

Byzantine Generals

A Lecture in CE Freshman Seminar Series:
Ten Puzzling Problems in Computer Engineering



About This Presentation

This presentation belongs to the lecture series entitled “Ten Puzzling Problems in Computer Engineering,” devised for a ten-week, one-unit, freshman seminar course by Behrooz Parhami, Professor of Computer Engineering at University of California, Santa Barbara. The material can be used freely in teaching and other educational settings. Unauthorized uses, including any use for financial gain, are prohibited. © Behrooz Parhami

Edition	Released	Revised	Revised	Revised	Revised
First	May 2007	Apr. 2008	Apr. 2009	Apr. 2010	Apr. 2011

Reminder on ECE 1's Theme and Direction

A puzzling problem:

- ☞ looks deceptively simple, but ...
- ☞ appears very difficult, or even impossible, but is readily tamed with the appropriate insight

Topics thus far:

- Easy, Hard, Impossible (Collatz conjecture)
- Placement and Routing (houses & utilities)
- Satisfiability (making change)
- Cryptography (secret message)
- Byzantine Generals (liars and truth-tellers)

Many engineering problems are puzzle-like (especially in CE)

Each lecture starts with puzzles that we try to solve together

I introduce you to CE problems that are related to the puzzles

Topics for the 2nd half:

- Binary Search (counterfeit coin)
- Task Scheduling (Sudoku)
- String Matching (word search)
- Sorting Networks (rearranging trains)
- Malfunction Diagnosis (logical reasoning)



Do you realize how dead you will be if she realizes *this* is your weekly "bar meeting"?

The Island of Liars and Truth-Tellers



Setting for puzzles in the next few slides:
You are on an island populated by two tribes.
Members of one tribe consistently lie.
Members of the other tribe always tell the truth.
Tribe members can recognize one another,
but you can't tell them apart.



You run into a man on the island and ask him whether he is a truth-teller. A blaring siren prevents you from hearing his answer. You inquire, "Sorry, did you say you're a truth-teller?" He responds: "No, I did not." To which tribe does the man belong?

He is a liar

You meet a woman on the island. What single (yes/no) question can you ask her to determine whether she is a liar or a truth-teller?

If I asked you whether you were a liar, what would your answer be?

Meeting Two People on the Island

You meet two people *A* and *B* on the island. *A* says, “Both of us are from the liars tribe.” Which tribe is *A* from? What about *B*? **A: Liar, B: TT**

You meet two people, *C* and *D* on the island. *C* says, “Exactly one of us is from the liars tribe.” Which tribe is *D* from? **D: Liar**

You meet two people *E* and *F* on the island. *E* says, “It is not the case that both of us are from the truth-tellers tribe.” Which tribe is *E* from? What about *F*? **E: TT, F: Liar**


You meet two people, *G* and *H* on the island. Each of the two makes a statement. Which tribes are *G* and *H* from?
G says: “We are from different tribes.”
H says: “*G* is from the liars tribe.” **G: TT, H: Liar**

Meeting Three People on the Island

You meet three people *A*, *B*, *C*. You ask *A*, “How many among you are truth-tellers?” You don’t hear her answer, so you ask *B*, “What did she just say?” “She said one,” he replies. Then *C* adds, “Don’t believe him, he is lying!” What can you say about *B*’s or *C*’s tribe?

Solution by
exhaustive search

Answer: *B* is a liar,
C is a truth teller



Some Variations on the Theme

Many variations of these puzzles exist

Liars who lie selectively; for example, in answer to every other question or on certain days of the week

Inhabitants of another island lie consistently on Tuesdays, Thursdays, and Saturdays, and they tell the truth on the other four days of the week. You have forgotten what day of the week it is, so you ask a passerby. “Saturday,” he answers. “And what day will it be tomorrow?” you inquire. “Wednesday,” he replies. Can you tell what day it is today?

Thursday

Truth-tellers and nay-sayers: You are allowed only yes/no questions. One group of people answer truthfully and the other always answer “no”

How can you determine whether a person is a truth-teller or a nay-sayer?

Ask any question whose correct answer is “yes”



Liars, Randoms, and Truth-Tellers



Setting for more complex puzzles:

You are on an island populated by three tribes.
Members of one tribe always tell the truth.
Members of the second tribe choose to tell the truth or lie, completely at random.
Members of the third tribe consistently lie.
Tribe members can recognize one another, but you can't tell them apart.

