

The steady-states of splitter networks

Basile Couëtoux¹, Bastien Gastaldi², Guylain Naves³

Abstract: We introduce splitter networks, which abstract the behavior of conveyor belts found in the video game Factorio. Based on this definition, we show how to compute the steady-state of a splitter network. Then, leveraging insights from the players community, we provide multiple designs of splitter networks capable of load-balancing among several conveyor belts, and prove that any load-balancing network on n belts must have $\Omega(n \log n)$ nodes. Incidentally, we establish connections between splitter networks and various concepts including flow algorithms, flows with equality constraints, Markov chains and the Knuth-Yao theorem about sampling over rational distributions using a fair coin.

1 Introduction

The transportation of materials or data within various networks represents an inexhaustible source of mathematical problems, which has led to almost as many solutions, theories and algorithms. These advancements have brought about significant improvements across diverse fields including supply chain management, logistics, network optimization. Transportation also serves as a central component in numerous games, as evidenced by the transportation category on BoardGameGeek which lists almost two thousand games [4]. In Factorio [24], a video game published in 2020 by Wube Software, players must mine natural resources to feed a rocket-building factory on an hostile planet. A major part of the gameplay involves the movement of resources within the factory, employing various mechanism: robotic arms, conveyor belts, drones or trains.

In this work, we study the conveyor belts of Factorio. An item placed on a belt will move at a constant speed toward the end of the belt, until it reaches that end, or is blocked by an item preceding it. Belts in Factorio can be combined using a splitter, connecting one or two incoming belts to one or two outgoing belts. A splitter takes items from the incoming belts and places them on the outgoing belts, trying to split the flow as fairly as possible between the incident belts, while maximizing the throughput. Given the scale of a typical Factorio game, players frequently encounter the need to balance the loads across multiple belts, and the community has devised numerous efficient networks to address this load-balancing problem.

An intriguing aspect of Factorio is its encouragement for players to construct vast systems of automation, requiring intensive planning and optimization. Ultimately, the limiting factor arises from the CPU load generated by game state updates. Consequently, players are incentivized to prioritize resource efficiency, particularly concerning gameplay elements that entail frequent computations such as splitters. This motivates the minimization of the number of splitters in load-balancing networks.

Our goal is two-fold: first we model the steady-state of a network of splitters. The network of conveyor belts is abstracted as a directed graph, with nodes corresponding to splitters and arcs to belts. A steady-state is a throughput function on the arcs; a circulation with additional constraints to capture the fact that splitters are fair and locally optimizing. We present two polynomial-time algorithms for computing a steady-state in a splitter network. An analogy is made with two classical maximum-flow algorithms: the blocking-flow algorithm [8] and the push-relabel algorithm [10]. In contrast to maximum flows, the primary challenge arises when a belt reaches full capacity, as its supplying splitter may no longer stay both fair and maximizing. In that case, the splitter is allowed to become unfair, but that decision changes the constraints applied to the flow, making the problem fundamentally non-convex. In a second part, we showcase various load-balancing network designs sourced from the Internet, formalizing concepts defined by the players community. Furthermore, we prove that those designs approach optimality. Specifically, we

¹Aix-Marseille Université, CNRS, LIS, Marseille, France, basile.couetoux@univ-amu.fr

²Télécom SudParis, Institut Polytechnique de Paris, Evry, France, bastien.gastaldi@telecom-sudparis.eu

³Aix-Marseille Université, CNRS, LIS, Marseille, France, guylain.naves@univ-amu.fr, corresponding author.




Figure 1: 3 splitter networks given with capacities and associated steady-states. Splitter will be represented by circle vertices, terminals by square vertices. Each terminal is tagged by its capacity, and each arc by its throughput. Saturated arcs are bolder than fluid arcs.

prove that any balancing network on n belts must have $\Omega(n \log n)$ splitters, by exhibiting a relation with the problem of sampling the uniform distribution over a set of n elements using only a fair coin. The core design is the Beneš network, a circuit-switching network well-known in the field of telecommunication [1, 2].

The blocking-flow-like algorithm relies on finding circulations with equality constraints. A *circulation* on a directed graph is a flow without any excess at any vertex. Given a directed graph (G, A) , we denote $\delta^+(v)$ and $\delta^-(v)$ the sets of outgoing and incoming arcs incident to a vertex v . Let $\mathcal{C}^=$ be a partition of A such that for each part $C \in \mathcal{C}^=$, there is some vertex v with $C \subseteq \delta^+(v)$. The $\mathcal{C}^=$ -circulation problem is to decide whether there is a non-zero circulation f that is constant within each part. While this problem can easily be solved using linear programming, we require a good characterization of graphs admitting a $\mathcal{C}^=$ -circulation. Additionally a polynomial-time algorithm is needed to either construct a $\mathcal{C}^=$ -circulation or identify an obstacle that prevents its existence. The algorithm relies on the computation of a stationary distribution of an auxiliary graph. In contrast, solving maximum integral flow problems with additional equality constraints is known to be NP-hard [17], even when the partition is exactly the sets of leaving arcs of each vertex [23, 18].

Sorting networks [12] and Beneš networks have topologies similar to splitter networks, with nodes of in-degree and out-degree 2. In microfluidics, mixing graphs are used to produce droplets of specific concentration, using devices that produces two identical droplets from two droplets of any concentration [7]. The concentration values on the arcs are subject to equality constraints similar to those of splitter networks, but without a maximizing constraint. The topology of splitter networks is nonetheless more general than these examples, as splitter networks may have directed cycles, those being necessary in particular to achieve load-balancing with an arbitrary number of outputs.

In an answer to a question on the mathematics section of `stackexchange`, David Ketcheson attempted to model and compute the throughputs of splitter networks [11]. Rather than binary categorizing each belt as full or not, each arc is assigned a density and a velocity. The density will be monotonically increasing, and the velocity monotonically decreasing during the run of the algorithm, until a steady-state is reached. In fact the velocity increases only after the density reaches its maximum at one. Therefore this description is equivalent to our solution, which involves a throughput function and a set of full belts. Unfortunately his algorithm does not always terminate, and its solutions do not satisfy that splitters use their incoming belts fairly. Ketcheson also gave a procedure, albeit non-polynomial, to determine whether a network (not necessarily load-balancing) may limit throughput. However, this procedure is applicable only to networks without directed cycle. In [15], Leue modeled splitter networks using Petri nets, and uses model checking to check the load-balancing properties of some small networks.

The Factorio community is very active and creative. Players have designed load-balancing networks of various sizes, with efficient embeddings into the grid while respecting the constraints of the game. Additionally, they have developed general methods for constructing arbitrary large load-balancing net-

works. They introduced the concept of balancing networks, along with the more robust properties of being throughput unlimited or universal, and subsequently designed networks that exhibit these characteristics. A notable example is the universal balancer presented by *pocarski* [20], although it uses non-fair splitters too; our universal balancer only uses fair splitters. They also discovered the relationship with Beneš networks. Factorio-SAT [21] is a project that uses a SAT-solver to find optimal embeddings of splitter networks in the grid. The project *VeryFactory* uses a SAT-solver to check various load-balancing properties of splitter networks [14]. Factorio belts are actually sufficiently complex to be Turing-complete [16]. There are many implementations of various devices inside Factorio, ranging from raytracers to programming language interpreters, using the diverse set of available gameplay mechanisms. Factorio has been the inspiration for several other academic works [22, 19, 5, 6, 9].

The rest of this section presents an overview of the main concepts and results of this work. Then [Section 2](#) is dedicated to the \mathcal{C}^- -circulation problem, which is a requirement for the next [Section 3](#), that focuses on the two algorithms to compute a steady-state. [Section 4](#) contains the constructions and proofs of the load-balancing designs. [Section 5](#) formalizes the proof for the lower bound on the number of splitters in a load-balancing network. In [Section 6](#), we define networks that simulate arcs of arbitrary rational capacities. In [Section 7](#), we investigate splitter networks when each splitter may receive priorities, therefore one incident arc over the others. Using priorities, we will define in [Section 8](#) a balancer with a number of fair splitters closer to our lower bound. Finally in [Section 9](#) we will present some perspectives.

1.1 Splitter networks and their steady-states

We start by modeling networks of conveyor belts and splitters by directed graphs, where each single belt is an arc, and each splitter is a node (thus abstracting the length of the belts).

Definition 1. A *splitter network* is a directed graph G (with possible loops or parallel arcs) whose vertex set can be partitioned into three sets $V(G) = I \uplus S \uplus O$ where

- (i) I is the set of *inputs*, and $d^+(i) = 1$, $d^-(i) = 0$ for any input i ;
- (ii) O is the set of *outputs*, and $d^-(o) = 1$, $d^+(o) = 0$ for any output o ;
- (iii) S is the set of *splitters*, and $d^-(s) = d^+(s) = 2$ for any splitter s .

We will use the word *flow* to informally describe the material transported by the network, and *throughput* for the amount of flow going through the arcs. Our work aims to understand the throughputs inside a splitter network at steady state, when some maximum throughputs are forced on its inputs and its outputs, which are respectively the sources and sinks of the flow passing through the network. To this end we will consider capacity functions on the input and output. A capacity c on an input means that the input has an incoming flow of throughput c . The input will try to push that much into the network, but no more. A capacity c on an output means that the output will accept a maximum throughput of c . We consider that the maximum throughput of any arc is 1, with all belts being identical.

A splitter can be described using two operational rules. The first rule, which takes precedence, is to maximize the amount of flow that goes through it. The secondary rule is to be fair. A splitter is fair relatively to its outgoing arcs: it tries to push as much flow onto each of them. It is also fair relatively to its incoming arcs: it tries to pull as much flow from each of them. As the maximization rule takes precedence, it will not be fair when being unfair leads to higher throughput. For instance, consider the network in [Figure 1 \(a\)](#), depicting a network with a single splitter. As one of the output has a lower capacity, it pushes more flow toward the other output, thereby maximizing the total throughput, while still being as fair as possible as it minimizes the difference of throughputs on its outgoing arcs.

The throughput of an arc may reach a limit when its head is an output with a low capacity. For example in [Figure 1 \(a\)](#), an output of capacity 0.4 acts as a bottleneck. In other cases the head of an arc is a splitter, which itself is limited by what its outgoing arcs can accept. For example in [Figure 1 \(b\)](#), as all the outputs have reached their capacities, the splitter cannot accept more flow, even if the bottom input could provide even more flow. In terms of conveyor belts, some belts will initially receive more items that




Figure 2: An example of steady-state in a moderately small network, and the reverse network with its steady-state obtained by reversal. Notice that the reversed steady-state satisfies rule **R8** but not rule **R8S**.

they can deliver, causing them to fill up. Once full, they can only accept from upstream as much as they deliver downstream, which may in turn limit throughputs upstream. We say that such belts are *saturated*.

The output capacities are not the only factor that limit the total throughput and create bottlenecks. This can be observed in **Figure 1 (c)**. There, the rightmost splitter tries to be fair and send some of the flow back to the left. The leftmost splitter also tries to be fair, thus accept the flow coming from the right. This results in the stabilization into the given throughputs. This example illustrates that the throughput is not globally maximized, contrary to the expectation of a total throughput of 1 for this network. Instead, it is only 0.5.

The following definition formalizes the notions of capacity, throughput and saturations, as well as the behaviour of splitters related to the flow going through the network in a steady state.

Definition 2. Let $G = (I \uplus S \uplus O, E)$ be a splitter network, and let $c : I \cup O \rightarrow [0, 1]$ be the maximal capacities of each input and output node. A *steady-state* for (G, c) is a pair (t, F) where

- R1 $t : E \rightarrow [0, 1]$ is the *throughput* function;
- R2 $F \subseteq E$ is the set of *fluid arcs*, $E \setminus F$ is the set of *saturated arcs*;
- R3 for each $i \in I$ with $\delta^+(i) = \{e\}$, $t(e) \leq c(i)$ and moreover if $e \in F$ then $t(e) = c(i)$;
- R4 for each $o \in O$ with $\delta^-(o) = \{e\}$, $t(e) \leq c(o)$ and moreover if $e \notin F$ then $t(e) = c(o)$;
- R5 for each $s \in S$, with $\delta^-(s) = \{e_1, e_2\}$ and $\delta^+(s) = \{e_3, e_4\}$, $t(e_1) + t(e_2) = t(e_3) + t(e_4)$;
- R6 for any $e_1, e_2 \in E$ with $\{e_1, e_2\} = \delta^-(s)$ and $e_1 \notin F$, $t(e_1) \geq t(e_2)$;
- R7 for any $e_1, e_2 \in E$ with $\{e_1, e_2\} = \delta^+(s)$ and $e_1 \in F$, $t(e_1) \geq t(e_2)$;
- R8 for any $uv \in E \setminus F$ and $vw \in F$, $t(uv) = 1$ or $t(vw) = 1$.

Rules **R3** and **R4** say that the throughputs are limited at each input and each output, and moreover, an input pushes as much flow as allowed by its capacity on a fluid arc. Similarly an output absorbs as much flow as allowed by its capacity from a saturated arc. Rule **R5** imposes the conservation of flow. Rules **R6** and **R7** enforce the fairness constraints: a splitter consumes no less flow from a saturated arc than from another incoming arc. A saturated arc represents a belt that is full. Therefore, the splitter is not limited in how much flow it can pull from that arc, and thus cannot pull less than from the other incoming arc. Similarly it produces no less flow in a fluid outgoing arc than in another outgoing arc. In particular, if both incoming arcs are saturated, or if both outgoing arcs are fluid, they must have equal throughput, suggesting the following definition.

Definition 3. Given a splitter network $G = (I \uplus S \uplus O, E)$, and a set $F \subseteq E$ of fluid arcs, we say that two arcs $e, e' \in E$ are

- ▷ *in-coupled* if $e, e' \notin F$ and there is a splitter vertex $v \in S$ with $\delta^-(v) = \{e, e'\}$,
- ▷ *out-coupled* if $e, e' \in F$ and there is a splitter vertex $v \in S$ with $\delta^+(v) = \{e, e'\}$,
- ▷ *coupled* if they are in-coupled or out-coupled.

Finally rule **R8** imposes the maximization of the throughput by each splitter. Indeed, a saturated arc can provide more flow, while a fluid arc can absorb more flow. Thus, a steady-state cannot contain a saturated arc followed by a fluid arc. The only exception is when one of them already has a throughput of 1.

We will prove in [Section 3.4](#) that the definitions of splitter networks and steady-states exhibit a remarkable symmetry. By reversing each arc, exchanging the role of inputs and outputs, and complementing the set of fluid arcs, a steady-state is transformed into a steady-state of the reverse graph, as seen in [Figure 2](#).

For convenience, when defining or representing splitter networks, we will allow splitters with in-degree one or out-degree one (see [Figure 2](#) for instance). This is justified by the fact that if a splitter s has in-degree one, we can add a dummy input node i with capacity $c(i) = 0$. An arc from i to s can then be added, that will always remain fluid. Similarly if s has out-degree one, we can add a dummy output node o with capacity $c(o) = 0$, and an always-saturated arc from s to o . The throughputs on those arcs are forced to be 0. Therefore it does not induce any new constraint on the non-dummy arcs as rules **R6** and **R7** are clearly true for those arcs.

Additionally, for convenience, for any input $i \in I$ with outgoing arc e , we note $t(i) := t(e)$, and similarly for any output $o \in O$ with incoming arc e , $t(o) := t(e)$. We also extend the capacities to arcs by setting $c(e)$ to be either $c(i)$ if $e \in \delta^+(i), i \in I$, or $c(o)$ if $e \in \delta^-(o), o \in O$, or 1 otherwise.

1.2 Existence and computation of steady-states

Let F be a fixed set of fluid arcs. Then the set of possible throughput functions t of a steady-state (t, F) can be described as a polyhedron. Indeed, each of the rules **R1**, **R3**, **R4**, **R5**, **R6**, **R7** can be encoded by linear inequations. Rule **R8** is non-convex, but we will later introduce its slight strengthening, rule **R8S**. That stronger rule admits an encoding as a family of linear inequations. Thanks to linear programming, finding a steady-state thus reduces to finding a set of fluid arcs that admits a steady-state. Nevertheless, we still need to find F . We propose two algorithms to compute a steady-state, which relates to two families of maximum flow algorithm:

- ▷ a push-relabel-like algorithm, where we relax the conservation rule **R5**, thus defining a *pre-steady-state* by analogy with *pre-flows*. Given a set F , we use a linear program to compute an optimal pre-steady-state (t, F) (for some well-chosen objective), and prove that either (t, F) is a steady-state, or there is an arc $e \in F$ such that $(t, F \setminus e)$ is also a (non-optimal) pre-steady-state. Then after at most $|E|$ steps we get a steady-state;
- ▷ a blocking-flow-like algorithm, where we relax the rule **R3** on input capacities, removing the requirement that an input whose throughput is less than its capacity must have a saturated outgoing arc. This defines the notion of *sub-steady-state*. Given a set F , we solve a linear system to find a sub-steady-state t , and prove once again that either (t, F) is a steady-state or there is an arc $e \in F$ such that $(t, F \setminus e)$ is a sub-steady-state.

