



Graphs with optimal forwarding indices: What is the best throughput you can get with a given number of edges?

Frédéric Giroire, Stéphane Pérennes, Issam Tahiri

► To cite this version:

Frédéric Giroire, Stéphane Pérennes, Issam Tahiri. Graphs with optimal forwarding indices: What is the best throughput you can get with a given number of edges?. [Research Report] RR-8752, INRIA Sophia Antipolis; INRIA. 2015. hal-01172725

HAL Id: hal-01172725

<https://inria.hal.science/hal-01172725>

Submitted on 7 Jul 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Graphs with optimal forwarding indices: What is the best throughput you can get with a given number of edges?

Frédéric Giroire, Stéphane Pérennes, Issam Tahiri

**RESEARCH
REPORT**

N° 8752

June 2015

Project-Team Coati



Graphs with optimal forwarding indices: What is the best throughput you can get with a given number of edges?

Frédéric Giroire, Stéphane Pérennes, Issam Tahiri

Project-Team Coati

Research Report n° 8752 — June 2015 — 26 pages

Abstract: The (*edge*) *forwarding index* of a graph is the minimum, over all possible routings of all the demands, of the maximum load of an edge. This metric is of a great interest since it captures the notion of global congestion in a precise way: the lesser the forwarding-index, the lesser the congestion. In this paper, we study the following design question: Given a number e of edges and a number n of vertices, what is the least congested graph that we can construct? and what forwarding-index can we achieve? Our problem has some distant similarities with the well-known (Δ, D) problem, and we sometimes build upon results obtained on it. The goal of this paper is to study how to build graphs with low forwarding indices and to understand how the number of edges impacts the forwarding index. We answer here these questions for different families of graphs: *general graphs*, *graphs with bounded degree*, *sparse graphs with a small number of edges* by providing constructions, most of them asymptotically optimal. Hence, our results allow to understand how the forwarding-index drops when edges are added to a graph and also to determine what is the *best (i.e. least congested) structure with e edges*. Doing so, we partially answer the practical problem that initially motivated our work: If an operator wants to power only e links of its network, in order to reduce the energy consumption (or wiring cost) of its networks, what should be those links and what performance can be expected?

Key-words: graphs, forwarding index, routing, design problem, energy efficiency, extremal graphs

*This work has been partially supported by ANR project Stint under reference ANR-13-BS02- 0007, ANR program Investments for the Future under reference ANR-11-LABX-0031-01, ANR VISE, CNRS-FUNCAP project GAIATO, the associated Inria team AlDyNet, the project ECOS-Sud Chile.

RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Graphes à indice de transmission optimal : quel est le meilleur débit atteignable avec un nombre d'arêtes donné ?

Résumé : L'indice de transmission (arête) d'un graphe est le minimum, sur tous les routages possibles d'exactly une demande de connexion pour chaque couple de sommets, de la charge maximum d'une arête. Cette métrique est très importante car elle reflète exactement la notion de congestion globale du réseau: plus l'indice de transmission est petit, plus grande est la congestion. Dans cet article, nous étudions la question suivante : étant donné un nombre e d'arêtes et un nombre n de sommets, quel est le graphe le moins congestionné que l'on peut construire et quel est l'indice de transmission que l'on peut atteindre ? Le but de notre étude est donc de déterminer comment construire des graphes avec de petits indices de transmission et de comprendre comment le nombre d'arêtes d'un graphe influence son indice de transmission. Nous répondons à ces questions pour différentes familles de graphes: *graphes généraux*, *graphes à degrés bornés*, *graphes clairsemés avec peu d'arêtes* en donnant dans ces cas des constructions asymptotiquement optimales. Nos résultats permettent non seulement de comprendre comment l'indice de transmission évolue quand on ajoute des arêtes dans un graphe mais ils permettent aussi de déterminer les structures avec e arêtes les moins congestionnées. Ce faisant, nous répondons partiellement au problème pratique qui motiva initialement ce travail : si un opérateur souhaite n'utiliser que e liens de son réseau, ce afin de réduire sa consommation d'énergie (ou ses coûts de câblage), comment choisir ces liens et quelle performance peut on espérer ?

Mots-clés : graphes, indice de transmission, routage, congestion, problème de conception, efficacité énergétique, graphes extrémaux

1 Introduction

Given a graph $G = (V, E)$ with $n = |V|$ vertices, a *routing* R is a collection of paths connecting all the pairs of vertices of G . A routing R induces on every edge e a *load* that is the number of paths going through e . The *edge-forwarding index* (or simply the *forwarding index*) $\pi(G, R)$ of G with respect to R is then the maximum number of paths in R passing through any edge of G . In other words, it corresponds to the maximum load of an edge of G when R is used. So $\pi(G, R)$ measures how congested is the routing R , hence-fore it is important to design routings minimizing this index. The forwarding index $\pi(G)$ of a connected graph G is the minimum $\pi(G, R)$ over all splittable (fractional) routings R 's of G . (We will also sometimes consider non-splittable (integral) routing and denote the minimum load $\pi_i(G)$ in this case.) By definition the forwarding index of a graph measures its intrinsic congestion, so it is a parameter as essential, and arguably more important than simpler parameters such as the diameter or the average distance.

Problem. In this paper, our goal is to provide for a given number of vertices n and for a given number of edges k graphs with the minimum forwarding indices, or at least graphs with low forwarding indices. For a given n , we will study how the number of edges of a graph impacts its forwarding index. Formally, we define the following design problem:

MIN CONGESTED (n, k) -GRAPH: *Given $n, k \in \mathbb{N}$, find a graph $(G = V, E)$ with $|V| = n$ vertices and $|E| = k$ edges such that $\pi(G = V, E)$ is minimum. We will denote this number $\pi^*(n, k)$ (when $k < n - 1$, note that $\pi^*(n, k) = \infty$).*

Motivation. Our problem can be viewed as: for a given bound U on the forwarding index, find a spanner F of G with minimum number of edges such that $\pi(F) \leq U$ or reciprocally given a bound on the number of edges minimize $\pi(F)$.

First, to the best of our knowledge the problem of designing a (sub) graph with minimum forwarding index has not been studied when the main other constraint is the number of edges. Indeed, most of the results have been derived either for classical graphs and graphs families or have been considering other constraints, as example the bounded degree one. So even if a constraint such as the number of edges is both natural and of importance it has not been well studied so far. As example, one of our initial goal was to understand how the forwarding index drops from order n^2 for tree like graphs to order $n \log n$ for cubic graphs, and also to understand how adding a single edge can decrease significantly (or not) the forwarding index.

Second, the recent trend of “Energy Saving” has made our problem even more relevant in practice, especially for network operators willing to reduce the energy consumed by their networks. In fact, most of the network links consume a constant energy independently of the amount of traffic they are flowing. Therefore the only way to reduce the energy used by the network links is to turn some of them off, or more conveniently, put them on an idle mode. Outside the rush hours, several studies [CGM09, Chi08, Chi09, BAH08] show that a good choice of the links to turn off can lead to significant energy savings, while keeping the same communication quality. In the case where the throughputs from every node to every other node are of the same order, and where the capacities also lie in same small range, a good choice of those links amount to solve the problem of finding spanners of the network with low forwarding indices.

Related work. The forwarding-index was introduced by Chung & Al in 87 [Chung87], due to its importance this parameter has been studied quite extensively : on one side results have been given for different graph classes (e.g. random graphs [random-graphs], transitive and Cayley graphs [HMS89, Sol95] graphs with small numbers of vertices [small-graphs] and

well-connected graphs [**well-connected-graphs**]). On the other side deep relations with other expansion-related graph invariants have been established : Laplacian, Cheeger constant (see the survey [Moh97]), Sparsest cut [LeRa99] and the “geometry of graphs” [LLR94]. This notion has also been used to prove that some Markov chains mix fast using either canonical paths (routings) or “resistance” [Sinclair92]. See the recent survey [JM12] for a global view on the known results. The problem is also known as the *maximum concurrent flow problem* and its dual was probably first introduced in [ShMa90] in which the authors also discussed the relation with the network throughput, in [EpTa92] a simple oblivious packet routing algorithm achieving network stability for any rate λ with $\lambda\pi < 1$ was provided. Some variants: load on arcs for digraphs ([MaTu96]) load on the vertices have also been studied.

The edge forwarding index is strongly related to distance properties of the graph. Indeed a usual naive lower bound on π (Average distance Bound) is:

$$\pi(G = V, E) \geq \frac{\sum_{(u,v) \in V^2} D(u,v)}{|E|} = \frac{\overline{D_G}|V|^2}{|E|} = 2|V|\frac{\overline{D_G}}{\overline{d_G}},$$

where $D(u,v)$, $d(v)$, $\overline{D_G}$ and $\overline{d_G}$ denote respectively the distance function, the degree function, the average distance in G and the average degree in G . This indicates that solving our design problem is strongly related to finding graphs with small average distance. The Degree-Diameter problem or (Δ, D) -DESIGN PROBLEM is about finding the graph with degree Δ and diameter D with the maximum number of vertices (or reciprocally it is about minimizing the diameter of a Δ -regular graph). It is quite a complex problem and it has been studied extensively (see [MiSi13] for a recent survey). Even after 30 years of steady efforts, generic constructions are still very far from being optimal. So, since good (n, k) -graphs should resemble (Δ, D) graphs, we may expect our problem to be complex. But we can also hope to be able to use results about the (Δ, D) -problem in our context.

Contributions and plan of the paper

- In Section 2, we provide formal definitions and present some important properties of the forwarding index, in particular its dual formulation which involves metrics.
- In Section 3, we consider our design problem for general graphs, that is when the only design constraint is the number of edges. We characterize the graphs with minimum forwarding index. When the number of edges is $k(n-k)$, $k \in \mathbb{N}$, optimum graphs happen to have a simple structure since they are the complete bipartite graphs $K_{k, n-k}$. In between these values, the function $\pi^*(n, k)$ follows, rather surprisingly, a stepwise function (see Propositions 3.4 and 3.5).
- In Section 4, motivated by telecommunication networks, we study the case of bounded degree graphs. We provide almost optimal graphs for the different values of maximum degree Δ . We then focus on graphs with a small number of edges ($\Delta = 3$) as they correspond to the range of values for which the forwarding index greatly changes. We determine quite sharply how the minimum forwarding index behaves and evolves from $\Theta(n^2)$ to $\Theta(n \log n)$ when the number edges grows from $n-1$ to $n + \frac{n}{2}$. We also develop a method that allow us simplify the design problem by considering the *graph skeleton*.
- We then examine the case $e = n + k$ with a fixed small $k \in \{1, 2, 3\}$ in Section 5. We determine the minimum forwarding index exactly for any n . This is possible because the main structure of the graph, that we called skeleton is finite, so we can explore all of them and use weight arguments in order to deal with a finite problem. Some of the results, as example Proposition 5.4, are strikingly counter intuitive.