Three people from the island, one representing each tribe, come to visit. How can you identify who is from which tribe by asking only three yes/no questions? Each question must be directed at only one person, but you can ask the same person multiple questions.



Hint : There are 6 possibilities for P1, P2, P3: **LRT, LTR, RLT, RTL, TLR, TRL**

Additional hint: Ask the leftmost person whether letters corresponding to the other two appear in alphabetical order. Then, regardless of the answer, you will have one position which excludes a random person.

The Two Generals' Paradox

Troops led by two generals are camped on the outskirts of an enemy city

The generals can only communicate via messengers who must travel through enemy territory and are thus subject to delays or capture


The two generals have previously agreed on a plan of attack, but they must communicate to set up the attack time

Not attacking together has dire results


G1 decides to send the message, "Let's attack at noon tomorrow"

G1 will not attack before getting an acknowledgment from G2

G2 will not attack before making sure that his acknowledgment was received by G1 (because he knows G1 would not attack otherwise), so he waits for an acknowledgment of his acknowledgment



Liars, Randoms, and Truth-Tellers Stand for ...



Sites communicating with one another to reach an agreement (e.g., to select a coordinating site, often called “leader”)

Site status ...

Healthy: Gives the appropriate response to every message

Truth-teller

Crashed: Does not respond to any message

Quiet

Permanently failed: May respond identically to every message

Nay-sayer

Permanently failed: May give the wrong response consistently

Liar

Arbitrarily failed: May give an unpredictable response

Random

Maliciously failed: Gives a response that is calculated to do the maximum harm (adversary, worst-case failure)

Byzantine

The Byzantine Generals Problem


A general and $n - 1$ lieutenant generals lead n divisions of the Byzantine army camped on the outskirts of an enemy city

The n divisions can only communicate via messengers, who may be captured or arbitrarily delayed (due to the need to hide for a while)



We seek a scheme for the generals to agree on a common plan of action (attack or retreat), even if some of the generals are traitors who will do anything to prevent loyal generals from reaching agreement

The problem is nontrivial even if messengers are totally reliable



With unreliable messengers, the problem becomes very complex



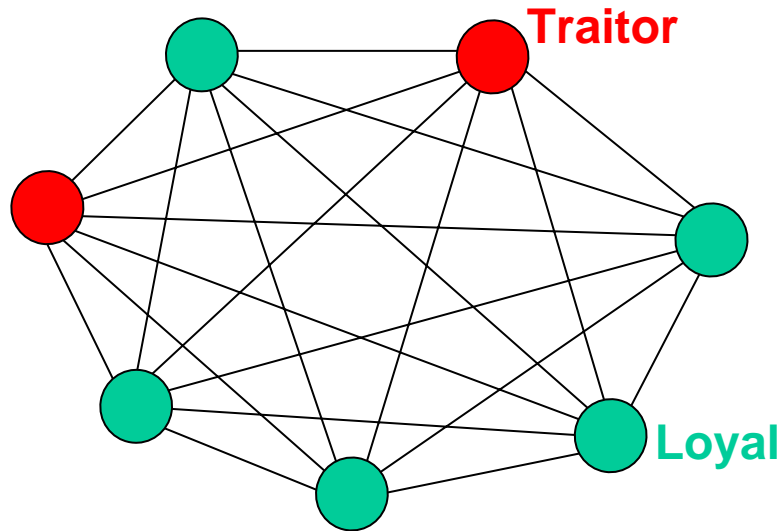
Byzantine Generals with Reliable Messengers



With f Byzantine failures, $\geq 3f + 1$ nodes needed to reach agreement



$3f + 1$ Generals Needed with f Traitors



By exchanging messages in multiple rounds, the $2f + 1$ loyal generals can eventually reach a common plan of action which matches the order of the commanding general, provided the latter is loyal

Some deem Byzantine faults very unlikely and not worth considering

In “The Real Byzantine Generals,” the authors show why Byzantine faults are real and must be treated in both hardware and software


<http://ieeexplore.ieee.org/iel5/9579/30281/01390734.pdf>

“If a designer spent 50 hours per week, 52 weeks per year, for 35 years staring at one system, that would be less than 10^5 hours . . . far short of typical avionics requirements.”

$$50 \times 52 \times 35 = 91,000$$

Without Malicious Faults, Voting Will Do

Data fusion:
Obtaining dependable results
from potentially incorrect,
inaccurate, or incomplete data



Centralized voting with majority rule

Approximate voting with imprecise inputs:

E.g., temperature readings of 78.2, ~~45.5~~, 79.1, 78.7, ~~21.2~~, ~~120.0~~, 77.6

Mean of reasonable inputs 78.4

Median of all inputs 78.2

Distributed voting: Same concept, provided erroneous values are seen identically by the fusion processes at all sites

Voting Comes in Many Flavors

Example: What time is it?

