

Hardware Acceleration of Background Modeling in the Compressed Domain

Stefan Popa, Danny Crookes, and Paul Miller

Abstract—In intelligent video surveillance systems, scalability (of the number of simultaneous video streams) is important. Two key factors which hinder scalability are the time spent in decompressing the input video streams, and the limited computational power of the processor. This paper demonstrates how a combination of algorithmic and hardware techniques can overcome these limitations, and significantly increase the number of simultaneous streams. The techniques used are processing in the compressed domain, and exploitation of the multicore and vector processing capability of modern processors. The paper presents a system which performs background modeling, using a Mixture of Gaussians approach. This is an important first step in the segmentation of moving targets. The paper explores the effects of reducing the number of coefficients in the compressed domain, in terms of throughput speed and quality of the background modeling. The speedups achieved by exploiting compressed domain processing, multicore and vector processing are explored individually. Experiments show that a combination of all these techniques can give a speedup of 170 times on a single CPU compared to a purely serial, spatial domain implementation, with a slight gain in quality.

Index Terms—Background subtraction, compressed domain, hardware acceleration, multicore, SSE, video surveillance.

I. INTRODUCTION

A KEY issue for intelligent video surveillance systems is that of scalability. Real-world deployment of video surveillance systems can involve hundreds, or even thousands, of cameras. Thus it is vital that video analytics are scalable with respect to both robust performance and also computability. Scalability is achieved for these systems by decomposing them into subsystems consisting of 8, 16, 32 or 64 IP cameras connected to a single server. Thus, we argue that scalability can be achieved with respect to video analytics, if we can process the data in real-time produced by one of these subsystems. With regard to computational efficiency, an important consideration is the hardware/software architecture of the video analytics processing. The main commercial approach to achieving computational efficiency is the use of the Texas Da Vinci Multimedia DSP chipset. In the video surveillance research community, the use of GPUs for implementation of video analytics has been investigated. However, a drawback

of this approach is the PCI bus bottleneck between the GPU and CPU. More recently, the availability of multicore CPUs and vectorization has provided another option with respect to computational efficiency.

Another consideration with respect to scalability and computational efficiency is that the video data are transmitted and stored in the compressed domain. However, most video analytic algorithms operate in the spatial domain; therefore, the video data has to be decompressed. From a computational viewpoint, decompression of, say, 64 video streams is nontrivial. Furthermore, the presence of compression artifacts can also reduce performance. One approach to reducing this computational overhead, therefore, is to implement the video analytics in the compressed domain.

The first step in video analytics is usually some form of foreground detection based on background modeling. The gold standard in this regard, both commercially and in research, is the mixture of Gaussians (MoG) algorithm [1]. Therefore, in this work we investigate the implementation of this algorithm in the compressed domain using multicore and vectorization. Specifically, we set ourselves the challenge: Can we perform foreground detection on 64 streams of video in real-time on a single chip without compromising performance? The work reported in this paper tries to answer this.

II. RELATED WORK

A. Background Modeling in the Compressed Domain

One of the most popular methods for extracting moving objects from a scene is background subtraction (BS). BS is one of the first low-level processing operations in virtually any intelligent video surveillance system, and it is the operation of identifying and segmenting moving objects in video frames by separating the still areas, called the background (BG), from the moving objects, called foreground (FG). Any BS algorithm first constructs a representation of the BG called the background model. Then, each subsequent frame is subtracted from this background model to give the resulting FG. Adaptive BS algorithms will also update the model along the sequence to compensate for eventual changes in illumination or other changes in the background [2].

The main advantage of including a BS stage in a video surveillance system is the increase in performance of subsequent higher-level video analysis [3].

As shown in [4], background modeling (BM) algorithms in the compressed domain can broadly be classified into three categories, depending on the compression features used: algorithms based only on Discrete Cosine Transform (DCT) coefficients

Manuscript received October 01, 2012; revised February 12, 2013 and May 25, 2013; accepted July 23, 2013. Date of current version September 06, 2013. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Frederic Dufaux.

The authors are with the School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Belfast, BT9 5BN, U.K. (e-mail: spopa01@qub.ac.uk; d.crookes@qub.ac.uk; p.miller@ecit.qub.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIFS.2013.2276753

[5]–[8]; algorithms based only on the motion vectors (MV) [9], [10]; and algorithms based on both [11].

Several studies have focused on BM algorithms using only DCT coefficients. In [5] the authors propose a framework that uses competing Hidden Markov Models over small neighborhoods, which are capable of maintaining a valid background model. In this case, the small neighborhoods are the JPEG 8×8 blocks of DCT coefficients. For this framework, the main challenge is background initialization, which requires a large amount of time. Another drawback is that this framework is only well suited for indoor surveillance. In [6] the author proposes a fast and efficient adaptive method for modeling the background, which uses two features extracted from each 8×8 block of DCT coefficients in a JPEG frame. The first feature is the first DCT coefficient (also called the DC coefficient), and the second feature is a weighted sum of a few DCT low frequency coefficients (also called the AC coefficients). A Gaussian distribution is used to model each of the two features. In addition, further processing is used to handle variations in the environment and improve the segmentation. By combining the model built around the two features with different characteristics (the DC coefficient is more sensitive to changes in illumination and the low frequency AC coefficients are more sensitive to changes in texture), the author claims that the model is robust to many of the problems linked with BS in video sequences representing outdoor scenes, such as gradual and sudden illumination changes, and small repetitive background movements. The authors of [7] use the same principle of modeling the background using two features extracted directly from each 8×8 block of DCT coefficients in a frame as part of a more complex tracking system. The main difference between the two is the choice of the second feature, which is represented by the first three AC coefficients in [7]. Also in [12]–[14], the Running Average algorithm with DCT coefficients is compared with algorithms in the spatial domain.

One of the most referenced papers addressing the problem of background modeling in the compressed domain, using only DCT coefficients, is [8]. Here, the authors present a framework composed of three algorithms for BM (Running Average, Median and Mixture of Gaussians) and a two-stage segmentation approach with pixel-level resolution. In the first stage, the background model is generated by one of the algorithms, which is further used to identify the block regions that are fully or partially occupied by a foreground object. In the second stage, the pixels from the partially occupied blocks are then processed (classified) in the spatial domain. Even though the processing is performed in the spatial domain for the partially occupied blocks, the authors prove through theoretical analysis and practical measurements that their algorithms in the compressed domain are several times faster than their spatial domain counterparts [15] and with better segmentation results.

B. Hardware Acceleration for Background Modeling

Several hardware-accelerated implementations for foreground detection and segmentation systems based on BM in the spatial domain have been developed. In [16], the authors implemented a BM algorithm based on an adaptive mixture of Gaussians (AGMM). Using an NVIDIA GPU with the

CUDA [17] they achieved $18\times$ acceleration compared with an implementation on an Intel multicore CPU using multi-threading. In the same paper, an implementation on IBMs Cell Broadband Engine Architecture (CBEA) achieved $3\times$ the acceleration compared with the same Intel multicore CPU benchmark. In [18], the authors implement a system that uses selective HOG-based features and BS based on a GMM for pedestrian detection. They report that, using an Intel i7 and a low budget NVIDIA GeForce 480GTX card, they are able to process 48 frames/second with high accuracy. Although the system does not achieve a significant speedup, the implementation is complex, and the BS represents only a part of the total computation. The authors of [19] implement only the AGMM algorithm. They report that the system can process close to 250 frames/second for high-resolution images, and close to 1600 frames/second for low-resolution images even when using a low-end NVIDIA GeForce 9600GT. Using the same GPU, they also report an acceleration of over $11\times$ compared with a reference CPU implementation.

