

Fibonacci Exposure Bracketing for High Dynamic Range Imaging

Mohit Gupta
Columbia University
New York, NY 10027
mohitg@cs.columbia.edu

Daisuke Iso
Columbia University
New York, NY 10027
iso@cs.columbia.edu

Shree K. Nayar
Columbia University
New York, NY 10027
nayar@cs.columbia.edu

Abstract

Exposure bracketing for high dynamic range (HDR) imaging involves capturing several images of the scene at different exposures. If either the camera or the scene moves during capture, the captured images must be registered. Large exposure differences between bracketed images lead to inaccurate registration, resulting in artifacts such as ghosting (multiple copies of scene objects) and blur. We present two techniques, one for image capture (Fibonacci exposure bracketing) and one for image registration (generalized registration), to prevent such motion-related artifacts. Fibonacci bracketing involves capturing a sequence of images such that each exposure time is the sum of the previous $N(N > 1)$ exposures. Generalized registration involves estimating motion between sums of contiguous sets of frames, instead of between individual frames. Together, the two techniques ensure that motion is always estimated between frames of the same total exposure time. This results in HDR images and videos which have both a large dynamic range and minimal motion-related artifacts. We show, by results for several real-world indoor and outdoor scenes, that the proposed approach significantly outperforms several existing bracketing schemes.

1. Introduction

High dynamic range (HDR) imaging is the process of capturing scenes with a larger intensity range than what conventional sensors can capture. Because HDR images faithfully capture details in both dark and bright parts of the scene, they are desirable in surveillance, astronomy, medical imaging, and more recently, even consumer photography. Exposure bracketing [9, 2] is the most popular technique for HDR digital imaging. The basic idea is to capture multiple images of the same scene with different exposures. While each captured image has a low dynamic range (LDR), a single HDR image is generated by merging the exposure-bracketed LDR frames. Because of its ease of implementation, bracketing for HDR is now available as a standard feature in most digital cameras, including cell-phones.

Despite its simplicity, exposure bracketing is not used in several real-world scenarios because it is prone to errors when there is scene or camera motion. In order to compensate for motion, the bracketed images are registered before merging into the HDR image. Registration is performed by computing motion information between adjacent frames.



Exponential bracketing and conventional registration Fibonacci bracketing and generalized registration

Figure 1. (Left) HDR image computed using conventional registration has blur and ghosting due to camera motion. Each of the three candle flames (inset) has multiple copies. (Right) HDR image computed using the proposed approach. Zoom in for details.

However, since the frames have different exposure times¹, they have different amounts of motion blur and noise. These differences cannot be removed by normalizing the image intensities by their exposure times². Because of this, even after normalization, image features are not preserved across frames, and motion information cannot be computed reliably. This results in artifacts, such as *ghosting* (multiple copies of scene objects), blur and distortions. An example is shown in Figure 1. Such artifacts often negate the quality enhancement that is brought about in the image by capturing a wide dynamic range. This presents a fundamental trade-off - while a large difference in image exposures is required to capture a wide intensity dynamic range, it also results in strong motion-related artifacts.

In this paper, we present new exposure bracketing and image registration techniques for handling scene and camera motion while also capturing a wide dynamic range (DR). The key idea is to compute motion information between sums of contiguous sets of frames. We call this *generalized registration*. This is different from conventional registration, where motion is computed between individual frames. We propose an exposure bracketing scheme called *Fibonacci bracketing* where the exposure times follow the Fibonacci property, i.e., each exposure time is the sum of the previous $N(N > 1)$ exposure times. Together, Fibonacci bracketing and generalized registration ensure that

¹In this paper, image exposure is changed only by varying the camera shutter-time. Variations in camera aperture and gain have also been used for capturing HDR images [14, 5].

²In all the simulations and experiments in the paper, wherever necessary, the captured images were normalized by their exposures before motion estimation in order to maintain brightness constancy between them.

motion is *always* estimated between frames of the same total exposure. The exposure times in a Fibonacci sequence grow exponentially, thus capturing a large DR as well. Figure 1 (Right) shows an image computed using the proposed techniques. It captures the wide DR of the scene and has negligible motion-related artifacts.

Hardware Prototype: For generalized registration, a sensor that allows exposure bracketing with a negligible inter-frame time-gap is required. Although most current cameras support exposure bracketing, there is often a large inter-frame gap (50 – 200ms). We implemented our techniques on a machine vision camera which is triggered externally using a micro-controller based circuit. This hardware setup achieves a small inter-frame gap of 0.1ms while exposure times are changed from one frame to the next. We show results for several real-world scenes, both indoors and outdoors, captured during different times of the day and having a wide range of motion characteristics. For the same time-budget, Fibonacci bracketing and generalized registration produce images of significantly higher quality as compared to existing techniques. We also extend our techniques to capture HDR video at up to 15 fps while adapting the bracketing sequence to scene brightness and motion.

Scope and Contributions: Our contributions are techniques for image capture and registration that mitigate artifacts due to *inter-frame* motion and different amount of motion blur between frames. These techniques do not reduce motion blur caused by *intra-frame* motion. We use existing deblurring methods to reduce blurring in the bracketed frames. Our techniques are robust to non-linear camera intensity response and small bit-depths, making them especially attractive for use in inexpensive cell-phone and point-and-shoot cameras. The proposed approach does not require any modifications to the optics. Because of its simplicity, our method is especially suited for implementation on compact cell-phone cameras, for which, low-light and low-dynamic-range are known problems. This makes our work particularly pertinent as cell-phone cameras are expected to dominate consumer imaging in the next five years.

