Transient Imaging

Egor Larionov

University of Waterloo

February 11, 2016

Contents

Abstract

This document provides a detailed description of the mathematics behind the time-of-flight (ToF) camera transient imaging approach presented by Lin et al. in their paper titled "Fourier Analysis on Transient Imaging by Multifrequency Time-of-Flight Camera" [3].

1 The Fundamentals

1.1 Configuration

The studied transient imaging system consists of a light source, an optical sensor, a signal generator and a phase shifter. The light source is modulated by the signal generator and the sensor correlates the raw optical signal with a phase shifted signal from the signal on the output. This setup is shown in Figure 1.1

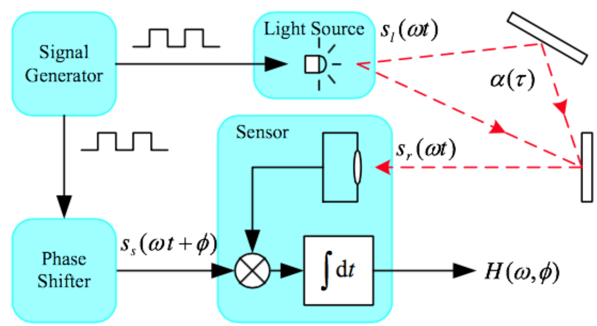

Figure 1.1: ToF Camera setup. Image taken from [3]

1.2 Definitions

In our mathematical formulation we will consider the transient image of a single pixel, since each pixel in the image is independent of the others.

Let f be the modulation frequency, that is the number of wavelengths per second. In our setup f typically ranges from 10MHz to 120MHz. Let $s_l(\omega t)$ be the light source output where $\omega := 2\pi f$. This function models the intensity of the light (a real value) produced at f pulses per second. Let P be the set of all undirected paths from the sensor to the light source. This set is scene dependent, as different materials and their configurations may affect the way light travels through the scene, so it is difficult if not impossible to see the structure of P. Consider the following functions on the set P:

$$\ell : P \to \mathbb{R}^+ \quad \text{is the } \textit{time} \text{ length of each light path}$$
$$a : P \to [0,1] \quad \text{is the attenuation of each light path}$$

Then let $S : \mathbb{R}^+ \to \mathcal{P}(P)$ be defined as

$$S(\tau) := \{p \in P : \ell(p) = \tau\}.$$

That is S is the set of all light paths with length[1] τ. Because paths of the same length are indistinguishable by our measurements, the attenuation of the light beam at time τ is given by

$$\alpha(\tau) := \int_{S(\tau)} a(p) d\mu_\tau \tag{1.1}$$

where μ_τ is some uniform probability measure on a possibly infinite dimensional space $S(\tau)$. In words, this means that

[1] We use the term length to really mean the time it takes light to travel a path of certain length.

the attenuation at time τ is the average attenuation over all light signals arriving at the sensor at time τ. This defines the transient image $\alpha(\tau)$ as the average attenuation over all light paths of length τ. It is difficult to define μ_τ exactly, let alone compute the scene dependent function α directly using equation (1.1).

However using the global illumination model, we can measure the sensor input at time t:

$$s_r(\omega t) := E_0 + \int_0^\infty \alpha(\tau) s_l(\omega(t-\tau))d\tau \qquad (1.2)$$

where $\alpha(\tau)$ is the attenuation average over all light paths p with $\ell(p) = \tau$; E_0 is the ambient illumination and $s_l(\omega(t-\tau))$ is the light signal travelled time τ and measured at time t.

To see why this is true, consider the direct illumination model where light arriving at the sensor could only come along a single light path from the source. In this model the sensor input at time t would be

$$s_r(\omega t) = a(p_0) s_l(\omega(t-\tau_0)) \qquad (1.3)$$

where $\tau_0 = \ell(p_0)$ is the length of the said light path and $a(p_0)$ is the attenuation of this path. Note that the signal arriving at the sensor at time t could only come from the signal emitted at time $t - \tau_0$.

