

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21

GPU-enabled pavement distress image classification in real time

Kristina Doycheva (Corresponding Author)

Ruhr-University Bochum

Universitätsstr. 150, Gebäude IC 6-69, 44801, Bochum, Germany

kristina.doycheva@rub.de

+49 (2) 234 32 24640

Christian Koch

University of Nottingham

Room B27 Coates Building, University Park, NG7 2RD, Nottingham, UK

christian.koch@nottingham.ac.uk

+44 (0) 115 84 68933

Markus König

Ruhr-University Bochum

Universitätsstr. 150, Gebäude IC 6-59, 44801, Bochum, Germany

koenig@inf.bi.rub.de

+49 (2) 234 32 23047

22 **ABSTRACT**

23 Pavement assessment is a crucial process for the maintenance of municipal roads. However, the
24 detection of pavement distress is usually performed either manually or offline, which is not only
25 time-consuming and subjective, but also results in an enormous amount of data being stored
26 persistently before processing. State-of-the-art pavement image processing methods executed on
27 a CPU are not able to analyze pavement images in real time. To compensate this limitation of the
28 methods, we propose an automated approach for pavement distress detection. In particular, GPU
29 implementations of a noise removal, a background correction and a pavement distress detection
30 method were developed. The median filter and the top-hat transform are used to remove noise
31 and shadows in the images. The wavelet transform is applied in order to calculate a descriptor
32 value for classification purposes. The approach was tested on 1549 images. The results show that
33 real-time pre-processing and analysis are possible.

34

35 **INTRODUCTION**

36 In recent years, the condition of municipal roads has deteriorated rapidly, leading to increased
37 fuel consumption, thus increased emissions and environmental pollution, and even greater
38 number of vehicle damages and traffic accidents [Spielman 2014]. To reduce the negative impact
39 of deteriorated roads on the driving quality, roads need to be maintained in good condition, for
40 example by repairing parts of the road surface where pavement distress, visible as cracks or
41 potholes, is present. For this purpose, knowledge about the exact location of pavement distress is
42 required and pavement assessment is an essential task [Orr 2015].

43 Several techniques for distress detection in asphalt pavement have been proposed in the last few
44 years. The most intuitive approach is manual observation, during which an expert makes notes
45 about the condition of the road by hand while walking over the road shoulder. The evaluation is
46 performed with the help of manuals specifying criteria for pavement assessment and rating
47 [NCHRP 2004]. There also exist methods which are based on the various types of pavement data
48 being collected, such as sensor data or images of the pavement surface. Sensor devices are often
49 utilized to measure parameters of the pavement surface. This approach is referred to as *sensor-*
50 *based pavement assessment*. On the other hand, visual data obtained by images or videos of the
51 pavement surface is also used for pavement assessment. The so-called *visual-based pavement*
52 *assessment techniques* analyze features of the images or video frames with respect to criteria
53 identifying the presence of distress. Visual-based pavement assessment techniques have been
54 widely applied recently, because they are less subjective and hazardous compared to manual
55 observations [Koch et al. 2015].

56 Furthermore, these techniques can be classified as purely manual, semi-automated or automated
57 based on the manner of processing the data. The observation by experts is an example of a purely

58 manual technique, while semi-automated and automated methods require only little or no human
59 intervention. Despite of the advances in automated pavement assessment in recent years, there is
60 still room for improvement. For example, video data is usually stored before it is actually
61 processed. Considering the length of the municipal road network in Germany, which is
62 approximately 610,000 km according to the German Association of Towns and Municipalities
63 [DStGB 2014], the amount of stored data is large (approx. 5 gigabytes per kilometer). To reduce
64 this amount of data, methods capable of analyzing the pavement surface in real time are required.
65 Such methods could be employed in order to store only those images on which distress had been
66 identified and discard all other images without distress, resulting in less memory requirements
67 and less subsequent processing time needed compared to the state-of-the-art case.

68 However, although the central processing unit (CPU) technology has evolved during the last
69 decade, modern CPUs are still not able to cope with the requirement of real-time execution of
70 related analysis methods, mainly due to the fact that image pre-processing is also needed. For
71 instance, noise removal as well as correction of non-uniform background illumination needs to
72 be applied to the images to enhance their quality in order to produce more accurate analysis
73 results.

74 Yet, the real-time processing requirement can be fulfilled by utilizing Graphics Processing Units
75 (GPUs). Applied not only for graphic operations, but also for computational tasks, GPUs have
76 proven their efficiency in diverse scientific fields in recent years [Owens et al. 2005].

77 In this work, GPUs were used to accelerate the pre-processing and the analysis of pavement
78 surface images for the purpose of real-time pavement defect detection. In particular, a noise
79 removal method, a shadow removal method and an approach towards pavement analysis based
80 on the wavelet transform were implemented and validated.

81 The next two sections provide information on state of practice and research concerning pavement
82 distress detection. Afterwards, GPUs are introduced. The approach is presented in thereafter, and
83 then the implementation is described. Performance tests were carried out to evaluate the
84 capability of the proposed implementation to process the images in real time. A case study was
85 performed to validate the approach and is described in the “Case Study” section. The paper
86 concludes with a summary of the main contributions and an outlook on future developments.

87

88 **STATE OF PRACTICE**

89 In the United States, the annual assessment and reporting of pavement conditions is currently
90 performed by transportation departments. For example, the New York State Department of
91 Transportation collects a variety of information about the pavement condition in cooperation
92 with the Federal Highway Administration (FHWA) [NYSDOT 2010]. A pavement surface rating
93 survey is conducted by a team consisting of a driver and a rater. The rater assesses the condition
94 of the pavement based on what is seen on the pavement and photographs of the pavement at each
95 rating point. As stated in New Yorks’s Pavement Condition Assessment Document [NYSDOT
96 2010], the rater should be experienced in condition survey procedures and possess knowledge of
97 road construction.

98 In Germany, the state of practice is similar. For example, in Bochum in 2013 seven teams with
99 15 employees have manually been assessing the pavement condition using portable computers
100 [Buske 2013]. The current data is entered in a database by extending very detailed road maps.
101 According to Carlos dos Santos [Buske 2013], this procedure is very laborious and one team
102 consisting of two employees can only assess two kilometers of road per day.

103 Obviously, the surveys are mostly conducted manually, but as technology improves, automated
104 assessment should become possible in the near future. For instance, a rule requiring rear
105 visibility technology in all new vehicles by May 2018 has been issued by the U.S. Department of
106 Transportation's National Highway Traffic Safety Administration (NHTSA) [2014]. This rule
107 has been issued in order to expand the required field of view for all passenger cars, trucks,
108 multipurpose passenger vehicles, buses, and low-speed vehicles with a gross vehicle weight of
109 less than 10,000 lbs. According to this rule, an area behind the vehicle which encompasses 5 feet
110 laterally from the longitudinal centerline of the vehicle and extends 20 feet rearward of the
111 vehicle's rear bumper must be visible to the driver.

