GPU-based Video Feature Tracking And Matching

Sudipta N. Sinha[1]**, Jan-Michael Frahm**[1]**, Marc Pollefeys**[1]**, Yakup Genc**[2]

[1] Department of Computer Science, CB# 3175 Sitterson Hall,
   University of North Carolina at Chapel Hill, NC 27599
[2] Real-time Vision and Modeling Department, Siemens Corporate Research,
   755 College Road East, Princeton, NJ 08540

Abstract   This paper describes novel implementations of the KLT feature tracking and SIFT feature extraction algorithms that run on the graphics processing unit (GPU) and is suitable for video analysis in real-time vision systems. While significant acceleration over standard CPU implementations is obtained by exploiting parallelism provided by modern programmable graphics hardware, the CPU is freed up to run other computations in parallel. Our GPU-based KLT implementation tracks about a thousand features in real-time at 30 Hz on 1024×768 resolution video which is a 20 times improvement over the CPU. It works on both ATI and NVIDIA graphics cards. The GPU-based SIFT implementation works on NVIDIA cards and extracts about 800 features from 640×480 video at 10Hz which is approximately 10 times faster than an optimized CPU implementation.

1 Introduction

Extraction and matching of salient 2D feature points in video is important in many computer vision tasks like object detection, recognition, structure from motion and marker-less augmented reality. While certain sequential tasks like structure from motion for video [18] require online feature point tracking, others need features to be extracted and matched across frames separated in time (eg. wide-baseline stereo). The increasing programmability and computational power of the graphics processing unit (GPU) present in modern graphics hardware provides great scope for acceleration of computer vision algorithms which can be parallelized [3, 11, 12, 14, 15, 16, 17]. GPUs have been evolving faster than CPUs (transistor count doubling every few months, a rate much higher than predicted by Moore's Law), a trend that is expected to continue in the near future. While dedicated special-purpose hardware or reconfigurable hardware can be used for speeding up vision algorithms [1, 2], GPUs provide a much more attractive alternative since they are

affordable and easily available within most modern computers. Moreover with every new generation of graphics cards, a GPU implementation just gets faster.

In this paper we present GPU-KLT, a GPU-based implementation for the popular KLT feature tracker [6, 7] and GPU-SIFT, a GPU-based implementation for the SIFT feature extraction algorithm [10]. Our implementations are 10 to 20 times faster than the corresponding optimized CPU counterparts and enable real-time processing of high resolution video. Both GPU-KLT and GPU-SIFT have been implemented using the OpenGL graphics library and the Cg shading language. While GPU-KLT works on both ATI and NVIDIA graphics cards, GPU-SIFT currently works only on NVIDIA but will be modified to also work with ATI cards in future. As an application, the GPU-KLT tracker has been used to track 2D feature points in high-resolution video streams within a vision based large-scale urban 3D modeling system described in [19].

Our work is of broad interest to the computer vision, image processing and medical imaging community since many of the key steps in KLT and SIFT are shared by other algorithms, which can also be accelerated on the GPU. Some of these are (a) image filtering and separable convolution, (b) Gaussian scale-space construction, (c) non-maximal suppression, (d) structure tensor computation, (e) thresholding a scalar field and (f) re-sampling discrete 2D and 3D scalar volumes. This paper is organized as follows. Section 2 describes the basic computational model for general purpose computations on GPUs (GPGPU). Sections 3 presents the basic KLT algorithm followed by its GPU-based implementation and experiments on real video and an analysis of the results obtained. Section 4 describes similar aspects of GPU-SIFT. Finally we present our conclusions in Section 5.

2 GPGPU Concepts

Modern programmable graphics hardware contains powerful coprocessors (GPUs) with a peak performance of hundreds of GFLOPS which is an order of magnitude higher than that of CPUs [21]. They are designed to independently process streams of vertices and fragments (pixels) in parallel. However their data parallel SIMD (single instruction multiple data) architecture also provides an abstraction for performing general purpose computations on GPUs (GPGPU) and for treating the GPU as a stream processor.

In the GPGPU framework, the fully programmable vertex and fragment processors perform the role of the computational kernels while video memory (framebuffers, textures etc.) provides it with a memory model (see Figure 1 for an overview of the graphics pipeline implemented in hardware). Texture mapping on the GPU is analogous to the CPU's random read-only memory interface while the ability to render directly into texture (off-screen rendering) provides a memory-write mechanism. However by virtue of its specialized design, the GPU has a more restricted memory model when compared to a CPU (scatter operations i.e. random memory writes are not allowed). Texture memory caches are designed for speed and prevent concurrent read and write into the same memory address. Thus distinct read and write textures must be used. They can be swapped after each render pass making the write texture available as input and vice versa (ping-pong rendering).
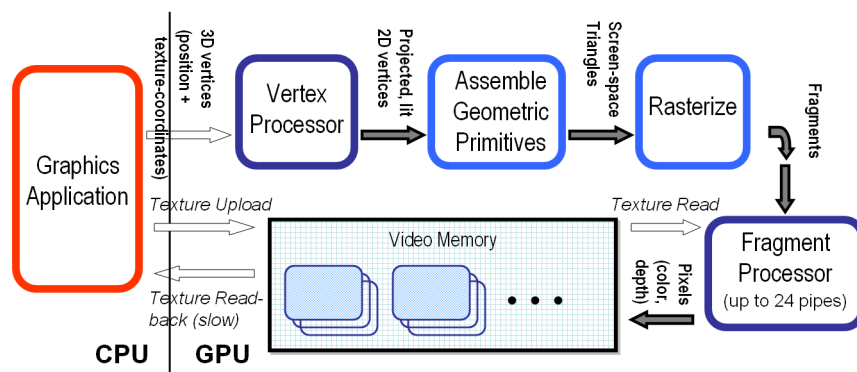
Fig. 1 Overview of the 3D Graphics Pipeline. The fragment processor and direct off-screen rendering capability is frequently used in GPGPU applications.