Note that the sensor reports the "cross correlation" of the arriving signal $s_r(\omega t)$ with a phase shifted, zero-mean signal, $s_s(\omega t + \phi)$. Over an exposure of time NT where $T = 1/f$ and N is some integer, the sensor reports the image given by

$$H(\omega, \phi) = \int_0^{NT} s_r(\omega t) s_s(\omega t + \phi)\, dt \qquad (1.4)$$

2 Calibration

In a real system, $s_l(\omega t)$ and $s_s(\omega t)$ are not known exactly, so we have to acquire the discrete version of $c(\tau, \omega, \phi)$ in some other way. Since $c(\tau, \omega, \phi)$ is independent of the scene, we can measure it once and use it for different scenes. Note that with the *direct illumination* model we can measure the correlation function directly by combining equations (1.4) and (1.3) to get:

$$\begin{aligned}
H_{di}(\omega, \phi) &= \int_0^{NT} s_r(\omega t) s_s(\omega t + \phi)\, dt \\
&= \int_0^{NT} \left(E_0 + a(p_0) s_l(\omega(t-\tau_0))\right) s_s(\omega t + \phi)\, dt \\
&= E_0 \underbrace{\int_0^{NT} s_s(\omega t + \phi)\, dt}_{=0} \\
&\quad + \int_0^{NT} a(p_0) s_l(\omega(t-\tau_0)) s_s(\omega t + \phi)\, dt \\
&= a(p_0) \int_0^{NT} s_l(\omega(t-\tau_0)) s_s(\omega t + \phi)\, dt \\
&= a(p_0) c(\tau_0, \omega, \phi).
\end{aligned}$$

Now consider the following identity:

$$\begin{aligned}
c(\tau, \omega, \phi) &= \int_0^{NT} s_l(\omega(t-\tau)) s_s(\omega t + \phi)\, dt \\
&= \int_0^{NT} s_l(\omega((t + \tau_0 - \tau) - \tau_0)) s_s(\omega t + \phi)\, dt \\
&= \int_0^{NT} s_l(\omega(t' - \tau_0)) s_s(\omega(t' - \tau_0 + \tau) + \phi)\, dt' \\
&= \int_0^{NT} s_l(\omega(t' - \tau_0)) s_s(\omega t' - \omega(\tau_0 - \tau) + \phi)\, dt' \\
&= c(\tau_0, \omega, \phi + \omega(\tau_0 - \tau)).
\end{aligned}$$

This implies that we can sample any value of $c(\tau, \omega, \phi)$. So to get the discrete value of $c(\tau_i, \omega_j, \phi_k)$, we need to measure $H_{di}(\omega_j, \phi_k + \omega_j(\tau_0 - \tau_i))$, that is:

$$c(\tau_i, \omega_j, \phi_k) = \frac{H_{di}(\omega_j, \phi_k + \omega_j(\tau_0 - \tau_i))}{a(p_0)}. \qquad (2.1)$$

It remains to estimate the value of $a(p_0)$. Note that by Beer-Lambert law, the incident intensity I_0 of light is attenuated as follows:

$$I(x) = I_0 e^{-\int_0^x \beta(x')dx'},$$

where $\beta(x)$ is the attenuation coefficient. What we call "attenuation" in this document, refers to the whole exponential function[2], that is

$$a(p_0) = e^{-\int_0^{x(p_0)} \beta(x')dx'},$$

where $x(p_0)$ is the distance length of path p_0. During calibration, if we shine the light source directly into the camera sensor, or perhaps place the sensor and light source close together and shine at a mirror such that the light bounces directly into the sensor, we may assume that the attenuation coefficient is constant along the light path, thus simplifying our attenuation expression to:

$$a(p_0) = e^{-\beta x(p_0)}$$

where β is now the attenuation coefficient of air. According to [2] this attenuation coefficient is on the order of 10^{-1} dB/km, meaning that given that $x(p_0) = 1$ metre, then

$$a(p_0) \approx 0.99990$$

which is sufficiently close to 1. Thus in the calibration step we use the approximation: $a(p_0) = 1$, giving us the calibration data:

$$c(\tau_i, \omega_j, \phi_k) = H_{di}(\omega_j, \phi_k + \omega_j(\tau_0 - \tau_i)). \qquad (2.2)$$

3 Fourier Analysis

In general, we use the global illumination model to describe the output of the sensor. This can be written by combining

[2] This is called *transmittance* in physics literature

equations (1.4) and (1.2):

$$
\begin{aligned}
H(\omega,\phi) &= \int_0^{NT} s_r(\omega t)s_s(\omega t+\phi)\ dt \\
&= \int_0^{NT}\left(E_0+\int_0^{\infty}\alpha(\tau)s_l(\omega(t-\tau))d\tau\right)s_s(\omega t+\phi)\ dt \\
&= E_0\underbrace{\int_0^{NT} s_s(\omega t+\phi)\ dt}_{=0} \quad \because s_s \text{ is zero-mean} \\
&\quad + \int_0^{NT}\!\!\int_0^{\infty}\alpha(\tau)s_l(\omega(t-\tau))s_s(\omega t+\phi)\ d\tau\ dt \\
&= \int_0^{\infty}\alpha(\tau)\underbrace{\int_0^{NT} s_l(\omega(t-\tau))s_s(\omega t+\phi)\ dt}_{=:\,c(\tau,\omega,\phi)}\ d\tau \\
&= \int_0^{\infty}\alpha(\tau)c(\tau,\omega,\phi)d\tau. \qquad (3.1)
\end{aligned}
$$

If we temporarily assume that the functions $s_l(\omega t)$ and $s_s(\omega t+\phi)$ are perfectly sinusoidal, it can be derived that, loosely speaking, the transient image is the inverse Fourier transform of the sensor output. Using this fact as motivation, we shall determine the inverse Fourier transform of the sensor output in terms of functions that we know, and attempt to extract the transient image from the resulting expression.