Seven students write the exact time
(hour and minute) on sticky notes

Sort the sticky notes on the board

Pick one of the following values:

Majority, if a majority exists


Plurality, if a plurality exists

Median of all the values proposed

Mean of all the values proposed


Mean of five values, after removing
the largest and smallest of the seven

Mean of three values, after removing
the 2 largest and 2 smallest values



Interval voting:

Each proposer supplies a range of values that is guaranteed to hold the correct value



Mathematics of Voting

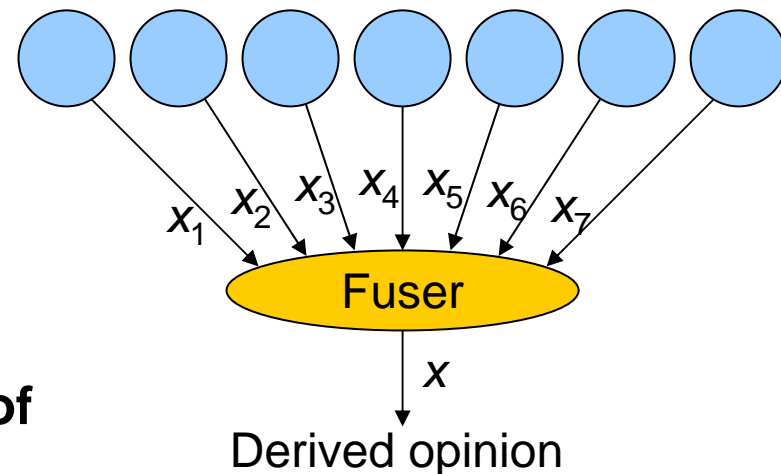
Voting studied in several fields:

Mathematics / Computing

Political science

Sociology (social choice theory)

Economics



No voting scheme is totally fool-proof

Regular voting: Candidates A1, A2, B

More qualified candidates A1 and A2 may split the votes, leading to the election of B (run-off helps in solving this problem, but creates others)


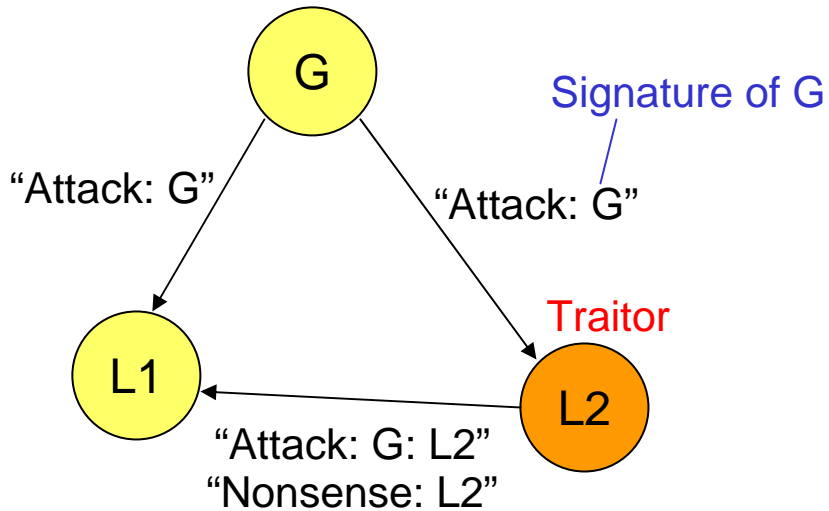
Approval voting

Vote for any number of candidates you like

Borda voting comes pretty close to an ideal voting scheme

Each participant ranks all candidates; tally votes by giving n points to each 1st-place choice, $n - 1$ points for 2nd place, ... , 1 point for n th place

With Signed Messages, Agreement is Easy



L2 may take two actions –

Forward the signed message:
This leads to correct outcome

Send a different fake message
that is recognized by L1 as fake
(loyal generals ignore messages
coming from known traitors)