- Last, in Section 6, we provide optimal cubic-graphs with *small number of vertices*, that is for $n \in [4, 22]$. Those graphs are not only interesting per se (and some structures again are surprising), but also because, as we shall see, their structure may be used as a *skeleton* to build good graphs with a few edges and arbitrarily size.

2 Definitions and Preliminary Results

Definition 2.1 (edge forwarding-index)

Fractional π : Given a graph $G = (V, E)$ the edge forwarding-index $\pi(G)$ is the solution of the next linear problem where \mathcal{P} denotes the set of paths of G , $\mathcal{P}(u, v)$ is the set of paths connecting u and v and $\lambda(P), P \in \mathcal{P}$ are the variables.

$$\begin{array}{llll} & \text{Minimize } \pi & & \\ \forall P \in \mathcal{P} & \lambda(P) & \geq & 0 \\ \forall e \in E & \sum_{P \in \mathcal{P}, e \in P} \lambda(P) & \leq & \pi \\ \forall (u, v) \in V^2 & \sum_{P \in \mathcal{P}_{uv}} \lambda(P) & \geq & 1 \end{array}$$

For the **integral case**, we require the variables $\lambda(P)$ to be integral and we denote the minimum load $\pi_i(G)$.

Weighted Case: in the case of a weighted graph (G, w) , $w(v)$ represents the number (or amount since $w(v)$ may not be integral) of vertices located at v , so we require $w(u)w(v)$ paths connecting u to v : $\forall (u, v) \in V^2, \sum_{P \in \mathcal{P}(u, v)} \lambda(P) = w(u)w(v)$.

First, let us say why we won't focus much on the integral forwarding index. First, it is "quite surely" NP-Complete to compute (to the best of our knowledge the result was proven by Saad, but only in the case of the vertex forwarding index [Saa93]). Thus, using it as main network parameter would not be so wise. Secondly, the integrity gap for the forwarding index problem is small for most practical values of π . This was proven by Raghavan and Thompson, who provided [RaTo87] a now standard randomized procedure that proves that $\pi_i(G) \leq \pi(G) + O(\sqrt{\log_2(|V||E|)})$. Since for most of the cases we will study, $\sqrt{\log_2(|V||E|)} = o(\pi(G))$ we will almost always be in the case for which $\pi_i(G) = (1 + o(1))\pi(G)$. So, we will mostly study the fractional edge forwarding index.

When one assumes uniform traffic demands between all ordered pairs of vertices at a given rate λ (i.e u send a packet to v with probability λdt) one can prove that whatever be the routing policy the maximum load over all edges will be at least $\lambda\pi(G = V, E)$; moreover there exists routing policies (best effort ones [EpTa92]) that ensure that if edges have throughput at least $\lambda\pi(G)$ then the network will be stable and the traffic routed. If we adopt a coarser point and view and look at demands as paths, a picture of the network at a given moment is a set of randomly chosen paths picked independently with probability λ and so each arc receive a load sharply concentrated around $\lambda\pi(G)$.

This makes the edge forwarding index of great practical importance since the maximum rate that a network can tolerate is $\frac{1}{\pi(G)}$.

So, our measure of goodness for a graph will be its edge forwarding index. This quantity is strongly related to distance properties of the graph. Indeed a usual naive lower bound on π is:

$$\pi(G = V, E) \geq \frac{\sum_{(u, v) \in V^2} D(u, v)}{|E|} = \frac{\overline{D(G)}|V|^2}{|E|} = 2|V| \frac{\overline{D(G)}}{\overline{d(G)}}$$

Where $D(u, v)$, $d(v)$, $\overline{D(G)}$ and $\overline{d(G)}$ denote respectively the distance function, the degree function, the average distance in G and the average degree of G . Moreover this bound is attained if and only if there exists a shortest paths routing that is balanced on the edges. This is the case graphs such as cycles, toruses and any other edge transitive graph (see [Sol95]). This indicates that solving our design problem is strongly related to finding a spanning subgraph with small average distance. The Degree-Diameter problem $((\Delta, D))$ design problem is about finding the graph with degree Δ and diameter D with the maximum number of vertices (or reciprocally it is about minimizing the diameter of a Δ -regular graph). It's quite a complex problem and it has been studied extensively (see [MiSi13] for a recent survey). Even after 30 years of steady efforts generic constructions are still very far from being optimal. So, since good (n, k) -graphs should resemble (Δ, D) graphs we may expect our problem to be complex but we can also hope to be able to use results about the (Δ, D) problem in our context.

Duality of linear programming implies that the relationship between distance and edge forwarding index is exact, but it holds for distorted metric. We let \mathcal{P} denote the set of all paths and \mathcal{P}_{uv} denote the set of all paths from u to v . As we already seen the $\pi(G)$ is the solution of the linear problem :

$$\begin{array}{llll} & \text{Minimize } \pi & & \\ \forall P \in \mathcal{P} & \lambda(P) & \geq & 0 \\ \forall e \in E & \sum_{P \in \mathcal{P}, e \in P} \lambda(P) & \leq & \pi \\ \forall (u, v) \in V^2 & \sum_{P \in \mathcal{P}_{uv}} \lambda(P) & \geq & 1 \end{array}$$

The dual of this problem is indeed a distance maximization problem, but first we need the following definition : Given a positive length function, $l : E \rightarrow \mathbb{R}^+$, we note $l(P)$, for $P \in \mathcal{P}$, the length of P . The dual of the edge forwarding index problem is then defined as:

$$\begin{array}{llll} & \text{Maximize } D_{tot} & & \\ \forall P \in \mathcal{P}_{u,v} & l(P) & \geq & l(u, v) \\ & \sum_{e \in E} l(e) & = & 1 \\ & \sum_{(u,v) \in V^2} l(u, v) & = & D_{tot} \end{array}$$

If we denote D_{tot}^* the maximum of the above problem, Farkas lemma (strong duality) implies that :

$$\pi(G = V, E) = D_{tot}^*(G = V, E)$$

In other words, the edge forwarding index is indeed an average distance on G but taken over a distorted metric. In many case two simple metrics often give the optimal value : either the uniform one (ie simply the usual graph distance) or a good cut. We say that a set $S \subset E$ is a cut if removing S disconnect the graph into two part A and $V \setminus A$ and we denote then S as $[A, V \setminus A]$. If we choose $l(e) = 0, e \notin S, l(e) = \frac{1}{|S|}$, we then have $d(a, b) \geq \frac{1}{|S|}$ for $a \in A, b \in V \setminus A$ and we get the following cut bound :

$$\forall S \subset V, \pi(G = V, E) \geq \frac{2|A||V \setminus A|}{|[A, V \setminus A]|}.$$

Due to Bourgain theorem [Bou85] (which claims that any metric is up to distortion $O(\log n)$ a positive linear combination of cut metrics, see also [LeRa99, LLR94]) these cut bounds are relatively tight since for any graph $G = (V, E)$ with n vertices:

$$\pi(G = V, E) \leq O(\log |V|) \times \text{Max}_{A \subset V} \frac{|A||V \setminus A|}{|[A, V \setminus A]|}$$

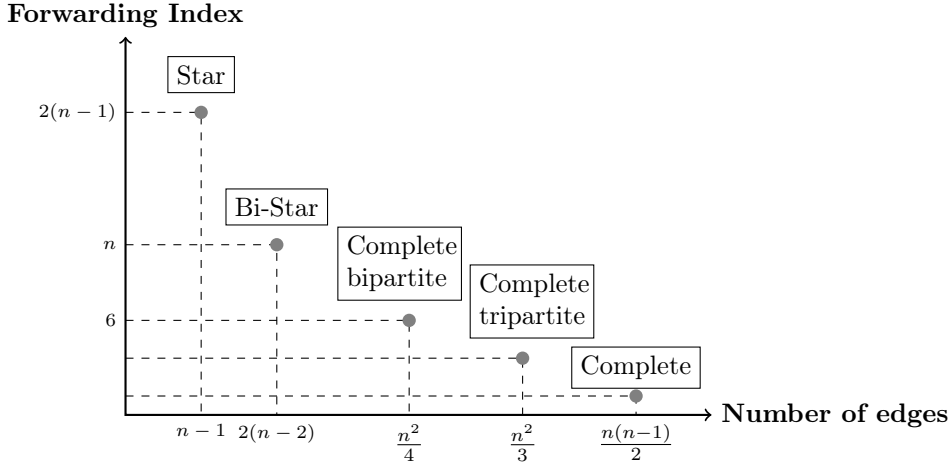


Figure 1: Forwarding indices of minimally congested graphs with n vertices as a function of their number of edges.

This actually implies that bounds provided by cut arguments are tight up to a factor of order $\log n$, and this gap is “tight” since for expander graphs the above bound is attained, i.e. $\pi(G) = O(\log |V|) \times \text{Max}_{A \subset V} \frac{|A||V \setminus A|}{|[A, V \setminus A]|}$.

Along the lines of this paper we will use dual arguments to provide lower bounds and guide our constructions (by somewhat fulfilling the complementary slackness conditions). It will turn out that the bounds provided by cuts and the uniform metric will be almost always sufficient.