First, let's define the following complex functions:

$$\tilde{H}(\omega) := H(\omega,0)+iH(\omega,\tfrac{\pi}{2}) \qquad (3.2)$$

$$\tilde{c}(\tau,\omega) := c(\tau,\omega,0)+ic(\tau,\omega,\tfrac{\pi}{2}) \qquad (3.3)$$

Then using equation (3.1), we can write:

$$\tilde{H}(\omega) = \int_0^{\infty}\alpha(\tau)\tilde{c}(\tau,\omega)d\tau.$$

Why exactly we use the modified definitions in equations (3.2) and (3.3) will become clear later.

Since $c(\tau,\omega,\phi)$ is a periodic function, it has a natural Fourier series:

$$c(\tau,\omega,\phi) = \tilde{A}_0(\omega)+\sum_{n=\pm1}^{\pm\infty}\tilde{A}_n(\omega)e^{-in(\omega\tau+\phi)},$$

where $\tilde{A}_n$ are the Fourier coefficients and we explicitly factor the phase dependence into the exponential. Note that

$$
\begin{aligned}
\tilde{A}_0(\omega) &= \frac{1}{T}\int_0^{T} c(\tau,\omega,\phi)e^{i(0)(\omega\tau+\phi)} \\
&= \frac{1}{T}\int_0^{T}\!\!\int_0^{NT} s_l(\omega(t-\tau))s_s(\omega t+\phi)\ dt\ d\tau \\
&= \frac{1}{T}\int_0^{NT} s_s(\omega t+\phi)\underbrace{\int_0^{T} s_l(\omega(t-\tau))\ d\tau}_{=\,C\ \text{(const.)}}\ dt \\
&= \frac{C}{T}\int_0^{NT} s_s(\omega t+\phi)\ dt \\
&= 0
\end{aligned}
$$

where in the third step, the integral over the period evaluates to a constant because it is translation independent[3], and the last step follows because $s_s(\omega t)$ is a zero-mean function.

Then it follows from equation (3.3) that

$$
\begin{aligned}
\tilde{c}(\tau,\omega) &= \sum_{n=\pm1}^{\pm\infty}\tilde{A}_n(\omega)e^{-in\omega\tau}+i\sum_{n=\pm1}^{\pm\infty}\tilde{A}_n(\omega)e^{-in(\omega\tau+\frac{\pi}{2})} \\
&= \sum_{n=\pm1}^{\pm\infty}\left(\tilde{A}_n(\omega)+i\tilde{A}_n(\omega)e^{-in\frac{\pi}{2}}\right)e^{-in\omega\tau} \\
&= \sum_{n=\pm1}^{\pm\infty}\tilde{B}_n(\omega)e^{-in\omega\tau} \qquad (3.4)
\end{aligned}
$$

where

$$\tilde{B}_n := \tilde{A}_n(\omega)(1+ie^{-in\frac{\pi}{2}}). \qquad (3.5)$$

For future derivations we compute the first two coefficients:

$$
\begin{aligned}
\tilde{B}_{-1} &= \tilde{A}_{-1}(\omega)(1+ie^{i\frac{\pi}{2}}) \\
&= \tilde{A}_{-1}(\omega)(1+i(i)) \\
&= \tilde{A}_{-1}(\omega)(1-1) \\
&= 0 \qquad (3.6)
\end{aligned}
$$

and

$$
\begin{aligned}
\tilde{B}_{1} &= \tilde{A}_{1}(\omega)(1+ie^{-i\frac{\pi}{2}}) \\
&= \tilde{A}_{1}(\omega)(1+1) \\
&= 2\tilde{A}_{1}(\omega). \qquad (3.7)
\end{aligned}
$$

It follows from (3.4) that

$$
\begin{aligned}
\tilde{H}(\omega) &= \int_0^{\infty}\alpha(\tau)\tilde{c}(\tau,\omega)d\tau \\
&= \sum_{n=\pm1}^{\pm\infty}\tilde{B}_n(\omega)\int_0^{\infty}\alpha(\tau)e^{-in\omega\tau}d\tau \\
&= \sum_{n=\pm1}^{\pm\infty}\frac{\tilde{B}_n(\omega)}{n}\int_0^{\infty}\alpha(\tfrac{\tau}{n})e^{-i\omega\tau}d\tau \qquad (3.8)
\end{aligned}
$$

where the last step follows by a change of variables such as $\tau' = \tau n$.