In order to implement an algorithm on the GPU, different computational steps are often mapped to different fragment programs. For each computational step, the appropriate fragment program is bound to the fragment processor and a render operation is invoked. The rasterization engine generates a stream of fragments and also provides a fast way of interpolating numbers in graphics hardware. Most GPGPU applications execute multiple fragment programs within multiple off-screen rendering passes. While pixel-buffers (pBuffers) exist on older graphics cards, recently frame-buffer objects (FBOs) were introduced, providing a simple and efficient off-screen rendering mechanism in OpenGL. Details about GPGPU programming are available in [20, 22].

Many computer vision algorithms map well into this parallel stream processing model. Image processing tasks which can process multiple pixels independently (eg. convolution) can be performed very fast by fragment programs (computation kernels) exploiting the high parallelism provided by multiple fragment pipes (upto 24 in modern cards). A large fraction of the GFLOPS dedicated to texture mapping in GPUs is non-programmable. While image processing applications can sometimes leverage this by using the bilinear interpolation of texture mapping, they also benefit from the 2D texture cache layouts designed for fast texture mapping.

Recently there has been a growing interest in the computer vision community to solve important computationally expensive problems like image registration [16], stereo and segmentation using graphics hardware. A correlation-based real-time stereo algorithm for the GPU was first proposed by [11] while more complex formulation of stereo [12, 13, 14] were implemented more recently. GPUs have been successfully used by [15, 17] to accelerate background segmentation in video, often used as a first step in many vision applications. A versatile framework for programming GPU-based computer vision tasks (radial undistortion, image stitching, corner detection etc.) was recently introduced by [3, 4].

3 KLT Tracking on GPU

3.1 The Algorithm

The KLT tracking algorithm [6,7] computes displacement of features or interest points between consecutive video frames when the image brightness constancy constraint is satisfied and image motion is fairly small. Assuming a local translational model between subsequent video frames, the displacement of a feature is computed using Newton's method to minimize the sum of squared distances (SSD) within a tracking window around the feature position in the two images.

Let $I(*, *, t)$ represent the video frame at time t. If the displacement of an image point (x, y) between time t and $t + \Delta t$, denoted by $(\Delta x, \Delta y)$ is small, then according to the brightness constancy constraint,

$$I(x, y, t + \Delta t) = I(x + \Delta x, y + \Delta y, t)$$

Let $\mathbf{x} = (x, y)^T$ and $\mathbf{v} = (\Delta x, \Delta y)^T$. In the presence of image noise r,

$$I(\mathbf{x}, t + \Delta t) = I(\mathbf{x} + \mathbf{d}, t) + r$$

KLT will compute the displacement vector $\mathbf{d}$ that minimizes the following error

$$r = \sum_W (I(\mathbf{x} + \mathbf{d}, t) - I(\mathbf{x}, t + \Delta t))^2$$

over a small image patch W. Approximating $I(\mathbf{x} + \mathbf{d}, t)$ by its Taylor expansion, one obtains the following linear system to estimate the unknown $\mathbf{d}$ where $\mathbf{G} = [\frac{\partial I}{\partial x} \ \frac{\partial I}{\partial y}]$ is the image gradient vector at position $\mathbf{x}$.

$$\underbrace{(\sum_W \mathbf{G}^T \mathbf{G})}_{\mathbf{A}} (\mathbf{d}) = \underbrace{\sum_W \mathbf{G}^T \Delta \mathbf{I}(\mathbf{x}, \Delta t)}_{\mathbf{b}} \tag{1}$$

Tomasi later proposed a variation of the KLT equation which uses both images symmetrically. This equation, derived in [8] is identical to Equation 1 except that here

$$\mathbf{G} = [\frac{\partial (I(*, t) + I(*, t + \Delta t))}{\partial x} \quad \frac{\partial (I(*, t) + I(*, t + \Delta t))}{\partial y}]$$

This symmetric version is used in our GPU implementation.

Feature to track are selected by finding image points where a saliency or cornerness measure

$$\mathbf{c} = min \, (eig \, (\sum_W [\frac{\partial I}{\partial x} \ \frac{\partial I}{\partial y}]^T [\frac{\partial I}{\partial x} \ \frac{\partial I}{\partial y}] \,) \,)$$

(the minimum eigen-value of the 2x2 structure tensor matrix obtained from gradient vectors) is a local maximum. It is evaluated over the complete image [6,7] and a subsequent non-maximal suppression is performed. The KLT algorithm is described in Figure 2.

```
KLT Tracking (ft_list, F₀ , F₁)  {

   (1) Build-Pyramid: builds multi-resolution
intensity and gradient pyramid from images F₀ , F₁

   (2) Track:

   For all pyramid levels from coarse to fine

        For multiple iterations

            For each feature f in ft_list

                compute coefficients of A and b as
   shown                    in Equation 1.

                solve A d = b

                evaluate d and update track of feature

   }
```

```
Re-select-Features (ft_list) {

   mask   = mask_out_region ( ft_list)

   c_map = evaluate corner-ness
            measure c over whole image

   // Perform non-maximal suppression

   pts     = find_features (#max_feats,
            mask, sort (c_map) )

   add_new_features (ft_list, pts)

   }
```

Fig. 2 Pseudo-code for the two fundamental routines in the KLT Tracking algorithm.