2. Related Work

Exposure Bracketing: One of the most widely used bracketing schemes for HDR imaging is the *exponential scheme* [9, 2], where a sequence of images with exponentially increasing shutter times are used. Kang *et al.* [6] proposed a bracketing sequence of alternating short and long exposures. In this approach, an HDR image is computed by registering three LDR images. Later, Zhang *et al.* [21] proposed capturing and registering a sequence of several very short and same exposure images. We propose using exposure times that have the Fibonacci property, i.e., each exposure is the sum of previous $N(N > 1)$ exposures.

Recently, there has been a lot of work in devising scene-adaptive exposure bracketing techniques [5]. These techniques attempt to maximize the signal-to-noise-ratio of the final HDR image by adapting the bracketing sequence to the scene’s brightness distribution. All these techniques assume the scene and camera to be static. In contrast, our goal is to

handle dynamic camera and scenes.

Post-processing for Ghost Removal: In order to remove the ghosting artifacts in HDR images, several post-processing techniques have been proposed [7, 3]. These methods attenuate the contribution of pixels belonging to moving objects in the final HDR image. While these approaches reduce ghosting, the moving objects may not have HDR content if the inter-frame motion is large. Recently, Sen *et al.* [15] proposed creating the HDR image by enforcing its consistency with the bracketed images in a patch-based optimization procedure. Our focus is different than the above post-processing techniques - *it is on acquiring images* so that ghosting artifacts can be prevented. The above techniques can be used in a complementary fashion to our approach in order to remove any residual artifacts.

Hardware Modifications: Several approaches have been proposed to increase the DR by making hardware (optical and electronic) modifications to the camera. These include using an array of neutral density filters [11] to spatially modulate light before reaching the sensor, splitting the light inside the camera using beam-splitters [19], and placing optical filters in front of the camera [13]. These systems requiring hardware modifications are often expensive and inaccessible to consumers. For most consumer cameras, especially the point-and-shoot and cell-phone ones, exposure bracketing remains the cheapest and the most viable HDR imaging option.

3. What are Good Exposure Bracketing Schemes for HDR Imaging?

Given a time-budget T for acquiring a single HDR image, an exposure bracketing sequence is defined as a set of frame exposures $E = \{e_1, e_2, \dots, e_K\}$ such that:

$$\sum_{i=1}^K e_i = T - (K - 1)\delta, \quad (1)$$

where δ is the inter-frame time gap due to sensor read-out delay. The captured exposure-bracketed LDR frames $\{f_1, f_2, \dots, f_K\}$ constitute the exposure stack. The maximum number of frames K is constrained by the maximum frame rate of the camera. For example, for an $F = 300$ frames-per-second camera and a time budget of $T = 120ms$, a maximum of $K = \frac{F * T}{1000} = 36$ LDR frames can be captured for a single HDR image.

Given a time-budget, there are infinite possible bracketing sequences. Which bracketing sequence achieves the highest quality HDR image? The dynamic range achieved by a bracketing scheme is given as [11]:

$$DR = \log \frac{I_{max} e_{max}}{I_{min} e_{min}}, \quad (2)$$

where e_{max} and e_{min} are the maximum and the minimum exposures in the bracketing scheme, respectively. I_{max} (determined by the sensor’s full well capacity) and I_{min} (determined by the sensor read noise) are the maximum and minimum signals detectable by the sensor. From Eq. 2, it is clear that in order to maximize the dynamic range, a bracketing scheme should have a large range of exposures so that

the ratio $\frac{e_{max}}{e_{min}}$ is maximized. The LDR frames in the exposure stack captured using such bracketing sequences will have large differences in exposures.

On the other hand, if there is camera or scene motion, large differences in LDR frame exposures can lead to incorrect registration. Why does this happen? This is because image registration techniques work best when both the source and the target images have similar features. However, because of different exposure times, images have different amount of motion blur and noise, and hence, image features are not preserved. Although normalizing the images with their exposure times maintains brightness constancy between them, *it does not remove the differences in motion blur*. This results in incorrect motion estimation.

This is illustrated in Figure 2. Images are simulated assuming the scene to be translating images patches of size 64×64 pixels. For each patch, two images with different exposure times are generated, using an affine image noise model. Image intensities are normalized by their exposures, and then dense optical flow is computed between them. The difference in the estimated flow and the ground truth flow gives the registration error. As shown, images with the same exposures have the minimum registration errors.

Qualitative comparison of existing bracketing schemes:

The exponential scheme [2] achieves good DR as the exposures grow exponentially. The alternating scheme [6] uses long and short exposures (ratio between the exposures is 16). This scheme achieves a moderate DR. Both exponential and alternating schemes are prone to registration errors due to large differences between consecutive exposures. The burst-of-short-exposures scheme [21] results in the smallest increase in the DR as all the exposures are the same. Since the burst scheme uses images of the same exposure, it is robust to registration errors.

Thus, in the context of exposure bracketing for HDR imaging, there is a fundamental tradeoff between the dynamic range and registration accuracy. To capture a large dynamic range, it is important to use a large range of exposures. However, large differences in exposures of images can result in strong registration artifacts. How can we create a bracketing scheme that achieves high dynamic range while minimizing the likelihood of registration artifacts?