For convenience, we define $\alpha_n(\tau) := \alpha(\frac{\tau}{n})$, and

$$b_n(\tau) := (\mathcal{F}^{-1}\tfrac{\tilde{B}_n}{n})(\tau) = \int_{\mathbb{R}}\tfrac{\tilde{B}_n(\omega)}{n}e^{i\omega\tau}\ d\omega. \qquad (3.9)$$

That is b_n is the inverse Fourier transform of $\tilde{B}_n/n$. We can

[3] A property of periodic functions.

now write the Fourier transform of $\tilde{H}$ as follows:

$$\begin{aligned}
\tilde{h}(\tau) &:= (\mathcal{F}^{-1}\tilde{H})(\tau) \\
&= \int_{\mathbb{R}} \sum_{n=\pm 1}^{\pm\infty} \frac{\tilde{B}_n(\omega)}{n} \int_0^\infty \alpha_n(\tau') e^{-i\omega\tau'}\, d\tau'\; e^{i\omega\tau}\, d\omega \\
&= \sum_{n=\pm 1}^{\pm\infty} \int_0^\infty \alpha_n(\tau') \underbrace{\int_{\mathbb{R}} \frac{\tilde{B}_n(\omega)}{n} e^{i\omega(\tau-\tau')}\, d\omega}_{b_n(\tau-\tau')}\, d\tau' \\
&= \sum_{n=\pm 1}^{\pm\infty} \int_0^\infty \alpha_n(\tau') b_n(\tau-\tau') d\tau' \\
&= \sum_{n=\pm 1}^{\pm\infty} (\alpha_n * b_n)(\tau) \\
&= (\alpha * b_1)(\tau) + \sum_{n=\pm 2}^{\pm\infty} (\alpha_n * b_n)(\tau)
\end{aligned} \tag{3.10}$$

where $(*)$ is the *convolution* operation, and by (3.6) and (3.9), the b_{-1} term vanishes.

Expression (3.10) tells us that the inverse Fourier transform of $\tilde{H}$ consists of the transient image α convolved with some windowing function b_1 mixed with some dilated components α_n convolved by other windowing functions b_n. If we had not used definitions (3.2) and (3.3), then α_{-1} wouldn't vanish, and we would be stuck with a significant source of error. As we will see next, it is possible to factor the term b_1 and remove the dilated components b_n to isolate the transient image α.

3.1 Limited Frequency Window

Note that we can only gather the data for $\tilde{H}(\omega)$ for a limited range of frequencies, namely $f \in [f_L, f_H]$ (recall that $\omega = 2\pi f$) where in our case, $f_L \approx 10$MHz and $f_H \approx 120$MHz. Now if we define:

$$R_L(\omega) := \begin{cases} 1 & \text{for } \omega \in [-2\pi f_L, 2\pi f_L] \\ 0 & \text{otherwise} \end{cases}$$

$$\text{and} \tag{3.11}$$

$$R_H(\omega) := \begin{cases} 1 & \text{for } \omega \in [-2\pi f_H, 2\pi f_H] \\ 0 & \text{otherwise} \end{cases}$$

Then what we actually measure for $\tilde{H}(\omega)$ will be of the form

$$\tilde{H}_c(\omega) = \tilde{H}(\omega)(R_H(\omega) - R_L(\omega)) \tag{3.12}$$

Then by the convolution property of the Fourier transform[4], we have that

$$\tilde{h}_c = \tilde{h} * (r_H - r_L) \tag{3.13}$$

where $\tilde{h}_c$, r_H and r_L are the Fourier transforms of $\tilde{H}_c$, R_H and R_L respectively.

This section concludes the model describing the systematic errors in the setup of this transient imaging system as prescribed in [3]. The systematic errors described here, are specific to taking the inverse Fourier transform of the gathered data (after data rectification described below). This means that an alternative method for extracting the transient image from the measured data (for instance using non-convex optimization as in [1]) may not be subject to these errors.

In summary, measuring the camera signal for phases 0 and $\pi/2$ as in equation (3.2), gives us the raw data $\tilde{H}_c(\omega)$, which we need to process in order to extract the transient image α.

