MAXIMUM FLOWS IN NETWORKS OF $\{0,1\}$-VALUED INFINITARY SUBMODULAR FUNCTIONS

GARRETT ERVIN

ABSTRACT. We characterize the $\{0,1\}$-valued submodular set functions f on a possibly infinite domain V in terms of pairs of filters on V. We show that such functions can be viewed as generalizations of edges in a directed graph, and prove the max-flow min-cut theorem for locally finite networks of such edges.

1. $\{0,1\}$-VALUED SUBMODULAR FUNCTIONS

Suppose V is a non-empty set. Let 2^V denote the powerset of V. A set function $f : 2^V \to \mathbb{R}$ is called *submodular* if for all $X, Y \subseteq V$ we have

$$f(X \cup Y) + f(X \cap Y) \leq f(X) + f(Y).$$

Submodular functions appear in many contexts and their theory has numerous applications, especially in combinatorial optimization. Typically when studying submodular functions the underlying set V is assumed to be finite. We are interested in studying submodular functions when V is infinite, though our results will apply to the case when V is finite as well.

For $X \subseteq V$ we let X^* denote the complement of X. Given a submodular function $f : 2^V \to \mathbb{R}$, the *dual function* $f^* : 2^V \to \mathbb{R}$ is defined by the rule

$$f^*(X) = f(X^*).$$

It is easy to verify that f^* is also submodular. Observe $(f^*)^* = f$.

A submodular function f on V is *increasing* if $X \subseteq Y$ implies $f(X) \leq f(Y)$, and *decreasing* if the inequality is reversed. Observe that f is increasing if and only if f^* is decreasing.

We are going to classify the $\{0,1\}$-valued submodular functions f on V. We first give some examples of such functions.

Examples 1.1.

1. The constant function $f \equiv 0$ is submodular.
2. The constant function $f \equiv 1$ is submodular.
3. Fix a non-empty subset $A \subseteq V$. Let 1_A denote the function defined by

$$\begin{aligned} 1_A(X) &= 1 \quad \text{if } X \cap A \neq \emptyset, \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

   Then 1_A is increasing and submodular.
4. Fix a non-empty subset $B \subseteq V$. Let 1_B^* denote the function defined by

$$\begin{aligned} 1_B^*(X) &= 1 \quad \text{if } B \not\subseteq X, \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

   Then 1_B^* is decreasing and submodular.

5. Suppose that A, B are non-empty subsets of V, not both equal to the same singleton. Let $1_{A\to B}$ be the function defined by

$$\begin{aligned} 1_{A\to B}(X) &= 1 \quad \text{if } X \cap A \neq \emptyset \text{ and } B \not\subseteq X \\ &= 0 \quad \text{otherwise.} \end{aligned}$$

Then $1_{A\to B}$ is submodular.

Proof. (1.) and (2.) are clear.

For (3.), fix $X, Y \subseteq V$, and consider $1_A(X)+1_A(Y)$ and $1_A(X\cap Y)+1_A(X\cup Y)$. If $1_A(X) = 1_A(Y) = 1$, then we always have $1_A(X)+1_A(Y) \geq 1_A(X\cap Y)+1_A(X\cup Y)$ since the right side of the inequality is at most 2. If exactly one of $1_A(X)$ and $1_A(Y)$ is 0, then one of X, Y does not intersect A, and hence neither does $X \cap Y$ so that $1_A(X \cap Y) = 0$, and it follows $1_A(X) + 1_A(Y) \geq 1_A(X \cap Y) + 1_A(X \cup Y)$ in this case as well, since both the left and right sides of the inequality are 1. If $1_A(X) = 1_A(Y) = 0$, then both X and Y miss A, and hence so do $X \cap Y$ and $X \cup Y$. It follows $1_A(X)+1_A(Y) \geq 1_A(X\cap Y)+1_A(X\cup Y)$ since both sides of the inequality in this case are 0. Hence 1_A is submodular. It is clearly increasing.

For (4.), notice that $1^*_B = (1_B)^*$.

For (5.), observe that if $1_{A\to B}(X)+1_{A\to B}(Y) = 1$ then one of X, Y either misses A or contains B. In the first case $X \cap Y$ also misses A, and in the second $X \cup Y$ contains B, so that $1_{A\to B}(X \cap Y) + 1_{A\to B}(X \cup Y)$ is at most 1. If $1_{A\to B}(X) + 1_{A\to B}(Y) = 0$, there are several possibilities. The first is that both X, Y miss A, the second that one of X, Y misses A and the other contains B, and the third is that both X, Y contain B. But in all of these cases, each of $X \cap Y$ and $X \cup Y$ either misses A or contains B, so that $1_{A\to B}(X \cap Y) + 1_{A\to B}(X \cup Y) = 0$ as well. Thus $1_{A\to B}(X) + 1_{A\to B}(Y) \geq 1_{A\to B}(X \cap Y) + 1_{A\to B}(X \cup Y)$ holds in every case. □

It will follow from our classification that if V is finite, then any submodular function f on V has exactly one of the five forms described in Examples 1.1. When V is infinite, the classification is similar, but the conditions of incidence with a set A and containment of a set B are replaced by positivity with respect to some filter $\mathcal{F}$ and membership in some filter $\mathcal{G}$.

Definition 1.2. A collection of subsets $\mathcal{F} \subseteq 2^V$ is a *filter* on V if the following hold:

i. (non-triviality) $V \in \mathcal{F}$, $\emptyset \notin \mathcal{F}$,
ii. (upward closure) if $A \subseteq B \subseteq V$ and $A \in \mathcal{F}$, then $B \in \mathcal{F}$,
iii. (closure under intersection) $A, B \in \mathcal{F}$ implies $A \cap B \in \mathcal{F}$.

For $\mathcal{F}$ a filter on V and $A \subseteq V$, we say that A is $\mathcal{F}$-*null* if $A^* \in \mathcal{F}$, $\mathcal{F}$-*positive* if A is not $\mathcal{F}$-null, and $\mathcal{F}$-*conull* if $A \in \mathcal{F}$.

Dual to the notion of a filter is the notion of an ideal.

Definition 1.3. A collection of subsets $\mathcal{I} \subseteq 2^V$ is an *ideal* on V if the following hold:

i. (non-triviality) $\emptyset \in \mathcal{F}$, $V \notin \mathcal{F}$,
ii. (downward closure) if $A \subseteq B \subseteq V$ and $B \in \mathcal{F}$, then $A \in \mathcal{F}$,
iii. (closure under union) $A, B \in \mathcal{F}$ implies $A \cup B \in \mathcal{F}$.

For $\mathcal{S} \subseteq V$ we write $\mathcal{S}^*$ for $\{X^* : X \in \mathcal{S}\}$. For a filter $\mathcal{F}$ on V, we have that $\mathcal{F}^*$ is an ideal on V, called the *dual ideal* to $\mathcal{F}$. And for any ideal $\mathcal{I}$ on V, $\mathcal{I}^*$ is a filter on V, the *dual filter* to $\mathcal{I}$.

For $\mathcal{I}$ an ideal on V and $A \subseteq V$, we say that A is *$\mathcal{I}$-null* if $A \in \mathcal{I}$, *$\mathcal{I}$-positive* if $A \notin \mathcal{I}$, and *$\mathcal{I}$-conull* if $A^* \in \mathcal{I}$. Thus the notions of null, positive, and conull are the same for both the filter and the ideal in a dual pair $\mathcal{F}, \mathcal{F}^*$. We express this informally by saying that the union of two null sets is null, and the intersection of two conull sets is conull.

A filter $\mathcal{U}$ on V is called an *ultrafilter* if for every $A \subseteq V$, either $A \in \mathcal{U}$ or $A^* \in \mathcal{U}$. An ideal $\mathcal{P}$ on V is a *prime ideal* if $\mathcal{P}^*$ is an ultrafilter. So, with respect to either an ultrafilter or its dual prime ideal, every subset of V is either null or conull.

Examples 1.4.