It is possible to conclude that, although there has been a significant amount of research in BM and higher level video analysis in the compressed domain, it is clear that the compressed domain has not been explored as extensively as the spatial domain with respect to algorithmic performance and hardware acceleration. To the best of our knowledge, there are no significant hardware accelerated implementations of BM algorithms in the compressed domain. This is the key novel contribution of the work reported in this paper.

Due to current industry practices, video surveillance footage is stored in MJPEG compressed format in order to be accepted as evidence in law courts. In addition, all surveillance video cameras currently available on the market can stream images in the MJPEG compressed format. For these reasons, the research will be based on it.

III. BACKGROUND

A. MJPEG/JPEG Compression Format

MJPEG is a video compression format that takes advantage only of the pixels redundancy within a frame (spatial pixels redundancy). Each frame in a MJPEG video stream is a JPEG encoded image, which is independent of any previous or future frames; therefore, the compressed domain data available for processing in MJPEG compressed video streams is the DCT coefficients of the current frame [20].

The JPEG image compression standard [21] defines different modes of operation. Among the different modes of operation the most representative mode is the baseline JPEG mode, which is described next for a gray-scale image.

First, the baseline JPEG encoder splits the image into 8×8 nonoverlapping blocks of pixels and the pixels intensity is shifted from a range of 0:255 to a range of $-127:127$ in order to evenly distribute the values around zero. Further, each block is almost independently encoded, as in Fig. 1, using the following steps.

The forward type II two-dimensional discrete cosine transform (2D-DCT type II) is first applied to the 8×8 block of pixels, resulting in an 8×8 block of DCT coefficients. The




Fig. 1. JPEG compression/decompression.

top-left frequency component in the transformed block is called the DC coefficient and it is equivalent to the average intensity value of the pixels in the block. Next, the 64 DCT coefficients are quantized using different quantization thresholds for different frequency components. This is the step that gives rise to the lossy nature of the JPEG compression. During quantization the high frequency coefficients are more often than not set to zero. By applying this transformation, the entropy encoding step will become much more effective, because it reduces the dynamic range of the DCT coefficients and, as a consequence, fewer bits are needed for representation. After quantization, the 64 DCT coefficients are rearranged in a zigzag order. In this way the low frequency coefficients, which, in contrast to high-frequency coefficients, are mainly nonzero, are grouped at the start of the array. The high frequency coefficients, which are usually zero, are grouped at the end of the array. This transform is called zigzag reordering. The DC coefficients of the neighboring blocks tend to be highly correlated. The JPEG compression algorithm exploits this property and predicts the value of the current DC coefficient based on the value of DC coefficient of the previous block. This transform is called DPCM DC (differential pulse-code modulation). After zigzag reordering and DPCM DC, the array containing the DCT coefficients usually turns out to have some nonzero values scattered between zeros. Using this property, the JPEG format encodes the nonzero values and the number of previous zeros in just a pair of numbers. This transform is called run-length encoding (RLE) of DCT coefficients. The entropy encoding transform assigns variable length codes to pairs of values created during RLE, so that less probable pairs are represented by longer encoded binary codes and more probable pairs are represented by shorter encoded binary codes. JPEG compression uses Huffman encoding for entropy encoding. At the end, the entropy encoded data, the quantization table and the Huffman table are combined with markers to create the final byte stream (data stream). The markers are nothing more than delimiters for the different component parts of the byte stream. During the decompression, the JPEG decoder performs the inverse of the operations previously described in reverse order: starting with entropy decoding and finishing with the inverse DCT transform and inverse level shifting. Due to quantization transform, the decompressed image is not identical to the original one.

Because of the computation required to decompress an MJPEG video stream, processing the DCT coefficients directly is attractive. However, extracting the DCT coefficients from the video stream still requires a significant amount of processing, though inverse DCT is not required. The process is called partial decompression. The compressed domain representation of each frame is a 2D matrix of blocks, where each block comprises the DCT coefficients for the 8×8 block of pixels.

B. Background Modeling Using Online Mixture of Gaussians

This section will introduce the Mixture of Gaussians (MoG) algorithm for which a hardware acceleration strategy will be presented later.

MoG is one of the most common BM algorithms in the literature. The algorithm received considerable interest from the moment it was introduced in [15]. The popularity of this algorithm is due to two important aspects. First, it can model backgrounds containing nonstatic objects, such as tree branches or bushes moving in the wind (multimodel background). The algorithm can handle multimodel backgrounds because it follows the evolution of a set of K Gaussian distributions simultaneously (where K is a small number). Second, a recursive (or online) formulation of the algorithm can significantly reduce its complexity. However, the algorithm remains complex and challenging. There are two notable drawbacks of MoG [3]. First, algorithm parameters require careful tuning. An extensive discussion about this and possible solutions can be found in [1]; due to the speed of our accelerated implementations (see the results section) and to the limited space in a paper we simply used the brute force method to find reasonable values for the parameters. Second, for the case where the scene remains stationary for too long, the variance of the Gaussian distributions will become too small and any sudden change in the global illumination will force most of the scene to be classified as foreground. This problem can be easily solved during implementation by using a minimum variance to cap the lower bound value of the variance; this measure prevents the model from over-fitting the data.

Further, we describe an unified view of the algorithms in [8], [15]. Considering the case of a sequence of gray-scale images, we have the following two situations. In the spatial domain, the input to a BM algorithm is, for each pixel position, a scalar value (one-dimensional feature vector) which is the pixel intensity at that location in the image. In the compressed domain, the input is, for each 8×8 block, a 64-dimensional feature vector which is obtained by concatenating the rows (or the columns) of the 8×8 block of DCT coefficients at that block location in the JPEG image. The underlying learning method used by the MoG algorithm to update the model parameters is not influenced by the type of feature vector used as input, and it can be applied to data coming from the spatial or compressed domain [8], [15]. The MoG algorithm uses a mixture of K Gaussian distributions to maintain a probability density function (PDF) for each feature vector (pixel intensity or block of DCT coefficients) in the current frame. This PDF is updated each frame.

Given the recent history $\{x_{t-n}, x_{t-n+1}, \dots, x_{t-1}\}$ of a feature vector x_t , the distribution of the current feature vector is:

$$P(x_t|\theta_t) = \sum_{k=1}^K \omega_{k,t} * \eta_k(x_t, \mu_{k,t}, \sigma_{k,t}^2) \quad (1)$$

where $\theta_t = \{\omega_{k,t}, \mu_{k,t}, \sigma_{k,t}^2\}_{k=1}^K$ is the model parameter vector and $\omega_{k,t}, \mu_{k,t}, \sigma_{k,t}^2$ are the weight, mean and variance of the k th Gaussian distribution, η_k , in the mixture and $(\sum_{k=1}^K \omega_{k,t}) = 1$. The Gaussians are D -dimensional, where D is equal to the size of the feature vector x_t . To simplify the computation, as in [8], [15], we make the following simplifications: $\Sigma_{k,t} = \sigma_{k,t}^2 * I$ (I is the unit matrix); although the Gaussians are D -dimensional, the variance is a scalar value (but the mean is a vector of length D).