4 Extracting the Transient Image

4.1 Data Rectification

First step in solving equation (3.10) for α is to factor out b_1 to isolate α in the sum. So let

$$\tilde{R} := \tilde{H}_c / \tilde{B}_1 \tag{4.1}$$

Then taking the inverse Fourier transform of $\tilde{R}$ gives us

$$r := \mathcal{F}^{-1}(\tilde{R}) = \mathcal{F}^{-1}(\tilde{H}_c) * \mathcal{F}^{-1}(1/\tilde{B}_1) = \tilde{h}_c * b_1^{-1}$$

by the convolution property, where

$$b_1^{-1} := \mathcal{F}^{-1}(1/\tilde{B}_1), \tag{4.2}$$

which represents the convolution inverse of b_1. Indeed, for any absolutely integrable function $f : \mathbb{R} \to \mathbb{R}$, it follows from the convolution property that $f = f * b_1 * b_1^{-1}$.

Thus we have effectively isolated α in equation (3.10). Observe:

$$\begin{aligned}
r &= \tilde{h}_c * b_1^{-1} \\
&= \tilde{h} * (r_H - r_L) * b_1^{-1} \\
&= \left(\alpha * b_1 + \sum_{n=\pm 2}^{\pm\infty} \alpha * b_n \right) * (r_H - r_L) * b_1^{-1} \\
&= \left(\alpha + \sum_{n=\pm 2}^{\pm\infty} \alpha_n * b_n * b_1^{-1} \right) * (r_H - r_L)
\end{aligned} \tag{4.3}$$

where the last step follows from the commutativity and distributivity of convolution. This process is called "data rectification". It now remains to recover the missing frequency spectra modelled by the term $(r_H - r_L)$, and remove the dilated components originating from the sum term,

$$\sum_{n=\pm 2}^{\pm\infty} \alpha_n * b_n * b_1^{-1}.$$

Note that in order to perform the operation in equation (4.1), we have to know $\tilde{B}_1$. This is determined at the calibration step as follows.

4.2 Calibration Data Fitting

It is not enough for us to know the values of the correlation function from (2.2). We must determine its Fourier coefficients in order to effectively use it in extracting the transient image as in (4.1), for instance. Thus we fit the calibration data to a truncated Fourier series:

$$A_0(\omega) + \sum_{n=1}^{n_0} A_n(\omega) e^{-i(n\omega(\tau+\tau_0) - \phi_n(\omega))}$$

[4] See Appendix A for a proof.

where τ_0 is the flying time of the shortest ray path, and n_0 is some small integer indicating the number of coefficients we fit the data to. This approximation gives us the approximate Fourier coefficients to the true correlation function as

$$\tilde{A}_n(\omega) \approx A_n(\omega)e^{i\phi_n(\omega)}$$

for $n > 0$. In this decomposition, A_n and ϕ_n are all real valued functions for $n > 0$ and A_0 is a complex function near the origin.

To fit the calibration data, the authors of [3] propose a gradient descent method.

This process gives us approximations for the first n_0 Fourier coefficients $\tilde{B}_n$ and hence b_n, by definitions (3.5) and (3.9). In particular we now have enough information to perform the "data rectification" step mentioned before.

4.3 Dilation Component Removal

The authors of [3] propose a procedure to remove the dilated components from equation (4.3). Their strategy is to subtract off the dilated components by iteratively subtracting a dilation of the known function r. This process does alleviate this source of error, but does not remove it completely, and leaves traces of high order error terms. The issue here is the limited frequency spectra imposed on the function r, which we will describe in detail in the next subsection.

First we will describe the algorithm for removing dilated components assuming no limited frequency window restriction. That is, we will remove the dilated components in (4.3) assuming no convolution with the $(r_H - r_L)$ term. So we start with the expression

$$r(\tau) = \alpha(\tau) + \sum_{n=2}^{n_0} (\alpha_n * b_n * b_1^{-1})(\tau), \tag{4.4}$$

where the higher order terms in the sum are 0 since we have fitted only n_0 of b_n terms to the correlation data, as described in the previous section. Then note that $\alpha(\tau) = 0$ for all $\tau \in [0, \tau_0)$ where τ_0 is the shortest flying time of light in the scene (e.g. from the light source to the closest object in the scene and back), because light hasn't reached the sensor yet at for those times. Then we note that

$$\alpha_n(\tau) = \alpha\left(\frac{\tau}{n}\right) = 0 \qquad \forall \tau \in [0, n\tau_0).$$

This way we can identify the domain of contribution from each term in (4.4). In particular, the term with α_2 may not be zero for $\tau \geq 2\tau_0$; the term with α_3 may not be zero for $\tau \geq 3\tau_0$ and so on. We can illustrate the contributions in the following diagram:

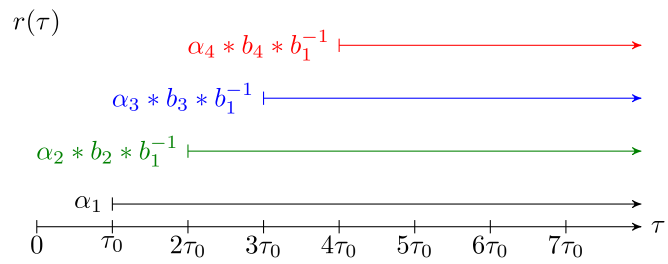

Figure 4.1: Contribution of each term on RHS of (4.4) to $r(\tau)$.