1. Suppose $\mathcal{F}$ is a filter on V. Let $1_\mathcal{F}$ be the function defined by
$$\begin{aligned} 1_\mathcal{F}(X) &= 1 \quad \text{if } X \text{ is } \mathcal{F}\text{-positive,} \\ &= 0 \quad \text{otherwise.} \end{aligned}$$
Then $1_\mathcal{F}$ is increasing and submodular.
2. Suppose $\mathcal{G}$ is a filter on V. Let $1^*_\mathcal{G}$ be the function defined by
$$\begin{aligned} 1^*_\mathcal{G}(X) &= 1 \quad \text{if } X^* \text{ is } \mathcal{G}\text{-positive,} \\ &= 0 \quad \text{otherwise.} \end{aligned}$$
Then $1_\mathcal{G}$ is decreasing and submodular.
3. Suppose $\mathcal{F}$ and $\mathcal{G}$ are filters on V, not both equal to the same ultrafilter. Let $1_{\mathcal{F}\to\mathcal{G}}$ be the function defined by
$$\begin{aligned} 1_{\mathcal{F}\to\mathcal{G}}(X) &= 1 \quad \text{if } X \text{ is } \mathcal{F}\text{-positive and } X^* \text{ is } \mathcal{G}\text{-positive,} \\ &= 0 \quad \text{otherwise.} \end{aligned}$$
Then $1_{\mathcal{F}\to\mathcal{G}}$ is submodular.

Proof. For (1.), $1_\mathcal{F}$ is increasing since being $\mathcal{F}$-positive is preserved under passing to a superset. For submodularity, fix $X, Y \subseteq V$. If exactly one of $1_\mathcal{F}(X)$ and $1_\mathcal{F}(Y)$ is 0, then exactly one of X and Y is $\mathcal{F}$-null, so that $X \cap Y$ is null and $X \cup Y$ is positive, i.e. $1_\mathcal{F}(X \cap Y) = 0$ and $1_\mathcal{F}(X \cup Y) = 1$. And if $1_\mathcal{F}(X) = 1_\mathcal{F}(Y) = 0$, then both X and Y are null, and hence so are $X \cap Y$ and $X \cup Y$, i.e. $1_\mathcal{F}(X \cap Y) = 1_\mathcal{F}(X \cup Y) = 0$. It follows that in all cases $1_\mathcal{F}(X \cap Y) + 1_\mathcal{F}(X \cup Y) \leq 1_\mathcal{F}(X) + 1_\mathcal{F}(Y)$.

For (2.), observe that $1^*_\mathcal{G} = (1_\mathcal{G})^*$.

For (3.), fix $X, Y \subseteq V$. If exactly one of $1_{\mathcal{F}\to\mathcal{G}}(X)$ and $1_{\mathcal{F}\to\mathcal{G}}(Y)$ is 0, then exactly one of X and Y is $\mathcal{F}$-null or $\mathcal{G}$-conull. In the first case $X \cap Y$ is $\mathcal{F}$-null, and in the second $X \cup Y$ is $\mathcal{G}$-conull. Thus in either case we have $1_{\mathcal{F}\to\mathcal{G}}(X \cap Y) + 1_{\mathcal{F}\to\mathcal{G}}(X \cup Y) \leq 1 = 1_{\mathcal{F}\to\mathcal{G}}(X) + 1_{\mathcal{F}\to\mathcal{G}}(Y)$.

If both $1_{\mathcal{F}\to\mathcal{G}}(X)$ and $1_{\mathcal{F}\to\mathcal{G}}(Y)$ are 0, then either X and Y are both $\mathcal{F}$-null, or both $\mathcal{G}$-conull, or one is $\mathcal{F}$-null and the other is $\mathcal{G}$-conull. In the first case, both $X \cap Y$ and $X \cup Y$ are $\mathcal{F}$-null; in the second, $X \cap Y$ and $X \cup Y$ are both $\mathcal{G}$-conull; and in the third, $X \cap Y$ is $\mathcal{F}$-null and $X \cup Y$ is $\mathcal{G}$-conull. Hence in all three cases $1_{\mathcal{F}\to\mathcal{G}}(X \cap Y) + 1_{\mathcal{F}\to\mathcal{G}}(X \cup Y) = 0$. $\square$

If $A \subseteq V$ is a non-empty subset of V, then the collection of sets $\mathcal{F}_A = \{X \subseteq V : A \subseteq X\}$ that contain A is a filter on V, called the *principal filter* determined by A. This filter is an ultrafilter (a *principal ultrafilter*) precisely when A is a singleton.

A set X is $\mathcal{F}_A$-positive if $A \cap X \neq \emptyset$. It follows that $1_A = 1_{\mathcal{F}_A}$. Similarly, if $B \subseteq V$ is non-empty, then $1^*_B = 1^*_{\mathcal{F}_B}$ and $1_{A\to B} = 1_{\mathcal{F}_A\to\mathcal{F}_B}$ (assuming A and B are not the same singleton). Thus the Examples in 1.4, along with the constant $\{0, 1\}$-valued functions, generalize Examples 1.1.

The following theorem is our classification of the $\{0,1\}$-valued submodular functions on V.

Theorem 1.5. Suppose that f is a $\{0,1\}$-valued submodular function on V. Then exactly one of the following holds:

i. $f \equiv 0$,
ii. $f \equiv 1$,
iii. $f = 1_{\mathcal{F}}$ for some filter $\mathcal{F}$ on V,
iv. $f = 1^*_{\mathcal{G}}$ for some filter $\mathcal{G}$ on V,
v. $f = 1_{\mathcal{F}\to\mathcal{G}}$ for some filters $\mathcal{F}$ and $\mathcal{G}$ on V, not both equal to the same ultrafilter.

Proof. Suppose first that f is non-constant and increasing. Let $\mathcal{I} = \{X \subseteq V : f(X) = 0\}$.

We claim that $\mathcal{I}$ is an ideal. Since f is non-constant and increasing we must have $f(\emptyset) = 0$ and $f(V) = 1$, so that $\emptyset \in \mathcal{I}$ and $V \notin I$. If $X \in \mathcal{I}$, and $Y \subseteq X$, then since f is increasing we must have $f(Y) \leq f(X) = 0$, so that $f(Y) = 0$ and $Y \in \mathcal{I}$. And if $X, Y \in \mathcal{I}$, so that $f(X) = f(Y) = 0$, then by submodularity we have $f(X \cap Y) + f(X \cup Y) = 0$, which gives $f(X \cap Y) = f(X \cup Y) = 0$. In particular, $X \cap Y \in \mathcal{I}$. Thus $\mathcal{I}$ is an ideal, as claimed.

Let $\mathcal{F} = \mathcal{I}^*$ be the dual filter. Then for $X \subseteq V$, the above gives $1_{\mathcal{F}}(X) = 1$ if and only if $X \notin \mathcal{I}$ if and only if $f = 1$, so that $f = 1_{\mathcal{F}}$.

If f is non-constant and decreasing, then f^* is non-constant and increasing, so that $f^* = 1_{\mathcal{G}}$ for some filter $\mathcal{G}$ on V. Hence $f = 1^*_{\mathcal{G}}$.

Finally, suppose that f is non-constant and neither increasing nor decreasing. We will need the following general claim.

Claim There is no triple of subsets $X \subseteq Y \subseteq Z$ of V such that $f(X) = f(Z) = 1$ and $f(Y) = 0$.

Proof. If there were such sets, then by the submodularity of f applied to the sets $X' = X \cup (Z \setminus Y)$ and Y, we have $f(X) + f(Z) \leq f(X') + f(Y)$. But $f(Y) = 0$ and $f(X) = f(Z) = 1$, so regardless of the value of $f(X')$ this inequality is false, a contradiction. $\square$

We claim that $f(\emptyset) = f(V) = 0$. Since f is non-constant, we cannot have $f(0) = f(V) = 1$, as this would yield some $\emptyset \subseteq X \subseteq V$ such that $f(X) = 0$, contradicting the claim. If $f(\emptyset) = 0$ and $f(V) = 1$, then since f is not increasing there must be some $X \subseteq Y \subseteq V$ such that $f(X) = 1$ and $f(Y) = 0$, again contradicting the claim. Dually, we cannot have $f(\emptyset) = 1$ and $f(V) = 0$. Thus $f(\emptyset) = f(V) = 0$, as claimed.

