Journal of Network and Computer Applications 43 (2014) 27–41

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

journal homepage: www.elsevier.com/locate/jnca

Energy Aware Two Disjoint Paths Routing

Gongqi Lin [a], Sieteng Soh [a], Kwan-Wu Chin [b], Mihai Lazarescu [a]

[a] Department of Computing, Curtin University of Technology, Perth, WA, Australia
[b] University of Wollongong, Northfields Avenue, Wollongong, NSW, Australia

ARTICLE INFO

Article history:
Received 5 June 2013
Received in revised form
27 February 2014
Accepted 7 April 2014
Available online 30 April 2014

Keywords:
Maximum link utilization
Energy savings
Reliability
Routing
Throughput
Two disjoint paths

ABSTRACT

Network robustness and throughput can be improved by routing each source-to-terminal (s_d, t_d) demand *d via* two disjoint paths (2DP). However, 2DP routing increases energy usage despite yielding lower link utilization and higher redundancy. In this paper, we address the problem of minimizing the energy usage of networks that use 2DP. Specifically, our problem, called Energy-Aware Two Disjoint Paths Routing (EAR-2DP), is to maximally switch off redundant links while guaranteeing at least $0 \le T \le 1.0$ fraction of all possible (s_d, t_d) 2DPs remain on and their maximum link utilization (MLU) is no greater than a configured threshold. We first prove that EAR-2DP is NP-complete. Then, we design a fast heuristic solution, called Two Disjoint Paths by Shortest Path (2DP-SP). We have extensively evaluated the performance of 2DP-SP on real and/or synthetic topologies and traffic demands with two link-disjoint paths (2DP-L) and two node-disjoint paths (2DP-N). Our simulation results show that 2DP-SP can reduce network energy usage, on average, by more than 20%, even for MLU below 50%. As compared to using Shortest Path (SP) routing, while reducing energy by about 20%, 2DP-SP does not significantly affect the path length of each (s_d, t_d) demand, even for MLU < 50%. Furthermore, almost 94.2% of routes produced by 2DP-SP have route reliability up to 35% higher as compared to SP and up to 50% of the routes are only 5% less reliable than those of 2DP routing without energy savings.

1. Introduction

Increasing the robustness and reliability of networks are critical management issues, especially given the recent emergence of real time applications such as video on demand and voice over Internet Protocol (Xiong et al., 2009). In addition, with optical fiber bandwidth and node capacity increasing exponentially, a broken fiber span or link failure can cause significant degradation in Quality of Service (QoS) (Suzuki and Tobagi, 1992). To improve resiliency to link/router failures, source to terminal (s_d, t_d) data packets for traffic demand *d* can be routed along multiple link disjoint paths (Guo et al., 2003). However, the authors of (Acharya et al., 2004) observe that while achieving reliability, a multiple disjoint paths scheme makes very inefficient use of network resource due to high path redundancy and low link utilization. Thus, employing multiple disjoint paths incurs higher energy usage, which further increases the overall energy expenditure of the Information and Communication Technology (ICT) sector. A number of studies estimate the energy consumption related to ICT itself to vary from 2% to 10% of the worldwide energy consumption (Pickavet et al., 2008). In particular, these studies reported that networking devices such as IP routers and line-cards consume the largest majority of energy (Tucker et al., 2008). Thus, effective energy-aware traffic engineering (Zhang et al., 2010) is needed to minimize over-provisioned resources.

This paper describes an NP-complete optimization problem that aims to reduce the energy usage of networks that support two disjoint paths (2DP) routing. Specifically, our problem aims to reduce network energy usage while ensuring the network maintains at least $0 \le T \le T_{max}$ fraction of all possible (s_d, t_d) 2DPs and each link's maximum link utilization (MLU) is at most $0 \le U_T \le 100\%$; T_{max} is the fraction of the total number of (s_d, t_d) paths that have at least one 2DP. Note that energy consumption is reduced by shutting down unnecessary links. Specifically, to provide energy efficient, fault-tolerant and high bandwidth to upper-layer applications, we route each (s_d, t_d) demand through its (s_d, t_d) 2DP, if one exists, while powering off idle/unused links to reduce energy. Moreover, we route demands whilst considering the MLU bound of U_T. This MLU requirement is usually done to absorb traffic burst, to ensure queuing delay at each router to be within a given bound (Zhang et al., 2010), and also to prevent network failures due to congestion and component failures (Kodialam et al., 2006). In summary, our contributions are twofold:

- First, we pose a problem, called Energy-Aware Two Disjoint Paths Routing (EAR-2DP), and prove the problem is NP-complete. Our problem is important in reducing energy usage in networks that use 2DP to improve fault tolerance and/or bandwidth/throughput.

http://dx.doi.org/10.1016/j.jnca.2014.04.004
1084-8045/

- Second, we propose a novel algorithm, called Two Disjoint Paths by Shortest Path (2DP-SP) to solve EAR-2DP. Our 2DP-SP identifies network links that can be powered-off under two constraints: the threshold of T and U_T. To the best of our knowledge, this is the first algorithm that jointly reduces the number of links in a wired network whilst satisfying the two constraints. Note that 2DP-SP can be used for applications that require two link-disjoint paths (2DP-L) or two node-disjoint paths (2DP-N) routing. Our extensive experiments confirm the efficiency of 2DP-SP and its impact on network performance, *i.e.*, reliability and delay.

The rest of the paper is organized as follows. Section 2 presents related work. Section 3 gives an overview of the EAR-2DP problem and a formal proof of its NP-completeness. Section 4 describes our heuristic algorithm, 2DP-SP, to solve the problem, its time complexity and implementation issues. Section 5 evaluates the performance of 2DP-SP using both real and/or synthetic topologies and data. Finally, Section 6 concludes the paper.

