

(Draft – not for distribution)

On the Expected Payment of Mechanisms for Task Allocation

Artur Czumaj* Amir Ronen†

Abstract

We study a representative task allocation problem called shortest paths: Let G be a graph in which the edges are owned by self interested agents. The cost of each edge is privately known to its owner. Let s and t be two distinguished nodes in G . Given a distribution on the edge costs, the goal is to design a mechanism (protocol) which acquires a cheap $s - t$ path.

We first observe that the class of generalized VCG mechanisms has desired monotonicity properties. We exploit this observation to obtain, under an independence assumption, expected payments which are significantly lower than the worst case bounds of [2, 4]. We then investigate whether these payments can be improved when there is competition among paths. Surprisingly, we give evidence to the fact that typically such competition hardly helps incentive compatible mechanisms. In particular, we show this for the celebrated VCG mechanism. We then construct a novel general protocol combining the advantages of incentive compatible and non-incentive compatible mechanisms. Under reasonable assumptions on the agents we show that the overpayment of our mechanism is very small.

1 Introduction

Situations in which a set of tasks is to be allocated among self interested parties appear in many economic domains. Among the examples are distributed supply chains, privatization of public sectors, communication and transportation networks, multi agent systems, and even the management of software components priced by their usage. Problems which stem from such domains give rise to many challenging questions on the border of mechanism design and computer science.

We study a representative task allocation problem called *shortest paths*: We are given a directed graph G and two special nodes s and t in it. Each edge e in the graph is owned by a self interested agent $i(e)$. The actual cost c_e of routing an object along an edge is *privately* known to its owner. The goal of the center is to acquire an $s - t$ path and pay as little money as possible. This goal is not shared by the participating agents. Each agent i selfishly wants to maximize her *own* profit which is the difference between her payment and her actual costs. That is, $p_i - \sum_{e \in P \cap i(e)} c_e$ where P denotes the chosen path. We assume that no agent owns a cut in the graph.

This problem has the advantage that on the one hand it has a simple combinatorial structure, but on the other hand, essential ingredients of task allocation are intimately related to properties of the underlying graph. For example, the competition in the “market” is related to the connectivity

*Department of Computer Science, New Jersey Institute of Technology, Newark, NJ 07102, USA. Research supported in part by NSF grant CCR-0105701. E-mail: czumaj@cis.njit.edu.

†Faculty of Industrial Engineering & Management, Technion, Israel. The author was supported in part by grant number 53/03 – 10.5 from the Israel Science Foundation. Email: amirr@ie.technion.ac.il

of the graph and dependencies among tasks or agents are related to the graph diameter. The problem has received a lot of attention in recent years [2, 4, 5, 7, 10, 11].

Following the standard approach of mechanism design, we focus on mechanisms (protocols) which adhere to three desired game theoretic properties: Incentive compatibility, individual rationality, and no free riding. (The formal definitions of these properties are given in Section 2.3.) We call such mechanisms *valid*.

The starting point of this paper are two recent pessimistic worst case results. [2], which pioneered the study of shortest path payments, shows that for every protocol which belongs to a large class of valid mechanisms, there exist cost vectors which force the mechanism to pay $\Omega(k)$ times the cost of the second cheapest path (where k denotes the number of edges on the winning path). This payment is unacceptably high. [4] generalizes this claim for every valid mechanism.

In light of these results we are interested in the **expected** payment of mechanisms for shortest paths. Given a distribution ϕ on the edge costs, our goal is to find a valid mechanism which minimizes the total payment¹. We focus on the class of generalized VCG mechanisms (GVCG). To date, GVCG is the only general known method for the construction of valid protocols. In particular, we are interested in the relation between the *structure* of the graph and the optimal payment.

We make two major additional assumptions on the problem. Firstly, we assume that the edge costs are independent and bounded, w.l.o.g in $[0, 1]$. Secondly, as in [2, 4], we assume that each agent owns only one edge. This assumption is mainly required in order to make the presentation of the theorems and the analysis more tractable. Our theorems essentially have analogs, at least in the case where agents do not own large portions of the graph.

Our main results are as follows. We first observe that the family of GVCG mechanisms adheres to desired monotonicity properties. Specifically, if a graph G contains a sub graph G' then the optimal GVCG mechanism for G has a (weakly) lower payment than the optimal GVCG mechanism for G' .

We then conclude that under mild conditions, it is possible to construct, for every graph G , a protocol which has an expected payment of $O(k \cdot \sqrt{k})$, where k denotes the diameter of the graph. This is significantly better than the $\Omega(k^2)$ worst case bound of [2, 4].

The number of disjoint $s-t$ paths of a graph measures, in some sense, the amount of competition in the "market". Naturally, we expect the optimal payment to decrease when the competition increases. Surprisingly, we give evidence that this is not so true. We study graphs with n disjoint paths, each of length k . For several natural cost distributions we show that the VCG payment of these graphs decreases extremely slowly as n grows. Specifically VCG payments are of order $\Omega(k \cdot \sqrt{k} / \sqrt{\log n})$. The analysis gives good reasons to conjecture that this also bounds the payment of every GVCG mechanism.

In general, the design of task allocation mechanisms, not necessarily valid, exhibit a tradeoff between the overall payment and the risk that no path will be acquired at all. For situations in which competition does not significantly reduce the payments of valid mechanisms, we suggest the following **combined** protocol: Choose a small set of disjoint paths S' and exclude it from the graph; execute a "risky" protocol on the remaining subgraph; if this protocol fails, execute a GVCG protocol on S' . The properties of combined mechanisms are discussed in Section 7. Under reasonable assumptions on the agents, there exist such mechanisms with payment of order $O(k)$. We believe that this method could be successfully applied to many task allocation problems.

Finally, in the appendix, we comment on the conceptual difficulty of analyzing protocols which are not incentive compatible.

¹Desirably, the construction uses as less knowledge of the distribution as possible.

Three recent optimistic results about shortest path payments are partially related to our work. [4] constructs an optimal Bayesian mechanism for shortest paths and shows that for many natural distributions, its payment can be significantly lower than the payment of the optimal valid mechanism. However, the Bayesian solution concept is arguable for complicated problems like shortest paths, and the Bayesian incentive compatibility of the proposed mechanism is dependent on the independence of costs and the assumption that each agent owns a single edge. [10] shows that, under fixed costs, and the assumption that each agent owns only one edge, the expected VCG payment of a random graph, is very small. This is shown, both for $G_{n,p}$ and for power-law graphs. It may be possible to unite these results with ours. Finally, [5] measures the average VCG payment of the Internet's Autonomous Systems graph. The overpayment measured there is relatively small.

2 The model

2.1 Setup

We study the following representative shortest path problem. We are given a directed graph G and two special nodes s and t in it. Each edge e in the graph is owned by a self interested agent $i(e)$. The actual cost c_e of routing an object along an edge is *privately* known to its owner. The goal of the center is to acquire an $s - t$ path and pay as little money as possible. This goal is not shared by the participating agents. Each agent i selfishly wants to maximize her *own* profit which is the difference between her payment p_i and her actual costs. That is, $p_i - \sum_{e \in P \cap i(e)} c_e$ where P denotes the chosen path. We assume that no agent owns a cut in the graph.

We denote the *type* of each agent i , by t_i . This is the vector of the actual costs of her edges.

Given a distribution ϕ on the costs of the edges, our goal is to construct a *mechanism* (protocol) which acquires an $s - t$ path, and pays, in expectation, as little money as possible. Obviously, we would like to use as less knowledge of the distribution as possible. Most of the protocols we consider use only the expected path costs.

We make the following additional assumptions on the problem.

Independent bounded costs We assume that the actual cost of each edge is independent of the others'. We assume that the costs are bounded, w.l.o.g. in $[0, 1]$. While independence is a strong assumption, it is common in economics and seems necessary to make the analysis more tractable. The assumption that the costs are bounded is natural to most applications.

Each agent owns one edge This assumption is needed only in order to make the presentation of the theorems and the analysis more tractable. In principle, our theorems have analogs, at least in the case where agents do not own large portions of the graph.

Finally, as in [2, 4], we implicitly assume that the center must acquire one of the paths. Our theorems have analogs when this is not the case.