Now observe that when we dilate r by 2, the contributions to r will also dilate, since

$$\alpha_1(\tau/2) = \alpha_2(\tau), \tag{4.5}$$

which is zero for $\tau \in [0, 2\tau_0)$. Similarly any contribution from α_k will be dilated and its support will be shifted to the right by 2. This can be visualized in the following diagram:

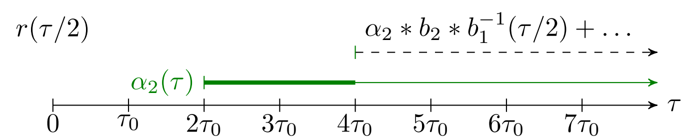

Figure 4.2: Contributions to $r(\tau/2)$, in which we use the **b**olded interval contribution of α_2 to remove from r.

Thus if we dilate r by k (to get $r(\tau/k)$) and subtract the result convolved with $b_k * b_1^{-1}$ from our original function r on the domain $\tau \in [k\tau_0, 2k\tau_0]$, for $k \in \{2, \dots, n_0\}$, we would, in effect, remove the error terms for parts of the domain:

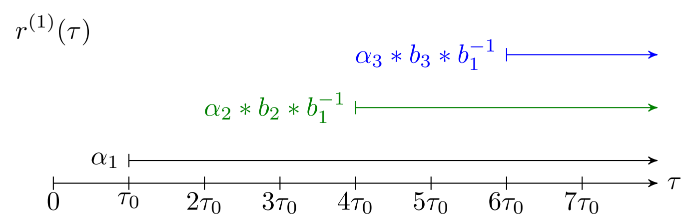

Figure 4.3: Result of one iteration of removing dilated components from r to get $r^{(1)}$.

In a similar fashion, we may compute $r^{(1)}(\tau/k)$ and subtract it from $r^{(1)}(\tau)$ on the domain $\tau \in [2k\tau_0, 4k\tau_0)$ for $k \in \{2, \dots, n_0\}$ to get $r^{(2)}$ as shown below:

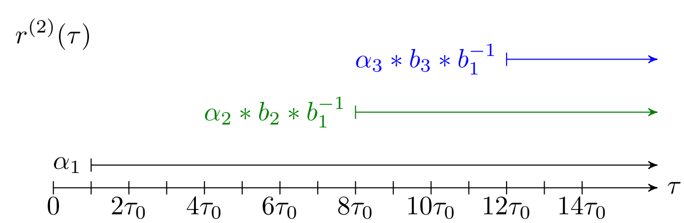

Figure 4.4: Another iteration, showing more of the domain for clarity.

It is enough to remove enough dilated components such that $r^{(k)}(\tau) = \alpha(\tau)$ for τ less than the desired duration of the transient image. Note that the required n_0 may depend on this duration, if one wants to remove the errors entirely.

Algorithm 1 summarizes this process, where we assume that n_0 is chosen such that $\tau' \in [0, n_0\tau_0)$, where τ' is the desired duration of the transient image. In addition, it's not necessary to compute any values outside of the $[0, n_0\tau_0)$ range.

Algorithm 1 Dilated Component Removal Algorithm

Input: n_0 and b_n computed from (3.9);
b_1^{-1} computed from (4.2);
r gathered from the data rectification step; and
τ_0 being the shortest flying time in the scene.
Output: $\alpha(\tau)$ for $\tau \in [0, n_0\tau_0)$.

```
1:  r^(0)(τ) ← r(τ) for all τ ∈ [0, n_0 τ_0).
2:  for all j = 0, ..., log_2(n_0) − 1 do
3:    r^(j+1)(τ) ← r^(j)(τ) for all τ ∈ [0, n_0 τ_0)
4:    for all k = 2, ..., n_0 do
5:      for all τ ∈ [(j+1)kτ_0, 2(j+1)kτ_0) do
6:        Interpolate value of r^(j)(τ/k)
7:        r^(j+1)(τ) ← r^(j+1)(τ) − r^(j+1) * b_k * b_1^(−1)(τ/k)
8:      end for
9:    end for
10: end for
```

In the process of dilating the intermediate functions $r^{(k)}$, we would also need to interpolate the discrete values of this function, which will contribute a small interpolation error to the result.