2. Related work

Kandula et al. (2005) propose a method to dynamically split traffic over multiple, not necessarily disjoint, paths. A load balancer (Kandula et al., 2005) is used to split the traffic of an (s_d, t_d) demand among its available paths with the objective of minimizing the maximum network utilization. The balancer selects the paths, starting from the shortest, to route the split traffic flows. In contrast, Izmailov and Niculescu (2002) propose using multiple link-disjoint paths to route evenly split flow to provide higher bandwidth protection against component failures. Specifically, their work involves selecting working/active and protection/alternative paths from the set of disjoint paths pre-computed for each (s_d, t_d) pair. However, reserving protection paths requires more resources that incurs low link utilization, and thus effectively wastes energy. Further, paper (Izmailov and Niculescu, 2002) does not address the case where working paths cannot not route split flow. In contrast, our work aims to switch-off links, to save energy, while satisfying a given bound on maximum link utilization and supporting applications that use Two Link-disjoint Paths (2DP-L) routing. Similar to the method in (Izmailov and Niculescu, 2002), our paper considers Equal-cost multipath (ECMP) (Multipath Issues in Unicast and Multicast Next-Hop Selection) over 2DP to offer higher bandwidth and fault-tolerance. ECMP is a commonly deployed technique where routers keep track of all shortest paths, and evenly split traffic amongst them (He and Rexford, 2008).

The problem of finding link/node disjoint paths can be viewed as a special case of the minimum cost flow problem as demonstrated in Suurballe and Tarjan (1984) and Bhandari (1994). In Suurballe and Tarjan (1984), Suurballe proposes an algorithm to find k node-disjoint paths with minimal total length using a path augmentation method. The basic idea of Suurballe's algorithm is to find 2DP-L based on the shortest path and a shortest augmenting path. In Xiong et al. (2009), the algorithm in Suurballe and Tarjan (1984) is extended to the Suurballe–Tarjan (S–T) algorithm for finding 2DP-L from one source node to n destination nodes using only a single *Dijkstra*-like computation. To find n pairs of disjoint paths, the S–T algorithm requires $O(m\log(1+m/n)n)$ time; where n is the number of destination nodes and m is the number of links.

Building on the S–T algorithm, Kar et al. (2002) and Kodialam and Lakshman (2000, 2002) develop new algorithms to find 2DP that can serve as active and backup paths for routing bandwidth guaranteed connections. In Liang (2001), Liang extended the S–T algorithm to find 2DPs between a source and a destination node with the following performance constraints: network load and routing cost.

Recently, some papers combine link/node-disjoint paths with QoS routing. However, these papers only consider bandwidth and/or delay as their QoS metrics (Kodialam and Lakshman, 2002; Bejerano et al., 2003). Guo et al. (2003) include reliability as a metric in their 2DP computation. They present a problem, called multiple constrained link-disjoint path pair (MCLPP), to find link/node-disjoint paths in multiple dimensions, and prove the problem is NP-Complete when multiple link metrics are used. However, all these works do not take energy expenditure into account.

Chiaraviglio et al. (2009) address the problem of finding the minimum set of routers and links that can accommodate a given traffic demand. The authors propose a heuristic algorithm that sorts routers according to their energy consumption and to power off redundant ones. Soteriou et al. (2004) studied the impact of energy in interconnection networks, and explored the design space for shutting down links. However, they did not investigate its impact on network performance, such as delay and link utilization. In a different work, Cuomo et al. (2011) propose an algorithm to reroute network flows during off-peak time periods to minimize energy usage. Their method iteratively finds a candidate link to turn-off. It first computes the shortest path for each traffic demand. Then for each router and each link, it calculates the number of times it is being used by shortest paths. Their method selects the least used link as a candidate link. However, they do not consider link disjoint paths.

Fisher et al. (2010) consider each core router connected by multiple physical cables that form one logical bundled link, and propose to turn-off redundant cables. They formulate the problem as an Integer Linear Program (ILP) and propose its linear-programming (LP) relaxation. However, their solution might re-route traffic demands through longer paths that incur delays beyond a tolerable limit, and push each link's utilization above an acceptable upper limit. Further, their approach may re-route a traffic demand through multiple non-link-disjoint paths; thus the solution is not useful for networks with applications that use two link-disjoint paths (Xiong et al., 2009).

Zhang et al. (2010) propose a power-aware traffic engineering algorithm, called GreenTE. They observed that network operators usually provide redundant network links and set a link utilization bound, *e.g.*, 40%, to ensure the network is fault-tolerant against network failures and congestion. However, similar to Fisher et al. (2010), generated multiple paths that are not required to be disjoint.

In Lin et al., (2012) reformulated the Multi-commodity Flow (MCF) model to consider delay and maximum link utilization. As the problem is NP-complete, they propose an efficient and effective heuristic, called MSPF, to solve this power-aware routing problem. They balance the energy and performance constraints, and use a simple mechanism to turn-off unnecessary links. However, their work does not consider disjoint paths. Lin et al. (2013) propose another TE solution, called SSPF, which re-routes each demand using only a single path. They show that single path routing is attractive for its simplicity as well as fast route and significant energy saving. Although, SSPF can be easily merged into existing protocol, *i.e.*, OSPF-TE or ISIS-TE, it routes each demand using only a single path and thus may fail to find a path with sufficient available bandwidth to support multimedia applications such as on-demand video delivery, which requires higher end-to-end bandwidth (Chen et al., 2006).

In contrast to the aforementioned works, this paper proposes a technique, called 2DP-SP, to reduce network energy while satisfying two QoS thresholds: fraction T of 2DPs and maximum link utilization U_T. 2DP-SP offers two key advantages. First, it uses traffic engineering to strike a balance between reducing network