2.2 Graph notation

Given a graph G we denote its $s - t$ path graph by G_P . We label each edge in the graph by its owner. We also allow dummy edges which we denote by a dummy agent. These agents do not get paid by the protocol. We say that $G \subseteq G'$ if G is a subset of G' . Intuitively, one would expect G' to have better payment properties than G . We are particularly interested in graphs which are composed of disjointed $s - t$ paths. We call such graphs *disjointed path graphs*. A disjointed path

graph composed of n disjointed paths, each of length k , is denoted by $G(n, k)$. Given integers $k_1 \leq k_2 \dots \leq k_n$, we let $G(k_1, k_2, \dots, k_n)$ denote the disjointed path graphs containing n paths such that each path i is of length k_i . For such graphs, the number of paths n corresponds to the amount of competition in the “marke”, and the k_i s correspond to the dependencies (complementarities) between the agents. Let $G(k_1, k_2, \dots, k_n)$ and $G(k'_1, k_2, \dots, k'_n)$ be two disjointed path graphs. We say that $G = G(k_1, k_2, \dots, k_n)$ is *better than* $G' = G(k'_1, k_2, \dots, k'_n)$, if $n' \leq n$ and $k_i \leq k'_i$ for each $i \leq k$. We denote this relation by $G \preceq G'$. In other words G contains more competition and less dependencies than G' . Intuitively, one would expect G to have better payment properties than G' . In the following sections we will see that this is not so obvious. If G' is a disjointed path graph and G a graph, we also say that G is better than G' , if G contains a disjointed path subgraph G'' which is better than G' . An example of a disjointed path graph can be found in Figure 1.

2.3 Mechanisms

The overall payment of a protocol for shortest paths depends not only on the actual costs (agent types) but mostly on the *behavior* of the agents during the execution of the protocol. For arbitrary protocols, this behavior is highly unpredictable. Thus, we take the usual approach of mechanism design theory and limit ourselves to protocols which adhere to standard desired game theoretic properties. A protocol is called *incentive compatible* or IC for short, if each agents has a dominant strategy. Such a protocol is called *individually rational* or IR for short, if agents who behave according to their dominant strategies always have non negative profits. Finally, we say that a protocol adheres to *no free riding* if only agents on the chosen path get paid. Most of this paper focuses on protocols which adhere to these three properties (*valid mechanisms*). In Section 7 we slightly relax the IC property.

A well known observation called the *revelation principle* shows that every valid mechanism has an equivalent valid mechanism in which the agents are simply required to declare their types. Of course, they may lie, but since truth-telling is dominant, rational agents will refrain from doing so. Based on the vector of declarations $d = (d_1, \dots, d_n)$, the mechanism computes the winning path $k(d)$, and the payments for each of the agent $p(d) = (p_1(d), \dots, p_n(d))$. Such a mechanism (revelation mechanism) is completely described by a pair of functions $(k(\cdot), p(\cdot))$.

Given a graph G and a revelation mechanism m we define the *payment* of m on G as $P_m(G) = E_{d \in \phi}[\sum_i p_i(d)]$. That is, the expected total payment when the types of the agents are drawn from the underlying distribution ϕ , and the agents behave truthfully. We denote by P_{opt} the expected payment of the optimal valid protocol.

Given a graph G with costs, an obvious lower bound on the payment of every valid protocol is the cost of the shortest path denoted by c_1 . We define the *overpayment* of a mechanism m as the ratio between its payment P_m and the expected cost of the shortest path $E[c_1]^2$.

2.4 Generalized VCG mechanisms

Perhaps the most celebrated positive result of mechanism design theory is the VCG method for the construction of IC mechanisms[14, 3, 6]. This method is very general and can be applied to many task allocation problems of interest. To date, VCG and its generalizations defined in this section, are the only general known methods for the construction of IC mechanisms. For some problems, they are provably the sole available methods (see [13, 9]).

²This definition is slightly different from the definition of Frugality in [2, 4]

We hereby define the VCG mechanism for shortest paths. Our definition corresponds to the valid variant of VCG which is due to Clarke [3]. Given a path selection algorithm $k(d)$, we define $k_{-i}(d)$ to be the result of the algorithm when the edges of agent i are omitted from the graph. Given a path P , we let $c_i(P)$ denote the cost of agent i , that is the sum of the actual costs of all her edges which belong to the path P . We now define the mechanism and describe some known properties of it.

Definition 2.1 (VCG mechanism) *A VCG mechanism for shortest paths is a revelation mechanism (k, p) such that:*

- *The allocation $k(d)$ is optimal*
- *The payment to each agent i is the difference between the optimal total cost when i 's edges are omitted from the graph, and the costs of the other agents on the chosen path. I.e. $p_i(d) = \sum_{j \neq i} c_j(k_{-i}(d)) - \sum_{j \neq i} c_j(k(d))$*

The VCG payment to each agent i is equal to the sum of her own expenses $c_i(k(d))$ and her “contribution” – the difference between the optimal cost with and without her edges. The total payment of VCG in an (n, k) -path graph has a particularly elegant form. Let c_1 and c_2 denote the lowest and the second lowest path costs. Let $\Delta = c_2 - c_1$. The total payment of VCG equals $c_1 + k \cdot \Delta$. As observed in [2], in many cases of interest, the bonus component $k \cdot \Delta$ dominates the overall payment. Applying VCG to a graph G and to its path graph G_P yields the same results. Note that when the agents rationally choose their dominant strategies and declare their true costs, VCG finds the optimal path. In fact VCG is the only valid mechanism with this desirable property.

It is known that VCG mechanisms can be generalized in several ways. Firstly, the path selection algorithm could be limited to minimize only on a subset of the possible paths. Secondly, it is possible to add, to the cost of each path P , a constant $F(P)$, and then apply VCG to the modified path costs. We call mechanisms obtained this way *non weighted generalized VCG (NWGVCG)*. Formally:

Definition 2.2 (NWGVCG mechanism) *A mechanism for shortest paths is called NWGVCG if there exists a subset S' of the paths, and a function $F : S' \rightarrow R$ such that:*

- *The allocation $k(d)$ is of the form $k(d) \in \arg \min_{p \in S'} y(p)$*
- *The payment of each agent i is given by $p_i(d) = c_i(k(d)) + y(k_{-i}(d)) - y(k(d))$*

where $y(p) = F(p) + \sum_j c_j(p)$.

NWGVCG mechanisms are known to be valid. An example of such a mechanism is given in Section 3.

Another known generalization of VCG is to apply a weighted version of it at the last stage of NWGVCG. We call such a mechanism *Generalized VCG (GVCG)*.

Note that each of the three methods defined in this section, generalizes its predecessor. In principle, whenever VCG is valid, then so is every GVCG mechanism.

Finally, another operator, which can be applied to mechanisms for shortest paths, is to bound the declared cost of some edges. The center then disqualifies every path which contains edges with declarations above their bounds. In general such bounds, if not too generous, either require strong knowledge of the distribution, or put the center in the risk of not acquiring any path at all. We sometimes mention how such bounds can reduce the overall payment.

3 Example: Augmenting a bad graph

Consider the $G = G(k, 2k)$ graph of Figure 1. Assume that G 's edge costs are independent and uniform in $[0, 1]$. The path costs of G are sharply concentrated around $k/2$ and k respectively. Thus, w.h.p $\Delta = \Theta(k)$. Consider a VCG protocol for shortest paths. The protocol pays w.h.p. a “bonus” of Δ to each agent on the shortest path. Therefore, although both path costs are of order $O(k)$, the overall payment of VCG is $\Omega(k^2)$. This phenomenon was first observed in [2]. [4] shows that for every valid mechanism, there exist cost vectors which force it to pay $\Omega(k)$ times the cost of the second cheapest path.




Figure 1: A bad $G(k, 2k)$ graph

Luckily, typically, the protocol can be designed given the underlying graph G and the distribution ϕ . In this case, it is not difficult to see that the G can be **augmented** such that the expected payment is only of order $k \cdot \sqrt{k}$. Consider the same graph except that k dummy edges with uniform cost distribution are added to the shorter path. Suppose that the designer conducts VCG on the augmented graph and then pays only to the non dummy edges. Clearly, from the agents' perspective, there is no difference whether the other edges are real or not. Therefore, the mechanism remains valid. Fortunately, w.h.p, Δ of the modified graph is of order $O(\sqrt{k})$ reducing the overall payment of the mechanism to $O(k \cdot \sqrt{k})$. It is not difficult to see that the protocol belongs to the NWGVCG family.

It is worth noting that the payment of VCG (as well as of any GVCG protocol) is not monotone with respect to natural properties of the underlying graph. If we remove k edges from $G(2k, 2k)$, the expected payment gets significantly worse.

4 Monotonicity properties of GVCG

As we saw in the previous section, valid protocols suffer from non monotone payments. Here we show that if the designer is allowed to optimize over NWGVCG (or GVCG) then this non-monotonicity is prevented. The fruits of this observation will be reaped in Sections 5 and 7.

Given a graph G , we define two operators on its path graph:

Elimination Eliminate paths from the path graph

Addition Add dummy edges of positive costs to paths of the path graph. (These edges are not paid by the mechanism.)