The algorithm given above is identical to the one presented in [3] and would work under our initial assumption that the limited frequency window is omitted. Now we can observe what happens if we apply this algorithm without this assumption. Now we remove the dilated components in equation (4.3). Define

$$\beta_n := \alpha_n * (r_H - r_L),$$

then equation (4.3) simplifies to

$$r(\tau) = \beta_1(\tau) + \sum_{n=2}^{n_0} (\beta_n * b_n * b_1^{-1})(\tau), \tag{4.6}$$

which is analagous to equation (4.4) with β_n replacing α_n. However, now our deduction in (4.5) is now invalid with β_n. Define $w : \mathbb{R} \to \mathbb{R}$ by

$$w(\tau) := (r_H - r_L)(\tau) \qquad \text{and}$$
$$w_{1/2}(\tau) := (r_H - r_L)(\tau/2)$$

for convenience, and observe that for our domain of interest, $\tau \in [2\tau_0, 4\tau_0)$, we have

$$\begin{aligned}
\beta_2(\tau) - \beta_1(\tau/2) &= \alpha_2 * w(\tau) - \alpha_1 * w(\tau/2) \\
&= \alpha_2 * w(\tau) - \int_{\mathbb{R}} \alpha_1(\tfrac{\tau'}{2}) w(\tfrac{\tau}{2} - \tfrac{\tau'}{2})\, d\tfrac{\tau'}{2} \\
&= \alpha_2 * w(\tau) - \tfrac{1}{2} \int_{\mathbb{R}} \alpha_2(\tau') w(\tfrac{1}{2}(\tau - \tau'))\, d\tau' \\
&= \alpha_2 * w(\tau) - \tfrac{1}{2} \alpha_2 * w_{1/2}(\tau) \\
&= \alpha_2 * (w - \tfrac{1}{2} w_{1/2})(\tau) \\
&\neq 0
\end{aligned}$$

This is exactly the operation needed to remove the dilated components, and it fails to do so in this case. This would take place in step 7 of Algorithm 1.

4.4 Frequency Spectra Recovery

Assuming that we have eliminated the dilated components completely (i.e. ignoring the error described in the previous section), we will be left with the function

$$\beta = \alpha * (r_H - r_L). \tag{4.7}$$

That is the transient image subject to errors due to a limited frequency window. It remains to solve equation (4.7) for α. Unfortunately we cannot simply convolve both sides of the equation by the convolution inverse of $r_H - r_L$ to solve this equation, because the inverse is undefined at the zeros of $r_H - r_L$. The authors of [3] don't have a solution for this problem but they propose a heuristic to solve a simpler problem, which assumes $R_H(\omega) = 1$ for all $\omega \in \mathbb{R}$. This assumption ignores the lack of large frequency data in $\tilde{H}_c$ and thus causes blurring in the solution $\alpha(\tau)$ as mentioned in [3].

Instead of solving (4.7), the heuristic approach solves the simplified equation

$$\beta = \alpha - \alpha * r_L \tag{4.8}$$

for α. When $\alpha * r_L$ is subtracted from α, the peaks in α are being sunk locally (see Figure 7 in [3] for a visualization). The idea is to iteratively add the missing part to β until the result is positive and thus resembles α as it is also positive on its domain (because α is the transient image). Given a small parameter $\theta > 0$, the algorithm can be described as follows:

1. Let $\beta_{old} := \beta$.

2. Compute

$$\beta_{pk}(\tau) := \begin{cases} \beta_{old}(\tau) & \text{if } \beta_{old}(\tau) > \theta \\ 0 & \text{otherwise} \end{cases}$$

3. Compute $\beta_{new}(\tau) := \beta(\tau) + \beta_{pk} * r_L(\tau)$.

4. If $\beta_{new}(\tau) > 0$ for all τ, then β_{new} is our final approximation to α; otherwise, let $\beta_{old} := \beta_{new}$ and repeat steps 2 to 4.

The parameter θ exists to prevent boosting noise into the result.

The authors of [3] claim that indeed $\beta < \alpha$, since $\alpha > 0$, however since α is convolved with r_L in the subtraction in (4.8), it is unclear whether this assertion holds. In fact, r_L is not positive on all of its domain, and although it is a real function[5], its first zero occurs at 25ns for $f_L = 10$MHz, meaning that depending on the shape of α, $\beta(\tau)$ may, in theory, fall above $\alpha(\tau)$ if $(\alpha * r_L)(\tau) < 0$.

