March 28, 2012

Revision Report for manuscript 1390

Dear Reviewers, Dear Editors,

thank you very much for the good reviews regarding our manuscript "Temporal Morphing for Adaptive SPH" indicating that the paper may be suitable for a fast track to Computer Graphics Forum. Please find below a detailed list of the changes we have made in response to each of the reviewer's points. A revised version of our paper and the old submission to Eurographics 2012 have been uploaded online.

We hope that our modifications now meet satisfactorily all points raised by the reviewers. We want to thank you and the reviewers very much again for their helpful comments.

Overall Review Summary from the Editors:

The reviewers were mostly positive about this paper, everyone agreed that the approach taken was novel and interesting, and deserves publication. They did, however, feel that it needs some additional work to be fully ready. The reviewer's fet that the paper's main weakness is the evaluation of the proposed method. Specifically, they would like to see the following:

(A) Give more real-world examples, more evaluation, comparison to other techniques, specifically [APKG07] and [SG11], along with a better quantitative description of speedups.

Changes: In accordance with the reviewers suggestions we have completely exchanged all prior scenes by the following real-world examples:

- a flooding of a valley, similar to [APKG07], in order to show the applicability of our temporal blending for a pre-defined high-resolution region. Additionally, we demonstrate the effect of a salt-diffusion in combination with a convective flux in between complex boundaries. We later utilize the scene in order to evaluate the scaling of the performance with respect to the number of particles for a SPH and a PCISPH simulation.

- a scene with two colliding fluid fronts between some pillars. We clarify the effect of damping which always accompany with larger particles. Additionally, we compare our method to a non-adaptive PCISPH and non-continuous resampling [APKG07]. All comparisons are made with respect to the introduced density error, the integration step size and the final rendering giving a detailed insight into all related techniques. We also present the limitation for an adaptive re-definition of high-resolution regions in case of strongly different particle sizes while giving reason for our future work. We agree that adaptivity due to internal activity is an important aspect which requires further research.

- We show some limitations concerning turbulent flow and/or due to a large diffusion layer by mixing coffee and cream in an Utah-Teapot. Similar to [APKG07], we demonstrate the adaptivity of our method to complex boundaries.

We provide a detailed performance table for the applied simulation steps and the overall computation times until convergence. Each timing is compared to a non-adaptive simulation. We also provide overall simulation times for a non-continuous adaptive PCISPH. We have added information about the memory consumption as well. Please note that the parameters which we are using are provided within in each Sec.. For example, $T_{\min}$ and $T_{\max}$ are defined in Sec.4.4.

Please note that we now strongly compare our method to [APKG07] and [SP09] as they are directly related. However, even if a comparison between a two-scale

SPH, as presented by [SG11], and single-scale SPH, as presented by [APKG07], would be great to have in the future, it is not directly necessary, as beside of performance we mainly target a consistent transition between resolution levels while keeping the sampling error to a minimum. As such, we address the issues regarding [SG11] in our related work section while we a compare our method to [SP09] and [APKG07] as they are much more important.

(B) Better explain what is happening in the video, and explain what different sized particles represent.

Changes: We apologize to the reviewers that the attached video was not as descriptive. Our new examples much better explain our temporal blending and better compare our techniques with the related work including non-adaptive as well as a non-continuous re-sampling.

(C) Improve the discussion of adaptive particle selection, particularly with respect to internal activity.

Changes: We changed the focus of our paper towards blend-sets. As such, the adaptive particle selection has been moved to the beginnig of the paper, section 3.2 and is now further clarified by additional screenshots presenting a teapot-example with internal activity including a diffusion layer. We also added more related work regarding the identification of high resolution regions.

(D) Clarify the relationship of the method with respect to SPH smoothing.

Changes: Firstly, we have removed the confusing and misleading remark, that temporal blending "extends the idea of smoothing in SPH to the time domain." Secondly, we thoroughly improved the description of our temporal blending while keeping the focus on blend-sets. Last but not least we changed the title to "Temporal Blending for Adaptive SPH" as it better describes what is going on.

(E) Better explain the relationship of the simulation step size with the error in the pressure term.

Changes: This is the part which we have changed the most. Thanks to the reviewers we figured that in it's last state it was not understandable. We have clearly improved and simplified the algorithm. The relation to the integration time-step is much more clear now. Instead of just blending with a constant linear step size we recompute the new blend-increment in each step with respect to a re-sampling error and with respect to the current integration step size.