3 Minimally congested graphs

In this section, we study the design of minimally congested graphs for given numbers of vertices n and edges e . We first give a trivial lower bound of $\pi^*(n, e)$, the minimum forwarding index of a (n, e) -graph. We then provide families of minimally congested graphs reaching this bound for some couples of values (n, e) , e.g. complete bipartite graphs $K_{i, n-i}$, complete k -partite graphs, or Kneser graphs, see Figure 1. These graphs are edge-transitive and of diameter 2. In particular, we show that $K_{i, n-i}$ ($i \in \mathbb{N}, i \leq \lfloor n/2 \rfloor$) are minimally congested $(n, i(n-i))$ -graphs with forwarding index $\pi^*(n, e) = 2(\frac{n(n-1)}{e} - 1)$. Last, we study the behavior of $\pi^*(n, e)$ when e varies between two “perfect” cases, from $i(n-i)$ to $(i+1)(n-(i+1))$. Surprisingly, π^* follows a step-wise function in the sense of Propositions 3.4 and 3.5 and jumps suddenly from $\pi^*(n, i(n-i))$ to $\pi^*(n, (i+1)(n-(i+1)))$.

Proposition 3.1 (Lower bound on $\pi^*(n, e)$.) *The forwarding index of an (n, e) -graph is lower bounded by:*

$$\pi^*(n, e) \geq \frac{2n(n-1)}{e} - 2.$$

Proof. The lower bound comes from the usual distance bound (uniform metric). In a graph with e edges, we find $2e$ couples of vertices at distance 1 (two per edge), and the $n(n-1) - 2e$ other couples are at distance at least 2. It follows that $\sum_{(u,v) \in V^2} D(u, v) \geq 2n(n-1) - 2e$. Since there are e edges $\pi^*(n, e) \geq \frac{\sum_{(u,v) \in V^2} D(u, v)}{e} \geq \frac{2n(n-1) - 2e}{e} = \frac{2n(n-1)}{e} - 2$. ■

Proposition 3.2 (Optimal (n, e) -graph) *An (n, e) -graph that is edge-transitive and of diameter 2 is optimal. Its forwarding index is*

$$\frac{2n(n-1)}{e} - 1.$$

Proof. The proof comes down to providing a (fractional) shortest path routing for these graphs loading uniformly every edge. Given any couple of nodes $(u, v) \in V$, we route the flow from u to v uniformly along all the shortest paths. Since the graph is edge transitive, the load on the edges is then from construction uniform and exactly equal to $\frac{\sum_{(u,v) \in V^2} D(u,v)}{e}$. Note that, we are in the nice but rare case in which a shortest path routing that loads uniformly the edges exists. ■

Corollary 3.1 (Families of optimal graphs) *The following families of graphs are optimal:*

- Complete bipartite graphs, giving:

$$\pi^*(n, i(n-i)) = \frac{2n(n-1)}{e} - 2, \quad i \in \mathbb{N}, i \leq \lfloor n/2 \rfloor.$$

- Turán graphs $T(n, r)$, for which r divides n (that is, complete multipartite regular graphs with r independent subsets of equal sizes), giving:

$$\pi^*\left(n, \frac{n}{2}\left(n - \frac{n}{r}\right)\right) = \frac{2n(n-1)}{e} - 2, \quad r \in \mathbb{N}, r \leq n.$$

- Kneser graphs $KN_{\nu, \kappa}$ for which $\kappa \geq \nu/3$ (Kneser graphs of diameter 2), giving:

$$\pi^*\left(\binom{\nu}{\kappa}, \frac{1}{2}\binom{\nu}{\kappa}\binom{\nu-k}{\kappa}\right) = \frac{2n(n-1)}{e} - 2, \quad \nu \in \mathbb{N}, \nu/3 \leq \kappa \leq \nu.$$

Proposition 3.3 (Integral Forwarding Index)

- Complete bipartite graphs are (almost) optimal, in the sense that, for $i \in \{1, 2, \dots, \lfloor n/2 \rfloor\}$, we have:

$$\pi_i^*(n, i(n-i)) \in \lceil \pi^*(n, i(n-i)) \rceil + \{0, 1, 2, 3, 4\}.$$

- Turán graphs $T(n, r)$, for which r divides n are (almost) optimal, in the sense that, for $i \in \{1, 2, \dots, \lfloor n/2 \rfloor\}$, we have:

$$\pi_i^*\left(n, \frac{n}{2}\left(n - \frac{n}{r}\right)\right) = \pi^*\left(n, \frac{n}{2}\left(n - \frac{n}{r}\right)\right) + \{0, 1, 2, 3, 4\}.$$

Proof. We provide an integral routing scheme for complete bipartite graphs $K_{i, n-i}$. We note V_i (resp. V_{n-i}) the side of the partition of size i (resp. of size $n-i$). For two nodes u and v on different sides of the partition, we directly route through the edge uv . For two nodes $u, v \in V_{n-i}$ (resp. $u, v \in V_i$), we route through the intermediary vertex $u+v \bmod i$ ($u+v \bmod n-i$). The load of this integral routing is then at most $\lceil \pi^* \rceil + 4$.

What is the load of the edge ux induced by two vertices in V_{n-i} ? The edge ux is used by paths between u and w , with w such that $u+w = u+v \bmod i$. There are at least $\lfloor (n-i)/i \rfloor$ and at most $\lceil (n-i)/i \rceil$ such nodes. This gives that the load of ux induced by nodes in V_{n-i} is at least $2\lfloor (n-i)/i \rfloor$ and at most $2\lceil (n-i)/i \rceil$. Similarly, the load of ux induced by nodes in V_i

is at least $2\lfloor i/(n-i) \rfloor$ and at most $2\lceil i/(n-i) \rceil$. Thus the load of this integral routing is then at most $\lceil \pi^* \rceil + 4$.

Similarly, we provide an integral routing for *Turán graphs*. For two nodes u and v in two different parts of the partition, we route them directly through edge uv . For two nodes u and v in the same part of the partition, we route through the intermediary vertex $u + v \bmod n - n/r$. For similar reasons as in the case of complete bipartite graphs, the load of this integral routing is then at most $\lceil \pi^* \rceil + 4$. ■

Since $\pi^*(n, e)$ decreases with e the above results implies that $\pi^*(n, e)$ evolves like $\Theta(\frac{2n^2}{e})$, but we don't know yet the precise behavior of $\pi^*(n, e)$ between two perfect cases (i.e. $e = i(i-k)$). As we shall prove this behavior is not a smooth linear decrease since it indeed proceeds with jumps occurring at values close to those perfect ones. First, we start studying the intermediary cases when e starts at $n-1$ ($\pi^*(n, e) = 2(n-1)$, optimal graph is a star) and grows to $e = 2(n-2)$ ($\pi^*(n, e) = n-2$, optimal graph is $K_{2, n-2}$). The next proposition shows that when e get larger than $n-1$, first π^* does not decrease significantly and stays around $2(n-1)$ then it jumps abruptly down to $n-1$ when e get close to $2(n-2)$. The proof that we give immediately extends then to larger values of e , but we prefer to give first detailed arguments for e in $[n-1, 2(n-2)]$.

Proposition 3.4

$$\begin{aligned} \forall e \in [n-1, 2(n-2) - o(n)] \quad & \pi^*(n, e) = 2(n-1) + o(n) \\ e = 2(n-2) \quad & \pi^*(n, e) = (n-1) + o(n) \end{aligned}$$

Proof.

In order to prove the result we consider some fixed arbitrarily small positive α and consider a graph G such that $\pi(G) \leq 2n(1-\alpha)$.

Remark first that a graph with low forwarding index must have an average distance close the one on $K_{k, n-k}$ which is 2. This implies the existence of one or many nodes with large degree (*hub nodes*) interconnecting the others that we shall call *leaves nodes*. This happens because only nodes with large degree generate many paths of length 2 per edge. So, to perform the analysis we fix a “large” number d_0 (to be fixed later) and we partition our graph into two sets :

- the leaf set L containing nodes with degree less than d_0 ,
- the hub set H containing nodes with degree $\geq d_0$.

Note first, that the hub set cannot be too big, since counting the edges we have $d_0|H| \leq 2|E|$ so:

$$|H| \leq \frac{2e}{d_0} \tag{1}$$

Then, we consider the set \mathcal{C} of connected components of the graph induced by the leaf set. For $C \in \mathcal{C}$ we denote $c = |C|$, $\delta(C) = |E \cap [C, H]|$ the size of the border of C (connecting C to H) and $e(C) = |E \cap [C, C]|$ the number of edges inside C . Last, we denote by $int(C)$ the set of connections from C to C that are routed using only edges inside C . More exactly, since we consider a routing that is fractional $int(C)$ is the sum of the weight of the paths routed inside C .

We shall prove that:

$$\forall C \in \mathcal{C}, \delta(C) + e(C) \geq 2c \tag{2}$$

The proposition 3.4 immediately follows since :

$$e \geq \sum_{C \in \mathcal{C}} (\delta(C) + e(C)) \geq \sum_{C \in \mathcal{C}} 2|C| = 2|L| = 2(n - |H|) \geq 2(n - \frac{2e}{d_0})$$

and since $e \leq 2n$ we conclude that

$$e \geq 2n(1 - \frac{4}{d_0}) \quad (3)$$

In order to prove equation 2, we estimate the load of the border of C :

- there are $2c(n - c)$ connections (between C and \overline{C}) that each load the border of C at last once.
- the amount of requests from C to C that are not routed fully inside C is $c^2 - \text{int}(C)$, and each load the border of C at least twice since it must leave C and re-enter.