The pre-steady-state algorithm is technically simpler but requires an LP-solver. The sub-steady-state only requires an algorithm to compute stationary distributions in directed graphs. We defer a complete presentation and proof of these algorithms to [Sections 3.2](#) and [3.3](#), and focus for now on explaining the sub-steady-state algorithm.

Definition 4. Given $G = (I \uplus S \uplus O, E)$ a splitter network with capacities $c : I \uplus O \rightarrow [0, 1]$, a *sub-steady-state* for (G, c) is a pair (t, F) satisfying **R1**, **R2**, **R4**, **R5**, **R6**, **R7** and the strong maximization rule **R8S**, and for any $i \in I$ and $e \in \delta^+(i)$, $t(e) \leq c(i)$.

The algorithm starts with the trivial sub-steady-state $(t : e \rightarrow 0, E)$, and will improve it iteratively until reaching a steady-state. At each iteration of the algorithm, we will be trying to increase the throughputs of the arcs without violating any rule. Unlike in maximum flows, we do not have the choice of which leaving arc to increase the flow on. Furthermore, rule **R8** forces each splitter to send as much flow forward as possible. A non-obvious consequence is that, when increasing the input capacities, throughputs can




Figure 3: Four examples of throughput changes at a single splitter, depending on which arcs are fluid.




Figure 4: Configurations of splitters and the corresponding vertex in the residual graph. The outgoing arcs from a vertex of the residual graphs are highlighted in red: notice that in a sub-steady-state, the throughputs on these arcs must be equal.

only increase on fluid arcs, and can only decrease on saturated arcs. This suggests a definition of the residual graph for the sub-steady-state (t, F) . Its vertex set is $\{z\} \cup S$, where z is obtained by identifying all the inputs and outputs into a single node. Its edge set contains some fluid arcs and the reverses of some saturated arcs.

Consider the splitters in [Figure 3](#). We examine what happens when we increase the throughput on edge e_1 by $+\varepsilon$, or in case (d) when we decrease $t(e_3)$ by ε . In case (a) , by rule [R7](#), the throughputs on the two leaving arcs must stay equal, hence both increases by $\varepsilon/2$. In case (b) , only the throughput of the fluid leaving arc e_4 can increase. In case (c) , both leaving arcs are saturated, the splitter cannot push more flow downward, hence it is forced to push back flow through its incoming saturated arcs. Thus $t(e_2)$ decreases while $t(e_1)$ increases, by no more than $(t(e_2) - t(e_1))/2$ because of rule [R6](#). Finally in case (d) , if we decrease $t(e_3)$, then $t(e_1)$ and $t(e_2)$ must decrease by half as much.

Case (c) presents a challenge due to rule [R6](#), which imposes $t(e_1) \leq t(e_2)$. When $t(e_1) = t(e_2)$, the throughput of e_1 cannot increase, and the throughput of e_2 cannot decrease. We say that e_1 and e_2 are *tight*. In such a case, removing e_1 from F is allowed by rule [R6](#). Fluid arcs e with $t(e) = c(e)$ or saturated arc with $t(e) = 0$ are also *tight*, since we cannot modify their throughput further. Then we define the edge-set of the residual graph to only contain non-tight fluid arcs and reverses of non-tight saturated arcs.

Due to the conservation rule [R5](#), any iterative change to the throughputs of the network must be in accordance with a circulation of the residual graph. Because of rules [R6](#) and [R7](#), some arcs are constrained to have the same throughput. Therefore the chosen circulation itself has similar constraints. This is illustrated in [Figure 4](#), where the arcs that have equal throughput are highlighted in the residual graph. As may be readily checked, those constraints are exactly set on the leaving arcs in the residual graph of each vertex corresponding to a splitter. As for the special vertex z , obtained from the identification of the inputs and the outputs, we may non-deterministically select one of its leaving arc. Then we force all other arcs leaving z to have zero flow, by removing those arcs from the residual graph. From the residual graph, we compute a circulation satisfying each equality constraint. First compute a stationary distribution of a random walk on the residual graph. Then assign to each arc the probability of being the next arc in a random walk from that distribution. This results in a so-called *stationary circulation* (see [Figure 9](#)). One must be careful if the residual graph is not strongly connected. Then either we can find a strongly connected subgraph induced by the leaving arcs of some subset of vertices, or the residual graph contains a sink (as in [Figure 10](#)). In the former case we can still find a circulation, while in the latter case, we will




Figure 5: A network having several steady-states. Any value for ε between 0 and $\frac{1}{2}$ gives a steady-state.

be able to remove some arc from F .

Once a circulation is found, we increase the throughput as much as possible. This process will result in the creation of at least one sink in the updated residual graph. We show that when the residual graph contains a sink, some arc can be safely removed from F and becomes saturated. This bounds the number of steps until the algorithm stops, when z itself becomes a sink. At this point, any arc leaving an input node is either at full capacity or is saturated. Hence rule **R3** is satisfied, (t, F) is a steady-state. Some additional details, not covered in this presentation, are provided in [Section 3.3](#). An example of run of the algorithm on a simple network is given in [Figures 9 to 14](#). Summarizing the discussion, we get:

Theorem 1. *There is an algorithm that given a splitter network $G = (I \uplus S \uplus O, E)$ with capacities $c : I \uplus O \rightarrow [0, 1]$, finds a steady-state (t, F) in time $O(|S|^2 + |S| \text{sd}(G_z))$, where G_z is the graph obtained by identifying $I \cup O$ into a single vertex z , and $\text{sd}(G_z)$ denotes the time to compute a stationary distribution on any orientation of a subgraph of G_z .*

Steady-states are not unique: a directed cycle with no input or output can have any constant throughput on all its arcs. [Figure 5](#) showcases a more interesting network, having one input, one output, and many possible steady-states. However, in this example, all steady-states have the same throughputs on the inputs and outputs. Is there a network with two steady-states having different throughputs on their inputs and outputs? We conjecture that this cannot happen: steady-states are unique up to minor modifications, as in [Figure 5](#). Those modifications would be adding or removing some arcs from F , and adding or subtracting a circulation from the residual graph that leaves the inputs and outputs unchanged.

1.3 Balancers

We now define load-balancing networks and their properties. The goal of a load-balancing network is to divide some input flow evenly between several output belts. In the simplest case, the output belts can receive an arbitrarily large flow (up to the capacity of the belt). In more general cases, some outputs may be restricted but we still want the flow to be divided as evenly as possible, without limiting the total throughput available. We distinguish three properties of load-balancing networks. The first of these properties considers networks where the output capacities are not constrained.

Definition 5. A splitter network $G = (I \uplus S \uplus O, E)$ is a *balancer* if for any $c : I \uplus O \rightarrow [0, 1]$ such that for each output $o \in O$, $c(o) = 1$, there is a steady-state (t, F) for (G, c) with t constant on $\delta^-(O)$. An (n, p) -balancer is defined as a balancer with $|I| = n$ inputs and $|O| = p$ outputs.

When $|I| = |O| = 2^k$, the *simple balancer of order k* is a balancer network. It can be defined recursively: a simple balancer of order $k + 1$ is made from two simple balancers of order k in parallel. We identify each pair of outputs with equal index from the two balancers, creating a new splitter whose leaving arcs go to new output nodes. The recursive process is highlighted by blue boxes in [Figure 6](#). A drawback of the simple balancer occurs when the output capacities are not uniformly 1. Then the balancing property is lost, as can be seen on the network in the left side of [Figure 6](#).

Another limitation of simple balancers is that the total throughput at steady-state is not as much as we could expect. A simple upper bound on the total throughput is $\min\{c(I), c(O)\}$. It is reasonable to




Figure 6: On the left, the simple balancer of order 3, with a steady-state that is not balanced when some output capacity is not 1. The capacity of each input (resp. output) is given at their left (resp. right). On the right, a simple balancer of order 2, with a steady-state with total throughput less than both the total input capacity and the total output capacity.




Figure 7: A Beneš network of order 4 with the recursive structure being made explicit.

expect from a load-balancing network to always reach that bound. However, simple balancers do not have this property, as shown by the example on the right side of Figure 6. Improving over the definition of simple balancer, the concept of throughput-unlimited balancer imposes a maximized global throughput.

Definition 6. A balancer $G = (I \uplus S \uplus O, E)$ is *throughput-unlimited* if for any $c : I \uplus O \rightarrow [0, 1]$, there is a steady-state (t, F) for (G, c) such that total throughput $t(I) = t(O)$ is maximized at $\min\{c(I), c(O)\}$.

Notice that it has to be balancing only when the output capacities are uniformly 1. Beneš networks are throughput-unlimited networks with $|O| = |I| = 2^k$. They can be described as gluing two simple balancers, where the second balancer is reversed, see Figure 7. Observe that Beneš networks are their own reverses.

On the negative side, Beneš network are still not balancing when output capacities are not uniformly 1, for instance one could extend the steady-state in the network on the left side of Figure 6 to a steady-state in a Beneš network with the same throughputs. This calls for a stronger property, that a network should be load-balancing and throughput-unlimited for any capacity function. This is the notion of universal balancer.

Definition 7. A splitter network $G = (I \uplus S \uplus O, E)$ is *universally balancing* if for each capacity $c : I \uplus O \rightarrow [0, 1]$, there is a steady-state (t, F) and $\alpha, \beta \in \mathbb{R}_{\geq 0}$ such that

- (i) for each input i , $t(\delta^+(i)) = \min\{c(i), \alpha\}$,
- (ii) for each output o , $t(\delta^-(o)) = \min\{c(o), \beta\}$.
- (iii) the total throughput $T := t(\delta^+(I))$ equals $\min\{c(I), c(O)\}$.

In [Section 4](#), we will show how to build a universal balancer with $|I| = |O| = 2^k$. From such a universal balancer, by ignoring any set of inputs and outputs (setting their capacities to 0), we can make balancers with arbitrary numbers of inputs and outputs. We will also prove that every balancer presented here contains $\Theta(n \log n)$ splitters where n is the number of inputs and outputs.

Proposition 1. *The number of splitters in the simple balancer, Beneš network and universal network of order k are respectively $S(k) = k \cdot 2^{k-1}$, $B(k) = (2k - 1) \cdot 2^{k-1}$, and $U(k) = (k + 1)2^{k+2}$.*

1.4 Lower bounds on the number of splitters

Our next goal is to provide an $\Omega((n + p) \log(n + p))$ lower bound on the number of splitters in a (n, p) -balancer. We begin with what may seem as an unrelated problem: sampling in a discrete probability distribution. Given a fair coin that can be tossed arbitrarily often, how to choose an outcome in $\{1, \dots, d\}$, with probabilities given by a distribution $\pi \in [0, 1]^d$? First, consider the case when $\pi(i)$ is a rational for each $i \in \{1, \dots, d\}$, say $\pi(i) = p_i/q$ where q is a common denominator. Then a sequence of coin tossing can be described as a (possibly infinite) binary decision tree, with each leaf labeled with a sampled value. Here we present a construction of such a tree. Start from a single vertex, which serves as the root. Grow the tree in repeated iterations. At each iteration, add two children to every unlabelled leaf. As soon as the deepest level of the tree contains at least q leaves, label p_i of these leaves with i , for each $i \in \{1, \dots, d\}$. Once labeled, each leaf becomes definitive and will not grow anymore. The process goes on by once again growing the unlabelled leaves, as long (possibly infinitely) as some unlabelled leaf exists. After the tree is completed, the tree can be optimized using a simple trick repeated multiple times. If at any depth d , two leaves share a common label, move them under a common parent, then replace these two leaves with a single leaf at depth $d - 1$ bearing the same label. This process can be generalized to irrational probabilities, and gives a sampling algorithm that minimizes the number of coins tossed:

Theorem 2 ([13]). *Let $\pi \in [0, 1]^d$ a discrete probability distribution (so $\mathbb{1}\pi = 1$). Then the minimum expected number of coin tosses necessary to sample an element with probability distribution π is $\sum_{i=1}^d \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k(\pi_i)$. This minimum is achieved by a binary decision tree where at each depth k and for each $i \in \llbracket 1, d \rrbracket$, the number of leaves with label i is $\text{binary}_k(\pi_i)$.*

Consider a splitter network, and think of the flow as discrete, arbitrarily small items. An item enters the network from some input, then meets splitters repeatedly until reaching an output. When an item arrives at a splitter with both outgoing arcs being fluid, it will continue on any of the two outgoing arcs, without preference for one over the other because the splitter is fair. It implies that, from the perspective of this single item, the splitter network behaves like a coin-tossing network, with each splitter corresponding to a coin toss. If the network is a balancer, the sampled distribution is the uniform distribution on O .

Formally, when all the arcs remains fluid, increasing a single input capacity from 0 to 1 results in a non-decreasing throughput on each arc. Because all arcs are still fluid, the sub-steady-state algorithm performs a single iteration. Therefore the increase in throughputs follows a single stationary circulation. As illustrated on [Figure 8](#), it is obtained from the embedding of a binary decision tree T onto the splitter network. The increase in throughput on an arc e is the sum of probabilities of the edges mapped to e . Furthermore, in a balancer network, the increase of throughput is the same on every output. This implies that, as we progressively increase each input capacity from 0 to 1, each binary decision tree must uniformly sample from O .

In each binary decision tree, label each edge e with the probability of its usage during sampling. The sum of these labels represents the expected number of tosses, and can be bounded as shown in [Theorem 2](#).




Figure 8: The infinite decision tree (in blue) used to sample uniformly over a three-element set $\{a, b, c\}$ can be embedded from any input into a $(3, 3)$ -balancer. Moreover, the sum of the probabilities of the 3 trees, one from each input, will be at most one on any arc, which shows that this network is indeed a simple balancer.

When mapped into the splitter network, for an arc e , the sum of these labels on each edge of the tree mapped to e is the additional throughput on e . By summing over all the binary decision tree, we get that the sum of all labels is at most the number of outgoing arcs of all splitters, that is $2|S|$. Applied on balancers, it yields:

Theorem 3. *Let $G = (I \uplus S \uplus O, E)$ be an (n, p) -balancer, such that when all input capacities are 1, the steady-state has no saturated arc. Then*

$$|S| \geq \frac{1}{2}|I||O| \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k \left(\frac{1}{|O|} \right)$$

For a balancer with $|I| = |O| = 2^k$, since $\sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k \left(\frac{1}{|O|} \right) = \frac{k}{2^k}$, we get a lower bound of $k2^{k-1}$ splitters, matching the value of $S(k)$. Therefore the simple balancer of order k is optimal among all balancer networks without any saturated arcs in their steady-states. By extending this argument to steady-states with saturated arcs, we can remove that restriction, albeit at the cost of halving the lower bound.