Let $\mathcal{I} = \{X \subseteq V : Y \subseteq X \Rightarrow f(Y) = 0\}$. We claim that $\mathcal{I}$ is an ideal. We have $\emptyset \in \mathcal{I}$ since $f(\emptyset) = 0$, and $V \notin \mathcal{I}$ since f is non-constant. Downward closure of $\mathcal{I}$ follows from the definition of $\mathcal{I}$, and closure under union follows from submodularity of f. Thus $\mathcal{I}$ is an ideal. Let $\mathcal{F} = \mathcal{I}^*$ denote the dual filter.

Dually, let $\mathcal{J}$ denote the ideal $\{X \subseteq V : Y \subseteq X \Rightarrow f^*(Y) = 0\}$, and let $\mathcal{G} = \mathcal{J}^*$. Observe that $X \in \mathcal{G}$ if and only if $X^* \in \mathcal{J}$, if and only if every subset $Y^* \subseteq X^*$ has $f^*(Y^*) = 0$, if and only if every superset $Y \supseteq X$ has $f(Y) = 0$ (since $f^*(Y^*) = f(Y)$).

We claim the $f(X) = 1$ if and only if X is $\mathcal{F}$-positive and X^* is $\mathcal{G}$-positive, i.e. that $f = 1_{\mathcal{F}\to\mathcal{G}}$. Suppose first that $f(X) = 1$. Then $X \notin \mathcal{I}$ by definition of $\mathcal{I}$, so that X is $\mathcal{F}$-positive. The $\mathcal{G}$-positivity of X^* is equivalent to the condition $X \notin \mathcal{G}$. But by the previous paragraph, $X \in \mathcal{G}$ implies in particular $f(X) = 0$. Hence $X \notin \mathcal{G}$, as desired.

Conversely, suppose X is $\mathcal{F}$-positive (so that $X \notin \mathcal{I}$) and X^* is $\mathcal{G}$-positive (so that $X \notin \mathcal{G}$). Then by definition of $\mathcal{I}$ there exists a subset $Y \subseteq X$ such that $f(Y) = 1$, and by our observation above there is a superset $Z \supseteq X$ such that $f(Z) = 1$. But then by the claim $f(X) = 1$, as desired. Note that it cannot be that $\mathcal{F}$ and $\mathcal{G}$ are equal to the same ultrafilter, as then $X \notin \mathcal{I}$ is equivalent to $X \in \mathcal{G}$, so that by what we have just shown this would give $f \equiv 0$.

We have shown that any non-constant $\{0,1\}$-valued submodular function f on V is of exactly one of the forms $1_{\mathcal{F}}$, $1^*_{\mathcal{G}}$, and $1_{\mathcal{F}\to\mathcal{G}}$, as desired. □

It is well-known that if V is finite, then the only filters on V are principal filters (i.e. have the form $\mathcal{F}_A$ for some non-empty $A \subseteq V$). Thus by Theorem 1.5, the examples in 1.1 exhaust the possibilities for $\{0,1\}$-valued submodular functions when V is finite.

2. Maximum flows in networks of filters

Our next goal is to define a generalized notion of directed graph in which the concept of a *directed edge* (i.e. a pair (a, b) of vertices from V) is replaced by the notion of a *filter edge*, which we define to be a pair $(\mathcal{F}, \mathcal{G})$ of filters on V. The motivation for this definition is that in a directed graph, an edge (a, b) is on the outgoing boundary of a set of vertices $X \subseteq V$ precisely when $1_{a\to b}(X) = 1$, i.e. when $a \in X$ and $b \notin X$. We will define a filter edge $(\mathcal{F}, \mathcal{G})$ to be on the boundary of X if $1_{\mathcal{F}\to\mathcal{G}}(X) = 1$, i.e. if X is $\mathcal{F}$-positive and X^* is $\mathcal{G}$-positive. By Theorem 1.5 this notion of edge is as general as possible, if by "edge" we mean something with a boundary indicator function that is $\{0,1\}$-valued and submodular. We call a collection of filter edges a *filter graph*.

In the proof of the max-flow min-cut theorem for finite directed graphs, the submodularity of the edge boundary function is used a crucial way. We will show that in a filter graph the edge boundary function is also submodular, and the natural analogue of the max-flow min-cut theorem holds.

We say *the* max-flow min-cut theorem, but there are many versions of the theorem. Roughly speaking, the version we will prove applies to the situation in which every edge has capacity 1, but we allow multiple sources and sinks, and also allow units of mass to flow out to infinity. A capacitated version of the theorem can also be proved, but for simplicity of presentation we will not do so here.

Before proving our max-flow min-cut theorem for filter graphs, we first prove its analogue for the corresponding notion of directed hypergraph. Though this will be subsumed by the theorem for filter graphs, the statement (and proof) of the hypergraph version will serve to motivate its more general form.

2.1. Max-flow min-cut for directed hypergraphs.

A *directed hyperedge* is an ordered pair $e = (A, B)$ of non-empty finite subsets $A, B \subseteq V$. The set A is the *outgoing side* or *tail* of e, and B is the *incoming side* or *head*. The cardinality of e is defined to be $|A \cup B|$, and denoted by $|e|$.

If $A = B$ then we also call e an *undirected hyperedge*, or simply *hyperedge*. Typically if e is an undirected hyperedge then $|e| \geq 2$ (i.e. e is not a loop). If

$A = B = \{a, b\}$, $a \neq b$, then we also call e an *undirected edge*. And if $A = \{a\}$ and $B = \{b\}$, $a \neq b$, then we also call e a *directed edge*.

A *directed hypergraph* G is a set of directed hyperedges (we view all of our graphs as being defined on the background vertex set V, and so identify a graph with its edge set). If G consists solely of hyperedges, or directed edges, or undirected edges, then we also call G an undirected hypergraph, or directed graph, or undirected graph, respectively. A directed hypergraph G is *finite* if G is a finite set, and *locally finite* if for every $x \in V$, the set $\{(A, B) \in G : x \in A \text{ or } x \in B\}$ consisting of edges with which x is incident, is finite. Every finite graph is also locally finite.

For $X \subseteq V$, we say that a directed hyperedge $e = (A, B)$ is on the *boundary* of X if $X \cap A \neq \emptyset$ and $B \not\subseteq X$, or equivalently if $1_{A\to B}(X) = 1$. When e is a directed edge, this coincides with the usual definition of e being on the outgoing edge boundary of X, and when e is an undirected edge, the edge boundary. We call $1_{A\to B}$ the *indicator* of the edge e, and also denote it by 1_e.

Definition 2.1.1. Suppose G is a locally finite directed hypergraph. Its *edge boundary function* ∂_G is the sum of its edge indicators:

$$\partial_G = \sum_{e\in G} 1_e.$$

For a finite subset $X \subseteq V$, $\partial_G(X)$ computes the size of the edge boundary of X (which will always be finite when G is locally finite). When G is fixed or understood, we also write $\partial_G(X)$ as $\partial(X)$.

It follows immediately from the definition of submodular function that a finite sum of submodular functions is submodular. Thus the edge boundary function ∂ in a locally finite graph G is submodular on finite subsets of V, in the sense that if $X, Y \subseteq V$ are finite sets of vertices then $\partial(X) + \partial(Y) \geq \partial(X \cup Y) + \partial(X \cap Y)$, since only a finite number of the edge indicators 1_e contribute to the values of $\partial(X), \partial(Y), \partial(X \cup Y)$, and $\partial(X \cap Y)$.

We will be interested in paths through our hypergraphs. These paths will either flow out to infinity, or terminate at a sink.

