Byzantine Generals

Wolfgang Mulzer

Let G be a set of n generals. Each general $g \in G$ has a *choice* $c_g \in \{A, R\}$, and the generals must reach a *common consensus* c^* that is close to the majority vote. Unfortunately, there are t *traitors* among the generals. We describe how the loyal generals can reach a common agreement despite the presence of traitors. For this, each loyal general executes the following algorithm. The local variable `me` stores the id of the current general. The associative arrays `rC` and `mC` contain the current general's views on the other generals' choices. The function `majority` receives a multi-set of choices and returns the most popular choice among them (breaking ties in favor of R).

1 $\texttt{rC}[\bot] \leftarrow c_{\texttt{me}}$
2 **for** $i := 0, \dots, t$ **do**
3   **for** *all pairwise distinct generals* $g_1 \to \cdots \to g_{i+1}$ *in* $G \setminus \{\texttt{me}\}$ **do**
4         $\texttt{send}\big(\texttt{me}, g_{i+1}, \texttt{rC}[g_1 \to \cdots \to g_i]\big)$
5   **for** *all pairwise distinct generals* $g_1 \to \cdots \to g_{i+1}$ *in* $G \setminus \{\texttt{me}\}$ **do**
6         $\texttt{receive}\big(g_{i+1}, \texttt{rC}[g_1 \to \cdots \to g_{i+1}]\big)$
7 **for** $i := t+1, \dots, 0$ **do**
8   **for** *all pairwise distinct generals* $g_1 \to \cdots \to g_i$ *in* $G \setminus \{\texttt{me}\}$ **do**
9         $\texttt{mC}[g_1 \to \cdots \to g_i] = \texttt{majority}\big(\texttt{rC}[g_1 \to \cdots \to g_i] \cup \texttt{mC}[g_{\to} \cdots \to g_i \to ?]\big)$
10 $c^* \leftarrow \texttt{mC}[\bot]$

Algorithm 1: The Byzantine Generals Algorithm.

Lemma 1. *Let $n \geq 3t+1$. The Byzantine Generals algorithm has the following two properties:*

(a) For all $i = 1, \dots, t+1$ and for all pairwise distinct generals $g_1 \to \cdots \to g_i$: if g_i is loyal, then for all loyal generals $g \notin \{g_1, \dots, g_i\}$, we have:

$$g.\texttt{mC}[g_1 \to \cdots \to g_i] = g.\texttt{rC}[g_1 \to \cdots \to g_i] = g_i.\texttt{rC}[g_1 \to \cdots \to g_{i-1}].$$

(b) For all $i = 1, \dots, t$ and for all pairwise distinct generals $g_1 \to \cdots \to g_i$: if $g_1, \dots, g_i$ are all traitors, then for all loyal generals $g, g' \notin \{g_1, \dots, g_i\}$, we have:

$$g.\texttt{mC}[g_1 \to \cdots \to g_i] = g'.\texttt{mC}[g_1 \to \cdots \to g_i].$$

Proof. We begin with property (a). The proof is by reverse induction on i. First, let $i = t+1$. In this case, by Line 9 from Algorithm 1, we have for any loyal general $g \notin \{g_1, \dots, g_{t+1}\}$,

$$g.\texttt{mC}[g_1 \to \cdots \to g_{t+1}] = g.\texttt{rC}[g_1 \to \cdots \to g_{t+1}] = g_{t+1}.\texttt{rC}[g_1 \to \cdots \to g_t],$$

since g_{t+1} is loyal. Next, we perform the inductive step from $i+1$ to i. Since g_i is loyal, general g_i sends the same value $g_i.\texttt{rC}[g_1 \to \cdots \to g_{i-1}]$ to all generals $g' \notin \{g_1, \dots, g_i\}$ in Line 4 of Algorithm 1. Thus, using the inductive hypothesis, for all loyal generals $g' \notin \{g, g_1, \dots, g_i\}$, we have

$$g.\texttt{mC}[g_1 \to \cdots \to g_i \to g'] = g.\texttt{rC}[g_1 \to \cdots \to g_i \to g'] = g_i.\texttt{rC}[g_1 \to \cdots \to g_{i-1}] = g.\texttt{rC}[g_1 \to \cdots \to g_i].$$

Since $n \geq 3t+1$ and since $i \leq t$, the set $G \setminus \{g, g_1, \dots, g_i\}$ contains at least t loyal generals and at most t traitors. Thus, according to line 9 in Algorithm 1,

$$g.\texttt{mC}[g_1 \to \cdots \to g_i] = \texttt{majority}\big(g.\texttt{rC}[g_1 \to \cdots \to g_i] \cup g.\texttt{mC}[g_1 \to \cdots \to g_i \to ?]\big) = g.\texttt{rC}[g_1 \to \cdots \to g_i].$$

This concludes the proof of (a), and we continue with the proof of property (b). Again, we use reverse induction on i. For the base case, let $i = t$. Since $g_1 \to \cdots \to g_t$ are all traitors, and since there are only t traitors overall, all generals in $G \setminus \{g_1, \dots, g_t\}$ are loyal. Thus, the multisets

$$g.\texttt{rC}[g_1 \to \cdots \to g_t] \cup g.\texttt{mC}[g_1 \to \cdots \to g_t \to ?]$$

and

$$g'.\texttt{rC}[g_1 \to \cdots \to g_t] \cup g'.\texttt{mC}[g_1 \to \cdots \to g_t \to ?]$$

are identical, so $g.\texttt{mC}[g_1 \to \cdots \to g_t] = g'.\texttt{mC}[g_1 \to \cdots \to g_t]$, as claimed. Next, we perform the inductive step from $i+1$ to i. By (a), we have

$$g.\texttt{mC}[g_1 \to \cdots \to g_i \to g'] = g.\texttt{rC}[g_1 \to \cdots \to g_i \to g'] = g'.\texttt{rC}[g_1 \to \cdots \to g_i]$$

and

$$g'.\texttt{mC}[g_1 \to \cdots \to g_i \to g] = g'.\texttt{rC}[g_1 \to \cdots \to g_i \to g] = g.\texttt{rC}[g_1 \to \cdots \to g_i]$$

Let $h \in G \setminus \{g_1, \dots, g_i, g, g'\}$. If h is loyal, then again by (a), we have

$$\begin{aligned} g.\texttt{mC}[g_1 \to \cdots \to g_i \to h] = g.\texttt{rC}[g_1 \to \cdots \to g_i \to h] &= h.\texttt{rC}[g_1 \to \cdots \to g_i] \\ &= g'.\texttt{rC}[g_1 \to \cdots \to g_i \to h] = g'.\texttt{mC}[g_1 \to \cdots \to g_i \to h]. \end{aligned}$$

If h is a traitor, then $g_1, \dots, g_i, h$ are all traitors, and by the inductive hypothesis, we have

$$g.\texttt{mC}[g_1 \to \cdots \to g_i \to h] = g'.\texttt{mC}[g_1 \to \cdots \to g_i \to h] =$$

By line 9 from Algorithm 1, it follows that $g.\texttt{mC}[g_1 \to \cdots \to g_i] = g'.\texttt{mC}[g_1 \to \cdots \to g_i]$, as claimed. □

Satz 2. *Suppose that $n \geq 3t+1$. For any two loyal generals $g, g' \in G$, Algorithm 1 ensures that $g.mC[g'] = c_{g'}$, $g'.mC[g] = c_g$, and $g.mC[h] = g'.mC[h]$, for any $h \in G \setminus \{g, g'\}$.*

Proof. This is a direct consequence of Lemma 1, by setting $i = 1$. □