

Title: SCM-BP: An Intelligent Buffer Management Mechanism for Database in Storage Class Memory
JIDM

Dear Cristina Dutra,

Since we are submitting a revised version of our paper "SCM-BP: An Intelligent Buffer Management Mechanism for Database in Storage Class Memory", we are sending you our response to the reviewers' comments. We would like to take this opportunity to thank the valuable comments of the Reviewers.

- **Reviewer #1:** "There is a nice paper by stonebraker (OLTP through the looking glass) that shows that quite often OLTP database can fit within main memory. So, in this case, things are quite different and this offers the possibility for future work".
 - The indicated paper has been included as bib reference.
 - The following text can be found in the 4th. Paragraph of Section 6
 - Some authors advocate that quite often OLTP databases can fit within main memory [Harizopoulos et al. 2008}. Of course, such a feature reduces the impact of the buffer management component on DBMS performance. Nonetheless, we claim that this is not the case of OLTPs running on very large databases. Furthermore, for main memory databases, at least the log file has to reside in nonvolatile memory and the recovery process may take too much time, which is not reasonable for many OLTP applications.
- **Reviewer #1:** "There are too many knobs to tune. Typically, there is too much argumentation when too many parameters have to be set ...".
 - The following text can now be found in the 3rd. paragraph of Section 6
 - "Another improvement is to define a way to autonomously and dynamically adapt database parameters (e.g. read and write thresholds -- *r-threshold* and *w-threshold* -- and the increasing rate of read and write regions -- *rri* and *wri*) when a workload varies over the time."
- **Reviewer #1:** "The intuition behind the formula for WSF should be given".
 - The following text has been inserted in the 6th. Paragraph of Page 9 (Section 4.1)
 - **"The intuition behind this function is the following: When the miss ratio of read region raises up and at same time the number of write operations in the write region drops down (or does not increase in a given time window), the function should be applied. It is important to note that the higher the miss ratio in the read region, the lower the hit ratio and the higher swap-in operation frequency.** The return value provided by *wfs* is the quantitative indicator for augmenting (or not) the read region."
- **Reviewer #1:** "Algo2: lines 3-4 and 7-8 are different in their essence? Because they are different in their expression, so this makes you think that they are different. Plz. fix or clarify. Fig 6b. If there is space, split the figure in pairs: currently this is really hard to follow, esp. on a b/w printout"
 - Algorithm 2 has been fixed
 - Fig. 6 has been enlarged, and we believe it is now more readable.

- **Reviewer #2:** “porque os autores não realizaram comparação com alguma proposta que use "SCM-oriented Policies"”?
 - Infelizmente, não tivemos acesso aos códigos fontes das "SCM-oriented policies". Adicionalmente, mesmo os pseudo-códigos apresentados em alguns artigos não fornecem elementos suficientes para garantir que a implementação estaria correta. Portanto não pudemos implementá-las. Todavia, isto está devidamente planejado como próximo passo na pesquisa, como declarado na seção 6
 - “Extending comparative evaluation w.r.t. SCM-oriented Policies is also planned.”
- **Reviewer #2:** “como o mecanismo é proposto no contexto de BD, seria interessante mostrar, nos experimentos qual foi o SGBD usado e quais as operações de leitura e escrita foram executadas. Ainda: os algoritmos foram implementados como uma camada externa ao SGBD ou como parte dele?
 - Como declarado no quarto parágrafo da Section 5: “The experiments have been implemented as external layer to a DBMS. Nevertheless, we intend to implement SCM-BP in PostgreSQL.”
- **Reviewer #3:** “I have no negative points. In the following I point some suggestion to improve the paper. I have a question related to baseline used in the experiments. LRU is the baseline. But, in the related work (section3.2), authors point some work that are derivation of LRU. Why do you choose LRU? Could LRU derivation have better results? Considering the experimental results, could you statistically compare the results instead of just to show percentage number for the perceptual improvements? Does it make sense?
 - We have not implemented LRU-based policies (Section 3.2), because we had not access to their source code. However, **we summarize the discussion on SCM-oriented Policies with a qualitative comparison, which is depicted in Table II.** Besides, “extending comparative evaluation w.r.t. SCM-oriented Policies is also planned.”