A *sink* is a non-empty set $A \subseteq V$ with associated indicator 1_A. A *directed hypergraph with sinks* is a pair (G, S) where G is a directed hypergraph on V and S is a set of sinks on V. We think of a sink $A \in S$ as being on the boundary of a set $X \subseteq V$ if $1_A(X) = 1$. The boundary function ∂ of a hypergraph with sinks (G, S) is defined as

$$\partial_{(G,S)} = \sum_{e\in G} 1_e + \sum_{A\in S} 1_A.$$

The graph (G, S) is *locally finite* if every vertex $x \in V$ is incident with only finitely many edges from G and finitely many sinks from S. If (G, S) is locally finite, then $\partial(X)$ is finite whenever $X \subseteq V$ is finite. Moreover, ∂ is submodular on finite subsets of V in the same sense as before, as on such sets ∂ reduces to a finite sum of submodular functions.

A *terminating path* in a directed hypergraph with sinks (G, S) is a sequence p of the form $x_1, e_1, x_2, e_2, \ldots, x_{n-1}, e_{n-1}, x_n, A$, where the $x_i \in V$ are vertices, the $e_i \in G$ are edges, and $A \in S$ is a sink. Writing $e_i = (A_i, B_i)$, we insist that $x_1 \in A_1$, and $x_n \in A$, and for $1 < i < n$ that $x_i \in B_{i-1} \cap A_i$. Finally, we insist that $e_i \neq e_j$ for $i \neq j$ (but allow $x_i = x_j$ for $i \neq j$ assuming the other conditions are satisfied). The *length* of such a path is $n - 1$, i.e. the number of edges appearing in the path (so a terminating path of length 0 is a pair x, A where x is a vertex belonging to the sink A). The vertex x_1 is the *initial vertex*, and we say that p *begins at* x_1.

A *partial path* (of length $n-1$) is a sequence p of the form $x_1, e_1, x_2, e_2, \ldots, x_{n-1}, e_{n-1}, x_n$ satisfying the same conditions as a terminating path, but without the terminal sink A.

A *non-terminating path* (or *infinite path*) in a directed hypergraph (with or without sinks) is a sequence $x_1, e_1, x_2, e_2, \ldots$ of vertices x_i and edges e_i, with x_1 on the outgoing side of e_1, and each x_i for $i > 1$ on the incoming side of e_{i-1} and outgoing side of e_i. As in a terminating path, vertices may be repeated along the path but edges may not.

A *complete path* is either a terminating or infinite path. A *path* refers to either a complete or partial path. Two paths p_1 and p_2 are *edge disjoint* if the edges appearing in p_1 are distinct from the edges appearing in p_2. A *flow* through a directed hypergraph G, possibly with sinks, is a set P of complete, pairwise edge disjoint paths through G.

We are interested in the following problem.

Problem 2.1.2. (Max flow problem) Suppose G is a locally finite directed hypergraph, possibly with sinks, and $X \subseteq V$ is finite. What is the maximum size of a flow P through G consisting of paths p with initial vertices in X?

If P is a flow all of whose initial vertices belong to X, we say that P *begins in* X. Given a finite set of vertices X, let $d(X)$ denote the maximum size of a flow beginning in X. The max flow problem is to determine $d(X)$. There is an obvious upper bound.

Definition 2.1.3. Suppose G is a locally finite directed hypergraph, possibly with sinks, and ∂ is its boundary function. Given a finite set of vertices $X \subseteq V$, define

$$c_G(X) = \min\{\partial(Y) : X \subseteq Y \text{ and } Y \text{ is finite}\}.$$

We call $c_G(X)$ the *capacity* of X.

When G is fixed or understood, we write $c_G(X)$ as $c(X)$. If $X \subseteq X'$ are finite sets of vertices, we say that X' *witnesses* or *realizes* the capacity of X if $\partial(X') = c(X)$.

Lemma 2.1.4. For a given G, the capacity function c is increasing and submodular on finite subsets of V.

Proof. Clearly c is increasing by definition. For submodulartiy, fix finite subsets $X, Y \subseteq V$. Choose finite sets $X' \supseteq X$ and $Y' \supseteq Y$ witnessing the capacity of X and Y respectively. We have:

$$\begin{array}{rcll} c(X \cup Y) + c(X \cap Y) & \leq & \partial(X' \cup Y') + \partial(X' \cap Y') & \text{(definition of } c) \\ & \leq & \partial(X') + \partial(Y') & \text{(submodularity of } \partial) \\ & = & c(X) + c(Y) & \text{(witnessing capacities)} \end{array}$$

as desired. $\square$

Suppose we have G, a finite set of vertices X, and a finite $X' \supseteq X$ witnessing the capacity of X. Recall that $\partial(X')$ counts the number of edges on the boundary of X' plus the number of sinks with vertices in X'. Observe that we must have $d(X) \leq c(X) = \partial(X')$: if P is a flow beginning in X, every path $p \in P$ must either terminate at some sink in X' or eventually cross the boundary of X', and since the paths in P are edge disjoint, any two paths that cross the boundary must cross along different edges.

The max-flow min-cut theorem says that this bound can be realized, i.e. that $d(X) = c(X)$. We will prove the theorem by using point masses to trace the paths that appear in the max flow P.

Definition 2.1.5. A *mass assignment* is a map $\mu : V \to \mathbb{N}$.

Given a mass assignment μ, we extend it to finite subsets of V by defining $\mu(X) = \sum_{x \in X} \mu(x)$ for all finite $X \subseteq V$.

In other words, a mass assignment is a non-negative, integer-valued finitely additive measure on (the finite subsets of) V. In particular, a mass assignment μ is *modular*, i.e. satisfies the identity

$$\mu(X \cap Y) + \mu(X \cup Y) = \mu(X) + \mu(Y),$$

for all finite $X, Y \subseteq V$.

Definition 2.1.6. Suppose that G is a locally finite directed hypergraph, possibly with sinks, and μ is a mass assignment on V. We say that μ is *feasible* for G if $\mu(X) \leq c(X)$ for all finite $X \subseteq V$.

If $\mu(x) = n$ for a given vertex x, we think of μ as assigning n point masses to x. Given a finite set of vertices X, we are going to index the paths p in our maximum flow P beginning from X by point masses assigned to vertices in X. We first need to show that we can find mass assignments that saturate the capacity of X. We will use several basic lemmas about submodular functions.

Lemma 2.1.7. Suppose f is non-negative and submodular on the finite subsets of V. Suppose $X, Y \subseteq V$ are finite such that $f(X) = f(Y) = 0$. Then $f(X \cup Y) = f(X \cap Y) = 0$.

Proof. By the submodularity and non-negativity of f, we have

$$0 \leq f(X \cup Y) + f(X \cap Y) \leq f(X) + f(Y) = 0,$$

which gives $f(X \cup Y) = f(X \cap Y) = 0$. $\square$

The lemma says that the sets assigned 0 by a non-negative submodular function form a lattice of sets.

Definition 2.1.8. Suppose G is a locally finite directed hypergraph, possibly with sinks, and μ is a feasible mass assignment for G. A finite $X \subseteq V$ is *saturated* by μ if $\mu(X) = c(X)$.

Lemma 2.1.9. If $X, Y \subseteq V$ are finite and saturated by μ, then so are $X \cup Y$ and $X \cap Y$.

Proof. Since μ is modular, so is $-\mu$, as modularity is preserved under taking negatives. In particular, $-\mu$ is submodular. Thus $c - \mu$ is submodular, being a sum of submodular functions. Since μ is feasible, $c - \mu$ is non-negative. The saturation of X and Y is equivalent to the assertion that $(c - \mu)(X) = (c - \mu)(Y) = 0$. The conclusion follows by Lemma 2.1.7. $\square$

Definition 2.1.10. Suppose μ is feasible for G. The *relative capacity function* c_μ is defined by

$$c_\mu(X) = \min\{\partial(Y) - \mu(Y) : X \subseteq Y \text{ and } Y \text{ is finite}\}.$$

Lemma 2.1.11. For a given G and feasible μ, the relative capacity function c_μ is increasing and submodular on finite subsets of V.