The following two propositions characterize the family of NWGVCG mechanisms.

Proposition 4.1 (Soundness) *Given a graph G , applying a sequence of elimination and addition operators on the path graph and then conducting VCG, yields a NWGVCG mechanism.*

Proof: Immediate from the definitions. The value of $F(p)$ equals the sum of the costs of the dummy edges added to p . \square

It is not difficult to see that *any* NWGVCG mechanism could be viewed as follows: The “forbidden” paths are eliminated from the path graph; dummy edges with a total weight of $F(p)$ are added to every allowed path P ; finally, a VCG mechanism is applied to the modified path graph. The dummy edges do not get paid by the mechanism.

Proposition 4.2 (Completeness) *Every NWGVCG mechanism can be obtained by applying a sequence of elimination and addition operators to the path graph and then conducting VCG.*

Proof: Clearly, adding the same positive constant to all paths changes neither the allocation nor the payment functions. Thus, the additive function $F(p)$ can be assumed non-negative. The proof then follows immediately from the definitions and the discussion above. \square

We sometimes add dummy edges with probabilistic costs. This does not change the basic game theoretic properties of the mechanism.

Consider two graphs G and G' such that $G' \subseteq G$. Propositions 4.1 and 4.2 imply that for any NWVCG mechanism m' applied to G' , there exists an equivalent NWVCG mechanism applied to G . We simply eliminate all the paths which are not in G'_p from G_p and then apply m' . Therefore:

Theorem 1 (Monotonicity of NWVCG) *Let $G' \subseteq G$ be two graphs. Let m and m' denote the optimal NWVCG mechanism for G and G' respectively. Then $P_m(G) \leq P_{m'}(G')$.*

Corollary 4.3 *Let G' be a disjointed path graph, and G a graph such that $G' \preceq G$. Then $P_{opt \in NWVCG}(G) \leq P_{opt \in NWVCG}(G')$*

By the same arguments, similar theorems hold for GVCG as well. The above theorems show that in order to bound the optimal payment on a given graph, one only needs to find a **sub-graph** of it with good payment properties!

5 Protocols with good payment properties

We now show that, under a mild condition, it is possible to construct for every graph G , a protocol with an expected payment significantly better than the worst case bounds of [2, 4].

Theorem 2 *Let G be a graph such that $G \preceq G(2, k)$, then there exists a NWVCG protocol m , such that $P_m(G) \leq O(k \cdot \sqrt{k})$.*

Proof: According to theorem 4.3 it is enough to prove this for $G' = G(2, k)$. By adding up to k edges of fixed costs less than or equal to 1 to each path we can equalize the expected costs of the two paths. Therefore by theorem 7 we have that $E[\Delta] = O(\sqrt{k})$. Thus, $P_{VCG}(G') \leq E[c_1] + (k \cdot \sqrt{k}) = O(k \cdot \sqrt{k})$. \square

It is worth noting that it is possible to augment the graph by adding just one edge to the path with the lower expected cost. Consider a (k_1, k_2) graph. If $k_2 - k_1$ is large, it is often better, instead of equalizing the path costs, to leave a small advantage to the path with the fewer edges p_1 . In this case, w.h.p. p_1 wins, so only k_1 agents get the bonus Δ . Therefore, if the graph contains a path of length k and another path which is not extremely long, it is possible to obtain a payment of order $O(k \cdot \sqrt{k})$. We omit the proof due to space constraints.

We comment that in graphs in which all (low) path costs are proportional to their lengths, the above mechanism has an overpayment of $O(\sqrt{k})$. This is in contrast to the worst case $\Omega(k)$ bound of [2, 4]. On the other hand, graphs in which there are many low cost edges and only a few high

cost ones are bad. In these graphs, the light edges are "free riding" the heavy ones gaining the same bonus. In such cases, a combined protocol (see definition 7.1), will likely solve the problem. Another (risky) option is to bound the declared costs.

Theorem 7 immediately implies that path competition is at least not too harmful. That is:

Theorem 3 $P_{VCG}G(n, k) \leq O(k \cdot \sqrt{k \cdot \log n})$.

Theorem 7 uses the union bound and henceforth shows that w.h.p. the difference between every two paths is $O(\sqrt{k})$. Therefore, similar results hold at least when agents own small subsets of the graph.

6 Competition between paths

The results of section 5 show that it is possible to obtain an overpayment of $O(\sqrt{k})$. Unfortunately, for many practical economic applications, any overpayment greater than a small constant is not tolerable. Consider for example a government privatizing a public sector. An overpayment of, say 2, can be much worse than leaving the sector centrally controlled, and can lead to serious welfare loss for the public. A natural question is whether competition can help in reducing the payment. In this section we examine the value of competition between paths. Specifically, we study the overall payment in $G(n, k)$ -graphs. This payment is intimately related to the behavior of the difference Δ between the costs of the lowest and the second lowest paths, as a function of k and n . Surprisingly, we show, both theoretically, and experimentally, that for many natural independent identical cost distributions (iid costs), Δ declines in a rate proportional to $1/\sqrt{\ln n}$. This rate is extremely slow meaning that a huge number of edges are needed in order to significantly reduce the overall payment. This bound immediately implies a lower bound of $\Omega(k \cdot \sqrt{k}/\sqrt{\log n})$ on the payment of VCG mechanisms. The analysis reinforces the conjecture this bound applies to all GVCG mechanisms. Subsection 6.3 explains why we believe that for most cost distributions, the decline of Δ will be very slow.

6.1 Bernoulli IID Costs

Consider the case where each edge has an independent cost of 1 with probability p and 0 otherwise. Theorem 10 shows that asymptotically, Δ tends to $\Omega(k \cdot \sqrt{k}/\sqrt{\log n})$. Therefore:

Theorem 4 $P_{VCG}G(n, k) = \Omega(k \cdot \sqrt{k}/\sqrt{\log n})$.

Note that Δ remains the same if the distribution is shifted by a constant. We strongly believe that the same bound holds for *every* NWVCG mechanism. Note that there are distributions in which Δ increases with m (at least up to some value). For such distributions, assigning high identical weight to some paths and zero weight to the others has the helpful effect of reducing the amount of parallelism. Thus, the analysis should take into account the structure of specific distribution. We also conjecture that the bound holds for *every* GVCG mechanism. Formally:

Conjecture 5 *Consider graphs with Bernoulli IID edges costs. Then, there exist a constant $c > 0$ s.t. for every GVCG mechanism m , and for n and k large enough, $P_{VCG}G(n, k) \geq c \cdot k \cdot \sqrt{k}/\sqrt{\log n}$.*

This conjecture might be difficult to prove, in particular for weighted VCG mechanisms.

6.2 Experimental results

The results of the previous sections are asymptotic. This section experimentally studies the behavior of the total payment of VCG with a reasonable value of $k = 10$. Figure 2 shows the overall payment of several distributions with n ranging from 2 to 16. (Note that the horizontal axis is $\log n$.) The edge costs are all distributed between 1 and 5. We have considered 3 distributions: Uniform, Bernoulli with $p = 1/3$ and Bernoulli with $p = 2/3$. The decline of the overall payment is indeed very slow. For concreteness, in the uniform distribution, the difference between a graph with 20 edges and one with 160 is less than 30%! All graphs exhibit an overpayment of at least 1.5. We have conducted many experiments with different parameters and observed similar results.




Figure 2: Expected payment of $G(n, 10)$ graphs

6.3 General Independent Distributions

The central limit theorem implies that under mild conditions, the cost of each path can be approximated by a normal distribution. The theory of extreme values in Statistics (see e.g. [12]), shows that for most IID path cost distributions, Δ will indeed decrease very slowly. In particular for identical gaussian path costs, the decline rate is asymptotically $\Theta(1/\sqrt{\log n})$. (Note that this means that theorem 4 holds for IID gaussian edge costs.) At least intuitively, for non IID path costs, the distributions with the high variance should dominate the low variance ones. Therefore there is a good reason to believe that Δ declines very slowly for many distributions of interest. Unfortunately, general theorems, in particular about non identical distributions, seem hard to prove and presumably require novel tools for the analysis of extreme values.

We comment that it is possible to give examples of IID distributions in which Δ decreases fast, but also examples in which it increases, at least up to some level of n .

7 Combined protocols

In general, the design of task allocation mechanisms, exhibits a tradeoff between the total payment and the chances that no path will be acquired at all. As we saw, there are natural situations in which competition is not helpful for valid protocols. For such situations it is possible to **combine** a “risky” protocol with one which always acquires a path. There are many possible ways of combining such protocols. This section describes one such method.