Consider the various configurations of fluid and saturated arcs incident to a splitter, illustrated in [Figure 3](#). If a splitter has two fluid outgoing arcs, any additional flow is evenly distributed between the two outputs, akin to the probabilities of a coin toss. If a splitter has two incoming saturated arcs, by rule [R8](#), its outgoing arcs are saturated or at full capacity. In an augmenting circulation, the throughput on those arcs may only decrease by the same quantity by rule [R6](#): the splitter still acts as a coin toss, but on the flow that is pushed back. Otherwise, a positive change in throughput on an incoming arc will be followed by an increase on a single outgoing fluid arc or a decrease on a single incoming saturated arc. Similarly a negative change of throughput on an outgoing saturated arc will impact only one other arc. Any additional unit of flow entering the splitter would be routed deterministically. Therefore, in the embedding of a binary decision tree into the splitter network, a node cannot be mapped to such a vertex, and no coin toss occurs here. Thus any splitter, depending on which of its incident arcs are fluid, acts as either a coin toss or a deterministic router. Thus, even in the presence of saturated arcs, we can embed a binary decision tree, by mapping each edge to a directed path in the residual graph. The inner nodes of any such path are *deterministic* splitters, while its extremities are *tossing* splitters. By [Corollary 3](#),

the throughput on each arc increases until it becomes saturated, then decreases. Therefore its throughput varies by at most 2 during the whole algorithm. This limits the extent to which an arc can be utilized by the embeddings of binary decision trees, leading us to the following conclusion:

Theorem 4. *Let $G = (I \uplus S \uplus O, E)$ be an (n, p) -balancer. Then*

$$|S| \geq \frac{1}{4}|I||O| \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k \left(\frac{1}{|O|} \right)$$

We will formalize this discussion and prove the theorem in [Section 5](#)

2 The $\mathcal{C}^=$ -circulation problem

Let $G = (V, E)$ be a directed graph. We consider a partition $\mathcal{C}^= \subseteq 2^E$ of the arcs, such that each $C \in \mathcal{C}^=$ is a subset of the outgoing arcs $C \subseteq \delta^+(v)$ of some vertex $v \in V$. Two arcs $e', e'' \in E$ are considered $\mathcal{C}^=$ -coupled (or just *coupled* when no ambiguity arises) if there is $C \in \mathcal{C}^=$ such that e', e'' belongs to C (an arc is coupled to itself). An arc e' is *single* if $\{e'\} \in \mathcal{C}^=$, indicating that it is not coupled to another arc. The tuple $(G, e, \mathcal{C}^=)$, where $e \in E$, is then an instance of the $\mathcal{C}^=$ -circulation problem: find a circulation in G that is non-zero on e and is constant on any set of $\mathcal{C}^=$. It can be expressed as finding a vector x with $x_e > 0$ satisfying the following linear system:

$$\begin{cases} x(\delta^+(v)) - x(\delta^-(v)) = 0 & (v \in V) \\ x_e - x_{e'} = 0 & (e, e' \in C \in \mathcal{C}^=) \\ x \geq 0 \end{cases} \quad (1)$$

Theorem 5. *An instance $(G = (V, E), ts, \mathcal{C}^=)$ of the $\mathcal{C}^=$ -circulation problem has a solution x with $x_{ts} > 0$ if and only if there is no set $S \subseteq V \setminus \{t\}$ with a partition $S = S_0 \uplus S_1 \dots \uplus S_k$ where $s \in S_0$ and for any arc $e \in \delta^+(S_i)$, there is an arc e' coupled to e (possibly $e = e'$) such that $e' \in E[S_i, S_j]$ and $i < j$.*

Proof. Suppose $S = S_0 \uplus S_1 \dots \uplus S_k$ exist and let x denote a solution to (1). We prove by induction from k to 0 that for all $i \in \llbracket 0, k \rrbracket$, $x(\delta^+(S_i)) = 0$. This will imply that $x_{ts} = 0$ as $ts \in \delta^-(S_0)$. By assumption, $\delta^+(S_k) = \emptyset$, and because x is a circulation, $x(\delta^-(S_k)) = 0$, proving the base case. Let $i \in \llbracket 1, k-1 \rrbracket$ and suppose that we have for all $j \in \llbracket i+1, k \rrbracket$, $x(\delta^+(S_j)) = 0$. Let $uv \in \delta^+(S_i)$, then there is an arc $uw \in E[S_i, S_j]$ with $j > i$, where uw is either equal to or coupled with uv , hence by the induction hypothesis $x_{uw} = x_{uv} = 0$. Thus $x(\delta^+(S_i)) = 0$. Again, by x being a circulation, we get that $x(\delta^-(S_i)) = 0$.

Now suppose that there is no $x \in \mathbb{R}_{\geq 0}^E$ satisfying (1) and $x_{ts} > 0$. By Farkas lemma, this implies that the following linear system has a solution $y \in \mathbb{R}^V$:

$$\begin{aligned} y_s - y_t &\geq 1 \\ \sum_{uv \in C} y_v - y_u &\geq 0 \quad (C \in \mathcal{C}^=) \end{aligned}$$

We may partition V into $S_{-l}, S_{-l+1}, \dots, S_k$, the equivalence classes defined by the relation $u \sim v$ if $y_u = y_v$, ordered based on increasing y -values, and such that $s \in S_0$. Because $y_t < y_s$, $t \notin S_0 \cup S_1 \dots \cup S_k$.

Let $i \in \llbracket 0, k \rrbracket$ and let $uv \in \delta^+(S_i)$. Let $C \in \mathcal{C}^=$ with $uv \in C$. Then $\sum_{uw \in C} y_w - y_u \geq 0$ implies that either $y_w = y_u$ for all $uw \in C$, which is not the case as $uv \in \delta^+(S_i)$ hence $y_u \neq y_v$, or there is some $uw \in C$ with $y_w > y_u$. But then $w \in S_j$ with $j > i$, proving that the partition S satisfies the desired property. \square

Corollary 1. *An instance $(G = (V, E), ts, \mathcal{C}^=)$ of the $\mathcal{C}^=$ -circulation problem where G is strongly connected has a solution x with $x_{ts} > 0$.*

Proof. By contraposition, consider an instance lacking a non-zero circulation, and let $S_0 \cup \dots \cup S_k$ be the partition given by [Theorem 5](#). Then $\delta^+(S_k)$ is empty, indicating that G is not strongly connected. \square

We can compute efficiently a $\mathcal{C}^=$ -circulation in the strongly connected case. We do so by restricting the problem to a spanning subgraph in which the classes in $\mathcal{C}^=$ coincides with the leaving arcs of each vertex.

Lemma 1. *There is an algorithm that, given an instance $(G = (V, E), ts, \mathcal{C}^=)$ of the $\mathcal{C}^=$ -circulation problem with G strongly connected, find a feasible solution x with $x_{ts} > 0$ in time $O(|E| + \text{sd}(G))$, where $\text{sd}(G)$ is the complexity of computing a stationary distribution on any subgraph of G .*

Proof. Let T denote a maximal in-arborescence of G rooted at t . Let

$$E' := \{e \in E : e \text{ is coupled to some arc in } T \cup \{ts\}\}.$$

By construction, t is reachable from any vertex in (V, E') . Because $ts \in E'$, s and t are in the same strongly connected component of (V, E') . Let $H = (V_H, E_H)$ be the strongly connected component of (V, E') containing s and t . Because any vertex in T has out-degree at most 1, and by our choice of E' , for any $v \in V$, $\delta^+(v) \cap E' \in \mathcal{C}^=$. Let $\mathcal{C}_H^= := \{C \in \mathcal{C}^= : C \subseteq E_H\}$.

We can rewrite the problem of finding a $\mathcal{C}_H^=$ -circulation on H , by setting $y_v := x_e$ for $e \in \delta^+(v)$:

$$\begin{aligned} \sum_{uv \in \delta^-(v)} y_u &= d^+(v)y_v, \\ y &\geq 0, \end{aligned}$$

and then setting $\pi_v = d^+(v)y_v$:

$$\begin{aligned} \sum_{uv \in \delta^-(v)} \frac{\pi_u}{d^+(u)} &= \pi_v, \\ \pi &\geq 0. \end{aligned}$$

We recognize the last system as the one defining a stationary distribution of a Markov chain over H , provided that we add the constraint $\mathbb{1}\pi = 1$. As H is strongly connected, the stationary distribution exists and is nowhere zero. Thus we obtain a non-zero $\mathcal{C}_H^=$ -circulation for H . We extend it into a non-zero $\mathcal{C}^=$ -circulation for G by setting $x_e = 0$ for any edge not in E_H . \square

The strongly connected case serves as a basis for an algorithm solving the general problem.

Lemma 2. *There is an algorithm that, given an instance $(G = (V, E), ts, \mathcal{C}^=)$ of the $\mathcal{C}^=$ -circulation problem, in time $O(|E| \log^4 |V| + \text{sd}(G))$, either find a feasible solution x with $x_{ts} > 0$ or correctly assert that none exists. Here $\text{sd}(G)$ is the complexity of computing a stationary distribution on any subgraph of G .*

Proof. We start by computing a strongly connected subgraph H of G containing ts and ensuring that for any $C \in \mathcal{C}^=$, either $C \subseteq E(H)$ or $C \cap E(H) = \emptyset$. If no such H exists, then the $\mathcal{C}^=$ -circulation instance does not admit a non-zero solution. If H exists, we reduce the problem of finding x to computing a stationary distribution in H .

To compute H , we use a dynamic decremental single-sink reachability algorithm with sink t . If ever t becomes unreachable from s , then we conclude that the instance does not admit a non-zero solution. The algorithm keeps a queue of vertices to delete, initially the vertices from which t is not reachable. While the queue is non-empty, it takes a vertex from it and delete its incident arcs and all the arcs coupled to them. As soon as it detects that some vertex v cannot reach t , it queues that vertex v for deletion. When the queue is empty, remove all the vertices not reachable from s and return the remaining subgraph H .

First, suppose that this algorithm fails to build H because t becomes unreachable from s . Each time a subset of vertices becomes disconnected from t , we label it S_i (with decreasing index i , adjusting the value of the starting index at the end of the algorithm). When such a component appears, all its leaving arcs must have been removed, which means they are coupled to an arc entering some already removed component. Hence we get the sequence $S_0, S_1 \dots S_k$ proving that a non-zero $\mathcal{C}^=$ -circulation cannot exist.

Suppose now that the algorithm returns a subgraph $H = (V_H, E_H)$. By construction, for any $C \in \mathcal{C}^=$, either $C \subseteq E_H$ or $C \cap E_H = \emptyset$. Indeed, if any arc of C is removed during the main loop, then C is removed in the same iteration. If some arc of C is removed during the final phase, when vertices unreachable from s are removed, then the vertex v for which $C \subseteq \delta^+(v)$ must not be reachable from s and thus C was removed.

Next we claim that H is strongly connected. By construction, for any vertex v , there is a path from t to v starting with ts . Moreover, notice that before removing the vertices unreachable from s , there was a path from v to t . Thus any vertex on this path is also reachable from s and is kept during the last phase. Hence t is reachable from v , proving that H is strongly connected. We then apply [Lemma 1](#) to compute x , and extend it to G by setting $x_e = 0$ for any arc not in H . The complexity follows by using the decremental single-sink reachability algorithm from [\[3\]](#). \square

The next lemma gives a sufficient condition for the existence of a non-zero $\mathcal{C}^=$ -circulation, when we may choose any arc to be positive instead of the specific arc ts as above.

Lemma 3. *Given a connected directed graph $G = (V, E)$ and a partition $\mathcal{C}^=$ of E that is a refinement of the out-incidencies of G , then*

- \triangleright either there is a non-zero $\mathcal{C}^=$ -circulation in G ,
- \triangleright or there is a vertex $v \in V$ with $\delta^+(v) = \emptyset$,

and we can find one or the other in time $O(|E| + \text{sd}(G))$.

Proof. One can find in time $O(|E|)$ a strongly connected component X of G that is a sink component: $\delta^+(X) = \emptyset$. If $|X| = 1$ we are done, because $v \in X$ is a sink. Now, assume that $|X| > 1$.

Because for any $C \in \mathcal{C}^=$, there is a vertex v with $C \subseteq \delta^+(v)$, for any arc $e \in E[X]$, any arc coupled to e is also in $E[X]$. Hence, a $\mathcal{C}^=$ -circulation x of $G[X]$ can be extended to a $\mathcal{C}^=$ -circulation on G by setting $x_e = 0$ for each $e \notin E[X]$. Such a $\mathcal{C}^=$ -circulation on $G[X]$ exists and can be computed in time $O(|E| + \text{sd}(G))$ by [Lemma 1](#). \square

3 Computing a steady-state

In this section, we give two algorithms to compute a steady-state in a splitter network.

3.1 A stronger maximization rule

The algorithms use a slightly stronger property than rule [R8](#):

[R8S](#) for any arcs $uv \in E \setminus F$ and $vw \in F$, $t(vw) = 1$.

Clearly rule [R8S](#) implies rule [R8](#). The two rules are actually almost equivalent:

Lemma 4. *If (t, F) is a steady-state, then there is a steady-state (t, F') that satisfies rule [R8S](#), with $F \subseteq F'$.*

Proof. By induction on the number of splitters on which rule [R8S](#) is not true. Let $uv \notin F$, $vw \in F$ with $t(vw) < 1$, by rule [R8](#) $t(uv) = 1$. Let $u'v$ be the arc in-coupled to uv . If $u'v$ is fluid, then $(t, F \cup \{uv\})$ is a steady-state as rule [R6](#) is checked on v and rule [R7](#) is checked on u . If $u'v$ is saturated, by rule [R6](#), $t(u'v) = 1$, then $(t, F \cup \{uv, u'v\})$ is a steady-state for similar reasons. Notice that in both cases, the modification preserves rule [R8S](#) on any splitter for which it held. Hence the number of those splitters increases by one. \square

3.2 An push-relabel-like algorithm to compute steady-states

In analogy with a pre-flow in a push-relabel max-flow algorithm, we relax the conservation rule **R5**, to define:

Definition 8. Given a splitter network $G = (I \uplus S \uplus O, E)$ with capacities $c : I \uplus O \rightarrow [0, 1]$, a *pre-steady-state* for (G, c) is a pair (t, F) satisfying **R1**, **R2**, **R3**, **R4**, **R6**, **R7**, and **R8S**, and such that for each splitter $s \in S$,

$$t(\delta^+(s)) \leq t(\delta^-(s)).$$

The next claim which can be readily checked, asserts the existence of a simple pre-steady-state.

Claim 1. *The pair $(t, E(G))$ is a pre-steady-state for (G, c) , where*

$$t : uv \rightarrow \begin{cases} c(u) & \text{if } u \in I, \\ 0 & \text{otherwise} \end{cases}$$

We proceed by iteratively improving upon this initial pre-steady-state. Let $G = (I \uplus S \uplus O, E)$ be a splitter network with capacities $c : I \uplus O \rightarrow [0, 1]$. Let $F \subseteq E$ be a set of fluid arcs. If (t, F) is a pre-steady-state, then t is a feasible solution to the following linear program:

$$\left| \begin{array}{ll} \max t(\delta^-(O)) & \text{subject to} \\ 0 \leq t \leq 1 & (e \in E) \\ t(is) \leq c(i) & (is \in \delta^+(I)) \\ t(so) \leq c(o) & (so \in \delta^-(O)) \\ t(wv) \leq t(uv) & (uv \in E \setminus F, wv \in E) \\ t(vw) = 1 & (uv \in E \setminus F, vw \in F) \end{array} \right. \quad \begin{array}{ll} t(\delta^+(s)) \leq t(\delta^-(s)) & (s \in S) \\ t(is) = c(i) & (is \in \delta^+(I) \cap F) \\ t(so) = c(o) & (so \in \delta^-(O) \setminus F) \\ t(uw) \leq t(uv) & (uv \in F, uw \in E) \end{array} \quad \text{(PSS)}$$

Conversely, the following lemma is immediate:

Lemma 5. *Given $G = (I \uplus S \uplus O, E)$ a splitter network, $c : I \uplus O \rightarrow [0, 1]$ a capacity function and $F \subseteq E$ a set of fluid arcs, for any feasible solution t of **(PSS)**, (t, F) is a pre-steady-state of (G, c) .*

Our algorithm is based on the next lemma.

Lemma 6. *Let t^* be an optimal solution to **(PSS)** such that $\sum_{e \in F} t(e) - \sum_{e \in E \setminus F} t(e)$ is maximized. Then either (t^*, F) is a steady-state, or there is an arc $e \in F$ such that $(t^*, F \setminus \{e\})$ is a pre-steady-state.*

Proof. Assume that (t^*, F) is not a steady-state: there is a vertex $s \in S$ with $t^*(\delta^-(s)) - t^*(\delta^+(s)) > 0$. Let $\{e_1, e_2\} = \delta^-(s)$ with $t^*(e_1) \leq t^*(e_2)$, and $\{e_3, e_4\} = \delta^+(s)$ with $t^*(e_3) \leq t^*(e_4)$.

Case 1: there is a fluid leaving arc $e \in \{e_3, e_4\} \cap F$ such that e is in-coupled to an arc $e' \in E \setminus F$ and $t^*(e) = t^*(e')$. Then $(t^*, F \setminus \{e\})$ is a pre-steady-state, as rules **R6** and **R8S** are clearly still satisfied.

Case 2: there is a fluid leaving arc $e \in \{e_3, e_4\} \cap F$ such that $e \in \delta^-(o)$ for some output $o \in O$, and $t^*(e_3) = c(o)$. Then $(t^*, F \setminus \{e\})$ is a pre-steady-state, as rules **R4** and **R7** are clearly still satisfied.

We may now assume that the fluid leaving arcs do not check the conditions for cases **1** and **2**. Consequently, increasing the output $t^*(e)$ to $t^*(e) + \varepsilon$ for some sufficiently small ε do not break any rule on the destination node of e .

Case 3: e_3 and e_4 are both fluid, with $t^*(e_3) = t^*(e_4) < 1$. Then, for some $\varepsilon > 0$ sufficiently small, $(t^* + \varepsilon \chi_{\{e_3, e_4\}}, F)$ is a pre-steady-state, because we increase the throughput on both arcs uniformly, preserving rule **R7**. achieving a better objective than t^* . The throughput of this pre-steady-state improves over t^* , a contradiction.

Case 4: exactly one of e_3, e_4 is fluid with throughput strictly less than one, and we may assume it is e_4 by rule **R7**. Then, for some $\varepsilon > 0$ sufficiently small, $t^* + \varepsilon\chi_{\{e_4\}}$ is a better solution than t^* , again a contradiction.

We may now assume that $e_3 \notin F$ or $t^*(e_3) = 1$, and that $e_4 \notin F$ or $t^*(e_4) = 1$.

Case 5: $e_2 \in F$. Then $(t^*, F \setminus \{e_2\})$ is a pre-steady-state, as rules **R6** and **R8S** are both satisfied.

We may now assume that $e_2 \in E \setminus F$.