The algorithm has two stages. The first stage is the binary classification of the current feature vector as BG or FG according to the background model; the second stage is the online update of the parameters using the current feature vector. The binary classification stage takes place as follows. The algorithm selects the first B most significant distributions to be part of the background model, where:

$$B = \arg \min_b \left(\left(\sum_{k=1}^b \omega_{k,t} \right) > T \right) \quad (2)$$

and T is a user-defined threshold ($0 < T < 1$). For the current feature vector x_t , a matching distribution \hat{k} is sought using the following equation:

$$(x_t - \mu_{k,t-1})^T (x_t - \mu_{k,t-1}) < f * \sigma_{k,t-1}^2 \quad (3)$$

where f is a deviation. If more than one matching distribution is found, the one closest to the feature vector is selected. This is done using the following equation:

$$\hat{k} = \min_k (x_t - \mu_{k,t-1})^T (x_t - \mu_{k,t-1}) / \sigma_{k,t-1}^2 \quad (4)$$

The current feature vector x_t , is classified as BG if there is a matching distribution \hat{k} and this matching distribution is part of the B selected distributions to represent the background ($\hat{k} \leq B$). Otherwise, x_t is classified as FG.

Next, the online update stage takes place as follows. If a matching distribution \hat{k} is found, the parameters of that distribution are updated as follows:

$$\begin{aligned} \omega_{\hat{k},t} &= (1 - \alpha) * \omega_{\hat{k},t-1} + \alpha \\ \mu_{\hat{k},t} &= (1 - \rho) * \mu_{\hat{k},t-1} + \rho * x_t \\ \sigma_{\hat{k},t}^2 &= (1 - \rho) * \sigma_{\hat{k},t-1}^2 + \rho * (x_t - \mu_{\hat{k},t})^T (x_t - \mu_{\hat{k},t}) \end{aligned} \quad (5)$$

where α is a user-defined learning rate ($0 < \alpha < 1$) and ρ is a calculated learning rate:

$$\rho \approx \frac{\alpha}{\omega_{\hat{k},t}} \quad (6)$$

The parameters of the other distributions ($k \neq \hat{k}$) are updated as follows:

$$\omega_{k,t} = (1 - \alpha) * \omega_{k,t-1} \quad \mu_{k,t} = \mu_{k,t-1} \quad \sigma_{k,t}^2 = \sigma_{k,t-1}^2 \quad (7)$$

If no matching distribution is found, the one with the smallest weight ($\omega_{k,t-1}$) is replaced by a new one with the following parameters:

$$\omega_{k,t} = \omega_0 \quad \mu_{k,t} = x_t \quad \sigma_{k,t}^2 = \sigma_0^2 \quad (8)$$

where ω_0 is a small weight and σ_0^2 a large variance.

After updating the parameters, the weights are renormalized, so that they still add up to 1 and the distributions are ordered in descending order of their significance, $s_{k,t}$, where:

$$s_{k,t} = \frac{\omega_{k,t}}{\sigma_{k,t}^2} \quad (9)$$

The reordering is necessary to ensure that only the relevant distributions are selected as part of the background model. A distribution is considered to be relevant when it has a large weight (evidence) and a small variance.

C. Hardware Acceleration

As mentioned in the introduction, real-world video surveillance systems can include deployments of a large number of cameras, which in turn generate a large amount of video data. Processing all this data in real-time requires a huge amount of computation. Modern CPUs commonly have multiple cores (typically 4–8) capable of parallel processing, where each core is equipped with short vector units capable of handling multiple data elements simultaneously. To handle the computation in a scalable and efficient manner, it is necessary to take advantage of all the resources available on the modern CPUs. To do so, the use of specific programming APIs and technologies is needed.

To use such CPUs to their full capacity, let us imagine a very simple image processing operation. The subtraction of two images, in a classic scenario, requires iterating over all elements in the two images and subtracting the pixel intensities at each location, using one core. By using, for example, a CPU with 4-cores and 4-ways SIMD processing units, the images are first split in four and the individual parts processed in parallel on the 4-cores; then each part is further split into vectors of four pixels which are processed simultaneously using the SIMD units. These optimizations can in theory give a speedup of $16\times$ compared to a purely serial implementation.

At present, OpenMP [22] is the industry standard for parallel programming of shared memory multiprocessors. Multicore CPUs are considered to be shared memory multiprocessors. OpenMP uses the shared memory model to represent the interaction between threads and the memory, and uses the fork/join execution model to manage multithreaded execution [23]. The shared memory model has the following characteristics. One or more threads can run on a processor/core, but all threads have access to the shared memory. Data can exist as shared (in which case it can be accessed by all threads) or private (in which case it can be accessed only by the thread that owns it). Data transfers between threads and the private or shared memory is transparent, but the synchronization is mostly explicit (i.e. it is the programmer's responsibility). The fork/join execution model has the following characteristics. The processing starts with a single thread, called the master thread, and it runs sequentially until a parallel region construct is encountered. When this happens, the master thread will create a team of parallel threads (fork), and will divide the execution of the code inside the parallel region among the team of threads. After a team of threads finishes the code inside the parallel region, they synchronize and terminate (join), leaving the master thread to continue the program sequentially. From

the programmer's point of view, OpenMP is composed of a set of compiler directives, a library of support functions and a set of environment variables.

Furthermore, each core commonly has both scalar units and short vector units, so vectorization is a form of parallelism available at the core level and the individual CPU cores can achieve vectorization independently. The vector units are single instruction, multiple data (SIMD) processing units and can process multiple data elements using the same instruction (data-level parallelism). The use of SIMD units should accelerate the calculations by a factor equal to the length of the unit (for example, a 4-way SIMD processing unit should in theory accelerate the calculations by a factor of four). On the Intel and AMD CPU micro-architectures the vector processing units use vector registers of 128 bits, which are able to handle the following data types: integer (usually 8×16 -bits or 4×32 -bits) and float (4×32 -bits or 2×64 -bits). This technology is called Streaming SIMD Extensions Technology (SSE) [24]. Some modern compilers can auto-vectorize scalar code to some extent, but usually just for simple cases. For complex cases, programmers have to explicitly use the vector data types and write code in a SIMD manner. From the programmer's perspective, SSE programming [24] can be done in C/C++ using compiler intrinsics (which are compiled to one or a small set of assembly instructions), or directly in assembly language.

There are three important aspects that anyone should be aware of when developing algorithms using SSE [24]–[27]: the alignment of data in memory, the layout of data in memory, and the instructions available. The best performance for accessing data in memory using the SSE instructions is achieved when the memory address of the accessed data (usually a 4-element vector) is divisible by 16 (16-byte aligned). The performance impact of misaligned memory accesses and methods to avoid them have been previously studied in [27].

When vectorizing code which is processing an array of structured objects (e.g. an array of grouped Gaussian parameters), the recommended layout for data is to have a separate array for each component of the structure (e.g. an array of means, an array of variances, etc.) This is called Structure of Arrays (SoA). An alternative layout is the more 'natural' (from an object oriented point of view) layout of a single array of structures. This is called Array of Structures (AoS). The advantage of SoA is that it keeps the homogeneous data components together in memory. This makes it easy to load and process them with the same instruction (or set of instructions). The performance impact of using different data memory layouts has been previously studied in [26]. For example, an RGB image using SoA layout is represented by three memory areas, where each memory area stores the data of one channel for all pixels. In contrast, using AoS layout, just one memory area is required, but each pixel occupies three consecutive memory locations, one for each channel (R, G, B).

