Int J Comput Vis (2013) 104:315–342
DOI 10.1007/s11263-012-0571-2

A Distributed Mincut/Maxflow Algorithm Combining Path Augmentation and Push-Relabel

Alexander Shekhovtsov · Václav Hlaváč

Received: 30 November 2011 / Accepted: 4 September 2012 / Published online: 22 September 2012

Abstract We propose a novel distributed algorithm for the minimum cut problem. Motivated by applications like volumetric segmentation in computer vision, we aim at solving large sparse problems. When the problem does not fully fit in the memory, we need to either process it by parts, looking at one part at a time, or distribute across several computers. Many MINCUT/MAXFLOW algorithms are designed for the shared memory architecture and do not scale to this setting. We consider algorithms that work on disjoint regions of the problem and exchange messages between the regions. We show that the region push-relabel algorithm of Delong and Boykov (A scalable graph-cut algorithm for N-D grids, in CVPR, 2008) uses $\Theta(n^2)$ rounds of message exchange, where n is the number of vertices. Our new algorithm performs path augmentations inside the regions and push-relabel style updates between the regions. It uses asymptotically less message exchanges, $O(\mathcal{B}^2)$, where $\mathcal{B}$ is the set of boundary vertices. The sequential and parallel versions of our algorithm are competitive with the state-of-the-art in the shared memory model. By achieving a lower amount of message exchanges (even asymptotically lower in our synthetic experiments), they suit better for solving large problems using a disk storage or a distributed system.

Keywords Maximum flow · Minimum cut · Distributed · Parallel · Push-relable · Augmenting path · Large scale

A. Shekhovtsov (✉) · V. Hlaváč
Center for Machine Perception, Czech Technical University in Prague, Technicka 2, 16627 Praha 6, Czech Republic
e-mail: shekhole@fel.cvut.cz

V. Hlaváč
e-mail: hlavac@fel.cvut.cz
url: http://cmp.felk.cvut.cz

1 Introduction

Minimum s–t cut (MINCUT) is a classical combinatorial problem with applications in many areas of science and engineering. This research was motivated by the wide use of MINCUT/MAXFLOW problems in computer vision, where large sparse instances need to be solved. We start by a more detailed overview of models and optimization techniques in vision, where the MINCUT problem is employed and give examples of our test problems.

1.1 MINCUT in Computer Vision

In some cases, an applied problem is formulated directly as a MINCUT. More often, however, MINCUT problems in computer vision originate from the Energy minimization framework (maximum a posteriori solution in a Markov random field model). Submodular Energy minimization problems completely reduce to MINCUT (Ishikawa 2003; Schlesinger and Flach 2006). When the energy minimization is intractable, MINCUT is employed in relaxation and local search methods. The linear relaxation of pairwise Energy minimization with 0-1 variables reduces to MINCUT (Boros et al. 1991; Kolmogorov and Rother 2007) as well as the relaxation of problems reformulated in 0-1 variables (Kohli et al. 2008). Expansion-move, swap-move (Boykov et al. 1999) and fusion-move (Lempitsky et al. 2010) algorithms formulate a local improvement step as a MINCUT problem.

Many applications of MINCUT in computer vision use graphs of a regular structure, with vertices arranged into an N-D grid and edges uniformly repeated, e.g., 3D segmentation models illustrated in Fig. 1(c), 3D reconstruction models, Fig. 1(b). Because of such regular structure, the graph itself need not be stored in the memory, only the edge capacities, allowing relatively large instances to be