(F) Discuss damping artifacts due to non-uniform sampling.

Changes: We have included a discussion and a visual impression of damping in the result section, Please see Fig.8 and Fig.10 and the accompanying video.

(G) Improve the exposition of the paper according to the reviewers comments, paying particular attention to Reviewer 5's comments about Equations 7, 8 and 11, and Section 6.

Changes: Please see comments of Reviewer 5.

Reviewer 1:

Comment 1: "The paper is mostly well-written, where it suffers most is in the lack of evaluation given to the proposed framework." "I would have liked to see more examples that show where other methods fail, but this one does not." "The method seems to work well, but a few more real-world examples wouldn't hurt, especially if those are made in conjunction with comparisons to existing techniques."

Answer: Thank you very much Reviewer 1. Our three new examples are much more helpful to describe the concept of our temporal blending and to understand the dynamics as shown in the video. You now will find a lot more examples in the paper along with comparisons to other techniques. E.g. in the results section of the paper. Please see our answer to comment A in the summary.

Comment 3: "...The video doesn't really accomplish this, and Figures 8 and 9, while useful, don't really lend much in this respect."

Answer: Hopefully, our current video fills this gap. Please see also our answer to comment B in the summary.

Reviewer 2:

Comment 4: "Section 7: (feature surface than S) - please add "a certain distance S" as it is unclear what is "S" at this point"

Answer: We have completely restuctured and rewritten those parts of the paper. We explain high-resolution areas in Sec.3.2 together with references to existing related work. Please see also our answer to comment C in the summary.

Comment 5: "Section 8.1: Last sentence in this section is very unclear. Please rewrite it."

Answer: According to your comment, we have completely rewritten our result section. Please see our answer to comment A in the summary.

Comment 6: "The paper states that the proposed temporal morphing approach extends the smoothing concept of SPH. I think that this is statement is not quite correct and might confuse the reader."

Answer: Many thanks. We have removed this misleading remark form our paper. Please see our answer to comment D in the summary.

Comment 7: "The paper states that a scheme for estimating the simulation and blending step size based on predicted error in pressure term is proposed. While the choice of blending step is explained in the paper, the relation of the choice of the simulation step to the introduced error in pressure term is unclear. The simulation step is related to the commonly used CFL condition and is very similar to the one presented in [DC99] (although they bound the simulation step using acceleration and not the forces, but these two quantities are strongly interconnected)."

Answer: Thank you very much for this comment. Due to your hints we figured that this section was not understandable at all. See Sec.4.4 in our paper. Please see also our answer to comment A in the summary.

Comment 8: "The paper would benefit from additional examples and comparison with the non-adaptive methods and the adaptive methods which use the immediate particle replacement in terms of the differences in actual behavior of the liquid."

Answer: Thank you Reviewer 2, the examples which we now provide are way more meaningful. Please see our answer to comment A in the summary as well.

Reviewer 3:

Comment 9: "The proposed method is fully GPU based and computes up to 50 time steps per second for half a million particle." "The timings of 50 updates per seconds for 300k particles is really impressive. Therefore, I wonder why only few frames of one scene are provided in the accompanying video. With this performance, it should be easy to set up more scenes that show different aspects which clearly reflect the benefit and the functionality of the approach."

Answer: We apologize for this mistake. Our paper has been in a rather unfinished state. Clearly, as shown in the Table 1 with 300k particles our temporal blending is able to simulate with a speed of 50ms per time-step on the GPU. Again, apologies for this mistake.

Comment 10:"However, the capability of the method can not be assessed due to the rather poor evaluation." "The evaluation of the method is poor. The provided results are not demonstrating the benefit and quality of the method." "The main motivation of adaptively sampling particles is to improve the performance. Consequently, the perfomance benefit should be carefully evaluated. Unfortunately, timings are only given for one time-step. A detailed analysis of how much time each sub-step takes would be helpful. A comparison of the computation time is only given for the CBD dam scene without stating the minimum number of particles for adaptive SPH." "Furthermore, the problem of damping artifacts resulting from coarser resolution as described in [SG11] is not addressed by the authors." "In total, the approach is only applied to two scenarios, namely the corner-breaking dam (CBD) and a bunny scene." "In my opinion the bunny scene, at least as presented, is not an optimal choice for demonstrating the different aspects of the contribution." "Implementing adaptive SPH is quite challenging due to many error sources. Therefore, all parameters, possible resolutions and limitations should be known. In my opinion, even the serial version of this method might be hard to reproduce by a non-expert, particularly since it is not clear if the approach works as claimed." "The presented results fail to demonstrate the different aspects of the contribution. A discussion of damping artifacts coming from non-uniform sampling as analyzed in [SG11] should be provided."