The SSE vector instructions can be grouped into categories [24], such as: load/store, arithmetic, logic, comparisons, miscellaneous (min/max, type conversions, rounding, cache control, etc.) and application specific (fast-block difference, etc.) and they are closely tied to the type of the vector elements. For most float vector instructions, SSE offers a corresponding

scalar version. These should not be confused with the classic (x86) scalar instructions because they operate on the vector registers (even though they are applied only on the first element of the vector register). Instructions for different data types, and even different instructions for the same data type, can be serviced by different parts of the processor. In fact, modern CPUs can issue up to six instructions per clock cycle, which increases the chances of executing instructions in parallel if there are no data dependencies between the operands of consecutive instructions. SSE is IEEE floating-point compliant, except for the reciprocal and square root reciprocal instructions. Division is one of the most expensive SSE instructions. The two instructions have been introduced to speed up algorithms involving division where high precision is not the first requirement. Another feature of the SSE instruction set is that it can handle conditional code without having to individually process each element in a vector register. This is possible by using a combination of logical SSE instructions or the newer *blend* instruction. The *blend* instruction takes two vector registers and selects their elements according to a mask. Usually the mask comes from a vector comparison instruction.

IV. ACCELERATING THE MIXTURE OF GAUSSIANS IN THE COMPRESSED DOMAIN

To accelerate multistream background subtraction, there are two forms of parallelism to be exploited: multicore (at the high level) and vector processing (low level). Our approach to exploiting multiple cores is straightforward: each input video stream is processed by a separate thread, with no need for communication with other threads. We do not partition the data for a video stream across multiple threads. So there is no significant programming effort involved in exploiting multicore parallelism. Each thread performs the partial decompression followed by background modeling on a complete video stream.

To exploit vector processing, our approach is to vectorize the processing of Gaussian mixtures, rather than processing image blocks in parallel. Since it is sufficient to use a mixture of four Gaussians, and given the size of SSE vector registers, it is appropriate to vectorize the processing of the Gaussians. Further, since each Gaussian has three parameters (plus sometimes a fourth index—see below), we also have the option of using either a SoA approach or, if need be, an AoS approach.

One further possible dimension for increasing the amount of parallelism is to reduce the precision of the Gaussian components from 32 bit float to 16 bit fixed point. This would in theory provide a further speed up of $2 \times$. However, having explored this, it was discovered that the increased complexity of the coding for 16 bit representation meant that the overall performance was actually slower than for processing 32 bit floats. In what follows, we therefore do not consider the question of precision any further.

Because of the complexity of SSE programming, the remainder of this section focuses on the SSE accelerated implementation of the MoG algorithm in the compressed domain.

As previously specified, the MoG algorithm in the compressed domain classifies and models the PDF of an entire block of DCT coefficients simultaneously using a mixture of




Fig. 2. Data layout used for the model parameters in the compressed domain.

K Gaussian distributions. Therefore, the parameters for one Gaussian distribution in a mixture are: one weight (which represents the probability of a block to match that particular distribution, and is calculated based on previous observations), one variance and a set of D-means (one for each DCT coefficient in the block). The number of Gaussian distributions in a mixture has been set to a fixed value of four ($K = 4$). This is (more than) adequate algorithmically, and matches the hardware constraints imposed by the size of the vector registers ($128 - \text{bits} = 4 \times 32 - \text{bits}$). One variation which we have introduced to the standard algorithm (introduced to map better onto the vector registers) is that we always maintain four Gaussians, instead of dynamically varying the number up to four. Also, in order to avoid some expensive memory copy operations in the later steps of the algorithm, a new index parameter, ID, is added to each distribution (see step 6 below for more details). These IDs will be used to access the correct set of means in the means memory area corresponding to a particular mixture. The IDs of the distributions in a mixture are initialized to (0, 1, 2, 3). The representation (layout) of the background model parameters in memory is SoA (Fig. 2).

Before processing a DCT block, the weights, IDs and variances of all distributions in the associated mixture are loaded into three vector registers. Due to these loads the CPU will also automatically load these values into the level 1 (L1) cache memory. The four sets of means corresponding to a mixture are loaded only when needed in the later steps. The parameters are saved back in the memory after their values have been updated and the process is repeated for the next block. The result of the assessment of the current block is a value which represents the binary classification of the block as FG or BG. In practice, the update steps for a block are as follows.

Step 1. Identify how many distributions make up the background model, according to (2). The result is an index, idx_1 , with a value between zero and three. The distributions are already ordered according to their significance. The algorithm loads the weight of one distribution at a time (using SSE scalar load—*movss*), adds it to a sum (using SSE scalar add—*addss*) and compares the sum with the matching threshold T (using SSE scalar compare—*comiss*).

Step 2. Identify the closest distribution to the current block using (3) and (4). The result of this step is an index, idx_2 , with value between zero and three if the current block matches any of the distributions, or -1 otherwise. The algorithm computes the distances between the current block and all the distributions in the mixture simultaneously using vector instructions. The computation of one of

the distances is equivalent to processing a scalar product between a vector and itself, which is a classic reduction problem, where the vector is the result of the difference between the current DCT block and the means set of one of the distributions. The distances can be processed simultaneously because the DCT block is read only once, and used for all distances. These distances are then stored in memory. Further, using the SSE scalar instructions, a short algorithm is implemented to find the index of the closest distribution based on the already computed distances (which is a typical “minimum” problem).

Step 3. Classify the current block of DCT coefficients as BG or FG using the indices computed in step 1 (idx_1) and step 2 (idx_2). The pixel is classified as BG if $idx_2 \geq 0$ and $idx_1 \geq idx_2$.

Step 4. Update all the distribution parameters. How this is done depends on the index computed in step 2 (idx_2), which identifies the closest distribution. There are two cases:

Step 4.a. If there is a matching distribution ($idx_2 \geq 0$), then its parameters are updated according to (5) and (6). The parameters of the remaining distributions are updated according to (7). Specific SSE programming techniques are used to avoid “branching” code. First, the weights and variances are updated using vector instructions twice, once according to (5) and once according to (7). Then, the index from step 2 (idx_2) is used to build the mask used by the vector Blend instruction (see Section III-C for details) to select the correct updated parameters for each distribution. This procedure adds more computation but avoids “branching” code. Note that the update of the means is done last and is highly efficient, because it is performed on the full length of the vector registers (see SoA layout, Section III-C).

Step 4.b. If there is no matching distribution, the distribution with the smallest weight is replaced by a new one according to (8). For this, the distribution containing the smallest weight is identified (again using SSE scalar instructions) and then its parameters are reinitialized. An important note for step 4(a and b) is that the selection of the correct set of means associated with a particular distribution in a mixture is done using the ID parameter which we have added earlier to each distribution (see step 6).

Step 5. Renormalize the weights to add up to 1. For this, specific SSE programming techniques are used in order to handle the horizontal operation (sum) on the vector register containing the weights. The vector division necessary for renormalization (*divps*) has been replaced with a vector multiplication (*mulps*) between the weights register and the vector reciprocal (*rcpps*) of the previously calculated sum.

Step 6. Reorder the distributions based on their significance. Here, few implementation details deserve mention. The significance values, $s_{k,t}$, are computed according to (9) using SSE vector instructions (here the vector division has been replaced). Further (Fig. 3), to prepare for reordering all distributions in a mixture, a 4×4 matrix transposition operation is required. This transformation will reorganize the parameters of the unsorted distributions from

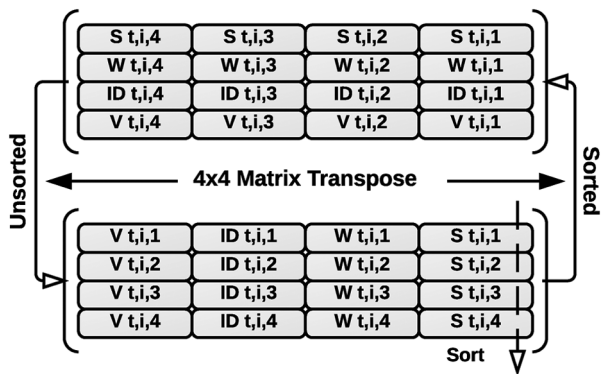


Fig. 3. Data transformations and sorting.