112

113 **STATE-OF-RESEARCH METHODS FOR VISION-BASED PAVEMENT DISTRESS** 114 **DETECTION**

115 **Pre-processing**

116 In order to guarantee accurate analysis results, pre-processing operations are applied to the
117 pavement images. An issue related to distress in pavement images is the existence of noise.
118 Varadharajan et al. [2014] calculated the blur magnitude of the images and selected only images
119 for which the blur magnitude was below a certain threshold value. Gaussian smoothing was
120 applied by Li et al. [2014] for denoising.

121 **Median filter**

122 The most commonly applied method for noise removal is median filtering [Lokeshwor et al.
123 2013, Radopoulou and Brilakis 2014]. The median filter is an order-statistics filter used very
124 often for noise reduction [Gonzalez and Woods 2006]. It introduces less blurring to the image
125 than linear filters of the same size and it is particularly effective in the presence of salt-and-

126 pepper noise. Experimental results have shown that the median filter has a good performance in
127 gray and RGB images [Ahmed et al. 2015]. The median filter replaces the value of the pixel on
128 which the kernel is centered by the median value of the gray levels in the neighborhood of that
129 pixel. To apply the median filter, the gray level values of the pixels in the neighborhood
130 including the value of the pixel itself are sorted in an ascending or descending order. Then, the
131 value in the middle of the sorted sequence is taken and assigned to the pixel in the center of the
132 kernel. Yet, the median filter is characterized by a high computational cost. The computational
133 complexity for sorting n values, a basic step within median filtering, with efficient sorting
134 algorithms is $O(n \cdot \log n)$.

135 Another problem related to pavement images is the non-uniform background illumination.
136 Commonly, the images are taken under various lighting conditions because of different weather
137 conditions or varying times of day. This results in a non-uniform background illumination and
138 lets shadows exist in the images. Since most of the analysis methods are based on the assumption
139 that distress pixels, such as crack pixels, have a darker intensity than pixels belonging to the
140 undamaged background, non-uniform background illumination could induce misleading results.

141 Several methods to handle this problem have been proposed. Varadharajan et al. [2014] selected
142 for the analysis only images taken under good weather conditions (i.e., when the weather was
143 overcast or mostly cloudy). However, the selection of the images is also a manual and time-
144 consuming process and all images have to be stored before the analysis can begin. Zou et al.
145 [2012] presented a geodesic shadow-removal algorithm which is able to preserve the cracks in
146 the images while removing shadows in the background. Cheng and Miyojim [1998] proposed an
147 image enhancement algorithm which corrects non-uniform background illumination by dividing
148 the image into rectangular windows. For each window, the average light intensity is calculated

149 and multipliers are generated for all pixels based on the window average intensity and a common
150 base intensity.

151 **Top-hat transform**

152 The top-hat transform [Gonzalez and Woods 2006] with a larger structuring element can be used
153 to estimate the background and subtract it from the image. It has been shown [Jähne and
154 Haussecker 2000; Solomon and Breckon 2010; Wu et al. 2008] that the top-hat transform can be
155 used for mitigating illumination gradients and producing evenly illuminated images without
156 shading variations. It is useful for enhancing details in the presence of shading. Opening the
157 image with a structuring element large enough so that it does not entirely fit within the details,
158 here within the distress area, produces an estimate of the background across the image. By
159 subtracting the background (i.e. the opening) from the original image, an image with more
160 uniform background can be obtained.

161 The opening $f \circ b$ of an image f by a structuring element b is denoted as

$$f \circ b = (f \ominus b) \oplus b \quad (1)$$

162 where \ominus and \oplus denote erosion and dilation, respectively. Erosion and dilation are morphological
163 operations that consist in convoluting an image with a kernel called structuring element
164 [Gonzalez and Woods 2006]. In case of dilation, the maximal gray level value overlapped by the
165 structuring element anchored at a certain pixel in the image is used to replace the value of this
166 pixel. As a result of the dilation, bright regions within the image become larger. Hence, the
167 operation is called dilation. In case of erosion, the minimal value is used, resulting in bright
168 valued areas getting thinner in a manner similar to erosion in geomorphology and geology.

169 As in the case with the median filter, the main drawback of the top-hat transform is its
170 computational complexity. The size of the structuring element required to preserve the edges or
171 details in the images leads to a vast number of pixels being considered for each anchor point.

172 **Image analysis**

173 A range of methods for distress detection in pavement images has been proposed in recent years.
174 Most of them have been specifically developed for particular types of distress, such as cracks,
175 potholes or patches. The role of digital image processing as a tool for pavement distress
176 evaluation was described by Georgopoulos et al. [1995]. A critical assessment of available
177 distress segmentation methods for crack detection and classification was presented by Tsai et al.
178 [2010].

179 Cracks are the most common distress type and, consequently, the majority of the methods
180 presented recently consider cracks. An automatic crack detection system was proposed by
181 Oliveira & Correia [2013]. The system is capable of crack type characterization and a
182 methodology for the assignment of crack severity levels was introduced. Subirats et al. [2006]
183 used wavelet transforms for crack detection, while Vivekanandreddy et al. [2014] utilized Hough
184 transforms for this purpose. Morphology-based methods have also been applied. For example,
185 Tanaka and Uematsu [1998] suggested black pixel extraction, saddle point detection, linear
186 feature extraction and connecting processing for crack detection in road surface images. Fang et
187 al. [2014] presented a crack detection technology based on an improved K-means algorithm.

188 Zou et al. [2012] built a crack probability map using tensor voting to enhance the connection of
189 crack fragments. After sampling a set of crack seeds from the crack probability map, minimum
190 spanning trees are defined from a graph model of these seeds and recursive tree-edge pruning is
191 applied to identify cracks. Li et al. [2014] classified image pixels into two categories: pixels that

192 belong to cracks and pixels that do not belong to cracks. Then, they applied Otsu's segmentation
193 method to separate the foreground from the background. The images containing cracks are
194 afterwards classified to distinguish between linear and alligator cracks using binary trees and
195 back propagation neural networks. Varadharajan et al. [2014] also adopted machine learning
196 approaches. Considering images, which can contain cars, traffic signs and buildings, they
197 segmented the ground plane out from the rest of the image and calculated feature descriptors
198 based on the color and texture of the pixels. Using data annotated by humans, they trained a
199 support vector machine capable of classifying the images. Moussa and Hussain [2011] used
200 machine learning, namely support vector machines, and applied graph cut segmentation to
201 segment an image into crack and background pixels. They extracted seven features from a binary
202 vector created after segmentation. The features were used to classify the crack type as transverse
203 cracking, longitudinal cracking, block cracking, or alligator cracking. In addition, they also
204 proposed an approach to calculate the extent and severity of the crack. An algorithm based on the
205 Gabor filter was proposed by Salman et al. [2013]. After convolution with the filter, the real
206 component of the resulting image was thresholded and a binary image was obtained. Huang and
207 Xu [2006] divided the image into cells for classification purposes. Each cell was classified as a
208 crack or non-crack cell depending on its contrast.