So the total load l on the border of C is such that:

$$l \geq 2c(n - c) + 2(c^2 - \text{int}(C)) = 2cn \left(1 - \frac{\text{int}(C)}{cn}\right)$$

We now distinguish two cases :

- If $\text{int}(C) < \alpha cn$ (which holds for $c = o(n)$), then l is at least $2cn(1 - \alpha)$. So if $\delta(c) \leq c$ the load of one edge is at least $\frac{l}{\delta(c)} \geq 2n(1 - \alpha)$ and $\pi(G) \geq 2n(1 - \alpha)$. which is a contradiction. So $\delta(C) \geq c + 1$, and since C is connected $e(c) \geq c - 1$ and $e(c + \delta(c)) \geq 2c$. and equation 2 holds.
- If $\text{int}(C) \geq \alpha nc$. Using the Moore bound, we remark that for a given vertex of C there are at most $2d_0(d_0 - 1)^{l-1}$ connections that can be routed fully inside C along path of length l . This implies that the average length of an internal connection is at least $\log_{d_0} \left(\frac{\text{int}(C)}{c}\right)$. So the total load of the internal connections on the edges inside C is at least $\text{int}(C) \log_{d_0} \left(\frac{\text{int}(C)}{c}\right) \geq \alpha nc \log_{d_0}(\alpha n)$. The total load cannot exceed $\pi(G)e(C)$ and since $\pi(G) \leq 2n$, we must have

$$\alpha n \log_{d_0}(\alpha n) \leq \pi(G)e(C) \leq 2ne(C)$$

This implies :

$$\frac{e(C)}{c} \geq \frac{\alpha \log_{d_0}(\alpha n)}{2}$$

The above function diverges to $+\infty$ when n grows, so for n large enough (namely $n \geq \frac{1}{\alpha} d_0^{4/\alpha}$) we get $e(C) \geq 2c$ and equation 2 is satisfied.

To summarize, if we look at a graph for which $\pi(G)$ is only a tiny bit smaller than $2n$ (i.e. $2n(1 - \alpha)$ for some $\alpha > 0$) and wish to prove that it has at least $2n(1 - \beta)$ edges (for some $\beta > 0$) we simply pick $d_0 \geq \frac{4}{\beta}$ and fix n to be $\geq \frac{1}{\alpha} d_0^{4/\alpha}$. Then equation 2 is true, equation 3 follows and $e(G) \geq 2n(1 - \frac{4}{d_0}) \geq 2n(1 - \beta)$

■ The result can be extended to larger values of e ($e = n + k$ with $k = o(n)$), see Proposition 3.5. . It is almost immediate. Note first that for $e = xn$ with $x = o(n)$ we have $\pi^*(n, xn) \geq \frac{2n^2}{xn} + o(n) = \frac{2n}{x}(1 + o(1))$. What we show here is that $\pi^*(n, xn)$ actually behave like $\frac{2n}{\lfloor x \rfloor}$. The proof is indeed the same, so we do not repeat the argumentation.

Proposition 3.5 *For any fixed $k \in \mathbb{N}$:*

$$\forall e \in [kn, (k + 1)n - o(n)] \quad \pi^*(n, e) = \frac{2n}{k} + o(n)$$

Proof. The proof for the case $k = 1$ is given in proposition 3.4, the general argument is almost the same up to cosmetic changes (mostly re-introducing the missing value k which is 1 in the proof of proposition 3.4).

Again we consider a graph G such that $\pi(G) \leq \frac{2}{k}n(1 - \alpha)$. and like in the case $k = 1$ we consider a “large” degree d_0 , the leaf set L and the hub set H . We still have

$$|H| \leq \frac{2e}{d_0} \quad (4)$$

Then, we consider the set \mathcal{C} of connected components of the graph induced by the leaf set, and for a component C we shall prove that

$$\forall C \in \mathcal{C}, \delta(C) + e(C) \geq (k + 1)c \quad (5)$$

Like in the case $k = 1$ the proposition 3.5 follows since :

$$e \geq \sum_{C \in \mathcal{C}} (\delta(C) + e(C)) \geq (k + 1)(n - |H|) \geq (k + 1)(n - \frac{2e}{d_0})$$

and since $e \leq (k + 1)n$ we conclude that

$$e \geq (k + 1)n(1 - \frac{2(k + 1)}{d_0}) = (k + 1)n(1 - o(1)) \quad (6)$$

In order to prove equation 5, we estimate the load of the border of C : the total load l on the border of C is still :

$$l \geq 2cn \left(1 - \frac{\text{int}(C)}{cn} \right)$$

We now distinguish two cases :

- If $\text{int}(C) < \alpha cn$ then l is at least $2cn(1 - \alpha)$ and if $\delta(C) \leq kc$ then the load of one edge is at least $\frac{l}{\delta(C)} \geq \frac{2}{k}n(1 - \alpha)$ and $\pi(G) \geq \frac{2}{k}n(1 - \alpha)$ (contradiction). So $\delta(C) \geq kc + 1$, and $e(C) + \delta(C) \geq kc + 1 + (c - 1) = (k + 1)c$. and equation 2 holds.
- If $\text{int}(C) \geq \alpha nc$. we use the exact same Moore bound and conclude (like in the case $k = 1$) that

$$\frac{e(C)}{c} \geq \frac{\alpha \log_{d_0}(\alpha n)}{2}$$

The above function diverge to $+\infty$ when n grows, so for n large enough (namely $n \geq \frac{1}{\alpha} d_0^{\frac{2(k+1)}{\alpha}}$) we get $e(C) \geq (k + 1)c$ and equation 2 is satisfied.

To summarize, if we look at a graph for which $\pi(G)$ is only a tiny bit smaller than $\frac{2}{k}n$ (i.e. $\frac{2}{k}n(1 - \alpha)$ for some $\alpha > 0$) and can to prove that it has at least $2n(1 - \beta)$ edges (for some any $\beta > 0$) as long as n is large enough (take $d_0 \geq \frac{4}{\beta}$ and fix n to be $\geq \frac{1}{\alpha} d_0^{\frac{2(k+1)}{\alpha}}$). Then equation 5 is true, equation 6 follows and $e(G) \geq (k + 1)n(1 - \frac{4}{d_0}) \geq (k + 1)n(1 - \beta)$

■




Figure 2: Forwarding indices of minimally congested graphs with n vertices as a function of their number of edges.

4 Bounded degree graphs with low edge forwarding index

In the preceding section, we provided somewhat optimal families of graphs. This solves the question of minimally congested graphs in the general case. We now study graphs with a constraint on the degree (Δ will denote the maximum degree). The motivation comes from telecommunication & real interconnection networks for which the node degree is often small, see for example [sndlib, GPT15]. In this section, we consider first the general case for $\Delta \geq 3$ ($\Delta = 2$ is trivial) and we succeed in determining how the forwarding index drops from $\pi(n, e) = n^2/4$ to $\frac{2}{3}n \log_2 n$ when the average degree raises from 2 to 3. So, we focus on graphs with a *small number of edges*, namely graphs with average degree $\bar{\Delta} \in [2, 3[$, that is when $e \in [n, \frac{3}{2}n]$, and we study the transition of $\pi(n, e)$ from $\frac{n^2}{4}$ to $\Theta(n \log n)$ when the number of edges e raises from $n-1$ to $\frac{3}{2}n$.

4.1 Graphs with bounded degree Δ : some remarks.

For $\Delta = 3$, when $e = \frac{3n}{2}$, graphs such like the shuffle exchange provide deterministic generic constructions for which $\pi(G) \leq n \log_2 n$ (this is a folk result for people studying network throughput, one may see [JM12]). Since using the Moore bound (that bound claims by direct counting that the average distance in a Δ bounded degree graph is of order $\log_{\Delta-1}(n)$, see as example [MiSi05]) one can prove that $\pi^*(n, \frac{3n}{2}) \geq \frac{2}{3}n \log_2 n(1+o(1))$ the lower and upper bounds match up to factor of $\frac{2}{3}$. Moreover we shall prove that random cubic graphs are almost optimal since with high probability they are such that $\pi(G) = \frac{2}{3}n \log_2 n(1+o(1))$. Moreover for larger values of Δ de Bruijn graphs and their variants provide Δ -regular graphs whose forwarding index is of the right order (see Figure 2). So when the degree is bounded by Δ , the value of $\pi(n, \frac{\Delta}{2}n)$ is relatively well understood (see [Chung87, Xu2012]), and structures close to the optimal are obtained using de Bruijn graphs or slight variants of it. Indeed, on the one hand, the Moore

bound implies that :

$$\pi^*(n, \frac{\Delta}{2}n) \geq \frac{2}{\Delta} n \log_{\Delta-1} n(1 - o(1)).$$

On the other hand, for de Bruijn graphs, one has (see [Chung87, Xu2012])

$$\pi(n, \frac{\Delta}{2}n) \leq \frac{2}{\Delta} n \log_{\lfloor \frac{\Delta}{2} \rfloor} n.$$

The argument that provides the above bound for the de Bruijn graph with degree $\Delta = 2d$ and d^n vertices, is quite simple since it exists in this graph an integral routing that is uniform on the edges and that connects each couple of vertices with a path of length exactly n . This length is only a constant factor larger than the minimum average distance predicted by the Moore bound, hence the ratio between the above upper and lower bound is at most 3 and decreases with Δ .

So our purpose is to understand what is happening between two well understood situations : $e = n - 1, \pi^*(n, e) = \frac{n^2}{4}$ and $e = \frac{3}{2}n, \pi^*(n, e) = \Theta(n \log n)$ that is when e evolves in $[n, \frac{3}{2}n]$, in other words we shall study the evolution of $\pi^*(n, e)$ when the number of edges e raises from $n - 1$ to $\frac{3}{2}n$.