Proof. Again, c is clearly increasing. For submodulartiy, fix finite subsets $X, Y \subseteq V$. Choose finite sets $X' \supseteq X$ and $Y' \supseteq Y$ witnessing the relative capacity of X and Y respectively, i.e. so that $c_\mu(X) = \partial(X') - \mu(X')$ and likewise for Y, Y'. We have:

$$\begin{array}{rcl} c_\mu(X \cup Y) + c_\mu(X \cap Y) & \leq & \partial(X' \cup Y') + \partial(X' \cap Y') - \mu(X' \cup Y') - \mu(X' \cap Y') \\ & \leq & \partial(X') + \partial(Y') - \mu(X' \cup Y') - \mu(X' \cap Y') \\ & = & \partial(X') + \partial(Y') - \mu(X') - \mu(Y') \\ & = & c_\mu(X) + c_\mu(Y) \end{array}$$

where in passing from the second to third line we applied the modularity of μ. □

For a vertex $x \in X$, we write 1_x for $1_{\{x\}}$, the *point mass at x*.

Lemma 2.1.12. Suppose G is a locally finite directed hypergraph, possibly with sinks, and μ is a feasible mass assignment for G. If X is a finite set of vertices such that $c_\mu(X) \geq 1$, then there is a vertex $x \in X$ such that the assignment $\mu' = \mu + 1_x$ is feasible for G. Moreover, $c_{u'}(X) = c_u(X) - 1$.

Proof. If there is no such x, then for every $x \in X$ we can find a finite set of vertices Y_x such that $(\mu + 1_x)(Y_x) > \partial(Y_x)$. Since $\mu(Y_x) \leq \partial(Y_x)$ by feasibility, it must in fact be that $\mu(Y_x) = \partial(Y_x)$ and that $(\mu+1_x)(Y_x) = \mu(Y_x)+1$, so that $x \in Y_x$. That is, Y_x is a μ-saturated set containing x. But then $Y = \bigcup_{x \in X} Y_x$ is also μ-saturated by Lemma 2.1.9, i.e. $\partial(Y) - \mu(Y) = 0$, and moreover $X \subseteq Y$. But this contradicts $c_\mu(X) \geq 1$.

Thus we can find x such that $u' = u + 1_x$ is feasible for G. We then clearly have $c_{\mu'}(X) = c_\mu(X) - 1$, since for every finite $Y \supseteq X$ we have $\partial(Y) - \mu'(Y) = \partial(Y) - \mu(Y) - 1$. □

Lemma 2.1.13. Suppose G is a locally finite directed hypergraph, possibly with sinks, and $X \subseteq V$ is finite. Then there is a feasible mass assignment μ for G such that $\mu(X) = c(X)$.

Proof. By induction, using Lemma 2.1.12. □

We now prove the max-flow min-cut theorem for directed hypergraphs with sinks. It will follow from the theorem that if X is a finite set of vertices in such a graph, then there is flow P beginning in X of size $c(X)$.

The max-flow min-cut theorem we actually prove below gives one step in the inductive process of building the paths that constitute a maximum flow P. Roughly, it says that if μ is a feasible mass assignment for a directed hypergraph (G, S) with sinks, and x is a vertex in G containing a unit mass, then either we can find an edge $(A, B) \in G$ with $x \in A$ such that we can move the unit mass on x to some $y \in B$ and remove the edge, retaining feasibility of the assignment in the resulting graph (i.e. the mass crosses the edge and uses its capacity), or we can find a sink $A \in S$ with $x \in A$ that we can remove along with the unit mass from x, retaining feasibility (i.e. the mass has reached its sink, so we use the sink's capacity and remove the mass). Once proved, we can apply the theorem inductively to move units of mass forward from their initial vertices in a given set X, until they either reach a sink and exit or flow out to infinity. Since in this process we remove edges as point masses move across them, the paths these masses trace are edge disjoint, i.e. constitute a flow.

For μ a mass assignment and x a vertex with $\mu(x) \geq 1$, we write $\mu_{x\to y}$ for the mass assignment that moves a unit of mass from x to y, i.e. $\mu_{x\to y} = \mu - 1_x + 1_y$.

Given (G, S) a directed hypergraph with sinks and $(A, B) \in G$, $(G - (A, B), S)$ denotes the graph with the edge (A, B) removed. Given a source $A \in S$, we write $(G, S - A)$ for the graph less the sink A.

Theorem 2.1.14. (Max-flow min-cut) Suppose (G, S) is a locally finite directed hypergraph with sinks and μ is a feasible mass assignment for G. Suppose further we have $x \in V$ with $\mu(x) \geq 1$.

Let $(A_1, B_1), (A_2, B_2), \ldots, (A_k, B_k), A_{k+1}, A_{k+2}, \ldots, A_n$ list the edges $(A_i, B_i) \in G$ for which $x \in A_i$, followed by the sinks $A_j \in S$ with $x \in A_j$.

Then either there is an edge (A_i, B_i), $1 \leq i \leq k$ and a vertex $y \in B_i$ such that the assignment $\mu_{x\to y}$ is feasible for $(G-(A_i, B_i), S)$, or there is a sink A_j, $k+1 \leq j \leq n$ such that $\mu - 1_x$ is feasible for $(G, S - A_j)$.

Proof. Let $\mu' = \mu - 1_x$. Denote (G, S) by G. For $1 \leq i \leq k$ let G_i denote G minus (A_i, B_i) and for $k + 1 \leq i \leq n$ let G_i denote G minus A_i. Let ∂_i denote the boundary function for G_i.

We first claim that for some $i \leq n$, μ' is feasible for G_i. If not, then for every $i \leq n$ we can find a finite Y_i such that $\partial_i(Y_i) < \mu'(Y_i)$. Since $\partial(Y_i) \geq \mu(Y_i)$ by feasibility, the only way this inequality is possible is if $\partial(Y_i) = \mu(Y_i)$, $\partial_i(Y_i) = \partial(Y_i) - 1$, and $\mu'(Y_i) = \mu(Y_i)$. These conditions imply that Y_i is saturated by μ in G, intersects A_i, and does not contain x. But then $Y = \bigcup_i Y_i$ is also saturated by μ in G and does not contain x. Moreover, Y intersects every A_i, $1 \leq i \leq n$. Since these list the sinks and edge tails with which x is incident, it follows that adding x to Y will not increase the boundary size of Y, as $\partial(Y)$ already counts any edges that might be on the boundary of a set containing x. That is, $\partial(Y \cup \{x\}) = \partial(Y)$, which equals $\mu(Y)$ by saturation. But $\mu(Y \cup \{x\}) > \mu(Y)$ as $\mu(x) \geq 1$, which gives $\mu(Y \cup \{x\}) > \partial(Y)$, contradicting the feasibility of μ.

Thus we cannot find such a Y_i for every $i \leq n$. Let I denote the collection of indices i for which such a Y_i exists, and let $Y = \bigcup_{i\in I} Y_i$ be the union of these sets. Then Y is saturated by μ in G, does not contain x, and intersects every A_i for $i \in I$.

Let $J = \{i \leq n : i \notin I\}$ denote the indices for which there is no such Y_i. We have just argued that J is non-empty. It follows from above that for every $i \in J$ we have that μ' is feasible for G_i. If there is $i \in J$ so that $k + 1 \leq i \leq n$, then we are done, since we have shown $\mu' = \mu - 1_x$ is feasible for $G_i = (G, S - A_i)$. So assume there is no such i.

We now claim there is $i \in J$ (with $i \leq k$ necessarily) and $y \in B_i$ such that $\mu' + 1_y = \mu_{x\to y}$ is feasible for G_i. If not, then for every $i \in J$ and every $y \in B_i$ we can find $Z_{i,y}$ such that $(\mu' + 1_y)(Z_{i,y}) > \partial_i(Z_{i,y})$. Since μ' is feasible for G_i, it follows that $Z_{i,y}$ is saturated by μ' in G_i, and moreover contains y. But then $Z_i = \bigcup_{y\in B_i} Z_{i,y}$ is also μ'-saturated in G_i, and moreover contains B_i as a subset. Since $B_i \subseteq Z_i$, we have $1_{A_i\to B_i}(Z_i) = 0$, so that $\partial_i(Z_i) = \partial(Z_i)$, which gives $\mu'(Z_i) = \partial(Z_i)$, i.e. Z_i is saturated in the original graph G by μ'. It follows that $x \notin Z_i$, since otherwise we would have $\mu(Z_i) = \mu'(Z_i) + 1 > \partial(Z_i)$, contradicting the feasibility of μ in G.