209 Compared to cracks, approaches towards patch detection in pavement images are fewer in
210 number. Radopoulou and Brilakis [2014] applied morphological operations to segment out patch
211 regions. Texture information was also used to generate feature vectors of both intact and patch
212 regions. Cafiso et al. [2006] applied a clustering method to analyze pavement images with
213 respect to patches.

214 Koch and Brilakis [2011] proposed a method for pothole detection in asphalt pavement images.
215 They first used histogram shape-based thresholding to segment an image into defect and non-
216 defect regions. The potential pothole shape was approximated based on morphological thinning
217 and elliptic regression. An improved method capable of tracking potholes in subsequent frames
218 is presented in [Koch et al. 2013]. Buza et al. [2013] also employed image processing and
219 spectral clustering for identification and rough estimation of potholes. In addition, they estimated
220 the surface of the potholes. Yu and Salari [2011] introduced an approach for pothole detection
221 and severity management based on laser imaging. The proposed algorithm also analyses the
222 severity of the pothole.

223 Methods exist capable of identifying pavement distress in general. Some of them, namely multi-
224 resolution texture analysis techniques using wavelet, ridgelet, and curvelet-based texture
225 descriptors, were compared in [Nejad and Zakeri 2001]. The curvelet-based method
226 outperformed all other multi-resolution techniques for pothole distress, while the ridgelet-based
227 yielded the most accurate results for cracks.

228 Most of the presented methods were developed solely for a specific type of distress. Since the
229 idea of this work is to roughly assess the condition of the pavement surface, methods capable of
230 detecting all types of distress need to be investigated. Thereby, it is not important whether the
231 methods distinguish between the different distress types, but rather if they are suitable for
232 parallel implementation. In order to enable real-time distress detection, we considered only
233 methods which achieved good results for all types of distress and do not require many
234 computational steps that depend on each other.

235 **Wavelet transform for pavement distress detection**

236 In this work, we chose a method based on the wavelet transform for pavement distress detection
237 and evaluation as it fulfills the requirements mentioned above. The method was proposed by
238 Zhou et al. [2006] and tested on 81 images. According to the developers of the method, it
239 achieved 100% reliability for these 81 images. Initially applied for signal processing, the wavelet
240 transform is used to decompose an image into a set of different-frequency components. Based on
241 the frequency, the components are arranged in groups called subbands. The subband components
242 are calculated by applying low pass (L) and high pass (H) digital filters to the image. (The
243 original image can be reconstructed from the wavelet components.) After one pass of the filters,
244 the image is decomposed into four subbands: three detail subbands (HL, LH, HH), and one
245 approximation subband (LL), whereby each subband has a width of $\frac{1}{2}$ of the original image
246 width and a height of $\frac{1}{2}$ of the original image height. The detail subbands contain detail
247 components with different orientation. HL contains the horizontal, LH the vertical, and HH the
248 diagonal components. An example of an image before application of the wavelet-transform is
249 presented in Figure 1. The horizontal details of the crack image are represented in the horizontal
250 subband HL. The approximation subband is further decomposed into four subbands. In this way,
251 different levels of decomposition can be achieved. In Figure 2, the 3-level wavelet transform is
252 presented. The LL_3 subband contains approximation coefficients and is most similar to the
253 original image before applying the wavelet transform.

254 Several wavelet families, i.e. sequences of functions that are performed to transform an image
255 into the wavelet domain, exist. The most commonly used are the Haar wavelet [Haar 1910] and
256 the Daubechies wavelet [Daubechies 1990]. The Haar wavelet is highly suitable for parallel (or
257 GPU) implementation. Hence, it was chosen for the real-time detection of pavement distress in
258 this work. The Haar transform is based on a technique called *averaging and differencing*

259 [Mulcahy] which only makes use of the simple mathematical operations addition, subtraction
260 and division by two. First, the average sum and the average difference of each pair of neighbor
261 elements in a row of the image are calculated. The sum is stored as a coefficient in the L
262 subband, while the difference is stored in the H subband. This step is performed for all rows of
263 the image. Afterwards, the same step is performed column-wise for all vertical neighbors in the
264 image. The horizontal and vertical step can be combined and executed at once, as shown in
265 Figure 3, where A, B, C, and D denote pixels and the corresponding wavelet coefficients are
266 highlighted in the transformed “image” on the right.

267 When applying the wavelet transform on pavement images, Zhou et al. observed that a
268 homogeneous background is transformed into the approximation subband, while distress is
269 represented in the detail subbands. Considering the latter observation, Zhou et al also developed
270 three statistical criteria for distress detection: standard deviation of wavelet coefficients (STD),
271 high-frequency energy percentage (HFEP), and high-amplitude wavelet coefficient percentage
272 (HAWCP). STD and HAWCP correctly detected all the distresses in the images. However, 2.6%
273 of the images which actually do not contain distress were incorrectly isolated by STD as distress
274 images, while HAWCP did not isolate any image wrongly. Hence, HAWCP is used in the work
275 presented in this paper.

276 HAWCP is calculated only at the first level of the wavelet transform, which results in a reduced
277 number of required wavelet transform operations. HAWCP represents a measure of the number
278 of wavelet coefficients in the detail subbands that are larger than a threshold used as an index for
279 pavement distress. To calculate HAWCP, first the wavelet modulus M is obtained as

$$M(p, q) = [HL^2(p, q) + LH^2(p, q) + HH^2(p, q)]^{\frac{1}{2}} \quad (2)$$

280 where (p, q) is the position of the coefficient in the corresponding subbands.

281 Then, the modulus is binarized according to Equation (3):

282

$$D(p, q) = \begin{cases} 1 & \text{if } M(p, q) \geq C_{th} \\ 0 & \text{if } M(p, q) < C_{th} \end{cases} \quad (3)$$

283 where D is the binarized modulus and C_{th} is a threshold value estimated by wavelet thresholding.

284 Finally, HAWCP is calculated as

$$HAWCP = \sum_{p=0}^{W/2} \sum_{q=0}^{H/2} D(p, q) / \left(\frac{W}{2} \frac{H}{2} \right) \quad (4)$$

285 where W and H represent the width and height of the image, respectively.

286 The HAWCP value ranges between 0 and 1 (or 0% and 100 %), where a value near 0 indicates a
287 good pavement surface, and high HAWCP values represent pavement distress.

288

289 **GRAPHICS PROCESSING UNITS**

290 During the last few years, GPUs have emerged as powerful computational hardware available at
291 low prices [Owens et al. 2005]. The utilization of GPUs for general-purpose computing
292 (GPGPU) has gained interest among developers of non-graphical applications. Often combined
293 with a CPU, GPUs are used to accelerate scientific, analytics, engineering, consumer or
294 enterprise applications [Nvidia Corporation 2015]. While CPUs are remarkably suitable for
295 control-intensive applications, such as searching or sorting, due to branch predictions, data-
296 intensive applications like image processing are appropriate for GPUs [Gaster et al. 2013].

297 The most common GPU programming frameworks are the Compute Unified Device Architecture
298 (CUDA) and the Open Computing Language (OpenCL). CUDA was developed by Nvidia and
299 supports only Nvidia devices, while OpenCL can be executed on diverse platforms produced by
300 different vendors, such as AMD, Intel, Nvidia, and others. OpenCL was developed by the

301 Khronos Consortium in 2008 and is often referred to as the *industry standard for heterogeneous*
302 *computing* [Khronos OpenCL Working Group 2013].