In addition, suppose that $\theta = 0$ and after some number of iterations n, we have that $\beta^n_{new} > 0$, where β^n_{new} is the result of the algorithm after nth iteration. Then observe that after the next iteration, we will have

$$\beta^{n+1}_{new} = \alpha - \alpha * r_L + \beta^n_{new} * r_L.$$

Also, after another iteration we will have

$$\begin{aligned}\beta^{n+2}_{new} &= \beta + \beta^{n+1}_{new} * r_L \\ &= \alpha - \alpha * r_L + (\beta + \beta^n_{new} * r_L) * r_L \\ &= \alpha - \alpha * r_L + \alpha * r_L - \alpha * r_L * r_L + \beta^n_{new} * r_L * r_L \\ &= \alpha - \alpha * r_L * r_L + \beta^n_{new} * r_L * r_L.\end{aligned}$$

It follows from the convolution property that $r_L * r_L = r_L$ since $R_L^2 = R_L$. So we have that

$$\beta^{n+2}_{new} = \alpha - \alpha * r_L + \beta^n_{new} * r_L = \beta^{n+1}_{new},$$

which indicates that after β^n_{new} reaches a positive value on all of its domain, then further iterations will have no effect on the result. This is the reason why the authors in [3] stop the algorithm when $\beta_{new} > 0$ on its domain. Furthermore there is no guarantee that $\alpha * r_L = \beta^n_{new} * r_L$ at this point, so this iteration does not approach the target solution, α.

5 Acknowledgements

This was work was done for a seminar course on computer graphics supervised by professors Stephen Mann and William Cowan. I'd like to thank all the students who took the course with me for diligently sitting through my lectures, especially Andrew Tinits, who donated his time to making helpful corrections in these notes.

References

[1] Heide, F., Hullin, M. B., Gregson, J., and Heidrich, W. Low-budget transient imaging using photonic mixer devices. *ACM Trans. Graph. 32*, 4 (July 2013), 45:1–45:10.

[2] Kim, I. I., McArthur, B., and Korevaar, E. J. Comparison of laser beam propagation at 785 nm and 1550 nm in fog and haze for optical wireless communications. vol. 4214, pp. 26–37.

[3] Lin, J., Liu, Y., Hullin, M. B., and Dai, Q. Fourier analysis on transient imaging by multifrequency time-of-flight camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2014), p. to appear.

[5] A real and even periodic function, such as R_L, has a real (inverse) Fourier transform.

A Proof of the Convolution Property of the Fourier Transform

The convolution property of the Fourier transform states that the Fourier transform of the convolution of two functions in time space is equal the product of the individual Fourier transforms of these functions. We use $\mathcal{F}$ to denote the Fourier transform:

$$\mathcal{F}(f)(\omega) := \int_{\mathbb{R}} f(t)e^{-i\omega t}dt$$

for an absolutely integrable[6] function $f : \mathbb{R} \to \mathbb{R}$. For the remainder of this section, assume that f and g are two absolutely integrable functions, where $F(\omega) := \mathcal{F}(f)(\omega)$ and $G(\omega) := \mathcal{F}(g)(\omega)$ are their Fourier transforms.

Proposition A.1 (Convolution Property)**.**

$$\mathcal{F}(f * g) = \mathcal{F}(g)\mathcal{F}(f)$$

or with the frequency parameter,

$$\mathcal{F}(f * g)(\omega) = F(\omega)G(\omega).$$

Proof.

$$\begin{aligned}\mathcal{F}(f * g)(\omega) &= \int_{\mathbb{R}} f * g(t)e^{-i\omega t}dt \\ &= \int_{\mathbb{R}}\int_{\mathbb{R}} f(t')g(t-t')dt'\, e^{-i\omega t}dt \\ &= \int_{\mathbb{R}} f(t')\int_{\mathbb{R}} g(t-t')e^{-i\omega(t-t')}dt\; e^{-i\omega t'}dt' \\ &= \int_{\mathbb{R}} f(t')\underbrace{\int_{\mathbb{R}} g(t)e^{-i\omega t}dt}_{G(\omega)}\; e^{-i\omega t'}dt' \\ &= G(\omega)\int_{\mathbb{R}} f(t')e^{-i\omega t'}dt' \\ &= G(\omega)F(\omega)\end{aligned}$$

□

The fourth equality in the proof above follows from the translation independence of an integral of a periodic function.

[6] It is only a *sufficient* condition for a function to be absolutely integrable, in order for its Fourier transform to exist. Recall that f is absolutely integrable if and only if $\int_{\mathbb{R}} |f(t)|dt < \infty$.