4. Generalized Registration for Exposure Bracketed Image Sequences

Consider a sequence of exposure bracketed images that are to be registered using optical flow. Our key idea is that instead of directly estimating the flow between individual frames, we estimate flow between sums of two contiguous sets of frames. We call the flow between sums of frames as *generalized flow* and the process of estimating generalized flow as *generalized registration*. Flow between individual frames is then computed by scaling the generalized flow. This is illustrated in Figure 3. In the next section, we will show that with the correct choice of exposure sequence, generalized registration allows computing flow between sums of frames having the same total exposure, while also achieving a high dynamic range.

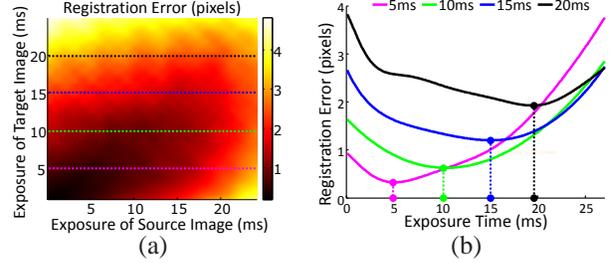


Figure 2. **Illustration of the iso-exposure advantage.** 400 image patches were extracted from HDR images. Images were simulated assuming the scene to be a translating patch. For each patch, two images with different exposure times were generated, normalized, and optical flow was computed between them. The difference in the computed flow and the ground truth flow gives the registration error. (c) 2D plot of the mean errors. The exposures of the source and the target images are plotted on the X and Y axes, respectively. (d) Error plots for four different target exposures. **Images with the same exposures have the minimum registration errors.**

Let the exposure bracketed frames be $\{f_1, \dots, f_i, f_{i+1}, \dots, f_K\}$, with exposure times $\{e_1, \dots, e_i, e_{i+1}, \dots, e_K\}$. In order to compute flow o_i^{i+1} between f_i and f_{i+1} , we first make two adjacent, contiguous sets of frames around f_i :

$$S_i^s = \{f_{i-n_s+1}, f_{i-n_s+2}, \dots, f_i\}, \quad (3)$$

$$S_i^t = \{f_{i+1}, f_{i+2}, \dots, f_{i+n_t}\}, \quad (4)$$

where the superscripts s and t stand for source and target. Number of frames in S_i^s and S_i^t is n_s and n_t , respectively. The *generalized frames* \mathcal{F}_i^s and \mathcal{F}_i^t for both sets are defined as the sum of the individual frames:

$$\mathcal{F}_i^s = f_{i-n_s+1} + f_{i-n_s+2} + \dots + f_i, \quad (5)$$

$$\mathcal{F}_i^t = f_{i+1} + f_{i+2} + \dots + f_{i+n_t}. \quad (6)$$

Suppose the inter-frame time gap δ between consecutive frames is negligible and the camera has a linear intensity response. Then, the frame \mathcal{F}_i^s is the same as the single frame $\widetilde{\mathcal{F}}_i^s$ that the camera would have captured had it exposed for the sum of exposure times $e_{i-n_s+1} + e_{i-n_s+2} + \dots + e_i$ ³. \mathcal{F}_i^t is related to $\widetilde{\mathcal{F}}_i^t$ in a similar manner. Let the flow between \mathcal{F}_i^s and \mathcal{F}_i^t be $\tilde{o}_{[i-n_s+1, i]}^{[i+1, i+n_t]}$ (generalized flow). The subscript and the superscript denote the first and the last frames in the sets S_i^s and S_i^t , respectively. Assuming that the flow vectors are linear within the duration of capture of the generalized frames \mathcal{F}_i^s and \mathcal{F}_i^t , the flow o_i^{i+1} is computed as⁴:

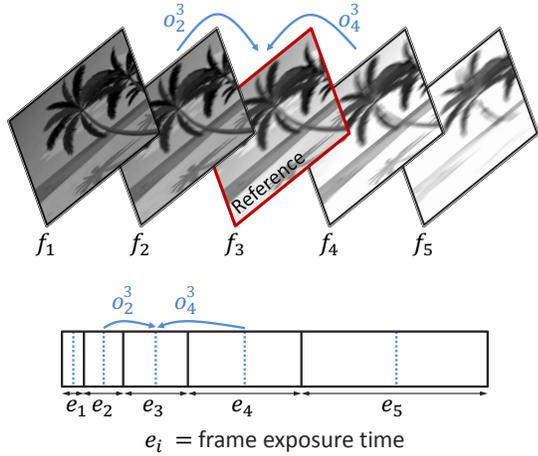
$$o_i^{i+1} = \frac{e_i + e_{i+1}}{i+n_t} \tilde{o}_{[i-n_s+1, i]}^{[i+1, i+n_t]}. \quad (7)$$

Two examples with $n_s = 2, n_t = 1$ ($\tilde{o}_{[1,2]}^{[3]}$) and $n_s = 1, n_t = 2$ ($\tilde{o}_{[5]}^{[3,4]}$) are illustrated in Figure 3. The

³The standard deviation of the effective read noise for \mathcal{F}_i^s is $\sqrt{n_s}$ times that of $\widetilde{\mathcal{F}}_i^s$; the photon noise is the same for \mathcal{F}_i^s and $\widetilde{\mathcal{F}}_i^s$.

⁴In general, if a large number of frames are added to create the generalized frames, the flows within a generalized frame may not be linear, especially for very long exposures. As shown in the next section, our approach requires adding only a few frames (1 – 3). This allows approximating the flows by linear vectors for a wide range of scenes and motions.