303 In OpenCL, a single host is defined that is responsible for the coordination of code execution on
304 one or more devices [Gaster et al. 2013]. The host also interacts with the environment external to
305 the OpenCL program, for example with the user. The device can be a CPU, a GPU, a digital
306 signal processor (DSP), or another processor supported by OpenCL. Streams of instructions
307 called *kernels* (not to be confused with convolution kernels) are executed on the device. A
308 portion of the code, called *host program*, runs on the host and defines kernels or collections of
309 kernels that are submitted to the devices by issuing a command for execution. An instance of the
310 kernel is executed for each point of an index space in parallel.

311 The kernels operate on the values of memory objects. Five distinct memory regions are defined
312 in OpenCL, namely host memory, global memory, constant memory, local memory and private
313 memory. They are used for different purposes. For example, global memory can be accessed by
314 all kernel instances in contrast to local and private memory.

315 Stürmer et al. [2012] and Sharma and Vydyanathan [2010] proposed GPU implementations of
316 the wavelet transform. However, in both cases the wavelet coefficients of the wavelet transform
317 are calculated at all decomposition levels. The method proposed by Zhou requires only the
318 values of the first wavelet decomposition level. Therefore, the computational overhead due to
319 unnecessary further decomposition should be eliminated for the purpose of real-time pavement
320 distress detection. Moreover, the computation of the HAWCP criterion could also be carried out
321 on GPU, as shown in this paper.

322

323 **PROBLEM STATEMENT AND OBJECTIVES**

324 Despite of the advances in vision-based pavement distress detection, gaps still exist in research
325 which we try to address in this paper. First, pavement assessment is usually carried out either
326 manually or by using special dedicated vehicles. Second, the data acquired for pavement distress
327 detection is mostly processed offline, which results in a huge amount of data being stored
328 persistently.

329 To address the aforementioned problems, the following two research questions have to be
330 answered:

- 331 1. How can we automate the pavement distress detection process, while using inexpensive
332 vehicles?
- 333 2. How can we reduce the amount of data saved for offline processing?

334

335 **APPROACH**

336 This paper addresses the issues described previously by presenting an approach which is founded
337 on common vehicles. Instead of using dedicated vehicles, the idea pursued hereby is to use
338 vehicles which drive daily on the roads, such as buses and taxis. Nowadays, such vehicles are
339 equipped with built-in cameras, for example backup cameras, which can be used not only to
340 support the driver while parking, but also for other tasks, particularly in this case for road distress
341 detection.

342 In order to address the second research question, we propose online processing of pavement
343 images in real-time. With the aim of reducing storage consumption, only images which contain
344 distress will be stored, while images of good pavement surface will be discarded directly after
345 they have been taken and processed. However, to enable real-time pavement distress detection
346 while driving, either methods which do not require a long execution time need to be developed

347 or existing methods should be enhanced or implemented for faster architectures. In this work,
348 GPUs are utilized to enhance the performance of existing pavement image pre-processing and
349 analysis methods. As a result, real-time pavement distress detection is possible.

350 The approach proposed here is presented in Figure 4. To remove the noise, the images are first
351 convolved with a median filter. Second, the top-hat transform is applied to produce a more
352 uniform background. The third step in the pipeline is transforming the image into the wavelet
353 domain. Then, the high-amplitude wavelet coefficient percentage is calculated. HAWCP is used
354 as a descriptor for classification. Based on a previously generated classification model, the image
355 is classified as a *good pavement image* or an *image containing distress*. This classification model
356 is created in advance using existing machine learning algorithms. To this end, training images
357 are acquired and manually labeled and a data mining tool is used to induce general rules that map
358 pavement images to the two aforementioned categories. Currently, all steps, except
359 classification, are implemented on GPU. An example of a processed image is presented in Figure
360 5.

361

362 **IMPLEMENTATION**

363 An overview of the implementation is depicted in Figure 6. First, the input image data that is
364 initially located only on the host (CPU) needs to be transferred to the device (GPU). For this
365 purpose, the image data is copied into a global memory buffer on the device. A kernel performs
366 median filtering on this data and the result (denoised image) is saved in another memory buffer
367 on the device. Then, a top-hat transform kernel is executed. The latter is used to correct the
368 background of the image and the result is also saved in a buffer on the device. The wavelet
369 transform and the calculation of the HAWCP descriptor are combined in one pavement analysis

370 kernel. The wavelet coefficients are stored in local memory to achieve better performance. The
371 HAWCP descriptor value is saved in global memory and, at the end, transferred to the host. In
372 the current implementation, this value is submitted to a third-party learning machine called
373 WEKA [Witten et al. 2011] and the image is classified based on a classification model generated
374 by the learning machine with the help of the HAWCP values of training images.

375 **Median Filter**

376 There exist several implementations of the median filter on GPUs [Banger and Bhattacharyya
377 2013, Intel Corporation 2012]. Both implementations provide very good results in terms of
378 performance enhancement. Since an Intel GPU is used for testing in this work, we adopted the
379 implementation proposed by Intel. It uses partial bitonic sorting to perform median filtering.

380 **Top-hat transform**

381 **Naïve implementation**

382 The top-hat transform is performed by subtracting the opening of an image from the input. The
383 opening is obtained by dilating the eroded image. Since there are no global synchronization
384 barriers among different workgroups in OpenCL, at least two kernels are required for the GPU
385 implementation of the top-hat transform. To guarantee that the erosion is completed for all pixels
386 in the image, it is defined in its own kernel. After the kernel had been executed, a dilation kernel
387 can be started. The last operation in the top-hat transform (i.e. the subtraction of the opening
388 from the original image) can also be performed in the dilation kernel. The erosion and dilation
389 kernels are implemented in a manner similar to the median filter. However, instead of computing
390 the median value of the neighborhood, the minimal and maximal value are taken. This
391 implementation is presented in Figure 6.

392 **Separable filter implementation**

393 Two-dimensional convolution operations can, in some cases, be separated into two one-
394 dimensional filters, namely a horizontal and a vertical filter. The horizontal filter is first applied
395 to the image row by row. Then, the vertical filter is applied column-wise to the result of the
396 horizontal convolution. The separable convolution is associative, so the one-dimensional filters
397 can be applied in reverse order. Separating the single 2D convolution into two 1D convolutions
398 usually results in reduced execution time even on the CPU when the convolution is executed
399 sequentially. This performance improvement can be explained if we look at Equations 5 and 6.
400 For example, for a rectangular image convolution kernel, the 2D convolution requires a total of

$$(K*L)*(M*N) \quad (5)$$

401 pixel accesses, where K and L denote the width and height of the convolutional kernel,
402 respectively, and M and N represent the width and height of the image, respectively.