Case 6: $e_1 \in F$ and $t^*(e_1) = t^*(e_2)$. Then $(t^*, F \setminus \{e_1\})$ is a pre-steady-state, as rules **R6** and **R8S** are satisfied.

Case 7: $e_1 \in E \setminus F$. Then $t^* - \varepsilon\chi_{\delta^-(s)}$ is a better solution than t^* for some $\varepsilon > 0$ sufficiently small, which leads to a contradiction.

Case 8: $t^*(e_1) < t^*(e_2)$, $e_1 \in F$, $e_2 \notin F$. Then, for some $\varepsilon > 0$ sufficiently small, $t^* - \varepsilon\chi_{\{e_2\}}$ is a better solution to (**PSS**) than t^* , again a contradiction. \square

This leads to the iterative algorithm of repeatedly solving the LP and removing an arc from F , that stops after at most $|E| + 1$ iterations, and solves $|E| + 1$ linear programs of polynomial size in the worst case.

Notice that the proof still holds with additional constraints $t(e) \geq l_e$ for any fluid arc e , and $t(e) \leq u_e$ for any saturated arc e . This is because the algorithm tries to increase the throughput on fluid arcs, and decrease it on saturated arcs. Therefore, as long as the algorithm can be initialized with a feasible solution, we can find a steady-state that also checks these constraints. The feasible solution is a pre-steady-state that already satisfies those additional constraints. This implies the following corollary.

Corollary 2. *Given $G = (I \uplus S \uplus O, E)$ a splitter network, $c : I \uplus O \rightarrow [0, 1]$ a capacity function and (t_0, F_0) a pre-steady-state for G, c . Then there exists steady-state (t, F) for G, c such that*

- (i) $F \subseteq F_0$;
- (ii) for each arc $e \in F_0$, $t(e) \leq t_0(e)$;
- (iii) for each arc $e \in E \setminus F$, $t(e) \geq t_0(e)$.

Proof. Find a steady-state with the additional constraints $t(e) \leq t_0(e)$ for each $e \in F_0$, and $t(e) \geq t_0(e)$ as long as e is a fluid arc. \square

3.3 A blocking-flow-like algorithm to compute steady-states

Let (t, F) be a sub-steady-state for a splitter network $G = (I \uplus S \uplus O, E)$ with capacities $c : I \uplus O \rightarrow [0, 1]$. Similar to the blocking-flow algorithm, we employ a notion of residual graph. We notice that throughput values may be bounded by rules **R6**, which governs the throughputs entering a splitter, and **R7**, which bounds the throughput at the outputs. Let s be a splitter, with $\delta^-(s) = \{e, e'\}$, where $e \in F$ and $e' \in E \setminus F$, and $t(e) = t(e')$, then $t(e)$ is lower-bounded and $t(e')$ is upper-bounded by rule **R6**. We say that e is *upper-tight* and e' is *lower-tight*. We also say that a fluid arc e is *upper-tight* when $t(e) = c(e)$, and a saturated arc e is *lower-tight* when $t(e) = 0$. Thus, an arc is tight when its throughput cannot be modified without violating some rule. Furthermore, rules **R6** and **R7** imposes some arcs to have equal throughput: out-coupled fluid arcs, or in-coupled saturated arcs. Any modification to such an arc imposes a modification to its coupled arc. Therefore, we say that a fluid arc is *loose* if none of its out-coupled fluid arcs is upper-tight. A saturated arc is *loose* if none of its in-coupled saturated arcs is lower-tight.

Definition 9. Let $G = (I \uplus S \uplus O, E)$ be a splitter network, with capacities $c : I \uplus O \rightarrow [0, 1]$ and a sub-steady-state (t, F) . The *residual graph* for (t, F) is the pair $(H, C^=)$ where $H = (V_H, E_H)$ is a graph defined by

- ▷ $V_H := \{z\} \cup S$ where z is a new vertex.
- ▷ for each loose fluid arc $uv \in F$, $\phi(u)\phi(v) \in E_H$,
- ▷ for each loose saturated arc $uv \in E \setminus F$, $\phi(v)\phi(u) \in E_H$,

where $\phi(u) = z$ if $u \in I \cup O$, $\phi(u) = u$ otherwise.

We denote ρ the natural bijection from the loose arcs in E to E_H . For each loose $uv \in E$, let $C_{uv} := \{\rho(u'v') : u'v' \in E \text{ loose and coupled with } uv\}$ and $\mathcal{C}^- := \{C_{uv} : uv \in E \text{ loose}\}$. Notice that \mathcal{C}^- forms a partition of E_H , and for all $C \in \mathcal{C}^-$, there is a vertex $v \in V_H$ with $C \subseteq \delta_H^+(v)$. The connected component of H containing z will be called its *main component*.

Coupled arcs must have the same throughput, thus their throughput can only be changed by an equal amount. We define \mathcal{C}^- to be the partition of the arcs into the sets of coupled arcs. Our next result is that a \mathcal{C}^- -circulation of the residual graph fills the role of an augmenting flow for the sub-steady-state.

Definition 10. For a sub-steady-state (t, F) , we define $\psi(t, F)$ by

$$\psi(t, F) = |\{e \in E \setminus F : t(e) = 0\}| - |\{e \in F : t(e) < c(e)\}| \quad (3)$$

Lemma 7. Let $G = (I \uplus S \uplus O, E)$ be a splitter network with capacities $c : I \uplus O \rightarrow [0, 1]$. Let x be a non-zero circulation on the residual graph H of a sub-steady-state (t, F) . Then there is a value $\lambda > 0$ such that (\bar{t}, F) is a sub-steady-state, where

$$\bar{t}(e) := \begin{cases} t(e) & \text{if } e \in E \text{ is not loose,} \\ t(e) + \lambda x_{\rho(e)} & \text{if } e \in F \text{ is loose,} \\ t(e) - \lambda x_{\rho(e)} & \text{if } e \in E \setminus F \text{ is loose.} \end{cases}$$

Moreover, either $\psi(\bar{t}, F) > \psi(t, F)$, or there is an arc $e \in F$ such that $(\bar{t}, \bar{F} := F \setminus \{e\})$ is a sub-steady-state with $\psi(\bar{t}, \bar{F}) > \psi(t, F)$.

Proof. We compute $\lambda > 0$ to be the maximum value such that $(\bar{t} := t + \lambda x, F)$ is a sub-steady-state. A positive λ exists because the support of x is the set of loose arcs of G , thus choosing λ small enough will not break any sub-steady-state rule. Notice that by definition of \bar{t} , the contribution of any arc e to ψ cannot decrease, thus we need only find one arc whose contribution increases. By the maximality of λ , there is a loose arc in (t, F) that is no longer loose in (\bar{t}, F) , that is:

- ▷ either there is an arc $e \in F$, $t(e) < \bar{t}(e) = c(e)$, hence $\psi(\bar{t}, F) > \psi(t, F)$;
- ▷ or there is an arc $e \notin F$, $t(e) > \bar{t}(e) = 0$, hence $\psi(\bar{t}, F) > \psi(t, F)$;
- ▷ or there is a vertex $v \in S$ with $\delta^-(v) = \{e, e'\}$, $e \in F$, $e' \in E \setminus F$ with $t(e') > t(e)$ and $\bar{t}(e') = \bar{t}(e)$ (e and e' stop being loose). By rule **R8S**, arcs in $\delta^+(v)$ are either saturated or have throughput 1 in t and in \bar{t} . Then $(\bar{t}, F \setminus \{e\})$ is a sub-steady-state as rules **R6** and **R8S** are satisfied, and $\psi(\bar{t}, F \setminus \{e\}) > \psi(t, F)$, as the contribution of e increases. \square

Lemma 8. Let (t, F) be a sub-steady-state for a splitter network $G = (I \uplus S \uplus O, E)$ with capacities $c : I \uplus O \rightarrow [0, 1]$, and let H be its residual graph. If the main connected component of H contains a sink $v \in S$, then there is an arc $e \in \delta_G^+(v)$ such that $(t, F \setminus \{e\})$ is a sub-steady-state.

Proof. First, if an edge e is fluid with $t(e) < c(e)$ but e is not in H , then e is not loose. By definition of loose, e is out-coupled to a fluid arc e' that is upper-tight. Therefore, e' is in-coupled to a saturated arc $e'' \notin F$ with $t(e') = t(e'')$. In that case, $(t, F \setminus \{e'\})$ is a substeady-state (checking rules **R6** and **R8S**). Similarly, if e is saturated such that $t(e) > 0$ but the reverse of e is not in H , then e is in-coupled to a lower-tight arc e' . e' is in-coupled to a fluid arc e'' with $t(e') = t(e'')$. In that case, $(t, F \setminus \{e''\})$ is a substeady-state.

Consider a sink v in H . We may now assume that for any $e \in \delta_G^+(v)$, either $t(e) = c(e)$ or $e \notin F$, and for any $e \in \delta_G^-(v)$, $t(e) = 0$ or $e \in F$. Let $e \in \delta_G^-(v)$ such that $t(e)$ is maximum. If $e \in F$, then $(t, F \setminus \{e\})$ is a sub-steady-state, as rules **R6** and **R8S** are satisfied. If $e \notin F$, then $t(e) = 0$, and by rule **R6**




Figure 9: Starting from a trivial sub-steady-state, we compute a residual graph and a stationary circulation in this graph (the two vertices marked z should be identified). Then we increase the throughputs accordingly, as much as possible without violating a sub-steady-state rule, by adding $\lambda = 6$ times the circulation at which point some edge reaches its capacity (see Figure 10).

$0 = \sum_{e \in \delta_G^-(v)} t(e) = \sum_{e \in \delta_G^+(v)} t(e)$, hence $\delta_G^+(v) \subseteq E \setminus F$. Because v is not isolated, the arc e' out-coupled to e must be loose. If e' is saturated, it is in-coupled to the lower-tight arc e , contradicting the fact that e' is loose. Then e' is fluid, with $t(e') = 0$, and $(t, F \setminus \{e'\})$ is a sub-steady-state, as rules R6 and R8S are satisfied on v and rule R7 is satisfied on the origin of e' . \square

As in an augmenting-path algorithm, we will repeatedly build the residual graph, and deduce either that the current pair (t, F) is a steady-state, or find a better solution.

Lemma 9. *There is an algorithm that, given a splitter network $G = (I \uplus S \uplus O, E)$ with capacities $c : I \uplus O \rightarrow [0, 1]$, and a sub-steady-state (t, F) , either checks that (t, F) is a steady-state, or find another sub-steady-state (\bar{t}, \bar{F}) such that $\psi(\bar{t}, \bar{F}) > \psi(t, F)$ and that runs in time $O(n + \text{sd}(G_z))$, where $\text{sd}(G_z)$ is the complexity of finding a stationary distribution on any residual graph for G .*

Proof. We apply Lemma 3 to the main component of the residual graph H , from which we get one of two cases.

Case 1: there is a non-zero circulation. By Lemma 7, the sub-steady-state can be improved by this circulation into a new sub-steady-state (\bar{t}, \bar{F}) with $\psi(\bar{t}, \bar{F}) > \psi(t, F)$.

Case 2: there is a sink v in the main component of H . If $v \notin S$, v is the special vertex z obtained by identifying the vertices of $I \cup O$, then every arc in $\delta_G^+(I)$ is saturated or not loose, which implies that rule R3 is checked, hence (t, F) is a steady-state. If $v \in S$, by Lemma 8 there is a fluid arc $e \in \delta_G^-(v) \cap F$ such that $(t, F \setminus \{e\})$ is a sub-steady-state. \square

Proof of Theorem 1. The algorithm repeatedly applies Lemma 9 from the initial sub-steady-state $(t : e \rightarrow 0, F)$. Each iteration takes time $O(|E| + \text{sd}(G))$ and increases ψ by one. As ψ is initially $-|E|$, and is at most $|E|$, this bounds the number of iterations by $2m$, from which we get the claimed complexity. \square

3.4 Reverse splitter networks and consequences

We present some consequences of the algorithms and of the following fact that splitter networks can be reversed. Because splitters are fair on their input as well as on their output, splitter networks and steady-states display a useful symmetry. For an arc set X , we denote $\overleftarrow{X} := \{vu : uv \in X\}$ the set of reverse arcs of X .

Definition 11. Let $G = (I \uplus S \uplus O, E)$ be a splitter network, we denote \overleftarrow{G} the reverse of G , defined as the splitter network $(O \uplus S \uplus I, \overleftarrow{E})$, where the inputs and outputs are interchanged.




Figure 10: We compute a new residual graph, which does not contain the arc with throughput 1, since this arc cannot increase. Then the existence of a sink prevents us to find a stationary circulation in this residual graph (the two vertices marked z should be identified). We remove from F the incoming arc to the sink with highest throughput, and go to the next iteration.




Figure 11: We compute a new residual graph, including the reverse of the newly saturated arc. Then we compute a stationary circulation in this residual graph. This leads to an improved sub-steady-state (taking $\lambda = \frac{3 \cdot 26}{4 \cdot 14}$, when another edge reaches throughput 1, see [Figure 12](#))




Figure 12: The next residual graph contains a sink, thus we remove one more arc from F . This yields the residual graph on the right. We compute the stationary circulation in this residual graph. Then we improve the sub-steady-state according to this circulation (here $\lambda = \frac{4}{7}$, when the incoming arcs of the central vertex become tight, see [Figure 13](#))




Figure 13: Again the residual graph contains sinks, that we remove by making some arcs saturated. This yields the residual graph on the right, free of any sink. We compute a stationary circulation in this residual graph, and use it to improve the sub-steady-state (taking $\lambda = \frac{9}{14}$, when the incoming arcs of the leftmost splitter become tight, see Figure 14)




Figure 14: After removing one more arc from F , z becomes a sink in the residual graph. It implies that (t, F) is a steady-state. The algorithm stops.




Figure 15: Increasing the input capacities may result in reduced effective throughputs on some inputs, as if they are in competition. Similarly, increasing the output capacities may also result in reduced effective throughputs on some outputs.

Lemma 10. *Let $G = (I \uplus S \uplus O, E)$ be a splitter network with capacities, $c : I \uplus O \rightarrow [0, 1]$, and (t, F) a steady-state for (G, c) . Then $(t, \overleftarrow{E} \setminus \overleftarrow{F})$ is a steady-state for (\overleftarrow{G}, c) .*

Proof. We check that each steady-state-defining rule is satisfied. Rules **R1**, **R2**, **R5** and **R8** can be readily checked. Rules **R3** and **R4** translates into each other, as do rules **R6** and **R7**. \square

Observe that the strong maximization rule **R8S** does not translate well into \overleftarrow{G} , hence the reverse of a steady-state does not necessarily check rule **R8S**. But by **Lemma 4**, there is a steady-state $(t, \overleftarrow{E} \setminus \overleftarrow{F'})$ of (\overleftarrow{G}, c) with $F' \subseteq F$ that satisfies rule **R8S**.

A careful look at the sub-steady-state algorithm shows that, during its resolution, each edge $t(e)$ starts in F with $t(e) = 0$, then $t(e)$ increases, before possibly being removed from F , then $t(e)$ decreases.

Corollary 3. *If (t, F) is a sub-steady-state for (G, c) , then there exists an steady-state (\bar{t}, \bar{F}) with $\bar{F} \subseteq F$, and for each $e \in E$, if $e \in \bar{F}$ then $\bar{t}(e) \geq t(e)$, and if $e \in E \setminus \bar{F}$, $\bar{t}(e) \leq t(e)$.*

If we increase the input capacities to \bar{c} , a steady-state (t, F) for (G, c) becomes a sub-steady-state for (G, \bar{c}) . Hence,

Corollary 4. *If (t, F) is a steady-state for (G, c) , and \bar{c} be a capacity function with $\bar{c}(u) \geq c(u)$ for each input or output $u \in I \cup O$. There exists a steady-state (\bar{t}, \bar{F}) for (G, \bar{c}) such that:*

- (i) if $\bar{c}(o) = c(o)$ for each output $o \in O$, then $\bar{t}(e) \geq t(e)$ for each $e \in \delta^-(O)$;
- (ii) if $\bar{c}(i) = c(i)$ for each input $i \in I$, then $\bar{t}(e) \geq t(e)$ for each $e \in \delta^+(I)$.

Proof. (t, F) is a sub-steady-state, which we can improve using the algorithm from **Theorem 1**. Let (\bar{t}, \bar{F}) be the resulting steady-state. Notice that if an arc $e \in \delta^-(O)$ is saturated, by rule **R4**, $t(e) = c(e)$, hence e is not loose. Thus the throughputs of saturated arcs incident to outputs cannot decrease. This proves (i). Then (ii) follows by taking the reverse network and applying (i) to the reverse steady-state, thanks to **Lemma 10**. \square

Thus increasing the input capacities will not decrease the output throughputs. Notice that increasing the input capacities may lead to the decrease in the throughput of some of the inputs, as can be seen in the example of **Figure 15**.