Answer: Thank you very much for this detailed list of changes. Due to your remarks we could clearly improve our result section and have exchanged all examples of the old manuscript. We include a discussion of damping along with longer simulation times as well as a detailed anaylsis of the simulation steps. We now better compare our approach to a non-adaptive simulation as well as the adaptive method as proposed by Adams et al.. We have added all parameters which are relevant for a temporal blending. For further details please see the result section and our answer to comment A in the summary.

Comment 11:"In my opinion, Section 7 which desribes the criterion for splitting and merging is not motivated well enough while the results do not reflect the described procedure." "I could not get how and why the diffusion layer is computed."

Answer: We have completely renewed the parts dealing with an adaptive identification of high-resolution regions including a figure describing the diffusion layer (See Sec.3.2 of our paper). Please see our answer to comment C in the summary.

Comment 12:"The blend sets generated in the video do not match with what is described in Sec. 7 or with the right image of Figure 1." "First of all, the size ratio of particles that are blended to those that are not blended is not intuitive. Maybe, the authors are drawing the blend sets much larger than other particles?" "Only the bunny is shown in the accompanying video while for the CBD scene, only one frame is illustrated in the paper." "In the video, I am missing a view which shows the non-uniform particles like they are illustrated in the right image of Figure 1. Showing the scenes whith such a view would definitely help to assess the quality of the method. Also the sequence should be shown longer and without changing the camera. Furthermore!, the video is missing a comparison with the non-adaptive method, in order to demonstrate that nearly no momentum is lost or added by the proposed method." "Moreover, both scenes described in the video are using 500k particles. Showing the scaling of the proposed method with the number of particles would be also helpful to assess the capability of the approach."

Answer: We have removed the bunny scene and instead implemented three new scenes, as described in comment B. Our video now shows all three scenes for longer while giving a better inside to our temporal blending and providing a side-by-side comparison to non-adaptive as well as nin comparison with a on-continuous re-sampling.

Comment 13:"The following references are highly related to the topic and should be cited and discussed: Solenthaler et al. 2011 "Two-Scale Particle Simulation" Raveendran et al. 2011 "Hybrid Smoothed Particle Hydrodynamics" Ihmsen et al. 2011 "A Parallel SPH Implementation on Multi-core CPUs"

Answer: We have included the references of Solenthaler et al. as well as the reference of Ihmsen et al. as they present a completely particle-based method. Thanks a lot for these references. Although we agree that Raveendran et al. improve the simulation speed with their method, they present a hybrid method combining their particles with an eulerian simulation grid. Instead, we rather stick to the most relevant related work.

Reviewer 4:

Comment 14:"Not even the adaptivity of particles is defined more flexible and still does only depend on the distance to the surface." "Also in general I'm not strongly convinced by the adaptive particle selection. This is not very flexible and just depends on the distance to the fluid surface. However, this does not take into account activity inside the fluid which could be missed by large particles!"

Answer: Please see our answers to comment A and comment B in the summary. We agree with you that an adaptive particle sampling based on the internal dynamics based on convection requires further investigation. We support a criterion for the internal diffusive flux. Usually, there is often a high correlation between physical and geometric complexity. As such we have shifted our focus on the temporal blending. We refer to this aspect in our future work at the end of the result section.

Comment 15: "The presentation is quite good, but some of the equations are not described and not easy to discover. " "For example the concentration, this is not well motivated and described and not easy to find where it is referenced from. Also the impulse correction part comes sudden and is not well rounded in the text." "The formulas and references are not as easy to find, read and understand as in some other SPH papers."

Answer: Thank you very much. For a description of the diffusion we now refer to some of our related work. Additionaly, we have included several result figures plus our video which shows the effect of diffusion. Additionally, we have included a figure in order to better describe why we need an impulse based correction of initial particle positions. Please see also the comments of Reviewer 5.

