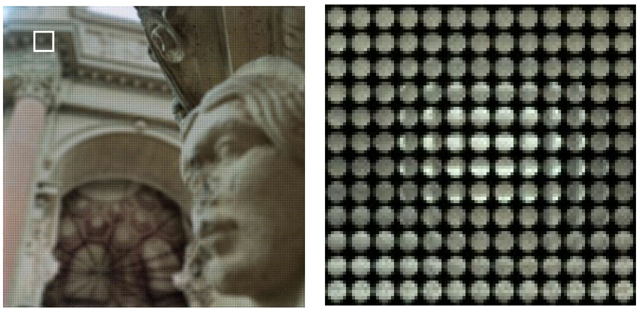

Figure 2 Plenoptic image (left) and detail from the white square in the plenoptic image (right).

Figure 3 Left: Images from the focal stack focused on the face (up) and building (down). Right: Details from the images on the left. Image size is 256×256.

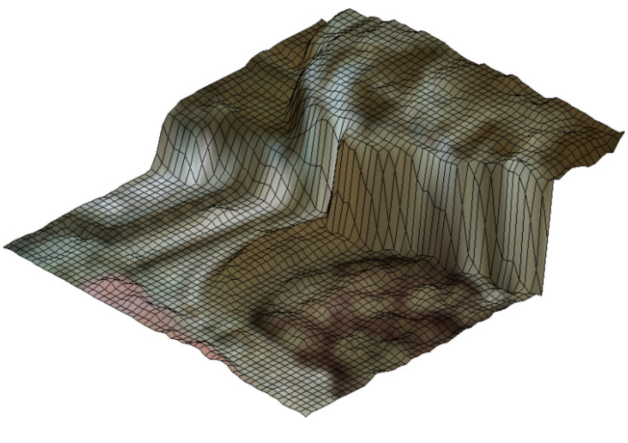

Figure 4 Depth estimation from the plenoptic image

5.1 Results for the Focal Stack

In Figure 3 we see two refocused photographs from the focal stack and a detailed region from each of them. No data preprocessing or postprocessing has been done. The size of the images on the left is 256×256. The results show that the transform produces high quality refocused images.

5.2 Results for Depth Estimation

An advantage of the DFST is that its analytic nature produces straightforward generalizations to other focal stack of interest like the modified laplacian focal stack, where each image contains the modified laplacian focus measure [9]. This focus measure is usually used in the depth from focus problem. The estimated depth from this approach can be seen in Figure 4.

6. CONCLUSIONS

In this paper we have presented a new discretization of the Photography operator based on the trigonometric interpolation of the lightfield and an approximation of the integration process. Based on that discretization we have defined the discrete focal transform whose main features are its algebraic exactness, geometric fidelity, speed and replication of the continuous properties. Computational results show the validity of the approach.

ACKNOWLEDGMENTS

This work has been supported by “Programa Nacional de Diseño y Producción Industrial (DPI 2006-07906) of the “Ministerio de Educación y Ciencia de España” and the “Fondo Europeo de Desarrollo Regional (FEDER)”

REFERENCES

[1] J. Berent and P.L. Dragotti, "Plenoptic Manifolds: Exploiting Structure and Coherence in Multiview Images," IEEE Signal Processing Mag., vol. 24, pp. 34-44, Nov. 2007.
[2] R. Ng, “Fourier slice photography,”. ACM Trans. Graph. 24(3) pp. 735-744, 2005.
[3] E. Adelson and J. Wang, “Single lens stereo with plenoptic camera”, IEEE trans. on PAMI, vol. 14, nº2, pp. 99, 1992.
[4] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, P. Hanrahan, "Light Field Photography with a Hand-Held Plenoptic Camera", Stanford Tech Report CTSR 2005-02, 2005.
[5] The Cafadis Camera: International Patent number PCT/ES2007/000046 (WIPO WO/2007/082975)
[6] A. Averbuch, R. Coifman, D. Donoho, M. Israeli, Y. Shkolnisky, I. Sedelnikov. “A framework for discrete integral transformations II - the 2D discrete Radon transform”. SIAM Journal on Scientific Computing, vol. 30(2), pp: 785-803 (2008).
[7] D. H Bailey and P. N. Swartztrauber, “The fractional Fourier transform and applications”, SIAM Review 33(3), pp.389-404, (1991).
[8] J. Meng, GPU implementation of Light-Field Photography, downloaded from: http://www.cs.virginia.edu/~jm6dg/
[9] S.K. Nayar Y. Nakagawa, “Shape from focus”, IEEE trans. on PAMI Vol. 16 nº 8 pp: 824 - 831, 1994.