4 Design of balancer networks

We first define formally the simple balancer of order k , for $k \geq 0$, whose recursive definition was suggested in **Section 1.3**, then prove that it is a balancer. For $k = 0$ and $k = 1$, the networks without splitter and with a single splitter respectively are universal balancers. We denote $\llbracket l, u \rrbracket$ the set of integers i with $l \leq i \leq u$.

Definition 12. For any integer $k \geq 2$, the *simple balancer of order k* is the network splitter $G = (I \uplus S \uplus O, E)$ where:

- ▷ $I := \{i_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$ and $O := \{o_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$,
- ▷ $S := \{s_{(l,j)} : (l,j) \in \llbracket 0, k-1 \rrbracket \times \llbracket 0, 2^{k-1} - 1 \rrbracket\}$,
- ▷ $E := \{i_j s_{(0,j/2)}, s_{(k-1,j/2)} o_j : j \in \llbracket 0, 2^k - 1 \rrbracket\} \cup \bigcup_{l \in \llbracket 0, k-2 \rrbracket} E_l$,
- ▷ $E_l := \left\{ (s_{(l,j)} s_{(l+1,j)}, s_{(l,j)} s_{(l+1, j \oplus 2^l)}) : j \in \llbracket 0, 2^{k-1} - 1 \rrbracket \right\}$.

where \oplus is the bitwise exclusive or.

Proposition 2. *For any $k \geq 1$, the simple balancer of order k is a balancer.*

Proof. Let $G = (I \uplus S \uplus O, E)$ be the simple balancer of order k , $c : I \uplus O \rightarrow [0, 1]$ with $c(o) = 1$ for each output $o \in O$. For each arc $e = s_{(l,j)} s_{(l+1,j')}$, we define $J_{(l,j)}$ to be the set of inputs i such that there is a directed path from i to e . Clearly, we have that $J_{(l,j)} = J_{(l-1,j)} \uplus J_{(l-1, j \oplus 2^{l-1})}$ and $J_{(l,j)}$ is the set of integers j' such that the binary representations of j and $j'/2$ coincides on the $k-1-l$ heavier bits. Then, set

$$t(s_{(l,j)}) := \frac{1}{2^{l+1}} \sum_{\alpha \in J_{(l,j)}} c(i_\alpha)$$

and $F := E$. It can be easily checked that (t, F) is a steady-state where each output has the same throughput. \square

4.1 Throughput-unlimited balancers

We now define formally the Beneš network, and prove its property of being throughput-unlimited.

Definition 13. A *Beneš network* of order k is a splitter network $G = (I \uplus S \uplus O, E)$ where

- ▷ $I := \{i_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$ and $O := \{o_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$,
- ▷ $S := \{s_{(l,j)} : l \in \llbracket 0, 2k-2 \rrbracket, j \in \llbracket 0, 2^{k-1} - 1 \rrbracket\}$,
- ▷ $E := \{i_j s_{(0,j/2)}, s_{(2k-2,j/2)} o_j : j \in \llbracket 0, 2^k - 1 \rrbracket\} \cup \bigcup_{l \in \llbracket 0, 2k-3 \rrbracket} E_l$,
- ▷ $E_l := \left\{ (s_{(l,j)} s_{(l+1,j)}, s_{(l,j)} s_{(l+1, j \oplus 2^{l'})}) : j \in \llbracket 0, 2^{k-1} - 1 \rrbracket \right\}$ where $l' = \max\{k-2-l, l-k+1\}$.

The arcs between levels l and $l+1$ represent swapping a bit of weight w , where w varies from $k-2$ to 0 then to $k-2$ again. Alternatively, a recursive definition of a Beneš network of order $k+1$ consists in placing two Beneš networks of order k in parallel. Then identify the inputs from both networks, index by index, into a splitter, and proceed similarly for the outputs (see [Figure 7](#)). Because the latter half of a Beneš network is a simple balancer, it is itself a simple balancer.

Proposition 3. *Beneš networks are balancer.*

Proof. This follows from the fact that the subgraph of splitters $s_{(l,j)}$ for $l \in \llbracket k-1, 2k-2 \rrbracket$ with their incident arcs is a simple balancer. \square

Proposition 4. *Beneš networks are throughput-unlimited.*

Proof. Let $G = (I \uplus S \uplus O, E)$ be the Beneš network of order k , and $c : I \uplus O \rightarrow [0, 1]$. We may assume that $c(I) \geq c(O)$ by the reversibility property of splitter networks [Lemma 10](#).

We define a new splitter network $G_L = (V_L, E_L)$, by removing O and all vertices $s_{(l,j)}$ with $l \geq k$, and adding 2^k new outputs O' , with arcs $s_{(k,j/2)} o'_j$ for each $j \in \llbracket 0, 2^k - 1 \rrbracket$. G_L is a simple balancer of order k between I and O' . Symmetrically we define $G_R = (V_R, E_R)$, a simple balancer of order k , by removing I and all vertices $s_{(l,j)}$ with $l < k-1$, and adding a set I' of new inputs.

Consider $c'_L : I \cup O' \rightarrow [0, 1]$ defined by $c'_L(I) = c(I)$ and $c'_L(O) = \mathbb{1}$. By [Proposition 2](#), there exists a steady-state (t'_L, F'_L) for G_L, c_L , such that for each $e \in \delta^-(O)$, $t'(e) = c(I)/2^k$, and $F'_L = E_L$.

Next consider $c_R : I' \cup O \rightarrow [0, 1]$ defined by $c_R(I') = 1$ and $c_R(O) = c(O)$. The reverse of G_R is a simple balancer of order k . Therefore by [Proposition 2](#), there exists a steady-state $(\overleftarrow{t}_R, \overleftarrow{F}_R)$ for \overleftarrow{G}_R, c_R , such that for each $\overleftarrow{e} \in \delta^-(I')$ of the reverse graph, $t(\overleftarrow{e}) = c(O)/2^k$ and $\overleftarrow{e} \in \overleftarrow{F}_R$. Reversing the network and the steady-state, by [Lemma 10](#), $(t_R : e \rightarrow \overleftarrow{t}_R(\overleftarrow{e}), F_R = E_R \setminus \overleftarrow{F}_R)$ is a steady-state in G_R . Furthermore, for each arc $e \in \delta^+(I')$, $t_R(e) = c(O)/2^k$ and $e \in E_R \setminus F_R$ is saturated.

We plan to build a steady-state using t'_L and t_R . However when $c(I) > c(O)$, this would lead to positive excess on splitters $s_{(k-1,j)}$. To avoid this, we decrease the throughputs on G_L with the following procedure. Let $c_L : I \uplus O' \rightarrow [0, 1]$ denote a capacity function with $c_L(I) = c(I)$ and $c_L(o) = c(O)/2^k$ for each $o \in O'$. Then define a throughput function \bar{t}_L , where $\bar{t}_L(e) = c_L(o)$ for each $o \in O'$ and $\{e\} = \delta^-(o)$, and $\bar{t}_L(e) = t'_L(e)$ for any other arc $e \notin \delta^-(O')$. Then (t'_L, E_L) is a pre-steady-state for (G_L, c_L) , with positive excess on splitters $s_{(k-1,j)}$ when $c(I) > c(O)$. Using the pre-steady-state algorithm and [Lemma 6](#), there exists a steady-state (t_L, F_L) for G_L, c_L . Furthermore, let $e \in \delta^-(O')$. If e is fluid, then by [Corollary 2](#), $c(e) = c(O)/2^k$. If e is saturated, then by rule [R4](#), $c(e) = c(O)/2^k$ also.

Finally, we define $t(e) = t_L(e)$ if $e \in E \cap E_L$, and $t(e) = t_R(e)$ if $e \in E \cap E_R$. Let $s = s_{(k,j)}$ for some $j \in \llbracket 0, 2^k - 1 \rrbracket$. Then $t(\delta^-(s)) = c(O)/2^{k-1} \geq t'(\delta^+(s)) = c(O)/2^{k-1}$. Let $F = F_L \cup F_R$. Then the arcs in $\delta^+(s)$ are saturated. Because rules [R6](#), [R7](#) and [R8](#) hold on s and on each splitter, (t, F) is a steady-state, with t uniform on $\delta^-(O)$. \square

4.2 Universal balancer

Before giving a design for universal balancer, we first define a network that behaves like a universal balancer as long as the input and output capacities are at most $1/2$.

Definition 14. The *half-universal network of order k* is a continous splitter network $G = (I \uplus S \uplus O, E)$ where

- ▷ $I := \{i_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$ and $O := \{o_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$;
- ▷ $S := S_B$;
- ▷ $E := \{uv \in E_B : u, v \in S_B\} \cup F \cup \{i_j s_{(0,j)}, s_{(2k,j)} o_j : j \in \llbracket 0, 2^k - 1 \rrbracket\}$;
- ▷ $F := \{s_{(2k,j)} s_{(0,j)} : j \in \llbracket 0, 2^k - 1 \rrbracket\}$;
- ▷ $(I_B \uplus S_B \uplus O_B, E_B)$ is the Beneš network of order $k + 1$, and $S_B = \{s_{(l,j)} : l \in \llbracket 0, 2k \rrbracket, j \in \llbracket 0, 2^{k+1} - 1 \rrbracket\}$.

See the left side of [Figure 16](#). The half-universal network is its own reverse. Its design includes loopback connections from each output back to each input. When the throughput of an output reaches its capacity, any excess throughput can be redirected back through one of these loopback arc, to the entry of the network. From there, it flows again to different outputs. Hence, the loopback arcs prevent the occurrence of saturated arcs inside the Beneš network, as long as at least one output has not reached its maximum capacity. Consequently, the balancing property of the Beneš network is also preserved. This is done at the cost of reducing by half the throughput, as we send back half of the flow. This limitation justifies the name half-universal, and our universal balancer will use two half-universal networks in parallel.

Proposition 5. Let $G = (I \uplus S \uplus O, E)$ be the half-universal splitter network, and let $c : I \uplus O \rightarrow [0, 1/2]$. Then there is a steady-state (t, F) for (G, c) and $\alpha, \beta \in \mathbb{R}_{\geq 0}$ such that

- (i) for each input i , $t(\delta^+(i)) = \min\{c(i), \alpha\}$,
- (ii) for each output o , $t(\delta^-(o)) = \min\{c(o), \beta\}$.
- (iii) the total throughput $T := t(\delta^+(I))$ equals $\min\{c(I), c(O)\}$.

Proof. We may assume that $c(I) \leq c(O)$ by [Lemma 10](#), because the half-universal network is its own reverse. We build a steady-state with the stated properties. Let $\beta \geq 0$ be such that $\sum_{o \in O} \min\{c(o), \beta\} = c(I)$. Observe that $\beta \leq \max_{o \in O} c(o) \leq \frac{1}{2}$. Let f_j be the loopback arc $s_{(2k,j)} s_{(0,j)}$ and $J' := \{j \in \llbracket 0, 2^k - 1 \rrbracket : c(o_j) < \beta\}$.

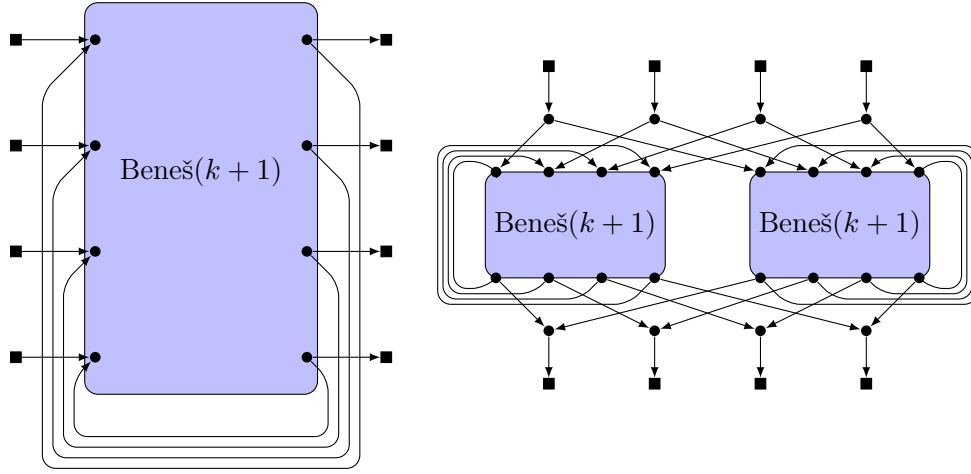


Figure 16: A schematic representation of a half-universal network on the left, and of a universal network on the right.

For $j \in \llbracket 0, 2^k - 1 \rrbracket$, set

$$\begin{aligned} t(\delta^-(o_j)) &:= \min\{c(o_j), \beta\}, \\ t(\delta^+(i_j)) &:= c(i_j), \\ t(f_j) &:= 2\beta - \min\{c(o_j), \beta\} \leq 1. \end{aligned}$$

Let (t^*, E_B) be the steady-state of the Beneš network of order k , when the input capacities are given by the values of $t(\delta^+(i_j))$ and $t(f_j)$ (for $j \in \llbracket 0, 2^k - 1 \rrbracket$) and the output capacities are uniformly 1. For any arc e in the Beneš subnetwork, set $t(e) := t^*(e)$.

We claim that $(t, E \setminus \{\delta^-(o_j) : j \in J'\})$ is a steady-state. Indeed, by [Proposition 4](#), the flow going through the Beneš subnetwork is

$$t(I) + t(F) = c(I) + \sum_{o \in O} 2\beta - \min\{c(o), \beta\} = \sum_{o \in O} 2\beta.$$

By [Proposition 3](#), the flow entering any node $s_{(2^k, j)}$ must be 2β , which equals the leaving flow. Moreover, [Rule R7](#) is clearly satisfied on node $s_{(2^k, j)}$ as $t(f_j) \geq t(\delta^-(o_j))$, concluding the proof. \square

Definition 15. The *universal network of order k* is the splitter network built from two disjoint copies of the half universal network of order k , by pairwise identifying the inputs and outputs from both copies, and adding 2^k input nodes and 2^k output nodes, each with one arc linking it to one of the identified vertices.

The right side of [Figure 16](#) illustrates the construction of a half-universal network. Again, universal networks are their own reverse.

Theorem 6. *The universal network of order k is universally balancing.*

Proof. Let $c : I \uplus O \rightarrow [0, 1]$ be a capacity function on the universal balancer G of order k . Consider a half-universal balancer $H = (I_H \uplus S_H \uplus O_H, E_H)$ of order $k - 1$, let $c_H : I_H \uplus O_H \rightarrow \llbracket 0, \frac{1}{2} \rrbracket$ be a capacity function defined by $c_H(i_j) = \frac{1}{2}c(i_j)$ and $c_H(o_j) = \frac{1}{2}c(o_j)$. Let (t_H, F_H) a steady-state for (H, c_H) given by [Proposition 5](#). Then one can define a steady-state (t, F) for (G, c) by using t_H on both sides of G and completing t to the arcs incident to output and input nodes, and easily check that it has the expected properties. \square




Figure 17: Adding loops on a simple balancer to reduce the number of inputs and outputs does not necessarily preserve the balancing property.

4.3 Balancers of all sizes

We extend the construction of balancers to general (n_i, n_o) -balancers. Given n_i and n_o , two positive integers representing respectively the number of inputs and outputs, let $k \in \mathbb{N}$ such that $m := \max\{n_i, n_o\} \leq 2^k$. Let $G = (I \uplus S \uplus O, E)$ be the universal network with 2^k inputs and 2^k outputs. Remove any surplus inputs or outputs to achieve a network with n_i inputs and n_o outputs. This can be accomplished by setting their capacities to 0. We obtain a universally balancing network with n_i inputs and n_o outputs.

Another method to decrease the number of outputs and inputs would be to add loopback arcs from the surplus output to the surplus input. This would avoid the usage of dummy terminals. Loopback arcs are also useful to reduce the number of inputs and outputs in non-universal balancer. However, preserving the balancer property when adding loopback arcs requires some care. Indeed, in a simple balancer, balancing is achieved when the capacities of the outputs are uniformly 1. Therefore it is necessary to ensure that the loopback arcs remain fluid. A careless implementation of this process may result in a network failing to be a balancer, as illustrated by [Figure 17](#).

5 Lower bounds for the size of balancers

We begin this section by establishing [Proposition 1](#), followed by the derivation of a lower bound on the number of splitters in a balancer.

Proof of [Proposition 1](#). This follows directly from the definitions for $S(k)$ and $B(k)$. A universal network of order k consists in its two Beneš subnetworks of order $k + 1$, plus $2 \cdot 2^k$ additional splitters to connect the inputs and outputs to each half-universal network. In total, we get

$$U(k) = 2 \cdot (2k + 1)2^k + 2 \cdot 2^k = (k + 1) \cdot 2^{k+2} \quad \square$$

As a consequence, the number of splitters of any balancer discussed in [Section 4](#) is $O(n \log n)$ where $n = \max\{|I|, |O|\}$. We proceed to establish a corresponding lower bound.