Fig. 4. Data layout used for the model parameters in the spatial domain.

SoA to AoS format. The operation can be done using only vector registers (without accessing the memory) and some specific SSE instructions for reordering the data in these registers. For the actual sorting we use a 4-ways sorting network [28] which offers an optimal number of comparisons. The sorting network implementation requires SSE scalar instructions for handling the conditions and vector instructions to swap a full register at a time (*movps*). After sorting, another 4×4 matrix transposition operation is necessary to switch the parameters back to SoA format. This implementation turns out to be highly efficient because no memory operations are required and the instructions involved in the transposition are low latency. But the most important optimization here is the introduction of the ID parameter to each distribution. The introduction of this parameter eliminated the necessity to move around the large sets of means associated with the D -dimensional Gaussian distributions during the sorting operation.

For benchmarking reasons an SSE optimized version of the MoG algorithm in the spatial domain was also implemented. Here, the parameters for one Gaussian distribution in a mixture are: one weight, one mean and one variance per pixel. The ID parameter is not necessary anymore. Due to the same hardware constraints the number of Gaussian distributions in a mixture has also been set to four ($K = 4$). The representation (layout) of the background model parameters in memory remains SoA (Fig. 4), as in processing in the compressed domain. Although the SSE implementations of the MoG algorithm for the two domains seem very similar, a few differences are visible straight away. The compressed domain implementation is using a very compact memory representation of the model parameters, which is possible because it keeps the means of all distributions of a mixture associated with a block of DCT coefficients in a contiguous memory block. This makes their update, during step 4,

much more efficient. Also, the compressed domain implementation has to maintain a lot fewer parameters. As already mentioned in [8], for a 8×8 pixel block, the algorithm uses $192 * K$ parameters in the spatial domain compared with just $66 * K$ in the compressed domain (K represents the number of distributions in a mixture). The ID parameter is not included in these numbers because it does not require updating.

V. EXPERIMENTS AND RESULTS

In this section we describe experiments aimed at answering the question: “Can we obtain significant acceleration of the MoG algorithm in the compressed domain, while maintaining foreground detection performance?”

Section V-A evaluates the acceleration obtained with partial decompression, multicore and vector processing. Section V-B is concerned with comparing the foreground detection of the MoG algorithm in the compressed domain against the standard implementation in the spatial domain.

A. Partial Decompression and Hardware Acceleration for Background Modeling

1) *Partial Decompression*: In the spatial domain, a frame is represented as a 2D matrix of pixels. To get access to the spatial pixel data in a video frame, the video stream has first to be fully decompressed. In the compressed domain, a frame is represented as a 2D matrix of blocks of DCT coefficients. This representation of a frame can be extracted from a video stream without fully decompressing the stream. Before any computation can begin, the video sequences need to be partially decompressed (if processing in the compressed domain) or fully decompressed (if processing in the spatial domain).

To evaluate our implementation of partial decompression, its performance was compared with five different full decompression implementations. For the latter, three different libraries were used: the standard Libjpeg v8c, Intel’s Libjpeg-based library (IJG v6b), and Intel’s proprietary library UIC which uses the Intel Performance Primitives Library (IPP). Two different compilers were also used: Microsoft’s C++ Compiler (MC), and Intel’s C++ Compiler (IC). Our MJPEG partial decoder makes use of the IPP library for Huffman decoding. The tests were carried out by simulating four camera streams, by replicating the input video sequence.

To test the performance on “big data”, we used a video sequence with a high resolution, called PETS2001 (Fig. 5), which comprises 2688 frames, each of size 768×576 pixels. All images were first loaded into memory (decompression times do not include the time to read the images from the hard drive). The CPU used for these tests is an Intel i7 960 4-core CPU. For the multithreaded implementation, OpenMP was used to launch 4 threads, mapped to the 4 cores.

Table I shows the decompression times per frame, for the five full decompression versions, and for the partial decompression. We can draw the following conclusions from the table.

The fastest full decompression version is UIC. The use of IPP approximately doubles the performance. The speed of partial decompression is 4.3 times that of the fastest full decompression version. At a frame rate of 25 frames per second, the time for partial decompression equates to the ability to process 210



Fig. 5. PETS2001 video sequence.

TABLE I
DECOMPRESSION SPEED RESULTS PER FRAME

	1 Core	4 Cores
Libjpeg v8c (MC)	8 ms	1.82 ms
IJG v6b (MC)	7.2 ms	1.6 ms
IJG v6b (IC)	7 ms	1.54 ms
IJGmod v6b (IC & IPP)	3.72 ms	0.86 ms
UIC (IC & IPP)	3 ms	0.82 ms
Partial Decoder (IC & IPP)	0.84 ms	0.19 ms

simultaneous video streams (not allowing for the input of the streams to memory). When processing fewer video streams (e.g. 64), this allows more time for processing tasks such as background modeling and subtraction.

2) *Hardware Acceleration for Background Modeling*: Experiments on the PETS2001 test video sequence (Fig. 5) showed that, averaged over the entire video sequence, only 14–16% of the 64 DCT coefficients in a block are nonzero, and all DCT blocks have 32 or fewer nonzero DCT coefficients. The number of nonzero DCT coefficients per block depends on two factors. First, the quantization table used during compression; a stronger quantization table will increase the number of zero DCT coefficients. Second, the complexity and level of activity in the scene; for example, in the PETS2001 video sequence, the sky areas can be encoded with just one nonzero DCT coefficient per block, the grass areas are typically encoded with 6–10 nonzero DCT coefficients per block, and the regions covered by the cars are encoded with over 10 DCT coefficients per block. In general, blocks containing many edges or complex texture patterns are encoded using a larger number of DCT coefficients. These experiments led us to test the MoG background modeling algorithm in the compressed domain using 32 and 16 coefficients per block.

Table II presents the execution times for the spatial domain implementation versus compressed domain implementation for three different numbers of DCT coefficients per block (64, 32, 16). For each case, it shows execution times for a single core and a 4-core implementation, first using scalar (nonvector) and then using vector implementations of the algorithm. The tests were run on the same PETS2001 video sequence and under the same conditions.

The results show that the multicore, vectorized versions of the compressed domain implementation using 16 DCT coefficients outperform the purely serial spatial domain version by an impressive factor of 170 \times .

The overall speedup is the result of the combination of several different optimization, which all complement each other. By processing in the compressed domain instead of the spatial

TABLE II
PROCESSING SPEED RESULTS PER FRAME

		1 Core	4 Cores
Spatial Domain	Scalar	23.342 ms	6.485 ms
		1x	3.6x
	Vector	17.493 ms	4.714 ms
		1.3x	4.9x
Compressed Domain - 64	Scalar	11.858 ms	3.082 ms
		1.9x	7.5x
	Vector	0.930 ms	0.616 ms
		25x	37.9x
Compressed Domain - 32	Scalar	5.854 ms	1.551 ms
		3.9x	15x
	Vector	0.566 ms	0.342 ms
		41.2x	68.2x
Compressed Domain - 16	Scalar	2.927 ms	0.781 ms
		7.9x	29.8x
	Vector	0.438 ms	0.137 ms
		53.3x	170.3x

domain, the number of modeling parameters is immediately reduced by almost a third. Also, the number of coefficients per block can be reduced from 64 to 16 while maintaining segmentation quality (see Section V-B). Together, these algorithmic optimizations contribute a speed up of approximately 8 \times .

