

Exploiting the power of multiplicity: a holistic survey of network-layer multipath

Junaid Qadir, Anwaar Ali, Kok-Lim Alvin Yau, Arjuna Sathiseelan, Jon Crowcroft

Abstract—The Internet is inherently a multipath network—for an underlying network with only a single path connecting various nodes would have been debilitatingly fragile. Unfortunately, traditional Internet technologies have been designed around the restrictive assumption of a single working path between a source and a destination. The lack of native multipath support constrains network performance even as the underlying network is richly connected and has redundant multiple paths. Computer networks can exploit the power of multiplicity—through which a diverse collection of paths is resource pooled as a single resource—to unlock the inherent redundancy of the Internet. This opens up a new vista of opportunities promising increased throughput (through concurrent usage of multiple paths) and increased reliability and fault-tolerance (through the use of multiple paths in backup/ redundant arrangements). There are many emerging trends in networking that signify that the Internet’s future will be unmistakably multipath, including the use of multipath technology in datacenter computing; multi-interface, multi-channel, and multi-antenna trends in wireless; ubiquity of mobile devices that are multi-homed with heterogeneous access networks; and the development and standardization of multipath transport protocols such as MP-TCP.

The aim of this paper is to provide a comprehensive survey of the literature on network-layer multipath solutions. We will present a detailed investigation of two important design issues, namely the *control plane problem* of how to compute and select the routes, and the *data plane problem* of how to split the flow on the computed paths. The main contribution of this paper is a systematic articulation of the main design issues in network-layer multipath routing along with a broad-ranging survey of the vast literature on network-layer multipathing. We also highlight open issues and identify directions for future work.

Index Terms—Multipath, Diversity, Routing, Load Balancing.

I. INTRODUCTION

It is common for communication networks to be built with significant redundancy for the purpose of fault tolerance and reliability. This results in the availability of multiple paths that can be used for connecting network nodes. These multiple paths can potentially be resource pooled for the purpose of routing and forwarding packets, and can be used either simultaneously or in a backup configuration, for increased reliability, end-to-end throughput, network efficiency, and fault tolerance. Unfortunately, despite the widespread path diversity

Junaid Qadir is an Assistant Professor at the Electrical Engineering Department of SECS, NUST. Anwaar Ali is a postgraduate student of Electrical Engineering at the Electrical Engineering Department of SECS, NUST. Kok-Lim Alvin Yau is an Associate Professor at the Faculty of Science and Technology, Sunway University, Malaysia. Arjuna Sathiseelan is a Senior Research Associate at the Computer Laboratory, University of Cambridge. Jon Crowcroft is the Marconi Professor of Networked Systems at the Computer Laboratory of the University of Cambridge.

available on the Internet, traditional Internet technologies have focused mainly on single-path routing for the sake of its simplicity and lower overhead. This has led to the artificial lockup of Internet’s capacity, lower fault tolerance and reliability, and inflexible support for quality of service (QoS). This is even in contrast to telephone networks, which have traditionally adopted multiple-path routing¹ since it leads to better reliability and greater customer satisfaction (due to the lesser probability of call blocking) [2].

Notwithstanding the lack of native multipath support on the current Internet, there is a growing convergence in the Internet community that multipath routing and forwarding, or *multipathing* in short, will play a big role in the future Internet by removing the artificial constraints of single-path routing. As we shall discuss next, single-path based approaches are deficient on many accounts.

A. Deficiencies of Shortest Single-Path Routing

There are two fundamental drawbacks that traditional routing techniques based on shortest single-path routing suffer.

Firstly, it potentially limits the network throughput since only one path is accessible per source-destination pair. It has been shown that single-path routing is constraining optimal routing solutions, and that elegant optimal routing solutions can be developed if multiple paths with flexible splitting are utilized [3]. The problem of joint optimization of routing and congestion control which belongs to the category of the general multi-commodity flow (MCF) problem is considered in [4] for both single-path and multi-path configurations. It is shown that this problem is *non-convex* and NP-hard when the user can utilize a *single path* only. On the other hand, if all available paths can be used with *source-based multipath* routing, the problem is convex and admits an elegant solution. Wang et al. [4] go on to investigate the performance loss of being restricted to a single path. This performance penalty of not using multiple paths is referred to as ‘the cost of not splitting’ [3], [4]. This performance penalty, although hard to determine (being an NP-hard problem), is known to be strictly positive.

Secondly, it is susceptible to routing oscillations since any changes in link lengths can cause abrupt traffic shifts. It

¹Traditional circuit-switched public switched telephone networks (PSTN) use a type of multipath routing known as *hierarchical routing* or *alternative path routing*. Each exchange typically maintains precomputed routing tables with multiple alternative routes to any given destination with a fixed first-choice route, and one or more alternative routes hierarchically arranged for if the first-choice route is blocked, from which the exchange selects the route at the time of the call depending on network congestion level [1].

is well-known in the literature that adaptive, shortest-path routing protocols—especially those that rely on link delays—can result in oscillatory behavior and instability leading to performance degradation [5], [6]. An example of an early delay-based adaptive routing algorithm is the distance-vector ARPANET routing protocol [7] that was later found out to be unstable [8], [9].

Optimal routing based on the rigorous mathematical theory of the MCF problem was proposed to eliminate several disadvantages of single-path routing by optimizing an average delay-like performance measure [1]. *Optimal routing is essentially a multipath routing framework.* In optimal routing, traffic from a source-destination pair is split at strategic points allowing gradual shifting of traffic among alternative paths. Optimal routing aims to send traffic exclusively on paths that are shortest with respect to some link lengths computed on the basis of flows carried on those links. In optimal routing, for low input traffic, a source-destination pair tends to use only one path (which is the fastest in terms of packet transmission time), while for increasing input traffic, additional paths are utilized to avoid overloading the shortest path.

B. Why the Internet’s Future is Unmistakably Multipath?

There are five main trends that indicate that the future Internet is going to be intimately intertwined with multipathing.

1) *Datacenters Trend:* With the increasingly central role of *datacenters (DCs)* in the modern Internet, multipathing is becoming an explicitly designed feature of modern architecture. Unlike the traditional networking topologies that have evolved spontaneously, DCs are highly engineered architectural artifacts that have intrinsically incorporated multipathing for increased performance and reliability.

2) *Wireless Trend:* In wireless communication, multipath plays an important role at various layers of the protocol stack (including the transport, network, and physical layers). Apart from the importance of multipathing at the transport and network layers due to the multihoming of modern mobile devices (such as smartphones), multipath also plays a critical role at the physical layer. Once considered as an antagonist of efficient wireless communications, due to the deleterious effects of multipath fading in which the reflected echoes of the transmission cancel each other destructively affecting successful reception, multipath is now considered advantageous at the physical layer with the recent advances in adaptive signal processing and antenna techniques in multiple input and multiple output (MIMO) technologies. To keep up with the growing bandwidth requirements of modern applications, almost all of the emerging wireless technologies rely on the modern trend of equipping wireless nodes with MIMO antennas. With multiple antennas relying on the extra spatial degrees of freedom, the ability to successfully use multiple paths has become critical to the success of modern wireless networks.

3) *Multiple Wireless NICs Trend:* Modern wireless network technologies have also made it feasible to equip network nodes with multiple network interface cards (NICs) tuned

to orthogonal channels. The ready availability of commodity NICs, and the availability of modern networks to support multiple interfaces, has opened up the possibility of the successful exploitation of multipath technology to improve network performance.

4) *Multihoming with Heterogeneous Technologies:* Multihoming is a technique that has been traditionally used by server farms on the Internet for reliability and high performance. Analysis of multihoming has shown significant performance benefits that accrue with multihoming, with a careful selection of three or more Internet service providers (ISPs) providing performance improvement, in terms of transfer throughput and round-trip-time (RTT), of upto 25% [10]. While traditionally multihoming has been relevant only for critical servers, it has now become a mainstream with billions of mobile devices now multihomed with heterogeneous access technologies (such as Wi-Fi/ 3G). A wireless device (such as a smartphone) can exploit its interfaces to heterogeneous technologies in order to improve its reliability, throughput, and fault tolerance. In recent times, there have been multipath transport-layer initiatives that are built specifically to incorporate such multihoming. In particular, the standardization of MP-TCP [11] in RFC6824 has led to mainstream applications of multipath: e.g., the Siri intelligent personal assistant on Apple’s iOS7 uses MP-TCP to increase efficiency and reliability.

5) *Path Diversity in ISPs:* Teixeira et al. have studied the availability of *path diversity* (the availability of multiple possible paths between two communicating entities) in ISPs [12]. It was shown that a typical large ISP² has a significant path diversity. The abundant availability of path diversity does not necessarily imply that the best paths are adopted though—in fact, statistics show that at any time 30 to 80 percent of the Internet traffic chooses a single suboptimal path even though alternative low-loss lesser-delay paths exist [14]. With the inherent path diversity in ISPs, there is a strong interest in the service provider community to exploit multiple paths [15].

C. Benefits of Multipathing

We provide a description of some of the major applications and features of multipathing below:

1) *Resource Pooling:* Resource pooling is a broad Internet architectural design principle—focusing on improving resource efficiency by presenting a collection of resources as a single pooled resource—that has many manifestations. For example, some of the most disruptive innovations on the Internet—packet switching, content-delivery networks (CDN), peer-to-peer networks (P2P), cloud computing—are based on resource pooling [16]. The packet switching technology, which acts as a foundation of the Internet, is essentially a resource pooling technique. To put things in perspective, circuit switching techniques use fixed multiplexing schemes in which dedicated and isolated non-pooled circuits are utilized. In contrast, statistical multiplexing—the underlying technique

²Empirical calculations were performed on a tier-1 Sprint ISP network using heuristics for inferring path diversity based on active measurements using the *Rocketfuel* [13] toolset.

of packet switching—achieves resource pooling by allowing a burst of traffic on a single circuit to use spare capacity on other circuits. CDNs and P2P networks utilize resource pooling to ensure efficient network operations. P2P protocols such as BitTorrent use swarming downloads that let receivers pool together multiple peers as a data source thereby using multipath to pool the network paths to these peers. Cloud computing relies on pooled resources gathered at central DCs to support the computing requirements of multiple tenants through virtualization. Multipathing has been envisioned as ‘*Packet Switching 2.0*’ [16] since it allows an analogous resource pooling benefit at the network layer. In particular, multipath routing can pool together a number of distinct links to provide the abstraction of a single unified network resource. With the abstraction of a single resource pool, the chances of a demand being refused due to one’s resources being utilized reduces.

2) *Load Balancing*: With resource pooling, a collection of resources acts as a single virtual resource. The biggest promise of resource pooling is load balancing that allows congestion control techniques for diffusing congestion over resources by equitably distributing the load amongst the elements of a resource pool. This helps to avoid situations where a certain network resource unnecessarily acts as a bottleneck although the network has alternative resources that can help to relieve the congestion [17], [18]. Load balancing can be done over resources (e.g., the load balancing over multiple access links and paths as done in multihoming and multipath transport respectively) or can be temporal (i.e., the load is balanced between peak and non-peak hours to exploit the typical diurnal network traffic pattern). Load balancing over time shifts the transmission of data from peak time to off-peak time and can help in reducing costs such as the commonly used 95 percentile pricing scheme [19].

3) *Efficiency*: Another major benefit of resource pooling through multiple paths is efficiency. It is well known in the literature that the Internet traffic is bursty with a pronounced correlation with a large variance of the traffic volume over long time periods [20], [21]. Allocating fixed amount of resources to individual circuits is inefficient due to the bursty nature of the Internet traffic, which can result in largely idle resources or resources that are insufficiently provisioned. One important benefit associated with resource pooling is that it leads to increased efficiency. Resource pooling, or statistical multiplexing of shared resources, suits the Internet traffic well. This is because the law of averaging (or the law of large numbers) does not apply to individual flows. Resource pooling is useful since it exploits the statistical regularity of the aggregate of individual bursty circuits. This results in efficient utilization of resources without extensive over-provisioning.

4) *Reliability and Fault Tolerance*: With the migration of critical businesses to the Internet, the reliability and availability of the Internet has become extremely important. Fault tolerance mechanisms for networks without multipath support are unwieldy with a link/ node failure potentially resulting in significant recovery times. The resource pooling benefit of multipathing allows the underlying network to gracefully

handle loss of capacity or the failure of individual links/ paths through diffusing this anomaly by shifting traffic to the working members of the resource pool. This allows enhanced fault tolerance and helps to build more resilient networks.

Network communication can utilize multipathing to enhance its *reliability* at various layers [22]. At the network layer, we can have multihomed servers that have different ISP connections; another option is to utilize backup paths with traffic being switched to an alternative path in the case of failure of the primary path. At the physical layer, techniques such as antenna diversity can be used to combat fading effects in wireless communication.

Multipathing can also be used to facilitate network security. More specifically, the use of multiple radio interfaces has been proposed for achieving greater data *confidentiality* by encrypting data and splitting it into multiple parts that are then transmitted using different physical channels [23]. Multipathing can also be used to mitigate malicious adversaries that drop packets intentionally on a path by circumventing the adversary using an alternative disjoint path if available.

5) *Higher Throughput*: A great number of modern high-performance Internet technologies utilize multipathing in one way or another. At the physical layer, MIMO technology is revolutionizing the wireless industry by promising significantly enhanced communication rates. The technique of space-division multiple access (SDMA), used by multi-user MIMO (MU-MIMO), allows multiple transmitters to send—and multiple receivers to receive—separate signals simultaneously over the same band. The improvement offered through multipathing at the link layer and the network layer is dramatic in environments where the traffic is concentrated on only a few links while other perfectly fine links are not used at all (e.g., in bridged networks where the spanning tree protocol is used). Other high-performance proposals include multipath-TCP (MP-TCP) at the transport layer that allows the pooling of multiple links into a pooled connection to support higher speed: for example, there has been a recent demonstration of a transfer achieving a goodput of 50 Gbps using MP-TCP with nodes equipped with 6-10 Gbps commodity NICs [24].

6) *Flexibility due to Pluralism and Multiplicity*: In recent literature [25], the metaphors of “the hedgehog and the fox”³—popularized by Isaiah Berlin in an essay of the same name—are being used to highlight the importance of flexibility that accrues with pluralism. For example, it has been shown that a forecaster who systematically entertains a variety of possibilities outperforms a rigid one-dimensional forecaster fixated on a single potential interpretation [26]. There is also a wealth of literature that supports the claim of “*the wisdom of crowds*” [27].

Many of the benefits of multipathing are brought about by the ‘*diversities*’ it offers. The benefits of diversity have been established in a broad spectrum of fields ranging from communications, biology, economics, to sociology. Intuitively speaking, diversity is useful in a variety of fields since it allows hedging in uncertain situations to minimize risk and thereby

³Based on a reported saying from the ancient Greek poet Archilochus who had said: “*the fox knows many things but the hedgehog knows one big thing*”.

downplays the negative effects of uncertainty.

7) *Support for Multimedia Applications*: The modern rise of high bandwidth multimedia applications can benefit significantly from the resilience offered by multipath techniques. While the traditional packet-switched Internet is inherently a best-effort network where a single path is mostly in use, multiple paths can coexist at the same time. As scarce bandwidth and packet loss are more significant problems for today's streaming media because of its increased transmission requirements, concurrent transmission of multimedia applications over multiple paths can be adopted to resolve the aforementioned problems.

8) *Role in the Future Internet Architectures*: It is anticipated that multipathing will play a major role in the future Internet architectures and next-generation wireless and mobile networks. It is already typical for mobile phones to be multihomed with heterogeneous connectivity (e.g., Wi-Fi and 3G). The upcoming 5G standard (expected to come to fruition around 2020) is expected to integrate and resource pool a number of diverse access technologies [28] with multipathing playing its part in seamless handovers and interworking of heterogeneous technologies. With the resource pooling of cellular and Wi-Fi networks, the possibilities of Wi-Fi offloading and cellular onloading emerge. With Wi-Fi offloading, the performance and scalability of mobile networks can be improved by offloading data traffic to the Wi-Fi networks [29]. With cellular onloading, on the other hand, home broadband/Wi-Fi connections are augmented with cellular data services for bandwidth aggregation [30]. Apart from wireless networks, multipathing is also useful in emerging networking configurations (such as DCs and clouds) and in other upcoming Internet architectures such as P2P, CDN, software-defined network (SDN), and content-centric network (CCN) [31], [32], [33].

9) *Context-aware Services*: It should be noted that different networking applications have different needs. In particular, various application configurations can coexist: e.g., the first application prefers a low-delay path, the second prefers a high-throughput connection, and another prefers a reliable connection. Using multipathing, context-aware services are in line with the application requirements.

D. Challenges in Implementing Multipath Routing

Even though the future of the Internet promises to be intertwined with multipath routing, there are numerous challenges that need to be surmounted. A major challenge in implementing Internet-wide multipath routing is to achieve scalability due to the resulting considerable computational and storage overhead. Specifically, multipath routing entails significant overhead both in the control as well as the data plane. Control plane overhead refers to the need of additional bandwidth and processing resources for exchanging path information, while the data plane overhead refers to the large forwarding table and the higher memory requirement associated with forwarding entries corresponding to an increased number of paths. Another stumbling block is the lack of complete control of any organizations over the entire end-to-end path. We will

discuss more open and current research issues later in the paper in section IV.

E. Organization and Contributions of This Paper

We organize this paper around five fundamental design issues that relate to network-layer multipathing. These designs issues are presented below in the form of design questions:

- Q1. *Route computation*: How to compute the multiple paths to be used?
- Q2. *Route metric*: What routing metrics should be used to construct the 'best' multiple paths?
- Q3. *Load balancing/ Flow splitting*: How to have a traffic flow use the constructed paths?
- Q4. *Number of paths*: How many paths should be constructed and used?
- Q5. *Concurrent usage*: Should the multiple paths be used concurrently or as backups?

The overall organization of this paper is illustrated in Figure 1. The main questions that we analyze, and present solutions for, have already been presented. In Section II, we discuss various techniques that can be used for the construction and selection of multiple paths (Q1) and also the metrics that are used for facilitating route construction (Q2). In Section III, we tackle the issue of how to use the multiple constructed paths (Q3). In particular, we discuss issues relating to load balancing, flow splitting, and multipath congestion control. We discuss the choice of how many paths to use (Q4) in Section III-B, and whether to adopt backup or concurrent multipathing (Q5) in Section III-B. In Section IV, we present potential open research issues in the field of multipathing at the network layer followed by our main conclusions in Section V.

The main contribution of this work is a comprehensive survey of network-layer multipath routing in which apart from our network-layer focus we also provide a balanced and holistic cross-layer coverage of multipathing (recognizing the fact that many important issues in multipathing—such as resource pooling, congestion control, flow splitting—involve cross-layer interaction. While there exist quite a few survey papers on multipathing protocols, these papers are typically confined to particular networking configurations (such as wireless sensor networks [34], [35], [36], mobile ad hoc networks [37], [38], [39], etc.) We will present a more general paper that will comprehensively describe all the important multipath proposals/ ideas and will also tie them up with new advances in transport-layer multipathing. With the increasing importance of multipathing on the Internet, there is a strong need to have such a paper to benefit practitioners and researchers in the networking community. This paper is timely due to the huge interest in this domain (emerging from new research/standardization work on multipath protocols, such as MP-TCP). The success of higher-layer multipath protocols depends critically on support from lower-layer protocols (in particular, the network-layer protocol).

II. HOW TO CONSTRUCT ROUTES?

On the Internet, the task of the construction of routes is handled by the *control plane*, while the task of forwarding

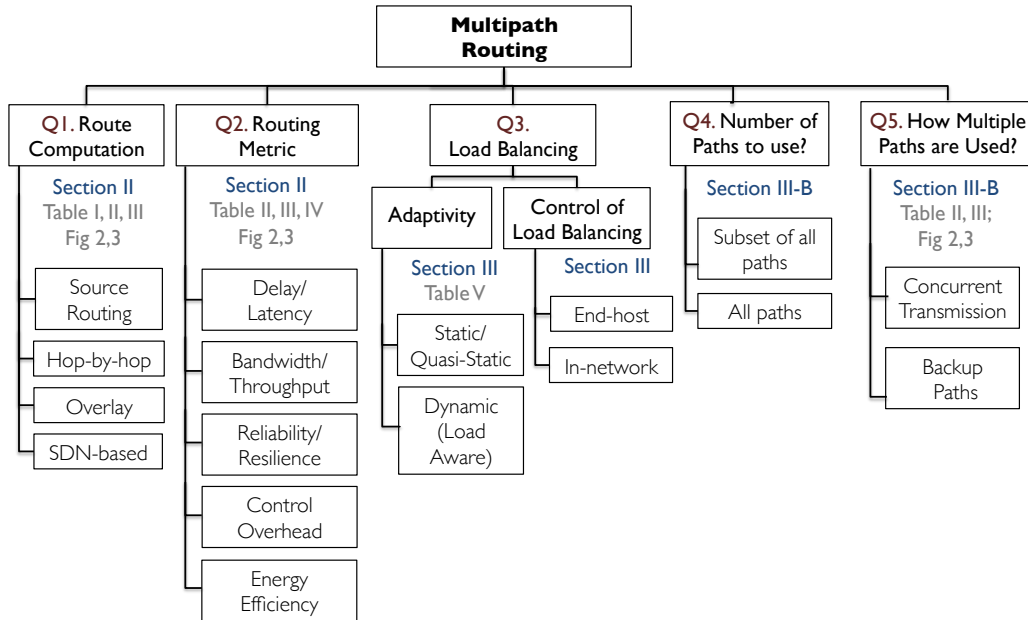


Fig. 1: Various network-layer multipathing design choices (and their coverage in this paper).

the packets based on the computed routes is performed by the Internet’s *data plane*. The Internet’s control plane is typically managed through dynamic routing protocols. Generally speaking, it is desirable for a routing protocol to (i) choose ‘optimal’ paths, (ii) be robust to dynamic network topology in order to avoid instability in the form of oscillatory route flaps, and (iii) minimize routing overheads in terms of memory required, messages exchanged, etc. There are several choices in the design of multipath routing protocols—including centralized or distributed routing, static or dynamic routing, source-based or hop-by-hop routing—corresponding to different tradeoffs with respect to complexity and resource usage.

In this section, we focus on the various techniques proposed in the literature for the computation of routes. These techniques include “source routing” in which the end-to-end route is computed and specified by a source and “hop-by-hop routing” in which network routers exchange control packets to establish state that is then used for hop-by-hop forwarding. Broadly speaking, the various multipath routing techniques can be classified into four categories of: (i) source routing, (ii) hop-by-hop routing, (iii) overlay routing, and (iv) SDN-based routing. In the following subsections, we discuss each of these categories in this order. A summary of the multipath proposals presented in this section—along with a description of their route computation technique, routing metrics adopted, and specific networking configuration assumed—is presented in Tables I, II, and III.

A. Source Routing

In many cases, the network is not well placed to know about the optimality of a route. The source host/ user, on the other hand, knows about the application’s unique requirements and can judge about the appropriateness of a route. The source is also best suited to deal with any failures that may arise along a route. This motivates *source routing* in which the source

host, or edge router, defines the end-to-end path to be used for forwarding packets with the potential to improve network performance, reliability and user satisfaction.

In the source routing paradigm, the source explicitly embeds information in a data packet about the path that the packet will traverse before reaching its destination [47]. Such a scheme proceeds typically by first discovering routes through the exchange of control messages and then embedding the discovered route information in the data packet’s header. The routers along the explicitly defined path then forward the data packet using the information stored in the data packet’s header. Although source routing provides maximum flexibility, it is important to point out some difficult challenges that it faces [93] including (i) limited control of intermediate autonomous system (AS); (ii) scalability problems due to the need of having the source to compute and define routes; and (iii) efficiency and stability problems that may arise from selfish source routing.

In the next two subsections, we will describe source routing proposals, initially discussing general solutions (developed mainly for wired networks) in Section II-A1, and then discussing source routed multipath solutions specific to wireless networking in Section II-A2.