4.2 Case $e \in [n, \frac{3}{2}n]$ for $\Delta \leq 3$, lower bound

In this section, we provide a lower bound on the forwarding indices of graphs with $e \in [n, \frac{3}{2}n]$ and $\Delta \leq 3$

Proposition 4.1 *If G is a $(n, n + k)$ graph with $\Delta = 3$ then $\pi(G) \geq \frac{(n-2k)^2}{3k}(\log(3k/2) - O(\log \log(k)))$.*

First let us define two important quantities \mathbf{W}_1 (resp. \mathbf{W}_2) which intuitively are the maximum weights (i.e number of vertices) that can lie inside a pending tree (resp. inside a subdivided edge) of the graph.

$$\begin{aligned} \pi(G) &= 2\mathbf{W}_1(n - \mathbf{W}_1) \\ \pi(G) &= 2\mathbf{W}_2(n - \mathbf{W}_2)/2 = \mathbf{W}_2(n - \mathbf{W}_2) \end{aligned}$$

The next lemma tells how one can construct graph with forwarding-index $\pi(G)$

Lemma 4.1 (skeleton) *Any n a min-congested graph with $n + k$ edges can be constructed as follows :*

- Take a cubic graph (with potentially multiple edges) with at most $x = 2k$ vertices, put at most 5 new edge-node on the edges and affect a total weight $n - x$ to the edge nodes such that i) each edge-node has weight lesser than \mathbf{W}_1 , ii) each contains weight lesser than \mathbf{W}_2 .
- If instead we subdivide the edges once, and use weights lesser than \mathbf{W}_2 the smallest forwarding index we can achieve is a lower bound on $\pi^*(n, k)$

4.2.1 proof of lemma 4.1

Proof. Eliminating degree 1 nodes. Let π_0 be the forwarding index of our graph. First we deal with degree 1 nodes. When a node is a pending leave we identify it with its father and we consider that the father has weight 2. Doing this procedure inductively, we collapse trees into nodes of some weight, and two nodes u, v with respective weights $w(u), w(v)$ have to be connected with $2w(u)w(v)$ paths. Since tree have a high load, the maximum size of a tree obtained running

this procedure is limited. Let $2\mathbf{W}_1(n - \mathbf{W}_1) = \pi_0$, then if a tree with weight \mathbf{W}_1 is attached to the network with a single edge, we have

$$\pi(G) \geq \mathbf{W}_1(n - \mathbf{W}_1) = \pi_0$$

So the weights are upper bounded by \mathbf{W}_1 and the procedure terminate. At the end we are left with nodes with degree at least 2 and possible weights. Moreover only nodes with degree 2 can have a weight ≥ 1 . We call this transformed graph G_1 . When performing the above transformation (replacing a tree by a weighted node) the value $k = e - n$ is not modified since we remove one edge and one node. It follows that G_1 has k extra edges (n_1 vertices and $n_1 + k$ edges). The procedure can be reverted since from a solution with weighted nodes of degree 2, such that $\forall v \in V, w(v) \leq \mathbf{W}_1$ we can derive a solution without weight by replacing a node with weight $w(v)$ by a node with weight 1 attached to a balanced binary tree with size $w(v) - 1$. This is possible since the load inside the tree (which is entirely determinate) will be always less than $2w(v)(n - w(v)) \leq \pi_0$ (here we obviously assume $\mathbf{W}_1 < n/2$)

It follows that we can simply consider graphs with nodes of degree 2, 3 with a possible bounded weight w on the nodes of degree 2 and such that $\forall v \in V, w(v) \leq \mathbf{W}_1$.

Eliminating degree 2 nodes. The graph G_1 has n_1 vertices and $n_1 + k$ edges, so it has less than $2k$ vertices of degree 3. Such vertices are connected by paths formed by degree 2 nodes that are potentially weighted (if a tree was initially attached in G at those nodes). Since the edge border of such a set has size 2, the sum of the weights of the nodes forming a subdivided edge is again bounded by a value \mathbf{W}_2 that satisfies :

$$\mathbf{W}_2(n - \mathbf{W}_2) = \pi_0 \quad (\text{LB1})$$

We may now replace paths with total weight w with an edge subdivided once with a *node-edge* with weight w and we obtain the graph G_2 , and assertion (ii) of the lemma is proven.

But unlike in the case of trees, the operation cannot always be reverted.

- Indeed if we try to emulate an edge with weight w by attaching a tree with weight w to a unique node that we place on the edge this will fail to produce a graph with forwarding-index lesser than π_0 whenever $w \geq \mathbf{W}_1$.
- if two nodes u, v with weight w_u, w_v are adjacent on a subdivided edge we can replace them safely by a single node with weight $w_u + w_v$ as long as $w_u + w_v \leq \mathbf{W}_1$. So we cannot always preserve the forwarding index using a single edge-node per edge. Indeed, since we can merge adjacent vertices whenever $w_u + w_v \leq \mathbf{W}_1$ we can always ensure that after the reduction we find at most $2\lceil \frac{\mathbf{W}_2}{\mathbf{W}_1} \rceil \leq 5$ per edge-nodes per edge.

This proves assertion (i). ■

Proof of proposition 4.1 : Counting distances of the skeleton

Definition 4.1 (skeleton) Given a $n, n + k$ graph, we call the graph G_2 of lemma 4.1 the skeleton of G , and we will denote it $S[G]$, from construction :

- it has at most $2k$ vertices, $\Delta = 3$,
- its forwarding index is at most $\pi(G)$,
- each of its edges “contains a weight” lesser than \mathbf{W}_2 and the sum of those weights is $\overset{\text{def}}{\mathbf{ns}} = n - 2k$.

Proposition 4.1: If G is a $(n, n+k)$ graph with $\Delta = 3$ then

$$\pi(G) \geq \pi^*(n, n+k) \geq \frac{(n-2k)^2}{3k} (\log(3k/2) - O(\log \log(k))).$$

Proof. In order to derive a lower bound on $\pi(G)$ we only need to prove one for $\pi(S[G])$. To do that, we simply count the total distance between the nodes located inside the edges (represented by weights). For this we use uniformly the *uniform metric with $l(e) = 1$ for any edge e of the skeleton*, more formally we consider the subdivided skeleton (with a single edge-node per edge, this forms a bipartite graph with $6k$ edges) and each edge is then given length $\frac{1}{2}$. Then we lower bound the total / average distance using the Moore bound while taking into account the fact that we find at most \mathbf{W}_2 nodes inside an edge. For sake of simplicity we only sum on the distances between the $\mathbf{ns} = n - 2k$ nodes (that we call *edge-nodes*) appearing inside the edges and we ignore the $2k$ nodes of the skeleton (but they can be taken into account using three different and similar Moore bounds).

Then, from a given node edge we find at most $s(h)$ edge nodes node edges at distance h with $s(0) = 1, h \geq 1, s(h) \stackrel{\text{def}}{=} 2^{h+1}$ (note that the -at most- 4 edge-nodes adjacent to an edge-node are indeed at distance 1 since they are connected by two edges with length $\frac{1}{2}$, and so $s(1) \leq 4$). Hence, at distance $\leq i$, we find at most $b(i) \stackrel{\text{def}}{=} \sum_{h=0}^i s(h) \leq 2^{i+2}$ edge nodes. Since we are trying to minimize the total distance we can also assume that each such edge-node has weight \mathbf{W}_2 . The sum of the distances is then $\mathbf{W}_2 \sum_{k=0,1,\dots} h \cdot s(h)$ but we have to truncate the above sum at some value h_0 that is the distance at which the weight \mathbf{ns} has been reached.

So h_0 is such that:

$$\frac{\mathbf{ns}}{\mathbf{W}_2} \in [b(h_0 - 1), b(h_0)]$$

that is $\frac{\mathbf{ns}}{\mathbf{W}_2} \in [2^{h_0+1}, 2^{h_0+2}]$. So we have:

$$h_0(\mathbf{ns}) \stackrel{\text{def}}{=} \lfloor \log_2 \frac{\mathbf{ns}}{\mathbf{W}_2} \rfloor$$

We can then lower bound the sum of the distances of the connections involving x (with weight $w(x)$) by:

$$2w(x) \times \mathbf{W}_2 \left(\sum_{i \in [0,1,\dots,h_0]} i s(i) + h_0 \cdot \left(\frac{\mathbf{ns}}{\mathbf{W}_2} - b(h_0) \right) \right), \text{ with } h_0 = h_0(\mathbf{ns})$$

Intuitively, due to the exponential nature of the series, almost all the connections are at distance $\sim h_0(\mathbf{ns})$. More precisely if we count the average distance as being $h_0(\mathbf{ns})$ we must subtract one for all the nodes at distance $\leq h_0(\mathbf{ns}) - 1$, and one again for the nodes at distance $< h_0(\mathbf{ns}) - 2$ and so on. The proportion of nodes at distance $\leq h_0(\mathbf{ns}) - i$ is less than $\frac{1}{2^i}$, so we overestimate our estimation by at most $\mathbf{ns} \sum \frac{1}{2^i} \leq n_{ns}$. So the average distance is at least $h_0(\mathbf{ns}) - 1$ and so the sum of the distances is at least :

$$\mathbf{ns}^2 (h_0(\mathbf{ns}) - 1) = \mathbf{ns}^2 (\lfloor \log_2 \frac{\mathbf{ns}}{\mathbf{W}_2} \rfloor - 1)$$

Since this load must be affected to the $3k$ edges of the skeleton, we conclude using the distance bound that :

$$\pi(G) \geq \frac{\mathbf{ns}^2}{3k} (\lfloor \log_2 \frac{\mathbf{ns}}{\mathbf{W}_2} \rfloor - 1) \quad (\text{LB1})$$

Now, recall that \mathbf{W}_2 is upper bounded since :

$$\pi(G) \geq \mathbf{W}_2(n - \mathbf{W}_2) \rightarrow \pi(G) \geq \frac{n\mathbf{W}_2}{2} \quad (\text{LB2})$$

So the bound (LB1) decreases when \mathbf{W}_2 grows and (LB2) decreases from $n^2/4$ when \mathbf{W}_2 decreases, so we get our actual bound by increasing \mathbf{W}_2 till (LB1) and (LB2) meet, which happens roughly when $\mathbf{W}_2 \sim \frac{2\mathbf{ns}}{3k} \log_2(\frac{3k}{2})$. Actually we need to solve :

$$\frac{\mathbf{ns}}{2} \mathbf{W}_2 = \frac{\mathbf{ns}^2}{3k} (\lfloor \log_2 \frac{\mathbf{ns}}{\mathbf{W}_2} \rfloor - 1)$$

So we take as saddle value:

$$W_0 = \frac{2\mathbf{ns}}{3k} \log(\frac{3k}{2})$$

Then either $\mathbf{W}_2 \geq W_0$ and $\pi \geq \frac{\mathbf{ns}^2}{3k} \log(\frac{3k}{2})$ (from LB2) or $\mathbf{W}_2 \leq W_0$ then $\frac{\mathbf{ns}}{\mathbf{W}_2} \geq \frac{3k/2}{\log(3k/2)}$ and from LB1 we conclude also that $\pi(G) \geq \frac{\mathbf{ns}^2}{3k} (\log_2(3k/2) - O(\log \log(k)))$. ■

Remark 4.1 As we will see in Proposition 4.3 , this bound is actually quite tight, especially when $k \ll n$, and the forwarding index is not actually given by the maximum load of an edge-cut but is due to the average distance, indeed if we put $2n/3k$ nodes on each subdivided edge, the edge cuts provide only lower bound of n^2/k .

