# A Simple Byzantine Agreement Protocol

Jonathan Katz*

October 23, 2013

Berman, Garay, and Perry [1] give a simple BA protocol for single-bit inputs with polynomial complexity and optimal resilience. The protocol involves running the following *phase-king* subroutine with parties $P_1, \ldots, P_{t+1}$ successively playing the role of the king.

**Round 1** Each party $P_i$ sends their input $v_i$ to all other parties.

$P_i$ then sets $C_i^b := 1$ (for $b \in \{0,1\}$) iff at least $n-t$ parties sent it the bit $b$.

**Round 2** Each party $P_i$ sends $C_i^0$ and $C_i^1$ to all other parties. Let $C_{i\to j}^b$ denote the relevant value received by $P_j$ from $P_i$.

Each party $P_i$ sets $D_i^b := \left|\left\{ j : C_{j\to i}^b = 1 \right\}\right|$. If $D_i^1 > t$, it sets $v_i := 1$; otherwise, it sets $v_i := 0$.

**Round 3** The king $P_k$ sends $v_k$ to all parties. Each party $P_i$ then updates their input as follows: If $D_i^{v_i} < n-t$ then set $v_i$ equal to the value the king sent to $P_i$; otherwise, leave $v_i$ unchanged.

We begin with two lemmas about the phase king (sub-)protocol.

**Lemma 1** *Let $t < n/2$, and assume all $n-t$ honest parties begin the phase-king subroutine holding the same input $b$. Then all honest parties terminate that subroutine with the same output $b$.*

**Proof** Since all honest parties begin with input $b$, in the first round each honest party receives $b$ from at least $n-t$ parties, and receives $1-b$ from at most $t < n-t$ parties. So each honest $P_i$ sets $C_i^b := 1$ and $C_i^{1-b} := 0$. It follows that in round 2, each honest $P_i$ has $D_i^b \geq n-t > t$ and $D_i^{1-b} \leq t$, and $v_i = b$ at the end of that round. Since $D_i^{v_i} = D_i^b \geq n-t$ for an honest $P_i$, all honest parties ignore the value sent by the king and terminate the phase-king subroutine with output $b$. ■

**Lemma 2** *Let $t < n/3$. If the king is honest in some execution of the phase-king subroutine, then the outputs of all honest parties agree at the end of that subroutine.*

**Proof** An honest king sends the same value $v_k$ to all parties. So the only way agreement can possibly fail to hold is if some honest party $P_i$ does not set their input to the king's value, i.e., if $D_i^{v_i} \geq n-t$. We claim that if there exists an honest party $P_i$ for whom $D_i^{v_i} \geq n-t$, then $v_i = v_k$ and so agreement holds anyway. To see this, consider the two possibilities:

- **Case 1: $\mathbf{v_i = 1}$.** Since $D_i^1 \geq n-t$ we have $D_k^1 \geq n-2t > t$, and so $v_k = 1$ as well.

*jkatz@cs.umd.edu. Department of Computer Science, University of Maryland.

- **Case 2: $\mathbf{v_i = 0}$.** The fact that $D_i^0 \geq n - t$ implies $D_k^0 \geq n - 2t > t$. So at least one honest party $P_j$ sent $C_{j \to k}^0 = 1$ to $P_k$, implying that at least $n - t$ parties sent the bit '0' to $P_j$ in round 1 and consequently at most $t$ parties sent '1' to $P_j$ in round 1. But then any honest party received a '1' from at most $2t < n - t$ parties in round 1, and so any honest party $P_i$ has $C_i^1 = 0$. It follows that each honest party, and $P_k$ in particular, has $D_k^1 \leq t$; we conclude that $v_k = 0$ as desired.

■

**Theorem 1** *The above protocol achieves Byzantine agreement for any $t < n/3$.*

**Proof** Say all honest parties begin holding the same input. Then Lemma 1 implies that none of the honest parties ever change their input value in any of the phase-king subroutines, and so in particular they all terminate with the same output.

In any other case, we know that there must be at least one execution of the phase-king subroutine in which the king is honest. Following that execution, Lemma 2 guarantees that all honest parties hold the same input. Lemma 1 ensures that this will not change throughout the rest of the protocol.

■

# References


[1] P. Berman, J. Garay, and K. Perry. Bit Optimal Distributed Consensus. In *Computer Science Research*, pp. 313–322, Plenum Publishing Corporation, 1992.