Conventional Registration



Generalized Registration

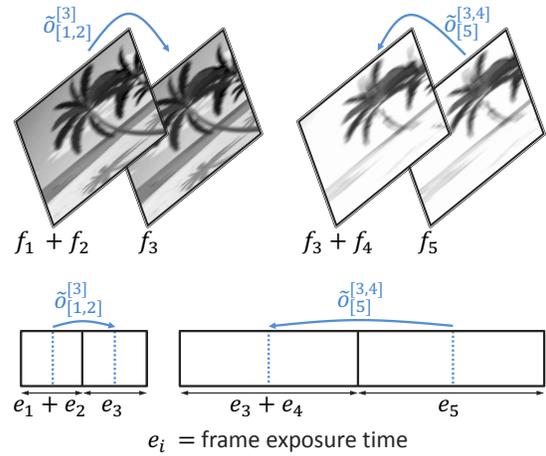


Figure 3. **Conventional versus generalized registration.** f_i are the bracketed LDR frames, e_i are the exposure times and o are the flows between frames. **(Left)** In conventional registration, flow is computed between individual frames. The flows are then used to register all the frames to a reference frame. Differences in exposures of the frames result in registration artifacts. **(Right)** We propose generalized registration where flow is computed between summations of frames. By choosing the exposure times appropriately, generalized registration ensures that the flow is always computed between sums of frames with the same total exposure times. Flows between individual frames are computed by scaling the generalized flows \tilde{o} .

individual flows are computed as $o_2^3 = \frac{e_2+e_3}{e_1+e_2+e_3} \tilde{o}_{[1,2]}^{[3]}$ and $o_4^3 = \frac{e_3+e_4}{e_3+e_4+e_5} \tilde{o}_{[5]}^{[3,4]}$.

5. Fibonacci Exposure Bracketing

In this section, we propose an exposure bracketing scheme that exploits generalized registration to ensure that optical flow is *always* computed between frames of the same exposure time. To formalize this, we define the *iso-exposure property* for an exposure sequence:

Definition 1 An exposure sequence $\{e_1, e_2, \dots, e_K\}$ has order (n_s, n_t) iso-exposure property if $\forall i \in [2 \dots K-1]$, there exist n_s and n_t such that
$$\sum_{j=i-n_s+1}^i e_j = \sum_{j=i+1}^{i+n_t} e_j.$$

If an exposure sequence has order (n_s, n_t) iso-exposure property, it is possible to make generalized frames F_i^s and F_i^t (using Eqs. 5 and 6) for every pair of adjacent frames f_i and f_{i+1} so that F_i^s and F_i^t have the same total exposure. Since the flow between f_i and f_{i+1} is computed by first estimating the flow between F_i^s and F_i^t (Eq. 7), the iso-exposure property ensures that the flow is always estimated between frames of the same exposure.

How should the parameters n_s and n_t be chosen? Both n_s and n_t should be as small as possible so that (a) flow vectors within the generalized frames can be assumed to be linear, and (b) the gaps created in F_i^s and F_i^t due to inter-frame gap δ are minimized. Moreover, $n_s - n_t$ should be minimized to ensure that the effective noise of both F_i^s and F_i^t is similar. Thus, the first natural choice is the (1, 1) iso-exposure sequence. In such a sequence, all the images have the same exposure. The burst-of-short-exposures [21] is an example. While such a sequence can minimize registration errors, it cannot capture a wide dynamic range.

The next higher order is (2, 1). Order (2, 1) iso-exposure property is achieved if every exposure e_i is equal to the sum of two previous exposures, i.e., $e_i = e_{i-1} + e_{i-2}$. This is the property of the *Fibonacci sequence* of numbers, a series well studied in number theory [8]. The sequence 1, 2, 3, 5, 8, 13, ... is the canonical Fibonacci sequence. We call the bracketing scheme with exposures forming a Fibonacci sequence as *Fibonacci exposure bracketing*.

What is the dynamic range achieved by Fibonacci bracketing? The *DR* is determined by the ratio of consecutive exposures [11]. The ratios of consecutive numbers in a Fibonacci sequence approaches (in the limit) $\phi = \frac{1+\sqrt{5}}{2}$, the *golden ratio*. This is a well-known result in number theory [8]. Thus, Fibonacci sequences behave like exponential sequences with a growth factor $G = \phi$. While it may appear that the *DR* of Fibonacci bracketing (DR_{fib}) is small due to a relatively small growth factor, it turns out that DR_{fib} is *always within a small additive constant of the maximum achievable dynamic range*.

Lemma 1 For any given time-budget T , the dynamic range achieved by Fibonacci bracketing DR_{fib} is within 1.39 stops of the maximum achievable dynamic range DR_{max} .

Proof 1 Consider an exposure sequence where the ratio of successive exposures is $\phi = \frac{1+\sqrt{5}}{2}$. This sequence follows the Fibonacci property. Suppose the time-budget T is equal to the sum of all exposures. Then, T is the sum of an exponential series with e_{min} as the first (minimum) exposure:

$$T = e_{min} \frac{(\phi^K - 1)}{\phi - 1} < e_{min} \frac{\phi^K}{\phi - 1}, \quad (8)$$

where K is the number of frames. The *DR* of Fibonacci sequence is given by substituting $e_{max} = e_{min} \phi^{K-1}$ in Eq. 2:

$$DR_{fib} = \log \left(\frac{I_{max}}{I_{min}} \phi^{K-1} \right). \quad (9)$$

The maximum dynamic range D_{max} for a time-budget T is achieved with two exposures, e_{min} and $T - e_{min}$:

$$DR_{max} = \log\left(\frac{I_{max} T - e_{min}}{I_{min} e_{min}}\right) < \log\left(\frac{I_{max} T}{I_{min} e_{min}}\right) \quad (10)$$