Repeating this argument for every $i \in J$ we find sets $Z_i, i \in J$ that are saturated by μ' in G and do not contain x. Hence $Z = \bigcup_{i\in J} Z_i$ is also saturated by μ' in G and does not contain x. Therefore $Y \cup Z$ is also saturated by u' in G and does not contain x. But u and u' agree on sets not containing x, so that $\mu(Y \cup Z) = \mu'(Y \cup Z) = \partial(Y \cup Z)$, i.e. $Y \cup Z$ is saturated by μ in G.

By construction, for every $i \leq n$ we have that $Y \cup Z$ either intersects A_i or contains B_i. It follows that adding x to this set will not increase the size of its boundary, i.e. $\partial(Y \cup Z \cup \{x\}) = \partial(Y \cup Z)$. But since $\mu(x) \geq 1$ we have $\mu(Y \cup Z \cup \{x\}) > \mu(Y \cup Z) = \partial(Y \cup Z) = \partial(Y \cup Z \cup \{x\})$, contradicting the feasibility of μ.

Thus there must be $i \in J$ and $y \in B_i$ for which $\mu_{x \to y}$ is feasible for G_i. We are done. □

Corollary 2.1.15. If G is a locally finite directed hypergraph, possibly with sinks, and $X \subseteq V$ is a finite set of vertices, then there is a flow P beginning in X of size $c(X)$.

Proof. By Lemma 2.1.13, we can find a feasible mass assignment μ for G such that $\mu(X) = c(X)$. If we iteratively apply Theorem 2.1.14 to the units of mass assigned to X by μ, then the paths traced by these point masses (which either eventually reach a sink and exit, or flow out to infinity) constitute a flow P of size $c(X)$ beginning in X. □

2.2. **Max-flow min-cut for filter graphs.** We now develop, in close analogy with the previous section, a generalized notion of a directed hypergraph that we call a *filter graph*, and prove the analogue of our max-flow min-cut theorem for locally finite filter graphs.

We first recall some basic facts about filters and ultrafilters that we will need for our results. We assume for now that V is an infinite set; when V is finite our work will reduce to that from the previous section. Let βV denote the collection of ultrafilters on V. For each set $A \subseteq V$, let $O_A \subseteq \beta V$ denote the collection of ultrafilters $\{\mathcal{U} \in \beta V : A \in \mathcal{U}\}$ that concentrate on A.

Lemma 2.2.1.

1. For every filter $\mathcal{F} \subseteq 2^V$, there is an ultrafilter $\mathcal{U} \in \beta V$ such that $\mathcal{F} \subseteq \mathcal{U}$. We say $\mathcal{U}$ *extends* $\mathcal{F}$.
2. If $\mathcal{F} \neq \mathcal{G}$ are distinct filters on V, then there is an ultrafilter $\mathcal{U}$ extending $\mathcal{F}$ that does not extend $\mathcal{G}$, i.e. $\mathcal{F} \subseteq \mathcal{U}$, but there is $A \in \mathcal{G}$ which is $\mathcal{U}$-null.
3. The collection $\langle O_A : A \subseteq V \rangle$ forms a basis for a topology on βV. Under this topology, βV is a totally disconnected, 0-dimensional, compact Hausdorff space. In particular, each basic open set O_A is clopen in this topology.

Roughly speaking, we are going to generalize our work from the previous section by replacing subsets $A, B, \ldots$ of V by filters $\mathcal{F}, \mathcal{G}, \ldots$ on V, and vertices (or point masses) $x, y, \ldots$ in V by ultrafilters $\mathcal{U}, \mathcal{V}, \ldots$ on V. This is a natural generalization, in the sense that if V is finite, then all filters $\mathcal{F}$ and ultrafilters $\mathcal{U}$ are principal, i.e. are determined by subsets $A \subseteq V$ and points $x \in V$, respectively.

More precisely, define a *filter edge* to be a pair $e = (\mathcal{F}, \mathcal{G})$ of filters on V. If $X \subseteq V$, the edge $(\mathcal{F}, \mathcal{G})$ is on the *boundary* of X if $1_{\mathcal{F} \to \mathcal{G}}(X) = 1$. We also write 1_e for $1_{\mathcal{F} \to \mathcal{G}}$. We say $\mathcal{F}$ is the *outgoing side* of $(\mathcal{F}, \mathcal{G})$, and $\mathcal{G}$ the *incoming side*. A *sink* is a filter $\mathcal{F}$ on V with indicator $1_{\mathcal{F}}$, and a set $X \subseteq V$ is *incident* with the sink if $1_{\mathcal{F}}(X) = 1$.

A *filter graph* is a pair (G, S) where G is a set of filter edges and S is a set of sinks. We will often refer to a filter graph by its edge set G. The *boundary function*

for G is defined as
$$\partial_G = \sum_{e \in G} 1_e + \sum_{\mathcal{F} \in S} 1_{\mathcal{F}}.$$
As before, when G is understood, we write ∂ for ∂_G.

If $\mathcal{U}$ is an ultrafilter on V, then its indicator $1_{\mathcal{U}}$ is the *ultrafilter measure* associated to $\mathcal{U}$. Since for an ultrafilter $\mathcal{U}$, we have that either $X \in \mathcal{U}$ or $X^* \in \mathcal{U}$ for every $X \subseteq V$, ultrafilter measures are finitely additive measures, and in particular are modular: $1_{\mathcal{U}}(X \cup Y) + 1_{\mathcal{U}}(X \cap Y) = 1_{\mathcal{U}}(X) + 1_{\mathcal{U}}(Y)$ for all $X, Y \subseteq V$.

Ultrafilter measures $1_{\mathcal{U}}$ play the role for filter graphs that point masses 1_x played for directed hypergraphs above.

Definition 2.2.2. A *mass assignment* μ is a non-negative integer weighted sum of ultrafilter measures, i.e.
$$\mu = \sum_{\mathcal{U} \in \beta V} n_{\mathcal{U}} 1_{\mathcal{U}},$$
where each $n_{\mathcal{U}} \in \mathbb{N}$.

If $e = (\mathcal{F}, \mathcal{G})$ is a filter edge and $\mathcal{U}$ is an ultrafilter, then if $\mathcal{U}$ extends $\mathcal{F}$ (respectively $\mathcal{G}$) we say $\mathcal{U}$ lies on the outgoing side of e (respectively, incoming side). We use the same language for the ultrafilter mass $1_{\mathcal{U}}$. Observe that if $\mathcal{U}$ lies on the outgoing side of e and $X \in \mathcal{F}$ we have $1_{\mathcal{U}}(X) = 1$; in this case we think of the mass $1_{\mathcal{U}}$ as being concentrated in every set in $\mathcal{F}$. Similarly, for a sink $\mathcal{F}$ we say $\mathcal{U}$ lies in $\mathcal{F}$ if $\mathcal{U}$ extends $\mathcal{F}$.

The notion of locally finite for filter graphs involves a technical condition without an analogue in the corresponding definition for hypergraphs.

Definition 2.2.3. A filter graph (G, S) is *locally finite* if for every ultrafilter $\mathcal{U}$ on V, the following conditions are satisfied:

- the number of edges $(\mathcal{F}, \mathcal{G}) \in G$ with $\mathcal{U} \supseteq \mathcal{F}$ is finite,
- the number of sinks $\mathcal{F} \in S$ with $\mathcal{U} \supseteq \mathcal{F}$ is finite,
- there is a set of vertices $A \subseteq V$ with $A \in \mathcal{U}$ such that for every edge $(\mathcal{F}, \mathcal{G}) \in G$ such that $\mathcal{F} \not\subseteq \mathcal{U}$, A is $\mathcal{F}$-null, and likewise for every sink $\mathcal{F} \in S$ with $\mathcal{F} \not\subseteq \mathcal{U}$, A is $\mathcal{F}$-null.