The subsequent individual contributions of the two forms of parallelism are: vectorization which contributes a further speedup of 6.7 \times (this is greater than the expected 4 \times maximum because the low level of programming required leads to reduced internal data transfers) and the use of multicore (four core) contributes a further 3.2 \times .

With a frame rate of 25 frames per second, the optimized implementation of the algorithm in the compressed domain using 16 DCT coefficients per block equates to the ability to perform background modeling on 292 video streams in real-time. This is equivalent to processing one frame every 0.14 ms. Further, by using the results from the previous section (where it was shown that a JPEG image can be partially decompressed in around 0.19 ms) it can be seen that, by using the same implementation, it is possible to perform combined partial decompression and background modeling on an equivalent of 122 video streams in real-time. The entire analysis is based on video streams with frames of 768 \times 576 pixels in size in order to provide relevant results.

B. Quantitative and Qualitative Analysis

For quantitative evaluation we calculate for every frame the number of foreground pixels correctly classified as foreground, called true positives (TP); the number of background pixels in a frame misclassified as foreground pixels, called false positives (FP); the number of background pixels in a frame classified as background, called true negatives (TN) and the number of foreground pixels in a frame misclassified as background, called false negatives (FN). From these we calculate, the total error (TE):

$$TE = FN + FP, \quad (10)$$

TABLE III
QUANTITATIVE ANALYSIS OF THE ALGORITHMS USING WALLFLOWER DATASET

Algorithm	Error Type	Problem Type							TE*	TE
		MO	TOD	LS	WT	C	B	FA		
MoG(s)	FN	0	877	1917	1730	783	1720	2502	9912	23752
	FP	0	275	12592	302	1544	774	588		
MoG(c)-64	FN	0	601	256	901	111	1861	1812	9271	25281
	FP	0	141	15754	780	1514	280	1270		
MoG(c)-32	FN	0	601	256	1083	111	1961	1812	9579	25653
	FP	0	141	15818	706	1642	252	1270		
MoG(c)-16	FN	0	601	128	618	111	1965	1807	9287	25297
	FP	0	205	15882	753	1578	192	1457		

the probability of detection (P_D) and false alarm (P_{FA})

$$P_D = \frac{TP}{TP + FN} \quad P_{FA} = \frac{FP}{TN + FP}, \quad (11)$$

the precision (P) and recall (R)

$$P = \frac{TP}{TP + FP} \quad R = \frac{TP}{TP + FN}, \quad (12)$$

the similarity (S) and $F1$ score

$$S = \frac{TP}{TP + FP + FN} \quad F1 = \frac{2 * P * R}{P + R}. \quad (13)$$

The value range of these metrics is between zero and one. For P_D (also called *Recall*), *Precision*, *Similarity* and $F1$, [12]–[14], higher values represent greater accuracy. For P_{FA} , smaller values represent greater accuracy.

Ideally, an evaluation of the above methods should use datasets with video resolutions typical of modern cameras (e.g. 768×576). However, there are certain standard datasets which are commonly used in the literature. These are usually low resolution (e.g. 160×120 or around this size). While processing in the compressed domain is much better suited to larger frame sizes, for familiarity we have selected the Wallflower dataset and a small collection of video sequences with dynamic background for the quantitative and qualitative analysis.

1) *Wallflower Dataset*: This dataset contains seven different scenes, which cover many of the possible scenarios that can arise when studying algorithms for foreground detection: displaced background (MO); gradual change in the global illumination (TOD); sudden change in the global illumination (LS); multimodal background (WT); low contrast between foreground and background (C); fast moving objects during background initialization (B); and slow moving object appearance during initialization (FA). The dataset contains one manually segmented frame per video sequence.

The algorithms compared are: MoG(s) which is the algorithm implementation in the spatial domain and MoG(c)-64, 32, 16 which is the algorithm implementation in the compressed domain using 64, 32 and 16 DCT coefficients per block. The parameters are presented in Table IV where SD are the parameters

TABLE IV
PARAMETERS SET FOR WALLFLOWER DATASET

		k	T	f	α	ω_0	σ_0^2
All Sequences	SD	4	0.86	4	0.002	0.05	4096
	CD	4	0.86	4	0.002	0.05	36

for spatial domain and CD are the parameters for compressed domain (see Section III-B for the meaning of the parameters). The MoG(c) implementation requires a higher initial variance because it operates at the block level. The parameter values are the same for all seven video sequences and they have been selected to generate optimal classification results over all video sequences, not for individual sequences. Also, they have not been changed during the tests. To remove random noise, post-processing has been applied using morphological filtering using a small diamond kernel with a diameter of three pixels.

Clearly all algorithms fail for the LS sequence as also reported in [29]; so for comparison, instead of TE , we use TE^* which represents the total error without the errors generated by the LS video sequence. The results are presented in Table III for all seven video sequences. From the table, it can be seen that MoG(c)-64, 32, 16 all have a smaller TE^* than MoG(s). Also, our algorithms have similar TE and TE^* values with the ones reported in [29]. Fig. 6 presents the qualitative results which show a good correspondence between the segmentation masks obtained with the MoG(c)-64, 32, 16 and those for MoG(s). Obviously, the former have a blocky appearance due to the fact we model the background at the block level and also the images are low resolution. However, we argue that the foreground definition is sufficient for further analysis such as tracking. This is supported by the quantitative results in Table III.

2) *Dynamic Background Dataset*: Although useful for benchmarking our results against previous work, we believe that the Wallflower dataset is not particularly suited to emphasizing the advantages of processing in the compressed domain. The drawbacks are the low resolution of the video frames (160×120 pixels, corresponding to only 20×15 blocks of DCT coefficients) and the fact that the video sequences contain either very small foreground objects (no more than 3–4 blocks of 8×8 pixels) or large foreground objects that cover most




Fig. 6. Wallflower dataset results. From first to last row: test frames, ground truth for test frames, results for MoG algorithm in the spatial domain and compressed domain using 64, 32, and 16 DCT coefficients.

of the video frame. For video sequences which contain small foreground objects, the misclassification of a few blocks has a significant impact in terms of FN and FP pixels. For small video frames, which contain a large object, the misclassification of blocks at the foreground/background boundary can also have a big impact in terms of FP pixels. This is also an issue with many other well-known test sequences. Also, we are particularly interested in sequences containing dynamic texture backgrounds caused either by sensor or background motion. In particular, we are currently working on a project concerned with developing an intelligent surveillance system for public transport [30].

For these reasons we further evaluated the algorithm performance using three other video sequences containing dynamic texture backgrounds. The first video sequence has a dynamic texture background of water waves [31], with a foreground of a floating bottle and with a resolution of 320×240 pixels. The first eighty frames show just the background water motion, whilst the subsequent thirty frames contain the bottle floating on the water. The second video sequence is a particularly challenging surveillance scene captured inside a moving bus, also with a resolution of 320×240 pixels. Here, the initial ninety-four frames show the inside of the upper bus saloon whilst in motion and with no passengers. In addition to fast illumination changes, those parts of the scene corresponding to windows contain dynamic textures due to the moving background outside the

bus. The remaining one hundred and six frames show a passenger coming into the visual field of the camera and taking a seat on the bus. Many of the wallflower challenges are contained within this one dataset. The bus sequence and its ground truth can be downloaded from [32]. Ground truth is available for all frames for both video sequences. The third sequence contains a scene of the sea and a path. After a period a person walks into the scene along the path. The resolution for this video sequence is 160×128 pixels, and ground truth is provided for only twenty of the frames containing the subject.