Remark 4.2 One can actually replace \mathbf{ns} by n in Proposition 4.1, for that we need to compute the total average distance, that is we only need to include in our counting the distances between the following couples of nodes : (skeleton, edge – node), (edge – node, skeleton), (skeleton, skeleton) we can use then 3 Moore bounds with the same behavior (all based on the same geometric serie) and conclude that the total average distance is still $\max(\log_2(n/\text{wtwo}), \log n)$ which allows to prove that for $k \leq n/2$

$$\pi^*(n, k) \geq \frac{n^2}{3k} (\log k - O(\log \log(k)))$$

and this latest bound is still tight when $k = \frac{n}{2}$.

4.3 Construction of minimally-congested graph with degree $\leq \Delta$

Our construction simply reverts the previous operation and builds graphs with few extra edges from good skeletons.

Definition 4.2 Given a graph, we construct $\text{Sub}(G, \mathbf{W})$ as follows: we subdivide each edge e by adding one node x_e and we then attach a binary tree with weight \mathbf{W} on x_e .

Lemma 4.2 Let G be a Δ -regular graph with x vertices, and let $H = \text{sub}(G, \mathbf{W})$ then $\pi(H) \leq \text{Max} \{ \pi(G)(\frac{\Delta \mathbf{W}}{2} + 1)^2 + \mathbf{W}(\frac{\Delta \mathbf{W}}{2} + 1)x, \mathbf{W}((\frac{\Delta \mathbf{W}}{2} + 1)x - \mathbf{W}) \}$

Proof. Indeed we shall route in $\text{sub}(G, \mathbf{W})$ almost like in G , but one half of the \mathbf{W} nodes attached on $x_e, e = (a, b)$ will first reach a and the other half will reach b . Since vertices of G have degree Δ they receive weight $\frac{\Delta \mathbf{W}}{2}$ (plus their own), so we route on G with weights scaled up by a factor $\frac{\Delta \mathbf{W}}{2} + 1$ this induces a load of $(\frac{\Delta \mathbf{W}}{2} + 1)^2 \pi(G)$. We then need to count the load induced by the part of the paths that reach the nodes of type x_e . Since the total weight

in the graph is $(\frac{\Delta W}{2} + 1)x$ this is less than $\frac{2W(\frac{\Delta W}{2} + 1)x}{2} = W(\frac{\Delta W}{2} + 1)x$ and so the total is $(\frac{\Delta W}{2} + 1)^2 \pi(G) + W(\frac{\Delta W}{2} + 1)x$. ■

To our surprise, we could not find the following result in the literature, moreover in the recent survey [Xu2012] the best bounds for cubic graphs were provided by shuffle exchange graphs, and more generally, for bounded degree graphs the best bounds known are derived using de Bruijn graphs. Those bounds are rather good since they differ from the lower bound only by a relatively small (always lesser than 2) constant factor. But indeed random regular graph are asymptotically optimal.

Proposition 4.2 *There exist cubic regular graphs such that $\pi(G) = \frac{2}{3}n \log_2(n)(1 + o(1))$, and Δ -regular graphs with $\pi(G) = \frac{2}{\Delta}n \log_{\Delta-1}(n)(1 + o(1))$.*

Proof. Our proof is based on the existence of an almost balanced routing using “paths” (more properly we should say walks) of length $\log_2(n)(1 + o(1))$ between all the pairs of nodes. Indeed we look at *non backwarding walks* originating from a node, those are sequence of edges which do not contains a cycle of length 2 (doing $a \rightarrow b$ and then $b \rightarrow a$) but we may repeat nodes or even edges. We find exactly $3 \cdot 2^{l-1}$ non backwarding paths of length l originating from a vertex in any cubic graph, it is those paths that are counted in the classical Moore bound. If we define $a_l(u, v)$ as the number of non backwarding path of length l from u to v we know from matrix theory (rapid mixing and Perron Frobenius theorem) that $a_l(u, v)/3 \cdot 2^{l-1}$ converges to $\frac{1}{|V|}$ (i.e. $a_l(u, v) \sim \frac{3 \cdot 2^{l-1}}{|V|=n}$, so for l large enough we may expect to find almost the same number of non backwarding walks of length l connecting any vertex to any other one.

To prove it and to upper bound the value of l for which the mixing happens, we use the classical argument given to prove that the diameter of a cubic random regular graph is less than $\log_2(n) + \Theta(\log \log n)$. We look at two spheres of radius $r = \frac{\log_2 n}{2} + a \log \log n$ (where a is a positive real), one centered at vertex u and the other at vertex v . They have size $\log^a n \sqrt{n}$ so we find on average $\sim 2 \log^a n \sqrt{n} \cdot \frac{\log^a n}{\sqrt{n}} = 2 \log^{2a} n$ ($\mu = 2 \log^{2a} n$) connections between those 2 balls, and since these connections behave like Bernoulli trials we can use the Chernoff bound to claim that with high probability the actual number of connections is close to the expected one. What we want to claim is that with high probability we find almost $\mu = 2 \log^{2a} n$ walks with length $2r + 1$ between u and v , and in order to conclude that this good event happens simultaneously for any of n^2 couple of vertices we will pick this probability to be higher than $1 - \frac{1}{n^2}$ (i.e. the probability of a bad event will then be less than $n^2 \cdot \frac{1}{n^2} < 1$).

The Chernoff bound tells us that for a sum $X(\omega)$ of independent Bernoulli trials with mean μ :

$$\text{Prob}[X(\omega) \geq (1 - \delta)\mu] \leq \exp(-\delta^2 \mu / 2)$$

to defeat the n^2 couples we need $\exp(-\delta^2 \mu / 2) < \frac{1}{n^2}$, and so $\delta^2 \mu / 2 \geq 2 \log_2 n$ and so $\delta = \sqrt{\frac{4 \log_2 n}{\mu}}$, so we only need $\mu \gg \log_2(n)$ (the expected number of connections) so we can pick $a = 1$ for which $\mu = 2 \log^2(n)$ and we conclude that : with positive probability, for any couple of vertices the number of paths with length $\log_2 n + 2 \log \log n$ is equal to the average number of paths (up to a factor $1 - \delta_0 = 1 - \sqrt{\frac{2}{\log_2 n}}$).

This proves that we find almost the same amount of non backwarding walks of length $l_0 = \log(n) + 2 \log \log n + 1$ between any pair of nodes (their number is not actually very important). Now we simply route on all the paths of length l_0 , and we do it being pessimistic and assuming that there are only $(1 - \delta_0)3 \cdot 2^{l_0-1} \cdot \frac{1}{n}$ such paths between each pair of nodes. So we send $\frac{1}{(1-\delta_0)} \frac{n}{3 \cdot 2^{l_0-1}}$ unit of flow on each such path. By construction the routing is balanced on the edges. The total load of the flow coming from a given source is then $\frac{1}{(1-\delta_0)} \frac{n}{3 \cdot 2^{l_0-1}} \times 3 \cdot 2^{l_0-1} \times l_0 =$

$\frac{1}{1-\delta_0} \times nl_0$, summing on all the sources we get $\frac{1}{1-\delta_0} \times n^2 l_0$, since $l_0 = (1 + o(1)) \log(n)$, and $\frac{1}{(1-\delta_0)} = (1 + o(1))$ we get a total load of $L = (1 + o(1))n^2 \log_2 n$ and since this load is balanced on the edges we get a forwarding index of $L/|E| = \frac{2}{3}(1 + o(1))n \log_2 n$.

The proof for a general Δ is the same. \blacksquare

Remark 4.3 *Note that the fair shortest path routing (in which each shortest path carries the same flow) is probably better and for small values of n it may even be significantly better, but we don't have currently a good method to evaluate its load and proving that so doing we get a better load. Probably the forwarding index of random cubic graph is $\frac{2}{3}n \log_2 n + \Theta(n)$, but we proved only a weaker result. Moreover the value of n for which our $(1 + o(1))$ becomes smaller than the $\frac{3}{2}$ are relatively high (order of 1000).*

Proposition 4.3 *There exist $(n, e = n + k)$ cubic graphs such that $\pi(G) \leq \frac{n^2}{3k} \log_2(k)(1 + o(1))$.*

Proof. We apply Lemma 4.2, where we pick for G a cubic graph with $v = 2k$ vertices and $e = 3k = v + k$ edges and minimum $\pi(G)$. Since we want to build a graph with n vertices, we pick \mathbf{W} such that $(\frac{\Delta \mathbf{W}}{2} + 1)2k = n$, i.e., $\mathbf{W} = \frac{n}{\Delta k} - \frac{2}{\Delta}$. Then Lemma 4.2 claims that $\pi(H) \leq \pi(G) \left(\frac{n}{2k}\right)^2 + \frac{n^2}{\Delta k}$. Since the best we can currently choose for G is such that $\pi(G) \sim \frac{2}{3}2k \log_2(2k)$ (random cubic graphs, from proposition 4.2), the first term always dominates (for n large enough) and we have: $\pi(G) \leq \frac{n^2}{3k} \log_2(k)(1 + o(1))$. \blacksquare

5 Edge forwarding index of cubic ($\Delta = 3$) graphs with few extra edges: $e = n + k$

When k is large, we provided in Section 4 asymptotically matching upper and lower bounds on the minimum congestion. This implies that $\pi^*(n, n + k)$ behaves like $\Theta(\frac{n^2}{k} \log \frac{n}{k})$ when both k and n are large. So, in order to get a complete picture of the situation, we still need to understand the case of $(n, n + k)$ graphs when k is fixed. In this section, we answer this question, that is we solve the MIN-CONGESTION DESIGN PROBLEM, for graphs with arbitrary n , but small values of k .

