


Homework 2

The role of performance models in understanding the performance of collective communications

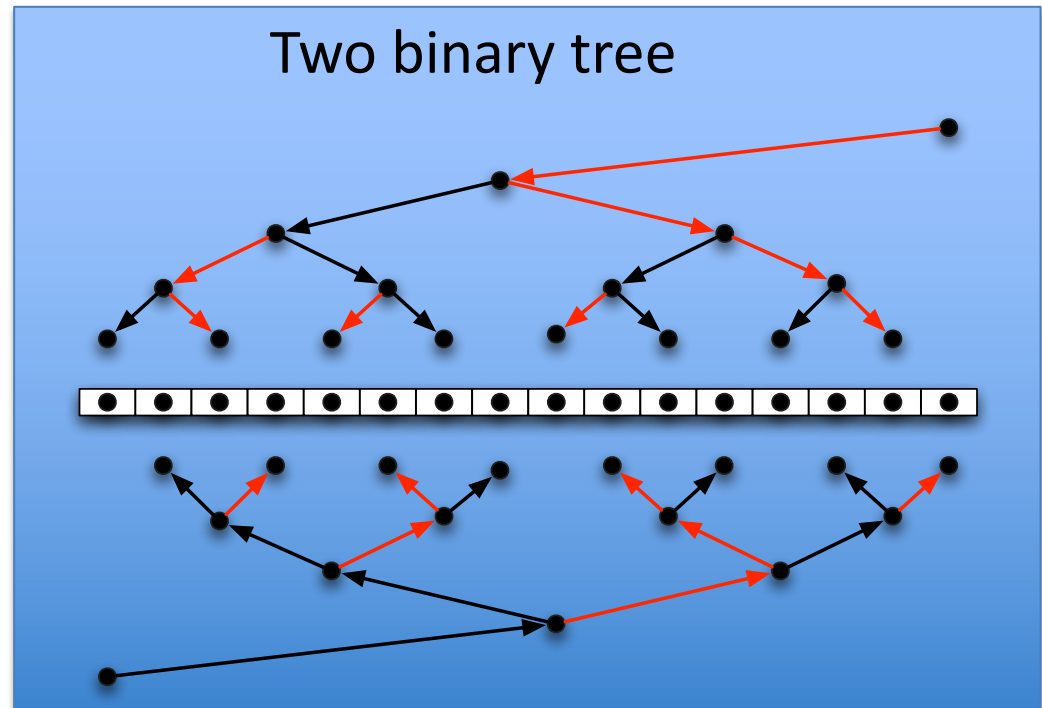
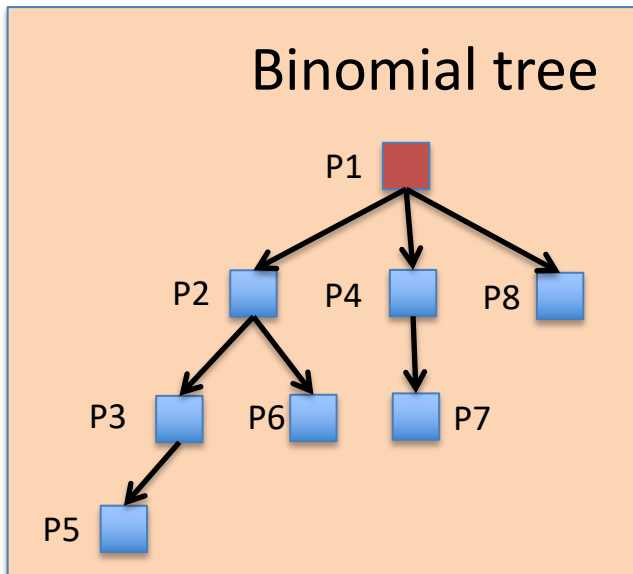

Deadline: January 29 2020

Broadcast

- A broadcast collective communication is a rooted communication pattern where a data is transferred from a single location toward multiple other locations. Upon completion all the locations hold identical content.
- Such a collective communication can be implemented using different communication patterns. We are interested in their properties, based on the underlying network characteristics.
- Let's take 3 patterns: a chain, a binomial tree and a two-binary tree (depicted on the next page).



Three popular collective topologies



Two-tree algorithms for full bandwidth broadcast, reduction and scan
Peter Sandersa, Jochen Specka, Jesper Larsson Träff
<http://www.sciencedirect.com/science/article/pii/S0167819109000957>

Q1: Using the LogP/LogGP model evaluate the cost of a simple broadcast over each of the previous 3 topologies.

- Consider a fixed number of participants and two types of messages (small and very large*). Order the three topologies based on the performance of the broadcast algorithm using the LogP parameters from ¹

Q2: Consider a pipelined approach to the broadcast operation, where the original message is split in several size-identical fragments. The broadcast of each fragment is pipelined with respect of all previous fragments.

- Consider again a fixed number of participants and two types of messages (small and very large*). Order the three topologies based on the performance of the broadcast algorithm under same constraint as above (¹)
 - Consider now the case of an intermediate message and vary the number of participants. Answer the previous question again.
- relationship between the number of potential fragments and the number of participants

¹ <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.95.3022&rep=rep1&type=pdf>

Extra question

- The reduction operation is somehow the opposite of the broadcast: originally each participant has a data, at the end the root will have the composition of all data from all participants (composition via a commutative/associative op).
- The major difference compared with the broadcast (in addition to the change of direction of the flow of data), is the operation applied between all data
 - Considering this extra operation, model and evaluate the cost of a reduce call (answer the same questions from the previous slide).