Since the linearity assumption is only valid for a small displacement d, a multi-resolution KLT tracker is often used in practice for handling larger image motion. It first tracks at coarse resolutions and then refines the result in finer resolutions. Multiple iterations are performed at each resolution for better accuracy. Due to camera motion and occlusion, features tracks are eventually lost; hence new features must be re-selected from time to time to maintain a roughly fixed number of features in the tracker.

3.2 GPU Implementation Details

GPU-KLT maps various steps of the tracking algorithm to different fragment programs. Every video frame is uploaded to video memory where it is smoothed and its multi-resolution pyramid of image intensity and gradients is constructed. The tracking is done on every frame using the image pyramids corresponding to the current and previous frames. Feature re-selection is performed once in every k frames to keep a roughly constant feature count in the tracker. The value of k was set to 5 for all our experiments but this generally depends on camera motion and the number of lost features.

Implementation Strategies: RGBA floating point textures were used for storage on the GPU. This is supported on most modern GPUs. Section 3.3 discusses the precision that was required by our implementation on different hardwares. The multi-resolution image pyramid and the associated gradient vectors are represented by a set of RGBA textures where different channels are used for the intensity and gradient magnitudes. A second set of identical image pyramid textures is needed during the construction of the image pyramid on the GPU (as explained below). The corner-ness map is represented by a pair of textures; one for partial sums and the second for the final values. The feature list table is represented by a $m \times n$ texture where m stands for the maximum feature count while n stands for (#tracking
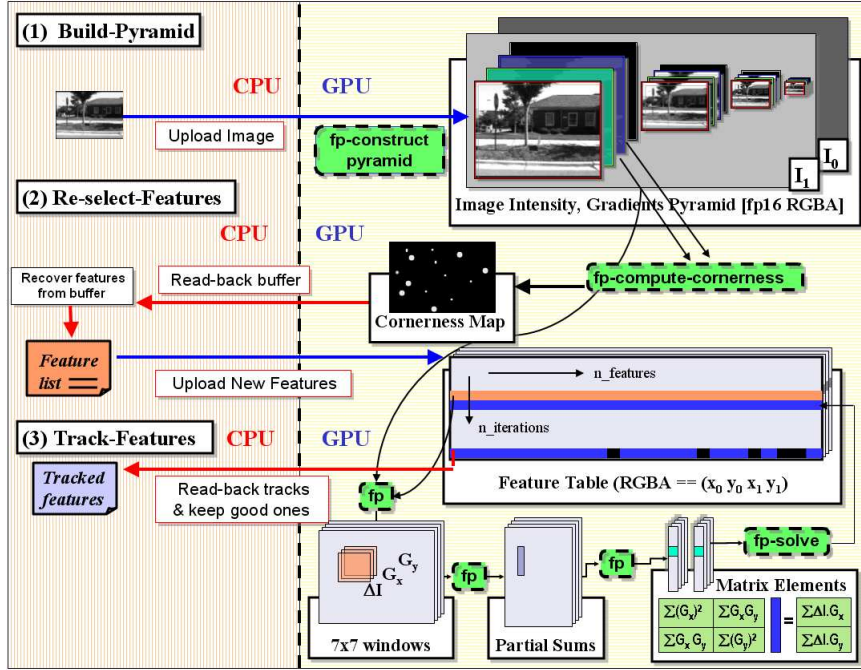
Fig. 3 Overview of the steps in the GPU-KLT implementation.

iterations)×(# pyramid levels). Three other texture units are used for computing and storing intermediate values computed during tracking and computing the elements of matrix **A** and vector **b** (refer Equation 1).

Build-Pyramid: The multi-resolution pyramid of the image intensity and its gradients are computed by a series of two-pass separable Gaussian convolutions performed in fragment programs. The fragment program uses OpenGL's multiple texture coordinates (TEXCOORD0 . . . TEXCOORD7) to read a row or column of pixels. The 1D convolution filter kernel size limited to 7, accounts for most practical values of σ. Since fragment programs support vector operations, the blurred pixel and the gradient magnitudes are computed simultaneously. The second set of textures are used to store the results of the row convolution pass and are subsequently read by the column convolution pass. Since the tracker requires information for only two video frames, textures for two image pyramids are allocated and a pointer indicating the current frame alternates between the two.

Track: KLT tracking performs a fixed number of tracking iterations at each image resolution starting with the coarsest pyramid level. Each tracking iteration constructs a linear system of equations in two unknowns for each interest point (see Equation 1), **A d = b** and directly solves them to update the estimated displacement. This is done in four steps by four fragment programs on the GPU. First a fragment program bilinearly interpolates intensity and gradient magnitudes in 7×7 patches around each KLT feature in the two images and stores them in a temporary texture. While NVIDIA provides hardware support for bilinear inter-

polation of floating point textures, a fragment program is required to do this on ATI. Various quantities evaluated at 7×7 image blocks are added in two passes; first computing partial row sums followed by a single column sum. The second and third fragment program evaluates all the six elements of the matrix **A** and the vector **b** and writes them into a different texture for the next fragment program to use. Finally Equation 1 is solved in closed form by the fourth fragment program which writes the currently tracked position into the next row in the feature table texture. The invocation of these four fragment programs corresponds to a single tracking iteration in the original algorithm (see Figure 2).