Substituting Eq. 8 into above, we get:

$$DR_{max} < \log\left(\frac{I_{max} \phi^K}{I_{min} \phi - 1}\right). \quad (11)$$

Subtracting Eq. 9 from Eq. 11, we get:

$$DR_{max} - DR_{fib} < \log\left(\frac{\phi}{\phi - 1}\right) < 1.39. \quad (12)$$

Thus, Fibonacci bracketing achieves both - a high dynamic range (close to maximum achievable) as well as robustness to registration errors. For example, a Fibonacci sequence constructed with a total time budget of $33ms$ and a minimum exposure of $0.1ms$ has 11 frames. This sequence results in a DR increase of 6.94 stops over a single LDR image. With the same time-budget and minimum exposure, the maximum achievable DR increase is 8.32 stops.

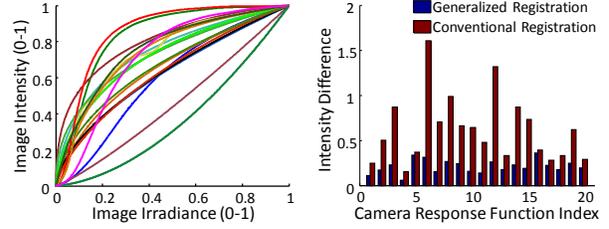
Robustness to non-linear camera response functions:

Several sensors have non-linear intensity response, especially most cell-phone cameras. If the two images have different exposures and the sensor has a non-linear response, the intensities of a scene point (after scaling by the exposures) are different. This can lead to strong registration errors. The maroon colored bars in Figure 4 (b) show the average intensity difference $D_{conv}(R)$ for conventional registration, corresponding to some typical camera responses R shown in Figure 4 (a). The expression for $D_{conv}(R)$ is derived in the technical report available at [1].

On the other hand, in the proposed approach, flow is always computed between two frames of the same total exposure, making it robust to non-linearities in the response functions. The blue-colored bars in Figure 4 (b) represent the intensity difference $D_{gen}(R)$ for Fibonacci bracketing and generalized registration. The differences are significantly smaller, thus making it possible to achieve good results without calibrating the camera's response curve. Moreover, the proposed approach can be especially useful in exposure-fusion based techniques, where the captured images are directly merged without radiometric calibration and computing an intermediate HDR image [10].

Higher order iso-exposure sequences: So far, we have considered order (2, 1) iso-exposure sequences. Next, we discuss higher order sequences. It is not always possible to make (N, M) order sequences for $M > 1$ while ensuring non-negative and non-decreasing exposures. In this paper, we consider sequences with only $(N, 1)$ order iso-exposure property for different values of N .

A sequence of exposures has the $(N, 1)$ order iso-exposure property if the exposure times are from an *order-N Fibonacci* sequence (or *N-bonacci* sequence). In an *N-bonacci* sequence, each number is the sum of previous N numbers. Examples are the *tribonacci* ($N = 3$) and *tetranacci* ($N = 4$) sequences. It turns out that all



(a) Camera responses (b) Normalized intensity errors
Figure 4. **Robustness of Fibonacci bracketing to non-linear camera response functions.** (a) 20 response functions (from [4]). (b) Average intensity difference between pixels in source and target images for each response function. Differences for Fibonacci bracketing based generalized registration is always less than that of conventional registration.

N -bonacci sequences behave similar to the *Fibonacci* sequence; the ratio of consecutive numbers approaches a constant in the limit [16]. The ratios, called the N -bonacci constants ϕ_N , lie between 1 and 2, i.e., $1 \leq \phi_N < 2$ for all N . For example, $\phi_3 = 1.84$ (tribonacci) and $\phi_4 = 1.93$ (tetranacci). Thus, all N -bonacci sequences behave like exponential sequences with a growth factor $G = \phi_N$, $1 \leq G < 2$. By following the same steps as in Proof 1, it can be shown that the DR achieved by tribonacci and tetranacci sequences is within 1.14 and 1.06 stops, respectively, of the maximum achievable DR. In Section 6, we show results of tribonacci exposure bracketing.

6. Hardware Prototype and Results

Several consumer cameras support exposure bracketing. It is possible to capture a sequence of images while varying the exposure time. However, there is a time-lag of approximately $50 - 200ms$ between successive frames. On the other hand, while there is negligible time gap between successive images of a video stream captured by a video camera, it is not possible to change exposure time during capture. For generalized registration, ideally, a sensor that allows varying exposures with a negligible inter-frame gap is required. One way of achieving this is to develop a video camera with a programmable timing control unit, so that successive images can have different exposures.

While we have not developed such an image sensor, we have emulated it by using a machine vision Miro M310 camera. By triggering the camera externally with pulses generated from an Arduino controller based circuit, it is possible to achieve a negligible inter-frame time gap ($0.1ms$) while varying exposure from one frame to the next. Our setup is shown in Figure 5.

Results: Figure 6 shows the result of Fibonacci bracketing and generalized registration for an outdoor night scene. The time-budget for capturing a single HDR image was set to $120ms$. The minimum exposure time was $0.3ms$. The Fibonacci exposure sequence obeying these constraints is $\{0.3, 0.49, 0.8, 1.3, 2.12, 3.45, 5.63, 9.17, 14.95, 24.37, 39.72\}ms$. Each exposure is the sum of the previous two; the ratio between successive exposures is 1.63. The sum of all the exposures is $102.3ms$. For comparisons, we use an exponential sequence with growth factor $G = 4$. The sequence

is $\{0.3, 1.2, 4.8, 19.2, 76.8\}ms$. The total exposure time is again $102.3ms$. In order to mitigate *intra-frame motion blur*, the LDR images were deblurred using a recent method [18] that can handle spatially-varying motion blur. The images for exponential bracketing were normalized by their exposure times before registration. Optical flow was computed using the technique proposed in [17].