Proof of [Theorem 4](#). Let $G = (I \uplus S \uplus O, E)$ be a balancer network. Let $c : I \uplus O \rightarrow [0, 1]$ be a capacity function. We initialize a capacity c' with $c'(i) = 0$ for each $i \in I$ and $c'(o) = c(o)$ for each $o \in O$. We then incrementally increase the capacity of each input i , from 0 to $c(i)$, recalculating a new steady-state at each step. Increasing an input capacity transforms a steady-state into a sub-steady-state, therefore each

iteration performs a series of augmentations. Denote by f_1, \dots, f_m the sequence of augmenting circulations computed until reaching the final steady-state.

As G is a balancer, and each intermediary sub-steady-state is a steady-state for some choice of input capacities, each augmenting circulation f_k is uniform on $\delta^-(O)$. By construction f_k is non-zero on at most one input arc $e_k \in \delta^+(I)$. The support of f_k is defined as the subgraph of all arcs e such that $f_k(e) > 0$. We will assume that the support of each circulation intersects $\delta^+(I)$ (and thus increases the global throughput), as other circulations will not contribute to the proof and thus can be ignored.

Consider the support of f_k on G . Reverse all saturated arcs, to ensure that f_k is a flow on this graph. Then contract each arc $uv \neq e_k$ with $|\delta^+(u)| = 1$, by identifying u and v and removing uv . This action eliminates all the vertices with out-degree one, except for one input i_k . The resulting graph is denoted H_k , and the flow g_k , the restriction of f_k on H_k , represents a flow from i_k to O . Let δ_k^+ and δ_k^- denote the incidence functions of H_k . The following properties hold:

- Claim 2.** (i) each remaining splitter s in H_k has out-degree 2;
(ii) for each vertex v in H_k , g_k is constant on $\delta_k^+(v)$;
(iii) g_k is constant on $\delta_k^-(O) \cap F$.

Proof. (i) and (ii) follow from the fact that each vertex other than i_k in the residual graph has out-degree at most 2 and its outgoing arcs are coupled. This last fact follows by a simple case analysis on which incident arcs are saturated.

By construction, f_k is an augmenting circulation between two steady-states of a balancer. Consequently, f_k is the difference of two throughput functions that are both uniform on $\delta^-(O)$. Therefore f_k is itself uniform on fluid arcs of $\delta^-(O)$, and so is g_k , proving (iii). \square

We proceed by constructing a directed arborescence rooted at i_k , which may be infinite. This is achieved by establishing a parent-child relation among the walks originating from i_k in H . A walk w is a parent of a walk w' if the length of w' is one plus the length of w , and w is a prefix of w' . This defines an arborescence (T, E_T) whose root r is the empty walk and each node is a walk. Additionally, there is a natural morphism ϕ from each walk $w \in T$ to its end vertex, mapping each parent-child pair to a directed path in H from $\phi(\text{parent})$ to $\phi(\text{child})$. Observe that nodes in the arborescence, whose end vertices are splitters, possess two children. In contrast, nodes whose end vertices are outputs have no children. Let us introduce the function $p : E_T \rightarrow [0, 1]$, where $p(w, w')$ denotes the probability of a random walk originating from i_k having w' as a prefix. Specifically, because the tree is binary, $p(w, w') = 2^{-\text{depth}(w)}$. For a splitter s , let W_s be the set of all walks whose end vertex is s . The expected number of occurrences of s in a random walk starting from i_k , is $\Lambda(s) = \sum_{w' \in W_s} p(ww')$.

Using p , we construct a flow on the graph H_k . For each arc a , let $\phi^{-1}(a) \subseteq E_T$ be the set of edges e in the arborescence such that $a \in \phi(e)$. Then the flow value $f'_k(a)$ on an arc a is defined by $f'_k(a) = \sum_{e \in \phi^{-1}(a)} p(e)$. Due to its construction, f'_k satisfies the conservation rules on each splitter, hence is a flow from i_k to O . Observe that the augmenting flows are defined by a linear system. For each node $v \in S \cup \{i_k\}$ of H_k , this linear system contains one variable, representing the amount of flow on each outgoing arc, and one constraint, representing the conservation rule. The constraints are linearly dependant, as their sum is zero, but have rank one less than the number of variables. Hence the solution space has dimension 1. Therefore f_k and f'_k are identical up to a scaling factor: $f_k = f_k(e_k) \cdot f'_k$. Because f_k is uniform on $\delta^-(O)$, so is f'_k . This implies that (T, E_T) is the binary decision tree of a sampling process over the uniform distribution on O . By [Theorem 2](#), in the residual graph,

$$\sum_{s \in S} \sum_{e \in \delta_k^-(s)} f_k(e) = f_k(e_k) \sum_{s \in S} \Lambda(s) \geq f_k(e_k) |O| \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k(|O|).$$

Summing over all augmenting circulation f_1, \dots, f_m , we obtain:

$$\sum_{k=1}^m \sum_{s \in S} \sum_{e \in \delta_k^-(s)} f_k(e) \geq |I||O| \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k(|O|).$$

We conclude the proof by analysing the contribution of each arc to the left-hand side of the inequality. Because of [Corollary 3](#), a fluid arc entering a splitter contributes up to a value of 1, until it reaches a throughput of 1. Then it may contribute as a saturated arc to the splitter from which it is a leaving arc, again by a total amount of at most 1, until it reaches a throughput of 0. Each leaving arc of a splitter contributes at most 1 when saturated, and each entering arc of a splitter contributes at most 1 when fluid. Therefore, the total contribution is at most $4|S|$, yielding:

$$|S| \geq \frac{1}{4} |I||O| \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k(|O|).$$

In the case when the final steady-state contains only fluid arcs, we do not count the contribution of saturated arcs. This leads to [Theorem 3](#):

$$|S| \geq \frac{1}{2} |I||O| \sum_{k \in \mathbb{N}} \frac{k}{2^k} \text{binary}_k(|O|). \quad \square$$

6 Simulating an arbitrary capacity

We present a way to build, for any rational $r \in [0, 1] \cap \mathbb{Q}$, a continuous splitter network with one input vertex i and one output vertex o , such that, when $c(i) = 1$, then the throughput of any equilibrium of this network is $\min\{r, c(o)\}$. Those networks can be used as gadgets to simulate arcs with rational capacities in larger networks, thanks to the additional property that the arc leaving i may be set as saturated if and only if its throughput is r or $c(o)$.

Let $r = \frac{p}{q} \in \mathbb{Q}$ be a positive rational with $0 < p < q$. Let $k \in \mathbb{N}$ be such that $2^{k-1} < q \leq 2^k$, $k \geq 1$.

We construct a splitter network from a complete binary out-arborescence of depth k from a source splitter s . The throughput entering s is fairly split among the 2^k leaves. Then we partition the leaves into three sets P, Q, R , with $|P| = p$, $|Q| = q - p$ and $|R| = 2^k - q$. Route all the flow from R back to the root of the tree, by adding an in-arborescence from these leaves to s . Then a random walk from s ends in P with probability $\frac{p}{q}$, and in Q with probability $1 - \frac{p}{q}$. If we send all the flow from P to an output, and all the flow from Q back to the input, as illustrated in [Figure 18](#), we get a splitter network with maximal throughput r .

Informally, the arc s' entering the root of the main arborescence serves as the bottleneck of that network. Its throughput, when capped at 1, equals the sum of the input throughput and the throughputs from $q-p$ leaves. By the rule of conservation [R5](#), the input throughput equals the output throughput, which is the sum of throughputs of p leaves. Therefore the throughput x at each leaf satisfies $px + (q-p)x = t(ss')$, that is $x = \frac{t(ss')}{q}$. It follows that, when $t(ss')$ is capped at 1, the input throughput is $\frac{p}{q} = r$. We do not prove precisely this result because this construction has two main shortcomings:

- ▷ it only holds as long as $\frac{p}{q} \geq \frac{1}{2}$. Otherwise, the tree collecting the flow back from the $p - q$ leaves to the input would become saturated, while we would rather have that the arc is' is saturated.
- ▷ the size of this construction is linear in q (hence exponential in the size of the encoding of the capacities).

Optimizing over that solution would overcome these shortcomings, but we will rather directly define a correct network with an optimal number of splitters (up to a multiplicative constant). Before introducing an efficient design, we consider another construction based on the binary expansion of r . Because $r = \frac{p}{q}$ is rational, its binary expansion is ultimately periodic: there exist binary words $x, y \in \{0, 1\}^*$ with $\text{binary}(r) = 0.xy^\omega$. Let $l = |x| + |y|$. We define the splitter network $G'_r := (V := I \uplus S \uplus O, E := E_0 \cup E_r \cup E_{1-r})$ with




Figure 18: An inefficient network to simulate a rational capacity of $\frac{p}{q}$, where $2^{k-1} < q \leq 2^k$. This only works when $\frac{p}{q} \leq 1 - \frac{p}{q}$, because we expect the arc leaving the input node to saturate as soon as the throughput reaches $\frac{p}{q}$.

- ▷ $I := \{i\}$, $O := \{o\}$,
- ▷ $S := \bigcup_{j=1}^l \{u_j, v_j, w_j\}$,
- ▷ $E_0 := \{iv_l, v_1u_1, w_lo, u_lu_{|y|+1}\} \cup \bigcup_{j=1}^{l-1} \{u_ju_{j+1}, v_{j+1}v_j, w_jw_{j+1}\}$,
- ▷ $E_r := \{u_iv_i : \text{for each } i \in \llbracket 1, l \rrbracket \text{ with } (xy)_i = 0\}$,
- ▷ $E_{1-r} := \{u_iv_i : \text{for each } i \in \llbracket 1, l \rrbracket \text{ with } (xy)_i = 1\}$.

Then G_r is obtained by contracting any splitter s with $d^+(s) = d^-(s) = 1$ into a single arc. Figure 19 shows an example of network G_r with $r = \frac{169}{504}$. One way to understand this design is to unfold the loop made by $u_{l-1}u_{|y|+1}$, making the u , v and w lines extend infinitely to the right; this gives a network that simply split the flow following the binary expansion of r , each vertex u_i distributing 2^{-i} unit of flow. This construction seems more compact than the construction based on trees, but can still have a large number of splitters, as l can be as large as $q-1$. Also, it still needs to be adapted for some values of r to make sure that the only saturated arcs are the arcs $v_{i+1}v_i$. Namely, it is sufficient to choose a periodic representation where the period starts with a 1.

Our best construction relies on combining the main ideas of the two previous constructions:

- split the flow into 2^k equal chunks, and make three groups of size p , $q-p$, 2^k-q , so that we get three flows of size $\lambda p 2^{-k}$, $\lambda(q-p)2^{-k}$ and $1 - q2^{-k}$. One flow is output, one loops back to the head of the bottleneck arc, and one loops back to the tail of the bottleneck arc.
- use the binary representations of each group size to create the three flows.

We first give the construction for $r = \frac{p}{q} \geq \frac{1}{2}$, and we will later show how to extend it to any rational value.

For an integer $n \in \mathbb{N}$, denote $\text{binary}(n)$ the binary representation of n , and $\text{binary}_i(n)$ the value of the bit of weight 2^i in that representation.

Definition 16. Let $p_1, p_2, \dots, p_l, k \in \mathbb{N}_{>0}$ be integers with $\sum_{i=1}^l p_i = 2^k$. Let $T_{p_1, p_2, \dots, p_l, k}$ be a binary out-arborescence where each leaf is labeled by an integer in $\llbracket 1, l \rrbracket$, such that for any label $i \in \llbracket 1, l \rrbracket$ and any depth $d \in \llbracket 1, k \rrbracket$, the number of leaves with label i at depth d is $\text{binary}_{k-d}(p_i)$.

Notice that $T_{p_1, p_2, \dots, p_l, k}$ always exists, but is not necessarily unique. We prove its existence by showing

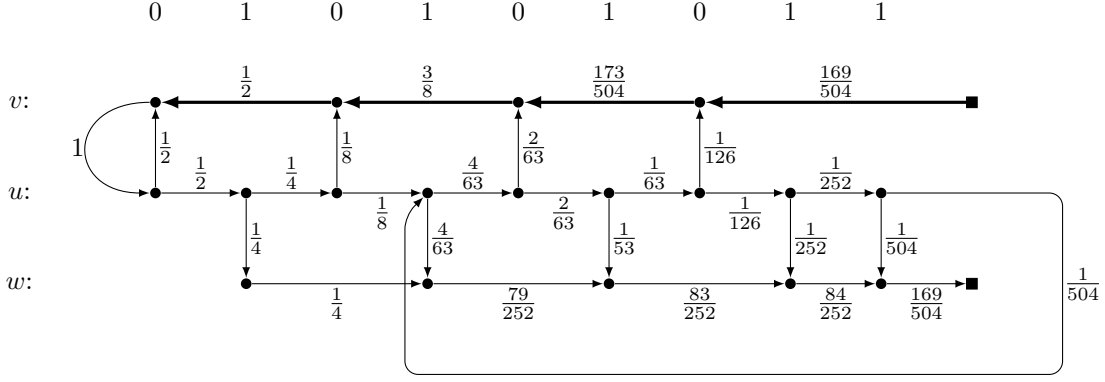


Figure 19: A network with maximal throughput $r = \frac{169}{504} = \frac{169}{8 \cdot 63}$, built from the binary representation of r : $0.010(101011)^\omega$.

that the number n_d of inner nodes at depth $d \in \llbracket 0, k \rrbracket$ is

$$n_d = 2^d - \sum_{i=1}^l \left\lfloor \frac{p_i}{2^{k-d}} \right\rfloor.$$

The proof is by induction on d , where $n_0 = 1$. Let $d \in \llbracket 1, k \rrbracket$, and assume that $n_{d-1} = 2^{d-1} - \sum_{i=1}^l \left\lfloor \frac{p_i}{2^{k-d+1}} \right\rfloor$. Then the number of nodes at depth d is twice n_{d-1} , to which we subtract the leaves at depth d . This yields

$$\begin{aligned} n_d &= 2n^{d-1} - \sum_{i=1}^l \text{binary}_{k-d}(p_i) \\ &= 2 \cdot \left(2^{d-1} - \sum_{i=1}^l \left\lfloor \frac{p_i}{2^{n-d+1}} \right\rfloor \right) - \sum_{i=1}^l \text{binary}_{k-d}(p_i) \\ &= 2^d - \sum_{i=1}^l 2 \left\lfloor \frac{p_i}{2^{n-d+1}} \right\rfloor + \text{binary}_{k-d}(p_i) \\ &= 2^d - \sum_{i=1}^l \left\lfloor \frac{p_i}{2^{k-d}} \right\rfloor. \end{aligned}$$

As $\sum_{i=1}^l p_i \leq 2^k$, n_d is non-negative, hence $T_{p_1, p_2, \dots, p_l, k}$ exists.

Definition 17. Given $r = \frac{p}{q} \in \mathbb{Q}$, with $\frac{1}{2} \leq \frac{p}{q} < 1$, $\gcd(p, q) = 1$, and $k \in \mathbb{N}_{>0}$ such that $2^{k-1} < q \leq 2^k$. Let $T = T_{p, q-p, 2^k-q, k}$ with root r . We define the splitter network $G_{p/q} := (V := I \uplus S \uplus O, E)$ where

- ▷ $I := \{i\}$, $O = \{o\}$ and $S = \{r'\} \uplus V(T)$,
- ▷ $E := E(T) \uplus E(P_1) \uplus E(P_2) \uplus E(P_3) \uplus \{ir', r'r, t_1o, t_2r', t_3r\}$,
- ▷ for each label $\gamma \in \{1, 2, 3\}$, P_γ is a path connecting all the leaves with label γ in increasing depth order, and t_γ is the end vertex of P_γ .

The construction is illustrated in [Figure 20](#) and [Figure 21 \(a\)](#).

Lemma 11. $G_{p/q}$ has a steady-state (t, F) with throughput $t^* = \min \left\{ c(i), c(o), \frac{p}{q} \right\}$. Moreover

- (i) if $t^* = \frac{p}{q}$, we may take $F = E(G_{p/q}) \setminus \{ii'\}$;
- (ii) if $t^* = c(i)$, we may take $F = E(G_{p/q})$;




Figure 20: A network of size $O(\log q)$ with maximum throughput $\frac{29}{55}$. The throughput values shown on the figure should be scaled by a factor $\frac{64}{55}$ to get the value at equilibrium. The lines P_1, P_2, P_3 are given by the binary encoding of $p = 29 = 011101$, $q = 26 = 011010$ and $2^k - q = 9 = 001001$. **Figure 21** (a) gives another representation of the same network.

(iii) if $t^* = c(o)$, we may take $\{ii', o'o\} \subseteq F$.