403 When the 1D horizontal convolution is performed, the number of pixel accesses is only

$$K*(M*N) \quad (6)$$

404 for the 1D vertical convolution it is

$$L*(M*N) \quad (7)$$

405 If we execute these convolutions consecutively, we obtain

$$(K + L)*(M*N) \quad (8)$$

406 pixel accesses.

407 Theoretically, this leads to an improvement factor of

$$K*L/(K+L) \quad (9)$$

408 Since the top-hat transform is based on erosion and dilation, it can be implemented as a
409 combination of consecutive horizontal and vertical filters. An overview of the improved
410 implementation is presented in Figure 7, in analogy to Figure 6.

411 Still, the number of sorting/search operations required to find the minimum or maximum element
412 in the one-dimensional filters is also lower than in case of the two-dimensional convolution. This
413 allows for improvement factors even greater than expressed in Equation 9.

414 **Wavelet transform and HAWCP**

415 The wavelet kernel is executed for each group of four adjacent pixels in the image. For example,
416 if we consider Figure 3, the same computations would be performed in parallel for the groups (A,
417 B, E, F), (C, D, G, H), (I, J, M, N), and (K, L, O, P). The detail coefficients (i.e. LH, HH, and
418 HH) are calculated using addition and subtraction. Then, the modulus at the certain position is
419 calculated according to Equation 2. The value of the modulus is compared to the threshold value
420 and if it exceeds it, the HAWCP value is incremented. Atomic operations are used to increment
421 the HAWCP value. A schematic of the implementation is presented in Figure 8.

422

423 **PERFORMANCE EVALUATION**

424 To evaluate the computational speed-up achieved by implementing the median filter, the top-hat
425 transform and the wavelet transform on GPU, performance tests were carried out. The objective
426 pursued was to measure the time required to execute the different pavement distress detection
427 steps on different architectures and to compare them. In particular, a sequential version of the
428 methods executed on a CPU, an OpenCL parallel version executed on the same CPU, the
429 OpenCL version executed on an integrated GPU, and the OpenCL implementation executed on a
430 discrete GPU were compared. In case of the OpenCL implementations of the median filter and
431 the top-hat transform, both the times for the 2D and for the separable convolution were
432 measured. As recommended in [Intel Corporation 2013], the same set of operations was wrapped
433 in the sequential and OpenCL implementations in order to make sure that the observed code

434 patterns are as similar as possible. Moreover, to guarantee accurate results, the methods were
435 invoked on 1000 images and the average value of all the 1000 executions was taken for
436 performance evaluation.

437 Profiling events were used to measure the OpenCL execution time. The data transfer time (i.e.
438 the time required to write data to the device or read data from the device) and the kernel
439 execution time were tracked separately due to the following two reasons. First, both the data
440 transfer time and the kernel execution time are highly dependent on the hardware. The time
441 needed to transfer data between a host and an integrated GPU is usually much lower than the
442 time required to transfer the same data between the host and a discrete GPU. Second, if we
443 consider Figure 4, it is obvious that only the input image data and the HAWCP results need to be
444 transferred between the host and the device. All other intermediate results are saved in memory
445 buffers on the device. Thus, only the kernel execution times are relevant for the overall
446 performance evaluation of the real-time pavement assessment approach.

447 The OpenCL initialization time, i.e. the time required to create a program, a context, command
448 queues, the kernels, and set the kernel arguments, is also not considered, because these
449 initialization steps are executed only once at application startup and are not repeated for each
450 frame or image that has to be processed.

451 The following hardware was used for the performance evaluation tests: a 2.10 GHz Intel Core i7-
452 4600 CPU, an integrated Intel HD Graphics 4400 GPU, and a dedicated Nvidia Tesla C2070
453 GPU. In addition, the approach was tested on images of different sizes, namely 256x256,
454 512x512, 1024x1024, and 2048x2048 pixels, because universal rear view cameras have different
455 resolutions. Resolutions of 500x500 pixels are common nowadays, but vehicle manufacturers
456 have already developed rear view cameras with 1,300,000 pixels [Nissan Motor Corporation

2014]. The speed-up achieved by implementing the approach on GPUs was computed, . This speed-up is defined as shown in Equation 10.

$$\text{Speed-up} = \text{Sequential C++ time} / \text{Best OpenCL time} \quad (10)$$

Data transfer

The data transfer time differs depending on what kind of device is used. The time required to transfer the image data to the integrated Intel GPU and the dedicated Nvidia GPU are illustrated in Figure 9. The transfer to the discrete GPU is significantly slower than the transfer to the integrated GPU for large images.

The difference between the times required to transfer the HAWCP value of a single image is not so considerable, because only one value needs to be transferred.

Median Filter

In our work, we used a median filter with a square structuring element of a size 3x3. The execution times in milliseconds are shown in Table 1.

Top-hat transform

The top-hat transform was tested with a structuring element of a size 10x10. The performance evaluation results are presented in Table 2 in milliseconds. For all image sizes, the separable implementation executed on the dedicated Nvidia GPU was the fastest one. In contrast to the median filter, a considerable performance improvement was achieved by using separate horizontal and vertical filters.

Wavelet transform and HAWCP

The wavelet transform execution time, including the time required to calculate the HAWCP descriptor, is presented in Figure 10. The operations were executed approximately 109 times faster on the Nvidia GPU compared to the sequential CPU. As shown in Figure 10, the

479 calculation takes more than 8 milliseconds when executed sequentially, which makes it
480 unsuitable for real-time processing of videos taken at high speeds. In contrast, all GPU
481 implementations require less than one millisecond, so there is sufficient time for pre-processing
482 operations.

483 **Overall enhancement**

484 To compare the execution of the different implementations on the CPU and the two GPUs, the
485 total execution times were calculated. As can be seen in Figure 11, in case of an image size of
486 2048x2048, the data transfer time is approximately 0.72 milliseconds, which is about 50% of the
487 total execution time. However, the Nvidia execution still significantly outperforms all other
488 implementations.

489 The total execution times for all image sizes are shown in Table 3. The speed-up calculated
490 according to Equation 10 is also presented. In case of the Nvidia GPU, the total execution time is
491 below 1.5 milliseconds. Theoretically, this allows processing more than 650 images per second.

492

493 **CASE STUDY**

494 To validate the approach, a case study was conducted. A road segment located in Bochum,
495 Germany, was chosen for validation due to the presence of parts of the road with a good
496 pavement surface and parts with pavement distress. The length of the road segment is
497 approximately 24 kilometers. The road segment includes different types of pavement. An
498 example of two different road surface textures is presented in Figure 14. To collect video data, a
499 Basler acA2040-180kc camera was mounted on a rear door back carrier. As a variety of rear
500 view cameras and vehicles exist, there are different ways and positions to mount the cameras.
501 While license mounted cameras are easy to install on the existing license plate, surface mounted

502 cameras are commonly mounted higher and would be a better choice for larger vehicles
503 [Rearview Camera Reviews]. The setup of the camera in this case study tries to imitate state-of-
504 the-art rear view camera setups as far as possible. The position and orientation of the camera are
505 presented in Figure 12. The camera is capable of acquisition with a frame rate of up to 180
506 frames per second, which are currently not achievable by rear view cameras. However, we
507 anticipate that in the near term vehicle manufacturers will use rear view cameras with even
508 higher frame rates. The pitch angle of the camera is approximately -70 degrees, which is almost
509 perpendicular to the road surface. The camera is placed at a height of 1.16 m above the road
510 surface.