The best-exposed LDR image chosen from the Fibonacci sequence (image with the maximum number of pixels in the intensity range $[0.07 - 0.93]$) contains saturated regions. The HDR image computed without registering the LDR images is blurred due to camera and scene motion. The HDR image computed using exponential bracketing and conventional registration has artifacts due to large differences in the exposures. The HDR images were tone-mapped using the photographic tone-reproduction operator [12].

Comparisons with existing bracketing schemes: Figure 7 shows comparisons of Fibonacci bracketing with the burst (of short exposures) scheme [21] and the alternating (long and short exposure) scheme [6]. The same time-budget of $120ms$ was used for all three schemes. To ensure that the scene was approximately the same for every scheme, all the input images were captured within a short duration ($380ms$). For the burst scheme, 36 frames were captured, each with an exposure of $0.3ms$ for bright scenes (top row), and $3ms$ for dark scenes (bottom row). Since the frame rate of the camera is 300 fps, 36 LDR frames could be captured within $120ms$. For the alternating scheme, 22 frames with alternating exposures of $0.3ms$ and $10ms$ were captured. The frames were normalized before registration.

The alternating scheme suffers from strong registration artifacts because of the large exposure differences and cannot reconstruct mid-tones of the scenes (table and flowers). Images captured using the burst scheme have a low dynamic range. Although the bright regions are faithfully captured (sky, candle flames), the images have low signal-to-noise-ratio in the dark regions. With the same capture time, HDR images created using Fibonacci bracketing have a significantly better quality. **For more results and comparisons, please see the project web-page [1].**

Results of tribonacci bracketing: Figure 8 shows HDR results computed using tribonacci exposure bracketing for the same scenes as in Figure 7. In this case, flow is computed between a frame and the sum of three previous frames. The ratio of consecutive exposures in a tribonacci sequence is $\phi_3 = 1.84$, with the minimum exposure of $0.3ms$. A total of 9 images were used, giving a dynamic range increase of $1.84^8 \approx 131$ times over a single LDR image.

Comparison between conventional and generalized registration: Figure 9 shows a comparison between conventional and generalized registration. For both, the same LDR frames were used (11 frames of a Fibonacci bracketing sequence). For conventional registration, image intensities were normalized. Conventional registration does not exploit the iso-exposure property of Fibonacci bracketing. This results in incorrect registration and ghosting artifacts. In contrast, generalized registration produces a ghost-free image. **Evaluating the effect of non-linear camera response**

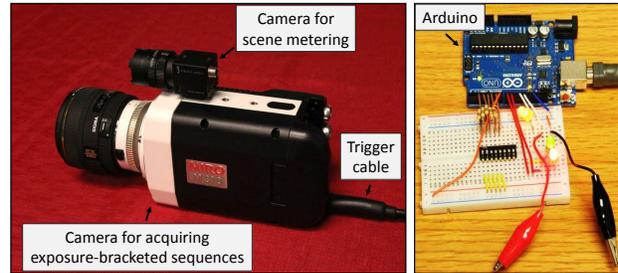


Figure 5. **Image acquisition setup.** (Left) The Miro M310 camera used to capture exposure bracketed images. Point-Grey Flea3 camera was used for scene-metering to capture HDR video (Section 7). (Right) Bracketing was performed using an external trigger generated by an Arduino microprocessor.

function: In order to emulate a non-linear response, we applied a γ -curve ($\gamma = 2.2$) on the input LDR images. Figure 10 shows the comparison for an outdoor scene. Because of differences between exposures of consecutive images, there are strong distortions for the exponential scheme. In contrast, as discussed in Section 5, the proposed approach is robust to non-linear camera response.

Comparisons with different registration and image merging techniques: Several techniques have been proposed to register and merge differently exposed images [20, 15]. We compared with the techniques of Ward [20] and Sen *et al.* [15]. We compared four cases: (a) Exponential-bracketing + conventional-registration, (b) Exponential-bracketing + Ward-registration, (c) Exponential-bracketing + Sen-method, and (d) Our method.

The average SNR (on 40 simulated image-sequences) for the four methods are (a) 37dB, (b) 35dB, (c) 35dB and (d) 40dB. We implemented the Ward method ourselves. For Sen *et al.* method, we used the authors' code available on their website.

7. Capturing HDR Video

In this section, we extend the proposed techniques to capture HDR videos. The bracketing sequence is changed according to scene characteristics (intensity and motion) as they vary during video capture. As discussed in Section 5, all N -bonacci sequences lend themselves to generalized registration. These sequences are defined by their growth factor G , which varies between 1 and 2, i.e. $1 \leq G < 2$. At one extreme is the sequence with $G = 1$, where all the exposures are the same. Since it has only short exposures, this sequence should be used only to capture bright scenes with large motion and relatively small dynamic range. On the other hand, sequences with larger G values (e.g., Fibonacci and tribonacci) have a wide range of exposures, and are more suitable for capturing scenes with a wide dynamic range and small/moderate motion.

