The coordinated attack problem
from Jim Gray 1978 presented in:

Knowledge and Common Knowledge in a Distributed Environment*

Joseph Y. Halpern

IBM Almaden Research Center
San Jose, CA 95120

Yoram Moses

Department of Applied Mathematics
The Weizmann Institute of Science
Rehovot, 76100 ISRAEL

Abstract: Reasoning about knowledge seems to play a fundamental role in distributed systems. Indeed, such reasoning is a central part of the informal intuitive arguments used in the design of distributed protocols. Communication in a distributed system can be viewed as the act of transforming the system's state of knowledge. This paper presents a general framework for formalizing and reasoning about knowledge in distributed systems. We argue that states of knowledge of groups of processors are useful concepts for the design and analysis of distributed protocols. In particular, *distributed knowledge* corresponds to knowledge that is "distributed" among the members of the group, while *common knowledge* corresponds to a fact being "publicly known". The relationship between common knowledge and a variety of desirable actions in a distributed system is illustrated. Furthermore, it is shown that, formally speaking, in practical systems common knowledge cannot be attained. A number of weaker variants of common knowledge that are attainable in many cases of interest are introduced and investigated.

*This is a revised and expanded version of a paper with the same title that first appeared in the *Proceedings of the 3rd ACM Conference on Principles of Distributed Computing, 1984*. It is essentially identical to the version that appears in *Journal of the ACM* **37**:3, 1990, pp. 549–587. The work of the second author was supported in part by DARPA contract N00039-82-C-0250.

> Two divisions of an army are camped on two hilltops overlooking a common valley. In the valley awaits the enemy. It is clear that if both divisions attack the enemy simultaneously they will win the battle, whereas if only one division attacks it will be defeated. The divisions do not initially have plans for launching an attack on the enemy, and the commanding general of the first division wishes to coordinate a simultaneous attack (at some time the next day). Neither general will decide to attack unless he is sure that the other will attack with him. The generals can only communicate by means of a messenger. Normally, it takes the messenger one hour to get from one encampment to the other. However, it is possible that he will get lost in the dark or, worse yet, be captured by the enemy. Fortunately, on this particular night, everything goes smoothly. How long will it take them to coordinate an attack?

We now show that despite the fact that everything goes smoothly, no agreement can be reached and no general can decide to attack. (This is, in a way, a folk theorem of operating systems theory; cf. [Gal79, Gra78, YC79].) Suppose General A sends a message to General B saying "Let's attack at dawn", and the messenger delivers it an hour later. General A does not immediately know whether the messenger succeeded in delivering the message. And because B would not attack at dawn if the messenger is captured and fails to deliver the message, A will not attack unless he knows that the message was successfully delivered. Consequently, B sends the messenger back to A with an acknowledgement. Suppose the messenger delivers the acknowledgement to A an hour later. Since B knows that A will not attack without knowing that B received the original message, he knows that A will not attack unless the acknowledgement is successfully delivered. Thus, B will not attack unless he knows that the acknowledgement has been successfully delivered. However, for B to know that the acknowledgement has been successfully delivered, A must send the messenger back with an acknowledgement to the acknowledgement .... Similar arguments can be used to show that no fixed finite number of acknowledgements, acknowledgements to acknowledgements, etc. suffices for the generals to attack. Note that in the discussion above the generals are essentially running a *handshake* protocol (cf. [Gra78]). The above discussion shows that for no k does a k-round handshake protocol guarantee that the generals be able to coordinate an attack.

In fact, we can use this intuition to actually prove that the generals can never attack and be guaranteed that they are attacking simultaneously. We argue by induction on d — the number of messages delivered by the time of the attack — that d messages do not suffice. Clearly, if no message is delivered, then B will not know of the intended attack, and a simultaneous attack is impossible. For the inductive step, assume that k messages do not suffice. If $k+1$ messages suffice, then the sender of the $(k+1)^{\text{st}}$ message attacks without knowing whether his last message arrived. Since whenever one general attacks they both do, the intended receiver of the $(k+1)^{\text{st}}$ message must attack regardless of whether the $(k+1)^{\text{st}}$ message is delivered. Thus, the $(k+1)^{\text{st}}$ message is irrelevant, and k messages suffice, contradicting the inductive hypothesis.