Fig. 7 shows relevant ROC curves for each video sequence, created using parameters from Table VI (SD—spatial domain, CD—compressed domain). To generate the ROC curve, f was varied, and for each value the average P_D and P_{FA} was calculated over all frames in the sequence. Analysis of Fig. 7 shows that the performance of MoG(c) is better than that of MoG(s) for all three sequences (identical curves to that of MoG(c)-16 were obtained for MoG(c)-64 and -32, however, for clarity only the former is shown). The performance is markedly superior for the bottle sequence and less so for the bus and sea sequences.

Table V shows further quantitative analysis of the two video sequences. For this analysis, we have selected an optimal f value using the ROC curve ($f = 10$ for bottle video sequence, $f = 24$ for bus video sequence, $f = 8$ for sea video sequence), and calculated the average of the different metric values over all the frames for which ground truth is available. Table V




Fig. 7. ROC curves for: bottle (left), bus (middle), sea (right) video sequences.

TABLE V
QUANTITATIVE ANALYSIS OF THE ALGORITHM USING BOTTLE, BUS AND SEA VIDEO SEQUENCES

Algorithm	Avg. Bottle Seq.					Avg. Bus Seq.					Avg. Sea Seq.				
	Sim.	F1	Prec.	Rec.	P _{FA}	Sim.	F1	Prec.	Rec.	P _{FA}	Sim.	F1	Prec.	Rec.	P _{FA}
MoG(s)	0.17	0.29	0.76	0.18	0.08	0.40	0.57	0.80	0.45	0.05	0.66	0.80	0.88	0.73	0.03
MoG(c)-64	0.50	0.67	0.82	0.56	0.02	0.47	0.64	0.80	0.54	0.03	0.65	0.78	0.88	0.70	0.03
MoG(c)-32	0.50	0.67	0.82	0.56	0.02	0.47	0.64	0.80	0.53	0.03	0.65	0.79	0.89	0.70	0.03
MoG(c)-16	0.50	0.67	0.81	0.57	0.02	0.47	0.64	0.80	0.53	0.03	0.65	0.79	0.89	0.71	0.03

TABLE VI
PARAMETERS SET FOR DYNAMIC TEXTURES DATASET

		k	T	α	ω_0	σ_0^2
Bottle Sequence	SD	4	0.82	0.006	0.05	4096
	CD	4	0.82	0.006	0.05	36
Bus & Sea Sequence	SD	4	0.82	0.001	0.05	4096
	CD	4	0.82	0.001	0.05	36

shows that the MoG(c) gives superior results to MoG(s) for all metrics for the first two sequences. For the sea sequence, the metric values are roughly the same for both algorithms. Fig. 8 shows the qualitative analysis of three frames from each video sequence. It can be clearly seen that the *FPs* obtained with MoG(s), due to the moving background outside the bus window, are clearly reduced for MoG(c). For the bottle and sea sequences, the *FPs* have been removed completely. On the other hand the foreground object definition for MoG(c) appears block-like due to the fact that detection is performed over 8×8 neighborhoods.

Figs. 7 and 8 emphasize the fact that the compressed domain algorithm seems to better model dynamic texture backgrounds. This is possible because the compressed domain version of the algorithm takes into account pixel neighborhoods, not just isolated pixels. Also, this makes the algorithms more robust against changes in the local illumination. Another important observation is that we can use just 16 DCT coefficients per block with no loss in detection performance.

VI. CONCLUSIONS

The implementation of background modeling presented in this paper uses compressed domain processing with a reduced

number of 16 DCT coefficients, plus hardware acceleration, to achieve a speedup of up to $170\times$ compared to a purely serial implementation of the same algorithm in the spatial domain. At 25 fps, this speed up equates to the ability to segment in real-time the equivalent of around 122 video streams (including the time for partial decompression), which exceeds our initial target of 64 simultaneous video streams.

The overall speedup is the result of the combination of several different optimizations, which all complement each other. By processing in the compressed domain instead of the spatial domain, the number of modeling parameters is immediately reduced by almost a third. Also, the number of coefficients per block can be reduced from 64 to 16 while maintaining segmentation quality. Together, these algorithmic optimizations contribute a speed up of approximately $8\times$.

The subsequent individual contributions of the two forms of parallelism are: vectorization which contributes a further speedup of $6.7\times$ (this is greater than the expected $4\times$ maximum because the low level of programming required leads to reduced internal data transfers) and the use of multicore (four core) contributes a further $3.2\times$.

In the spatial domain implementation, the use of the SIMD units provided only $1.3\times$ speed up, because of the low computation to data ratio. This shows the interesting result that operating in the compressed domain can increase the potential for hardware acceleration.

From the point of view of programming effort, the use of multicore required minimal effort. On the other hand, the use of SIMD processing was by far the most difficult task. It required considerable programming effort and a steep learning curve because of the low level, architecture dependent nature of the programming. The effort in implementing the compressed domain processing lay somewhere between these two steps, once the partial decompression code was available.




Fig. 8. Bottle (left), bus (middle), and sea (right) dataset results. From top to bottom: test frames, ground truth for test frames, results for MoG algorithm in the spatial domain and compressed domain using 64, 32, 16 DCT coefficients.

For the case of background modeling at least, the speed up obtained did not come at the expense of a loss in segmentation quality when moving from the spatial to the compressed domain.

Future work will extend the foreground detection in the compressed domain to human detection and tracking for a multi-camera system.

REFERENCES

[1] T. Bouwmans, F. E. Baf, and B. Vachon, "Background modeling using mixture of gaussians for foreground detection—A survey," *Recent Patents Comput. Sci.*, vol. 1, no. 3, pp. 219–237, Nov. 2008.

[2] M. Cristani, M. Farenzena, D. Bloisi, and V. Murino, "Background subtraction for automated multisensor surveillance: A comprehensive review," *EURASIP J. Adv. Signal Process.*, vol. 2010, pp. 1–24, 2010.

[3] H. Hassanpour, "Video frame's background modeling: Reviewing the techniques," *J. Signal Inf. Process.*, vol. 2, no. 2, pp. 72–78, 2011.

[4] S. Primechaev, A. Frolov, and B. Simak, "Scene change detection using DCT features in transform domain video indexing," in *Proc. 14th Int. Workshop Systems, Signals and Image Processing, 2007 and 6th EURASIP Conf. Focused on Speech and Image Processing, Multimedia Communications and Services*, 2007, pp. 369–372.

[5] M. Lamarre and J. Clark, "Background subtraction using competing models in the block-DCT domain," in *Proc. 16th Int. Conf. Pattern Recognition, 2002*, 2002, vol. 1, pp. 299–302.

[6] S. Schwartz, "A transform domain approach to real-time foreground segmentation in video sequences," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing, 2005 (ICASSP '05)*, 2005, vol. 2, pp. 685–688.

[7] L. Dong and S. Schwartz, "DCT-based object tracking in compressed video," in *Proc. 2006 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2006)*, pp. 665–668.

[8] W. Wang, J. Yang, and W. Gao, "Modeling background and segmenting moving objects from compressed video," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 18, no. 5, pp. 670–681, May 2008.

[9] J. Meng, Y. Juan, and S. Chang, "Scene change detection in a MPEG compressed video sequence," in *Proc. IS&T/SPIE Symp.*, Feb. 1995, vol. 2419, pp. 1–12.