5.1 The skeleton approach, complexity

From the results of Section 4, we know that $(n, n + k)$ graphs are constructed from a cubic skeleton on which are attached trees with size u . So, when k is small, we may enumerate all the possible skeletons (like we enumerated all the cubic graphs) and determine for each the best way to attach trees. Attaching trees means determining for each edge $e \in E$ the size $\alpha(e)$ of the tree that we attach in the edge. Hence, we want to find the best weight repartition $\alpha : E \rightarrow \mathbb{N}$ that satisfies $\sum_{e \in E} \alpha(e) = n$ and $\forall e \in E, \alpha(e) \leq w_{max}$, where by best we mean with the smallest forwarding index. So, finding the best way to subdivide edges means solving a problem of the following flavor:

Definition 5.1 (Best Mass Repartition) *Given a graph G and a maximum weight w_0 find a weight function $w : V \rightarrow \mathbb{R}^+$ with $\forall v \in V, w(V) = 1, w(v) \leq w_0$ such that $\pi(G, w)$ is minimum.*

Note first that, since $\pi(G, \lambda w) = \lambda^2 \pi(G, w)$, using weight α for vertex $v \in V$ actually means attaching a tree with size αn at v ; secondly on the skeleton we don't put weight directly on the vertices of the skeleton, instead we introduce at most 3 vertices on each edge to represent the

original graph structure. And it is on those vertices that we perform a weight repartition. So, we solve a slightly more general problem, since the function w must be zero on a subset of V (the nodes of the original graph).

Due to LP-duality, and since $\pi(G, w)$ can be expressed as the maximum distance between node pairs for a suitable dual metric l , finding a solution with cost $\leq C$ for the Best Mass Repartition problem can also be expressed as determining a weight function that fulfills all the dual constraints:

Problem 5.1 Find $w : V \rightarrow \mathbb{R}^+, w(V) = 1, w(v) \leq w_0$ such that for any edge metric $l : E \rightarrow \mathbb{R}^+, l(E) = 1 \sum_{(u,v) \in V^2} w(u)w(v)d_l(u, v) \leq C$.

So, at the core of the problem, we find the following average distance minimization:

Definition 5.2 (Average Distance Minimization) Given a graph G and a maximum weight w_0 find a weight function $w : V \rightarrow \mathbb{R}^+$ with $w(V) = 1, \forall v \in V, w(v) \leq w_0$ that minimizes the average (equiv. total) distance $\sum_{(u,v) \in V^2} d(u, v)w(u)w(v)$

This problem is about minimizing a quadratic form under linear constraints. If D is the distance matrix associated to the graph, we can express it compactly as:

$$\begin{aligned} \text{Minimize: } & w^\top D w \\ & \mathbf{0} \leq w \leq w_0 \times \mathbf{1} \\ & w \cdot \mathbf{1} = 1 \end{aligned} \tag{obj}$$

Such problems are usually concave and NP-Hard (see as example [PaVa91]), but they can be solved efficiently when they are convex which happens when D is positive semi-definite. In our particular case, the problem is still NP-Hard.

Proposition 5.1 The Average Distance Minimization problem is NP-hard.

Proof. We reduce the problem to max Clique. We consider a graph G set $w_0 = \frac{1}{k}$, let Opt be the minimum of obj, we remark that $Opt = \frac{1}{k^2}k(k-1) = 1 - \frac{1}{k}$ if and only if G contains a clique with size k . since it is the only case for which we can find k nodes that are pairwise adjacent. So if we can solve the average distance minimization problem we can find the maximum clique of G . ■

The above result suggests that finding Best Mass Repartition for a skeleton is NP-Hard. We shall see in the next section that the natural intuition that indicates that the uniform repartition should be optimal for vertex-transitive graphs is also wrong. Indeed, uniform solutions are always optimal for convex problems on symmetric structures, but mass repartition is somewhat concave.

We hope to have provided clues indicating that, even given a skeleton, there is little hope to algorithmically determine the best weight function and let open the following question :

Open Problem 5.1 Is the Best Mass Repartition problem NP-Hard ? If so, is it APX-hard ? Can the Min Average Distance problem be approximated with a polynomial algorithm?

5.2 Optimal $(n, n - 1 + k)$ cubic graph for $k = 0, 1, 2, 3$

5.2.1 $(n, n - 1)$ graphs: Trees

When $k = 0$ and $e = n - 1$, the network is a tree with max degree $\Delta = 3$. The case of degree Δ trees is trivial since for such trees, considering the most balanced cut, we get: $\pi(T) \geq$

$2\Delta(\Delta - 1) \left(\frac{n}{\Delta}\right)^2$ and this value is attained using a balanced Δ -ary tree or a subdivided Δ -star with branches with equal size $\frac{n}{\Delta}$. So, for $\Delta = 3$, we have:

$$\pi^*(n, n-1) = 2\Delta(\Delta - 1) \left(\frac{n}{\Delta}\right)^2 = 2\frac{(\Delta - 1)}{\Delta}n^2 = \frac{4}{3}n^2.$$

5.2.2 (n, n) graphs: Trees + $(k = 1)$ extra edge

In this case, the first intuition is that the cycle C_n should be the optimal structure. Recall that $\pi(C_n) = \frac{n^2}{4}$ when n is even, and $\pi(C_n) = \frac{n-1}{2} \frac{n+1}{2}$ when n is odd (indeed $\pi(C_n) = \lceil \frac{n-1}{2} \rceil \lfloor \frac{n+1}{2} \rfloor$). The cycle is the only 2 connected structure but it is not the min-congested one since some graphs with bridges do have lesser congestion.

Proposition 5.2

$$\pi^*(n, n) = \frac{12}{49}n^2$$

Proof. First we prove the lower bound and then we provide a construction.

We consider an (n, n) -graph G with optimal forwarding-index Opt . Since G has only a single extra edge there is only one possible skeleton: a cycle C . Let us denote V_C the vertex set of the cycle. Each vertex $u \in V_C$ receives a proportion of the total weight $\stackrel{\text{def}}{=} w_c(u)n$ and so $\sum_{u \in V_C} w_u = 1$. Let w_0n be the maximum weight used on a node of the cycle, then we have two following bounds on Opt :


$$Opt \geq 2w_0(1 - w_0)n^2 \tag{LB1}$$

$$Opt \geq \frac{1}{2} \left(\frac{1 - w_0}{2} \right) \left(\frac{1 + w_0}{2} \right) n^2 \tag{LB2}$$

The bound LB1 immediately follows by considering the bridge attaching the tree with maximum weight to the rest of the graph, the load of that bridge is then $2w_0n(n - w_0n)$. For LB2, we consider the 2-cut on the cycle structure that is the most balanced. When we remove or add a node u to one side of the cut, we increase or decrease the size of the cut-sides by $+/-w_u$. So, if $\forall u \in V_C w_u \leq w_0$, we are certain to find a cut such that the size of side lies within $[\frac{n}{2} - w_0n/2, \frac{n}{2} + w_0n/2]$. Then, we find at least $A = 2(n/2 - w_0n/2)(n/2 + w_0n)$ connexions crossing two links and we get LB2 : $Opt \geq \frac{1}{2}(n/2 - w_0n/2)(n/2 + w_0n)$. Solving $LB1 = LB2$ we find $w_0 = \frac{1}{7}$ and $LB1 = LB2 = \frac{12}{49}n^2$ and this implies that $\max(LB1, LB2) \geq \frac{12}{49}n^2$. Now we provide a graph G_7 with $\pi(G) = Opt = \frac{12}{49}n^2$: we simply take the cycle C_7 and on each vertex we attach a tree (any tree is suitable) with $\frac{n}{7}$ nodes, see Figure 3. The routing on G_7 is trivial since, on the C_7 skeleton, we route like on the cycle C_7 and, inside the trees, the routing is uniquely determined. From construction, the trees have size $\frac{1}{7}n^2$. So, the load inside the trees is less than Opt . On the cycle skeleton the load is $(\frac{n}{7})^2\pi(C_7) = \frac{12}{49}n^2$.

Note that, to get an almost optimal construction, we simply need a cycle with length $o(n)$ on which we attach 7 trees with size $(1 + o(1))\frac{n}{7}$. ■

Remark 5.1 Note that there is a price for 2-connectivity or Hamiltonicity, since we get $\pi^*(n, n) = \frac{12}{49}n^2$ if we allow bridges and $\pi^*(n, n) = \frac{1}{4}n^2 = (1 + \frac{1}{48})\frac{12}{49}n^2$ if we restrict the solutions to be 2-connected or Hamiltonian structures.

Figure 3: Generic Optimal Solution for n edges.**5.2.3 $(n, n + 1)$ graphs: Trees + $(k = 2)$ extra edges****Proposition 5.3**

$$\pi^*(n, n + 1) = \frac{2}{9}n^2.$$

Proof. First, we prove the lower bound. Again, we consider the potential skeletons:


- a) 2 cycles connected by a bridge, formally the skeleton is indeed a pair of double edge connected by a bridge.
- b) A cycle with a chord, more exactly the skeleton is a set of 3 parallel edges e_1, e_2, e_3 .

Case (a) If we consider the lightest side of the bridge we can route all the internal demands and reach the border of the set with a load lesser than the one on the bridge edge. This is due to the fact that the load of a bridge separating into the graph into $[S, \bar{S}]$ is $2|S||\bar{S}|$, in other words when a tree contains less than half of the nodes the maximum load is always reached on the edge that attach it to the rest of the graph. So in case (b) we can remove the extra edge inside the light side of the bridge and so we cannot beat the bound for (n, n) graphs.

Case (b) For $e \in \{e_1, e_2, e_3\}$, let $w(e)$ be the total weight attached on the arc e , w.l.o.g assume than e_1 has the maximum weight, then $w(e_1) \geq \frac{1}{3}$ and since $w(e_1) \leq \frac{1}{2}$ the load of the 2-cut that separate e_1 from e_2, e_3 is at least $2\frac{1}{3}n \times \frac{2}{3}n = 2\frac{2}{9}n^2$; and since it is covered by only two edges we conclude that $\pi(G) \geq \frac{2}{9}n^2$.

A possible construction (see Figure 4) is then to use $\forall i \in \{1, 2, 3\}$ a path P_i of length $n/3$ for e_i , then one can cover all the request using 3 cycles of size $\frac{2n}{3}$ ($P_i \cup P_j, i \neq j$). An edge appears then into 2 cycles so the load is at most $2\pi(C_{2n/3}) = 2\frac{(2n/3)^2}{4} = \frac{2n^2}{9}$. Note that in this analysis we route twice the demand that are internal to each path P_i . Alternatively we can also connect 2 degree three vertices with 3 path of length 4 and attach one tree of size $\frac{n}{6}$ on each of the 6 internal nodes with degree 2.