At the end of (*#max-iterations*)×(*#pyramid-levels*) tracking iterations, the final feature positions (the last row in the feature table texture) are read back to the CPU along with two other values per feature - Δ **d**, the final tracking update of the iterative tracker and **res**, the SSD residual between each initial and tracked image patch. An inaccurate feature track is rejected when its Δ **d** and **res** exceeds the respective thresholds. While KLT originally performs these tests after every tracking iteration, GPU-KLT skips them to avoid conditional statements in fragment programs for speed. This however forces it to track all **N** (=*#max-features*) features for all the iterations. Hence GPU-KLT's running time depends on **N** and not the number of valid features being tracked.

GPU-KLT performs tracking completely on the GPU contrary to [3] who builds the matrices (Equation 1) on the GPU using fragment programs, performs a readback and then solves a stacked linear system on the CPU. Multi-resolution, iterative tracking is ruled out in their case due to the CPU-GPU tranfer bottleneck. [3] also does not compare CPU and GPU implementations for accuracy and timings. Our multi-resolution, iterative tracker handles larger image motions than [3] and performs accurate tracking in real-time on high-resolution video.

Re-select-Features: The KLT corner-ness map is computed in two passes. The first pass computes the 2×2 structure tensor matrix at each pixel. The values in a 7×7 window centered at every pixel are added using partial row sums followed by a column sum. The minimum eigen value of the resulting matrix is stored in the corner-ness map. During feature re-selection, the neighborhood of existing features is invalidated and early Z-culling is used to avoid computations in these image regions. Early Z-culling works as follows. In the first pass, a binary mask is created by rendering $t \times t$ quads (where t is the desired minimum distance between features) for every existing valid feature. The depth test in the graphics pipeline is disabled while depth writing is enabled and this binary mask is loaded into the graphics hardware's depth buffer. In the next pass, depth writing is disabled while the depth test is enabled. With early Z-culling hardware support, fragments corresponding to invalidated pixels are not even generated when the corner-ness map is being computed. Finally a corner-ness map with sparse entries is read back to the CPU. Non-maximal suppression is done on it to find new additional features to track. Using the GPU for invalidating image regions before computing the corner-ness map makes this final step on the CPU much faster.
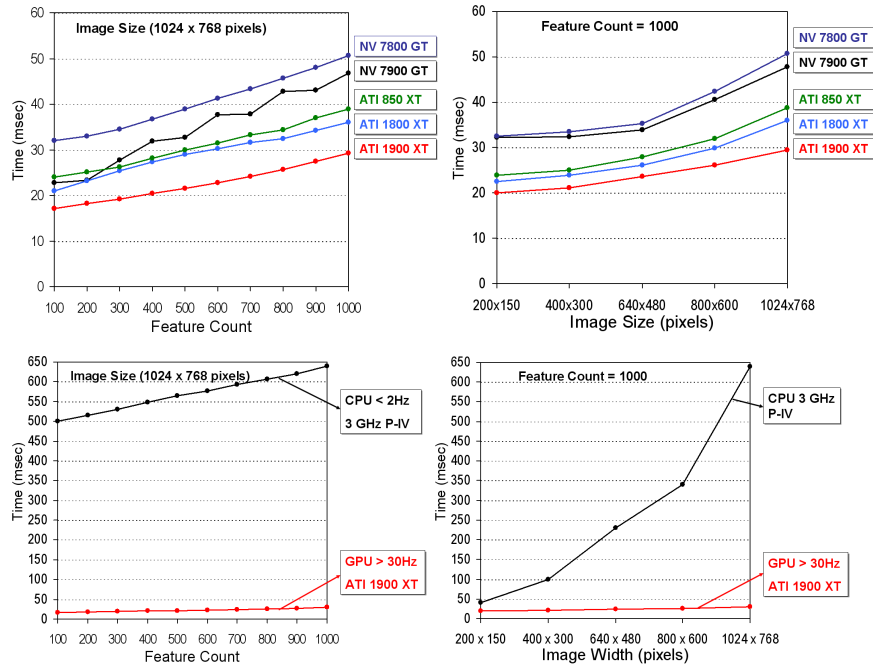
Fig. 4 A comparison of GPU-KLT timings on various graphics cards (Center) and between GPU and CPU (Below).

3.3 Results

To evaluate the performance of GPU-KLT, tests were performed on various ATI (850XT, 1800XT, 1900XT) and NVIDIA (7800GTX, 7900GTX) graphics cards. These tests showed an improvement of one order of magnitude in speed over a standard KLT implementation [9]. A $20\times$ speedup over the CPU version was observed on a ATI 1900XT, where GPU-KLT tracks 1000 features in 1024×768 resolution video at 30 Hz. The performance measurements are shown in Figure 4. The evaluation shows that currently all ATI graphic cards outperform the tested NVIDIA graphics cards. This is due to the precision required for solving Equation 1 within a fragment program. The required 32 bit floating point precision is always provided by ATI cards even when the storage textures have only 16 bit floating point precision. In contrast to ATI, NVIDIA cards could only provide 32 bit precision computations in the fragment programs when the allocated textures too had 32 bit precision. This increased the memory bandwidth during processing on NVIDIA cards and explains their lower speeds. Furthermore, the measurements in Figure 4 show that GPU-KLT is bandwidth limited on all tested graphics cards and its computational complexity linearly depend on the number of features as well as on the number of pixels in the images.