[10] Y. Chen and I. Bajic, "Compressed-domain moving region segmentation with pixel precision using motion integration," in *IEEE Pacific Rim Conf. Communications, Computers and Signal Processing, 2009 (PacRim 2009)*, pp. 442–447.

[11] F. Porikli and F. Bashir, "Compressed domain video object segmentation," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 20, no. 1, pp. 2–14, Jan. 2010.

[12] S. C. Huang, "An advanced motion detection algorithm with video quality analysis for video surveillance systems," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 21, no. 1, pp. 1–14, Jan. 2011.

[13] F. C. Cheng, S. C. Huang, and S. J. Ruan, "Scene analysis for object detection in advanced surveillance systems using laplacian distribution model," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 41, no. 5, pp. 589–598, Sep. 2011.

- [14] S. C. Huang and F. C. Cheng, "Motion detection with pyramid structure of background model for intelligent surveillance systems," *ELSEVIER Eng. Applicat. Artif. Intell.*, vol. 25, pp. 1338–1348, Feb. 2012.
- [15] C. Stauffer and W. Grimson, "Adaptive background mixture models for real-time tracking," in *Proc. 1999 IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, pp. 246–252, (Cat. PR00149) IEEE Comput. Soc.
- [16] M. Poremba, Y. Xie, and M. Wolf, "Accelerating adaptive background subtraction with GPU and CBEA architecture," in *Proc. 2010 IEEE Workshop Signal Processing Systems*, Oct. 2010, pp. 305–310.
- [17] "CUDA C Programming Guide" NVIDIA [Online]. Available: <http://developer.nvidia.com/cuda/nvidia-gpu-computing-documentation>
- [18] D. Weimer, S. Kohler, C. Hellert, K. Doll, U. Brunsmann, and R. Krzikalla, "Gpu architecture for stationary multisensor pedestrian detection at smart intersections," in *Proc. 2011 IEEE Intelligent Vehicles Symp. (IV)*, Jun. 2011, no. Iv, pp. 89–94.
- [19] V. Pham, P. Vo, V. T. Hung, and L. H. Bac, "GPU implementation of extended gaussian mixture model for background subtraction," in *Proc. 2010 IEEE RIVF Int. Conf. Computing and Communication Technologies, Research, Innovation, and Vision for the Future (RIVF)*, Nov. 2010, pp. 1–4.
- [20] I. E. Richardson, *Video Codec Design: Developing Image and Video Compression Systems*, 2002.
- [21] *Information technology - Digital compression and coding of continuous-tone still images—Requirements and guidelines.*, ISO/IEC 10918-1:1993(E) and CCITT Recommendation T.81 (1992 E), ISO, 1993 [Online]. Available: <http://www.w3.org/Graphics/JPEG/itu-t81.pdf>
- [22] "OpenMP Specifications" OpenMP-Committee [Online]. Available: <http://www.openmp.org/mp-documents/OpenMP3.1.pdf>
- [23] B. Chapman, G. Jost, and R. V. D. Pas, *Using OpenMP, Portable Shared Memory Parallel Programming*, 2007.
- [24] "Intel 64 and IA-32 Architectures and Software Development Manuals" Intel [Online]. Available: <http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>
- [25] A. Klimovitski, *SSE/SSE2 Toolbox—Solutions for Real-Life SIMD Problems*, 2001 [Online]. Available: http://www.thomasdideriksen.dk/misc/SIMD/sse_in_real_applications.pdf
- [26] W. Eckhardt and A. Heinecke, "An efficient vectorization of linked-cell particle simulations," in *Proc. 9th Conf. Computing Frontiers (CF '12)*, 2012, p. 241.
- [27] M. Alvarez, E. Salami, A. Ramirez, and M. Valero, "Performance impact of unaligned memory operations in SIMD extensions for video codec applications," in *Proc. IEEE Int. Symp. Performance Analysis of Systems and Software (ISPASS)*, 2007, pp. 62–71.
- [28] T. Furtak, J. N. Amaral, and R. Niewiadomski, "Using SIMD registers and instructions to enable instruction-level parallelism in sorting algorithms," in *Proc. 19th Ann. ACM Symp. Parallel Algorithms and Architectures (SPAA '07)*, 2007, p. 348.
- [29] K. Toyama, J. Krumm, B. Brumitt, and B. Meyers, "Wallflower: Principles and practice of background maintenance," in *Proc. ICCV'99*, 1999, pp. 255–261.
- [30] P. Miller, W. Liu, C. Fowler, H. Zhou, J. Shen, J. Ma, J. Zhang, W. Yan, K. McLaughlin, and S. Sezer, "Intelligent sensor information system for public transport-to safely go . . .," in *Proc. 7th IEEE Int. Conf. Advanced Video and Signal Based Surveillance*, 2010, pp. 533–538.
- [31] J. Zhong and S. Sclaroff, "Segmenting foreground objects from a dynamic textured background via a robust kalman filter," in *Proc. 9th IEEE Int. Conf. Computer Vision*, 2003, Oct. 2003, vol. 1, pp. 44–50.
- [32] "Bus Video Sequence" QUB [Online]. Available: <http://qub.ac.uk>



Stefan Popa received the B.Sc. degree in robotics from the Technical University of Bucharest, in 2006, and the M.Sc. degree in computer and electronic security from Queen's University Belfast, in 2011. He is currently studying for a Ph.D. degree at the Institute for Electronics, Communications and Information Technology (ECIT), Queens University Belfast. His current research interests are image processing in the compressed domain, high performance video processing, and video surveillance.



Danny Crookes was appointed to the Chair of Computer Engineering in 1993 at Queens University Belfast, and was Head of Computer Science from 1993 to 2002. He is Director of Research for Speech, Image and Vision Systems at the Institute for Electronics, Communications and Information Technology (ECIT) at Queens University Belfast. His current research interests include the use of novel architectures (especially GPUs) for high-performance image processing. He has applied expertise in language design, optimizing compilers and software generators, plus software tools for hardware description and architecture generation, to the goal of developing high level software tools to enable rapid development of real-time video processing systems. He is currently involved in projects in automatic shoeprint recognition (ESPRC), speech separation and enhancement (EPSRC) and processing of 4D confocal microscopy imagery (sponsored by INI and Andor Technologies). He has published over 200 scientific papers in journals and international conferences, and has presented tutorials on parallel image processing at several international conferences.



Paul Miller is a Senior Lecturer in the School of Electronics, Electrical Engineering and Computer Science at Queens University Belfast (QUB). He is also Research Director of the Intelligent Surveillance Systems group in the 30 M funded Centre for Secure Information Technology (CSIT). He has led several major research and development projects supported by industry and government, with a total project value of over 8 M. He has published over 70 papers in image and video analysis, including a best paper award for his work on object recognition.

Previously, he worked as a senior research scientist at the Defence, Science and Technology Organization, Australia where he led a team providing science and technology advice on unmanned aerial surveillance systems. Before that he worked as a research fellow at QUB. He received the Ph.D. degree in optical image processing from QUB in 1989. Since returning to academia, he has continued to work in video analytics for both defence and civilian CCTV applications, and also biomedical image analysis. His research interests include image restoration, segmentation, multicamera tracking, and gender/age profiling of subjects in video. More recently, his main research interest has been in the convergence of cyber and physical security through the development of a framework for a unified approach to securing both cyber and physical spaces. Specifically, he has been investigating graph-based data mining approaches that identify anomalous behavior using both physical and network access control databases. He is also investigating the use of uncertainty reasoning for enhanced situation awareness across both the physical and cyber space of enterprises. During his academic career, he has constantly worked in close collaboration with industry, including both multinational and SME companies.