■

Figure 4: Optimal graphs with $n + 1$ edges

5.2.4 $(n, n + 2)$ graphs: Trees + $(k = 3)$ extra edges. Non-uniform subdivision of the K_4

The next result result is rather surprising since intuitively a uniform (or at least symmetric) subdivision of the K_4 should provide an optimal solution. But a phenomena similar to the one we already met in the case $k = 1$ (the C_7) happens again in a slightly more complex way.

Proposition 5.4

$$\pi^*(n, n + 2) = \frac{20}{11^2} n^2$$

Proof. We first prove the lower bound in lemma 5.1, and then we provide a matching upper bound (construction) in lemma 5.2. ■

We first prove the lower bound :

Lemma 5.1

$$\pi^*(n, n + 2) \geq \frac{20}{11^2} n^2$$

Proof.


We look at the possible skeletons, recall that the maximum degree is 3 and that when forming the skeleton (strictly) we remove nodes with degree 2, 3 it mean that we have either a cycle + 2 chords which means either (a) a subdivision of K_4 or (b) 2 double edges connected by a double bridge or 2 cycles connected by two edges which indeed is b again.

- **Case b:** In this case the total load of the double bridge is $2 \frac{n}{2} \frac{n}{2}$ for 2 edge which implies that $\pi(G) \geq \frac{n^2}{4} > \frac{20}{11^2} n^2$.
- **Case a:** First remark that the maximum weight of an attached tree is at most $\frac{1}{11}n$, otherwise $\pi(G) \geq 2 \times \frac{1}{11}n \times \frac{10}{11}n/1 = \frac{20}{11^2} n^2$ and we have proved the lower bound.

Now, let us call a an $[a, b]$ cut a cut such that the two sides have total weight in $[a, b]$, we want to show that there exists an $[\frac{5}{11}n, \frac{6}{11}n]$ cut with border 3, this will imply that

$$\pi(G) \geq 2 \frac{\frac{5}{11} \times \frac{6}{11}}{3} = \frac{20}{11^2} n^2$$

For that case we consider for each node $v_i, i \in [3]$ of the skeleton the cut formed by v_i its three adjacent edges with the weight they carry. Each of those 4 cuts has a border of size 3. Since each edge appears in exactly two cut, when summing the weight of those cut we

Figure 5: Optimal K_4 subdivision with $n + 3$ edges.

count twice the total weight. So at least one such cut has weight at least $\frac{n}{2}$. Now, while the cut weight is larger than $\frac{6n}{11}$ we decrease the weight of the cut by removing step by step the last tree attached on one of the three subdivided edges. The maximum weight of a tree is $\frac{n}{11}$, so at each step, we decrease the weight by an amount lesser than $\frac{1}{11}n$. It follows that must at some point obtain an $[\frac{5}{11}n, \frac{6}{11}n]$ cut.

■

To complete the proof, we only need to provide a construction. The lower bound argument almost enforces the following construction.

Definition 5.3 ($K_{4,sub}(n)$) *The graph $K_{4,sub}(n)$ is obtained by subdividing 5 edges of K_4 twice and one edge once, thus we add 11 new nodes. Then, we attach a tree with weight $\frac{n}{11}$ on each new node, see Figure 5.*

Lemma 5.2

$$\pi^*(n, n + 2) \leq \pi(K_{4,sub}(n)) = \frac{20}{11^2}n^2$$

Proof. The claim can be checked by computing the forwarding index of a weighted graph with 15 nodes, 4 with weight 0 and 11 with weight 1. This finite graph is obtained by replacing each attached trees of $K_{4,sub}(n)$ by a vertex with weight $\frac{n}{11}$. Last, we know that we can set this weight to 1 and multiply the result by $(\frac{n}{11})^2$. To conclude, there exists a routing achieving this bound.

■

To complete this case we briefly study the uniform subdivision of the K_4 which intuitively should have been optimal. Equivalently we know that we may as well determine the forwarding-index of the graph $K_{4,uni}$ that we define as :

Definition 5.4 ($K_{4,uni}(n)$) *The graph $K_{4,uni}(n)$ is obtained by subdividing the 6 edges of K_4 twice, thus we add 12 new nodes. Then we attach a tree with weight $\frac{n}{12}$ on each new node (equiv. we set the weight to $\frac{n}{12}$ for new nodes and to 0 for the nodes of the original K_4).*

Proposition 5.5

$$\pi(K_{4,uni}(n)) = \frac{n^2}{6}$$




Figure 6: The broadcast tree from a node (used in the proof of Theorem 5.5).

Proof. The lower bound comes from the $[n/2, n/2]$ cut with border 3 that we prevented to appear in the non-uniform case (i.e. take a vertex of the original K_4 plus all the internal nodes of the subdivided paths adjacent to it), this cut implies that $\pi \geq 2 \times \frac{n}{2} \times \frac{n}{2} / 3 = \frac{n^2}{6}$.

The graph is transitive on the vertices with weight $\frac{n}{12}$ (we can ignore the other ones) but it is not edge-transitive, indeed there are exactly two orbits for the edges, 12 edges of type (A) connect new nodes to original nodes, and 6 edges of type (B) connect new-nodes (they appear in what was the middle of one of the initial K_4 edges). We use the tree of figure 6, and rotate and translate it and get something uniform on each class. For this tree the load of the (B) edges is 12 since we will have 12 trees the total load on (A) edges will be 12^2 and since it will be balanced on those 6 edges, each will receive load $\frac{12^2}{6} = 24$. Similarly, for the (A) edges the load for one tree is 24, so we get total load 2×12^2 on 12 edges. So the load of the routing is 24. Now each unit of flow actually represent $\left(\frac{n}{12}\right)^2$ connections, we scale up by that factor and $24 \times \left(\frac{n}{12}\right)^2 = \frac{n^2}{6}$. ■

6 Graphs with a small number of vertices ($\Delta = 3$)

We have seen in Sections 4 and 5 the importance of having good skeletons to build graphs with low forwarding indices. In Table 1 on page 26, we present graphs with a small number of vertices which have the minimum possible forwarding indices. These graphs can serve as skeletons to build families of graphs with an arbitrary number of vertices. In some cases, optimality is easy to prove using:

- the Moore bound. In a cubic graph, and for a given vertex, the number of vertices that are at distance $0, 1, 2, 3, \dots$, are respectively, at most $1, 3, 6, 12, \dots$. When those bounds are reached for all the vertices of a cubic graph, the latter minimizes $\mathcal{L} = 2|V| \frac{\overline{D(G)}}{d(G)}$ among all the graphs with the same size and with degree 3. When the graph is optimal for the Moore bound and is edge-transitive, its forwarding index is minimum. This is the case for $n = 6, 14$;
- cut arguments, for $n = 4, 8, 10$.

In other cases, ($n = 12, 16, 18$), the generic arguments fail to provide matching upper and lower bounds. We had to check all the possible cubic graphs ([**cubics**]).

6.0.1 Consequences for unbounded n but a few edges

All those graphs can be used as skeletons, as example if one wishes to get a good $(n, 6)$ graph with $e = n + 6$ edges one can simply pick the Petersen graph as skeleton and apply lemma 4.2. We use the uniform weight function $\mathbf{W} = \frac{n}{15}$ and using the generic routing of the lemma we get : $\pi(n, 5) \leq \pi(G) \leq 10 \left(\frac{n}{10}\right)^2 + 2 \frac{n}{30} \times \frac{14n}{15} = \frac{n^2}{10} + \frac{14n^2}{225}$. This may be potentially improved by computing the exact forwarding index of the so defined weighted graph (that has only 15 vertices).

Solving the best mass repartition problem would allow us to go quite further, but currently we have no clue about what is the best repartition even for a small structure. It is certainly possible to repeat what we did for 0, 1, 2, 3 extra edges, but the difficulty shall increase considerably each time we add one edge, finding a method that would scale more than considering cases by “hand” is certainly interesting.

7 Conclusion

In this paper, we provided a basic understanding of the interplay between the forwarding-index of a graph and its number of extra-edges. Our bounds are mostly asymptotically tight and explain as example how the transition happens between highly congested graphs (Trees, Paths, ...) to cubic regular graphs which have much lower congestion.

Some results, like the step-like behavior in Proposition 3.4 or irregular optimal structures, are also *fun*, since they are unexpected. Last, we believe that our work opens many questions:

- **Small cases:** In the case of a few extra-edges, we stopped at 3 extra edges (and even in those cases the proofs are not immediate). So, it may be interesting to go further and to understand if optimal graphs with k extra-edges are built using an optimal cubic graph with $\frac{k}{2}$ vertices (we determined such graphs till $k = 22$). As example: is the family of optimal graphs with 5 extra edges built using the Petersen and subdividing it properly? And, if so, how do we find the best subdivision (we saw the the uniform subdivision is not always optimal).
- **Construction from skeletons:** Given a skeleton, we do not know how to affect weights in order to minimize the forwarding-index of the resulting graph. That problem can be expressed as a quadratic non convex problem and we conjecture that it is NP-Complete.










$n = 4, \pi = \mathcal{L} = 2$ \mathcal{L} given by a cut	$n = 6, \pi = \mathcal{L} = 4.66\dots$ \mathcal{L} given by the Moore bound	$n = 8, \pi = \mathcal{L} = 8$ \mathcal{L} given by a cut
		
$n = 10, \pi = \mathcal{L} = 10$ \mathcal{L} given by a cut	$n = 12, \pi = \mathcal{L} = 14.26\dots$ \mathcal{L} found by brute force	$n = 14, \pi = \mathcal{L} = 18$ \mathcal{L} given by the Moore bound
		
$n = 16, \pi = \mathcal{L} = 22$ \mathcal{L} found by brute force	$n = 18, \pi = \mathcal{L} = 26.66\dots$ \mathcal{L} found by brute force	$n = 20, \pi = 30.84\dots, \mathcal{L} = 30$
		

Table 1: Small cubic graphs with minimum edge forwarding index



**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-0803