GPU-KLT was also tested qualitatively for tracking accuracy. This evaluation is in general difficult to perform as it would require ground truth tracks. To our knowledge there is no standard data set for such an evaluation. Hence we compared
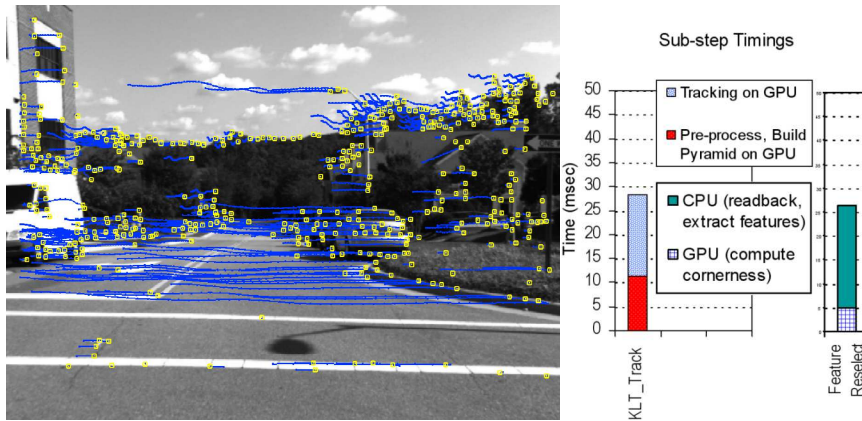
Fig. 5 (Top Left) A frame from a 1024 × 768 resolution video shows features being tracked in real-time using GPU-KLT. (Top Right) GPU-KLT sub-step timings.

GPU-KLT and the standard KLT to each other. Due to the different orders of operations and the different ALU's in the GPU and the CPU the results are in general not equal. We tested the tracking inside an application for camera pose estimation [19] using the quality of the estimated camera poses as the criteria for tracking accuracy. It showed that both trackers provide in general the same quality of tracks. Thus we conclude that GPU-KLT provides an order of magnitude of speedup over CPU implementations while maintaining the same tracking quality. Our open source implementation is available at `http://cs.unc.edu/~ssinha/GPU_KLT`

4 SIFT Feature Extraction on GPU

4.1 The Algorithm

The Scale Invariant Feature Transform (SIFT) [10] algorithm is a popular candidate for extraction of interest points invariant to translation, rotation, scaling and illumination changes in images. It first constructs a Gaussian scale-space pyramid from the input image while also calculating the gradients and difference-of-gaussian (DOG) images at these scales. Interest points are detected at the local extremas within the DOG scale space. Once multiple keypoints have been detected at different scales, the image gradients in the local region around each feature point are encoded using orientation histograms and represented in the form of a rotationally invariant feature descriptor. The details are described in [10].

4.2 GPU Implementation Details

The construction of the Gaussian scale space pyramid is accelerated on the GPU using fragment programs for separable Gaussian convolution. The intensity image, gradients and the DOG values are stored in a RGBA texture and computed in the
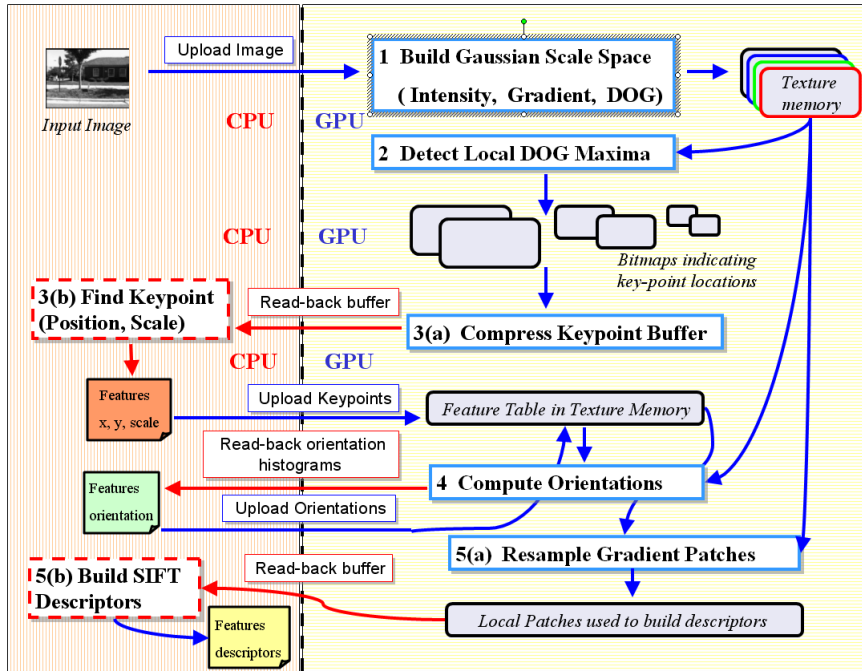
Fig. 6 Overview of the steps in the GPU-SIFT implementation.

same pass using vector operations in fragment programs. Blending operations in
graphics hardware are used to find local extremas in the DOG pyramid in parallel
at all pixel locations. The Depth test and the Alpha test is used to threshold these
keypoints; The local principal curvatures of the image intensity around the key-
point is inspected; this involves computing the ratio of eigenvalues of the 2×2
structure tensor matrix of the image intensity at that point. The keypoint locations
are implicitly computed in image-sized, binary buffers, one for each scale in the
pyramid. A fragment program compresses (a factor of 32) the binary bitmap into
RGBA data, which is readback to the CPU and decoded there.