1) Source Routed Multipathing (General/ Wired Solutions):

To accompany this section, we present a comprehensive listing of general multipathing solutions (designed primarily for wired networks) in Table II. This table serves as a convenient summary to complement the description of various hop-by-hop multipathing routing solutions in this section. In addition, the table also includes information about works that have used other routing schemes (source routing, overlay routing, and SDN-based routing), and we will be revisiting this table later in this paper. We note here that the classification is approximate since not all works can be cleanly pigeonholed into only one

TABLE I: Frameworks for Determining Multiple Paths (Route Selection)

<i>Project</i>	<i>Year</i>	<i>Description</i>
<i>Source Routing</i>		
Split Multipath Routing (SMR) [40], [41]	1999, 2000	Applies a probing-based mechanism to achieve a balanced load over multiple constructed paths.
Multipath MPLS [42]	2001	Constructs multiple paths in polynomial time subject to a constraint (i.e., number of hops and paths).
Feedback based routing (FBR) [43]	2003	Performs routing decision at the access router with two routes computed for each network prefix using dynamic quality monitoring based on structural information disseminated by the transit routers .
BANANAS [44]	2003	Enables multipath routing in conventional routing protocols (e.g., OSPF and BGP) using global identifiers for abstraction.
Single Hop Source Routing (SOSR) [45]	2004	Uses randomly chosen intermediaries to route packets indirectly towards the eventual destinations in order to recover from failures which occur near the destinations.
SPREAD [46]	2004	Divides a secret message into multiple parts and sends them over multiple independent paths towards the destination in order to improve secrecy and security.
Routing Deflections [47]	2006	Provides path diversity by allowing routers to redirect packets based on tagged information provided by end hosts, subject to the compliant of the ISP routing policy.
Joint Multi-channel Multi-path Control (JMM) [48]	2007	Achieves a balanced distribution of the contending input traffic at different times over different channels and different paths through cross-layer interaction between link and network layers.
New Internet Routing Architecture (NIRA) [49]	2007	Provides path diversity by enabling a packet, in a source controlled manner, to switch among multiple ISPs enroute to its destination.
Path Splicing [50]	2008	Constructs a path by combining slices of multiple routing trees, which serve as backups paths, en route to an eventual destination to improve reliability.
Pathlet Routing [51]	2009	Performs routing deflection at the domain level of a hierarchical network. Each domain advertises a few paths or <i>pathlets</i> used by the a source from another domain.
Slick Packets [52]	2011	Embeds alternate route information is embedded in the packet's header in a source controlled fashion to avoid failures. This scheme requires minimal network routing state because of the source routing feature.
AMIR [53]	2012	Establishes multiple paths around the primary paths in hierarchical networks.
<i>Hop-by-Hop Routing</i>		
Dispersity routing [54]	1975	Sends distinct data packets over multiple paths for load balancing, or sends duplicated data packets over multiple paths for improved reliability and resiliency.
Geographical Power Efficient Routing (GPER) [55]	2004	Uses probabilistic multipathing to achieve path diversity. The nodes are rendered to choose their transmission ranges to conserve energy.
Loop-free Invariant (LFI)-based Approaches [56], [57], [58]	1999, 2001	Incorporates multipath routing in traditional link state-based and distance vector-based algorithms to find loop-free multiple paths using LFI condition. <i>Diffusing computation</i> is used in [58].
Hashing-based ECMP [59]	2000	Selects paths from a set of <i>equal cost multiple paths</i> in order to split flows for load balancing. It comprises of direct hashing and table-based hashing methods.
Srinivas and Modiano [60]	2003	Uses 2-link-disjoint and k-node-disjoint algorithms to conserve energy.
Packet-caching [61], [62]	2003, 2010	Circumvents route breakdowns by caching/ storing (a fixed number of) data packets along with alternate route information at nodes in a distributive fashion.
Interference Minimization Approaches (M-MPR [63], [64]; Maimour [65])	2003, 2008	Constructs multiple paths in an incremental fashion keeping in view the congestion and interference effects (using the interpath interference and disjointness metrics). M-MPR incorporates FEC and selective forwarding.
AODV-enhanced Routing Schemes (AOMDV [66]; Ye et al. [67]; MP-AODV [68]; NDMR [69])	2001, 2003, 2004, 2004	AOMDV establishes multiple link-disjoint paths in an on-demand manner; Ye et al. use reliable nodes, in terms of energy resources, to establish multiple paths; MP-AODV selects primary paths using received signal strength, and alternative paths through heartbeat message exchanges; NDMR incorporates multipathing into traditional AODV and DSR, and uses alternative paths as backups.
Biologically Inspired Routing [70], [71]	2005	Uses biologically inspired approaches (e.g., ant colony optimization) to establish and maintain multiple paths.
Security-Enhanced Approaches (SEIF, SMBR, H-SPREAD)	2006, 2011	Enables nodes to take independent decisions to address joint fault and intrusion tolerance.
Enhanced BGP [72], [73]	2007, 2009	Extends the traditional BGP routing protocol to establish and maintain loop-free multiple paths in the face of multiple failures in hierarchical networks.
Maximum Alternative Routing Algorithm (MARA) [74]	2009	Constructs directed acyclic graphs (DAGs) that (i) maximize the min. connectivity; (ii) maximize the min. max. flows; and (iii) maximize the min. max. flows as an extension of shortest-path routing.
Huang and Fang [75]	2009	Selects multiple parent nodes in an opportunistic manner to establish multiple paths even though nodes may have unknown speeds in vehicular ad-hoc networks.
Network coding based reliable disjoint and braided multipath routing protocol [76]	2010	Uses braided and disjoint multiple paths in a distributed fashion in conjunction with network coding for reliability and energy conservation purposes.
Yet another multipath routing (YAMR) [77]	2010	Establishes multiple paths, and provides concurrent transmission with minimal overhead in hierarchical networks.
Beltagy et al. [78]	2011	Uses multiple paths to reduce negative effects from and to licensed users in cognitive radio networks.
<i>Overlay Routing</i>		
Deroute [14]	1999	Proposes an edge-based routing scheme for the Detour overlay network comprising geographically distributed routes for better end-to-end performance.
Resilient Overlay Network (RON) [79]	2001	Focuses on equipping end hosts and applications with the ability to switch between direct routing (through the underlying network) or indirect routing (using the overlay network).
Multipath Video Streaming on Overlay Networks [80]	2004	Proposes a link correlation-based QoS metric along with path correlation model for multipathing using an overlay network for efficient video streaming.
Multipath multiple description video streaming [81]	2005	Selects a set of multiple paths based on <i>multiple description coding</i> models.
BARON [82]	2008	Proposes a bandwidth-aware routing scheme for overlay networks to meet application bandwidth requirements and improve the overall throughput performance.
<i>SDN-Based</i>		
Hedera [83]	2010	Proposes an efficient, scalable, adaptive and central scheduling method, with a universal scope of active flows in view, for multistage switching processes in datacenters.
Micro-TE [84]	2011	Implements central controllers using OpenFlow, as well as uses short-term traffic predictability and multipath routing, to schedule flows in a datacenter environment.
Rain Man [85]	2012	Proposes a layer two (L2) multipath packet forwarding algorithm for a SDN-based firmware architecture for datacenters.
B4 (Google) [86]	2013	Proposes a centralized TE solution that utilizes OpenFlow/ SDN principles along with multipath forwarding to improve network efficiency (i.e., link utilization and average network utilization).
Wireless Mesh Software-defined Networks (wmSDN) [87]	2013	Achieves traffic balancing and selects gateways in an intelligent round robin fashion for Internet connectivity.
Software Defined VANETs [88]	2014	Establishes multiple paths to address route failures, as well as for the transmission of high priority flows, in VANETs.

of our broadly defined categories.

Due to the lack of flexible support for traffic engineering in IP-based networks, the multiprotocol label switching (MPLS), with its rich support for traffic engineering (TE), has become ubiquitously deployed in ISPs. A lot of work has been done on

efficient optimized multipath TE. Lee et al. [121] proposed a constrained multipath TE scheme for MPLS-based networks to help find optimal paths that meet the provided traffic demand as well as the constraints (such as maximum hop count and node/ link preferences). The proposed scheme calculated the

TABLE II: Summary of General Multipath Routing Proposals (Designed Mainly For *Wired Networks*)

Work	Year	Route Construction	Paths	Metric(s)	Description
Dispersity Routing [54]	1975	Hop-by-Hop	Concurrent	Delay, Reliability, Load Balancing, Link Utilization	Disperses traffic over multiple paths either in a redundant or non-redundant manner.
OSPF-OMP [89], [90]	1999	Hop-by-Hop	Concurrent	Delay, Load Balancing	Provides a hashing-based OSPF extension to support optimized multipathing for load balancing.
ECMP [91]	2000	Hop-by-Hop	Concurrent	Delay, Load Balancing	Forwards packets based on a CRC16 hash to one of the nodes forming equal-cost paths towards a destination.
Multipath MPLS [42]	2001	Source Routing	Concurrent	Link Utilization, Load Balancing	Constructs multiple paths in polynomial time subject to a constraint (i.e., number of hops and paths).
Feedback-based (FBR) [43]	2003	Source Routing	Backup	Delay	Enables an access router to compute two routes for each network prefix based on the structural information disseminated by transit routers.
BANANAS [44]	2003	Source Routing	Concurrent	Resource Utilization	Provides a framework that creates <i>global identifier</i> to enable multipathing in the Internet.
SOSR [45]	2004	Source Routing	Backup	Overhead, Reliability	Presents a <i>single hop source routing (SOSR)</i> to avoid failures which occur near the destinations in the Internet.
MONET [92]	2005	Overlay Routing	Backup	Overhead, Delay, Reliability	Provides multiple backup paths through proxies to websites for enhanced web availability.
Routing Deflections [47]	2006	Source Routing	Concurrent	Overhead, Reliability	Provides an ISP that constructs alternate paths or <i>deflections</i> based on the <i>tags</i> that users put in their packets.
MIRO [93]	2006	Interdomain Negotiations	Backup	Reliability	Provides interdomain negotiations while establishing paths by improving upon BGP and source routing.
R-BGP [72]	2007	Hop-by-Hop	Backup	Reliability	Provides alternate paths to overcome transient failures by modifying BGP.
NIRA [49]	2007	Source Routing	Backup	Reliability, Overhead	Enables a user to choose the sequence of ISPs for its packets in order to switch among multiple routes.
Path Splicing [50]	2008	Source Routing	Backup	Reliability	Allows <i>slices</i> of different routing trees to be spliced together for reliable delivery of a packet towards a destination.
D-BGP and B-BGP [94]	2009	Hop-by-Hop	Backup	Overhead, Reliability	Provides a path diversity-aware routing protocols in which an advertisement method is proposed for multiple paths.
Pathlet Routing [51]	2009	Source Routing	Backup	Reliability	Allows a source to concatenate and advertise segments of paths or <i>pathlets</i> in packets while constructing an end-to-end path.
Multipath BGP [73]	2009	Hop-by-Hop	Concurrent	Resource Utilization, Reliability, Overhead	Uses multiple paths in a concurrent fashion and preserves loop-freeness by modifying BGP.
YAMR [77]	2010	Hop-by-Hop	Backup	Reliability, Overhead	Constructs a set of multiple paths to avoid (at least) an interdomain link failure.
Multi-path BGP [95]	2011	Hop-by-Hop	Concurrent	Resource Utilization	Proposes two multipathing approaches at the interdomain level.
AMIR [53]	2012	Source Routing	Concurrent	Reliability	Constructs alternative paths around the primary path using topology information from neighboring ASes.

flow split ratio for the multiple paths. The evaluation of the proposed scheme showed superior results in utilization, traffic volume, and overhead compared to the conventional shortest path routing scheme. In another work, Seok et al. [42] proposed to integrate MPLS traffic engineering with source multipath routing subject to a constraint on the total number of hops and the total number of paths. The presented algorithm calculated the flow split ratio by using hashing over the calculated multiple paths as per the traffic demands dictated by an ingress-egress pair of a network. A dynamic traffic engineering approach is proposed to tackle the issue of congestion as the heuristic presented in this work protects a link from the overwhelming effect of huge traffic demands by minimizing the maximum utilization of a link and using the aspect of multipathing for load balancing.

Approaches to Avoid Failure/ Performance Bottlenecks

It was shown in [14] that upto 80% of the default routing paths on the Internet are suboptimal with an alternative route offering potentially lower loss rate. This has motivated the development of numerous *source-routed* routing protocols that can help route around the performance bottleneck/ failed entity.

Gummadi et al. [45] proposed a source routing-based approach named “*scalable one-hop source routing*” (SOSR) to recover from Internet path failures by routing indirectly through randomly chosen intermediaries. It has been shown in the previous work that last-hop and end-host failures account for a disproportionately high number of path failures. Such failures are especially hard to recover as there is a limited number of network elements to route through as the route nears the destination. SOSR recovers from such failures by routing packets indirectly using a small set of *randomly chosen* network intermediaries. When a network failure occurs, then

the source probes the destination through the randomly chosen set of intermediary nodes. If such a probe is successful, then the data is routed via this indirect route; otherwise, it is concluded that the destination is now unreachable. In this technique, since the choice of the intermediary router is made randomly from a preselected set, the continuous maintenance of backup routes is avoided. This allows SOSR to outperform other overlay or backup path techniques that require continuous path maintenance and are more expensive in terms of control overhead.

Using the concept of *path splicing* [50]—which can be applied in both interdomain and intradomain routing scenarios—a packet can traverse across segments or *slices* of different routing trees (based on the discretion of the end hosts). In path splicing, multiple routing trees are established with a packet able to switch from one routing tree to another at any hop/ edge. This concept has been shown to improve network resilience.

Interdomain Source-Routed Multipath Routing

To manage a global-scale routing, the approach adopted on the Internet is to separate routing into intradomain routing and interdomain routing, where the former focuses on “optimal” routing (defined in the terms of some optimization metrics) within a single AS, while the latter focuses on ensuring loop-free policy compliant routing between multiple ASes. Since most traffic on the Internet traverses across multiple domains, interdomain multipathing is of great practical importance. The defacto interdomain routing protocol of the Internet is *Border Gateway Protocol (BGP)*. BGP allows different ASes to exchange reachability information that is then used by BGP for computing loop-free policy-compliant routing.

Although BGP is well entrenched in the Internet ecosystem, it is widely known that most routes chosen by BGP are not

TABLE III: Summary of Multipath Routing Proposals for *Wireless Networks*

Work	Year	Route Construction	Paths	Metric(s)	Description
Wireless Sensor Networks (WSNs)					
Disjoint and Braided [96]	2001	Hop-by-Hop	Backup	Energy, Reliability, Overhead	Proposes a reliable, non-disjoint or <i>braided multipathing</i> scheme that performs 50% better than disjoint multipathing.
Selective Forwarding [63], [64]	2003	Hop-by-Hop	Concurrent	Energy, Throughput	Proposes a meshed multipath routing protocol with selective forwarding.
MMSPEED [97]	2006	Overlay Routing	Concurrent	Delay, Reliability	Proposes a probabilistic multipath selection scheme under reliability and QoS constraints (i.e., delay).
H-SPREAD [98]	2006	Source Routing	Concurrent	Reliability	Proposes a hybrid multipath method where node/link-disjoint paths are maintained on each sensor node.
EBMR [99]	2006	Deterministic (Defined by BS)	Backup	Energy	Enables a BS to discover, maintain and switch among multiple paths keeping in view the nodal residual energy.
Totally Disjoint Paths [100]	2006	Hop-by-Hop	Concurrent	Interference, Throughput	Proposes and studies a 2-path routing scheme in a single and a multiple source-destination pair scenarios.
GPFR [101]	2007	Hop-by-Hop	Backup	Energy	Enables nodes to decide their transmission ranges distributively for power efficiency using probabilistic multipathing.
DMPR [102]	2007	Hop-by-Hop	Concurrent	Load Balancing, Throughput	Constructs two (colored) trees, which are rooted at the sink node. Two paths are formed from a source via each tree, which is link/node disjoint.
SEEM [103]	2007	Deterministic (Defined by BS)	Backup	Energy, Throughput, Overhead, Reliability	Enables BS to discover and maintain multiple routes, as well as switch among them, in order to conserve energy and enhance security.
Multicast Multipath [104]	2007	Hop-by-Hop	Concurrent	Delay, Bandwidth	Proposes three tree-based on-demand and distributed multipath routing methods under multicast QoS constraints.
Ming et al. [105]	2007	Hop-by-Hop	Concurrent	Energy, Delay, Overhead	Presents a scalable, distributed, reliable and node-disjoint multipath routing protocol to achieve efficient load balancing.
Maximally Radio Disjoint Routing [65]	2008	Hop-by-Hop	Concurrent	Energy, Bandwidth, Interference	Constructs multiple paths in an incremental way to reduce interference in multimedia WSNs.
Multiconstrained QoS Routing [106]	2008	Hop-by-Hop	Concurrent	Delay, Reliability	Constructs multiple paths based on a probabilistic link-state method under delay and reliability QoS constraints.
REER [107]	2009	Hop-by-Hop	Both	Energy, Delay, Reliability	Uses backup and concurrent multipathing approaches to conserve energy and enhance robustness.
NC-RMR [76]	2010	Hop-by-Hop	Concurrent	Energy, Reliability	Uses both disjoint and braided multiple paths, with network coding, in a per-hop fashion for energy conservation and reliability enhancement.
SEIF [108]	2011	Hop-by-Hop	Concurrent	Energy, Reliability	Proposes a distributed, BS independent and one-way hash chains for multipathing.
IM2PR [109]	2014	Hop-by-Hop	Concurrent	Energy, Reliability	Presents an interference-minimized multipath routing protocol (IM2RP) that aims to discover minimum interfering paths.
Mobile ad-hoc networks (MANETs)					
On-demand [40]	1999	Source Routing	Backup	Overhead	Modifies DSR to support multipathing.
MSR [41]	2000	Source Routing	Concurrent	Load Balancing	Uses a probing-based mechanism to provide multipath support in DSR for load balancing.
SMR [110]	2001	Source Routing	Concurrent	Delay, Overhead, Load Balancing	Constructs two maximally disjoint paths, with one being the smallest delay path, in an on-demand fashion.
AOMDV [66]	2001	Hop-by-Hop	Backup	Reliability	Modifies AODV to construct multiple link-disjoint paths using the advertised hop count information.
MP-DSR [111]	2001	Source Routing	Concurrent	Bandwidth, Reliability	Provides a QoS-aware multi-path dynamic source routing protocol for MANETs.
Ye et al. [67]	2003	Hop-by-Hop	Concurrent	Reliability	Modifies AODV to support multiple node-disjoint paths by introducing a few <i>controlled reliable nodes</i> in the topology.
CHAMP [61]	2003	Hop-by-Hop	Concurrent	Reliability, Resource Utilization	Circumvent frequent link breakdowns by caching (a few) data packets along with alternate route information at nodes distributively.
Min. Energy Paths [60]	2003	Hop-by-Hop	Concurrent	Energy	Provides two algorithms to construct 2-link-disjoint and k-node-disjoint paths to conserve energy in dense and sparse networks, respectively.
SPREAD [46]	2004	Source Routing	Concurrent	Reliability	Sends data in the form of <i>shares</i> over multiple independent paths to ensure confidentiality.
Serial Multiple Disjoint Trees	2004	Source Routing	Concurrent	Robustness	Provides real-time efficient video streaming using multiple paths for unicast and multicast traffics.
Multicast Routing (Serial MDTMR) [112]					
MP-AODV [68]	2004	Hop-by-Hop	Backup	Reliability	Provides adaptive multipathing in AODV based on signal strength received at nodes.
NDMR [69]	2004	Hop-by-Hop	Backup	Delay, Reliability	Constructs node-disjoint multiple paths for a source-destination pair.
Chen et al. [113]	2005	Hop-by-Hop	Concurrent	Delay, Reliability, Load Balancing	Presents an analytical model for multipath routing in MANETs for optimal load balancing, as well as to address delay and reliability issues.
Multipath OLSR [114]	2005	Hop-by-Hop	Concurrent	Energy, Load Balancing	Modifies OLSR to support multipathing.
AntHocNet [70], [71]	2005	Hop-by-Hop	Concurrent	Delay, Reliability, Jitter	Provides a hybrid multipathing algorithm with reactive path setup, as well as proactive path probing and maintenance, based on an <i>ant-colony optimization framework</i> .
SDMR [115]	2008	Source Routing	Concurrent	Load balancing, Bandwidth, Interference, Path Closeness	Enables a source node to calculate multiple spatially disjoint paths based the information extracted from route reply packets from the destination node.
Multipath Routing Reliability [116]	2008	Overlay Routing	Backup	Reliability	Analyzes the reliability of multipath routing protocols based on an overlay graph concept framework.
RAMP [62]	2010	Hop-by-Hop	Concurrent	Delay	Jointly addresses the issues of multipath routing and deferential delay.
Mobile Cloud [117]	2014	SDN-based	Concurrent	Reliability	Uses a central controller in a SDN fashion to perform multipath routing in MANET.
Wireless Mesh Networks (WMNs)					
Weighted Interference Multipath (WIM) [118]	2007	Source Routing	Concurrent	Interference, Delay, Reliability	Selects multipath based on <i>weighted mutual-interference (WIM)</i> metric in a multi-channel multi-radio environment.
Joint Multichannel Multipath (JMM) [48]	2007	Source Routing	Concurrent	Throughput	Splits traffic among different channels at different times over multiple paths.
wmSDN [87]	2013	SDN-based	Backup	Load Balancing, Resource Utilization	Selects Internet gateways of a mesh network in a round robin fashion to balance outgoing traffics using a central controller.
Cognitive Radio Networks (CRNs)					
Beltagy et al. [78]	2011	Hop-by-Hop	Backup	Reliability, Path Closeness	Constructs multiple routes to minimize the effects of PU activity based on a <i>closeness</i> metric.
Vehicular Ad-hoc Networks (VANETs)					
Huang et al. [75]	2009	Hop-by-Hop	Concurrent	Delay, Interference, Reliability	Constructs node-disjoint multiple paths based on the mutual interference of paths.
LIAITHON [119]	2012	Hop-by-Hop	Concurrent	Load Balancing, Interference, Path Closeness	Provides a receiver-initiated video streaming multipath solution in which nodes comprising the multiple paths are chosen so that they have minimum coupling/closeness.
Software defined VANETs [88]	2014	SDN-based	Backup	Reliability	Deploys VANETs in a software-defined environment for reliability enhancement and power management using a central controller.
Personal Area Networks (PANs)					
IEEE 802.15.4 [120]	2011	Hop-by-Hop	Concurrent	Delay, Reliability	Chooses a parent node opportunistically to enable multipathing by combining IEEE 802.15.4's MAC layer <i>cluster-tree</i> method with RPL's <i>destination-oriented directed acyclic graph (DODAG)</i> .

optimal in terms of performance, cost, or reliability. There are many challenges involved in integrating source routing, multipathing, and global Internet routing [15]. A major impediment being the rigid inflexible policies of various ASes for interdomain routing [15]—e.g., it is typical for BGP to ignore potential paths suggested by a source routing scheme [15]—which makes the control of the end-to-end routing behavior difficult for end hosts or even any single AS. There are ongoing efforts to devise new techniques to adapt BGP to implement source routing while supporting the provisioning of multiple paths through ASes [72], [73], [94]. With various trends on the Internet (such as multihoming) nudging the Internet towards multipathing, there is also an increased interest in adapting interdomain routing protocols for multipath. Different works have attempted to develop new techniques based on the legacy interdomain routing protocols that we describe next.

Feedback-based routing (FBR) [43] is an alternative approach to interdomain routing. In contrast to BGP, this approach differentiates routing information into its structural components (that deals with the existence of links) and dynamic components (which refers to the quality of these links). In this approach, the core routers of the Internet are only responsible for disseminating structural information, while the routing decisions are taken at the edge based on dynamic end-to-end performance measurements in addition to the structural information. FBR is multipath-based since each access router computes two paths to every advertised network prefix (that differs from each other as much as possible) based on the structural information it receives. The access router then monitors the route quality, and uses the better path to carry its traffic. In case the access router finds that both the routes have an infinite RTT, then two new routes are computed. Since the core/transit routers do not compute routes (and only disseminate structural information), FBR scales better than BGP.

The caveat of source routing is the overhead involved in carrying the routing information in the packet itself. Various approaches have been proposed for this. LIPSIN proposes putting a *Bloom filter* into data packets in source routing style as a forwarding fabric for multicast traffic for publish/subscribe Internets [122]. MPLS can be used for explicit routing, which is similar to routing, by using labels avoiding the overhead of carrying the information of the entire path in the packets. The *BANANAS* architectural framework [44], on the other hand, works by encoding paths using a short hash (called the PathID) which is a sequence of globally-known identifiers.

In the *BANANAS* framework [44]—which can be deployed in both intradomain or interdomain settings—is an attempt to make conventional routing protocols (such as OSPF, ISIS and BGP) capable of exploiting the path diversity of the Internet. The need of specifically constructing multiple paths is eliminated by generating a short hash called a *PathID* using a globally accessible sequence of identifiers. A path is encoded as PathID using a hash on a sequence of identifiers (or any non-null permutation of such sequence) that can be the IDs of routers, link interfaces, link weights, AS numbers (used in the case of interdomain routing), etc. The global

significance of these identifiers distinguishes *BANANAS* from the conventional MPLS and ATM architectures. Routers that are multipath-enabled can calculate a PathID with the global scope of the above-mentioned identifiers in view. This feature makes the routers independent and autonomous with respect to the setup of multiple paths without the need of *explicit* control signaling protocol for a specific path setup.

Multipath Interdomain Routing (MIRO) [93] is a BGP-based interdomain multipath routing framework that is backward compatible with BGP and incrementally deployable. In MIRO, the default routes are learnt through the BGP protocol, and arbitrary pairs of domains can perform AS-level negotiations to add additional paths tailored to their specific needs. MIRO provides more flexible control of traffic to the intermediate ASes with reasonable overhead. This is done to promote path diversity and improve path resilience in interdomain routing. The main features of MIRO include (i) AS-level path selection; (ii) negotiation support for alternate routes between two arbitrary ASes (that may not be adjacent); and (iii) support for policy-driven export of alternate routes by the responding AS during negotiation (through which transit AS have greater control over which traffic enters their networks). Although MIRO is similar to overlay routing in that it uses tunnels, it differs from overlay schemes because it allows paths to be selected in the underlay through the cooperation of intermediate ASes. The difference between MIRO and most of the other source routing works is that, apart from providing end hosts or the edge routers with routing control, MIRO allows control over path selection by intermediate ASes.