We firstly define a non valid protocol which we call *descending mechanism with reserved price* R (or R -descending for short). The protocol is as follows. Each agent submits an initial proposal

for the price she wants if her edge is acquired. At each stage of the protocol there is a provisional winning path and the agents, in parallel, are given a chance to reduce their proposals. The mechanism halts when no agent does so. If the asked cost of the provisional winning path is greater than R , no path is acquired.

This protocol (see Appendix A) has many ex post equilibria. In some of them, no path is acquired, even if there are paths which cost less than R . There is no normative way to quantify which equilibrium will be selected without additional assumptions on the behavior of the agents. We say that an equilibrium is *myopic* if whenever the actual cost of a non winning path is less than the provisional shortest cost, at least one of the agents will lower her proposal. Myopic behavior seems reasonable, in particular when there are many disjoint paths. Clearly, if there are paths with actual cost below R then there exist myopic equilibria. Also, in such equilibria, the payment never exceeds the cost of the cheapest path which is disjoint to the winner. Similarly to Section 4, elimination and addition operations can be applied to the path graph before the protocol is executed. These operations may sometimes significantly reduce the expected value of Δ .

Now, consider the following combined mechanism.

Definition 7.1 (combined NWVCG R -descending mechanism) *Given a reserved price R , a combined NWVCG R -descending mechanism is a mechanism which is obtained as follows.*

- Choose a small set S' of disjoint paths
- Execute the R -descending protocol on the rest of the graph (possibly apply elimination and addition operations beforehand)
- If the above protocol fails to acquire any path, apply an optimal NWVCG mechanism to S'

Theorem 6 *For $n > 2$, let $G \preceq G(n, k)$. Then, there exists a combined NWVCG R -descending mechanism m which has the following properties:*

- The mechanism always acquires a path
- For every ex post equilibrium (in which agents who have dominant strategies choose them), the expected payment is at most $O(k \cdot \sqrt{k})$
- There exist myopic equilibria, and for every myopic equilibrium, the total payment is never higher than k .

Proof: Note that whenever the descending protocol succeeds, the payment is at most R . Therefore defining R to be k guarantee the last part of the theorem. According to theorem 2 excluding two disjoint paths ensures an the expected payment of $O(k \cdot \sqrt{k})$ in case the descending protocol fails. \square

Many variants of the above protocol are possible. In particular, the reserved price may be less than k . Note that at least intuitively, the chance that the protocol will terminate in a “cheap” equilibrium increases with n . This is because it suffices that only a fraction of the agents will behave myopically. Note also that the design of the protocol hardly requires knowledge of the underlying distribution. We believe that such combined mechanisms can be successfully applied to a wide range of task allocation problems.

Acknowledgments: We thank Moshe Tennenholtz for many inspiring discussions, Ishay Weissman for his help with extreme value theory, and Inbal Ronen for commenting on earlier drafts of this paper.

References

- [1] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. In *Proceedings of The Thirty-Fifth Annual ACM Symposium on Theory of Computing (STOC 2003)*. ACM, 2003.
- [2] Aaron Archer and Éva Tardos. Frugal path mechanisms. *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 991–999, 2002.
- [3] E. Clarke. Multipart pricing of public goods. *Public Choice*, 18:19–33, 1971.
- [4] Edith Elkind, Amit Sahai, and Ken Steiglitz. Frugality in path auctions. To appear, 2003.
- [5] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, and Scott Shenker. A bgp-based mechanism for lowest-cost routing. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, pages 173–182. ACM Press, 2002.
- [6] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [7] John Hershberger and Subhash Suri. Vickrey prices and shortest paths: What is an edge worth? In *Proceedings of The 42rd Annual IEEE Symposium on the Foundations of Computer Science (FOCS 2001)*, pages 252 – 259. IEEE, 2001.
- [8] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, March 1963.
- [9] Ron Lavi, Noam Nisan, and Ahuva Mualem. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2003)*, 2003.
- [10] Milena Mihail, Christos Papadimitriou, and Amin Saberi. On certain connectivity properties of the internet topology. In *Proceedings of The 44rd Annual IEEE Symposium on the Foundations of Computer Science (FOCS 2001)*. IEEE, 2003.
- [11] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166–196, 2001. Extended abstract appeared in proceedings of the Thirty First Annual ACM symposium on Theory of Computing (STOC99).
- [12] Sidney Resnick. *Extreme Values, Regular Variation and Point Processes*. Springer-Verlag, 1987.
- [13] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences*, pages 321–349. North-Holland, 1979. Papers presented at the first European Summer Workshop of the Econometric Society.
- [14] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

A On the difficulty of analyzing non-incentive compatible protocols

For valid protocols, the assumption that agents follow their dominant strategies is natural. Therefore, valid protocols could be measured and ranked according to their (truthful) expected payments. Having such a measure is essential for the design of mechanisms. In this section we demonstrate that it is not so clear what the right measure of non valid mechanisms is. (See also [1].)

Consider a simple $G(n, 2)$ graph ($n \geq 4$). Suppose that in two of the paths each edge has a high cost h , and the cost of all other edges is zero. Consider, e.g. the descending mechanism of Section 7. Clearly, there are ex post equilibria with zero payment. Even more so, as n grows, we expect, at least intuitively, the total payment to be very small. Unfortunately, there are also bad equilibria. Consider the case where the expensive agents ask for their actual price h and every cheap edge asks a high price of $3 \cdot h$. This is also an equilibrium since no cheap edge can individually lower its offer such that its path would win. Therefore, there is an arbitrary ratio between the payment of various equilibria. In complicated graphs, in particular when agents own sets of edges, it is not clear how to refine the various equilibria or predict the chances of each one to be selected. Therefore, finding a natural and tractable measure of non valid protocols seems challenging and important.

B Analysis of Δ of $G(n, k)$ Graphs

B.1 Preliminaries

In our analysis of the distribution of Δ in $G(n, k)$ graphs we shall frequently use three basic inequalities: $1 + x \leq e^x$, $1 + x \geq e^{x/(x+1)}$, and $1 + x > e^{x/2}$ (the last inequality holds for all $0 < x < 1$).

B.2 Basic estimation for the distribution of Δ — a general upper bound

In this subsection, we consider the case when there are n edge-disjoint paths, each of length k , the cost of each edge is an independent random variable taking the values in $[0, 1]$ and the expected total cost of every path is the same.

Let X_i be the random variables denoting the cost of the i th path, $1 \leq i \leq n$. Then, we want to consider the random variable $\Delta = \min_{1 \leq i \leq n} \{X_i\} - \min_{1 \leq i \leq n} \{X_i : \exists X_j : 1 \leq j \leq n, X_j \leq X_i \text{ \& } i \neq j\}$. Or in other words, if we denote by C_i to be the i th smallest cost among the paths, then $\Delta = C_2 - C_1$.

At the beginning, let us consider the simplest scenario when $n = 2$, in which case $\Delta = |X_1 - X_2|$. Then one can easily see that Δ is the absolute value of the sum of k independent distributed random variables Q_1, Q_2, \dots, Q_k such that for every i , $1 \leq i \leq k$, we have $Q_i \in [-1, 1]$ and $\mathbf{E}[\sum_{i=1}^k Q_i] = 0$. (Indeed, we set Q_i to be the difference between the cost of the i th edge in the first path and the cost of the i th edge in the second path.) Once we observe that, we can use Chernoff-Hoeffding bound to estimate the concentration of the value of Δ around its mean (which is 0). We use the classical Chernoff-Hoeffding bound in the following form:

Claim B.1 [8] *Let V_1, \dots, V_N be independent random variables taking values in the unit interval $[0, 1]$ and let $V = \sum_{i=1}^N V_i$. Then, for any $t > 0$,*

$$\Pr[|V - \mathbf{E}[V]| > t] < 2 \cdot e^{-2t^2/N} .$$

Let us consider the independent random variables V_1, \dots, V_k taking values in the unit interval $[0, 1]$ and such that $\mathbf{E}[V_i] = \frac{1}{2}$. (We can do the transformation by taking $V_i = \frac{1}{2} \cdot (Q_i + 1)$; then we will have $\Delta = |\sum_{i=1}^k Q_i| = |2 \cdot \sum_{i=1}^k V_i - k| = 2 \cdot |\sum_{i=1}^k V_i - k/2|$.) Then, clearly, $\mathbf{E}[\sum_{i=1}^k V_i] = k/2$, and hence we can use the claim above to obtain the following bounds:

$$\begin{aligned} \Pr[\Delta > t] &= \Pr\left[2 \cdot \left|\sum_{i=1}^k V_i - k/2\right| > t\right] = \Pr\left[\left|\sum_{i=1}^k V_i - k/2\right| > t/2\right] \\ &= \Pr\left[\left|\sum_{i=1}^k V_i - \mathbf{E}\left[\sum_{i=1}^k V_i\right]\right| > t/2\right] < 2 \cdot e^{-t^2/(2k)} . \end{aligned}$$

Next, we could extend the inequality above to deal with many paths. Indeed, the analysis above implies that for any pair $1 \leq i < j \leq n$, we have $\Pr[|X_i - X_j| > t] < 2 \cdot e^{-t^2/(2k)}$. Therefore, by the union bound we have

$$\Pr[\Delta > t] \leq \Pr[\forall 1 \leq i < j \leq n |X_i - X_j| > t] \leq \binom{n}{2} \cdot 2 \cdot e^{-t^2/(2k)} .$$

This immediately yields the following theorem.