Thus, we capture HDR video by changing the growth factor G as a function of the scenes intensities and the amount of motion. As discussed in the previous paragraph, G should be inversely proportional both to scene's brightness values and amount of motion. Moreover, since we aim to capture HDR video, G should vary smoothly as the scene changes. We use the following simple function that can be



Figure 6. **Comparison between exponential bracketing and Fibonacci bracketing.** (a) The best exposed LDR image contains saturated regions. (b) HDR image computed without registering LDR images is blurred due to camera and scene motion. (c) HDR image computed using exponential bracketing and conventional registration has strong registration artifacts. (d) HDR image obtained using the proposed Fibonacci bracketing and generalized registration techniques. **See the project web-page [1] for more results and comparisons.**

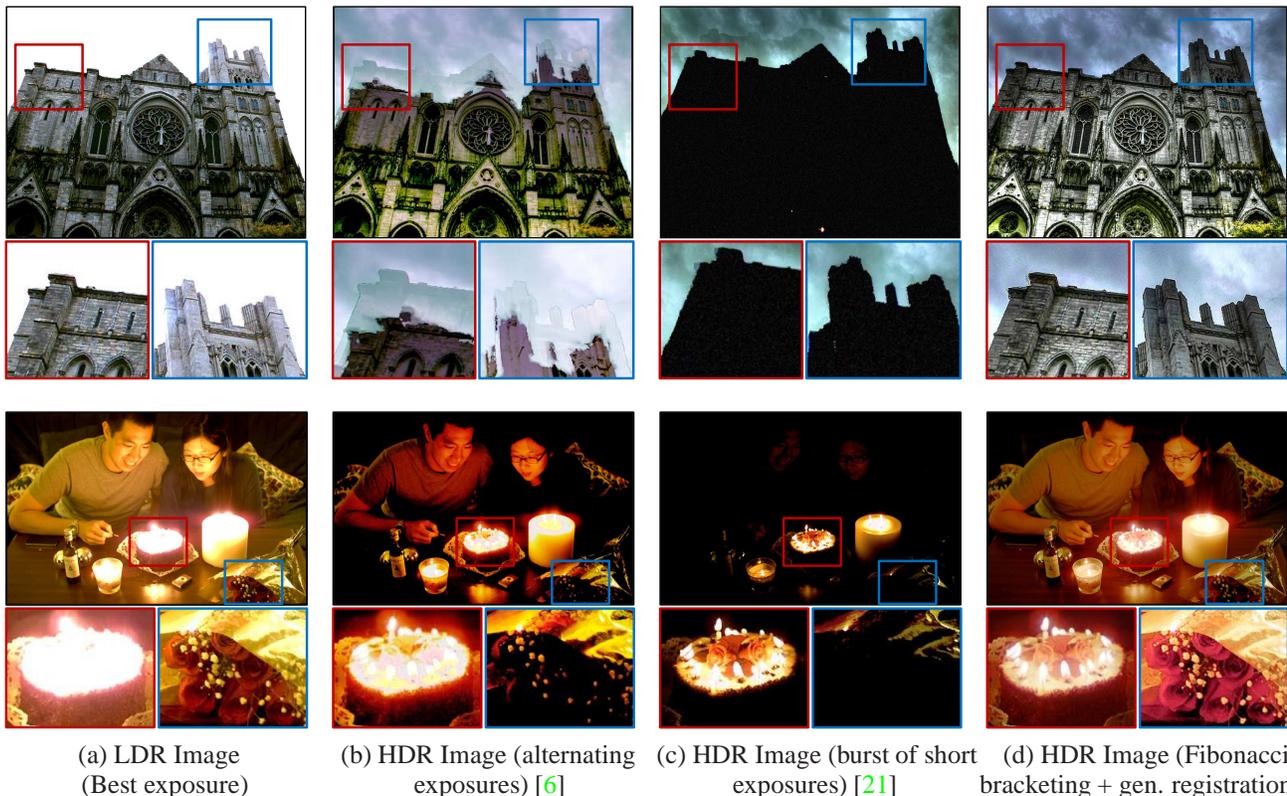


Figure 7. **Comparison between Fibonacci and two existing bracketing schemes.** (Top) Church on a cloudy day. (Bottom) Indoor birthday party. Both scenes have large dynamic range ($10^5 - 10^6$). (a) The best exposed LDR image has saturated regions (sky, candles). (b) The alternating (long and short exposure) scheme suffers from strong registration artifacts (church) and can not reconstruct mid-tones of the scenes faithfully (flowers, table). (c) Images captured using the burst (of short exposures) scheme have a low dynamic range, resulting in low quality in the dark regions. (d) HDR images created using Fibonacci bracketing and generalized registration.

computed sufficiently fast on commodity hardware:

$$G_k = 1 + \left((1 - \hat{I}_{k-1}) (1 - \hat{M}_{k-1}) \right), \quad (13)$$

where \hat{I}_{k-1} is the median intensity of the previous frame and \hat{M}_{k-1} is the mean motion between two previous frames. Motion is computed by computing correlation between 1-D projections of the two frames along rows and columns. Both \hat{I}_{k-1} and \hat{M}_{k-1} are normalized to lie in the range $[0, 1]$. We used a Miro M310 camera (see Figure 5) for capturing exposure bracketed images. A Point-Grey Flea3 camera was used for ‘scene-metering’ - intensity and motion information was computed on images captured by

the Flea3 camera for determining the bracketing sequence. Since image analysis and capture steps are parallelized and the total capture time for each bracketing sequence is about $60ms$, our system captures HDR video at 15 fps. **See the project web-page [1] for the videos and more results.**

8. Discussion and Limitations

In order to be used widely in consumer cameras, HDR techniques should handle motion. Ours is the *first* exposure bracketing scheme that is designed to deal with dynamic camera and objects. The proposed approach is robust to non-linearities in the camera response functions and low

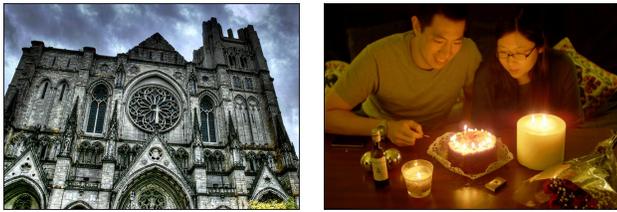


Figure 8. Results of tribonacci bracketing and generalized registration for the scenes shown in Figure 7.