Comment 16:"... the following very important reference is missing: Solenthaler et al. 2011 "Two-Scale Particle Simulation""

Answer: Thank you, this is really a interesting approach. We have included the reference of [SG11] in our Related Work.

Comment 17: "The reasons for not giving a high score are that the method does not clearly show practical improvements, in visual quality for example, to [APKG07] nor to [SG11] and the experimental results are very thin. There are only short running times shown for the comparison to direct replacement, and it seems that after a few frames this works quite well and thus could be faster as it is simpler. Also the speedup is not very well documented to the other approaches, some more numbers and tables would be adequate here." "Also memory consumption could be reported more accurately, e.g. for future work to compare to."

Answer: We totally agree with you, the results have not been descriptive. Please see our changes according to this comment in the comment A.

Reviewer 5:

Comment 19:"The video is not so convincing. I would have liked to see a more realistic scenario. What are the groups of large spheres? Are we seeing 2 levels of particle sizes? Why do these groups pop in immediately instead of being blended in and out over time?"

Answer: The video as submitted to the Eurographics 2012 was in a quite unfinished state. Please see our new video including all three scenes. Please see also our answer to comment B in the summary.

Comment 20:"Considering the introduced complexity, the speedup of this method vs. standard SPH shown in Figure 9 are quite moderate."

Answer: Unfortunately you are right. However, we still get a speed-up of factor 1.5 in most of our examples which will increase with higher numbers of particles. In the context of incompressible fluids the speed-up is higher due to the incompressiblity constraint. Still if we would enable an uncouple particle levels as done by [SG11] we could get higher frame rates, however introducing divergence problems.

Comment 21:"Page 3: "Pressure linearly depends on density..." This is just one EOS. There are others too."

Answer: We have corrected this mistake. Please see Section 3.1. in the new manuscript.

Comment 22:"Eqn (6). Below the equation you say "evaluated in the domain Omega. So shouldn't the equation have an integral, i.e. $E(x) = \int_{\Omega}$—...—?" "Eqn (8) is the central equation. This is the main part of the paper that the authors should fix. For me, this equation does not describe what is going on. A representation cannot be expressed by a volume (Ω) alone. It is all about different sets of particles. So on the rhs of Eqn (8) I would expect sums over two different sets of particles. Looking at the lhs and the second term on the right I read that Q(x,t) is computed as the integral over Q(x,t) which looks recursive to me."

Answer: Thank you for revealing the problems in the explanation of the temporal blending. We have removed this ambiguity. Instead of using the integral interpolant over fluid volumes we now use the summation interpolant over particle sets L and H representing a high-resolution and a low-resolution particle set. Please see Sec.4. of the revised paper for a detailed description.

Comment 23:"I think talking about how to handle a concentration attribute with diffusive flux (Eqn (4)) complicates the paper unnecessarily. It is not essential for the explanation of the basic idea and to show how it works."

Answer: We agree with you and we first thought of removing the diffusion from our description. However, as one contribution our algorithm is capable of handling diffusion which would introduce problems for non-consistent simulation methods, e.g. [SG11], in order to conserve the total amount of a soluble substance. Additionally, we improved on the description by adding several figures explaining the effect of diffusion.

Comment 24:" λ in Eq(7) is not defined"

Answer: We have now included the values for λ_F and for λ_v and we refer to Monaghan for a more detailed description. Please see Section 3.1.

Comment 26: "Eqn (11) is not clear to me either. Are the authors handling

$1 \leftarrow n$ and $n \rightarrow 1$ splits and merges only? So is x_i always a single particle and x_p a set of particles? In Figure 5 it seems that the sync goes both ways. How are the x_p's synced?"

Answer: Thank you for this comment. Due to the missing information we have included a figure which presents the split/merge structures and how resolution levels interact with each other. We also deal with splitting and merging in the video. With both we think the synchronization between blend-domains is much more clear now.

Comment 27:"The second half of the first paragraph of Section 6 was difficult to understand too. How exactly are the k_0, k_1 .. computed? What does "each is computed in an optimal neighborhood with only level-l particles" mean? What is an "optimal" neighborhood? Can these constants be computed independently?"

Answer: We have rewritten the explanation in order to clarify this implementation aspect. Please see Sec.5 of our manuscript.

Yours sincerely, the authors