Theorem 7 *For a set of n edge disjoint paths of length k each, where the cost of each edge is an independent random variable taking the values in $[0, 1]$ and the expected total costs of every paths are the same, it holds for any $t > 0$,*

$$\Pr[\Delta > t] < 2 \cdot \binom{n}{2} \cdot e^{-t^2/(2k)} .$$

In particular,

$$\mathbf{E}[\Delta] \leq O(\sqrt{k \ln n}) .$$

B.3 Basic claims about the distribution of X_i

In this subsection we present three basic properties of the distribution of the costs of the paths. We consider a scenario of independent paths of length k each, such that the cost of every edge has Bernoulli distribution with the probability of success $\frac{1}{2}$. Let X_i be random variable denoting the cost of the i th path. Clearly, since each edge on each path has Bernoulli distribution, X_i has independent binomial distributions with parameters k and $\frac{1}{2}$, which we denote by $\mathbb{B}(k, \frac{1}{2})$, that is, for every $t = 0, \dots, k$ and i , we have $\Pr[X_i = t] = \Pr[\mathbb{B}(k, \frac{1}{2}) = k] = \binom{k}{t} \cdot 2^{-k}$.

Our goal in this subsection is to investigate the distribution of X_i . We state the following three claims whose proofs are deferred to Subsection B.6.

Claim B.2 *For every $\lambda \leq k/2$, $\lambda = \Omega(\sqrt{k})$,*

$$\Pr[X_i = k/2 - \lambda] = \Theta\left(\frac{1}{\sqrt{k}}\right) \cdot e^{-\Theta(\lambda^2/k)} . \quad (1)$$

Claim B.3 *There is an integer $\lambda = \Theta(\sqrt{k})$ such that*

$$\Pr[X_i = k/2 - \lambda] \geq \frac{1}{e} \cdot \Pr[X_i = k/2] > \Pr[X_i = k/2 - (\lambda + 1)] . \quad (2)$$

Claim B.4 *If $s = \Omega(\sqrt{k})$ and $s \leq k/4$, then there is an integer $\lambda = \Theta(k/s)$ such that*

$$\Pr[X_i = k/2 - s - \lambda] \geq \frac{1}{e} \cdot \Pr[X_i = k/2 - s] > \Pr[X_i = k/2 - s - (\lambda + 1)] . \quad (3)$$

B.4 Analysis of binomial distribution for two paths of identical length

In this subsection we consider the basic case in which there are two edge disjoint paths consisting of k edges each. We use the model from the previous subsection and let X_1 and X_2 be random variables denoting the cost of the 1st and the 2nd path, respectively. With such notation we have $\Delta = |X_1 - X_2|$. Our goal is to investigate the distribution of Δ .

B.4.1 Expected value of Δ for two paths

We begin with the analysis of the expected value of Δ . In the previous subsection we showed that $\mathbf{E}[\Delta] = O(\sqrt{k})$, and in this subsection we show that that bound is asymptotically tight. For simplicity of arguments, let us assume k is even. It is known that $\Pr[X_1 = k/2] = \Pr[\mathbb{B}(k, \frac{1}{2}) = k/2] = \frac{1}{\Theta(\sqrt{k})}$. (This standard observation can be easily proven using, e.g., Stirling's formula, and is proven formally in the later parts of our analysis.) Furthermore, $\Pr[\mathbb{B}(k, \frac{1}{2}) = t]$ achieves the maximum values for $t = k/2$. Therefore, if $\Pr[X_1 = k/2] \leq c_1/\sqrt{k}$ for some positive constant c_1 , then we have $\Pr[X_1 = k/2 + t] \leq \Pr[X_1 = k/2] \leq c_1/\sqrt{k}$ for all non-negative t . Hence,

$$\Pr[X_1 < k/2 + \frac{\sqrt{k}}{4c_1}] \leq \frac{1}{2} + \sum_{t=0}^{\lceil k/(4c_1) - 1 \rceil} \Pr[X_1 = k/2 + t] \leq \frac{1}{2} + \frac{\sqrt{k}}{4c_1} \cdot c_1/\sqrt{k} = \frac{3}{4} .$$

Therefore, for some constant c , we must have $\Pr[X_1 \geq k/2 + c\sqrt{k}] \geq \frac{1}{4}$.

It is easy to see that $\Pr[X_2 \leq k/2] \geq \frac{1}{2}$. Therefore, using the constant c above,

$$\Pr[\Delta \geq c\sqrt{k}] \geq \Pr[X_1 \geq k/2 + c\sqrt{k}] \cdot \Pr[X_2 \leq k/2] \geq \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8} .$$

Notice that this implies immediately that

$$\mathbf{E}[\Delta] = \Omega(\sqrt{k}) ,$$

and thus, by the analysis from the previous subsection we have

$$\mathbf{E}[\Delta] = \Theta(\sqrt{k}) .$$

We summarize this result in the following theorem.

Theorem 8 *For two edge disjoint paths of length k each, where the cost of every edge is an independent Bernoulli random variable with the success probability $\frac{1}{2}$, it holds that*

$$\mathbf{E}[\Delta] = \Theta(\sqrt{k}) . \quad (4)$$

Comment: why we cannot achieve good “high probability bounds” It is tempting to hope that the bound for the expected value of Δ holds with high probability; to the contrary, in this paragraph we show that we cannot have $\Delta = O(\sqrt{k})$ with high probability.

Notice that Claim B.2 implies that for $s = \Omega(\sqrt{k})$,

$$\begin{aligned} \Pr[X_i \leq k/2 - s] &= \sum_{t=s}^{k/2} \Pr[X_i = k/2 - t] \geq \sum_{t=s}^{s+k/(2s)} \Pr[X_i = k/2 - t] \\ &\geq \frac{k}{2s} \cdot \Pr[X_i = k/2 - (s + k/(2s))] \approx \frac{k}{2s} \cdot \frac{1}{\sqrt{k}} \cdot e^{-\Theta((s+k/(2s))^2/k)} \\ &= \frac{\sqrt{k}}{s} \cdot e^{-\Theta(s^2/k)} . \end{aligned}$$

This implies that for all $\lambda = \Omega(1)$,

$$\Pr[X_i \leq k/2 - \sqrt{\lambda \cdot k}] \approx \frac{1}{\sqrt{\lambda}} \cdot e^{-\Theta(\lambda)} = e^{-\Theta(\lambda)} .$$

Finally, we can apply this bound to give some estimations for the probability distribution of Δ for two paths ($n = 2$). Indeed, since $\Pr[X_1 \geq k/2] \geq \frac{1}{2}$, the estimations above yield the following,

$$\Pr[\Delta \geq \sqrt{\lambda \cdot k}] = e^{-\Theta(\lambda)} .$$

This implies that in order to get the confidence in the value of Δ of at least $1 - p$, we can obtain only the following bound:

$$\Pr[\Delta \leq \sqrt{\ln \frac{1}{p} \cdot k/2}] \approx 1 - p .$$

Finally, let us notice that this bound is essentially identical to the upper bound obtained in Theorem 7; hence this bound is essentially tight.

B.5 Analysis for n edge-disjoint paths

In this subsection we extend the analysis from the previous subsection to an arbitrary number n of edge disjoint-paths. Our main technical tool is partitioning of the values of X_i s into blocks in which all values are essentially indistinguishable in the stochastic sense which later allows us to make a formal analysis of the expected values of C_1 and Δ .

B.5.1 Partitioning into blocks

Let us partition the set of integers s , $0 \leq s \leq k$, into blocks $\mathbb{B}_1, \mathbb{B}_2, \dots$ such that:

- integer s is contained in the r th block \mathbb{B}_r such that

$$e^{-r} \cdot \Pr[X_i = k/2] < \Pr[X_i = k/2 - s] \leq e^{-(r-1)} \cdot \Pr[X_i = k/2] .$$

From now on, let ℓ_r denote the smallest number in the r th block, that is, the smallest s such that

$$e^{-r} \cdot \Pr[X_i = k/2] < \Pr[X_i = k/2 - s] \leq e^{-(r-1)} \cdot \Pr[X_i = k/2] .$$

We shall also use $|\mathbb{B}_r|$ to denote the size of the r th block, \mathbb{B}_r .