Fibonacci bracketing + Conventional registration Fibonacci bracketing + Generalized registration
 Figure 9. Comparison between conventional and generalized registration. The same LDR images (11 frames of a Fibonacci bracketing sequence) were used for both cases. Conventional registration does not exploit the iso-exposure property of Fibonacci bracketing. This results in incorrect registration and strong ghosting artifacts. Generalized registration produces a ghost-free image.

sensor bit-depth, and require minimal modifications to use with existing sensors. Thus, our techniques are particularly suitable for implementation on inexpensive image sensors such as ones used in cell-phone cameras.

The technique of Fibonacci bracketing+generalized registration should be seen as a general ‘pre-conditioning’ step in HDR imaging, that enhances the accuracy of existing image-alignment/optical flow methods. Note that it is not a new alignment technique in itself.

Limitations: While the proposed approach significantly mitigates registration artifacts, it may not completely remove them. Our method shares the limitations of dense optical flow techniques (e.g., aperture problem), and hence may not perform reliably for textureless regions, occlusions and in the presence of highly non-rigid motion (such as fluid motion). For extremely fast motions, or large inter-frame time gaps, our technique may not produce a good result. In order to remove the residual artifacts, one of the post-processing techniques discussed in Section 2 may be used.

References

- [1] Project webpage. <http://www.cs.columbia.edu/CAVE/projects/FibonacciHDR/>. 5, 6, 7
- [2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH*, 1997. 1, 2, 3
- [3] O. Gallo, N. Gelfand, W. Chen, M. Tico, and K. Pulli. Artifact-free HDR imaging. In *IEEE ICCP*, 2009. 2
- [4] M. D. Grossberg and S. K. Nayar. What is the space of camera response functions? In *IEEE CVPR*, 2003. 5



HDR Image (exponential) HDR Image (Fibonacci)
 Figure 10. Evaluating the effect of non-linear camera response. A non-linear response was emulated by applying a γ -curve ($\gamma = 2.2$) on the input LDR frames. Because of differences between exposures of consecutive frames, there are strong distortions for exponential (bushes, walking path and people at the bottom) scheme. In contrast, the proposed approach is robust to non-linear response.

- [5] S. W. Hasinoff, F. Durand, and W. T. Freeman. Noise-optimal capture for high dynamic range photography. In *IEEE CVPR*, 2010. 1, 2
- [6] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High dynamic range video. *ACM Trans. Graph.*, 22(3), 2003. 2, 3, 6, 7
- [7] E. A. Khan, A. O. Akyuz, and E. Reinhard. Ghost removal in high dynamic range images. In *IEEE ICIP*, 2006. 2
- [8] M. Livio. *The Golden Ratio: The Story of Phi, the World's Most Astonishing Number*. Broadway Books, 2002. 4
- [9] S. Mann and R. W. Picard. On being ‘undigital’ with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proc. of IST*, 1995. 1, 2
- [10] T. Mertens, J. Kautz, and F. V. Reeth. Exposure fusion. In *Proc. Pacific Graphics*, 2007. 5
- [11] S. K. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. In *CVPR*, 2000. 2, 4
- [12] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. In *ACM SIGGRAPH*, 2002. 6
- [13] M. Rouf, R. Mantiuk, W. Heidrich, M. Trentacoste, and C. Lau. Glare encoding of high dynamic range images. In *IEEE CVPR*, 2011. 2
- [14] D. Schleicher and B. G. Zagar. High dynamic range imaging by varying exposure time, gain and aperture of a video camera. In *Proc. of IEEE Instrumentation and Measurement Technology Conference*, 2010. 1
- [15] P. Sen, N. K. Kalantari, M. Yaesoubi, S. Darabi, D. B. Goldman, and E. Shechtman. Robust patch-based HDR reconstruction of dynamic scenes. *ACM Trans. Graph.*, 31(6), 2012. 2, 6
- [16] W. R. Spickerman and R. N. Joyner. Binet’s formula for the recursive seq. of order K. *Fibonacci Quart.*, 22, 1984. 5
- [17] D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *IEEE CVPR*, 2010. 6
- [18] M. Tallon, J. Mateos, S. Babacan, R. Molina, and A. K. Saggelos. Space-variant blur deconvolution and denoising in the dual exposure problem. *Information Fusion*, 2012. 6
- [19] M. Tocci, C. Kiser, N. Tocci, and P. Sen. A versatile HDR video production system. *ACM Trans. Graph.*, 30(4), 2011. 2
- [20] G. Ward. Fast, robust image registration for compositing high dynamic range photographs from handheld exposures. *Journal Of Graphics Tools*, 8, 2003. 6
- [21] L. Zhang, A. Deshpande, and X. Chen. Denoising versus deblurring: HDR techniques using moving cameras. In *IEEE CVPR*, 2010. 2, 3, 4, 6, 7