The “new Internet routing architecture” (NIRA) [49] is a source-routed interdomain multipath routing protocol in which the end host has more flexible control over the sequence of ASes to be traversed by its packets. By providing the end-host control over the adopted end-to-end path, NIRA can help foster healthy ISP competition leading to enhanced services and improved end-to-end performance and reliability. NIRA provides users/end hosts with more control over the chosen routes without running a global link-state (LS) routing protocol, and thus NIRA can provide a user with a route with lower control overhead. With NIRA, a user only specifies a source address and a destination address to choose a specific interdomain path comprising multiple ASes, and the path is altered when the source or the destination addresses is changed. This technique is efficient because of its faster fail-over action and lower control overhead.

The concept of routing deflections [47] is applied at the level of domains in *pathlet routing* [51]. In pathlet routing, domains advertise multiple paths or *pathlets* to a source located in another domain. These pathlets afford the source with the ability to exploit the advertised path diversity through route deflections. More abstractly, pathlet routing allows source routing over a virtual topology. Using the concept of routing over virtualized topologies, network operators can declare their routing services expressively in a flexible fashion. In particular, using pathlet routing, the routing styles of BGP, source routing, MIRO, LISP and NIRA can be emulated [51].

AMIR is a multipath interdomain routing protocol proposed in [53] that establishes multiple paths and deploys source

routing at AS level. In this scheme, the multiple paths are established around a primary path that is usually constructed by a policy-compliant traditional interdomain routing protocol. This is achieved by negotiating path provisioning with neighboring ASes. The participating ASes are called *multipath service agents (MPSA)* and these agents impose charges for providing topological information for the computation of multiple paths through them. Data is routed over these ASes by a *multipath service forwarding agent (MSFA)* embedded at the edge routers of such ASes. Paths are set up and the data is forwarded in a source routing fashion using MSFAs of MPSAs. AMIR enhances application-level user experience with lower overhead incurred in path setup at the AS level.

2) *Source Routed Multipathing (Wireless Solutions)*: There are various source-routed multipath routing proposals for wireless networks—with the major bulk of work being done for mobile ad-hoc networks (MANETs) and wireless mesh network (WMNs). In this subsection, we will cover the wireless networking configurations of wireless sensor networks (WSNs), MANETs, and WMNs.

Source-Routed Multipath Protocols for WSNs:

WSN multipath routing protocols have rarely adopted source-routed solutions to avoid the scalability problem that may arise by overburdening the potentially energy constrained source node. The major bulk of multipath protocols proposed for WSNs are based on the hop-by-hop routing paradigm instead that will be covered in Section II-B. As an example of a source-routed WSN multipath solution, we will discuss the H-SPREAD [98] routing protocol that facilitates the construction of multiple node-disjoint paths towards the base station in a distributed manner. In such a setting, only a few paths can be compromised if some malicious nodes appear in the network. If the next-hop node of a packet at an intermediate node is unavailable, then this node redirects the packet (in a source routing manner) to an alternate next hop which is part of an available backup path, thus enhancing the reliability of a data transfer. It is only when there is no available alternate path that a packet is dropped at a node.

Source-Routed Multipath Protocols for MANETs:

Most of the routing protocols proposed for MANETs are reactive/ on-demand. On-demand routing protocols compute routes when they are needed for communication and are better suited to dynamic topologies such as MANETs. Such an approach also avoids the extra overhead associated with proactive routing (due to the periodic exchange of control packets). In on-demand routing, routes are discovered through the flooding of specific control packets by the source that are used to gather the topological information. This information is then collected using a reply packet from the destination. The legacy protocol for MANETs called *dynamic source routing (DSR)* [123] works on the same principle. After computing the route, a source node embeds the route information in a packet's header. These packets are then forwarded by the intermediate mobile nodes to the eventual destination.

In [40], Nasipuri et al. have modified DSR so that it can support the provisioning of disjoint multiple paths. Two flavors

of multipath routing are presented: in the first approach, only the source node is provisioned with multiple paths; while in the other, both source and intermediate nodes (that make up the primary path in the normal operation of DSR) are also provided with multiple disjoint paths around the primary path. This work makes an effort to reduce the control overhead incurred by the conventional DSR in the route discovery process. This control overhead is a consequence of flooding the query packets in the network. The reduction in the control overhead is due to the provisioning of multiple backup paths. So, in the case when the primary path fails, the traffic is switched to one of the backup paths without rediscovering the route all over again.

Another similar work based on modifications in DSR is called *multipath source routing (MSR)* that is proposed in [41]. Multiple paths are computed using the conventional DSR's route discovery and maintenance operations. In this work, a probing-based mechanism is used to achieve a balanced load over multiple constructed paths. This mechanism probes multiple paths for their RTT responses and the traffic is load balanced accordingly to control congestion and reduce end-to-end delay.

Two maximally disjoint paths are created to balance the input traffic over them. One of the paths constructed is the conventional shortest path. This process is done in an *on-demand* fashion. This work is called *split multipath routing (SMR)* [110] and it performs better than DSR in terms of end-to-end delay. Traffic is distributed over the multiple paths on a per-packet manner. This technique, although effective for load balancing and controlling congestion, can create a packet-reordering problem at the destination that makes its use with TCP problematic.

Multiple independent paths are deployed for the purpose of improving secrecy and security in [46]. This method is called *SPREAD*. In this method, a secret message is divided into multiple parts called *shares*. Then, using secret (threshold) sharing methods, these shares are routed over multiple independent paths towards the destination. An adversary node has to compromise at least a certain number of shares (a specific threshold number) out of the total to completely decipher the original message. To enhance the security further, the shares are forwarded over multiple independent paths towards a destination. In this case, an attack on all the paths must be made by salvaging at least the threshold number of shares to decode the original message. By this way, even if some of the nodes along the paths are compromised, the message as a whole is protected.

Physical distance between the discovered paths in a multipath routing protocol is critical in the case of wireless networks. This is because of the mutual interference between the paths that can degrade the communication quality. *Spatially disjoint multipath routing protocol (SDMR)* is proposed in [115]. A source node floods route request packets towards its destination. On the way back, in the Route Reply packets, the intermediate nodes, between the source and the destination, embed their link-state information (containing information about one-hop neighbors) in these packets. This helps the source node to obtain a partial graph of the whole network

and choose paths that have the largest *spatial disjointness*. The packets are then forwarded using source routing in a round robin manner over the discovered multiple spatially disjoint paths. The goal of this scheme is to reduce mutual interference between the discovered paths to gain the benefits of load balancing and bandwidth aggregation in MANETs.

Source-Routed Multipath Protocols for WMNs:

In [118], reliability and interference minimization are studied in a multiradio multichannel WMN. In this work, non-disjoint multiple paths are considered, and reliability is achieved through packet duplication over these paths. A metric, called *weighted interference minimization path selection metric (WIM)*, is introduced for interference minimization among the paths. This metric and the accompanying algorithm consider interference among paths and neighboring nodes using the same radio channel. It has been shown that this technique, where non-disjoint multiple paths are used, enhances reliability, interference minimization and end-to-end delay as compared to the techniques where only maximally disjoint paths are considered.

In [48], *joint multi-channel multi-path control (JMM)*, which is a cross-layer design involving the multi-channel link layer and the routing network layer, is proposed to enable multipath routing in WMNs. In this study, the time horizon is divided into slots, and then different channels are attributed to different time slots. By this way, the contending input traffic is balanced at different times over different channels and different paths. As a result, the end-to-end throughput performance is enhanced.

B. Hop-by-Hop Routing

In a hop-by-hop routing scenario, routing decisions are taken, as opposed to source routing, in a distributed or per-hop fashion. Each node along a path processes incoming packets and forwards them by consulting a routing table stored in its memory. Hop-by-hop routing is especially popular for multihop and infrastructure-less environment such as WSNs due to its decentralized and robust nature. Hop-by-hop is by far the most famous and widely adopted technique in IP networks. Link state (LS) and distance vector (DV) routing schemes are the most popular techniques that have been adopted in such routing scenarios. In the rest of this section, we review works that deploy this routing technique in different network types.

1) *Hop-by-Hop Multipathing (General/ Wired Solutions):* A *link state*-based algorithm is proposed by Vutukury and Aceves in [56] to calculate loop-free multiple paths with minimum dependance on the underlying physical network topology. In [57] and [58], Vutukury and Aceves present a *distance vector*-based algorithm to address the same problem of finding loop-free multiple paths. In all of these works, the multiple paths are provisioned based on a special condition called *loop-free invariant (LFI)* (which is described in [56]). When a routing algorithm satisfies this condition, then the loop-freeness at every instance is ensured.

In [74], a family of protocols called *maximum alternative routing algorithm (MARA)*, which is based on a directed

acyclic graph (DAG), is presented. Using the DAG approach ensures that loop-free paths are constructed. Based on a network topology, the DAG approach constructs a graph with the maximum number of multiple paths within a minimum amount of time subject to a throughput constraint in order to improve the reliability of communication.

Hashing enables a node to select a path from a set of available paths while forwarding packets along a path. Two hashing schemes, namely *direct hashing* and *table-based hashing*, are studied in [59] for the purpose of packet forwarding in a multipath routing scenario. In direct hashing, a hash function is performed on a packet's header, and the output is directly mapped to one of the paths from the set of constructed multiple paths. In table-based hashing, on the other hand, the output of a hash function is first put in one of the bins (with the total number of bins potentially differing from the total number of available paths) and the bins are then mapped to multiple paths. In the work of Cao et al. [59], five 16-bit cyclic redundancy check (CRC)-based direct hashing techniques and an XOR of source-destination IP-based and table-based hashing technique are developed for balancing the load over *equal-cost multiple paths (ECMP)*. ECMP is a simple method for provisioning multiple paths that has been incorporated into many conventional routing environments (such as OSPF, MPLS and ISIS) [89], [90]. ECMP is explained in detail in Section III-C1a. An interesting feature of table-based hashing is that it can be used to balance load over unequal path weights (as described in [59]) at the cost of requiring more states than direct hashing.

Dispersity routing enables concurrent transmission of data over multiple paths [54]. There are two different flavors of dispersity routing. A node can either send distinct data packets over multiple available paths for load balancing, or send duplicated data packets over multiple paths for improved reliability and resiliency.

R-BGP is an extended version of the traditional BGP routing in which multipath routing is applied to hierarchical networks comprising different ASes [72]. In addition to the primary paths, the additional multiple paths serve as backup paths, and all of them aim to maintain connectivity even in the face of multiple failures subject to the compliance with AS policy. In [94], an advertisement mechanism is proposed for BGP to advertise information regarding multiple paths at the AS level [94] while minimizing control overhead with the objective of providing path diversity as a means to circumvent failures. In [73], BGP is enhanced to calculate loop-free multiple paths.

Yet another multipath routing (YAMR) is a multipath routing protocol that constructs multiple paths and establishes concurrent transmission in hierarchical networks [77]. It renders reliable communication against any single interdomain link failure. Compared to BGP, it incurs lower control overhead by isolating failures and keeping the whole Internet oblivious of their occurrences. This technique of containing the failure is termed as *hiding*. This technique combined with multipathing helps YAMR to converge with a lower control overhead than BGP while demonstrating enhanced reliability.

2) *Hop-by-Hop Multipathing (Wireless Solutions)*: In this section, we will present a summary of hop-by-hop multipathing protocols that have been proposed for wireless networks.

Hop-by-Hop Multipath Protocols for WSNs:

Conventional single-path routing schemes put all traffic on a precomputed set of links, thus quickly draining the energy of the intervening network elements. In multipath routing techniques [105], on the other hand, network traffic is distributed over multiple paths with load balancing and concurrent transmission, thereby avoiding overburdening a single set of network elements and enhancing the network lifetime (a critically important feature in WSNs).

The provisioning of multiple paths in a wireless scenario is prone to intra- and inter-path interference. In [65], a radio disjoint multipathing method is proposed. This technique constructs multiple paths in an incremental fashion keeping in view the congestion and interference effects, so it helps to improve bandwidth gains among bandwidth-hungry multimedia applications. In the first instance, only one path is established; and upon receiving feedback on network congestion, other paths are established in an incremental fashion. However, when high interference is observed, sensor nodes may be put into passive mode to reduce the number of multiple paths, and with that the amount of interference as well. Overall, this technique improves throughput and energy efficiency but at the cost of higher control overhead.

Another technique called *meshed multipath routing* (M-MPR) with *selective forwarding* is proposed in [63], [64] to incorporate forward error correction (FEC) and selective forwarding into multipath routing for better reliability in order to improve throughput performance. In selective forwarding, a packet is forwarded over the path with the best downstream wireless channel quality, with acknowledgement driven retransmissions eliminated by using FEC mechanism, which facilitates efficient load balancing with lesser control overhead. It has been shown that M-MPR provides higher throughput than its *disjoint multipath* counterparts.

Yang et al. [76] have proposed a *network coding-based reliable disjoint and braided multipath routing* protocol. In the disjoint routing approach, multiple disjoint paths are created in a hop-by-hop fashion and routing information is stored locally on a node which does not take end-to-end path in view. A node groups its neighbors based on their respective distances from a destination node in terms of the number of hops hence aiding the load balancing process, in which each packet is forwarded after an election at a node with the path with lesser number of hops towards a destination being selected. Each node calculates a few logical paths as backups, which may not be completely disjoint, and hence are called braided multiple paths. Reliability is introduced in this work by the deployment of network coding at the source and the intermediate nodes. The encoded packets using this technique can only be decoded at the respective destination nodes. Simulation results show that multipathing with braided paths works better in terms of reliability and load balancing in WSNs. Braided multiple paths are considered in another study [96] for improving reliability

and energy efficiency. It is observed that the braided path technique can avoid 50% more isolated failures while also having a lesser control overhead when compared to disjoint path techniques.

A geographical approach known as *geographical power efficient routing* (GPER) for multipathing has been proposed for WSNs [55]. This protocol enables individual nodes to select their own transmission ranges depending upon their respective residual energy resources thus enhancing network lifetime. This act also helps in preserving the energy resources of a passive listening neighbor. As a further improvement, multiple paths are provisioned on a per-node basis. This approach is adopted so that the overall communication can be distributed over a wider area of a network and hence saving a particular set of network elements from being overwhelmed and deprived of their energy resources sooner than the rest of the network. It has been observed that by doing so enhances the network lifetime by up to 30%.

Although the traditional focus of WSN research has been on energy efficiency, the requirement of security is also very important—especially for environments such as military or healthcare. Stavrou et al. [124] have surveyed security-related issues. A threat model that covers the behaviors of malicious users is presented, and the secure multipath routing protocols for WSNs are categorized according to their specific operational capabilities in achieving security-related objectives.

The problem of joint fault and intrusion tolerance is considered in [108]. A distributed and in-network fault-intrusion tolerance protocol called *secure and efficient intrusion-fault tolerant protocol* (SEIF) is developed. This technique renders nodes to take decisions in regards to resilience and reliability that come with independence from base station. By this way, the time incurred to detect an intrusion is reduced. This security mechanism is incorporated into a multipath routing protocol called *sub-branch multipath routing protocol* (SMBR) to improve network discovery time and energy efficiency, as well as to reduce control overhead as only one message per node is required to construct a network topology. H-SPREAD [98] is a routing protocol for WSNs that facilitates the construction of multiple node-disjoint paths towards the base station in a distributed manner. In such a setting, only a few paths can be compromised if some malicious nodes appear in the network.

Hop-by-Hop Multipath Protocols for MANETs:

Ad-hoc on-demand multipath distance vector routing (AOMDV) [66] is one of the famous modifications of AODV. AOMDV provisions multiple link-disjoint paths between a source-destination pair in an on-demand manner. Another AODV modification is presented by Ye et al. in [67]. Reliable nodes, in terms of energy resources, are deployed in the network topology. These nodes are used to construct reliable multiple paths, along with other unreliable nodes, to ensure routes are more reliable and less prone to failures and outages in MANETs. An algorithm is presented to control the placement and trajectories of the reliable nodes so that these nodes are highly likely to be part of a path that a normal (or unreliable) node constructs. An important discovery made

in this work is that there can only be a few disjoint paths between a source-destination pair in a MANET—with the number of disjoint paths decreasing as the distance between communicating pairs increases. This discovery motivates the use of reliable nodes in such networks.

MP-AODV is an adaptive multipath routing scheme based on the conventional AODV routing scheme proposed in [68]. Primary paths are selected based on the received signal strength. Alternative paths are then constructed by the exchange of, so called, source generated *heartbeat messages* among nodes in a MANET. These messages help a source to gather the mobility information of nodes that make up the multiple (backup) paths towards a destination. One of the main objectives of this work is to enhance the reliability of communication in MANETs. Another work [69] presents *node-disjoint multipath routing* (NDMR) protocol that also makes use of backup paths, and improves upon the legacy protocols, namely *AODV* and *DSR*. In this work, the issue of reliability in terms of *packet delivery ratio* is addressed. Besides reliability, this protocol is also efficient in terms of end-to-end delay and control overhead.

A technique called *packet-caching* is proposed by Valera et al. in [61] to ameliorate congestion and packet losses with the provision of multiple paths. In this work, *caching and multipath* (*CHAMP*), which consists of cooperative packet caching and shortest multipath routing, is proposed to cater for route breakdown, which is a common yet important issue in MANETs. In *CHAMP*, each node has a buffer in which it stores/ caches data packets that it encounters during a normal operation. Information of alternate routes is also cached at the same node to be used when the downstream node (relative to this node) fails to forward a data packet or somehow malfunctions. If such a situation occurs, an upstream node forwards the dropped packet from its cache along an alternative path. This technique performs better than *DSR* and *AODV* in terms of end-to-end delay, packet delivery ratio and control overhead. In another work [62], *reliable adaptive multipath provisioning* (*RAMP*) is proposed for reliable adaptive multipath provisioning to meet bandwidth and differential delay constraints. The problem of differential delay is particularly important for the multipath routing scenario. When traffic is dispersed over multiple paths, then different paths may offer different amount of delays, and as a result the traffic reaches at a destination at different times. This difference in delay among multiple paths is called the differential delay. Traffic is allowed to traverse over multiple paths, which may be disjoint, so that communication outage due to a single link failure can be avoided.

An energy saving protocol for MANETs is proposed by Srinivas and Modiano in [60]. In order to conserve energy, 2-link-disjoint and k-node-disjoint algorithms are proposed for dense and sparse networks, respectively. A set of algorithms are developed to find node- or link-disjoint paths in a given network graph so that transmission energies of nodes are minimized. An interesting observation has been reported in this work, that is the incremental energy of each newly added link-disjoint path is decreasing due to the wireless nature of the underlying network.

There has also been work done on multipathing for a specific configuration of MANET known as vehicular ad-hoc networks (VANETs). In [75], a node-disjoint concurrent multipathing protocol is proposed for VANETs to improve the overall packet delivery ratio even though the relaying nodes may have unknown speeds. In another video streaming work for VANETs [119], *location-aware multipath video streaming* (*LIAITHON*) is proposed for urban vehicular networks. In this protocol, a receiver initiates a video streaming process by sending a request to a source. The source then transmits the video stream over the intermediate nodes that are well away from each other to reduce inter-path interference. The selection of such nodes is carried out using the geographical location information of these nodes (with respect to the source node).

There have also been biologically inspired routing protocols proposed in MANETs, particularly AntHocNet, which is an *ant colony optimization* (ACO)-based multipath routing protocol [70], [71]. The ACO-based AntHocNet protocol has both a proactive path probing component as well as a reactive path setup/maintenance component. It is shown in [70], [71] that AntHocNet outperforms AODV in terms of delay, jitter and packet delivery ratio, while having the drawback of being more expensive in terms of its control overhead.

Hop-by-Hop Multipath Protocols for CRNs:

In [78], a multipath routing scheme for cognitive radio networks (CRNs) is deployed by introducing a metric that tells about the *route closeness* of multiple paths, which depends on the physical distance among the multiple available paths. In CRNs, an unlicensed or secondary user (SU) can access the licensed radio spectrum if it is not being used by the licensed or primary users (PUs). Since the PU has higher priority over SUs, the arrival of PU precludes SUs from accessing the spectrum and makes the routing problem challenging. The paths with less closeness are preferred because such paths are less prone to the PU activities due to the greater distance among these paths.

Hop-by-Hop Multipath Protocols for PANs:

There have also been works done on multipathing for wireless personal area networks (WPANs). Multipathing is realized over the MAC layer of the low-rate WPAN (LR-WPAN) standard IEEE 802.15.4 in [120]. This realization is achieved by selecting multiple parent nodes in an opportunistic manner at a time. This is done by making IEEE 802.15.4's MAC layer's cluster tree to operate in harmony with the *destination oriented directed acyclic graph* (DODAG) as proposed in *routing protocol for low-power and lossy networks* (RPL) [125].

C. Overlay Routing

We have earlier seen that source routing equips end hosts with the ability to deal with any failures arising along a route itself. The source can thus either identify an end-to-end route itself as is done in the “source routing” paradigm, or it may use the “overlay routing” paradigm to choose some intermediate

hosts to route through in order to avoid performance and reliability problems that may exist on the direct route. An overlay network is essentially a virtualized logical network built on top of a physical network with tunnels interconnecting edge devices. The overlay network appears to the nodes connecting to it as a native network with the possible existence of multiple overlays on the same underlying physical infrastructure (which allows support for multi-tenancy). Overlay routing can be used to override suboptimal routing in the substrate networks, and can be used to support more flexible self-driven routing control of multihomed end hosts. In this section, we will discuss the overlay routing paradigm.

1) *Overlay Routed Multipathing (General/ Wired Solutions)*: Source routing is not supported in most of the Internet routers. The most common technique used to support multipath on the Internet is to maintain an overlay network for connecting nodes and supporting source routing. Overlay routing is especially popular on the Internet for supporting TE and multipathing.

Multihoming for overlay networking is deployed in [92], and a new *multihomed overlay network* (MONET) protocol is proposed. The purpose of this work is to increase the availability of websites to users by creating multiple paths to web proxies. A method called *way point selection* is used to choose a subset of available paths leading to the proxies in order to access websites. This method decides the order in which each path is chosen to forward packets towards a particular destination server depending upon the success statistics of the paths. In another work, Andersen et al. [79] proposed to enable an end user to switch between actual network paths and those constructed by an overlay network. This incentive is also available to the applications depending on their network requirements. In order to enhance bandwidth availability to meet application requirements, an overlay networking-based multipath solution is presented in [82].

Multimedia streaming is another important application in which reliability is important for seamless transmission. Multipathing reduces jitter and packet loss for efficient and reliable multimedia streaming. A coding technique called *multiple description coding* (MDC) [126], [127] is deployed to enable multipathing for streaming media. In this coding technique, streaming media from content distribution networks is divided into multiple independently decodable streaming flows [128]. Packets from these multiple streaming flows are then transmitted over multiple disjoint paths. Many works base their multipath techniques on the streaming media using MDC.

In another overlay routing scheme, Savage et al. proposed the *Detour* overlay network [14]—comprising geographically distributed routes interconnected using tunnels (which act as virtual point-to-point links) over the commodity Internet—to be used for better end-to-end performance by routing around performance bottlenecks. Any host using the Detour network directs its outbound traffic to the nearest Detour router with the packets taking an exit at a point close to the destination. Detour uses a dynamic multipath routing approach with automatic load balancing and congestion avoidance for high performance. The router and congestion is controlled at

the Detour routers—which are, technically speaking, still edge devices of the overlay network.

In [80], a quality of service (QoS) metric based on *link correlation* is presented for overlay networks. This metric, along with *path correlation*, is then used to select multiple paths for an MDC streaming video flow. It is observed that this overlay network-based technique improves SNR compared to the conventional link-disjoint methods by up to 3.2 dB. Another work in [81] makes use of overlay routing and MDC for efficient and reliable video streaming. Based on MDC, different multipath streaming models have been derived. Then, based on these models, an optimum strategy is given to select an optimum set of multiple paths. This technique also improves the SNR by up to 6.07 dB as compared to the single-path and link-disjoint multipathing.

Various researchers have studied how to support virtual overlay networks on top of the physical substrate networks, with a few works exploiting multipath for efficiently supporting virtual networks (VN) on the underlying substrate. Yu et al. [129] have studied the strategies that the substrate network can adopt to flexibly support virtual network embedding in which the virtual network requests are mapped to specific nodes/links in the underlying substrate network. Yu et al. simplify the general problem of virtual link embedding (which is known to be computationally intractable) by allowing the substrate network to (1) split a virtual link over multiple substrate links, and (2) migrate paths to periodically optimize the utilization of substrate network. The related problem of multi-domain VN embedding with limited information disclosure is studied by Dietrich et al. [130] who proposed a VN embedding framework that enables VN request partitioning scheme while assuming only limited information disclosure.

2) *Overlay Routed Multipathing (Wireless Solutions)*: In this section, we introduce various overlay multipath routing protocols proposed for wireless networks, specifically WSNs.