The third condition says that every for every ultrafilter $\mathcal{U}$, we can find a set of vertices A carrying the ultrafilter mass $1_{\mathcal{U}}$ that only activates the indicators in ∂ that are activated by $\mathcal{U}$. We leave it to the reader to verify that this condition is satisfied in every finite filter graph (i.e. every finite filter graph is locally finite).

For a filter graph (G, S), a set of vertices $X \subseteq V$ is *G-finitary* (or simply *finitary* if G is understood) if there are only finitely many filter edges $(\mathcal{F}, \mathcal{G})$ in G such that X is incident with $\mathcal{F}$ (i.e. such that $1_{\mathcal{F}}(X) = 1$), and finitely many sinks $\mathcal{F}$ in S such that X is incident with $\mathcal{F}$.

If X is finitary, then $\partial(X)$ is defined (i.e. finite), since only finitely many indicators from $\partial(X)$ are possibly nonzero on X. If X and Y are finitary, then it is not hard to see that $X \cap Y$ and $X \cup Y$ are finitary as well. Since edge indicators 1_e and sink indicators $1_{\mathcal{F}}$ are submodular, and finite sums of submodular functions are submodular, we have $\partial(X \cap Y) + \partial(X \cup Y) \leq \partial(X) + \partial(Y)$ for any pair of finitary sets X and Y. We express this by saying that ∂_G is submodular on G-finitary subsets of V.

We now develop the notions of capacity and feasibility for filter graphs, and prove the attendant lemmas. In many of these the main difference from the corresponding lemmas in the previous section is that the assumption of finite is replaced by the assumption of finitary.

Definition 2.2.4. Suppose G is a locally finite filter graph. Given a G-finitary set of vertices $X \subseteq V$, define

$$c_G(X) = \min\{\partial(Y) : X \subseteq Y \text{ and } Y \text{ is finitary}\}.$$

We call $c_G(X)$ the *capacity* of X.

Note that $c_G(X)$ is defined (i.e. finite) for all finitary X. When G is understood, we write $c(X)$ for $c_G(X)$. If $X' \supseteq X$ is finitary and $\partial(X') = c(X)$, we say as before that X' *witnesses* the capacity of X.

Lemma 2.2.5. For a given filter graph G, the capacity function c is increasing and submodular on finitary subsets of V.

Proof. Essentially the same as the proof of Lemma 2.1.4, with finite sets replaced by finitary sets. □

The following definitions and related lemmas are analogous to the corresponding ones from the previous section. Their proofs are essentially the same as before, so we leave them for the reader to verify.

Definition 2.2.6. Suppose that G is a filter graph and μ is a mass assignment. We say that μ is *feasible* for G if $\mu(X) \leq c(X)$ for all G-finitary $X \subseteq V$.

Definition 2.2.7. Suppose that G is a filter graph and μ is a feasible mass assignment for G. A finitary $X \subseteq V$ is *saturated* by μ if $\mu(X) = c(X)$.

Lemma 2.2.8. If $X, Y \subseteq V$ are G-finitary and saturated by μ, then so are $X \cup Y$ and $X \cap Y$.

Definition 2.2.9. Suppose μ is feasible for a filter graph G. The *relative capacity function* c_μ is defined for G-finitary X by

$$c_\mu(X) = \min\{\partial(Y) - \mu(Y) : X \subseteq Y \text{ and } Y \text{ is finite}\}.$$

Lemma 2.2.10. For a given filter graph G and feasible μ, the relative capacity function c_μ is increasing and submodular on G-finitary sets.

The next lemma is the filter graph analogue of Lemma 2.1.12. It says that if X is G-finitary and has excess capacity with respect to a given feasible μ, then we can add an ultrafilter mass to X while maintaining feasibility. This will allow us to saturate finitary subsets. The proof relies on the compactness of the space of ultrafilters βV.

Lemma 2.2.11. Suppose G is a filter graph and μ is a feasible mass assignment for G. If X is a finitary set of vertices such that $c_\mu(X) \geq 1$, then there is an ultrafilter $\mathcal{U}$ on V with $X \in \mathcal{U}$ such that the assignment $\mu' = \mu + 1_{\mathcal{U}}$ is feasible for G. Moreover, $c_{u'}(X) = c_u(X) - 1$.

Proof. Suppose there is no such ultrafilter $\mathcal{U}$. Then for every ultrafilter $\mathcal{U}$ containing X, i.e. every $\mathcal{U}$ in the clopen set $O_X \subseteq \beta V$, we can find a finitary $Y_{\mathcal{U}}$ such that $\mu'(Y_{\mathcal{U}}) > \partial(Y_{\mathcal{U}})$. It follows that $\mu(Y_{\mathcal{U}}) = \partial(Y_{\mathcal{U}})$ and $Y_{\mathcal{U}} \in \mathcal{U}$, which gives $\mathcal{U} \in O_{Y_{\mathcal{U}}}$. Hence $O_X \subseteq \bigcup_{\mathcal{U} \in O_X} O_{Y_{\mathcal{U}}}$, so that $\langle O_{Y_{\mathcal{U}}} : \mathcal{U} \in O_X \rangle$ is an open cover for O_X. Since O_X is a closed subset of the compact space βV, it is compact, and hence there is a finite subcollection $Y_{\mathcal{U}_1}, Y_{\mathcal{U}_2}, \ldots, Y_{\mathcal{U}_n}$ of the sets $Y_{\mathcal{U}}$ such that $O_X \subseteq \bigcup_{i \leq n} O_{Y_{\mathcal{U}_i}}$. This gives that if $\mathcal{U}$ is an ultrafilter with $X \in \mathcal{U}$, then there is an $i \leq n$ such that $Y_{\mathcal{U}_i} \in \mathcal{U}$. It follows that the set $Y = \bigcup_{i \leq n} Y_{\mathcal{U}_i}$ has the property that it belongs to every ultrafilter $\mathcal{U}$ containing X. In particular, by considering the principal ultrafilters $\mathcal{U}_x$ corresponding to points in $x \in X$, this gives that $X \subseteq Y$.

But Y is a finite union of μ-saturated, finitary sets, hence is itself μ-saturated and finitary. But the fact that $\partial(Y) - \mu(Y) = 0$ contradicts the condition that $c_\mu(X) = \min\{\partial(Z) - \mu(Z) : Z \supseteq X \text{ and } Z \text{ finitary}\} \geq 1$. □

It follows from the lemma that we can saturate finitary sets of vertices.

Corollary 2.2.12. Suppose G is a filter graph and $X \subseteq V$ is G-finitary. Then there is a feasible mass assignment μ for G such that $\mu(X) = c(X)$.

Proof. By induction, using Lemma 2.2.11. □

Finally, we prove the max-flow min-cut theorem for filter graphs. We note that in the case when V is finite, the theorem reduces to the max-flow min-cut theorem from the previous section.

If $\mathcal{V}$ is an ultrafilter and μ is a mass assignment, we say that μ assigns mass at $\mathcal{V}$ if in the sum

$$\mu = \sum_{\mathcal{U} \in \beta V} n_{\mathcal{U}} 1_{\mathcal{U}}$$

the coefficient $n_{\mathcal{V}}$ is at least 1.

If μ assigns mass at $\mathcal{U}$ and $\mathcal{V}$ is another ultrafilter, we write $\mu_{\mathcal{U} \to \mathcal{V}}$ for the assignment $\mu - 1_{\mathcal{U}} + 1_{\mathcal{V}}$.

Theorem 2.2.13. (Max-flow min-cut for filter graphs) Suppose (G, S) is a locally finite filter graph and μ is a feasible mass assignment for G. Suppose that $\mathcal{U}$ is an ultrafilter such that μ assigns mass at $\mathcal{U}$.

Let $e_1 = (\mathcal{F}_1, \mathcal{G}_1), e_2 = (\mathcal{F}_2, \mathcal{G}_2), \ldots, e_k = (\mathcal{F}_k, \mathcal{G}_k), \mathcal{F}_{k+1}, \mathcal{F}_{k+2}, \ldots, \mathcal{F}_n$ list the edges $e_i \in G$ for which $\mathcal{U} \supseteq \mathcal{F}_i$, followed by the sinks $\mathcal{F}_j \in S$ with $\mathcal{U} \supseteq \mathcal{F}_j$.