At this stage, a list of keypoints and their scales have been retrieved. Since
reading back the gradient pyramid (stored in texture memory) to the CPU is ex-
pensive, the subsequent steps in SIFT are also performed on the GPU. Gradient
vectors near the keypoint location are Gaussian weighted and accumulated inside
an orientation histogram by another fragment program. The orientation histogram
is read back to the CPU, where its peaks are detected. Computing histograms [3]
on the GPU is expensive and doing it on the CPU along with a small readback
is a little faster. The final step involves computing 128 element SIFT descriptors.
These consist of a set of orientation histograms built from 16×16 image patches in
invariant local coordinates determined by the associated keypoint scale, location
and orientation. SIFT descriptors cannot be efficiently computed completely on
the GPU, as histogram bins must be blended to remove quantization noise. Hence
we partition this step between the CPU and the GPU. We resample each feature's

gradient vector patch, weight them using a Gaussian mask using blending support on the GPU. The resampled and weighted gradient vectors are collected into a tiled texture block which is subsequently transferred back to the CPU and then used to compute the descriptors. This CPU-GPU partition was done to minimize data read-back from the GPU since transferring the whole gradient pyramid back to the CPU is impractical. Moreover texture re-sampling and blending are efficient operations on the GPU; hence we perform those steps there. This also produces a compact tiled texture block which can be transferred to the CPU in a single readback.

GPU-SIFT gains a large speed-up in the Gaussian scale-space pyramid construction and keypoint localization steps. The compressed readback of binary images containing feature positions reduces the readback data-size by a factor of 32. The feature orientation and descriptors computation is partitioned between the CPU and GPU in a way that minimizes data transfer from GPU to CPU. Overall a 8-10X speedup is observed compared to CPU versions.


Fig. 7 Features extracted using GPU-SIFT from an image pair.

4.3 Results

GPU-SIFT was implemented in OpenGL/Cg, used Pbuffers for off-screen rendering. A texture manager allocates and manages all the data in GPU memory within a single double-buffered PBuffer. In future we will replace PBuffers with FBOs. The current implementation runs only with NVIDIA hardware and was tested on the Geforce 7800GTX and 7900GTX cards. Figure 7 shows GPU-SIFT features extracted from video frames separated in time (camera undergoing rotation). Figure 8 compares timings between the CPU and GPU implementations for a range of image resolution and feature-count. The NVIDIA 7900GTX gave a 10X speedup over an optimized CPU implementation. GPU-SIFT running on the NVIDIA 7900GTX could extract about 1000 SIFT features from streaming 640 × 480 resolution video at an average frame-rate of 10 Hz. Figure 9 shows how

Image Size 640 x 480

Time (msec) vs Feature Count

Legend:
- NV 7800GTX PCI-e
- NV 7900GTX PCI-e
- CPU 3GHz P-IV

Timing includes 2-way CPU-GPU transfer time

Feature Count = 1000

Time (msec) vs Image Width (in pixels)

Legend:
- NV 7800GTX PCI-e
- NV 7900GTX PCI-e
- CPU 3GHz P-IV
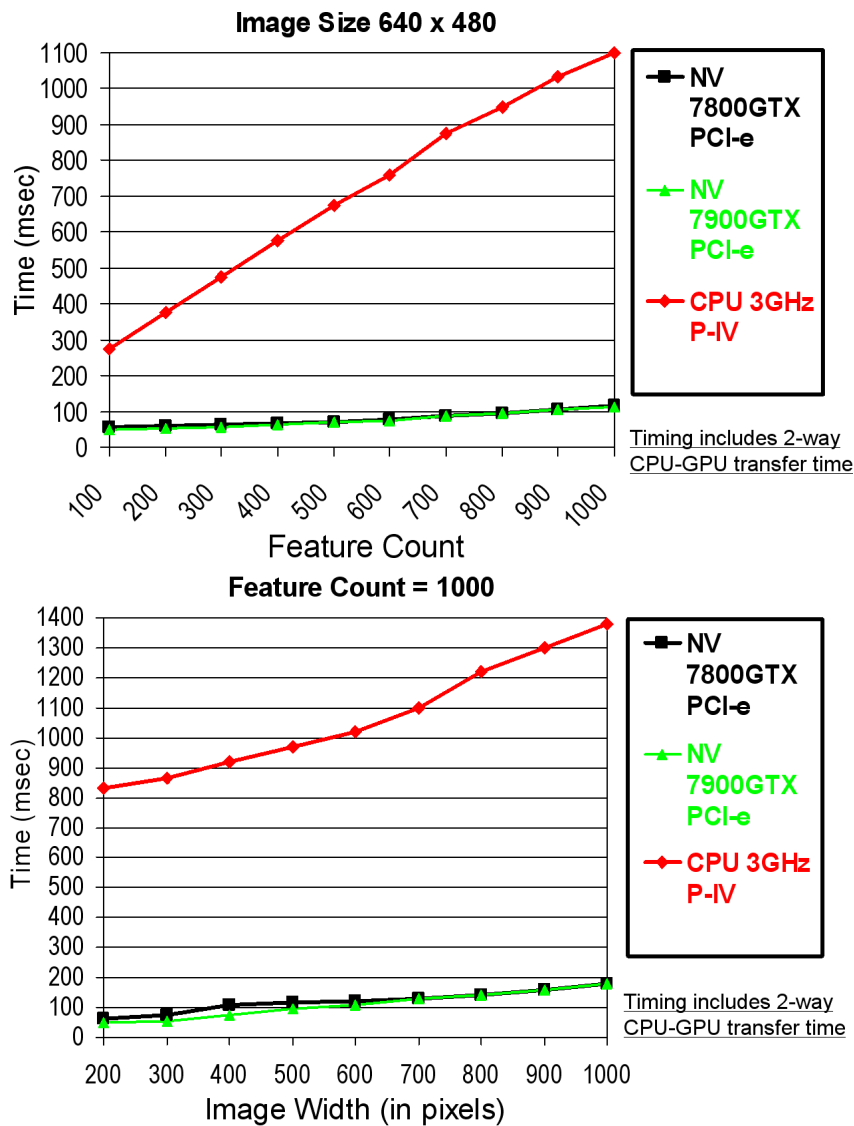
Timing includes 2-way CPU-GPU transfer time