Notice that the definition above and the fact that $\Pr[X_i = k/2] = \Theta(1/\sqrt{k})$ immediately imply the following claim.

Claim B.5 For any $s \in \mathbb{B}_r$, it holds that

$$\Pr[X_i = k/2 - s] = \frac{\Theta(1)}{\sqrt{k}} \cdot e^{-r} .$$

By Claims B.3 and B.4, we have the following lemmas:

Lemma B.6 The size of the block \mathbb{B}_1 is $|\mathbb{B}_1| = \ell_1 = \Theta(\sqrt{k})$.

For any $2 \leq r \leq c \cdot \sqrt{k}$, for certain constant c , the size of the r th block \mathbb{B}_r is equal to $|\mathbb{B}_r| = \ell_{r+1} - \ell_r = \Theta(k/\ell_r)$.

(The upper bound for $r \leq c \cdot \sqrt{k}$ in Lemma B.6 is caused by the fact that for certain large r , $\ell_r = \Omega(k)$, we will have $|\mathbb{B}_r| = 0$. Similar assumption should be stated explicitly in the further analysis, but for simplicity of presentation we will slightly abuse the notation and ignore it.)

The following nontrivial lemma is important for our further analysis.

Lemma B.7 If $s \in \mathbb{B}_r$, then $s = \Theta(\sqrt{k \cdot r})$. In particular, for any $r \geq 1$, $|\mathbb{B}_r| = \Theta(\sqrt{k/r})$.

Proof: Let us first fix r and let us find the smallest index $r + j$ for which $2 \cdot \ell_r \geq \ell_{r+j}$. Clearly, this is essentially equivalent to find j such that $\ell_{r+j} - \ell_r \approx \ell_r$. Next, by Lemma B.6, we have

$$\ell_{r+j} - \ell_r = \sum_{s=0}^{j-1} \ell_{r+s+1} - \ell_{r+s} = \Theta\left(\sum_{s=0}^{j-1} k/\ell_{r+s}\right) .$$

Next, since for any s , $0 \leq s \leq j-1$, we have $\ell_r \leq \ell_{r+s} \leq 2 \cdot \ell_r$, we clearly obtain

$$\ell_{r+j} - \ell_r = \Theta(k \cdot j/\ell_r) .$$

Therefore, since $\ell_{r+j} = \Theta(\ell_r)$, we obtain $j = \Theta(\ell_r^2/k)$. Next, we can show that $j \geq r$ what implies that $\ell_{r+j} = \Theta(\ell_r) = \Theta(\sqrt{k \cdot j}) = \Theta(\sqrt{k \cdot (r+j)})$. \square

Notice also the following simple claim.

Claim B.8 For any r , $1 \leq r \leq c \cdot \sqrt{k}$ for certain positive constant c , it holds that

$$\Pr[X_i = k/2 - s \text{ with } s \in \mathbb{B}_r] = \frac{\Theta(1)}{\sqrt{r}} \cdot e^{-r} .$$

Proof: Follows directly from Claim B.5 and Lemma B.6. \square

B.5.2 “Almost-geometric” random variables

Let us consider a sequence of identically distributed independent random variables Y_1, Y_2, \dots, Y_n such that for any i , $1 \leq i \leq n$, and for any $r \in \mathbb{N}_{\geq 1}$, we have

$$\Pr[Y_i = r] = \Theta\left(\frac{1}{\sqrt{r} \cdot e^r}\right) .$$

Notice that this immediately implies that for any r we have also

$$\begin{aligned} \Pr[Y_i \leq r] &= 1 - \Theta\left(\frac{1}{\sqrt{r} \cdot e^r}\right) , \\ \Pr[Y_i \geq r] &= \Theta\left(\frac{1}{\sqrt{r} \cdot e^r}\right) . \end{aligned}$$

Let $\mathbf{Y} = \max\{Y_1, \dots, Y_n\}$, that is, \mathbf{Y} is the random variable that takes the maximum value among Y_1, \dots, Y_n . Then, it is easy to see that for any $r \geq 1$, we have

$$\begin{aligned} \Pr[\mathbf{Y} = r] &= \sum_{i=1}^n \Pr[Y_i = r \text{ and } \forall_{j < i} Y_j < r \text{ and } \forall_{j > i} Y_j \leq r] \\ &= \sum_{i=1}^n \left(\Pr[Y_i = r] \cdot \prod_{j=1}^{i-1} \Pr[Y_j < r] \cdot \prod_{j=i+1}^n \Pr[Y_j \leq r] \right). \end{aligned}$$

Therefore, by our bounds for the distribution of Y_i above, we obtain

$$\Pr[\mathbf{Y} = r] = n \cdot \Theta\left(\frac{1}{\sqrt{r} \cdot e^r}\right) \cdot \left(1 - \Theta\left(\frac{1}{\sqrt{r} \cdot e^r}\right)\right)^{n-1} = \Theta\left(\frac{n}{\sqrt{r} \cdot e^r} \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right)\right)\right).$$

B.5.3 Expected value of C_1 for arbitrary n

Let $C_1 = \min\{X_1, \dots, X_n\}$. By the discussion in the two previous subsections, we have

$$\Pr[C_1 = k/2 - s \text{ for certain } s \in \mathbb{B}_r] = \sum_{s \in \mathbb{B}_r} \Pr[C_1 = k/2 - s] = \Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right) \cdot \exp\left(-\frac{\lambda_r \cdot n}{\sqrt{r} \cdot e^r}\right)$$

for certain $\lambda_r = \Theta(1)$ (from now on, we shall use λ_r in the bound above, where one should understand that λ_r depends (or might depend) on r but for any r it is lower and upper bounded by some absolute constants). It is easy to verify that this probability is maximized for $r = \ln n - \frac{1}{2} \ln \ln n + \Theta(1)$, in which case it is constant.

Notice, for example, that this implies that the expected number of the block in which C_1 is contained is $\Theta(\ln n)$.

Now, we proceed to the following lemma, (we notice that in our analysis we do not really need this lemma, but we feel it is important to fully understand the behavior of the system and the following analysis)

Lemma B.9 $\mathbf{E}[C_1] = k/2 - \Theta(\sqrt{k \cdot \ln n})$.

Proof: Let us first observe that $\Pr[C_1 \geq k/2]$ can be neglected in the analysis. Next, let us continue with the definition and some basic manipulations:

$$\begin{aligned} \mathbf{E}[C_1] &\approx \sum_{s=0}^{k/2} (k/2 - s) \cdot \Pr[C_1 = k/2 - s] = k/2 - \sum_{s=0}^{k/2} s \cdot \Pr[C_1 = k/2 - s] \\ &= k/2 - \sum_{s \in \mathbb{B}_1} s \cdot \Pr[C_1 = k/2 - s] - \sum_{r \geq 2} \sum_{s \in \mathbb{B}_1} s \cdot \Pr[C_1 = k/2 - s] \\ &= k/2 - \Theta(\sqrt{k}) \cdot e^{-\Theta(n)} - \sum_{r \geq 2} \Theta(\sqrt{k \cdot r}) \cdot \Theta\left(\frac{n}{\sqrt{r} \cdot e^r} \cdot \exp\left(-\frac{\lambda_r \cdot n}{\sqrt{r} \cdot e^r}\right)\right) \\ &= k/2 - \sqrt{k} \cdot e^{-\Theta(n)} - \Theta(\sqrt{k} \cdot n) \cdot \sum_{r \geq 2} e^{-(r + \lambda_r \cdot n \cdot e^{-r/\sqrt{r}})}. \end{aligned}$$

It is easy to see that we can neglect the term $\sqrt{k} \cdot e^{-\Theta(n)}$, and to give a good bound for $\mathbf{E}[C_1]$ it is now sufficient to take care of the sum $\sum_{r \geq 2} e^{-(r + \lambda_r \cdot n \cdot e^{-r}/\sqrt{r})}$. By differentiating, we can see that $r + \lambda_r \cdot n \cdot e^{-r}/\sqrt{r}$ is minimized when

$$\sqrt{r} \cdot e^r = \lambda_r \cdot n \cdot \left(1 + \frac{1}{2r}\right),$$

and hence, using identical arguments as those used before, $e^{-(r + \lambda_r \cdot n \cdot e^{-r}/\sqrt{r})}$ is maximized for $r = \ln n - \frac{1}{2} \ln \ln n + \Theta(1)$. Therefore,

$$\sum_{r \geq 2} e^{-(r + \lambda_r \cdot n \cdot e^{-r}/\sqrt{r})} = \Theta\left(e^{-(\ln n - \frac{1}{2} \ln \ln n + \Theta(1))}\right) = \Theta\left(\frac{\sqrt{\ln n}}{n}\right),$$

and thus

$$\begin{aligned} \mathbf{E}[C_1] &= k/2 - \sqrt{k} \cdot e^{-\Theta(n)} - \Theta(\sqrt{k} \cdot n) \cdot \sum_{r \geq 2} e^{-(r + \lambda_r \cdot n \cdot e^{-r}/\sqrt{r})} \\ &= k/2 - \sqrt{k} \cdot e^{-\Theta(n)} - \Theta\left(\sqrt{k} \cdot n \cdot \frac{\sqrt{\ln n}}{n}\right) \\ &= k/2 - \Theta(\sqrt{k \cdot \ln n}). \end{aligned}$$

□

B.5.4 Expected value of Δ for arbitrary n

Now we are ready to give a bound for the expected value of Δ for arbitrary n . Our analysis uses similar tools as those that we needed in the proof of Lemma B.9, and actually, intuitively, it almost immediately follows from Lemma B.9.