Overlay Multipath Protocols for WSNs

A multipath routing protocol called MMSPEED [97] is proposed to achieve QoS guarantee. In this protocol, multipath routing is deployed in a distributed overlay manner to improve reliability while satisfying the constraint of packet delivery speeds. Here, the packet delivery speed is equal to the ratio of one-hop average distance to one-hop average delay. Braided multiple paths are considered in [96] for the enhancement of reliability and energy efficiency. It has been observed that this technique works 50% better as compared to its *disjoint* counterparts in terms of isolated and patterned failure avoidance with low control overhead. The isolated failure here refers to those incidents in which nodes fail randomly and independently, while in the case of patterned failures, nodes belonging to a specific geographic region fail.

D. SDN-based Routing

In this section, we discuss SDN-based multipath routing proposals. SDNs allow networks to be programmable through the separation of the control plane and the data plane, as well as through the support of high-level control abstractions

TABLE IV: Routing Metric Based Classification of Multipath Routing Frameworks

Routing Metric	Sample Works in Various Networking Configurations
Multipath Specific Routing Metrics	
Path Disjointness/ Independence: To gain the benefits of multipath routing, it is important to choose paths that are, as much as possible, independent from each other. With disjoint/ diverse paths, routing is more resilient to failure, and the paths will not share a common bottleneck.	<i>General/ Wired:</i> [131], [132] ; <i>WSNs:</i> totally disjoint paths [100], maximally radio disjoint routing [65]; <i>MANETs:</i> SDMR [115], SMR [110]; <i>VANETs:</i> Node-disjoint-multipath [75], LIAITHON [119].
Path Closeness/ Correlation: Similar to the schemes that aim at building link/ node-disjoint paths, routing protocols can choose the metric of ‘path closeness’ or ‘path correlation’ to prefer diverse paths over paths that are correlated. Paths that are more dissimilar and independent tend to provide better performance.	<i>General/Wired:</i> [80]; [133]; <i>MANETs:</i> SDMR [115]; <i>VANETs:</i> LIAITHON [119]; <i>CRNs:</i> [78].
Interpath Interference Minimization: Routing protocols can be designed to prefer paths that minimize interpath interference. With reduced interpath interference, concurrent usage of resources can be performed and the overall routing performance improves.	<i>General/ Wired:</i> Dynamic Constrained Multipath Routing with MPLS [42] ; <i>WMNs:</i> weighted interference minimization (WIM) metric [118]; <i>WSNs:</i> [65], IM2PR [109].
Traditional Routing Metrics Used For Multipath	
Delay/ Latency: Delay refers to the average end-to-end latency experienced by a communicating source-destination pair. We are interested in minimizing latency which can be achieved by minimizing network congestion and alleviating bottlenecks. Solutions include: concurrent multipathing for dispersing the traffic to avoid bottlenecks, and using backup paths for more rapid route rediscovery after failure.	<i>General/Wired:</i> [58], [134]; <i>WSNs:</i> [97], [105]–[107]; <i>WMNs:</i> [118]; <i>MANETs:</i> [62], [66], [69]–[71], [110], [113], [135], [136]; <i>Other:</i> [75], [120], [137].
Reliability/ Resiliency: Reliability refers to the ability of the networking protocols to work consistently as required. Failures can be circumvented by using backup paths, and higher packet delivery ratio can be ensured through concurrent multipathing over diverse links.	<i>General/Wired:</i> [45], [49]–[51], [72], [77], [92]–[94]; <i>WSNs:</i> [41], [76], [97], [98], [106]–[108], [138]; <i>WMNs:</i> [118]; <i>MANETs:</i> [46], [61], [67], [69]–[71], [113], [135], [136]; <i>CRNs:</i> [78]; <i>VANETs:</i> [75].
Bandwidth/ Throughput: Throughput, or more informally bandwidth, refers to the amount of data that can be transferred within a period of time. The use of traditional single-path routing schemes can result in congestion hotspots that act as a performance bottleneck. Load balancing and dispersion of network traffic using multiple paths at the same time offer higher bandwidth and throughput gains.	<i>General/Wired:</i> [73]; <i>WSNs:</i> [41], [63]–[65], [100], [103], [105]; <i>WMNs:</i> [48], [139], [140]; <i>MANETs:</i> [136]; <i>CRNs:</i> [141], [142].
Control Overhead: Control overhead refers to the number of control messages that need to be exchanged to compute and maintain multiple paths. An efficient multipathing protocol minimizes control overhead.	<i>General/Wired:</i> [58]; <i>WSNs:</i> [41], [103]; <i>WMNs:</i> [140], [143]; <i>MANETs:</i> [40], [66], [110], [144]; <i>CRNs:</i> [142].
Energy Efficiency: Energy efficiency refers to the extent the multipath technique conserves energy resource. Such a metric is important for WSNs and MANETs in which energy is a limited and scarce resource. Both concurrent transmission and traffic engineering can increase network lifetime by dispersing the traffic over multiple paths.	<i>General/Wired:</i> [17], [18], [145]–[148] ; <i>WSNs:</i> [63]–[65], [76], [96], [99], [101], [103], [105], [107], [108], [149], [150]; <i>WMNs:</i> [151]; <i>MANETs:</i> [60], [114], [144]; <i>CRNs:</i> [137].

[152], [153]. The centralized decision making approach at the SDN controller, along with the development of new sophisticated control abstractions, allows convenient control and management of networks, and can be used to implement multipath routing.

1) *SDN-Routed Multipathing (General/ Wired Solutions):* Software-defined networking (SDN) is becoming increasingly popular for management of networks, especially in DC environments. SDN is often an enabling technique for the operations of DCs. Rain Man is a SDN-based firmware architecture proposed in [85]. This firmware has a layer two (L2) routing mechanism that enables efficient multipath routing in DC networking. Another work called Hedera, with a global view of active flows, presents an adaptive and efficient multipathing approach based on SDN concepts [83]. In [84], Micro-TE is proposed so that flows in a DC environment are scheduled using a central controller with the help of stochastic properties of network traffic and multipath routing. Full link utilization is achieved in B4 [86] using traffic engineering solution with the help of OpenFlow-enabled centralized controller and packet forwarding mechanism using multiple paths. OpenFlow is a protocol for SDNs through which the control plane of a central controller talks to (or programs) the data plane of the switches in the network.

2) *SDN-Routed Multipathing (Wireless Solutions):* In this section, we introduce various overlay multipath routing protocols proposed for the wireless networking configurations of MANETs, WMNs and VANETs.

SDN-based Multipath Protocols for MANETs

SDN concepts have been used in [117] to realize a MANET infrastructure. It has been shown that carrying out routing operation through SDN enhances the reliability of such networks.

Concurrent transmission of data packets has been studied as a use case to demonstrate that this technique increases packet delivery ratio at the expense of moderate control overhead increase.

SDN-based Multipath Protocols for WMNs

An interesting study called *wireless mesh software-defined networks (wmSDN)* is presented in [87]. In this study, a mesh network comprises wireless mesh routers (WMRs) that are IEEE 802.11-based and OpenFlow-enabled. The data plane of these WMRs is programmed via a centralized controller. Traffic balancing is carried out by the central controller using an algorithm called *gateway selection algorithm (GSA)*. In this technique, the WMRs act as gateways for mesh network to forward traffic out of the network, and these are selected for Internet connectivity in a round robin fashion, where the least used gateway is selected first and so on, instead of using the nearest gateway at all times to forward traffic out of this network. This is done in order to achieve an enhanced network utilization.

SDN-based Multipath Protocols for VANETs

A similar work is presented for the case of VANETs in [88]. In this work, routing decisions are controlled using a central controller in an SDN environment. Multipathing provides resiliency in the case of emergency scenarios, such as route failures, or for the high priority flows. To achieve this, the controller reserves a few paths dynamically as backups.

E. Routing Metrics for Multipath Routing

Different metrics have been used to quantify the performance and quality of multipath routes, including some that are direct extensions of the common routing metrics used in single-path routing (e.g., throughput, delay, congestion)

and others that are customized for the unique settings and requirements of multipath routing. We discuss both types next. A summary of the various metrics used for multipath routing is presented in Table IV. The choice of routing metrics used by various wired and wireless multipathing works is also summarized in Tables II and III.

1) Multipath Specific Metrics: The quality of multipath routing essentially depends on the independence and diversity of the paths chosen (so that an isolated fault/ failure is not able to affect all of the paths). We have divided multipath specific metrics into three types, namely those that measure: (i) path disjointedness/ independence; (ii) interpath interference; and (iii) path closeness/ correlation. These metrics are related to each other—e.g., if we select multiple paths that are disjoint, we also select paths with lesser closeness/correlation. The similarity of these metrics stems from the fact that they all measure correspondence to the qualitative ideal of maximally independent/diverse collection of paths that do not share a common bottleneck. Each of these metrics has been used in various network configurations including wired and wireless networks (as can be seen in Table IV).

Path Disjointedness/ Independence: As noted earlier, the key to benefiting from a multipath routing framework is to ensure that multiple paths support each other when needed. This is best ensured when all the paths are independent so that a congestion hotspot or a performance bottleneck is not able to simultaneously affect the different paths. Many proposed multipath routing works have utilized some variation of this metric for path computation to measure link- or node-disjoint paths that are maximally independent. Examples of works proposed for wired networks include [131], [132]; while works proposed for wireless networks include totally disjoint paths [100] and maximally radio disjoint routing [65] for WSNs; SDMR [115] and SMR [110] for MANETs, and node-disjoint multipath and LIAITHON [119] for VANETs.

Interpath Interference Minimization: Routing protocols can be designed to prefer paths that minimize interpath interference. With reduced interpath interference, concurrent usage of resources can be performed and the overall routing performance improves. A number of interpath interference minimizing multipath protocols have been proposed, both for wired networks (such as the MPLS-based dynamic constrained multipath routing approach [42]) as well as for wireless networks (such as the use of weighted interference minimization or WIM metric for WMNs [118] and the IM2PR protocol for WSNs [109]).

Path Closeness/ Correlation: Path closeness/ correlation is very similar to path disjointedness/ independence metric—and indeed, it can be thought of as its variant. With multipath routing, we seek to minimize the metrics of ‘path closeness’ or ‘path correlation’ to prefer multipath paths that are mutually more diverse. Such metrics have been used both for wired networks [80], [133] as well as for wireless networks (such as the SDMR [115] work for MANETs, and the work of Beltagy et al. for CRNs).

2) Traditional Routing Metrics: Some applications might be delay sensitive, while the others put emphasis on reliability. Here, we present different classes of multipath routing protocols whose design considerations revolve around a specific metric. The metrics being discussed in this section include: (i) delay, (ii) throughput, (iii) reliability, (iv) control overhead, and (v) energy efficiency.

a) Delay/ Latency: In many time-sensitive applications, the delay metric is of vital importance. In the case of WSNs, time sensitivity matters in the case of military, health and life-saving scenarios, such as a fire-detection network application. In such scenarios, using multiple paths concurrently can significantly reduce end-to-end delay. In concurrent multiple-path transmissions, which serve as congestion avoidance mechanism [89], [90], using efficient load balancing and flow-splitting techniques help to avoid bottlenecks contributing to lower end-to-end delay. Another aspect is in the usage of multiple paths as backups [69], [68], [66]. In this case, when a primary path fails, the traffic is shifted to one of the alternative paths from a pre-calculated set of backup paths. In this way, the amount of time for route rediscovery is reduced.

b) Bandwidth/ Throughput: Bandwidth efficiency, or efficient utilization of network resources, is among the most important benefits of multipathing. Concurrent multipathing, with efficient techniques of load balancing and flow splitting, distribute the network traffic across a network. By this way, the network traffic does not overburden a single set of nodes and links as opposed to the case of single-path routing [154], [54], [89]. The majority of existing state-of-the-art multipathing protocols utilize bandwidth/ throughput as part of the goals in network performance enhancement.

c) Reliability/ Resiliency: Packet delivery ratio is an important metric to measure reliability/ resiliency. A number of factors such as network congestion, interference, mobility make links/ nodes prone to failure. Ideally, network routing should be resilient such that in the case of the failure of the primary path, routing switches to an alternative backup path (if available) within a short period of time. In addition, using multiple paths in a concurrent way, the effects of bottlenecks and congestion can be reduced through packet duplication or redundancy. Many works discuss the issue of reliability, resiliency and packet delivery ratio in several networking applications [92], [93], [72]. In the case of general/ wired networks, multipath routing is mostly used to provide reliability in both intradomain and interdomain scenarios [49], [51]. Multipathing provides reliable communications in WSNs [97], [105], [106], CRNs [78], VANETs [75], IEEE’s 802.15.4 [120], and MANETs [46], [61], [67], [69], [113].

d) Control Overhead: Implementation of multipath routing is often more expensive than single-path routing due to the extra cost of control messaging imposed by discovering and maintaining more than a single path for a source-destination pair. A number of works have focused on minimizing control overhead associated with multipath routing techniques. An example is a multipath distance vector routing scheme called MDVA shown in [58]. In addition to computing loop-free




Fig. 2: Summary of route computation techniques for *wired* networks




Fig. 3: Summary of route computation techniques for *wireless* networks

paths, this scheme uses the technique of *diffusing computation*, where a node reports its distance information to an upstream node after calculating it based on the distance information from its downstream node, and neighbor-to-neighbor synchronization which is more efficient in terms of control message exchange compared to distributed Bellman-Ford (DBF) and topology broadcast (TOPB) algorithms; specifically, it improves the convergence time after link failures. The control overhead routing metric is highly relevant to energy efficiency. To avoid drainage of energy resources, it is desirable to contain the number of control messages that need to be exchanged. Researchers have focused on this problem and have proposed routing schemes to limit the control overhead [41], [103]. The routing metric of energy efficiency has been applied in WMNs [66], [110] and MANETs [66], [110].

e) Energy Efficiency: Energy is one of the main concerns in the design of routing protocols, particularly in power-constrained networking configurations such as WSNs and MANETs. Multipath routing, with the use of concurrent multiple paths, improves energy efficiency [96], [63], [64]. In a legacy single-path routing, the network resources of the primary path are drained much quicker because of the traffic load it bears. Using multipathing and concurrent transmission with efficient load balancing techniques, the network traffic spreads out ensuring that only a few particular network elements are not overburdened. This provides not only bandwidth and throughput benefits but also energy conservation.

F. Summary of Trends

In this section, we identify the main trends observed on the basis of papers presented in this survey (which in our view provide a balanced representative view of the overall multipath routing literature). The trends of the commonly adopted routing approaches/ choices are noted below and summarized visually in Figure 2 for wired networks, and in Figure 3 for wireless networks.

- Trend 1: Hop-by-hop routing is the most popularly used route computation technique, whereas overlay routing

is the least used (amongst those papers that we have surveyed).

- Trend 2: Considering the hop-by-hop route computation technique, we note that this multipath routing technique is especially popular in WSNs (compared to other network configurations). This is in contrast to other network types such as wired networks in which hop-by-hop routing and source routing are used a similar number of times. Hop-by-hop routing is, on the other hand, comfortably the most popular choice for WSNs.
- Trend 3: Multiple paths can be used in both concurrent/ backup configurations. It is observed (based on the works presented in this study) that wired/ general networks prefer backup and concurrent multiple paths almost equally, while in wireless scenarios, we see that the concurrent transmission option is adopted more often.

A multipath algorithm designer can consult Figures 2 and 3 to quickly learn which design choices have traditionally been most popular for the particular networking type or configuration being considered. As an example, if we are to design a multipathing method for WSNs, then the most popular choice is to adopt hop-by-hop routing with concurrent use of the multiple paths while focusing on the energy conservation metric.

III. HOW TO USE MULTIPLE PATHS?

A network-layer multipath framework subsumes both the control plane operations of multipathing (that involves construction of multiple paths) as well as the data plane operations of multipathing (that involves the use of the constructed multipath paths for forwarding the flow). After having spent a lot of time focusing on how routes are computed, we now turn to the (arguably more) important question of how to use these constructed paths. We note here that route computation task is typically considered a network-layer function, while the usage of resources—and in particular, congestion control through which end-user's transmitting rate is controlled to ensure that the overall network performance is not deteriorated due to

TABLE V: Summary of Flow Splitting/ Load Balancing Techniques

Approach	Year	Main Idea
Static/ Quasi-Static Approaches (In these approaches, the traffic split is <i>either static or is managed offline based on a predicted traffic demand matrix with changes occurring after a significantly long period of time</i>)		
Weighted round robin or Deficit round robin [155]	1995	Adopts an efficient <i>fair queuing</i> method to schedule contending traffic flows in an almost perfectly fair fashion with efficient running time and a complexity of $O(1)$ per packet processing time.
ECMP [156]	1998	Divides traffic evenly over multiple equal-cost paths. It is stable as it does not make dynamic adjustments to OSPF costs based on traffic loading. There are three splitting methods in ECMP: 1) Per-flow hashing; 2) Per-packet round robin; and 3) Dividing destination prefixes among the available next hops in forwarding table.
Valiant Load Balancing (VLB) [18], [17]	1990	Performs a decentralized 2-phase routing technique, which is agnostic to the traffic matrix. In the first phase, VLB—independent of the destination—redirects traffic to arbitrary intermediate routers. The intermediate routers then forward the traffic to the destination in the second phase.
OSPF-OMP [89]	1999	Extends OSPF, which can form multiple equal-cost paths and uses ECMP techniques to divide traffic equally over these paths, so that it utilizes opaque LSAs to distribute loading information to facilitate uneven splitting.
MPLS Optimized Multipath (MPLS-OMP) [90]	1999	Proposes a TE procedure to achieve load balancing across multiple label-switched paths in MPLS. In MPLS-OMP, interior gateway protocol (IGP) floods traffic load information according to OSPF-OMP and ISIS-OMP specifications.
OSPF/ IS-IS with Link Weight Adjustment [157]	2000	Explores near-optimal Internet TE approach under two constraints. Firstly, paths are established using shortest-path routing. Secondly, there is equal flow splitting across equal-cost paths (ECMP) associated with OSPF/ IS-IS for <i>known demands</i> . A local search heuristic is proposed for the NP-hard problem of optimizing IGP link weights given a particular traffic matrix.
Adapting OSPF/ IS-IS Weights in a Changing Environment [158], [159]	2002	Proposes a family of algorithms for efficient OSPF traffic engineering in a dynamic environment. If the traffic demand matrix changes significantly, weight settings are updated with just a few changes.
Sridharan et al. [160]	2005	Proposes a near-optimal traffic distribution approach based on OSPF/ IS-IS shortest-path routing—which selects a subset of next hops for each prefix using a simple heuristic with provable performance bounds—while preserving the existing protocols and infrastructure.
Fong et al. [161]	2005	Proposes techniques for improving OSPF including features that allow flows to traverse non-shortest paths (with a possible cost of occasional loops).
DEFT [162]	2007	Proposes a link-weight-based flow-splitting mechanism that allows routers to redirect traffic on non-shortest paths with an exponential delay on longer paths. DEFT provably improves the performance of OSPF and IS-IS.
PEFT [163]	2011	Enhances the link-state routing protocol with hop-by-hop forwarding by utilizing non-shortest paths. Like DEFT, it splits traffic on all paths with longer paths carrying an exponential penalty. PEFT is different from <i>link-based splitting</i> methods—such as DEFT [162] and Fong et al. [161]—in that it uses <i>path-based splitting</i> (as explained in [163]).
Dynamic (Load-aware) Approaches (In these approaches, the traffic split is <i>dynamically adapted to the changes in observed load</i>)		
MPLS-based Adaptive Traffic Engineering (MATE) [164]	2001	Uses probing messages and explicit congestion notification packets to solve multipath routing which is formulated as optimization problem with the objective function of minimizing congestion on highly utilized links in the network. It is an online framework.
TeXCP [165]	2005	Uses explicit congestion notification packets to perform flow splitting in order to balance load in real-time manner in response to actual traffic demands/ failures. It is an online distributed protocol.
Common-case Optimization with Penalty Envelope (COPE) [166]	2006	Establishes optimal multiple paths efficiently for expected network scenarios while providing worst-case guarantees. It is a hybrid of oblivious routing and prediction-based optimal routing—a class of TE algorithms.
DATE [167]	2007	Incorporates both congestion control and routing mechanisms to enable <i>edge routers</i> to split traffic for each source-destination pair over multiple paths. It is a cross-layer multipath routing framework.
REPLEX [168]	2006	Avoids oscillations which is inherent to load-sensitive routing. It is an online traffic engineering technique based on <i>Wardrop</i> routing policies [169]. Since it is not restricted to MPLS-based solutions, it is a more general solution compared to TeXCP and MATE.
Flowlet-aware Routing Engine (FLARE) [170]	2007	Splits a single TCP flow across multiple paths without resulting in packet disordering.
Hedera [83]	2010	Provides a centralized load-aware scheduler for datacenters.
Path RE-Feedback with Loss EXposure (PREFLEX) [171]	2010	Uses estimates of loss rate to balance congestion by automatically splitting traffic.
MicroTE [84]	2011	Exploits short-term, partial predictability of traffic matrix in order to adapt to traffic variations. It is a load-aware centralized TE solution designed for <i>datacenters</i> .
MP-TCP [172]	2011	Shifts traffic from one path to another; hence, it is a load balancing scheme for end hosts.
Adaptive Load Balancing Algorithm (ALBAM) [173]	2011	Performs multipath scheduling of mobile users who access heterogeneous wireless networks using multiple interfaces.
Hop-by-hop Adaptive Link-state Optimal (HALO) [174]	2013	Solves the TE intradomain routing problem through adaptation to traffic demands and computing link weights locally and optimally using the same information as OSPF. It is the <i>first</i> optimal link-state hop-by-hop routing algorithm.
TRUMP [175]	2013	Proposes four adaptive and optimal distributed solutions using decomposition techniques to exploit multiple paths.
Theoretical Analysis (Multipath algorithms for optimal flow splitting given a set of paths)		
Kelly et al. [176]	1998	Presents and analyzes two classes of rate control algorithms to set prices of utilized resources fairly: one algorithm from the perspective of user and another one from the network.
Wang et al. [177]	2001	Shows that shortest-path routing with positive link weights can be used to support arbitrary flow splitting to provide an arbitrary set of loop-free routes. This Internet TE problem is considered without using overlay networks.
Kelly et al. [178]	2005	Presents and analyzes a quick load sharing/balancing approach across multiple paths between a source-destination node pair based on end-to-end measurements.
Han et al. [179]	2006	Presents and analyzes the concept of <i>multipath TCP</i> where congestion control and multipath routing are studied jointly. Using overlay routers, which support source routing, multipathing in networks is enabled.
Theoretical Analysis Multipath algorithms for optimal (or near-optimal) decisions on <i>both</i> routing path and flow splitting		
Banner et al. [180]	2007	Provides <i>optimal</i> and polynomial-time algorithmic solutions to congestion control using <i>K-path routing (KPR)</i> and <i>restricted multipath (RMP)</i> problems, both of which are NP-hard, using acceptable approximation schemes.

network congestion—is considered a transport-layer function. The routing task is managed at the network layer through the interaction of routers while congestion control is managed at the transport layer at end hosts. These systems unfortunately work largely independent of each other resulting in an unwelcome tussle that arguably is holding back multipathing from achieving the immense benefits it can offer [171].

We note here that while the layered model is a good tool for conceptualizing networking functionality, many important aspects of the Internet—issues such as reliability, flow control, congestion control, QoS, and indeed resource pooling through multipath—cannot be pigeonholed into a single layer. The rigid bifurcation of multipath-related functionality into layers results in an antagonistic tussle in which each

layer attempts to balance traffic and perform resource pooling myopically according to its own needs to the detriment of the overall performance. Following the maxim “*layering is a good servant, but a bad master*”, we recognize that the overall success of multipathing requires solutions beyond those that are implemented at the network layer. In this paper, therefore, despite our focus on network-layer multipathing, we also describe related issues (such as flow splitting, load balancing, and multipath congestion control methods) that are meta-network-layer and typically considered in the territory of the transport layer.

In this section, we see how the multiple paths constructed by the routing schemes described in Section II can be used for forwarding the packets. We will see next how the processes of *flow splitting* and *traffic engineering* together contribute to the act of load balancing that is crucial to reaping the benefits of resource pooling with multipath.

Considering first a single flow, we see that the packets of a single flow can either use multiple paths by splitting the traffic to concurrently use the diverse paths, or alternatively, the multiple paths may be used in a backup configuration where the primary path is used until a node/ link failure along the primary path upon which the routing defaults to the backup paths. The purpose of *flow splitting* in concurrent configuration is to achieve the benefits of resource pooling using which a collection of multiple paths appears as a single virtual pooled link. An important question that arises is how to decide the ratio of flow splitting since both of the computed multiple paths may have very different properties (in terms of delay, bandwidth, congestion, etc.). There are various requirements that we expect from a reliable flow-splitting multipathing system: there should not be a need for per-flow state to ensure scalability; the system should be tunable, dynamic, and should support heterogeneity in bandwidths/ delay. In addition, the reordering of packets should be minimized (for which it is important to recognize distinct flows, and have mechanisms to mark packets as belonging to a particular flow that uses a certain sub-path). Coming now to multiple flows, and how their packets share the multiple paths available over the network, it becomes important to consider both the utilization of the network resources and the fairness of resource allocation. This is the prerogative of *traffic engineering* which is typically handled inside the network by ISPs.

There are five important questions relating to how the multiple paths are used that influence the organization of this section.