Fig. 8 GPU-SIFT timings compared with an optimized CPU implementation for a range of image-sizes and feature-counts. GPU-SIFT has a 10-12X speed-up.

different steps in the algorithm scale with input size. As the image resolution increases, scale-space construction and keypoint localization performed on the GPU dominates running time while as the feature count increases, more time is spent in computing feature descriptors on the CPU.
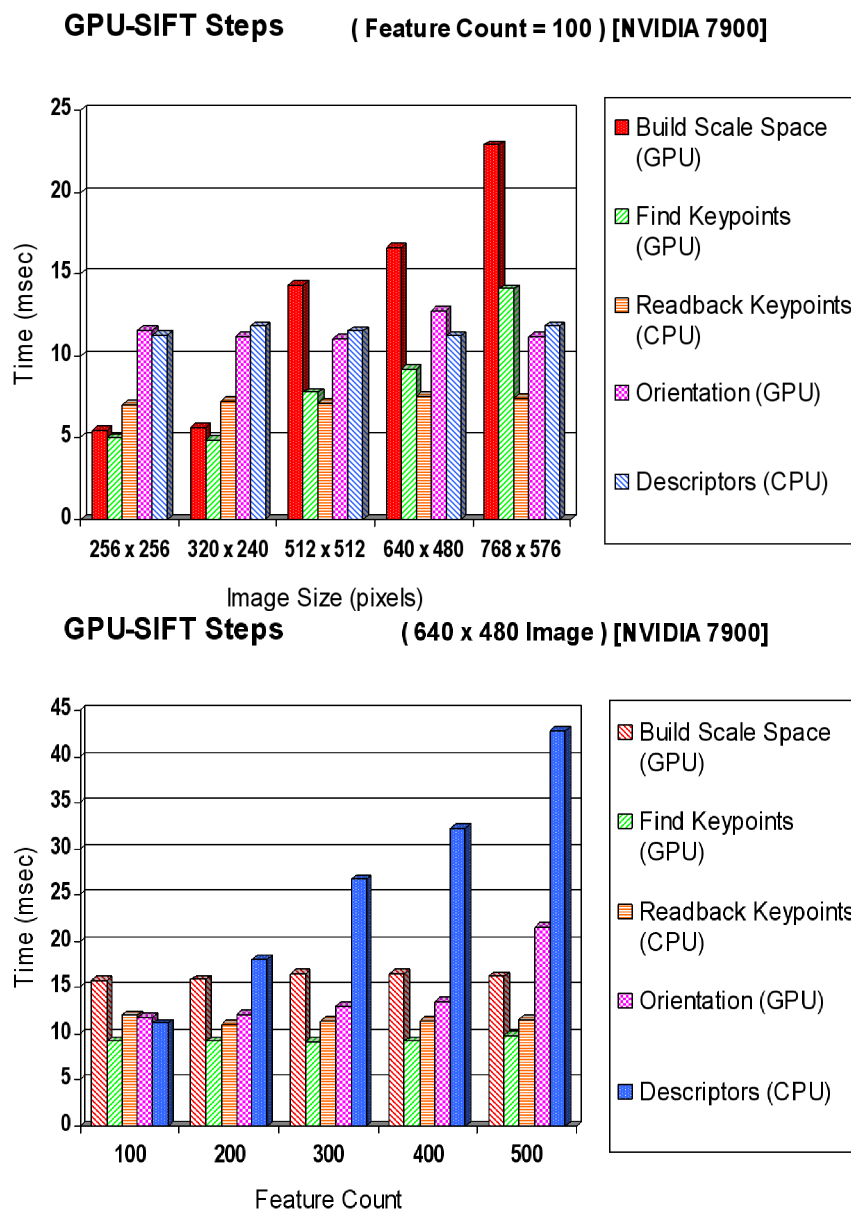
GPU-SIFT Steps (Feature Count = 100) [NVIDIA 7900]

GPU-SIFT Steps (640 x 480 Image) [NVIDIA 7900]

Fig. 9 Timings of steps in GPU-SIFT. For large feature-counts, computing descriptors on CPU dominates the running time while for large images, the GPU dominates.