511 In order to enable the validation of the applied methods, all images were saved. Under real
512 conditions, the images on which no distress was identified would be discarded and only images
513 on which pavement defects were detected would be saved. To test the classification, 1549 images
514 were selected. Both images of a good pavement surface as well as images containing cracks,
515 potholes and patches were considered (Figure 13).

516 The images were manually labeled and ten-fold cross validation was performed in order to get a
517 reliable error estimate. For this purpose, the data was split into ten approximately equal
518 partitions. Each of these partitions was used for testing once, while the remaining data was used
519 for training. Three algorithms were used for classification, namely the C4.5 [Quinlan 1993]
520 algorithm, Multilayer Perceptron [Witten 2011], and Rotation Forest [Rodriguez 2006]. The
521 results of the classification are presented in Table 4. The confusion matrix for the test images
522 classified with the Rotation Forest algorithm is presented in Table 5. The time required to test the
523 tree models on the training split was 0.02 seconds for C4.5, 0.66 seconds for Multilayer
524 Perceptron, and 0.14 seconds for Rotation Forest.

525 The 5% of the images that were classified incorrectly are 77 images in total. Out of them, 15
526 images without distress were classified as images containing distress (false positives). In Figure
527 14, an example of a correctly classified intact pavement image (left) and an intact pavement
528 image that was incorrectly classified as image containing distress (right) is presented.
529 Nevertheless, this is still a promising classification result, because the objective of the rough
530 distress detection stage described in this paper is to identify potential distress locations. In a
531 further step, these potential locations will be assessed in detail by more comprehensive
532 algorithms.

533 Vice versa, the other 62 images which actually contain distress were classified as distress free
534 images (false negatives), mainly because of the different types of road surfaces considered in the
535 case study. Consequently, the locations these images were acquired at would not be taken into
536 account within the fine analysis. In order to counteract such errors, the methodology presented
537 here will be extended by incorporating textural features.

538

539 **CONCLUSION**

540 Pavement condition assessment is a key component of pavement maintenance programs.
541 Currently, pavement distress is detected during observations by trained personnel and reported
542 manually. State-of-the-art automated methods for pavement distress detection utilize special
543 vehicles equipped with sensors and cameras and try to compensate the limitations of the manual
544 distress detection process. However, the need to reduce the amount of required memory to
545 capture all pavement related data is still present.

546 With the aim of enabling real-time pavement image processing and, thus, reducing the amount of
547 stored data, this paper proposed an approach based on graphics processing units (GPUs).

548 Specifically, GPU implementations of a noise removal, a background correction and a pavement
549 distress detection method were developed. In order to remove noise in the images and correct
550 their non-uniform background, the median filter and the top-hat were used. The wavelet
551 transform was applied in order to calculate a descriptor value for classification purposes. Based
552 on this value, the images were classified as good pavement images or images containing distress.
553 To compare the performance of the GPU implementations against sequential applications and to
554 validate the classification methodology, the approach was tested on 1549 images. The results
555 show that by exploiting the computational power of the GPU it is possible to do pre-processing
556 and analyze pavement images with a resolution of 2040 x 2048 pixels in real time. In addition, it
557 has been demonstrated that the wavelet transform can successfully be applied on pavement
558 images for the purpose of distress detection. Based on the high-amplitude wavelet coefficient
559 percentage descriptor, 95% of the images used for testing were classified correctly by the
560 Rotation Forest algorithm.

561 Yet, some images containing small cracks were incorrectly classified as good pavement images.
562 The approach presented in this paper can be improved by combining multiple descriptors to
563 obtain a more accurate representation of the distress. Future steps include the implementation of
564 other pavement distress detection techniques on the GPU, as well as the employment of Graphics
565 Processing Units for further pre-processing steps, such as the Bayer pattern de-mosaicing.

566

567 **ACKNOWLEDGMENT**

568 The authors gratefully acknowledge the support of this ongoing project by the German Research
569 Foundation (DFG) under grant KO4311/2-1.

570

571 **REFERENCES**

- 572 Ahmed, E. S. A., Elatif, R. E. A., and Alser, Z. T. (2015). Median filter performance based on
573 different window sizes for salt and pepper noise removal in gray and RGB images. *International*
574 *Journal of Signal Processing, Image Processing and Pattern Recognition*, 8 (10), pp. 343-352
- 575 Banger, R. and Bhattacharyya, K. (2013). *OpenCL Programming by Example*. Packt Publishing,
576 Birmingham, UK
- 577 Buske, K. (2013). Schlagloch-Kataster wird auf Vordermann gebracht. *DerWesten*, Online,
578 available at: [http://www.derwesten.de/staedte/bochum/schlagloch-kataster-wird-auf-vordermann-](http://www.derwesten.de/staedte/bochum/schlagloch-kataster-wird-auf-vordermann-gebracht-id8177431.html)
579 [gebracht-id8177431.html](http://www.derwesten.de/staedte/bochum/schlagloch-kataster-wird-auf-vordermann-gebracht-id8177431.html), Accessed on 11.01.2016
- 580 Buza, E., Omanovic, S., and Huseinovic, A. (2013). Pothole detection with image processing and
581 spectral clustering. *2nd International Conference on Information Technology and Computer*
582 *Networks*, Antalya, Turkey
- 583 Cafiso, S., Di Graziano, A., and Battiato, S. (2006). Evaluation of pavement surface distress
584 using digital image collection and analysis. *Seventh International Congress on Advances in Civil*
585 *Engineering*
- 586 Cheng, H.D., and Miyojim, M. (1998). Novel system for automatic pavement distress detection.
587 *Journal of Computing in Civil Engineering*, vol. 12, pp. 145-152
- 588 Daubechies, I. (1990). The Wavelet Transform, Time-Frequency Localization and Signal
589 Analysis. *IEEE Transactions on Information Theory*, vol. 36, no. 5, pp. 961 – 1005
- 590 DStGB (Deutscher Städte- und Gemeindebund). (2014). PKW-Maut für alle Straßen richtiger
591 Ansatz – Beteiligung der Kommunen an den Einnahmen unverzichtbar. Online, available at:
592 <http://www.dstgb.de/dstgb/Home/Pressemeldungen>, accessed on 03.12.2014

593 Fang, C., Zhe, L., and Li, Y. (2014). Images crack detection technology based on improved K-
594 means algorithm. *Journal of Multimedia*, 9 (6), pp. 822-828

595 Gaster, B. R., Howes, L., Kaeli, D.R., Mistry, P., and Schaa, D. (2013). *Heterogeneous*
596 *Computing with OpenCL: Revised OpenCL 1.2 Edition*. Morgan Kaufmann Publishers Inc. San
597 Francisco, CA, USA

598 Georgopoulos, A., Loizos, A., and Flouda, A. (1995). Digital image processing as a tool for
599 pavement distress evaluation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 50 (1),
600 pp. 23-33