Proof. Let t be the following throughput function:

$$t(e) := \begin{cases} \frac{2^{k-d}}{q} & \text{if } e \in T \text{ is from depth } d-1 \text{ to depth } d, \\ \frac{2^{k+d}}{q} \left\lfloor \frac{v(\gamma)}{2^d} \right\rfloor & \text{if } e \in P_\gamma, \gamma \in \{1, 2, 3\}, \text{ is from depth } d \text{ to depth } d' > d, \\ \frac{p}{q} & \text{if } e \in \{ir', t_1o\}, \\ 1 - \frac{p}{q} & \text{if } e = t_2r' \\ \frac{2^k}{q} - 1 & \text{if } e = t_3r \\ 1 & \text{if } e = r'r \end{cases}$$

where $v(1) = p, v(2) = q - p, v(3) = 2^k - q$. We check that $(t, E \setminus \{ir'\})$ is a steady-state when $c(i) = c(o) = 1$. As $2^{k-1} < q \leq 2^k$, the throughputs are between 0 and 1. The conservation rule **R5** is easily checked on every vertex, except vertices of P_γ . Let $u \in V(P_\gamma)$ at depth d . let wu the incoming arc on P_γ , with w at depth d' . Then $\text{binary}_{d'}(v(\gamma)) = \text{binary}_d(v(\gamma)) = 1$, and $\text{binary}_\delta(v(\gamma)) = 0$ for each $\delta \in \llbracket d'+1, d-1 \rrbracket$. This implies

$$2^{d'-d} \cdot 2^{d'} \left\lfloor \frac{v(\gamma)}{2^{d'}} \right\rfloor + 1 = \frac{1}{2^d} \left\lfloor \frac{v(\gamma)}{2^d} \right\rfloor.$$

Scaling this equation by a factor $\frac{2^k}{q}$ yields the conservation rule for u . Rule **R6** at r' requires that $\frac{p}{q} \geq 1 - \frac{p}{q}$, which holds by the assumption that $\frac{p}{q} \geq \frac{1}{2}$. The other rules can be readily checked. This steady-state remains valid as long as $\min\{c(i), c(o)\} \geq \frac{p}{q}$.

Suppose now that $\min\{c(i), c(o)\} \leq \frac{p}{q}$. When $c(i) \geq c(o)$, we build a pre-steady-state by decreasing $t(t_1o)$ to $c(o)$, generating some excess on t_1 , and removing t_1o from the fluid arcs. Then using the pre-steady-state algorithm, the flow will be push back to i through ir' , when ir' is saturated. Therefore there is a steady-state with total throughput $c(o)$ and ir', t_1o are saturated. When $c(i) < c(o)$, we start from $(t, E \setminus \{ir'\})$, take the reverse steady-state in the reverse network, then decrease its output capacity on output i to $c(i)$, saturating $r'i$ and inducing some excess at r' . Then we apply the pre-steady-state algorithm to find a steady-state. The first step of the algorithm will be to make rr' saturated, in order to

push back the excess at r' . From there all arcs are saturated. Running the algorithm to its termination, we obtain a steady-state for the reverse network, that we reverse into a steady-state for $G_{p/q}$ with only fluid arcs. \square

When $\frac{p}{q} < \frac{1}{2}$, we construct $G_{p/q}$ from $G_{2p/q}$. This adds at most $2 \log q$ additional splitters, thus the number of splitters of $G_{p/q}$ remains $O(\log q)$.

Definition 18. Let $\frac{p}{q} \in \mathbb{Q}_{>0}$ with $\frac{p}{q} < \frac{1}{2}$, we define the splitter network $G_{p/q} := (V := I \uplus S \uplus O, E)$ where

- $I = \{i\}$, $O = \{o\}$, $S = V(H)$,
- $E = \{ii', o'o, o'i'\} \uplus E(H)$,

where H is a copy of $G_{2p/q}$ with input i' and output o' .

We extend the proof of Lemma 11 to every rational $\frac{p}{q}$.

Proof. We prove it by induction on the number of recursive steps in the definition of $G_{p/q}$. The base case is given by the first proof of the lemma, when $\frac{p}{q} \geq \frac{1}{2}$. Let $\frac{p}{q} < \frac{1}{2}$, and suppose that $G_{2p/q}$ behaves as expected.

Case 1: $\frac{p}{q} \leq \min\{c(i), c(o)\}$. Let (t', F') be a steady-state on $G_{2p/q}$ when $c'(i') = 1 = c'(o')$, obtained by induction hypothesis. Then $\delta^+(i') \cap F' = \emptyset$. Let $t(ii') = t(o'o) = t(o'i') = \frac{p}{q}$ and $t(e) = t'(e)$ for any other arc e . Then, checking all the rules on i , o , i' , and o' , we get that $(t, F' \cup \{ii', o'i'\})$ is a steady-state, as expected.

Case 2: $c(i) \leq c(o)$ and $c(i) < \frac{p}{q}$. Let (t', F') be a steady-state for $G_{2p/q}$ and $c'(i') = c'(o') = 2c(i) < \frac{1p}{q}$, obtained by induction hypothesis. Then $F' = E(G_{2p/q})$, all arcs are fluid. We extend this steady-state, by setting $t(ii') = t(o'o) = t(o'i') = c(i)$ and $t(e) = t'(e)$ for any other arcs. Then $(t, E(G_{2p/q}))$ is a steady-state, with only fluid arcs.

Case 3: $c(o) < \min\{c(i), \frac{p}{q}\}$. Let (t', F') be a steady-state on $G_{2p/q}$ when $c'(i') = c'(o') = 2c(o) \leq \frac{2p}{q}$, obtained by induction hypothesis. Then $\delta^-(o')$ and $\delta^+(i')$ are saturated in this steady-state. We extend it by setting $t(ii') = t(o'o) = t(o'i') = c(o)$, $t(e) = t'(e)$ for any other arc e , and $F = F'$. Then (t, F) is a steady-state for (G, c) , with ii' and $o'o$ saturated, as expected. \square

The next theorem summarizes the previous results.

Theorem 7. For any rational $r = \frac{p}{q} \in \mathbb{Q}$ with $r \in [0, 1]$, there is a splitter network with maximum throughput $\frac{p}{q}$. This network may be used as a gadget to simulate an arc with capacity r . The size of this network is linear in the size of r when r is encoded either as a ratio or by its binary expansion (prefix and period).

The network $G_{p/q}$ exhibits a curious asymmetry. It operates by splitting the input flow carefully until reaching the exact throughput $\frac{p}{q}$. Intuitively $G_{p/q}$ acts on the input. The reversed network $\overleftarrow{G_{p/q}}$ is also a capacity-simulating network, with the same maximum throughput. Intuitively $\overleftarrow{G_{p/q}}$ acts on the output, by carefully grabbing amount of flows until reaching throughput $\frac{p}{q}$. $G_{p/q}$ and $\overleftarrow{G_{p/q}}$ share the same general behaviour, with two distinct strategies. Figure 21 shows an example of $G_{p/q}$ together with its reverse networks and the throughputs at maximum capacity.

We cannot extend these constructions to build network with an irrational maximum throughput.

Lemma 12. For any irrational $x \in [0, 1] \setminus \mathbb{Q}$, there is no splitter network with a single input i and a single output o , such that $c(i) = c(o) = 1$, the total throughput of a steady-state is x .

Proof. The maximal throughput t of a network under some capacities is an optimal solution to the linear program (PSS) for some set F of fluid arcs, hence t must be a convex combination of rational vectors. By maximality of t , all those rational vectors have the same total throughput, which is a rational. \square




Figure 21: Two networks with maximum throughput $\frac{29}{55}$, reverse from each other.

We conclude this section, by mentioning that the construction of a capacity-simulating network can be extended to design networks that distribute their incoming flow over their outputs following an arbitrary rational distribution. Considering a rational vector (p_1, \dots, p_l) , and $q \in \mathbb{N}$, where l is the required number of outputs and $\frac{p_o}{q}$ is the expected throughput for output o . Choose k with $2^{k-1} < q \leq 2^k$. We replace the single edge bottleneck by a bottleneck cut of size $s = \lceil \frac{p_1 + \dots + p_l}{q} \rceil$. Then we define $\frac{r_i}{q}$ for $i \in \llbracket 1, s \rrbracket$, the amount of flow that loop back after the bottleneck, and $\frac{q'}{q}$ the amount of flow that loop back before the bottleneck. We start from a binary out-arborescence $T = T_{p_1, \dots, p_l, q', r_1, \dots, r_s}$. As the maximum throughput may be larger than 1, we replicate the highest nodes of T , until each layer contains at least s nodes. For each of the s replicated roots, we add to incoming arc, one is in the bottleneck cut, the other is a loopback arc with throughput $\frac{r_i}{q}$. Figure 22 illustrates an example of distribution network. From there, it is possible to glue a simple balancer over the inputs to make the network throughput unlimited, and then to add even more loopback arcs, to make a *half-universal distribution network*, and duplicate it to define ultimately a *universal distribution network*. In such a network, the throughputs on fluid outputs arcs stay proportional to each other, following the required distribution.

7 Priority splitters

7.1 Definition and algorithms

Factorio allows players to assign priorities to splitters, altering their behavior. A splitter that prioritizes some output belt will send as much throughput to that belt, until saturating or reaching its capacity. Any excess flow will then be sent on the other output belt. Similarly, a splitter will pull its flow from the prioritized input belt, and use the other output belt only to complete its throughput.

We modify the definition of steady-states to accomodate priority splitters.

Definition 19. Let $G = (I \uplus S \uplus O, E)$ be a splitter network, $c : I \cup O \rightarrow [0, 1]$ a capacity function, $p^+, p^- : S \rightarrow E \cup \{\perp\}$ the so-called the output-priority and input-priority functions, such that for each $s \in S$, $p^+(s) \in \delta^+(s) \cup \{\perp\}$ and $p^-(s) \in \delta^-(s) \cup \{\perp\}$. A *steady-state* for (G, c, p^+, p^-) is a pair (t, F) where

P1 $t : E \rightarrow [0, 1]$ is the throughput function;




Figure 22: An splitter network defined from a rational vector $(\frac{27}{49}, \frac{33}{49}, \frac{11}{49}, \frac{37}{49})$, with $(r_1, r_2, r_3, r_4) = (49, 49, 11)$. The acyclic subgraph H , colored in blue, is obtained by replicating the highest layers of a tree. All blue arcs (in H) between two consecutive levels have equal throughput indicated on the right, except for the first two levels that appeared by replicating the top of the tree and form a simple balancer.

- P2 $F \subseteq E$ is the set of *fluid arcs*, $E \setminus F$ is the set of *saturated arcs*;
P3 for each $i \in I$ with $\delta^+(i) = \{e\}$, $t(e) \leq c(i)$ and moreover if $e \in F$ then $t(e) = c(i)$;
P4 for each $o \in O$ with $\delta^-(o) = \{e\}$, $t(e) \leq c(o)$ and moreover if $e \notin F$ then $t(e) = c(o)$;
P5 for each $s \in S$, with $\delta^-(s) = \{e_1, e_2\}$ and $\delta^+(s) = \{e_3, e_4\}$, $t(e_1) + t(e_2) = t(e_3) + t(e_4)$;
P6 for any $e_1, e_2 \in E$ with $\{e_1, e_2\} = \delta^-(s)$ and $e_1 \notin F$, if $p^-(s) = e_1$ then $t(e_1) = 1$ or $t(e_2) = 0$, and if $p^-(s) \neq e_2$ then $t(e_1) \geq t(e_2)$;
P7 for any $e_1, e_2 \in E$ with $\{e_1, e_2\} = \delta^+(s)$ and $e_1 \in F$, if $p^+(s) = e_1$ then $t(e_2) = 0$ or $t(e_1) = 1$, and if $p^+(s) \neq e_2$ then $t(e_1) \geq t(e_2)$;
P8 for any $uv \in E \setminus F$ and $vw \in F$, $t(uv) = 1$ or $t(vw) = 1$.

Compared to the definition of steady-states for (G, c) , all rules are identical except Rules **P6** and **P7** which generalizes Rules **R6** and **R7**. Therefore when p^- and p^+ are uniformly \perp , the two notions of steady-states coincide.

The extended definition of steady-state preserves the symmetry already observed, therefore if (t, F) is a steady-state for (G, c, p^+, p^-) , then $(t, \overleftarrow{E \setminus F})$ is a steady-state for $(\overleftarrow{G}, c, \overleftarrow{p}^-, \overleftarrow{p}^+)$, where $\overleftarrow{p}(e) = \overleftarrow{p}(e)$.

In the Rule **P7**, for $\delta^+(s) = \{e_1, e_2\}$, when $e_1 \in F$ and $p^+(s) = e_1$, the constraint $t(e_1) = 1$ or $t(e_2) = 0$ is not immediately representable as a linear constraint. However, in the context of iteratively solving the linear system, and removing arcs from F at each step, at the start of an iteration, if $e_1 \in F$ and $t(e_1) < 1$, we may force the constraint $t(e_2) = 0$. Hence $t(e_2)$ will stay null until e_1 becomes tight or saturated. Then we can remove the constraint $t(e_2) = 0$ as it is no longer induced by the definition of steady-state. The same principle applies to Rule **P6**. Therefore, the LP-based algorithms can be adapted to compute the steady-state of a splitter network with priorities. Likewise, the definition of residual graph can be adapted, to accomodate priorities.

Priorities allow to change the feasible steady-states. Therefore, the question of finding a steady-state optimizing the total throughput becomes relevant in the context of splitter network with priorities. Formally, the input of a throughput maximization problem consists in a splitter network (G, c) . Its output is a quadruple (p^+, p^-, t, F) such that (t, F) is a steady-state for (G, c, p^+, p^-) . The goal is to maximize the global throughput $t(\delta^+(I))$.

We study several possible scenarii for this optimization problem: in **Section 7.2** when we may choose all the priorities, in and out of each splitter; in **Section 7.3** we impose the in priorities and ask for maximizing out priorities; and finally in **Section 7.4** we impose some priorities arbitrarily, and must choose the other priorities.

7.2 Choose input and output priorities for each splitter

Proposition 6. *The throughput maximization problem is polynomial-time solvable when $c = \mathbb{1}$. Let $G = (I \uplus S \uplus O, E)$ be a splitter network. The maximum throughput equals the value of a maximum flow from I to O in G with unit capacities.*

Proof. Consider an integral maximum flow t from I to O . Because there is a minimum cut having the same value, the maximum throughput cannot exceed $t(\delta^+(I))$. We show that we can choose F , p^+ and p^- such that (t, F) is a steady-state in (G, c, p^+, p^-) .

For each splitter $s \in S$, with $\delta^+(s) = \{e_1, e_2\}$, we define

$$p^+(s) = \begin{cases} \perp & \text{if } t(e_1) = t(e_2); \\ e_1 & \text{if } t(e_1) = 0 \text{ and } t(e_2) = 0; \\ e_2 & \text{if } t(e_2) = 1 \text{ and } t(e_1) = 1. \end{cases}$$

Define $p^-(s)$ similarly. Finally, we define F by setting saturated arcs to be precisely the arcs e with $t(e) = 0$ that are reachable from I in the residual graph of the maximal flow t . This choice guarantees that Rule **P8** is satisfied. All the other rules are readily checked. \square




Figure 23: An illustration of the reduction from the partition problem to finding output priorities with maximal throughput.

7.3 Choose output priority for each splitter

Proposition 7. *Given $(G, c = \mathbb{1})$ and input priorities $p^- : S \rightarrow E$, the problem of finding a steady-state (t, F, p^+, p^-) with maximum throughput is polynomial-time solvable. The maximal throughput equals the value of a maximum flow from I to O in G with unit capacities.*

Proof. For an integral maximum flow t from I to O , and setting $p^+(s)$ as in the proof of [Proposition 6](#), (t, E) is a sub-steady-state, as can be readily checked. We initiate the sub-steady-state algorithm with (t, E) , yielding a steady-state (t', F) . Because the sub-steady-state cannot reduce the global throughput, and by the existence of a minimum cut of size $t(\delta^+(I))$, the throughput of (t', F) is precisely the value of the maximum flow. \square

Corollary 5. *Given $(G, c = \mathbb{1})$ and output priorities $p^+ : S \rightarrow E$, the problem of finding a steady-state (t, F, p^+, p^-) with maximum throughput is polynomial-time solvable. The maximal throughput equals the value of a maximum flow from I to O in G with unit capacities.*

Proof. Apply [Proposition 7](#) to the reverse splitter network. \square

In contrast, when we force each splitter to have an output priority distinct from \perp and the input capacities are not uniform, finding a maximal throughput is hard.

Proposition 8. *Let $G = (I \uplus S \uplus O, E)$ be a splitter network, capacities $c : I \cup O \rightarrow [0, 1]$ and input priorities p^- , with $c(o) = 1$ for each $o \in O$. The problem of finding output priorities $p^+ : S \rightarrow E$ and a steady-state (t, F, p^-, p^+) of maximum throughput is NP-hard, even when $p^-(s) = \perp$ for each splitter $s \in S$.*

Proof. We reduce the partition problem: given s_1, \dots, s_n with $\sum_{i=1}^n s_i = 2$, is there a subset $I \subseteq \llbracket 1, n \rrbracket$ with $\sum_{i \in I} s_i = 1$. The reduction is straightforward, and illustrated in [Figure 23](#). The n leftmost terminals have output capacities s_1, s_2, \dots, s_n respectively. s_0 and the outputs have capacity 1. The maximal throughput is 3 if and only if the partition instance has a solution. Indeed, when a partition exists, routing all the flows from terminals s_i with $i \in I$ up, and all the other flows down, provides a routing where no flow goes through u . Inversely, if any flow goes through u , it must go through v too. That flow will concur with the one unit flow between s_0 and t_0 , therefore forbidding the total throughput out of s_0 to be one. However, in order to avoid u to receive any flow, the output flow of each input s_i must be fairly split between the bottom and top parts of the splitter network, hereby defining a partition. \square




Figure 24: A network splitter with two capacity functions and their steady-states.

