GEORGETOWN UNIVERSITY

Nitin H Vaidya <nv198@georgetown.edu>

OPODIS19 notification for paper 54

1 message

OPODIS19 <opodis19@easychair.org> Fri, Oct 25, 2019 at 11:45 AM
To: Nitin Vaidya <nitin.vaidya@georgetown.edu>

Dear Nitin Vaidya,

Congratulations on your paper being accepted to OPODIS 2019!

Please find below the reviews for:

Exact Byzantine Consensus on Arbitrary Directed Graphs under Local Broadcast Model

Please do your best to take these comments into account when revising your paper for final publication. We will be in touch soon regarding the next steps in the process.

We look forward to seeing you in Neuchâtel in December!

Take care,
Seth and Roy

SUBMISSION: 54
TITLE: Exact Byzantine Consensus on Arbitrary Directed Graphs under Local Broadcast Model

---------- REVIEW 1 ----------
SUBMISSION: 54
TITLE: Exact Byzantine Consensus on Arbitrary Directed Graphs under Local Broadcast Model
AUTHORS: Muhammad Khan, Lewis Tseng and Nitin Vaidya

---------- Overall evaluation ----------
SCORE: 2 (accept)
----- TEXT:
This paper studies the classical Byzantine consensus problem in directed graphs in which the communication model is “local broadcast”. That is, communication links are not necessarily bi-directional, and a node cannot send different messages to different neighbors in a single round.

Under such model, the authors identify two equivalent conditions (namely, the SC condition and the NC condition) that are both necessary and sufficient for solving binary consensus. In particular, to prove sufficiency, a consensus algorithm is proposed, and its correctness if formally proved. However, this algorithm has runtime exponential in the number of participating nodes, and has made a somewhat strong assumption (all nodes know the entire network topology).

Technically, it seems the paper uses many techniques that are extensions of existing work, but some non-trivial adoption is needed. Moreover, sometimes, new tools need to be applied. I have not carefully checked every proof. But I have not found errors in the ones that I did read. So, I would assume the results are correct.

The paper is very well written and nicely polished. The proofs and the algorithm are not easy to describe, but the authors have done a good job of presenting the core ideas and intuitions, before giving out all the details.

Overall, I think this paper extends our understanding of Byzantine consensus, in a new setting. I recommend acceptance.

---------- REVIEW 2 ----------
SUBMISSION: 54
TITLE: Exact Byzantine Consensus on Arbitrary Directed Graphs under Local Broadcast Model
AUTHORS: Muhammad Khan, Lewis Tseng and Nitin Vaidya

---------- Overall evaluation ----------
SCORE: 2 (accept)
----- TEXT:
*** SUMMARY

The paper studies Byzantine agreement in directed graphs under the local broadcast model, which is motivated by wireless networks. Whenever a node broadcasts a message in this model it is received by all outgoing neighbors. This restriction also applies to the Byzantine nodes and hence weakens them significantly, as they can no longer send different messages to different processes.

The main contribution of the paper is a necessary and sufficient condition on the solvability of Byz agreement in this setting. The authors provide an algorithm and prove a lower bound that show that Byz agreement is solvable in a network with local broadcast, if and only if, for every partition (A,B), every node in B is reachable from some node in A by ≥f+1 vertex disjoint paths or vice versa, assuming f Byzantine failures.

*** Strengths:

- Local broadcast in directed networks is a well-motivated setting that is likely to be of interest to the community.
- Tight characterization of the network topology.

*** Weakness:

- algorithm iterates over all possible failure sets, i.e., has exponential time complexity.
- results are somewhat incremental over [13].
- only considers deterministic algorithms, while much of the recent progress on Byzantine agreement was achieved using randomization.

*** OTHER REMARKS:

- Footnote 1 can be omitted.

- Is "SC" an abbreviation?

- The reference to (a) in step (c) of algorithm 1 should be to (b).

---------- REVIEW 3 ----------
SUBMISSION: 54
TITLE: Exact Byzantine Consensus on Arbitrary Directed Graphs under Local Broadcast Model
AUTHORS: Muhammad Khan, Lewis Tseng and Nitin Vaidya

---------- Overall evaluation ----------
SCORE: 1 (weak accept)
----- TEXT:
The paper looks at the problem of Byzantine agreement in directed graphs. It define a condition on the topology and

shows that it is both necessary and sufficient. This is done by extending earlier results on arbitrary but undirected graphs.

The paper appears sound and is reasonably easy to follow. The problem is indeed important, but much work has already been published on this topic, and the significance of this added contribution is modest in my opinion (see reasons below). In other words, if one factors out all of the related work on this topic, this paper brings a non-null but still minor contribution.

The paper considers the local broadcast model, where a process can only communicate by broadcasting to its all neighbors, thus preventing equivocation (when a Byzantine process sends different versions of a message to different neighbors). All communication channels are however assumed to be FIFO. Although this is apparently inherited from previous work, this defeats the stated motivation of the paper of addressing wireless networks. Physically, such networks are prone to message loss and hence require retransmission. Managing retransmission typically relies on acknowledgments, which would require bidirectional channels. The paper should at least try to convince why assuming FIFO is reasonable (in particular for the sufficiency proof). It is true that omissions can be dealt with as Byzantine faults, but this comes at greater cost and effectively decreases the number of Byzantine faults being tolerated.

The proposed algorithm is combinatorial in the number of faults, because it has to execute as many rounds as combinations of fault sets. As such, the algorithm serves its purpose of proving the sufficiency of the condition in the FIFO model. However, it has little practical value except for systems with few processes and a very small number of Byzantine faults.

Some statements are erroneous (although this does not affect overall correctness, they should be corrected). On page 1, "in undirected graphs, each node has some path(s0 to each of the other nodes". This is of course wrong unless you state "in *CONNECTED* undirected graphs, ..." since the rest of the sentence is the very definition of a connected graph. Given the context, maybe the paper is trying to say that directed graphs can be weakly-connected? and this is what poses a challenge to adapt the algorithm in [13]?

The definition of condition NC and the proof in Figure 1 are very reminiscent of prior papers on this, with the Left and Right components and the faulty processes in the center. In particular, this is very similar to the condition found in "Iterative approximate byzantine consensus in arbitrary directed graphs" by Vaidya et al. in PODC 2012.

Some typos: p.1 "proof of correction" should be "proof of correctness"