Let us use a convention and write $C_i \in \mathbb{B}_r$ if $C_i = k/2 - s$ for certain $s \in \mathbb{B}_r$. In the previous subsections we showed that

$$\Pr[C_1 \in \mathbb{B}_r] = \Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right) \cdot \exp\left(-\frac{\lambda_r \cdot n}{\sqrt{r} \cdot e^r}\right)$$

for certain $\lambda_r = \Theta(1)$.

Let us notice that for any $0 \leq t < s \leq k/2$, we have

$$\begin{aligned} \Pr[C_1 = k/2 - s \text{ and } C_2 = k/2 - t] &= \sum_{i \neq j} \left(\Pr[X_i = k/2 - s] \cdot \Pr[X_j = k/2 - t] \cdot \right. \\ &\quad \left. \prod_{\ell < j, \ell \neq i} \Pr[X_\ell > k/2 - t] \cdot \prod_{\ell > j, \ell \neq i} \Pr[X_\ell \geq k/2 - t] \right), \end{aligned}$$

and for any $0 \leq s \leq k/2$, we have

$$\begin{aligned} \Pr[C_1 = C_2 = k/2 - s] &= \sum_{i=1}^{n-1} \sum_{j=i+1}^n \left(\Pr[X_i = k/2 - s] \cdot \Pr[X_j = k/2 - s] \cdot \right. \\ &\quad \left. \prod_{\ell < j, \ell \neq i} \Pr[X_\ell > k/2 - s] \cdot \prod_{\ell > j} \Pr[X_\ell \geq k/2 - s] \right). \end{aligned}$$

Furthermore, we have for any $r > t$,

$$\begin{aligned}
\Pr[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_t] &= \sum_{i=1}^n \left(\sum_{j=1}^{i-1} \Pr[C_1 = X_i \in \mathbb{B}_r \ \& \ C_2 = X_j \in \mathbb{B}_t] \right. \\
&\quad \left. + \sum_{j=i+1}^n \Pr[C_1 = X_i \in \mathbb{B}_r \ \& \ C_2 = X_j \in \mathbb{B}_t] \right) \\
&= \sum_{i=2}^n \sum_{j=1}^{i-1} \Pr[X_i \in \mathbb{B}_r] \cdot \Pr[X_j \in \mathbb{B}_t] \cdot \prod_{\ell=1}^{j-1} \Pr[X_\ell < \mathbb{B}_t] \cdot \\
&\quad \cdot \prod_{\ell=j+1}^{i-1} \Pr[X_\ell \leq \mathbb{B}_t] \cdot \prod_{\ell=i+1}^n \Pr[X_\ell \leq \mathbb{B}_t] + \\
&\quad + \sum_{i=1}^{n-1} \sum_{j=i+1}^n \Pr[X_i \in \mathbb{B}_r] \cdot \Pr[X_j \in \mathbb{B}_t] \cdot \prod_{\ell=1}^{i-1} \Pr[X_\ell < \mathbb{B}_t] \cdot \\
&\quad \cdot \prod_{\ell=i+1}^{j-1} \Pr[X_\ell < \mathbb{B}_t] \cdot \prod_{\ell=j+1}^n \Pr[X_\ell \leq \mathbb{B}_t] \\
&= n \cdot (n-1) \cdot \Theta\left(\frac{1}{\sqrt{r} e^r}\right) \cdot \Theta\left(\frac{1}{\sqrt{t} e^t}\right) \cdot \left(1 - \Theta\left(\frac{1}{\sqrt{t} e^t}\right)\right)^{n-2} \\
&= \Theta\left(\frac{n^2}{\sqrt{rt} \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right).
\end{aligned}$$

Similarly, “for $t = s$ ” we have

$$\Pr[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_r] = \Theta\left(\frac{n^2}{r \cdot e^{2r}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right)\right),$$

and hence we obtain the following claim.

Claim B.10 *For every $0 \leq t \leq s \leq k/2$, we have*

$$\Pr[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_t] = \Theta\left(\frac{n^2}{\sqrt{rt} \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right). \quad (5)$$

Next, we proceed with the analysis of the expected value of Δ .

We begin with a lower bound. Let us set $r = \ln n - \frac{1}{2} \ln \ln n + \Theta(1)$. Then we will have (in the analysis, for simplicity of presentation only, we do not “consider” the case $C_2 > k/2$, which can be essentially neglected because it is very unlikely to occur),

$$\begin{aligned}
\mathbf{E}[\Delta] &\geq \sum_{s \in \mathbb{B}_r} \sum_{\ell \in \mathbb{B}_{r-2}} (s - \ell) \cdot \mathbf{Pr}[C_1 = s \ \& \ C_2 = \ell] \\
&\geq \sum_{s \in \mathbb{B}_r} \sum_{\ell \in \mathbb{B}_{r-2}} |\mathbb{B}_{r-1}| \cdot \mathbf{Pr}[C_1 = s \ \& \ C_2 = \ell] \\
&= |\mathbb{B}_{r-1}| \cdot \mathbf{Pr}[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_{r-2}] \\
&\geq \Theta(\sqrt{k/r}) \cdot \Theta\left(\frac{n^2}{r \cdot e^{2r}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right)\right) \\
&= \Theta(\sqrt{k/\ln n}) .
\end{aligned} \tag{6}$$

Next, we continue with an upper bound for the expected value of Δ . We have,

$$\begin{aligned}
\mathbf{E}[\Delta] &\leq \sum_r \sum_{s \in \mathbb{B}_r} \sum_{\ell \in \mathbb{B}_r, \ell \leq s} (s - \ell) \cdot \mathbf{Pr}[C_1 = s \ \& \ C_2 = \ell] + \\
&\quad + \sum_r \sum_{t < r} \sum_{s \in \mathbb{B}_r} \sum_{\ell \in \mathbb{B}_t} (s - \ell) \cdot \mathbf{Pr}[C_1 = s \ \& \ C_2 = \ell] \\
&\leq \sum_r \sum_{s \in \mathbb{B}_r} \sum_{\ell \in \mathbb{B}_r, \ell \leq s} |\mathbb{B}_r| \cdot \mathbf{Pr}[C_1 = s \ \& \ C_2 = \ell] + \\
&\quad + \sum_r \sum_{t < r} \sum_{s \in \mathbb{B}_r} \sum_{\ell \in \mathbb{B}_t} \mathbf{Pr}[C_1 = s \ \& \ C_2 = \ell] \cdot \sum_{q=t}^r |\mathbb{B}_q| \\
&= \sum_r |\mathbb{B}_r| \cdot \mathbf{Pr}[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_r] + \\
&\quad + \sum_r \sum_{t < r} \mathbf{Pr}[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_t] \cdot \sum_{q=t}^r |\mathbb{B}_q| .
\end{aligned}$$

Now, let us remind the identity (5) and that $|\mathbb{B}_r| = \Theta(\sqrt{k/r})$. Then, we have,

$$\begin{aligned}
\mathbf{E}[\Delta] &\leq \sum_r |\mathbb{B}_r| \cdot \mathbf{Pr}[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_r] + \\
&\quad + \sum_r \sum_{t < r} \mathbf{Pr}[C_1 \in \mathbb{B}_r \ \& \ C_2 \in \mathbb{B}_t] \cdot \sum_{q=t}^r |\mathbb{B}_q| \\
&= \sum_{r \geq 1} \Theta(\sqrt{k/r}) \cdot \Theta\left(\frac{n^2}{r \cdot e^{2r}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right)\right) \\
&\quad + \sum_{r \geq 2} \sum_{1 \leq t < r} \Theta\left(\frac{n^2}{\sqrt{r} t \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right) \cdot \sum_{q=t}^r \Theta(\sqrt{k/q}) .
\end{aligned}$$