601 Gonzalez, R. C., Woods, R. E. (2006). *Digital Image Processing*. Prentice-Hall, Inc. Upper
602 Saddle River, NJ, USA

603 Haar, A. (1910). Zur Theorie der Orthogonalen Funktionensysteme. *Mathematische Annalen*,
604 vol. 69, pp. 948 – 956

605 Huang, Y, and Xu, B. (2006). Automatic inspection of pavement cracking distress, *J. Electron.*
606 *Imag.* 15 (1)

607 Intel Corporation (2012). Intel® SDK for OpenCL* - Median Filter Sample. Document Number:
608 325264-003US, Revision: 1.3

609 Intel Corporation (2013). Intel® SDK for OpenCL* Applications 2013 R2 Optimization Guide.
610 Document Number: 326542-003US

611 Jähne, B. and Haussecker, H. (2000). *Computer Vision and Applications: A Guide for Students*
612 *and Practitioners*. Academic Press

613 Khronos OpenCL Working Group (2013). The OpenCL Specification, Version: 2.0. Document
614 Revision 19

615 Koch, C. and Brilakis, I. K. (2011). Pothole detection in asphalt pavement images. *Advanced*
616 *Engineering Informatics*, vol. 25, no. 3, pp. 507–515.

617 Koch, C., Georgieva, K., Kasireddy, V., Akinci, B., and Fieguth, P. (2015). A review on
618 computer vision based defect detection and condition assessment of concrete and asphalt civil
619 infrastructure. *Advanced Engineering Informatics* 29, pp. 196–210

620 Koch, C., Jog, G. M., and Brilakis, I. (2013). Automated pothole distress assessment using
621 asphalt pavement video data. *Journal of Computing in Civil Engineering* 27(4), pp. 370-378

622 Li, L., Sun, L., Ning, G., and tan, S. (2014). Automatic Pavement Crack Recognition Based on
623 BP Neural Network. *Promet – Traffic&Transportation*, Vol. 26, No. 1, pp. 11-22

624 Lokeshwor, H., Das, L.K., and Sud, S.K. (2013). Method for automated assessment of potholes
625 cracks and patches from road surface video clips. *Procedia - Social and Behavioral Sciences*
626 104, pp. 312–321.

627 Moussa, G., and Hussain, K. (2011), A new technique for automatic detection and parameters
628 estimation of pavement crack. *4th International Multi-Conference on Engineering Technology*
629 *Innovation (IMETI 2011)*

630 Mulcahy, C. Image compression using the Haar wavelet transform. *Spielman Science and Math*
631 *Journal*

632 NCHRP (National Cooperative Highway Research Program). (2004), Automated Pavement
633 Distress Collection Techniques: A Synthesis of Highway Practice

634 Nejad F. M. and Zakeri, H. (2001). A comparison of multi-resolution methods for detection and
635 isolation of pavement distress. *Expert Systems with Applications*, vol. 38, pp. 2857 – 2872

636 NHTSA (National Highway Traffic Safety Administration) (2014). Federal Motor Vehicle
637 Safety Standards; Rear Visibility. Federal Register, the Daily Journal of the United States

638 Government, Online, available at: [https://www.federalregister.gov/articles/2014/04/07/2014-](https://www.federalregister.gov/articles/2014/04/07/2014-07469/federal-motor-vehicle-safety-standards-rear-visibility)
639 [07469/federal-motor-vehicle-safety-standards-rear-visibility](https://www.federalregister.gov/articles/2014/04/07/2014-07469/federal-motor-vehicle-safety-standards-rear-visibility), Accessed on 11.01.2016
640 Nissan Motor Corporation (2014). Nissan Motor develops the “Smart rearview mirror”, which
641 helps provide clear rearward visibility in various conditions. Online, available at:
642 http://www.nissan-global.com/EN/NEWS/2014/_STORY/140228-01-e.html, Accessed on
643 11.01.2016
644 Nvidia Corporation (2015). CUDA ZONE. Online available at:
645 <https://developer.nvidia.com/cuda-zone> Accessed 11.06.2015
646 NYSDOT (New York State Department of Transportation) (2010). Pavement condition
647 assessment. V2.0w
648 Oliveira, H. and Correia, P.L. (2013). Automatic road crack detection and characterization. *IEEE*
649 *Transactions on Intelligent Transportation Systems*, 14 (1)
650 Orr, S. (2015). Officials want the public to report pothole locations. *Evansville Courier&Press*,
651 Online, available at: [http://www.courierpress.com/news/local-news/officials-want-the-public-to-](http://www.courierpress.com/news/local-news/officials-want-the-public-to-report-pothole-locations_36253717)
652 [report-pothole-locations_36253717](http://www.courierpress.com/news/local-news/officials-want-the-public-to-report-pothole-locations_36253717), accessed on 09.06.2015
653 Owens, J. D., Luebke, D., Govindaraju, N., Harris, M., Krüger, J., Lefohn, A. E., and Purcell, T.
654 J. (2005). A survey of general-purpose computation on graphics hardware. *Eurographics, State*
655 *of the Art Reports*, August 2005, pp. 21-51
656 Quinlan, J. R. (1993). C4.5: programs for machine learning. Morgan Kaufmann Publishers
657 Inc. San Francisco, CA, USA
658 Radopoulou, S.C., and Brilakis, I. (2014). Improving patch distress detection using vision
659 tracking on video data. *Proceedings of the 21st International Workshop on Intelligent Computing*
660 *in Engineering*

661 Rearview Camera Reviews. The complete buyers guide for Rear View Cameras. Online,
662 available at: <http://rearviewcamerareviews.com>, Accessed on 11.01.2016

663 Rodriguez, J. J. & Kuncheva, L. I. (2006) Rotation Forest: A New Classifier Ensemble Method.
664 IEEE Transactions on Pattern Analysis and Machine Intelligence 28(10), pp. 1619-1630

665 Salman, M., Mathavan, S., Kamal, K., and Rahman, M. (2013). Pavement crack detection using
666 the Gabor filter. *16th International IEEE Conference on Intelligence Transportation System*
667 *(ITSC 2013)*, pp. 2039–2044

668 Sharma, B., and Vydyanathan, N. (2010). Parallel Discrete Wavelet Transform using the Open
669 Computing Language: a performance and portability study. *2010 IEEE Int. Symp. Parallel and*
670 *Distributed Processing, Workshops and Ph.D. Forum (IPDPSW)*, pp.1 – 8

671 Solomon, C. and Breckon, T. (2010). *Fundamentals of Digital Image Processing: A Practical*
672 *Approach with Examples in Matlab*, Wiley

673 Spielman, F. (2014). Chicago potholes trigger record number of damage claims. *Chicago Sun-*
674 *Times*, Online, available at: <http://chicago.suntimes.com/?p=167606>, Accessed on 09.06.2015

675 Stürmer, M., Köstler, H., and Rude, U. (2012). Fast wavelet transform utilizing a multicore-
676 aware framework. *PARA'10 Proceedings of the 10th international conference on Applied*
677 *Parallel and Scientific Computing*, vol. 2, pp. 313-323, Springer-Verlag Berlin, Heidelberg

678 Subirats, P., Dumoulin, J., Legeay, V., and Barba, D. (2006). Automation of pavement surface
679 crack detection using the continuous wavelet transform. International Conference on Image
680 Processing, Atlanta, USA

681 Tanaka, N. and Uematsu, K. (1998). A crack detection method in road surface images using
682 morphology. *IAPR Workshop on Machine Vision Applications*, Makuhari, Japan

683 Tsai, Y-C., Kaul, V., and Mersereau, R.M., (2010). Critical Assessment of Pavement Distress
684 Segmentation Methods, *J. Transp. Eng.*, 136(1), pp. 11–19.

