

This lecture is based on "The Byzantine Generals Problem," a classic paper by L. Lamport, R. Shostak, and M. Pease. It appeared in ACM Transactions on Programming Languages and Systems, Vol. 4, No. 3 (July 1982). Within Brown it can be found at http://delivery.acm.org/ 1 0 . 1 1 4 5 / 3 6 0 0 0 0 / 3 5 7 1 7 6 / p 3 8 2 - 1 a m p o r t . p d f ? key1=357176&key2=2878731721&coll=ACM&dl=ACM&CFID=58767845&CFTOKEN=877102 67. It is discussed in Coulouris et al. in Section 15.5.









This version of the problem can be solved by doing the Byzantine Generals Problem n times concurrently, once with each general as the commander and the others as lieutenants.



Here the commander is a traitor.



To check whether a traitor is in their midst, the lieutenants exchange messages stating what they heard from the commander.



Of course, it might not be the commander who's a traitor, but one of the lieutenants. Note that, from the point of view of the bottom-left general, there's no difference between this scenario and the one of the previous slide. He can't be sure, if he attacks, that there will be another general joining him. Though this isn't a formal proof, it should be convincing evidence that with three generals of whom at most one is a traitor, the Byzantine generals problem has no solution.





The last assumption is satisfied in *synchronous* systems.





Here we have four generals, with at most one of them a traitor. In this case, the traitor is the commanding general.



Again we have four generals, but the traitor is the rightmost general. The two lieutenant generals can't distinguish this situation from the one of the previous slide. However, in both cases, since the majority of messages received say attack, they can feel confident that if they attack, they will be joined by two other generals.







The assumptions here are that communication is synchronous and that messages are unsigned. We'll modify these assumptions soon.



The reference to Albanians comes from the original paper ...

















This is from the aforementioned paper by Lamport, Shostak, and Pease. It's easy to analyze, but it's not a well formed algorithm ...



This is from the aforementioned paper by Lamport, Shostak, and Pease.



In this slide we work out the notation we'll be using for [7,2] (and beyond).

The commander, C, sends its order (v_0) to the lieutenants. Lieutenant i stores it in ${}_iv_0^0$. The Lieutenants then send their values to the others. The value Lieutenant k receives from Lieutenant m is placed in ${}_kv_m^{m0}$. The intent is that the superscripts show the paths through the tree. Lieutenant 1 follows the order *majority*(${}_1v_0^0$, ${}_1v_2^{20}$, ${}_1v_3^{30}$); lieutenant 2 follows the order *majority*(${}_3v_0^0$, ${}_3v_1^{10}$, ${}_3v_2^{20}$).



Consider now BGP with 7 generals, two of whom might be traitors. If we're assured the commander is not a traitor, then each lieutenant can simply use the order received from the commander.



If the commander could be a traitor, then the lieutenants must check with one another what order was actually sent. The slide shows the result of lieutenants 2 and 5 communicating with the others. If none of the lieutenants are traitors, we need go no further (though all the other lieutenants must communicate the orders they received as well). But, of course, it's not known whether any of the lieutenants are traitors.



But if both the commander and one of the lieutenants may be traitors, then more work is required. Here lieutenant 2 has communicated the order it received from the commander to the others, as in the previous slide. However, since one of the lieutenants may be a traitor, no one trusts lieutenant 2 to have communicated the same order to each of the others. We now have another instance of BGP, this time with at most one traitor (since we're already assuming the commander is a traitor). Thus the other lieutenants communicate with one another the order they received from lieutenant 2 (effectively saying to one another "lieutenant 2 said that the commander said ..."). Of course we have to do this for each of the lieutenants, so the complete diagram gets rather large.



Here's an attempt at a correct algorithm for BGP. It's initially invoked by the commander. m is the maximum number of traitors. gens is the set of generals (initially including the commander). v is the order (value). path is the path taken from the root (i.e., from the commander), initially empty. sender is the ID of the invoker, where 0 is the commander.

The *sendmsg* routine sends a message to the general given by the first argument. That general is to execute the command (i.e., procedure) given in the second argument (BGP), with the following arguments.

Note that the recursion proceeds in a breadth-first rather than a depth-first manner. After a sequence of calls to sendmsg, the caller does not wait for the results of these invocations to complete, but for the results of the invocations of sibling subtrees. In other words, the invocations a general makes of BGP are to send messages to other generals. What the general waits for is for the other generals to send messages to it. This waiting is handled by the "when defined" statement, which waits for the set of variables given as arguments to be given values. In principle it could be implemented as a sequence of operations on a set of semaphores, one for each variable.





This algorithm is from "Cloture Votes: n/4-resilient Distributed Consensus in (T+1) Rounds," P. Berman and J. Garay, Mathematical Systems Theory, 26(1), 1993.



This presentation is taken from "Distributed Computing: Principles, Algorithms, and Systems," by Ajay D. Kshemkalyani and Mukesh Singhal, Cambridge University Press, 2008.















Here the commander cryptographically signs her messages and the lieutenants are required to send copies of the signed messages to the others. Thus a traitorous lieutenant cannot lie about what the commander sent.



Here both lieutenant generals realize the commander is a traitor. However, the two of them must nevertheless come to a consensus. If there isn't a clear majority for attack or retreat (as is the case here), then the convention is that they retreat. Thus they have a consensus.



Here one lieutenant has failed to respond. In a synchronous system, we can use a timeout to decide that a process has failed and thus, in this case, interpret the silence as a "retreat".