7.4 Choose output priorities for some splitters

When an arbitrary subset of priorities are imposed, the maximum throughput problem becomes intractable, stays so even when we only choose some of the output priorities and all other splitters act fairly.

Theorem 8. *Let $G = (I \uplus S \uplus O, E)$ be a splitter network, $c : I \cup O \rightarrow [0, 1]$ a capacity function, $p^- : S \rightarrow E \cup \{\perp\}$ input priorities for each splitter, and $p_0^+ : S' \rightarrow E \cup \{\perp\}$ be a partial output-priority function, where $S' \subseteq S$. The problem of finding a steady-state (t, F, p^-, p^+) of maximum throughput with p_0^+ being the restriction of p^+ to S' is NP-hard, even when p^- and p_0^+ are uniformly \perp and $c = \mathbb{1}$.*

The reduction is based on two gadgets, representing respectively the variables and the clauses of a 3-SAT formula. Before establishing the reduction, we study the two properties of the two gadgets. The variable gadget is based on the splitter network described in Figure 24. We will exploit the fact that its global throughput is not a convex function of the input capacities. Indeed, as depicted in Figure 24, for the input capacities $(1, 1, 0)$, $(\frac{1}{2}, 1, \frac{1}{2})$ and $(0, 0, 1)$, the global throughput are respectively $\frac{11}{4}$, $\frac{10}{4}$ and $\frac{11}{4}$. Therefore, by adding a splitter with undetermined priority, to provide the flow to the two opposite inputs, we incentivize a choice of priority different from \perp for that additional splitter. Implementing this modification yields the variable gadget, defined in Figure 25.

Proposition 9. *Depending on the choice of $p^+(x?)$, and assuming that each arc leaving the gadget is fluid, the throughput on the leaving arcs of the splitter network depicted in Figure 25 are either $(\frac{7}{8}, \frac{7}{8}, \frac{1}{2}, \frac{1}{2})$, or $(\frac{5}{8}, \frac{5}{8}, \frac{5}{8}, \frac{5}{8})$, or $(\frac{1}{2}, \frac{1}{2}, \frac{7}{8}, \frac{7}{8})$.*

Proof. This follows immediately from the steady-states depicted in Figure 24, considering the throughput on the leaving arc of $x?$ will be either 1, 0, or $\frac{1}{2}, \frac{1}{2}$, or 0, 1, depending on the choice of out-priority. \square




Figure 25: The variable gadget. The node $x?$ represents a splitter whose output priority is not predetermined. All other splitters s are fair: $p_0^+(s) = p^-(s) = \perp$. The leaving arcs indexed by x are connected to gadgets of clauses containing x positively, those indexed by \bar{x} are connected to gadgets of clauses containing x negatively. If there are not enough clauses, each superfluous arc has a new terminal as destination.

Hence, each variable gadget induces a choice, that translates into some arcs carrying a throughput of $\frac{7}{8}$ and others $\frac{1}{2}$. We propose a clause gadget that can accept a throughput of up to $\frac{9}{4} = \frac{1}{2} + \frac{7}{8} + \frac{7}{8}$, forbidding the three inputs from having all a throughput of $\frac{7}{8}$. The splitter network presented in Figure 26 features two parts: the upper part receives the flow from the variable gadget, and reduces it to an eighthth of its original value. Hopefully the remaining throughput should be at most $\frac{9}{32} < \frac{1}{2}$. The bottom part consists in a capacity simulating network, with capacity $1 - \frac{9}{32} = \frac{23}{32} > \frac{1}{2}$. Then we join this two flows toward the output t_1 in the splitter p . Increasing the input flow from the variable gadgets will increase the flow entering p from above. As $\frac{9}{32} < \frac{23}{32}$, by respecting Rule P6, the throughput on the arc entering p from the right would then decrease, and thus s_2 would provide strictly less than $\frac{23}{32}$.

Proposition 10. *For any throughput values f_1, f_2, f_3 on the three entering arcs of the clause gadget, depicted in Figure 26, there is a steady-state (t, F) for which the three entering arcs e_1, e_2, e_3 are fluid, with $t(e_i) = f_i$ for each $i \in \{1, 2, 3\}$. Moreover, $f_1 + f_2 + f_3 \leq \frac{9}{4}$ if and only if $t(\delta^+(s)) = \frac{23}{32}$.*

Proof. In order to show that the arcs entering the gadget are always fluid, we consider the case when the throughput on the three entering arcs e_1, e_2 and e_3 are uniformly 1. Then the throughput on each arc of the upper part of the gadget can be easily computed, and the throughput on the arc e entering p from above is $\frac{3}{8}$. This arc is fluid, and the other arc e' entering p from the right is saturated with throughput $\frac{5}{8}$. Therefore e_1, e_2 and e_3 are fluid in this steady-state.

More precisely, the throughput on e is $t(e) = \frac{f_1 + f_2 + f_3}{8}$, and the throughput on e' is $t(e') = \max\{\frac{23}{32}, 1 - t(e)\}$, with e' saturated when $t(e) > \frac{9}{32}$. The steady-state at the critical point, when $f_1 + f_2 + f_3 = \frac{9}{4}$, is illustrated in Figure 26. \square

Using the two gadgets, we prove that the throughput maximization problem is NP-hard.

Proof of Theorem 8. We reduce 3SAT, where each variable appears at most twice positively and twice negatively. Let ϕ be a set of 3-clauses on variables x_1, \dots, x_n . We define a splitter network, with one copy of the variable gadget for each variable in the formula, and one copy of the clause gadget for each of its clauses. A leaving arc labeled x (respectively \bar{x}) is identified with an entering arc of a clause gadget containing the literal x (respectively \bar{x}). Any superfluous leaving arc is redirected to a new distinct output node. The priorities of every splitter are set to \perp , except for the output priority of the splitter $x?$ from the variable gadget, for each variable x . The output priority of that splitter is free. Formally

much closer to the lower bound in regards to the number of fair splitters. Unfortunately, it also contains many splitters with non-fair output priorities, making it asymptotically larger than the simple balancer. The construction is nonetheless interesting, not only for its smaller number of fair splitters, but also in demonstrating that saturation may also be used in an effective way to balance the outputs.

The saturating $(2^k, 2^k)$ -balancer is composed of three parts, illustrated in [Figure 27](#). The first part contains the input and a half-grid of priority splitters, whose output priority are set to the same direction (vertical in the picture). This design ensures that the flow is pushed as much as possible to the top of the grid. Therefore the throughputs on the arcs leaving the half-grid, from top to bottom, are $1, 1, \dots, 1, t, 0, \dots, 0$, where $t = c(I) - \lfloor c(I) \rfloor$. In particular, if $c(I) \leq 2^{k-1}$, all the flow leaves the half grid from the top 2^{k-1} arcs, else the 2^{k-1} receive 2^{k-1} units of flow.

The next part is made of a simple balancer of order 2^{k-1} from the 2^{k-1} highest rows to the 2^{k-1} lowest rows. The last part is a column of 2^{k-1} fair splitters, whose out-neighbors consist in the network's outputs. From the outputs of the simple balancer to those fair splitters, 2^{k-1} arcs serves as the main bottleneck of this network.

The saturating balancer works under two possible regimes. When $c(I) \leq 2^{k-1}$, all the flow is pushed in the simple balancer, arrives balanced into the bottleneck arcs, and the is split evenly by the last column of splitters toward the output. Then every arc is fluid. The other regime happens when $c(I) > 2^{k-1}$. At $c(I) = 2^{k-1}$, all the arcs in the simple balancer and in the bottleneck reach throughput 1. Increasing the input capacity above 2^{k-1} , the arcs of the simple balancer saturates. The excess flow is pushed into the 2^{k-1} lowest rows, directly toward the bottleneck arcs. Because the bottleneck is already full, the incoming flow is pull backed by the simple balancer. At this point, the simple balancer, which is fully saturated, acts a balancer on the pulled back flow: the pulled back flow is sent back evenly to its input, from which it proceeds on the horizontal arcs of the 2^{k-1} highest rows, to the last column of splitters. Thus, the potential of each fair splitter in the simple balancer is fully used: as long as $c(I) < 2^{k-1}$ it splits the throughput fairly into its two outgoing fluid arcs. When $c(I) > 2^{k-1}$, it splits the pulled-back throughput evenly into its two incoming saturated arcs.

The saturating $(2^k, 2^k)$ -splitter contains $2^{k-1} + S(2^{k-1}) = (k+1)2^{k-2}$ fair splitters, which is closer to our lower bound of $k2^{k-2}$ splitters from [Theorem 4](#). The difference is explained by the last column of fair splitters, which are only used with fluid outgoing arcs. The total number of splitters is $\Theta(2^{2k})$, dominated by the half-grid. However, this part can be replaced by a network with the following property: the total throughput on the 2^{k-1} highest leaving arcs must be $\min\{c(I), 2^{k-1}\}$. Such a network can be defined with $o(2^{2k})$ priority splitters, but we do not know whether one exists with only $O(k2^k)$ priority splitters.

9 Perspectives

We formalized splitter networks and their steady-states, and presented various load-balancing designs. The ability to design universal balancers enables the simulation of networks with integral capacities: each arc is replicated according to its capacity, and each splitter is replaced by a universal balancer. A universal balancer is fair by the balancing property, and maximizing by the unlimited-throughput property, effectively generalizing splitters. Our definition of splitter network can also be extended to support arc capacities natively, with most of the proofs requiring only minor modifications.

Although our continuous model is convenient for modeling the expected throughput of splitter networks, Factorio's belt systems operates discretely. Therefore, the observed throughputs in Factorio's splitter networks are only approximations of those theorized by our model. Further investigation into the disparities between the discrete and continuous splitter networks is necessary to accurately apply our findings to Factorio.

We have let several questions unanswered. The most fundamental remains regarding the uniqueness of the steady-state throughput. While it is possible for a single splitter network to admit multiple steady-states, we have yet to encounter a network with two steady-states that yield different throughputs on their outputs.




Figure 27: A $(2^k, 2^k)$ -balancer utilizing priorities, which contains significantly less fair splitters than the simple balancer. In this example, we set $k = 4$. Fair splitters are represented by full nodes. In each priority splitter, a small arrows indicates which arcs are prioritized: the outgoing arc pointed to by the arrow's head, and the incoming arc aligned with the origin of the arrow. On the right, the half-grid pushes up to 2^{k-1} units of flow to the top rows. In blue, a simple balancer of order 2^{k-1} proceeds to most of the balancing. The red arcs form a bottleneck, when the total throughput exceeds 2^{k-1} . On the right, a last column of fair splitters distributes the flow to the outputs.

Our lower bounds for the number of splitters in balancers have a constant multiplicative gap across all designs, indicating they are not tight. For simple balancers of order k , this gap is closed when we forbid saturated arcs in the steady-state of the balancer. Consequently, leveraging saturation is necessary to further reduce the number of splitters in load-balancing networks. This is partially done in [Section 8](#), at the cost of introducing many priority splitters. Furthermore, it is worth investigating stronger lower bounds in the context of universal balancers.

Factorio allows splitters to be configured to prioritize either an outgoing arc, or an incoming arc. Utilizing this feature, the universal network described in [20] achieves a significantly smaller size compared to our design. Our technique still establishes a lower bound on the number of fair splitters. In general, what is the minimum size achievable for networks utilizing these more general splitters? We proved several complexity results for the problem of global throughput maximization, when we can choose which arcs to prioritize in each splitter or a subset of those splitters. However, the complexity of global throughput maximization when the input and output capacities are arbitrary is not fully settled yet.

As a last series of questions, consider a network whose steady-state, when all inputs and outputs have capacity 1, has no saturated arcs. If the augmenting flow from any single input is uniformly distributed across the outputs, then the network is a balancer. This provides a polynomial-time procedure for determining whether a network is a balancer, subject to the absence of saturation. Is it feasible to devise a general procedure to decide whether a splitter network is balancing, throughput unlimited or universal?

References

- [1] Václav E Beneš. “On rearrangeable three-stage connecting networks”. In: *The Bell System Technical Journal* 41.5 (1962), pp. 1481–1492.
- [2] Václav E Beneš. “Permutation groups, complexes, and rearrangeable connecting networks”. In: *Bell System Technical Journal* 43.4 (1964), pp. 1619–1640.
- [3] Aaron Bernstein, Maximilian Probst, and Christian Wulff-Nilsen. “Decremental strongly-connected components and single-source reachability in near-linear time”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 365–376.
- [4] BoardGameGeek. *Boardgame category: transportation*. <https://boardgamegeek.com/boardgamecategory/1011/transportation>.
- [5] Bonnie S Boardman and Caroline C Krejci. “Simulation of Production and Inventory Control using the Computer Game Factorio”. In: *ASEE 2021 Gulf-Southwest Annual Conference*. 2021.
- [6] Chase Covello, Hyunjang Jung, and Bryan C Watson. “Using graph theory to investigate the role of expertise on infrastructure evolution: A case study examining the game Factorio”. In: (2023).
- [7] Miguel Coviello Gonzalez and Marek Chrobak. “Towards a theory of mixing graphs: A characterization of perfect mixability”. In: *Theoretical Computer Science* 845 (2020), pp. 98–121. ISSN: 0304-3975.
- [8] Yefim Abramovich Dinitz. “The method of scaling and transportation problems”. In: *Studies in Discrete Mathematics, Moscow* (1973), pp. 46–57.
- [9] Shivam Duhan, Chengming Zhang, Wenyu Jing, and Mingqi Li. “Factory optimization using deep reinforcement learning AI”. In: (2019).
- [10] Andrew V Goldberg and Robert E Tarjan. “A new approach to the maximum-flow problem”. In: *Journal of the ACM (JACM)* 35.4 (1988), pp. 921–940.
- [11] Ketcheson David. *Mathematics Stackexchange: Belt Balancer problem (Factorio)*. <https://math.stackexchange.com/questions/1775378/belt-balancer-problem-factorio>.
- [12] D.E. Knuth. *The Art of Computer Programming: Sorting and Searching, Volume 3*. Pearson Education, 1998. ISBN: 9780321635785.

- [13] Donald Ervin Knuth and AC Yao. *The complexity of nonuniform random number generation*. *Algorithms and Complexity: New Directions and Recent Results*, ed. JF Traub. 1976.
- [14] Legnagi, Alessandro and Montini, Axel. *VeriFactory*. <https://github.com/alegnani/verifactory/>.
- [15] Andre Leue. “Verification of Factorio Belt Balancers using Petri Nets”. PhD thesis. Bachelorarbeit, Darmstadt, Technische Universität Darmstadt, 2021.
- [16] MatthaeusHarris. *Factorio belts are Turing-complete*. https://www.reddit.com/r/factorio/comments/lc25cx/factorio_belts_are_turing_complete/.
- [17] Carol A. Meyers and Andreas S. Schulz. “Integer equal flows”. In: *Operations Research Letters* 37.4 (2009), pp. 245–249. ISSN: 0167-6377.
- [18] Amandeep Parmar. “Integer programming approaches for equal-split network flow problems”. PhD thesis. Georgia Institute of Technology, 2007.
- [19] Sean Patterson, Joan Espasa, Mun See Chang, and Ruth Hoffmann. “Towards Automatic Design of Factorio Blueprints”. In: *arXiv preprint arXiv:2310.01505* (2023).
- [20] pocarski. *Universal 8-8: Perfectly Balanced, as All Things Should Be*. <https://web.archive.org/web/20230922022806/https://alt-f4.blog/ALTF4-27/>.
- [21] R-O-C-K-E-T. *factorio-SAT: Enhancing the Factorio experience with SAT solvers*. <https://github.com/R-O-C-K-E-T/Factorio-SAT>.
- [22] Kenneth N Reid, Iliya Miralavy, Stephen Kelly, Wolfgang Banzhaf, and Cedric Gondro. “The factory must grow: automation in Factorio”. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 2021, pp. 243–244.
- [23] K Srinathan, Pranava R Goundan, MVN Ashwin Kumar, R Nandakumar, and C Pandu Rangan. “Theory of equal-flows in networks”. In: *Computing and Combinatorics: 8th Annual International Conference, COCOON 2002 Singapore, August 15–17, 2002 Proceedings 8*. Springer. 2002, pp. 514–524.
- [24] Wube Software. *Factorio*. <https://www.factorio.com/>.