Firstly, how many of the available multiple paths should be used for transmitting the traffic flow? While higher number of paths can improve the performance, the management cost also increases with increased number of paths.

Secondly, an important question is whether the multiple paths should be used concurrently or in a backup configuration? Traditionally most techniques have used multiple paths in backup mode for its simplicity although increasingly concurrent usage is being adopted to improve performance.

Thirdly, should the flow splitting/ load balancing be static or dynamic? Since the congestion on different paths vary dynamically, the load balancing and flow splitting mechanisms

should also be ideally dynamic. However, in practice static flow splitting/ load balancing schemes are also used due to their simplicity and low overhead.

Fourthly, how can we use multiple paths to perform load balancing through the diffusion and control of congestion? In particular, traditional congestion control schemes need to be adapted for use in multipath scenarios.

Finally, an important aspect of load balancing is who gets to control the process? and where? Traditionally, there has been a tussle between end hosts and the network for network control with traffic engineering typically implemented in networks while congestion control implemented at the end hosts.

Apart from discussing these five questions (in the order presented) in their respective subsections, we will also present a subsection that summarizes the trends presented in this section.

A. How Many Paths to Use?

While multipathing can provide significant performance and reliability benefits, the computation and maintenance of multiple paths also entails costs. It has been shown that utilizing *all* the available paths provides only limited performance gains over an approach that utilizes a limited subset of all paths [3]. In fact studies have shown that, to obtain performance gain offered by multiple paths, one or two alternative paths, in addition to the traditional single path, can suffice [3], [181] with multiple reservation algorithms potentially providing a gain in terms of connection establishment time which is attractive for interactive applications [181].

In previous research, it has been shown through multiple studies that the bulk of benefits of multipathing can be reaped with only two paths, and that increasing the number of routes provide diminishing returns [3]. This finding is formally established in a widely cited paper by Mitzenmacher on randomized load balancing [182]. In another work, Akella et al. have also posited similar insights in an empirical study [10] focusing on multihoming. More recently, Merindol et al. have proposed a multipath routing technique that also relies on two precomputed paths.

Key and Towsley [183], [184] have similarly shown that not many paths are needed to reap the benefits of multipath. The chosen set of paths need to however be dynamically monitored for their quality to ensure that the best paths are chosen and retained while any poorly performing paths are flushed out and replaced. While the dynamic load balancing is done at a rapid time scale, the exploration of routes can take place at a slower rate. To reiterate, the network-layer functionality of computing routes and finding connectivity is done relatively slowly while the dynamic load balancing over the chosen paths is done through the end systems via the transport layer rapidly.

It must be pointed out the flow size is an important factor in the choice of how many paths to use, and in deciding how to prefer different access technologies in multi-homed systems. For instance, it has been shown in a study conducted by Deng et al. [185] that for short flows, there is no significant benefit of using multipath, and selecting the right access network technology is more crucial for high performance.

B. Concurrent or Backup Paths

In *single-path routing*, a router only maintains a single route to each destination, whereas in *multipath routing*, multiple routes are maintained. While most common configurations in multipath routing are to maintain a single primary path along with alternative paths, it is also possible to use multiple routes in a *backup* or *concurrent* fashion.

Both of these approaches have their own benefits. Depending upon a particular network application or a specific metric (e.g., reliability or load balancing), a routing algorithm adopts any one of these approaches. To ensure reliability alongside load balancing in a network, the concurrent multipath approach [111], [148] gives the desired results. On the other hand, a set of multiple backup paths are precomputed in *path failover* strategic algorithms [186], which are provisioned in order of a specific *path ranking* mechanism, in the case of failure of the primary path [66], [40]. Information about the concurrency of multiple paths is summarized for various routing protocols in Tables II and III.

Next, we discuss various aspects of backup and concurrent multipathing techniques in more detail in the following subsections.

1) *Backup Multipathing*: Routing protocols that make use of multipathing with multiple backup paths usually construct a primary shortest path in a source-destination pair along with a set of alternative paths. The alternative paths in the set are maintained and updated according to the network topology changes. In case when the primary path fails, then the traffic is quickly shifted to one of the backup paths from the set of precomputed alternative paths. In this sense, the time incurred in route rediscovery is avoided, thus contributing to more reliable, fault-tolerant, and resilient routing.

Legacy protocols, particularly AODV, are improved to support multiple backup paths in AOMDV [66]. Similarly, DSR, which is another well-known routing protocol is enhanced in [40] to improve the reliability of routing. In [68], multiple paths are provisioned in an adaptive manner by updating AODV to increase packet delivery ratio in networks. *Static-fail-over routing* (SFR) is another important protocol making use of alternative backup paths to circumvent network failures.

Another work [186] discusses the issue of reliability in multipath routing in the terms of *backup paths*. In this work, SFR is proposed to study the positive and negative aspects of legacy *fail-over routing* techniques.

2) *Concurrent Multipathing*: With modern networking applications having huge data traffic demands, network congestion has become a significant problem. In legacy single-path routing protocols, many potential paths in a source-destination pair are ignored. In concurrent multipathing, alternative paths are exploited to relieve network congestion while improving network performance, efficiency, and reliability. A large number of works employing multiple paths, along with intelligent load balancing techniques, have been proposed in the literature. We present a representative sample of these works (many of which are based on concurrent multipathing) in Tables II and III.

We note here that utilizing multiple paths in parallel is not a recent idea. This general idea has been proposed under various names in the literature, viz. traffic dispersion, striping, and inverse multiplexing. *Traffic dispersion* has a rich history in networking with Maxemchuk first proposing a traffic dispersing framework known as dispersity routing in 1975 [54]. This approach advocates spreading the traffic spatially (by dividing the original message into sub-messages and then transmitting them in parallel over disjoint paths of approximately equal length) to achieve load balancing and fault handling. The use of disjoint paths ensures that transmission errors on different paths are statistically independent. *Striping* pools multiple resources to meet capacity requirements even if individual links/ network interfaces, or hardware components, cannot satisfy the required capacity [187]. *Inverse multiplexing* refers to the phenomenon of aggregating channels to allow greater flexibility in selecting bandwidth [188].

The concept of *concurrent multipath transfer* is presented in the work [133]. This concept implies that a source-destination pair that uses multiple paths in a multihomed environment can reap the benefits in terms of better network utilization and data transfer resiliency. This work addresses the issue of *shared bottlenecks* in the computation of multiple concurrent paths. A metric based on *path correlation* is proposed to resolve this issue. This metric tells us about the correlation of two paths in terms of shared bottlenecks. Lower value of this correlation metric indicates that the paths become more uncorrelated or disjoint.

Another important load balancing technique that is popular at the link layer in recent multipath-based topologies for DCs is *Valiant Load Balancing* (VLB) [17], [145], [146] (covered in Section III-C1b in more detail). This technique has inspired many other multipath routing works to cater to the problem of congestion and inefficient utilization of network resources. An example of VLB-based work that uses uniform flow splitting is presented in [148] in which a design of network backbone is provided to provision multiple concurrent paths.

Recently, concurrent multipathing has become very popular and mainstream with the development of multipath transport protocols such as MP-TCP. The great promise of concurrent multipath transmission using transport protocols is the promise of true resource pooling using *coupled congestion control* [189] with which it will be possible to shift traffic away from congested paths. Recent fluid-flow models [178], [179] have shown the benefit of shifting traffic from congested to uncongested paths. In particular, it was shown that multipath congestion control can perform the task of load-dependent routing stably at shorter time scales (at the same scale of TCP congestion control that operates at the order of RTTs). In effect, multipath congestion control assumes a role typically associated with network-layer routing (of moving traffic to another path thereby avoiding congestion hotspots). This motivates the discussion of multipath congestion control in our paper whose focus is on network-layer multipathing solutions. In a work related to concurrent transport-layer multipathing, Wischik et al. have proposed the principle of equipoise—or equally balanced concurrent multipath transport—as a concurrent multipathing technique that allows better resource pool-

ing, resilience to faults on any subpath, and better handling of traffic surges [190].

C. Load Balancing Approaches for Multipath

Broadly speaking, load balancing and flow splitting approaches can be either static/quasi-static or dynamic/load-aware. The static/quasi-static methods are suited for network traffic matrices that do not change rapidly, while the dynamic/load-aware methods are suited for changing network traffic matrices. There are also a few theoretical works that describe optimal, or near-optimal, methods for load balancing and flow splitting. We have summarized major load balancing works, with their references, in Table V.

1) *Static/ Quasi-Static Load Balancing*: This approach is adopted when traffic matrices are relatively stable with changes occurring after a significantly long period of time. Flow splitting techniques for load balancing in such networks are managed in a static or offline manner. Most of the static load balancing approaches aim to fairly schedule the flows giving equal preference to individual links [155].

a) *Equal Cost Multipathing (ECMP)*: ECMP [91], [156], [191] is a multipath routing mechanism that has been employed in many routing protocols like OSPF and ISIS [89], [157]. Intuitively speaking, ECMP can be seen as a simple “bonding” technique that can load balance over equal-cost paths. A router can implement three main techniques while choosing a next-hop node to forward the traffic including the following.

- *Hashing*: In this technique, a *hash function* is applied to the header of each packet. The output of this hash function determines over which path this packet is forwarded to. In [59], hashing is proposed for the purpose of network traffic load balancing over multiple equal-cost paths. This work makes use of a 16-bit CRC over the five tuples of headers of incoming packets. Five direct and one table-based hashing methods are presented. This work also presents a solution for keeping the packets of a TCP flow in order when it is split over multiple paths because of hashing.
- *Per-Packet Round Robin*: Fair queuing methods are deployed on a per-packet fashion to provide network resources to multiple contending queues of data packets in an efficient manner. In [155], efficient and fast running algorithms for fair queuing, in which each flow that traverses across the network gets equal share of network resources, such as weighted round robin, are proposed. These can be utilized to load balance the incoming contending flows over the equal-cost multiple paths.
- *Destination Prefix-based Approach*: In this technique, different paths are attributed with different prefixes of a destination. The load is then split according to these prefixes [89].

Since ECMP deals with ‘equal-cost’ multipathing, it does not consider suboptimal routes although they can be potential candidates. *Constrained Shortest Path First Protocol (CSPF)* [192], which is an extended shortest path protocol, calculates

path that fulfills a certain set of constraints, such as minimum bandwidth, maximum number of hops, and maximum end-to-end delay requirements.

b) *Valiant Load balancing (VLB)*: VLB is an efficient parallel communication technique first presented by Valiant [17], [145], [146] for general-purpose computer processors. VLB is an attractive form of load balancing since it can be applied on a per-flow or on a per-packet fashion, and can help build networks that can efficiently and reliably support any traffic matrix. In the case of VLB, it is sufficient to know the capacity of the ingress router of a network instead of the knowledge of the whole network traffic matrix.

In VLB, each data flow is routed towards its ultimate destination in two stages: in the first stage, the data flow is routed to a randomly chosen intermediate router without the consideration of an underlined traffic matrix; the intermediate router routes the data flow to its ultimate destination in the second stage.

We explain the concept of the two-staged VLB operation with the help of a diagram shown in Figure 4. Consider a network of N nodes each with a capacity of R as shown in Figure 4. The two stages work in harmony as follows:

- *Stage 1*: In the first stage, the incoming traffic that flows through the *ingress point* is divided evenly, regardless of the eventual destination, among all the points in the network. As the capacity of the ingress point is R , and if it receives the traffic at this same rate, then the traffic is divided among the links giving $\frac{R}{N}$.
- *Stage 2*: In the second stage, each intermediary, which has received the traffic in the first stage, forwards the traffic towards the actual destination. As in the first stage, each node receives $\frac{R}{N}$ amount of traffic; so in this stage, each link is loaded with $\frac{R}{N}$ traffic by each node.

The combination of these two stages thus leads to the link capacity of $\frac{2R}{N}$ each for the normal/ efficient operation of VLB. Here, it can be noted that, if we only know the capacity of the ingress point (assuming the capacities of all the other points are equal), then there is enough information to perform load balancing, in a congestion-free way, on an incoming traffic load. However, an ingress router must not be overloaded. Uniform VLB—where incoming traffic is split equally at the ingress point over all the outgoing links—can promise 100% network throughput.

As an example application of VLB, Zhang et al. have used VLB to design a backbone network with fault tolerance [148]. It has been discussed that, by 10% over-provisioning of network elements in terms of nodes, a 50-node network continues to operate in a congestion-free manner with up to five node or link failures when VLB is applied to balance traffic load in such a network. Over-provisioning can be calculated by the ratio of an arbitrary number of failures to the total network nodes.

c) *Other Static Approaches*: MPLS with *optimized path* (MPLS-OMP) [90] is used for load balancing purpose for label-switched paths. TE, which is further explained in Section III-E1, is used to balance traffic load across the label-switched




Fig. 4: Load balancing Through VLB

paths. In the work of Fortz et al. [157], OSPF/ISIS is deployed with known traffic matrix to perform TE in a near-optimal fashion under the constraints of shortest-path routing and even split of flows over equal-cost multiple links. In the follow-up works for changing network traffic matrix, Fortz et al. in [158], [159] introduced a family of algorithms that adjust the link weights according to the changes in the traffic matrix. Here, OSPF/ISIS is used for the traffic engineering purpose again. The fact that link weights are altered according to the changes in the traffic matrix makes the works quasi-static in [158], [159]. In these works, the changes in the network traffic matrix are predicted for an intradomain routing. In another work [160], a TE solution for OSPF/ISIS shortest-path routing is deployed for flow splitting purposes. The proposed solution computes approximately optimal link loads that provide a provable performance bound. In [161]–[163], improvements for OSPF are presented so that non-shortest paths can be included for flow splitting. However, it is observed that a delay penalty that increases exponentially is incurred with the introduction of non-shortest paths.

2) *Dynamic (Load-aware) Load Balancing*: In this approach the TE solutions split the input traffic in an adaptive or online manner [164]–[166], [175]. The flow split is governed by the changes occurring in the network traffic matrix. We note that most of the works [164], [165], [167], which adopt this approach, address the problem of congestion control alongside balancing the input traffic load in a real-time manner.

Load-sensitive routing is prone to collisions. A general solution called REPLEX is proposed in [168] to dynamically balance the traffic load and also avoid collisions. Another issue that can arise when the traffic load is split over multiple paths is the disordering of packets of a TCP flow. FLARE [170] is proposed with real-time ability to split a TCP flow that does not result in disordering of TCP flow packets. MP-TCP [172] is another online flow distribution method that renders an end user to shift the load from one path to another according to the input load demands and congestion feedback from the network. HALO is proposed in [174] as the *first optimal LS intradomain routing protocol*. It also deploys TE solution for load-aware traffic load balancing. Adaptive load balancing algorithm for multipath scheduling (ALBAM) is proposed for heterogeneous wireless access technologies for mobile users.

DCs are an important application area in which multipath routing with efficient load balancing mechanism provides efficient and data intensive operation. Centralized control solutions, which manage and schedule the DC traffic flows, are presented in the literature. Hedera, which is a central load-aware scheduler, is proposed in [83] for the DCs. This solution schedules flows in a dynamic and adaptive manner to efficiently utilize network resources. Another work for the same purpose, which is called MicroTE, is presented in [172]. This technique is also load-aware and it uses short-term traffic matrix predictions to distribute the network traffic over multiple paths. Here, the short-term traffic, which has been shown to exist in DCs in this work, is used to infer unpredictable traffic in order to ameliorate congestion.

In [171], a mutualistic resource pooling architecture named PREFLEX was proposed in which the hosts and edge routes worked in unison to collectively share the burden, and conversely reap the accrued rewards, of balancing traffic over multiple paths. PREFLEX is called a mutualistic framework since it allows congestion control (performed at the transport layer at end hosts) and traffic engineering (performed inside the network) to coexist and evolve independently by exposing the apparent network preferences and the transport expectations.

As we shall see, transport-layer solutions to load balancing are increasingly becoming more popular [183] [193] [190] [194]. Multipath transport protocols like MP-TCP already incorporate congestion control features that can automatically shift traffic from congested paths onto uncongested paths [189]. As we discuss next, in many aspects transport-layer load balancing is very appealing.

3) *Theoretical Results of Optimal Flow Splitting*:

*“Somewhere ages and ages hence:
Two roads diverged in a wood, and I—
I took the one less traveled by,
And that has made all the difference.”—Robert Frost*

It has been shown in multipath routing literature [178] that, consistent with Frostian wisdom, taking the less-travelled-by path does indeed make all the difference in the context of multipath routing in communication networks. In multipath terminology, this means that multipath-capable flows should

be designed to shift all their traffic from congested paths to uncongested paths. Interestingly, it has been shown that this load shifting can be done stably by the end hosts at the transport layer at a rapid time scale (at the order of RTTs) [178]. In this section, we consider theoretical works that have analyzed congestion control and the benefits of flow splitting on multiple paths using fluid flow models.

In optimal flow splitting, we are provided a set of paths, and we have to decide on what is the most fair and effective way to split traffic on the given set of paths. Fairness is an important goal to strive for and it entails providing every flow a “fair” share of the available bandwidth whenever there is congestion at a bottleneck. The most common form of fairness has been max-min fairness in which all the connections get the same share of the bottleneck. Another common form is proportional fairness in which any change in the distribution of rates will result in the sum of the proportional rates being negative. The rate control of TCP—which is based on additive increase and multiplicative decrease (AIMD)—is proportionally fair. In weighted proportional fairness, on the other hand, each connection is associated with a price and the amount paid per rate (and not the rates) is made fair. As an example, a connection with a price of two will get the same rate as two connections with a price of one each.

It has been known through the work of Kelly et al. [176] that, in a network in which each user chooses a price that maximizes its utility, the system evolves to a state that maximizes the total network utility. Two classes of control algorithms—one from the user’s perspective and another one from the network’s perspective—are presented to set prices for the utilized resources in a fair manner.

In a follow-up work, Kelly et al. [178] considered the important problem of joint routing and rate control where it was demonstrated that load sharing among multiple paths, in a multipath routing scenario, can be achieved in a similar time scale in which rate control is performed. But this quick load sharing approach must be performed while keeping in view the network stability (responsiveness vs. stability). This is why, in this work, a method is proposed to analyze end-to-end stability in the joint routing and rate control scenario with the presentation of a sufficient condition for this purpose. This sufficient condition works in a distributed manner and deals with the stability of end-to-end algorithm for joint routing and rate control.

Han et al. have also considered the problem of joint multipath routing and congestion control with multipath-TCP [179]. This problem is studied based on an overlay network of routers that support source routing. Wang et al. in [177] have also considered arbitrary traffic split among multiple paths in a shortest-path routing scenario. But here, an overlay network is not considered in this TE approach. Links with positive weights have been used to support the distribution of arbitrary input traffic matrix.

D. Multipath Congestion Control

Apart from the issues of route computation and traffic engineering, another important issue is *congestion control*. In

practice, congestion on the Internet is passively controlled at the end hosts through TCP that observes and responds to packet losses. Alternatively, the network nodes can also actively manage their queues using active queue management (AQM) technology—using algorithms such as random early detection (RED) [195], explicit congestion notification (ECN) [196], or controlled delay (CoDel) [197]—to take preventive measures to avoid congestion.

The purpose of multipathing is to enable resource pooling. For effective resource pooling, it is important to develop the abstraction of a *single-pooled resource* from the collection of diverse resources. The key to the creation of this abstraction is effective load balancing. More relevant to multipathing, load balancing can be performed over multiple resources (e.g., the load is balanced over multiple links/ paths). With multipath resource pooling, the aim is to pool multiple paths and to provide the illusion of a single much more powerful path. This resource pooling can provide the benefits of higher throughput, improved fault tolerance, and the capacity to handle traffic surges.

There is a lot of ongoing work in the research community in developing efficient *multipath congestion control* [198], [184], [199]. This interest is further buoyed by the ready availability of path redundancy and diversity in modern networking environments such as datacenter networking. For maximum efficiency, there should be a strong emphasis on the construction of holistic multipathing frameworks that jointly optimize the various facets of network-layer multipathing.

1) *Coordinated or Uncoordinated Control*: With multipath congestion control, an important consideration is whether to use *uncoordinated* or *coordinated* control. Uncoordinated control can be implemented in the current Internet (running unmodified TCP) in a straightforward fashion. This makes uncoordinated control attractive for application-layer solutions, and explains the use of such schemes by P2P frameworks. With coordinated control, on the other hand, the coordinated controller can better balance load when the paths are fixed. Kelly and Voice [178] and Key et al. [183] have studied congestion management in multi-flow environments. They have shown that congestion response for each subflow should be coupled to ensure stable and efficient rate allocation between the subflows. This is important both when a path fails and also when a new path appears. This has been studied using tools of game theory by Key and Towsley [184] by the formulation of a path selection game which studies the choices a user should make about the set of throughput-optimal paths to use seeking to maximize their own net utilities. The ensuing Nash equilibria is shown to be welfare-maximizing social optima for the case of coordinator controller and for the restricted case of uncoordinated controller that do not exhibit RTT bias (unlike TCP which does). For uncoordinated controllers with RTT bias, the Nash equilibria may not coincide with the social welfare-maximizing optimal. Coordinated controllers have also been shown in [184] to improve upon greedy least-loaded resource selection as proposed by Mitzenmacher.

2) *Congestion Exposure*: Today’s traffic management protocols are typically suboptimal since they disregard (i) protocol

interactions between congestion control, routing, and traffic engineering; and (ii) the tussle that emerges from independent network *or* end-host control of congestion. The task of harmonizing congestion control, routing, and traffic engineering is challenging due to the different timescales they act on (short, medium, and long timescales respectively). In recent times, there is renewed interest in integrating routing and congestion control in a cost-effective fashion. The EU-funded Trilogy project [200]—named as tribute to the crucial trilogy of “routing, congestion control and cost effectiveness”—aimed at improving the quality of the internal functioning of the Internet through cost-effective resource pooling and integration of the routing and congestion control functionalities.

An important component of the Trilogy architecture aiming to plug a gap in the Internet’s architecture is Re-ECN (re-feedback of explicit congestion notification)—which is an economics-based congestion pricing resource management framework that holds users/ applications for the congestion they create. It has been shown in previous work that resource allocation through the traditional TCP method can be gamed/ exploited (e.g., a user can open multiple TCP flows to gain disproportionately high throughput at a bottleneck link) since the rate control is performed at the end hosts. The network operators have limited visibility into the congestion in the network and their focus has been traditionally on managing utilization. The traditional focus on “TCP-friendliness” [201] and the focus on equal flow rates have been challenged recently and their fairness and desirability have been questioned [202].

An important protocol in the Re-ECN framework is the Congestion Exposure (ConEx) protocol that can also lead to better integration of end-host applications and service provider, and harmonious interworking of various stakeholders on the Internet [202]. ConEx attempts to fairly “divide the resource pie” and mitigates the indiscriminate abuse of the Internet by an atypical form of traffic policing that only pushes back users who are limiting other users’ freedom. This is done by charging users according to the congestion they cause, rather than by the volume of traffic generated. The underlying idea here is that we should not limit the traffic volume generated by users unless the resource share of other users is being impacted. Thus, if there are more people who wish to send data, there is a higher price of transmission, and vice versa. Without such a system, the system can be exploited causing indiscriminate use of the Internet leading to tussles between ISPs and user applications such as P2P [202].

In ConEx, the packets carry an accurate prediction of the congestion they expect to cause downstream, which is visible to the ISPs/ network operators who can now act on the congestion signal being exposed to them by the end hosts [203]. This information can then be used for a variety of purposes including congestion policy, traffic engineering, and accountability. ConEx encourages well-behaved applications that aid load balancing (e.g., by delaying the transmission of traffic from congested times to uncongested times) and can help facilitate better network economic pricing models. This can allow ISPs to charge users for congestion caused while also providing users incentives to fairly consume resources

since they have to now pay for the congestion they are causing. There is also the scope of novel tussle-aware congestion management schemes that utilize information made available to policing of network congestion [204].

3) *New Developments in Congestion Control*: As an interesting example of a new congestion control algorithm, that involve greater interaction between the network and the end systems, consider Low Extra Delay Background Transport (LEDBAT)—a protocol used by Apple and BitTorrent for large background transfers that account for up to an approximate 15 to 20% of the the global Internet traffic—which works by limiting congestion induced by the LEDBAT flow itself in the network. There has also been work in machine-learning-based congestion algorithms that are customized for user and network preferences. In the ReMy work, proposed by Winstein et al. [205], the computer-designed endpoint algorithms meet the defined preferences and assumptions about the users, application traffic, and network much better than human-designed algorithms (even those algorithms that deploy code inside the network). In other works, ReMy searches for the best congestion-control algorithm and optimizes the expected objective over prior assumptions. This work promises that we can retain the end-to-end design philosophy and work with a simple “dumb” network, and have a computer-designed algorithm operates at end points and outperforms even the best in-network algorithm.

E. Which Layer/ Where to Perform Load Balancing/ Traffic Engineering?

For the success of a load balancing/ traffic engineering scheme on the Internet, it is important that the proposed framework holistically integrates the various layers of the Internet architecture. The network alone is incapable of managing the traffic since the routers do not have enough information about the end-to-end traffic flow and thus cannot make an information decision about the paths to be used. The end host, on the other hand, does not know about the network traffic and cannot ensure fairness of the network alone. The problem with performing congestion control at the transport layer is that the control process may not have enough time to kick in for small flows. In many cases, application-layer multipathing is best suited to serve the unique needs of the application and can be used to choose from which server to download (as is done in P2P applications such as BitTorrent that already inherently utilize multipath).