5 Conclusions

Both SIFT and KLT have been used for a wide range of computer vision tasks ranging from structure from motion, robot navigation, augmented reality to face recognition, object detection and video data-mining with quite promising results.

We have successfully ported these popular algorithms to the GPU. In both cases, strategies were developed for dividing computation between the CPU and GPU in the best possible way under the restrictions of the GPU's computational model. Our GPU implementations which exploited the parallelism and incredible raw processing power provided by today's commodity graphics hardware are considerably faster than optimized CPU versions. As new generation graphics cards evolve (faster than predicted by Moore's law), our implementations would run even faster. This now makes it possible to perform high quality feature tracking, interest point detection and matching on high resolution video in real-time on most modern computers without resorting to the need for special-purpose hardware solutions.

References

1. M. Bramberger, B. Rinner, and H. Schwabach, "An Embedded Smart Camera on a Scalable Heterogeneous Multi-DSP System". In Proceedings of the European DSP Education and Research Symposium (EDERS 2004), Nov 2004.
2. S. Klupsch, M. Ernst, S.A. Huss, M. Rumpf, and R. Strzodka. "Real time image processing based on reconfigurable hardware acceleration". In Proceedings of IEEE Workshop Heterogeneous Reconfigurable Systems on Chip, 2002.
3. J. Fung and S. Mann, "OpenVIDIA: parallel GPU computer vision", ACM MULTIMEDIA 2005, pp. 849–852.
4. J. Fung, S. Mann, "Computer Vision Signal Processing on Graphics Processing Units", Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2004), Montreal, Quebec, Canada, May 17-21, 2004, pp. V-93 - V-96.
5. M. Gong, A. Langille and M. Gong. "Real-time image processing using graphics hardware: A performance study", International Conference on Image Analysis and Recognition, pp.1217–1225, 2005.
6. C. Tomasi and T. Kanade, "Detection and Tracking of Point Features", Tech. Rept. CMU-CS-91132, Carnegie Mellon University, April 1991.
7. B.D. Lukas and T. Kanade, "An iterative image registration technique with an application to stereo vision". In Proceedings of the International Joint Conference on Artificial Intelligence, pp. 674-679, 1981.
8. S. Birchfield, "Derivation of Kanade-Lucas-Tomasi tracking equation", unpublished notes, 1997.
9. S. Birchfield, "KLT: An Implementation of the Kanade-Lucas-Tomasi Feature Tracker", http://www.ces.clemson.edu/ stb/klt, Nov 2005.
10. D.G. Lowe, "Distinctive image features from scale-invariant keypoints", IJCV, Vol. 60(2), pp. 91–110, November, 2004.
11. R. Yang and M. Pollefeys, "Multi-Resolution Real-Time Stereo on Commodity Graphics Hardware", CVPR, pp. 211-217, 2003.
12. C. Zach, H. Bischof and K. Karner, "Hierarchical Disparity Estimation With Programmable 3D Hardware", WSCG (International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision), Short Communications, pp. 275–282, Plzen, Slowakei, 2004
13. J. Woetzel and R. Koch, "Real-time multi-stereo depth estimation on GPU with approximative discontinuity handling", 1st European Conference on Visual Media Production, March, 2004.
14. P. Labatut, R. Keriven, J.-P. Pons, "A GPU Implementation of Level Set Multiview Stereo". International Conference on Computational Science (4) pp. 212-219, 2006.

15. R. Yang and G. Welch, "Fast image segmentation and smoothing using commodity graphics hardware", Journal of Graphics Tools, 7(4):91–100, 2002.
16. R. Strzodka, M. Droske, and M. Rumpf. "Image registration by a regularized gradient flow - a streaming implementation in DX9 graphics hardware". Computing, 73(4), pp. 373-389, 2004.
17. A. Griesser, S.D. Roeck, A. Neubeck and L.J.V. Gool. "GPU-Based Foreground-Background Segmentation using an Extended Colinearity Criterion". Vision, Modeling, and Visualization (VMV), 2005
18. M. Pollefeys, L.J.V. Gool, M. Vergauwen, F. Verbiest, K. Cornelis, J. Tops and R. Koch, "Visual Modeling with a Hand-Held Camera", IJCV, Vol. 59(3), pp. 207–232, 2004.
19. A. Akbarzadeh, J.-M. Frahm, P. Mordohai, B. Clipp, C. Engels, D. Gallup, P. Merrell, M. Phelps, S. Sinha, B. Talton, L. Wang, Q. Yang, H. Stewenius, R. Yang, G. Welch, H. Towles, D. Nistr and M. Pollefeys, Towards Urban 3D Reconstruction From Video, Invited paper, Third International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT 2006).
20. GPGPU: General-Purpose Computation on GPUs. http://www.gpgpu.org, 2004.
21. K. A. Bjorke, NVIDIA Corporation, "Image processing using parallel GPU units", Proceedings of SPIE Vol. 6065, January 2006.
22. Pharr, M. and Fernando, R. "GPU Gems 2: Programming Techniques for High-Performance Graphics and General-Purpose Computation (Gpu Gems)". Addison-Wesley Professional, 2005.