A King in two tournaments

Yehuda Afek * Eli Gafni † Nati Linial ‡

March 31, 2012

Abstract

According to Landau's theorem every tournament has a King, i.e., a vertex that can reach every other vertex in two steps or less. Here we extend this result by considering two not-necessarily-equal tournaments T_1, T_2 on the same vertex set. We show that there is a vertex which can reach all vertices via (i) A step in T_1, or (ii) A step in T_2 or (iii) A step in T_1 followed by a step in T_2.

1 Introduction

Directed graphs play a major role in modeling computer systems. This is particularly the case in the domain of distributed computing. In such a model, a directed edge (i, j) in a digraph indicates that processor P_i is sending a message to processor P_j. In what follows, communication is governed by an n-vertex tournament $T = ([n], E)$. Initially each processor P_i holds a data item m_i that is known only to it. In a round of communication processor P_i tells every processor P_j, with $(i, j) \in E$, everything it knows. An old theorem of Landau [2] asserts that every tournament has a vertex that 2-dominates all other vertices. Consequently, after two rounds of communication some data item m_ν is known to all processors P_j. Notice that likewise in a dual way, there is also a vertex μ such that, after two rounds processor P_μ gets to know all data items m_j.

In the study of certain fundamental aspects of distributed computing [1] the following question came up. What happens if different

*The Blavatnik School of Computer Science, Tel-Aviv University, Israel 69978. afek@tau.ac.il

†Computer Science Department, Univ. of California, LA, CA 95024, eli@cs.ucla.edu

‡School of Computer Science and Engineering, Hebrew University, Jerusalem 91904, Israel nati@cs.huji.ac.il

tournaments govern the communication in different rounds? As before, communication takes place in rounds, each of which is governed by an arbitrary tournament. Thus, we start again with each processor P_i holding its private data item m_i for $n \geq i \geq 1$. There is a sequence of tournaments $T_1 = ([n], E_1), T_2 = ([n], E_2), \ldots$. At round 1 the communication pattern is dictated by tournament T_1, at round 2 by T_2, etc. Do the consequences of Landau's theorem still hold in this more general setting? If so, how many rounds of communication are necessary to yield the same conclusions as above? We find it satisfactory and somewhat surprising that as in Landau's case, two rounds of communication suffice.

2 Two rounds suffice

Let $T_1 = ([n], E_1), T_2 = ([n], E_2)$ be two tournaments. Suppose that in a round of T_1 followed by a round of T_2 processor P_j gets to know data item m_i. We denote this by $i \Rightarrow j$. Clearly, this is equivalent to

$$i = j \;\vee\; (i,j) \in E_1 \;\vee\; (i,j) \in E_2 \;\vee\; \exists k \; s.t. \; (i,k) \in E_1 \text{ and } (k,j) \in E_2$$

It is useful to note that the negation of this condition $i \not\Rightarrow j$ is equivalent to

$$i \neq j \quad \wedge \quad (j,i) \in E_1 \quad \wedge \quad (j,i) \in E_2 \quad \wedge \quad \Gamma_2(j) \supset \Gamma_1(i) \tag{1}$$

where $\Gamma_\delta(x)$ is the set of out-neighbors of x in tournament T_δ.

Theorem 1. *Let $T_1 = ([n], E_1), T_2 = ([n], E_2)$ be two tournaments. Then there is $\nu \in [n]$ such that $\nu \Rightarrow j$ for every j.*

Proof. By induction on n. The statement is easily verified for $n = 3$. Let n be the smallest integer for which the theorem does not hold and let $T_1 = ([n], E_1), T_2 = ([n], E_2)$ be a counterexample. By minimality of n, for every $n \geq j \geq 1$ there is some $n \geq i \geq 1$ such that $i \Rightarrow_j k$ for every $k \neq i, j$, where $\Rightarrow_j$ indicates that the relation is defined with respect to tournaments $T_1 \setminus \{j\}$ and $T_2 \setminus \{j\}$. When this happens we say that i is *singled out* by j. Clearly $i \not\Rightarrow j$, or else the theorem holds with $\nu = i$, since $i \Rightarrow_j k$ clearly implies $i \Rightarrow k$. Consequently, no vertex is singled out more than once. We denote $\pi(j) = i$ and conclude that π is a permutation on $[n]$, since $\pi(j)$ is defined for every $n \geq j \geq 1$ and π is an injective mapping.

However, this is impossible. By Condition (1), every j satisfies $|\Gamma_2(j)| > |\Gamma_1(\pi(j))|$. But $\sum_j |\Gamma_2(j)| = \sum_j |\Gamma_1(\pi(j))| = \binom{n}{2}$, since π is a permutation. This contradiction completes the proof. $\square$

The same proof yields as well

Corollary 2. *Let $T_1 = ([n], E_1), T_2 = ([n], E_2)$ be two tournaments. Then there is $\mu \in [n]$, such that $i \Rightarrow \mu$ for every i.*

Proof. The same proof works. Just switch between T_1, T_2 and reverse all edges in the two tournaments.

□

References

[1] Y. Afek and E. Gafni, *Asynchrony from Synchrony.* Arxiv, http://arxiv.org/abs/1203.6096 , Janurary 2012

[2] H. Landau. On dominance relations and the structure of animal societies, III: The condition for score structure. *Bulletin of Mathematical Biophysics*, 15(2):143148, 1953.