685 Varadharajan, S., Jose, S., Sharma, K., Wander, L., and Mertz, C. (2014). Vision for Road
686 Inspection. *Proceedings of WACV 2014: IEEE Winter Conference on Applications of Computer*
687 *Vision*

688 Vivekanandreddy Navaneetha, D., Kammar, A. & Sowmyashree.B (2014). Hough transforms to
689 detect and classify road cracks. *International Journal of Engineering Research & Technology*,
690 3(6), pp. 1500 – 1505

691 Witten, I. H., Frank, E., and Hall, M. A. (2011) Data mining: practical machine learning tools
692 and techniques, Elsevier

693 Wu, Q., Merchant, F.A., and Castleman, K. R. (2008). *Microscope Image Processing*, Elsevier

694 Yu, X. and Salari, E. (2011). Pavement pothole detection and severity measurement using laser
695 imaging. *IEEE International Conference on Electro/Information Technology (EIT)*, Mankato,
696 USA

697 Zhou, J., Huang, P.S., and Chiang, F.-P. (2006). Wavelet-based pavement distress detection and
698 evaluation, *Opt. Eng.* 45 (2).

699 Zou, Q., Cao, Y., Li, Q., Mao, Q., Wang, S. (2012). CrackTree: automatic crack detection from
700 pavement images, *Pattern Recog. Lett.* 33 (3), pp. 227–238.

701

702

703

704

705

706 **List of Figures**

707

708 Figure 1: Pavement crack image

709 Figure 2: Three-level wavelet transform of the crack image

710 Figure 3: Calculation of the wavelet coefficients

711 Figure 4: Pavement distress image classification method

712 Figure 5: Image processing pipeline a) original image b) median filtered image c) top-hat

713 transformed image d) HAWCP value

714 Figure 6: An overview of the naïve GPU implementation

715 Figure 7: GPU implementation using one-dimensional filters

716 Figure 8: A schematic of the implementation of the wavelet transform and HAWCP calculation

717 on GPU

718 Figure 9: Data transfer times on different architectures

719 Figure 10: Wavelet transform and HAWCP execution time

720 Figure 11: Total execution time on the Nvidia GPU

721 Figure 12: Data acquisition vehicle

722 Figure 13: Examples of images acquired for training and testing

723 Figure 14: Correctly (left) and incorrectly (right) classified images of intact pavement surface

724

725

Table 1: Median filter execution times in milliseconds

	256x256	512x512	1024x1024	2048x2048
Sequential	14.3	57.936	230.758	889.876
OpenCL Intel CPU	0.108943	0.327316	1.22963	4.77966
OpenCL Intel GPU	0.013582	0.049675	0.193708	0.769058
Nvidia GPU	0.002747	0.010321	0.0399	0.156663

726

727

Table 2: Top-hat transform execution times in milliseconds

	256x256	512x512	1024x1024	2048x2048
Sequential	203.007	765.11	2980.79	11611
OpenCL Intel CPU Naïve	1.13034	5.03241	18.0581	76.2757
OpenCL Intel CPU Separable	0.431628	4.80577	15.9215	58.1489
OpenCL Intel GPU Naïve	0.584406	2.31147	8.23475	25.4106
OpenCL Intel GPU Separable	0.0851977	0.314928	1.25112	5.08258
Nvidia GPU Naïve	0.025724	0.0961443	0.370388	1.4927
Nvidia GPU Separable	0.00853265	0.0301136	0.11383	0.43868

728

729

Table 3: Total execution times of all implementations

	256x256	512x512	1024x1024	2048x2048
Sequential	217.407	823.436	3213.158	12509.4564
OpenCL Intel CPU	1.29138047	5.51308102	19.8314412	83.2016187
OpenCL Intel CPU Separable	0.57764077	5.22357002	17.2738532	63.4995187
OpenCL Intel GPU	0.6230764	2.44574345	8.64068146	26.7954523

OpenCL Intel GPU Separable	0.1264993	0.45310435	1.66696046	6.49687731
Nvidia GPU	0.03927566	0.14796738	0.58636053	2.4431657
Nvidia GPU Separable	0.02226623	0.08221098	0.33134483	1.3884667
Speed-up	9763.97715	10016.1317	9697.32345	9009.54728

730

731

Table 4: Results of the classification of the pavement images

Algorithm	Correctly classified in %	Precision	Recall
C4.5	95	0.949	0.950
Multilayer Perceptron	87	0.880	0.872
Rotation Forest	95	0.950	0.950


732

Table 5: Confusion matrix for the test images classified with the Rotation Forest algorithm

Image containing distress	Good pavement image	Classification outcome
		Actual condition
306	62	Image containing distress
15	1166	Good pavement image

734








A	B	C	D
E	F	G	H
I	J	K	L
M	N	O	P





$\frac{A+B+E+F}{4}$	$\frac{C+D+G+H}{4}$	$\frac{A-B+E-F}{4}$	$\frac{C-D+G-H}{4}$
$\frac{I+J+M+N}{4}$	$\frac{K+L+O+P}{4}$	$\frac{I-J+M-N}{4}$	$\frac{K-L+O-P}{4}$
$\frac{A+B-E-F}{4}$	$\frac{C+D-G-H}{4}$	$\frac{A-B-E+F}{4}$	$\frac{C-D-G+H}{4}$
$\frac{I+J-M-N}{4}$	$\frac{K+L-O-P}{4}$	$\frac{I-J-M+N}{4}$	$\frac{K-L-O+P}{4}$






















Execution time on the Nvidia GPU









List of Figures

Figure 1: Pavement crack image

Figure 2: Three-level wavelet transform of the crack image

Figure 3: Calculation of the wavelet coefficients

Figure 4: Pavement distress image classification method

Figure 5: Image processing pipeline a) original image b) median filtered image c) top-hat transformed image d) HAWCP value

Figure 6: An overview of the naïve GPU implementation

Figure 7: GPU implementation using one-dimensional filters

Figure 8: A schematic of the implementation of the wavelet transform and HAWCP calculation on GPU

Figure 9: Data transfer times on different architectures

Figure 10: Wavelet transform and HAWCP execution time

Figure 11: Total execution time on the Nvidia GPU

Figure 12: Data acquisition vehicle

Figure 13: Examples of images acquired for training and testing

Figure 14: Correctly (left) and incorrectly (right) classified images of intact pavement surface