Then either there is an edge e_i in the list and an ultrafilter $\mathcal{V} \supseteq \mathcal{G}_i$ such that the assignment $\mu_{\mathcal{U} \to \mathcal{V}}$ is feasible for $(G - e_i, S)$, or there is a sink $\mathcal{F}_i$ in the list such that $\mu - 1_{\mathcal{U}}$ is feasible for $(G, S - \mathcal{F}_i)$.

Proof. Let $\mu' = \mu - 1_{\mathcal{U}}$. For $i \leq k$, let G_i denote $(G - e_i, S)$ and for $k + 1 \leq i \leq n$, let G_i denote $(G, S - \mathcal{F}_i)$. Let ∂_i denote the boundary function for G_i.

We claim there is $i \leq n$ such that μ' is feasible for G_i. If not, then for every $i \leq n$ we can find a finitary Y_i such that $\partial_i(Y_i) < \mu'(Y_i)$. By the feasibility of μ, it follows that for every i we must have $\partial(Y_i) = \mu(Y_i)$, $\partial_i(Y_i) = \partial(Y_i) - 1$, and $\mu'(Y_i) = \mu(Y_i)$. These conditions imply that Y_i is saturated by μ in G, $1_{\mathcal{F}_i}(Y_i) = 1$, and $Y_i \notin \mathcal{U}$. Let $Y = \bigcup_i Y_i$. Then Y is finitary, and it follows that Y is also μ-saturated in G, $1_{\mathcal{F}_i}(Y) = 1$, and $Y \notin \mathcal{U}$ (each Y_i is $\mathcal{U}$-null since $\mathcal{U}$ is an ultrafilter).

By the local finiteness of G, we can find $A \in \mathcal{U}$ such that for any $\mathcal{F}$ which is a sink in G or the outgoing side of some edge in G, we have $1_{\mathcal{F}}(A) = 1$ if and only if $\mathcal{F} = \mathcal{F}_i$ for some $i \leq n$. It follows that $\partial(Y \cup A) = \partial(Y)$, since any indicator in ∂ activated by A is already activated by Y. It follows $\partial(Y \cup A) = \mu(Y)$. On the other hand, since $Y \notin \mathcal{U}$ but $Y \cup A \in \mathcal{U}$ we have $\mu(Y \cup A) > \mu(Y)$ since μ assigns mass to $\mathcal{U}$. This gives $\mu(Y \cup A) > \partial(Y \cup A)$, contradicting the feasibility of μ.

Thus we cannot find such a Y_i for every $i \leq n$. Let I denote the indices $i \leq n$ for which such a Y_i exists, and let $Y = \bigcup_{i \in I} Y_i$ be the union of these sets. Then arguing as above, Y is μ-saturated in G, $1_{\mathcal{F}_i}(Y) = 1$ for $i \in I$, and $Y \notin \mathcal{U}$.

Let $J = \{i \leq n : i \notin I\}$ denote the indices for which there is no such Y_i. We have shown that J is non-empty. It follows from above that for every $i \in J$, μ' is feasible for G_i. If there is $i \in J$ with $k+1 \leq i \leq n$, then we are done, since we have shown $\mu' = \mu - 1_{\mathcal{U}}$ is feasible for $G_i = (G, S - \mathcal{F}_i)$. So assume there is no such i.

We claim there is $i \in J$ (so $i \leq k$) and an ultrafilter $\mathcal{V} \supseteq \mathcal{G}_i$ such that $\mu_{\mathcal{U} \to \mathcal{V}} = \mu' + 1_{\mathcal{V}}$ is feasible for G_i. If not, then for every $i \in J$ and every ultrafilter $\mathcal{V} \supseteq \mathcal{G}_i$ (i.e. every $\mathcal{V}$ in the closed subspace $D_i = \{\mathcal{V} \in \beta V : \mathcal{V} \supseteq \mathcal{G}_i\}$), we can find a finitary $Z_{i,\mathcal{V}}$ such that $(\mu' + 1_{\mathcal{V}})(Z_{i,\mathcal{V}}) > \partial_i(Z_{i,\mathcal{V}})$. Since μ' is feasible for G_i, it must be that $Z_{i,\mathcal{V}}$ is μ'-saturated in G_i, and moreover $Z_{i,\mathcal{V}} \in \mathcal{V}$. Thus $\langle O_{Z_{i,\mathcal{V}}} : \mathcal{V} \in D_i \rangle$ is an open cover for D_i. Let $O_{Z_{i,\mathcal{V}_1}}, O_{Z_{i,\mathcal{V}_2}}, \ldots, O_{Z_{i,\mathcal{V}_n}}$ be a finite subcover. Then every ultrafilter $\mathcal{V} \supseteq \mathcal{G}_i$ contains one of the sets $Z_{i,\mathcal{V}_1}, \ldots, Z_{i,\mathcal{V}_n}$, and hence contains their union Z_i. It follows from Lemma 2.2.1 that $Z_i \in \mathcal{G}_i$. Since Z_i is a finite union of finitary sets that are μ'-saturated in G_i, Z_i is also finitary and μ'-saturated in G_i. Since $Z_i \in \mathcal{G}_i$, we have $1_{\mathcal{F}_i \to \mathcal{G}_i}(Z_i) = 0$, so that $\partial_i(Z_i) = \partial(Z_i)$, which gives $\mu'(Z_i) = \partial(Z_i)$, i.e. Z_i is saturated in the original graph G by μ'.

We now claim there is $i \in J$ (with $i \leq k$ necessarily) and $y \in B_i$ such that $\mu' + 1_y = \mu_{x \to y}$ is feasible for G_i. If not, then for every $i \in J$ and every $y \in B_i$ we can find $Z_{i,y}$ such that $(\mu' + 1_y)(Z_{i,y}) > \partial_i(Z_{i,y})$. Since μ' is feasible for G_i, it follows that $Z_{i,y}$ is saturated by μ' in G_i, and moreover contains y. But then $Z_i = \bigcup_{y \in B_i} Z_{i,y}$ is also μ'-saturated in G_i, and moreover contains B_i as a subset. Since $B_i \subseteq Z_i$, we have $1_{A_i \to B_i}(Z_i) = 0$, so that $\partial_i(Z_i) = \partial(Z_i)$, which gives $\mu'(Z_i) = \partial(Z_i)$, i.e. Z_i is saturated in the original graph G by μ'. It follows that $x \notin Z_i$, since otherwise we would have $\mu(Z_i) = \mu'(Z_i) + 1 > \partial(Z_i)$, contradicting the feasibility of μ in G. It follows that $Z_i \notin \mathcal{U}$, since otherwise we would have $\mu(Z_i) = \mu'(Z_i) + 1 > \partial(Z_i)$, contradicting the feasibility of μ in G.

Repeating this argument for each $i \in J$, we find finitary sets Z_i that are μ'-saturated in G with $Z_i \notin \mathcal{U}$. Hence their union $Z = \bigcup_{i \in J} Z_i$ is also finitary, μ'-saturated in G, and $Z \notin \mathcal{U}$. Therefore $Y \cup Z$ is also finitary, saturated by μ' in G, and $Y \cup Z \notin \mathcal{U}$. Since μ and μ' agree on sets not in $\mathcal{U}$, we have $\mu(Y \cup Z) = \mu'(Y \cup Z) = \partial(Y \cup Z)$ so that $Y \cup Z$ is μ-saturated in G.

By construction, for every $i \leq n$ we have that either $Y \cup Z$ is $\mathcal{F}_i$-positive or $Y \cup Z \in \mathcal{G}_i$. It follows that $\partial(Y \cup Z \cup A) = \partial(Y \cup Z)$ (where A is our set from above with $A \in \mathcal{U}$). But $\mu(Y \cup Z \cup A) > \mu(Y \cup Z)$ since μ assigns mass to $\mathcal{U}$, which gives $\mu(Y \cup Z \cup A) > \partial(Y \cup Z \cup A)$, contradicting feasibility. $\square$