Our discussion on transport-layer multipathing in this paper, despite our network-layer focus, is motivated from the convergence of the research community on the important role that congestion-controlled multipath transport will play in the future multipath-enabled Internet. It has been argued in literature, following the work of Kelly and Voice [178], that end systems using transport-layer functionality, rather than the networking nodes implementing network-layer protocols, are best placed to implement load balancing—particularly across network domains—on the Internet. The information below the transport layer is too coarse to dynamically control the traffic and can only enable coarse-grained functionality

(such as shifting of a few elephant flows away from hotspots within a network domain). Indeed, many application-layer protocols (especially P2P protocols such as BitTorrent) already perform multipath in their function. Implementing a generic multipath load balancing/ resource pooling functionality can help implement a new ‘narrow waist’ of the future Internet. This also naturally fits with the existing role of TCP in implementing congestion and rate control functionalities.

Broadly speaking, there are three main techniques for multipath resource pooling: (i) Routing-based traffic engineering; (ii) Application-based load balancing between multiple servers; and (iii) SDN-based approach to traffic engineering.

1) *Router-based Traffic Engineering (TE)*: Since it is common for large networks to offer significant path diversity, there is great interest in using TE to increase network efficiency. A common approach to implement flexible forwarding is to first classify packets (according to application requirements with packets for the same flow being similarly classified according to the ‘type of service’ field in IP packet header) and then to map packets to different paths by having an edge router examine the packet header, evaluate the path characteristics, and determine a suitable path for a certain class of traffic.

To forward packets using the alternative paths, routers that are part of the communication session establish a logical link with each other using tunneling. The tunneling approach can be utilized in MPLS or IP-in-IP tunnels. Although an extra header is needed for encapsulation; however, compared to the much longer prefix match of destination address, a label-based look up is an easier and faster approach. Using explicit routing, each router can specify a path using a packet header, or a label embedded inside the packet header. Tunneling is a useful approach for special applications that need end-to-end bandwidth guarantee. This approach is also employed in a way that a complete router-level path is specified along with IP options. Moreover, each intermediate node performs a label lookup to discover the outgoing link.

Traffic engineering using network-layer routing functionality is often performed at slow time scales and is not automated (often with a human in the control loop). Internet traffic engineering functionality, like most of the Internet’s control functionality, has emerged as a melting pot of numerous piecemeal fixes (such as the congestion control fix provided by Jacobson’s work in TCP; the traffic engineering fix provided by MPLS; and the policy routing fix provided by BGP).

The BGP protocol, developed mainly for advertising policy-compliant loop-free connectivity, has limited support for traffic engineering. The cumbersome process of establishing peering, provisioning links, as well as tuning routing and QoS policies makes traditional router-based traffic engineering unattractive. This, along with the coarse-grained traffic engineering functionality by protocols such as BGP, has motivated the development of other approaches (such as the two approaches—application-based and SDN-based, traffic engineering—that we discuss next).

2) *Application-based Load Balancing/ Traffic Engineering*: The traffic engineering decision can be made by the application itself at the edge host. As an example, P2P applications

like BitTorrent utilize multipath—but often in a fashion that is oblivious and impervious to the ISP (e.g., the peers are selected randomly without necessarily looking for peers that are local to the same ISP)—thereby impacting the efficiency and costs of ISPs. There has been some work done to allow the ISPs greater flexibility in managing the P2P applications. For instance, the P4P framework [206] aims to equip the provider with the ability to manage P2P applications by performing smarter peer selection, better traffic distribution, and higher transfer speeds, all at a lower cost to the ISP. P4P does this by allowing the ISPs to notify the clients about the preferred local clients. Similarly, the Ono system aims to “tame the torrent” of cross-ISP data generated by BitTorrent by modifying the BitTorrent client to prefer local peers [207].

While application-layer load balancing and traffic engineering can be quite effective, it can conflict with what the network operators are doing, causing an unfortunate tussle. In particular, application-layer multipathing, as performed in P2P applications, is suboptimal due to the lack of availability of topological information to the application layer at the end hosts. The Application-Layer Traffic Optimization (ALTO) framework [208] proposes to equip distributed Internet applications with sufficient topological information exposure to optimize resource utilization. With ALTO, P2P performance can be improved by allowing better-than-random peer selection allowing P2P traffic to align itself with ISP constraints.

3) *SDN-based Traffic Engineering*: In recent times, SDN architecture has emerged as a viable platform to facilitate QoS-aware and efficient traffic engineering. A popular use case of SDN is to use traffic engineering aided by a centralized network-optimizing vantage point in datacenters, campus networks, and in service provider networks.

The modern SDN trend—where network policies can be disseminated from a centralized controller to network elements such as switches—can help fundamentally transform TE as demonstrated most dramatically by Google with their traffic engineering system B4 [86]. B4 uses an OpenFlow-based SDN architecture that achieves nearly 100% traffic utilization when splitting application flows across multiple paths while meeting application priority/ demands. SDN has recently used by McCormick et al. [209] to provision real-time alpha fairness-based TE. The proposed algorithm executes in the range of milliseconds and can be used as an online tool.

An excellent survey for TE in SDN environment by Akyildiz et al. is presented in [210]. We note that there are four frontiers of TE for the modern networking age of SDN. As mentioned in [210], these four frontiers are: (i) flow management, (ii) fault/ failure tolerance, (iii) topology/ network policy update, and (iv) traffic analysis/ monitoring. This is depicted in Figure 5. Here, we briefly describe these frontiers with the relevant reference works, which treat these frontiers in the perspective of SDNs and DCs.

- *Flow Management*: In a SDN environment, each new type of flow is routed towards a SDN controller by an ingress switch. The controller then decides the path for this and the similar prospective flows. The controller disseminates this new path information to all the relevant switches.




Fig. 5: TE Frontiers for SDN (Source: [210]).

However, if there are a large number of new incoming flows, then it takes longer for a controller to decide new paths for each new category of flows, as well as to program the data planes of the relevant switches. This increases delay and can create problems, especially for the DCs, which have been solved using intelligent and innovative TE solutions. Hedera [83] is proposed to deal with these issues using a global view of active flows in the network. It creates non-conflicting paths for the flows as per the current routing and traffic demands in a real-time manner. Another similar work is Mahout [211], which is an end host-based DC traffic management solution. In this study, traffic management is carried out by detecting large flows or *elephant flows* at the buffers of the end hosts. Then, TE is applied to manage these flows using multiple paths with low control overhead.

- *Fault/ Failure Tolerance*: In a centralized control environment of SDN, failure recovery is a challenging task. In case of a failure, a controller must provide the affected switches or other network elements/ resources with a fail-over routing path(s). These paths can be considered as backup paths where the traffic is shifted once the primary path(s) fail. Provision of such type of alternative paths in a real-time fashion can be a challenging task for a SDN controller. Two types of failures can occur, so two types of failure recovery mechanisms are needed, namely (i) data-plane failure recovery and (ii) control-plane failure recovery. For data-plane failure recovery [212], restoration and protection operations for network elements, like nodes and links, must be performed; while for control-plane failure recovery [213], a single-point failure must be avoided. Due to the dependency on the controller in a SDN environment, backup and replication techniques for the controller failures are adopted.
- *Topology/ Network Policy Update*: Topology update in SDNs, unlike conventional networks, is related to the dissemination of network policy changes. The control plane of a SDN is responsible to plan for the updates in the network policy rules. The challenge is maintaining the *consistency* of any updates in the policy rules in the data planes of all switches so that they are operating under the new rules. Problem arises during the *update*

process of the data planes of the underlying switches. Active flows, which are still under the influence of earlier network policy, can get affected or completely discarded during this update process. This phenomenon degrades the QoS of a SDN-based system. So, it is desired to perform the topology update mechanism in a near real-time fashion so that only minimum degradation occurs. Network policy update process is a planned mechanism by a SDN controller. Per-packet/ per-flow consistent and non-conflicting updates are ensured by solutions presented in [214], [215]. In these works, the controller requires an ingress switch to label incoming packets or flows with a tag to tell the policy imposed on this packet so that the switches along a path can process accordingly. By this way, a single packet is processed according to only one policy—either the new or the earlier policy, but not both simultaneously.

- *Traffic Analysis/ Monitoring*: This frontier can be considered as a vital part for the whole TE mechanism in SDN. Accurate and efficient implementation of this frontier renders other frontiers to operate in a systematic cohort with high accuracy. In traffic analysis, different important observations—such as programming bugs, traffic data, network state and traffic/ data patterns—are taken. These measurements are important for timely detection of malfunctioning or the failures of network elements. Unfortunately, most of the current SDN-based systems still employ monitoring tools offered by the conventional IP domain systems. This approach often incurs huge delays, control overhead and inaccuracies in a SDN environment. A need for modern and SDN-related traffic monitoring tools is felt due to the aforementioned issues. Various traffic monitoring tools for SDN environment are available both from the industry, e.g, [86], [216], and the academia, e.g, [217], [218]. A few SDN-based traffic analysis/ monitoring works are discussed next as examples. MicroTE [84] is a load-aware TE solution that predicts short-term traffic matrix that helps in adapting to the changing TE patterns. PayLess is another work that collects aggregate statistics at packet, flow and port levels in a real-time manner with high accuracy and low control overhead [219].

F. Summary of Trends

In this section, we summarize the major trends in how multiple paths are used in wired and wireless networks. We refer to Figures 2 and 3 for a visual depiction of the design choices made by various multipath protocols proposed for wired and wireless networks, respectively. We observe that multiple paths can be used in both concurrent/ backup configurations. It is observed (based on the works presented in this study) that wired/ general networks use backup and concurrent multiple paths almost equally, while in wireless scenarios, we see that the concurrent transmission option is adopted more often. Increasingly, the control of flow splitting, congestion control, and load balancing is being performed at the transport layer and the recent development of multipath

transport protocols can be influential in strengthening this trend. With the recent proliferation of DCs, the use of dynamic/ load-aware flow splitting and SDN-based traffic engineering is also becoming more popular and will grow in importance in the future.

IV. OPEN AND CURRENT RESEARCH ISSUES

A. Pragmatic Deployable Multipathing

1) *More Flexible Support For Policies:* As the domain of applicability of multipathing expands with the growing popularity of multipath, the important problem of harmonizing the aim of efficient multipathing with the policies of the various stakeholders involved becomes more complex. The BGP protocol has traditionally provided limited support for traffic engineering. In particular, BGP does not natively facilitate common tasks such as balancing load across multiple links to a peering AS [220]. Due to the wide deployment of BGP, a clean slate approach to interdomain multipath routing that breaks away significantly from the normal operation mode of BGP is unlikely to be deployed due to the central role of BGP on the current Internet. There have been various works proposed to introduce flexible interdomain multipath routing but the problems are far from fully resolved.

In addition, multipathing, by its very nature, brings to bear the possibility of policy complexity and conflict (both locally as well as globally). For example, a local conflict can arise between the objectives of energy efficiency and minimizing price when using WiFi/ Bluetooth/ Zigbee vs. Cellular [221]; whereas example of a global conflict could be the need to respect the interdomain peering arrangements while performing multipathing efficiently. On a mobile device, multipathing brings about interesting tradeoffs: e.g., if the desire is to plainly maximize battery life, then the lowest-energy interface (Wi-Fi or 3G) will be used; else if the interest is in maximizing throughput, then it will make more sense to use MP-TCP-like protocol to concurrently utilize both the interfaces; finally, if the preference is for price minimization, then a host of factors including the applicable pricing (whether roaming or not) will have to be incorporated to decide. While there are a number of ways the end points could interact with the routing systems at the intradomain and interdomain levels to offset the policy tussle and ensure efficient multipathing, nothing concrete has been agreed upon yet, and this remains an open research issue.

2) *Seamlessly Handling Heterogeneous Paths:* By its nature, multipathing, especially in multi-homed devices/ networks, involves traversing heterogeneous technologies/ domains. As an example, Wi-Fi and cellular interfaces are pervasively available on modern wireless devices. Wi-Fi and cellular have very different characteristics: Wi-Fi typically exhibits stable RTTs with relatively high loss rates, while cellular has variable and large RTTs with low loss rates (with the large variable RTTs attributed to “bufferbloat” [222], [223] that refers to unnecessary provisioning of buffers throughout the network which hinders the congestion control process). The heterogeneous characteristics of the underlying technologies

on the Internet can cause performance issues. It has been shown that MP-TCP performance can suffer when Wi-Fi and cellular flows are coupled due to cellular bufferbloat which causes the problem of *flow starvation* and *idle restart* [224]. More research needs to be done on how to tame the inherent heterogeneity on the Internet so that multipathing can show its full benefits.

3) *Multipathing and Cluster-based Networks:* Despite the many advantages offered by multipath, the high cost of maintaining multiple paths impedes wide-scale adoption. A potential solution to this problem is clustering in which nodes are grouped into logical groups/ clusters, with a clusterhead elected to represent the group. Clustering has many routing applications such as its ability to limit the amount of control messages exchanged, and its ability to aid abstraction (e.g., the Nimrod routing architecture [225] uses clusters to abstract the internal topology, hiding away the details and only revealing additional details upon request). While clustering has been widely investigated, there is a lack of research work on supporting multipathing in clustered networks—particularly the effects of the underlying cluster structure (e.g., cluster size) and clustering operations (e.g., cluster maintenance) on multipathing. Further investigation can be pursued to investigate clustering as an effective solution to improve the scalability of multipathing.

4) *Interworking with Middleboxes:* Middleboxes are network entities that cause interruptions in end-to-end protocols like TCP. Middleboxes can be of various types and purposes, such as firewalls, network proxies, network address translators (NATs), and so on. These entities, unlike routers, are link layer entities that check the layer three information of a data packet. The proliferation of these middleboxes is one of the reasons for the ossification of the current Internet. With modern, often software-defined, networks and end hosts with potentially multiple interfaces/ addresses, the need for an end-to-end multipathing protocol is direly felt. The TCP has an inherent capability to support multipathing through the correct settings of the option fields in TCP packets. The major impediment in the way of realizing this multipath dream is middleboxes as they are tamed to work in an ossified fashion of processing packets assuming no prospective evolution. While making TCP to support multipath, it has to be considered that how the space of sequence numbers is to be used [226], [227]. Do we need a single space for all flows or a separate space will work just fine? If there are gaps in the sequence space of multiple TCP flows, then what might happen at the middleboxes? How will transmission, ordering and retransmission work in the presence of traditional middleboxes [226], [227]? Future research can be pursued to investigate these open questions.

While traditional middleboxes can create problems for multipath TCP, middleboxes can also come to the aid of MP-TCP. For example, Detal et al. have proposed a custom middlebox known as MiMBox that works as a protocol converter between MP-TCP and TCP [228]. Since MP-TCP requires modification of the end hosts, it suffers from the chicken-and-egg deployment problem; MiMBox can facilitate the transition of Internet’s transport towards MP-TCP.

B. Addressing Multipathing Holistically

The Internet’s control architecture is not a result of carefully planned organization of functionality. Such an “architectural soup of network controls” is ill equipped to support multipathing and there is a dire need of a holistic control framework for multipath routing. Some particular problems are described below.

1) *Cross-layer Multipath Support*: Besides the network layer, there has been work on exploiting the benefits of multipathing at other layers; hence it is important to deploy and exploit cross-layer optimization in order to develop effective Internet-wide multipath solutions. The network layer multipath solutions must interwork with application, transport and link layers as described below.

The interaction of the *application layer* and the *network layer* can help to create efficient multipath solutions. The development of overlay routing often requires cross-layer support. It is often the case that certain routing-related observations (such as communication failures at the application layer in which the network layer fails to detect) can only be suitably made at the end hosts. The end hosts can utilize the overlay networking technique at the application layer to implement multipath routing on top of today’s relatively inflexible routing system. It is also often the case that datacenters use load balancers to direct flows at the application layer to suitable servers keeping in view the network conditions. These observations motivate greater support for cross-layer interactions between techniques at the network and application layers.

Previous research has also demonstrated the *importance of the transport layer* in efficient multipathing. Kelly and Voice [178] have shown that, while the network layer is well suited to provide structural information, the load balancing task is best managed at the transport layer. In particular, the key constraint on the responsiveness of each router is the RTT of that route (the information of which is conveniently available at the transport layer). There is also a great need for increased *interaction between the transport layer and other layers*. For example, transport layer can work with network-layer multipathing to reorder out-of-order packets arriving at the destination. There is thus a strong motivation of greater interaction between the network and transport layers for efficient multipathing. With recent research progress in transport layer protocols (with works such as MP-TCP, SCTP, along with numerous cross-layer solutions [167] being proposed), the intersection of the network and transport layers for multipathing is expected to become a promising area ripe for further research.

Finally, exploiting the information at the *link layer* is also very important, especially for wireless networks where channel allocation and assignment are performed to encourage diversity [229], [230]. It is noted here that in wireless networks, transmissions on orthogonal channels can take place simultaneously, and thus orthogonal channels can be considered as logically distinct channels/ paths. Multipath protocols can increase their efficiency by exploiting the multiplicity of routes made available through the use of channel-diversity in multi-channel wireless networks [231]. In a similar fashion, the

interface diversity (or the availability of multiple radio network interface cards) of multi-radio wireless networks [232] can be exploited to improve the performance of multipath protocols.

2) *Mitigating the Tussle Between the Various Stakeholders*: Due to the fact that resource pooling mechanisms based on multipath are built on multiple layers and levels, and that a coherent framework addressing all aspects of multipathing (route computation, load balancing, congestion control, etc.) is lacking, it is important to ensure that the various resource pooling mechanisms work well in harmony and do not conflict. Towards this end, there needs to be greater synergy between the end systems, the end networks, and intermediate autonomous systems so that potential mechanisms/ policy conflicts (that can lead to suboptimal performance) can be avoided.

As an example, P2P applications such as BitTorrent perform load balancing at the application layer independent of the ISPs by preferentially retrieving data on uncongested paths. Such an arrangement could come close to optimizing the cost for congestion pricing. However, with AS peering managed by the ISPs with arbitrary pricing models, the conflicting policies of the end systems and the ISPs can lead to a significant performance loss. It has been shown that the ‘cost of anarchy’ resulting from the end systems and the ISPs adopting different metrics for congestion can potentially be arbitrarily high [233] [234].

While network-layer multipathing solutions will play an important part in the future multipath-capable Internet, efficient resource pooling of multiple paths will eventually require cross-layer support and a seamlessly integrated resource sharing framework. It has been argued convincingly that the traditional inability of the Internet to progress beyond a single-path paradigm was less rooted in the lack of routing solutions but more so in the inflexible resource sharing model of the Internet which had an inevitable tussle between the hosts and the network [235]. This tussle restricted the overall promise of resource pooling since the independent policies of the stakeholders often differ—usually at the expense of each other. Effectively, the various stakeholders are performing load balancing in a siloed fashion (independently and while trying to be inconspicuous to others). This situation is far from ideal since the interaction between the stakeholders will be at best a form of commensalism (a situation in which one entity benefits while the other parties remain unaffected) and at worst destructive for each other.

In the EU-funded Trilogy project [200], there has been an explicit attempt to acknowledge and manage the tussles between the different stakeholders in the network. Trilogy projects propose a new tussle-aware Future Internet architecture that intrinsically incorporates multipath routing and transport while integrating technical solutions with socio-economic factors. Trilogy project proposed four general design principles for the future Internet [236] [237]. Firstly, the *Information Exposure Principle* requires the exposure of sufficient information about resource usage to support efficient and timely resource allocation. Secondly, the *Separation of Policy and Mechanism Principle* proposes to separate higher-level pol-




Fig. 6: A depiction of the holistic multipathing architecture proposed by the Trilog future Internet architecture.

icy decisions from standardized implementation mechanisms. Thirdly, the *Fuzzy Ends Principle* proposes a mutualistic framework in which the endpoints can delegate some functions to the network. Finally, the *Resource Pooling Principle*—most relevant to multipathing—dictates that the network resources should be pooled effectively so that spare capacity can be used effectively to counter overload and faults.

The Trilog architecture (shown in Figure 6) serves as a candidate embodiment of these principles. The overall architecture consists of a packet forwarding layer and a transport services layer. The packet forwarding layer comprises two players: (i) the *reachability plane* which is responsible for hop-by-hop routing necessary for enabling network-wide reachability, and (ii) the *resource plane* which is responsible for rate control and for deciding how to perform resource allocation between different packets/ flows. The transport services layer, on the other hand, manages the functionality of reliability, flow control, and message framing transparently from the packet forwarding layer. The transport layer is implemented end-to-end and interacts with the application layer.

Multipath TCP and *Multipath Routing*, along with *Congestion Exposure*, comprise the three building blocks of the Trilog architecture. Multipath TCP can directly exploit multiple paths on multi-homed endpoints whereas multipath routing allows the computation of disjoint paths even for single-homed devices (multipath theory indicates that the benefits of multipath can be reaped even without complete path disjointness as long as there is no shared path bottleneck). A synergy of transport and network layer multipathing is thus proposed in a holistic multipath solution that can allow faster traffic engineering (now performed end-to-end at the timescale of RTTs). Facilitating this interaction is ‘congestion exposure’ which enables all nodes to see the congestion they cause by sending traffic. Congestion exposure also allows the operator to judge which paths are good and should be made visible to the end users, while also being able to see if the traffic engineering performed by the end hosts is satisfactory and not malicious.

C. Multipathing and Future Internet Architectures

As explained in Section I, we anticipate that multipathing will be a big part of the future Internet. Many of the future

Internet architectures utilize multipathing in their framework. In this section, we will describe the role multipathing plays in the ICN and SDN architectures.

1) *Multipathing in Information-Centric Networks*: The ICN architecture employs content caching as an intrinsic and ubiquitous part of the networking infrastructure. A key aspect of ICN is to decouple location and content thus allowing content to be retrieved chunk-by-chunk from any location that is convenient. ICN inherently supports multipath since the a node’s ‘Interest’ for a content is not sent to a single location but can be sent through several interfaces with the replies being recomposed for application use. In addition to resource pooling of end-to-end resources and paths (as is done by multipath TCP protocols), resource pooling can be done on the basis of in-network resources (as is proposed in the ICN context with the in-network resource pooling work proposed recently [238]) to improve reliability, end-to-end throughput, network efficiency, and flexibility.

2) *Using SDN for Multipath Provisioning*: More research needs to be conducted to discover how to best use the centralized paradigm offered by SDN, along with its more flexible control abstractions, for managing and provisioning multipath routing. SDN has been recently used by McCormick et al. [209] to provide real-time alpha-fairness based traffic engineering. The proposed algorithm executes in the range of milliseconds and can be used as an online tool. Using the SDN architecture to split the control and data planes can also enable rapid response to network failures [239]. In particular, in a SDN-based setting, multipath routing can not only improve performance and reliability, but can also forestall the need to dynamically recompute routes in response to failures.

D. Deploying Diversity Efficiently

One of the main benefits of multipathing is the pluralism and diversity afforded by the presence of independent links. It has been observed in a wide variety of settings that systems exploiting diversity and pluralism can outperform individual systems that do not rely on diversity [27], [240]. In this regard, the following are some important points to consider.

- *Quantifying the Available Diversity*. Various quantitative measures of diversity have been developed in various sciences including entropy measures, variation measures, distance measures, and attribute-based measures [240]. Although it is unlikely that a single universal “one-size-fits-all” quantitative definition can be defined, there is a strong need of quantitative measures of multipath diversity to build a rigorous theory on multipath routing. Some initial works on using diversity index in networking have been proposed [241], [242], but there is a lot of scope for further work.
- *How to Use the Degrees-of-Freedom Afforded by Multipath?* The various degrees-of-freedom afforded by multipath can be used differently by application designers. For example, if we consider wireless communication at the physical layer, the availability of multiple paths can be used for either *multiplexing* or for *reliability* [243].

Similar options can also be used at the network layer, and the network designers should carefully evaluate how to use the available diversities.

- *Decorrelated and Disjointed Links/ Paths.* To truly reap the benefits of multipath, the pluralism and diversity available through multiple paths should be leveraged. Care should be adopted to ensure the independence of the selected links/ paths.
- *Avoiding too much diversity.* The idea that diversity works well in moderation but can backfire when overdone is well known in common folklore. We know that “two heads are better than one” and also that “too many cooks can spoil the broth”. While increased path diversity typically results in better performance and more robustness, there is generally a diminishing marginal utility with increasing diversity, and too much diversity can even potentially cause performance degradation [240]. In particular, coordination and management overheads (such as the control overhead of managing a large number of routing paths) can become prohibitively large with greater diversity. The implications of this need to be studied in the context of network-layer multipath solutions.

E. Multipathing and Green Networking

The problem of energy efficient multipathing is an important concern for battery operated mobile devices. While mobile devices can increase their throughput by striping their connections over heterogeneous networks, this comes at a cost of higher energy consumption. With the proliferation of mobile devices, the incorporation of energy efficiency mechanisms into the design of multipath solutions has become extremely important. Early work on energy-efficient multipathing has already shown the potential for improving smartphone’s energy consumption [244], [245], [246], [247].