Notice that similarly as in our analysis for $\mathbf{E}[C_1]$, the value of the sum

$$\sum_{r \geq 1} \Theta(\sqrt{k/r}) \cdot \Theta\left(\frac{n^2}{r \cdot e^{2r}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right)\right)$$

is asymptotically determined by the value of the term with $r = \ln n - \frac{1}{2} \ln \ln n + \Theta(1)$. Therefore, we obtain,

$$\sum_{r \geq 1} \Theta(\sqrt{k/r}) \cdot \Theta\left(\frac{n^2}{r \cdot e^{2r}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{r} \cdot e^r}\right)\right) = \Theta(\sqrt{k/\ln n}) .$$

To analyze the second sum, we begin with some basic transformations:

$$\begin{aligned} & \sum_{r \geq 2} \sum_{1 \leq t < r} \Theta\left(\frac{n^2}{\sqrt{rt} \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right) \cdot \sum_{q=t}^r \Theta(\sqrt{k/q}) = \\ &= \sum_{r \geq 2} \sum_{1 \leq t < r} \Theta\left(\frac{n^2}{\sqrt{rt} \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right) \cdot (r-t+1) \cdot \mathcal{O}(\sqrt{k/t}) = \\ &= \Theta(\sqrt{k}) \cdot \sum_{r \geq 2} \sum_{1 \leq t < r} (r-t+1) \cdot \Theta\left(\frac{n^2}{t \cdot \sqrt{r} \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right) . \end{aligned}$$

Similarly as above, one can show that the sums take the maximum for $r = \ln n - \frac{1}{2} \ln \ln n + \Theta(1)$ and $t = r - \Theta(1)$, and therefore

$$\begin{aligned} & \Theta(\sqrt{k}) \cdot \sum_{r \geq 2} \sum_{1 \leq t < r} (r-t+1) \cdot \Theta\left(\frac{n^2}{t \cdot \sqrt{r} \cdot e^{r+t}}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{t} \cdot e^t}\right)\right) = \\ &= \Theta(\sqrt{k}) \cdot \Theta(1) \cdot \Theta\left(\frac{n^2}{\ln n \cdot \sqrt{\ln n} \cdot (n/\sqrt{\ln n})^2}\right) \cdot \exp\left(-\Theta\left(\frac{n}{\sqrt{\ln n} \cdot (n/\sqrt{\ln n})}\right)\right) = \\ &= \Theta(\sqrt{k/\ln n}) . \end{aligned}$$

This allows us to conclude with the main theorem.

Theorem 9 *For a set of n edge disjoint paths of length k each, where the cost of every edge is an independent Bernoulli random variable with the success probability $\frac{1}{2}$, it holds that*

$$\mathbf{E}[\Delta] = \Theta(\sqrt{k/\ln n}) .$$

B.6 Analysis of the distribution of X_i — proofs of Claims B.2–B.4

In this subsection we analyze the probability that any given $X_i = k/2 - s$, and in particular, we prove Claims B.2, B.3, and B.4. Let us first begin with the following simple identity.

$$\begin{aligned} \frac{\Pr[X_i = k/2 - s]}{\Pr[X_i = k/2 - s - \lambda]} &= \frac{\binom{k}{k/2-s}}{\binom{k}{k/2-s-\lambda}} = \frac{k!}{(k/2-s)! \cdot (k/s+s)!} \cdot \frac{(k/2-s-\lambda)! \cdot (k/2+s+\lambda)!}{k!} \\ &= \frac{(k/2-s-\lambda)! \cdot (k/2+s+\lambda)!}{(k/2-s)! \cdot (k/s+s)!} = \prod_{i=1}^{\lambda} \frac{k/2+s+i}{k/2-s-i} \\ &= \prod_{i=1}^{\lambda} \left(1 + \frac{2(s+i)}{k/2+s+i}\right) . \end{aligned}$$

Notice that for $s = 0$ we have

$$\begin{aligned} \frac{\Pr[X_i = k/2]}{\Pr[X_i = k/2 - \lambda]} &= \prod_{i=1}^{\lambda} \left(1 + \frac{2i}{k/2 + i}\right) \leq \prod_{i=1}^{\lambda} \left(1 + \frac{2i}{k/2}\right) \\ &\leq \prod_{i=1}^{\lambda} \exp\left(\frac{2i}{k/2}\right) = \exp\left(\frac{4}{k} \cdot \sum_{i=1}^{\lambda} i\right) = \exp\left(\frac{2\lambda \cdot (\lambda + 1)}{k}\right). \end{aligned}$$

Similarly, using the inequality $1 + x > e^{x/2}$ for all $0 < x < 1$, we have

$$\begin{aligned} \frac{\Pr[X_i = k/2]}{\Pr[X_i = k/2 - \lambda]} &= \prod_{i=1}^{\lambda} \left(1 + \frac{2i}{k/2 + i}\right) \geq \prod_{i=1}^{\lambda} \left(1 + \frac{2i}{k/2 + \lambda}\right) \\ &> \prod_{i=1}^{\lambda} \exp\left(\frac{i}{k/2 + \lambda}\right) = \exp\left(\frac{1}{k/2 + \lambda} \cdot \sum_{i=1}^{\lambda} i\right) = \exp\left(\frac{\lambda \cdot (\lambda + 1)}{k + 2\lambda}\right). \end{aligned}$$

These two bounds immediately imply Claims B.2 and B.3.

Next, let us consider the case $s \geq 1$. We have,

$$\begin{aligned} \frac{\Pr[X_i = k/2 - s]}{\Pr[X_i = k/2 - s - \lambda]} &= \prod_{i=1}^{\lambda} \left(1 + \frac{2(s+i)}{k/2 + s + i}\right) \leq \prod_{i=1}^{\lambda} \left(1 + \frac{2(s+\lambda)}{k/2}\right) \\ &= \left(1 + \frac{2(s+\lambda)}{k/2}\right)^{\lambda} \leq \exp\left(\frac{2(s+\lambda)\lambda}{k/2}\right) = \exp\left(\frac{4(s+\lambda)\lambda}{k}\right). \end{aligned}$$

Notice that this implies the following,

$$\frac{\Pr[X_i = k/2 - s]}{\Pr[X_i = k/2 - s - \lambda]} \leq e, \quad \text{for } \lambda \leq \frac{1}{2} \cdot (\sqrt{k + s^2} - s). \quad (7)$$

Similarly, using the inequality $1 + x > e^{x/2}$ for all $0 < x < 1$, we have

$$\begin{aligned} \frac{\Pr[X_i = k/2 - s]}{\Pr[X_i = k/2 - s - \lambda]} &= \prod_{i=1}^{\lambda} \left(1 + \frac{2(s+i)}{k/2 + s + i}\right) \geq \prod_{i=1}^{\lambda} \left(1 + \frac{2(s+i)}{k/2 + s + \lambda}\right) \\ &> \prod_{i=1}^{\lambda} \exp\left(\frac{2(s+i)}{k/2 + s + \lambda}\right) = \exp\left(\frac{\lambda}{k/2 + s + \lambda} \cdot \left(s + \frac{\lambda + 1}{2}\right)\right). \end{aligned}$$

This immediately implies that for $\lambda > 1 + \frac{1+k/2}{s-1}$ we have

$$\frac{\Pr[X_i = k/2 - s]}{\Pr[X_i = k/2 - s - \lambda]} > \exp\left(\frac{s\lambda}{k/2 + s + \lambda}\right) > e. \quad (8)$$

This implies the proof of Claim B.4.

B.7 The analysis of arbitrary Bernoulli distributions

Our analysis from this section can be extended to a more general case when the cost of every edge is an independent Bernoulli random variable with the success probability p , where p is an arbitrary constant, $0 < p < 1$. That is, with such notation we have $\Pr[X_i = s] = \binom{k}{s} \cdot p^s \cdot (1-p)^{k-s}$ for every i and $0 \leq s \leq k$.

There are only minor changes in the analysis in Subsections B.3–B.6 which need to be done. In principle, the only difference is the proof of Claim B.2, which remains true also for arbitrary $p = \Theta(1)$. Once this claim is proven, the rest of the analysis is essentially identical, and hence we can conclude with the following theorem.

Theorem 10 *For a set of n edge disjoint paths of length k each, where the cost of every edge is an independent Bernoulli random variable with the success probability p , where p is a constant, $0 < p < 1$, it holds that*

$$\mathbf{E}[\Delta] = \Theta(\sqrt{k/\ln n}) .$$