

Open Consensus

(Comments of Reviewer 2)

Romain Boichat

Svend Frølund

Rachid Guerraoui

18 May 2001

Major Changes

- The introduction has been revamped and the open consensus concept is now well motivated with the necessary background for readers that are not familiar with the subject. Moreover, we have described more clearly the advantage of open consensus by better introducing the broadcast application.
- Throughout the whole paper, we have cleared out all assumptions about some previous work, e.g., Lamport or Aguilera et al.
- We have summarised clearly in one subsection (4.3) the analytical advantages of using open consensus vs traditional consensus.

Specific Changes

- The referee says *The paper is based heavily on, or relies is strongly intertwined with, the earlier paper by Aguilera, Chen and Toueg. This is not unreasonable, since both papers deal with very similar problems, but the seeming reliance of this paper on that suggests that it needs a much more extensive review of that work. The emphasis on ACT00a, with numerous citations underlines its importance, yet the exposition of that work is limited. Throughout the paper, reference to, and comparison with, ACT00a is made, yet there is no point at which this is coherently presented.*

The referee was right, we have remodeled the introduction and have tried to make a clear presentation of traditional consensus (ACT00a). Now, Section 1.3.2 presents the differences between open consensus and traditional consensus. We compare the analytical performance of open consensus and traditional consensus in Subsection 4.3.

- The referee says *The assumption is made that the reader will necessarily be familiar with all aspects of the cited work (eg Lamport) and that reference to specific aspects of that is acceptable.*

Throughout the whole paper, we have cleared out all assumptions about some previous work, e.g., Lamport or Aguilera et al., and have introduced the necessary background such that no specific knowledge is necessary to understand the paper, i.e., the paper is self-contained.

- The referee says *The limited code descriptions at the end of the paper didn't seem to offer very much at all, and it wasn't obvious to me why they were included.*

We have kept the code description of Section 6 since we find them useful for the reader to have a better picture of the classes and actually see the modularity of the whole architecture.

- The referee says *p3, line 26. Rephrase to avoid the double negative.*

Done.

- The referee says *p4, line 18. loses - > loses*

Modified.

- The referee says *p4, line 24. loses – > loses*
Changed.
- The referee says *p4, line 25. what are the channel definitions of ACT00a?*
We have changed this part of Subsection 2.2 to better introduce the background of the channels.
- The referee says *p7, line 2. well-behaviour – > good behaviour?*
This was not english, we modified it to good behaviour.
- The referee says *p7, first line of figure: should this be v1' or v'1? It seems to conflict with the text.*
The referee is correct, it is v1' and we modified it.
- The referee says *p7, line 13: "a process decides another value that it proposed even if this value was not decided" seems not to make sense to me.*
Indeed, this sentence did not make sense, we modified it.
- The referee says *p7, line 15: even this – > even though this*
Done.
- The referee says *p8, line 8: should "proposee" be "proposer"?*
We kept the notion of proposee since it is commonly used in the literature of distributed algorithms.
- The referee says *p8, footnote: rotatcing – > rotating*
Done.