Apart from its importance for wireless networks, energy efficiency is also critically important for datacenter networking (which operates at the scale of huge warehouses, where even minor savings can add up significantly). More research needs to be conducted in the area of energy efficient “green multipathing” to use the power of multipathing while reducing the energy consumption on the Internet.

V. CONCLUSIONS

In this paper, we have surveyed the existing literature on network-layer multipathing. After describing the motivation and benefits of multipath in networks, and establishing that the future of the Internet is multipath based, we discussed two important problems that are related to multipathing, namely the issue of how to compute the multiple routes to be used, and how should these multiple routes be used by traffic flows. While our focus is on network-layer multipathing, our coverage is holistic as we also discuss issues such as congestion control, flow splitting, and resource pooling (which do not fall neatly into a single layer). In addition to highlighting the main problems in network-layer multipathing, and describing the various approaches proposed, we have provided a broad-ranging survey of multipath protocols in different kinds of

networks. Despite a vast amount of work on network-layer multipathing, many open questions remain; we conclude this paper by highlighting some important issues that require further investigation.

ACKNOWLEDGMENT

Arjuna Sathiseelan and Jon Crowcroft were funded under the EU 7th Framework Programme, grant agreement number 317756, Trilogy II.

REFERENCES

- [1] D. Bertsekas and R. Gallager, “Data networks. 1992,” *PrenticeHall, Englewood Cliffs, NJ*, 1992.
- [2] S. Keshav, *An engineering approach to computer networking: ATM networks, the Internet, and the telephone network.* Addison-Wesley Longman Publishing Co., Inc., 1997.
- [3] M. Wang, C. W. Tan, W. Xu, and A. Tang, “Cost of not splitting in routing: Characterization and estimation,” *Networking, IEEE/ACM Transactions on*, vol. 19, no. 6, pp. 1849–1859, 2011.
- [4] M. Wang, C. wei Tan, A. Tang, and S. Low, “How bad is single-path routing,” in *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*, Nov 2009, pp. 1–6.
- [5] Z. Wang and J. Crowcroft, “Analysis of shortest-path routing algorithms in a dynamic network environment,” *ACM SIGCOMM Computer Communication Review*, vol. 22, no. 2, pp. 63–71, 1992.
- [6] D. P. Bertsekas, R. G. Gallager, and P. Humblet, *Data networks.* Prentice-hall Englewood Cliffs, NJ, 1987, vol. 2.
- [7] J. M. McQuillan and D. C. Walden, “The ARPA network design decisions,” *Computer Networks (1976)*, vol. 1, no. 5, pp. 243–289, 1977.
- [8] J. M. McQuillan, I. Richer, and E. Rosen, “The new routing algorithm for the ARPANET,” *Communications, IEEE Transactions on*, vol. 28, no. 5, pp. 711–719, 1980.
- [9] A. Khanna and J. Zinky, “The revised ARPANET routing metric,” in *ACM SIGCOMM Computer Communication Review*, vol. 19, no. 4. ACM, 1989, pp. 45–56.
- [10] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman, “A measurement-based analysis of multihoming,” in *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications.* ACM, 2003, pp. 353–364.
- [11] C. Paasch and O. Bonaventure, “Multipath tcp,” *Commun. ACM*, vol. 57, no. 4, pp. 51–57, Apr. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2578901>
- [12] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, “In search of path diversity in ISP networks,” in *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement.* ACM, 2003, pp. 313–318.
- [13] N. Spring, R. Mahajan, and D. Wetherall, “Measuring ISP topologies with rocketfuel,” in *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4. ACM, 2002, pp. 133–145.
- [14] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, “The end-to-end effects of internet path selection,” in *ACM SIGCOMM Computer Communication Review*, vol. 29, no. 4. ACM, 1999, pp. 289–299.
- [15] J. He and J. Rexford, “Toward internet-wide multipath routing,” *Network, IEEE*, vol. 22, no. 2, pp. 16–21, 2008.
- [16] D. Wischik, M. Handley, and M. B. Braun, “The resource pooling principle,” *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 5, pp. 47–52, 2008.
- [17] L. G. Valiant, “A bridging model for parallel computation,” *Commun. ACM*, vol. 33, no. 8, pp. 103–111, Aug. 1990. [Online]. Available: <http://doi.acm.org/10.1145/79173.79181>
- [18] R. Zhang-Shen, “Valiant load-balancing: Building networks that can support all traffic matrices,” in *Algorithms for Next Generation Networks.* Springer, 2010, pp. 19–30.
- [19] R. G. Clegg, R. Landa, J. T. Araújo, E. Mykoniati, D. Griffin, and M. Rio, “Tardis: stably shifting traffic in space and time,” in *The 2014 ACM international conference on Measurement and modeling of computer systems.* ACM, 2014, pp. 593–594.
- [20] V. Paxson and S. Floyd, “Wide area traffic: the failure of poisson modeling,” *IEEE/ACM Transactions on Networking (ToN)*, vol. 3, no. 3, pp. 226–244, 1995.

- [21] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *Networking, IEEE/ACM Transactions on*, vol. 2, no. 1, pp. 1–15, 1994.
- [22] P. Bailis and K. Kingsbury, "The network is reliable," *Queue*, vol. 12, no. 7, p. 20, 2014.
- [23] T. Ye, D. Veitch, and J. Bolot, "Improving wireless security through network diversity," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 34–44, 2008.
- [24] C. Paasch, G. Detal, S. Barre, F. Duchene, and O. Bonaventure, "The fastest TCP connection with multipath TCP," Available: <http://multipath-tcp.org/pmwiki.php?n=Main.50Gbps>, 2013.
- [25] D. Kahneman, *Thinking, fast and slow*. Macmillan, 2011.
- [26] P. Tetlock, *Expert political judgment: How good is it? How can we know?* Princeton University Press, 2005.
- [27] J. Surowiecki, *The wisdom of crowds*. Random House LLC, 2005.
- [28] J. G. Andrews, S. Buzzi, W. Choi, S. Hanly, A. Lozano, A. C. Soong, and J. C. Zhang, "What will 5G be?" *arXiv preprint arXiv:1405.2957*, 2014.
- [29] M. A. P. Gonzalez, T. Higashino, and M. Okada, "Radio access considerations for data offloading with multipath tcp in cellular/wifi networks," in *Information Networking (ICOIN), 2013 International Conference on*. IEEE, 2013, pp. 680–685.
- [30] N. Vallina-Rodriguez, V. Erramilli, Y. Grunenberger, L. Gyarmati, N. Laoutaris, R. Stanojevic, and K. Papagiannaki, "When David helps Goliath: the case for 3G onloading," in *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*. ACM, 2012, pp. 85–90.
- [31] G. Carofiglio, M. Gallo, L. Muscariello, and M. Papalini, "Multipath congestion control in content-centric networks," *IEEE NOMEN*, vol. 13, 2013.
- [32] G. Rossini and D. Rossi, "Evaluating CCN multi-path interest forwarding strategies," *Computer Communications*, vol. 36, no. 7, pp. 771–778, 2013.
- [33] G. Tyson, S. Kaune, S. Miles, Y. El-khatib, A. Mauthe, and A. Taweel, "A trace-driven analysis of caching in content-centric networks," in *Computer Communications and Networks (ICCCN), 2012 21st International Conference on*. IEEE, 2012, pp. 1–7.
- [34] M. Radi, B. Dezfouli, K. A. Bakar, and M. Lee, "Multipath routing in wireless sensor networks: survey and research challenges," *Sensors*, vol. 12, no. 1, pp. 650–685, 2012.
- [35] K. Sha, J. Gehlot, and R. Greve, "Multipath routing techniques in wireless sensor networks: A survey," *Wireless personal communications*, vol. 70, no. 2, pp. 807–829, 2013.
- [36] M. Tarique, K. E. Tepe, S. Adibi, and S. Erfani, "Survey of multipath routing protocols for mobile ad hoc networks," *Journal of Network and Computer Applications*, vol. 32, no. 6, pp. 1125–1143, 2009.
- [37] S. Mueller, R. P. Tsang, and D. Ghosal, "Multipath routing in mobile ad hoc networks: Issues and challenges," in *Performance tools and applications to networked systems*. Springer, 2004, pp. 209–234.
- [38] W. Lou, W. Liu, and Y. Zhang, "Performance optimization using multipath routing in mobile ad hoc and wireless sensor networks," in *Combinatorial optimization in communication networks*. Springer, 2006, pp. 117–146.
- [39] G. Parissidis, V. Lenders, M. May, and B. Plattner, "Multi-path routing protocols in wireless mobile ad hoc networks: A quantitative comparison," in *Next Generation Teletraffic and Wired/Wireless Advanced Networking*. Springer, 2006, pp. 313–326.
- [40] A. Nasipuri and S. R. Das, "On-demand multipath routing for mobile ad hoc networks," in *Computer Communications and Networks, 1999. Proceedings. Eight International Conference on*. IEEE, 1999, pp. 64–70.
- [41] L. Wang, L. Zhang, Y. Shu, and M. Dong, "Multipath source routing in wireless ad hoc networks," in *Electrical and Computer Engineering, 2000 Canadian Conference on*, vol. 1. IEEE, 2000, pp. 479–483.
- [42] Y. Seok, Y. Lee, Y. Choi, and C. Kim, "Dynamic constrained multipath routing for MPLS networks," in *Computer Communications and Networks, 2001. Proceedings. Tenth International Conference on*. IEEE, 2001, pp. 348–353.
- [43] D. Zhu, M. Gritter, and D. R. Cheriton, "Feedback based routing," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 1, pp. 71–76, 2003.
- [44] H. T. Kaur, S. Kalyanaraman, A. Weiss, S. Kanwar, and A. Gandhi, "BANANAS: an evolutionary framework for explicit and multipath routing in the Internet," in *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 4. ACM, 2003, pp. 277–288.
- [45] P. K. Gummadi, H. V. Madhyastha, S. D. Gribble, H. M. Levy, D. Wetherall *et al.*, "Improving the reliability of internet paths with one-hop source routing," in *OSDI*, vol. 4, 2004, pp. 13–13.
- [46] W. Lou, W. Liu, and Y. Fang, "SPREAD: Enhancing data confidentiality in mobile ad hoc networks," in *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 4. IEEE, 2004, pp. 2404–2413.
- [47] X. Yang and D. Wetherall, "Source selectable path diversity via routing deflections," in *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4. ACM, 2006, pp. 159–170.
- [48] W.-H. Tam and Y.-C. Tseng, "Joint multi-channel link layer and multipath routing design for wireless mesh networks," in *INFOCOM, 2007*, pp. 2081–2089.
- [49] X. Yang, D. Clark, and A. W. Berger, "NIRA: a new inter-domain routing architecture," *Networking, IEEE/ACM Transactions on*, vol. 15, no. 4, pp. 775–788, 2007.
- [50] M. Motiwala, M. Elmore, N. Feamster, and S. Vempala, "Path splicing," in *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4. ACM, 2008, pp. 27–38.
- [51] P. Godfrey, I. Ganichev, S. Shenker, and I. Stoica, "Pathlet routing," in *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 4. ACM, 2009, pp. 111–122.
- [52] G. T. Nguyen, R. Agarwal, J. Liu, M. Caesar, P. Godfrey, and S. Shenker, "Slick packets," *ACM SIGMETRICS Performance Evaluation Review*, vol. 39, no. 1, pp. 205–216, 2011.
- [53] D. Qin, J. Yang, Z. Liu, H. Wang, B. Zhang, and W. Zhang, "AMIR: another multipath interdomain routing," in *Advanced Information Networking and Applications (AINA), 2012 IEEE 26th International Conference on*. IEEE, 2012, pp. 581–588.
- [54] N. F. Maxemchuk, "Dispersy routing in store-and-forward networks," Ph.D. dissertation, University of Pennsylvania, Philadelphia, 1975.
- [55] S. Wu and K. S. Candan, "Gper: Geographic power efficient routing in sensor networks," in *Network Protocols, 2004. ICNP 2004. Proceedings of the 12th IEEE International Conference on*. IEEE, 2004, pp. 161–172.
- [56] S. Vutukury and J. Garcia-Luna-Aceves, "A simple approximation to minimum-delay routing," in *Proc. of ACM SIGCOMM, 1999*.
- [57] —, "An algorithm for multipath computation using distance-vectors with predecessor information," in *Computer Communications and Networks, 1999. Proceedings. Eight International Conference on*. IEEE, 1999, pp. 534–539.
- [58] S. Vutukury and J. J. Garcia-Luna-Aceves, "MDVA: A distance-vector multipath routing protocol," in *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2001, pp. 557–564.
- [59] Z. Cao, Z. Wang, and E. Zegura, "Performance of hashing-based schemes for internet load balancing," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2000, pp. 332–341.
- [60] A. Srinivas and E. Modiano, "Minimum energy disjoint path routing in wireless ad-hoc networks," in *Proceedings of the 9th annual international conference on Mobile computing and networking*. ACM, 2003, pp. 122–133.
- [61] A. Valera, W. K. G. Seah, and S. Rao, "Cooperative packet caching and shortest multipath routing in mobile ad hoc networks," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, vol. 1. IEEE, 2003, pp. 260–269.
- [62] W. Zhang, J. Tang, C. Wang, and S. De Soysa, "Reliable adaptive multipath provisioning with bandwidth and differential delay constraints," in *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010, pp. 1–9.
- [63] S. De, C. Qiao, and H. Wu, "Meshed multipath routing with selective forwarding: an efficient strategy in wireless sensor networks," *Computer Networks*, vol. 43, no. 4, pp. 481–497, 2003.
- [64] —, "Meshed multipath routing: an efficient strategy in sensor networks," in *Wireless Communications and Networking, 2003. WCNC 2003. 2003 IEEE*, vol. 3. IEEE, 2003, pp. 1912–1917.
- [65] M. Maimour, "Maximally radio-disjoint multipath routing for wireless multimedia sensor networks," in *Proceedings of the 4th ACM workshop on Wireless multimedia networking and performance modeling*. ACM, 2008, pp. 26–31.
- [66] M. K. Marina and S. R. Das, "On-demand multipath distance vector routing in ad hoc networks," in *Network Protocols, 2001. Ninth International Conference on*. IEEE, 2001, pp. 14–23.
- [67] Z. Ye, S. V. Krishnamurthy, and S. K. Tripathi, "A framework for reliable routing in mobile ad hoc networks," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, vol. 1. IEEE, 2003, pp. 270–280.

- [68] P. Sambasivam, A. Murthy, and E. M. Belding-Royer, "Dynamically adaptive multipath routing based on AODV," in *Proc. 3rd Annual Mediterranean Ad Hoc Networking Workshop*, 2004.
- [69] X. Li and L. Cuthbert, "On-demand node-disjoint multipath routing in wireless ad hoc networks," in *Local Computer Networks, 2004. 29th Annual IEEE International Conference on*. IEEE, 2004, pp. 419–420.
- [70] G. Di Caro, F. Ducatelle, and L. M. Gambardella, "AntHocNet: an adaptive nature-inspired algorithm for routing in mobile ad hoc networks," *European Transactions on Telecommunications*, vol. 16, no. 5, pp. 443–455, 2005.
- [71] F. Ducatelle, G. Di Caro, and L. M. Gambardella, "Ant agents for hybrid multipath routing in mobile ad hoc networks," in *Wireless On-demand Network Systems and Services, 2005. WONS 2005. Second Annual Conference on*. IEEE, 2005, pp. 44–53.
- [72] N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs, "R-BGP: Staying connected in a connected world," in *NSDI*. USENIX, 2007.
- [73] I. Van Beijnum, J. Crowcroft, F. Valera, and M. Bagnulo, "Loop-freeness in multipath BGP through propagating the longest path," in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*. IEEE, 2009, pp. 1–6.
- [74] Y. Ohara, S. Imahori, and R. Van Meter, "MARA: Maximum alternative routing algorithm," in *INFOCOM 2009*. IEEE, 2009, pp. 298–306.
- [75] X. Huang and Y. Fang, "Performance study of node-disjoint multipath routing in vehicular ad hoc networks," *Vehicular Technology, IEEE Transactions on*, vol. 58, no. 4, pp. 1942–1950, 2009.
- [76] Y. Yang, C. Zhong, Y. Sun, and J. Yang, "Network coding based reliable disjoint and braided multipath routing for sensor networks," *Journal of Network and Computer Applications*, vol. 33, no. 4, pp. 422–432, 2010.
- [77] I. Ganichev, B. Dai, P. Godfrey, and S. Shenker, "YAMR: Yet another multipath routing protocol," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 5, pp. 13–19, 2010.
- [78] I. Beltagy, M. Youssef, and M. El-Derini, "A new routing metric and protocol for multipath routing in cognitive networks," in *Wireless Communications and Networking Conference (WCNC), 2011 IEEE*. IEEE, 2011, pp. 974–979.
- [79] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "The case for resilient overlay networks," in *Hot Topics in Operating Systems, 2001. Proceedings of the Eighth Workshop on*. IEEE, 2001, pp. 152–157.
- [80] Z. Ma, H.-R. Shao, C. Shen *et al.*, "A new multi-path selection scheme for video streaming on overlay networks," in *ICC*, 2004, pp. 1330–1334.
- [81] A. C. Begen, Y. Altunbasak, O. Ergun, and M. H. Ammar, "Multipath selection for multiple description video streaming over overlay networks," *Signal Processing: Image Communication*, vol. 20, no. 1, pp. 39–60, 2005.
- [82] S.-J. Lee, S. Banerjee, P. Sharma, P. Yalagandula, and S. Basu, "Bandwidth-aware routing in overlay networks," in *INFOCOM 2008. The 27th Conference on Computer Communications*. IEEE, 2008.
- [83] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: Dynamic flow scheduling for data center networks," in *NSDI*, vol. 10, 2010, pp. 19–19.
- [84] T. Benson, A. Anand, A. Akella, and M. Zhang, "MicroTE: fine grained traffic engineering for data centers," in *Proceedings of the Seventh Conference on emerging Networking EXperiments and Technologies (CoNeXT'11)*. ACM, 2011, p. 8.
- [85] B. Stephens, "Designing scalable networks for future large datacenters," Ph.D. dissertation, MS thesis, Rice University, 2012.
- [86] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu *et al.*, "B4: Experience with a globally-deployed software defined WAN," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. ACM, 2013, pp. 3–14.
- [87] A. Detti, C. Pisa, S. Salsano, and N. Blefari-Melazzi, "Wireless mesh software defined networks (wmsdn)," in *Wireless and Mobile Computing, Networking and Communications (WiMob), 2013 IEEE 9th International Conference on*. IEEE, 2013, pp. 89–95.
- [88] I. Ku, Y. Lu, M. Gerla, R. L. Gomes, F. Ongaro, and E. Cerqueira, "Towards software-defined vanet: Architecture and services," *Med-Hoc-Net*, 2014.
- [89] C. Villamizar, "OSPF optimized multipath (OSPF-OMP)," *Internet Draft*, <http://tools.ietf.org/html/draft-ietf-ospf-omp-01>, 1999.
- [90] —, "MPLS optimized multipath (MPLS-OMP)," 1999.
- [91] C. E. Hopps, "Analysis of an equal-cost multi-path algorithm," *IETF RFC 2992*, 2000.
- [92] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. N. Rao, "Improving web availability for clients with MONET," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation-Volume 2*. USENIX Association, 2005, pp. 115–128.
- [93] W. Xu and J. Rexford, *MIRO: multi-path interdomain routing*. ACM, 2006, vol. 36, no. 4.
- [94] F. Wang and L. Gao, "Path diversity aware interdomain routing," in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 307–315.
- [95] F. Valera, I. van Beijnum, A. Garca-Martnez, and M. Bagnulo, "Multi-path bgp: motivations and solutions," in *Next-Generation Internet Architectures and Protocols*, B. Ramamurthy, G. N. Rouskas, and K. M. Sivalingam, Eds. Cambridge University Press, 2011, pp. 238–256, cambridge Books Online. [Online]. Available: <http://dx.doi.org/10.1017/CBO9780511920950.013>
- [96] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly-resilient, energy-efficient multipath routing in wireless sensor networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 11–25, 2001.
- [97] E. Felemban, C.-G. Lee, and E. Ekici, "MMSPEED: multipath multi-speed protocol for qos guarantee of reliability and. timeliness in wireless sensor networks," *Mobile Computing, IEEE Transactions on*, vol. 5, no. 6, pp. 738–754, 2006.
- [98] W. Lou and Y. Kwon, "H-SPREAD: a hybrid multipath scheme for secure and reliable data collection in wireless sensor networks," *Vehicular Technology, IEEE Transactions on*, vol. 55, no. 4, pp. 1320–1330, 2006.
- [99] Y. Chen and N. Nasser, "Energy-balancing multipath routing protocol for wireless sensor networks," in *Proceedings of the 3rd international conference on Quality of service in heterogeneous wired/wireless networks*. ACM, 2006, p. 21.
- [100] S. Waharte and R. Boutaba, "Totally disjoint multipath routing in multihop wireless networks," in *ICC*, 2006, pp. 5576–5581.
- [101] S. Wu and K. S. Candan, "Power-aware single-and multipath geographic routing in sensor networks," *Ad Hoc Networks*, vol. 5, no. 7, pp. 974–997, 2007.
- [102] S. Ramasubramanian, H. Krishnamoorthy, and M. Krunk, "Disjoint multipath routing using colored trees," *Computer Networks*, vol. 51, no. 8, pp. 2163–2180, 2007.
- [103] N. Nasser and Y. Chen, "SEEM: Secure and energy-efficient multipath routing protocol for wireless sensor networks," *Computer Communications*, vol. 30, no. 11, pp. 2401–2412, 2007.
- [104] H. Wu and X. Jia, "QoS multicast routing by using multiple paths/trees in wireless ad hoc networks," *Ad Hoc Networks*, vol. 5, no. 5, pp. 600–612, 2007.
- [105] Y. Ming Lu and V. WS Wong, "An energy-efficient multipath routing protocol for wireless sensor networks," *International Journal of Communication Systems*, vol. 20, no. 7, pp. 747–766, 2007.
- [106] X. Huang and Y. Fang, "Multiconstrained QoS multipath routing in wireless sensor networks," *Wireless Networks*, vol. 14, no. 4, pp. 465–478, 2008.
- [107] B. Yahya and J. Ben-Othman, "REER: Robust and energy efficient multipath routing protocol for wireless sensor networks," in *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*. IEEE, 2009, pp. 1–7.
- [108] Y. Challal, A. Ouadjaout, N. Lasla, M. Bagaa, and A. Hadjidj, "Secure and efficient disjoint multipath construction for fault tolerant routing in wireless sensor networks," *Journal of Network and Computer Applications*, vol. 34, no. 4, pp. 1380–1397, 2011.
- [109] M. Radi, B. Dezfouli, K. A. Bakar, S. A. Razak, and T. Hwee-Pink, "Im2pr: interference-minimized multipath routing protocol for wireless sensor networks," *Wireless Networks*, pp. 1–17, 2014.
- [110] S.-J. Lee and M. Gerla, "Split multipath routing with maximally disjoint paths in ad hoc networks," in *Communications, 2001. ICC 2001. IEEE International Conference on*, vol. 10. IEEE, 2001, pp. 3201–3205.
- [111] R. Leung, J. Liu, E. Poon, A.-L. Chan, and B. Li, "MP-DSR: a QoS-aware multi-path dynamic source routing protocol for wireless ad-hoc networks," in *Local Computer Networks, 2001. Proceedings. LCN 2001. 26th Annual IEEE Conference on*. IEEE, 2001, pp. 132–141.
- [112] W. Wei and A. Zakhor, "Multipath unicast and multicast video communication over wireless ad hoc networks," in *Broadband Networks, 2004. BroadNets 2004. Proceedings. First International Conference on*. IEEE, 2004, pp. 496–505.
- [113] C. Chen, W. Wu, and Z. Li, "Multipath routing modeling in ad hoc networks," in *Communications, 2005. ICC 2005. 2005 IEEE International Conference on*, vol. 5. IEEE, 2005, pp. 2974–2978.

- [114] M. Kun, Y. Jingdong, and R. Zhi, "The research and simulation of multipath-OLSR for mobile ad hoc network," in *Communications and Information Technology, 2005. ISCT 2005. IEEE International Symposium on*, vol. 1. IEEE, 2005, pp. 540–543.
- [115] J. J. Galvez, P. M. Ruiz, and A. Skarmeta, "Spatially disjoint multipath routing protocol without location information," in *Local Computer Networks, 2008. LCN 2008. 33rd IEEE Conference on*. IEEE, 2008, pp. 570–571.
- [116] M. Caleffi, G. Ferraiuolo, and L. Paura, "A reliability-based framework for multi-path routing analysis in mobile ad-hoc networks," *International Journal of Communication Networks and Distributed Systems*, vol. 1, no. 4, pp. 507–523, 2008.
- [117] I. Ku, Y. Lu, and M. Gerla, "Software-defined mobile cloud: Architecture, services and use cases," *International Wireless Communications and Mobile Computing Conference (IWCMC)*, 2014.
- [118] J. W. Tsai and T. Moors, "Interference-aware multipath selection for reliable routing in wireless mesh networks," in *Mobile Adhoc and Sensor Systems, 2007. MASS 2007. IEEE International Conference on*. IEEE, 2007, pp. 1–6.
- [119] R. Wang, C. Rezende, H. S. Ramos, R. W. Pazzi, A. Boukerche, and A. A. F. Loureiro, "Liaithon: A location-aware multipath video streaming scheme for urban vehicular networks," in *Computers and Communications (ISCC), 2012 IEEE Symposium on*. IEEE, 2012, pp. 000436–000441.
- [120] B. Pavković, F. Theoleyre, and A. Duda, "Multipath opportunistic RPL routing over IEEE 802.15.4," in *Proceedings of the 14th ACM international conference on Modeling, analysis and simulation of wireless and mobile systems*. ACM, 2011, pp. 179–186.
- [121] Y. Lee, Y. Seok, Y. Choi, and C. Kim, "A constrained multipath traffic engineering scheme for mpls networks," in *Communications, 2002. ICC 2002. IEEE International Conference on*, vol. 4. IEEE, 2002, pp. 2431–2436.
- [122] P. Jokela, A. Zahemszky, C. Esteve Rothenberg, S. Arianfar, and P. Nikander, "Lipsin: line speed publish/subscribe inter-networking," in *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 4. ACM, 2009, pp. 195–206.
- [123] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile computing*. Springer, 1996, pp. 153–181.
- [124] E. Stavrou and A. Pitsillides, "A survey on secure multipath routing protocols in WSNs," *Computer Networks*, vol. 54, no. 13, pp. 2215–2238, 2010.
- [125] J. W. Hui, "The routing protocol for low-power and lossy networks (rpl) option for carrying rpl information in data-plane datagrams," *IETF RFC 6553*; <http://tools.ietf.org/html/rfc6553>, 2012.
- [126] J. G. Apostolopoulos, "Reliable video communication over lossy packet networks using multiple state encoding and path diversity," in *Photonics West 2001-Electronic Imaging*. International Society for Optics and Photonics, 2000, pp. 392–409.
- [127] J. Apostolopoulos, T. Wong, W.-t. Tan, and S. Wee, "On multiple description streaming with content delivery networks," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3. IEEE, 2002, pp. 1736–1745.
- [128] V. K. Goyal, "Multiple description coding: Compression meets the network," *Signal Processing Magazine, IEEE*, vol. 18, no. 5, pp. 74–93, 2001.
- [129] M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: substrate support for path splitting and migration," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 17–29, 2008.
- [130] D. Dietrich, A. Rizk, and P. Papadimitriou, "Multi-domain virtual network embedding with limited information disclosure," in *IFIP Networking Conference, 2013*. IEEE, 2013, pp. 1–9.
- [131] D. Sidhu, R. Nair, and S. Abdallah, "Finding disjoint paths in networks," *ACM SIGCOMM Computer Communication Review*, vol. 21, no. 4, pp. 43–51, 1991.
- [132] Y. Guo, F. Kuipers, and P. Van Mieghem, "Link-disjoint paths for reliable qos routing," *International Journal of Communication Systems*, vol. 16, no. 9, pp. 779–798, 2003.
- [133] J. Liao, J. Wang, T. Li, and X. Zhu, "Introducing multipath selection for concurrent multipath transfer in the future internet," *Computer Networks*, vol. 55, no. 4, pp. 1024–1035, 2011.
- [134] N. Vasić and D. Kostić, "Energy-aware traffic engineering," in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. ACM, 2010, pp. 169–178.
- [135] A. Tsigros and Z. J. Haas, "Multipath routing in the presence of frequent topological changes," *Communications Magazine, IEEE*, vol. 39, no. 11, pp. 132–138, 2001.
- [136] Y. Li, S. Mao, and S. S. Panwar, "The case for multipath multimedia transport over wireless ad hoc networks," in *Broadband Networks, 2004. BroadNets 2004. Proceedings. First International Conference on*. IEEE, 2004, pp. 486–495.
- [137] L. Gui, S. Zou, and X. Zhong, "Distributed best-route selection for multipath routing in cognitive radio ad hoc networks," *Electronics Letters*, vol. 48, no. 25, pp. 1630–1632, 2012.
- [138] C. Karlof and D. Wagner, "Secure routing in wireless sensor networks: Attacks and countermeasures," *Ad hoc networks*, vol. 1, no. 2, pp. 293–315, 2003.
- [139] N. S. Nandiraju, D. S. Nandiraju, and D. P. Agrawal, "Multipath routing in wireless mesh networks," in *Mobile adhoc and sensor systems (MASS), 2006 IEEE international conference on*. IEEE, 2006, pp. 741–746.
- [140] B. Radunovic, C. Gkantsidis, P. Key, and P. Rodriguez, "An optimization framework for opportunistic multipath routing in wireless mesh networks," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*. IEEE, 2008.
- [141] X. Wang, T. T. Kwon, and Y. Choi, "A multipath routing and spectrum access (mrsa) framework for cognitive radio systems in multi-radio mesh networks," in *Proceedings of the 2009 ACM workshop on Cognitive radio networks*. ACM, 2009, pp. 55–60.
- [142] B. Khanna and R. Ramanathan, "Interference-aware multipath routing in a cognitive radio ad hoc network," in *MILITARY COMMUNICATIONS CONFERENCE, 2011-MILCOM 2011*. IEEE, 2011, pp. 855–860.
- [143] S.-J. Lee and M. Gerla, "AODV-BR: Backup routing in ad hoc networks," in *Wireless Communications and Networking Conference, 2000. WCNC. 2000 IEEE*, vol. 3. IEEE, 2000, pp. 1311–1316.
- [144] K. Wu and J. Harms, "Performance study of a multipath routing method for wireless mobile ad hoc networks," in *Modeling, analysis and simulation of computer and telecommunication systems, 2001. Proceedings. Ninth International Symposium on*. IEEE, 2001, pp. 99–107.
- [145] L. G. Valiant and G. J. Brebner, "Universal schemes for parallel communication," in *Proceedings of the thirteenth annual ACM symposium on Theory of computing*. ACM, 1981, pp. 263–277.
- [146] L. G. Valiant, "A scheme for fast parallel communication," *SIAM journal on computing*, vol. 11, no. 2, pp. 350–361, 1982.
- [147] R. Zhang-Shen and N. McKeown, "Designing a predictable internet backbone with valiant load-balancing," in *Quality of Service-IWQoS 2005*. Springer, 2005, pp. 178–192.
- [148] —, "Designing a fault-tolerant network using valiant load-balancing," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*. IEEE, 2008.
- [149] S. Dulman, J. Wu, and P. Havinga, "An energy efficient multipath routing algorithm for wireless sensor networks," *University of Twente, Centre for Telematics and Information Technology (CTIT)*, 2003.
- [150] H. Liu, Z.-L. Zhang, J. Srivastava, and V. Firoiu, "PWave: A multi-source multi-sink anycast routing framework for wireless sensor networks," in *NETWORKING 2007. Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet*. Springer, 2007, pp. 179–190.
- [151] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: a survey," *Computer networks*, vol. 47, no. 4, pp. 445–487, 2005.
- [152] J. Qadir, N. Ahmed, and N. Ahad, "Building programmable wireless networks: an architectural survey," *EURASIP Journal on Wireless Communications and Networking*, vol. 2014, no. 1, p. 172, 2014. [Online]. Available: <http://jwcn.eurasipjournals.com/content/2014/1/172>
- [153] J. Qadir, N. Ahad, E. Mushtaq, and M. Bilal, "SDN, clouds, and big data: New opportunities," *12th International Conference on Frontiers of Information (FIT)*, 2014.
- [154] D. S. Phatak and T. Goff, "A novel mechanism for data streaming across multiple IP links for improving throughput and reliability in mobile environments," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2002, pp. 773–781.
- [155] M. Shreedhar and G. Varghese, "Efficient fair queueing using deficit round robin," in *ACM SIGCOMM Computer Communication Review*, vol. 25, no. 4. ACM, 1995, pp. 231–242.
- [156] J. Moy, "OSPF version 2," *IETF RFC 2328*, 1998.
- [157] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *INFOCOM 2000. Nineteenth Annual Joint*

- Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2000, pp. 519–528.
- [158] —, “Optimizing OSPF/IS weights in a changing world,” *IEEE journal on selected areas in communications*, vol. 20, no. 4, 2002.
- [159] B. Fortz, J. Rexford, and M. Thorup, “Traffic engineering with traditional IP routing protocols,” *Communications Magazine, IEEE*, vol. 40, no. 10, pp. 118–124, 2002.
- [160] A. Sridharan, R. Guerin, and C. Diot, “Achieving near-optimal traffic engineering solutions for current OSPF/IS-IS networks,” *IEEE/ACM Transactions on Networking (TON)*, vol. 13, no. 2, pp. 234–247, 2005.
- [161] J. H. Fong, A. C. Gilbert, S. Kannan, and M. J. Strauss, “Better alternatives to OSPF routing,” *Algorithmica*, vol. 43, no. 1-2, pp. 113–131, 2005.
- [162] D. Xu, M. Chiang, and J. Rexford, “DEFT: Distributed exponentially-weighted flow splitting,” in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*. IEEE, 2007, pp. 71–79.
- [163] —, “Link-state routing with hop-by-hop forwarding can achieve optimal traffic engineering,” *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 6, pp. 1717–1730, 2011.
- [164] A. Elwalid, C. Jin, S. Low, and I. Widjaja, “MATE: MPLS adaptive traffic engineering,” in *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3. IEEE, 2001, pp. 1300–1309.
- [165] S. Kandula, D. Katabi, B. Davie, and A. Charney, “TeXCP: Responsive yet stable traffic engineering,” in *Proc. ACM SIGCOMM*, 2005.
- [166] H. Wang, H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg, “COPE: traffic engineering in dynamic networks,” in *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4. ACM, 2006, pp. 99–110.
- [167] J. He, M. Bresler, M. Chiang, and J. Rexford, “Towards robust multi-layer traffic engineering: Optimization of congestion control and routing,” *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 5, pp. 868–880, 2007.
- [168] S. Fischer, N. Kammenhuber, and A. Feldmann, “REPLEX: dynamic traffic engineering based on wardrop routing policies,” in *Proceedings of the 2006 ACM CoNEXT conference*. ACM, 2006, p. 1.
- [169] J. G. Wardrop, “Some theoretical aspects of road traffic research,” in *ICE Proceedings: Engineering Divisions*, vol. 1, no. 3. Thomas Telford, 1952, pp. 325–362.
- [170] S. Kandula, D. Katabi, S. Sinha, and A. Berger, “Dynamic load balancing without packet reordering,” *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 2, pp. 51–62, 2007.
- [171] J. T. Araújo, M. Rio, and G. Pavlou, “A mutualistic resource pooling architecture,” in *Proceedings of the Re-Architecting the Internet Workshop*. ACM, 2010, p. 3.
- [172] A. Ford, C. Raiciu, M. Handley, S. Barre, J. Iyengar *et al.*, “Architectural guidelines for multipath TCP development,” *IETF, Informational RFC*, vol. 6182, pp. 2070–1721, 2011.
- [173] F. Zhong, C. K. Yeo, and B. S. Lee, “Adaptive load balancing algorithm for multi-homing mobile nodes in local domain,” in *Consumer Communications and Networking Conference (CCNC), 2011 IEEE*. IEEE, 2011, pp. 482–486.
- [174] N. Michael, A. Tang, and D. Xu, “Optimal link-state hop-by-hop routing,” in *ICNP*, 2013, pp. 1–10.
- [175] J. He, M. Suchara, M. Bresler, J. Rexford, and M. Chiang, “From multiple decompositions to TRUMP: Traffic management using multipath protocol,” in *Elsevier Computer Networks*, 2008.
- [176] F. P. Kelly, A. K. Maulloo, and D. K. Tan, “Rate control for communication networks: shadow prices, proportional fairness and stability,” *Journal of the Operational Research society*, vol. 49, no. 3, pp. 237–252, 1998.
- [177] Y. Wang, Z. Wang, and L. Zhang, “Internet traffic engineering without full mesh overlaying,” in *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2001, pp. 565–571.
- [178] F. Kelly and T. Voice, “Stability of end-to-end algorithms for joint routing and rate control,” *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 2, pp. 5–12, 2005.
- [179] H. Han, S. Shakkottai, C. Hollot, R. Srikant, and D. Towsley, “Multipath TCP: a joint congestion control and routing scheme to exploit path diversity in the internet,” *IEEE/ACM Transactions on Networking (TON)*, vol. 14, no. 6, pp. 1260–1271, 2006.
- [180] R. Banner and A. Orda, “Multipath routing algorithms for congestion minimization,” *IEEE/ACM Transactions on Networking (TON)*, vol. 15, no. 2, pp. 413–424, 2007.
- [181] I. Cidon, R. Rom, and Y. Shavitt, “Analysis of multi-path routing,” *IEEE/ACM Transactions on Networking (TON)*, vol. 7, no. 6, pp. 885–896, 1999.
- [182] M. Mitzenmacher, “The power of two choices in randomized load balancing,” *Parallel and Distributed Systems, IEEE Transactions on*, vol. 12, no. 10, pp. 1094–1104, 2001.
- [183] P. Key, L. Massoulié, and P. Towsley, “Path selection and multipath congestion control,” in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*. IEEE, 2007, pp. 143–151.
- [184] P. Key, L. Massoulié, and D. Towsley, “Path selection and multipath congestion control,” *Communications of the ACM*, vol. 54, no. 1, pp. 109–116, 2011.
- [185] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan, “Wifi, lte, or both?: Measuring multi-homed wireless internet performance,” in *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 2014, pp. 181–194.
- [186] M. Chiesa, I. Nikolaevkiy, A. Panda, A. Gurtov, M. Schapira, and S. Shenker, “Exploring the limits of static failover routing,” *arXiv preprint arXiv:1409.0034*, 2014.
- [187] H. Adishesu, G. Parulkar, and G. Varghese, “A reliable and scalable striping protocol,” in *ACM SIGCOMM Computer Communication Review*, vol. 26, no. 4. ACM, 1996, pp. 131–141.
- [188] P. H. Fredette, “The past, present, and future of inverse multiplexing,” *Communications Magazine, IEEE*, vol. 32, no. 4, pp. 42–46, 1994.
- [189] C. Raiciu, M. Handley, and D. Wischik, “Coupled congestion control for multipath transport protocols,” *draft-ietf-mptcp-congestion-01 (work in progress)*, 2011.
- [190] D. Wischik, C. Raiciu, and M. Handley, “Balancing resource pooling and equipoise in multipath transport,” <http://www.wischik.com/damon/Work/Research/mtcpa.pdf>, 2010.
- [191] C. E. Hopps and D. Thaler, “Multipath issues in unicast and multicast next-hop selection,” *IETF RFC 2991*, 2000.
- [192] K. Manayya, “Constrained shortest path first,” <http://tools.ietf.org/search/draft-manayya-constrained-shortest-path-first-00>, 2010.
- [193] M. Honda, Y. Nishida, L. Eggert, P. Sarolahti, and H. Tokuda, “Multipath congestion control for shared bottleneck,” in *Proc. PFLDNet workshop*, 2009.
- [194] Q. Peng, A. Walid, J. Hwang, and S. Low, “Multipath tcp: Analysis, design, and implementation,” *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.
- [195] S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoidance,” *Networking, IEEE/ACM Transactions on*, vol. 1, no. 4, pp. 397–413, 1993.
- [196] S. Floyd, “Tcp and explicit congestion notification,” *ACM SIGCOMM Computer Communication Review*, vol. 24, no. 5, pp. 8–23, 1994.
- [197] K. Nichols and V. Jacobson, “Controlling queue delay,” *Communications of the ACM*, vol. 55, no. 7, pp. 42–50, 2012.
- [198] C. Raiciu, D. Wischik, and M. Handley, “Practical congestion control for multipath transport protocols,” *University College of London Technical Report*, 2009.
- [199] D. Wischik, C. Raiciu, A. Greenhalgh, and M. Handley, “Design, implementation and evaluation of congestion control for multipath TCP,” in *Proceedings of the 8th USENIX conference on Networked systems design and implementation*. USENIX Association, 2011, pp. 8–8.
- [200] M. H. aBT Innovate, “The trilogy architecture for the future internet,” *Towards the Future Internet: A European Research Perspective*, p. 79, 2009.
- [201] S. Floyd and K. Fall, “Promoting the use of end-to-end congestion control in the internet,” *IEEE/ACM Transactions on Networking (TON)*, vol. 7, no. 4, pp. 458–472, 1999.
- [202] B. Briscoe, “Flow rate fairness: Dismantling a religion,” *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 2, pp. 63–74, 2007.
- [203] T. Moncaster, L. Krug, M. Menth, J. Araujo, and R. Woundy, “The need for congestion exposure in the internet,” Internet draft, IETF, Tech. Rep., 2009.
- [204] B. Briscoe, A. Jacquet, C. Di Cairano-Gilfedder, A. Salvatori, A. Soppera, and M. Koyabe, “Policing congestion response in an internet-network using re-feedback,” *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 4, pp. 277–288, 2005.
- [205] K. Winstein and H. Balakrishnan, “Tcp ex machina: computer-generated congestion control,” in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. ACM, 2013, pp. 123–134.

- [206] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4p: Provider portal for applications," in *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4. ACM, 2008, pp. 351–362.
- [207] D. R. Choffnes and F. E. Bustamante, "Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4, pp. 363–374, 2008.
- [208] J. Seedorf and E. Burger, "Application-layer traffic optimization (alto) problem statement," *IETF RFC 5693*, 2009.
- [209] B. McCormick, F. Kelly, P. Plante, P. Gunning, and P. Ashwood-Smith, "Real time alpha-fairness based traffic engineering," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 199–200. [Online]. Available: <http://doi.acm.org/10.1145/2620728.2620762>
- [210] I. F. Akyildiz, A. Lee, P. Wang, M. Luo, and W. Chou, "A roadmap for traffic engineering in sdn-openflow networks," *Computer Networks*, vol. 71, pp. 1–30, 2014.
- [211] A. R. Curtis, W. Kim, and P. Yalagandula, "Mahout: Low-overhead data-center traffic management using end-host-based elephant detection," in *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2011, pp. 1629–1637.
- [212] S. Sharma, D. Staessens, D. Colle, M. Pickavet, and P. Demeester, "Enabling fast failure recovery in openflow networks," in *Design of Reliable Communication Networks (DRCN), 2011 8th International Workshop on the*. IEEE, 2011, pp. 164–171.
- [213] P. Fonseca, R. Bennesby, E. Mota, and A. Passito, "A replication component for resilient openflow-based networking," in *Network Operations and Management Symposium (NOMS), 2012 IEEE*. IEEE, 2012, pp. 933–939.
- [214] M. Reitblatt, N. Foster, J. Rexford, and D. Walker, "Consistent updates for software-defined networks: Change you can believe in!" in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*. ACM, 2011, p. 7.
- [215] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker, "Abstractions for network update," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. ACM, 2012, pp. 323–334.
- [216] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven wan," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. ACM, 2013, pp. 15–26.
- [217] R. Wang, D. Butnariu, J. Rexford *et al.*, "Openflow-based server load balancing gone wild," 2011.
- [218] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "Flowvisor: A network virtualization layer," *OpenFlow Switch Consortium, Tech. Rep*, 2009.
- [219] S. R. Chowdhury, M. F. Bari, R. Ahmed, and R. Boutaba, "Payless: A low cost network monitoring framework for software defined networks," in *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2014.
- [220] N. Feamster, J. Borkenhagen, and J. Rexford, "Guidelines for interdomain traffic engineering," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 5, pp. 19–30, 2003.
- [221] N. Vallina-Rodriguez and J. Crowcroft, "Energy management techniques in modern mobile handsets," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 1, pp. 179–198, 2013.
- [222] J. Gettys and K. Nichols, "Bufferbloat: Dark buffers in the internet," *Queue*, vol. 9, no. 11, p. 40, 2011.
- [223] H. Jiang, Z. Liu, Y. Wang, K. Lee, and I. Rhee, "Understanding bufferbloat in cellular networks," in *Proceedings of the 2012 ACM SIGCOMM workshop on Cellular networks: operations, challenges, and future design*. ACM, 2012, pp. 1–6.
- [224] Y.-C. Chen and D. Towsley, "On bufferbloat and delay analysis of multipath tcp in wireless networks," *IFIP Networking 2014*, 2014.
- [225] I. Castineyra and M. Steenstrup, "The nimrod routing architecture," in *IETF RFC*. Citeseer, 1992.
- [226] M. Honda, Y. Nishida, C. Raiciu, A. Greenhalgh, M. Handley, and H. Tokuda, "Is it still possible to extend TCP?" in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 181–194.
- [227] C. Raiciu, C. Paasch, S. Barre, A. Ford, M. Honda, F. Duchene, O. Bonaventure, M. Handley *et al.*, "How hard can it be? designing and implementing a deployable multipath TCP," in *NSDI*, vol. 12, 2012, pp. 29–29.
- [228] G. Detal, C. Paasch, and O. Bonaventure, "Multipath in the middle (box)," in *Proceedings of the 2013 workshop on Hot topics in middle-boxes and network function virtualization*. ACM, 2013, pp. 1–6.
- [229] A. K. Mir, A. Akram, E. Ahmed, J. Qadir, and A. Baig, "Unified channel assignment for unicast and broadcast traffic in cognitive radio networks," in *LCN Workshops*, 2012, pp. 799–806.
- [230] M. J. Farooq, M. Hussain, J. Qadir, and A. Baig, "A game-theoretic spectrum allocation framework for mixed unicast and broadcast traffic profile in cognitive radio networks," in *Local Computer Networks (LCN), 2013 IEEE 38th Conference on*. IEEE, 2013, pp. 425–432.
- [231] M. Zeeshan, M. F. Manzoor, and J. Qadir, "Backup channel and cooperative channel switching on-demand routing protocol for multi-hop cognitive radio ad hoc networks (bcccs)," in *Emerging Technologies (ICET), 2010 6th International Conference on*. IEEE, 2010, pp. 394–399.
- [232] M. T. Hassan, E. Ahmed, J. Qadir, and A. Baig, "Quantifying the multiple cognitive radio interfaces advantage," in *Advanced Information Networking and Applications Workshops (WAINA), 2013 27th International Conference on*. IEEE, 2013, pp. 511–516.
- [233] T. Roughgarden and É. Tardos, "How bad is selfish routing?" *Journal of the ACM (JACM)*, vol. 49, no. 2, pp. 236–259, 2002.
- [234] D. Acemoglu, R. Johari, and A. Ozdaglar, "Partially optimal routing," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 6, pp. 1148–1160, 2007.
- [235] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: defining tomorrow's internet," in *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4. ACM, 2002, pp. 347–356.
- [236] B. Editor: Philip Eardley, "Trilogy: Re-architecting the internet. the impact of trilogy. ict-216372," http://www.trilogy-project.org/fileadmin/publications/Deliverables/D17_D18_-_The_Impact_of_Triology.pdf, 2011.
- [237] A. Ford, P. Eardley, and B. van Schewick, "New design principles for the internet," in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*. IEEE, 2009, pp. 1–5.
- [238] I. Psaras, L. Saino, and G. Pavlou, "Revisiting resource pooling: The case for in-network resource sharing," in *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*. ACM, 2014, p. 24.
- [239] M. Caesar, M. Casado, T. Koponen, J. Rexford, and S. Shenker, "Dynamic route recomputation considered harmful," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 2, pp. 66–71, 2010.
- [240] S. E. Page, *Diversity and complexity*. Princeton University Press, 2010.
- [241] S. Dolev, Y. Elovici, and R. Puzis, "Routing betweenness centrality," *Journal of the ACM (JACM)*, vol. 57, no. 4, p. 25, 2010.
- [242] V. Chellappan and K. M. Sivalingam, "Application of entropy of centrality measures to routing in tactical wireless networks," in *Local & Metropolitan Area Networks (LANMAN), 2013 19th IEEE Workshop on*. IEEE, 2013, pp. 1–6.
- [243] A. Lozano and N. Jindal, "Transmit diversity vs. spatial multiplexing in modern MIMO systems," *Wireless Communications, IEEE Transactions on*, vol. 9, no. 1, pp. 186–197, 2010.
- [244] C. Pluntke, L. Eggert, and N. Kiukkonen, "Saving mobile device energy with multipath tcp," in *Proceedings of the sixth international workshop on MobiArch*. ACM, 2011, pp. 1–6.
- [245] C. Raiciu, D. Niculescu, M. Bagnulo, and M. J. Handley, "Opportunistic mobility with multipath TCP," in *Proceedings of the sixth international workshop on MobiArch*. ACM, 2011, pp. 7–12.
- [246] C. Paasch, G. Detal, F. Duchene, C. Raiciu, and O. Bonaventure, "Exploring mobile/WiFi handover with multipath TCP," in *Proceedings of the 2012 ACM SIGCOMM workshop on Cellular networks: operations, challenges, and future design*. ACM, 2012, pp. 31–36.
- [247] Y.-S. Lim, Y.-C. Chen, E. M. Nahum, D. Towsley, and R. J. Gibbens, "How green is multipath tcp for mobile devices?" in *Proceedings of the 4th workshop on All things cellular: operations, applications, & challenges*. ACM, 2014, pp. 3–8.