On Optimum k-way Partitions with Submodular Costs and Minimum Part-size Constraints

Maurice Queyranne

Sauder School of Business at UBC,

and CNRS – Universidad de Chile

joint work with Flavio Guiñez (Universidad de Chile)

Workshop Modern Aspects of Submodularity

Georgia Tech, March 19-22, 2012

Overview

- Given integer $k \geq 2$, a ***k*-way partition**, in short, a ***k*-cut**, of a set V is a partition $\boldsymbol{P} = \{P_1, \ldots, P_k\}$ of V into k (nonempty) parts.
- Given a cost function $f : 2^V \to \mathbf{R}$, an optimum k-cut minimizes the total cost $f(\boldsymbol{P}) = \sum_{i=1..k} f(P_i)$
- Given a k-vector $\boldsymbol{s} = (s_1, \ldots, s_k)$ of positive integers, k-cut $\boldsymbol{P}$ is ***s*-size** if $|P_i| \geq s_i$ for all i

We consider the **computational complexity** of optimum ***s***-size k-cut problems defined by special cost functions:

- Graph k-cuts: **G *k*-cuts**
- Hypergraph k-cuts: **H *k*-cuts**
  and bounded-rank hypergraph k-cuts: **rank-*r* H *k*-cuts**
- Submodular costs: **S *k*-cuts**
  and symmetric submodular costs: **SS *k*-cuts**

Talk Overview

1. Graph and Hypergraph *k*-cuts
2. Main Solution Approaches
3. Why Does Submodularity Help?
4. Solved and Unsolved Problems

1. Graph k-cuts

Given graph $G = (V, E)$ and $S \subseteq V$, the **cutset**

$$\delta(S) = \{e \in E : e \cap S \neq \emptyset \neq e \cap (V \setminus S)\}$$

For edge costs $c_e \geq 0$, the **cut function** $c(S) = \sum_{e \in \delta(S)} c_e$ is symmetric (i.e., $c(S) = c(V \setminus S)$) and submodular

Graph k-cut (k-way cut, multiway cut) problem, G k-cut: given the weighted graph (G, c), find a k-cut $\boldsymbol{P}$ minimizing $c(\boldsymbol{P}) = \sum_{i=1..k} c(P_i)$

Therefore: G k-cut $\subset$ SS k-cut $\subset$ S k-cut

Theorem [Goldschmidt and Hochbaum, 1994]
G k-cut is NP-hard when k is part of the input

- We will focus on the case of fixed k ("small" k)

Hypergraph k-cuts

Given hypergraph $H = (V, E)$ and a k-cut $\boldsymbol{P}$ of V, let

$$\Delta(\boldsymbol{P}) = \{e \in E : \; e \not\subseteq P_i \text{ for any } i \in \{1,\ldots,k\}\}$$

the set of edges cut by $\boldsymbol{P}$

- H k-cut problem: given hypergraph $H = (V, E)$ and edge costs $c_e \geq 0$, find a k-cut $\boldsymbol{P}$ minimizing $h(\boldsymbol{P}) = \sum_{e \in \Delta(\boldsymbol{P})} c_e$

Hypergraph $H = (V, E)$ is **rank-*r*** if $|e| \leq r$ for all $e \in E$

- A rank-2 hypergraph is a graph
  - in this case $h(\boldsymbol{P}) = ½\, c(\boldsymbol{P})$, hence G k-cut = rank-2 H k-cut

Remark: the hypergraph cut function $c(S) = \sum_{e \in \delta(S)} c_e$

where $\delta(S) = \{e \in E : \; e \cap S \neq \emptyset \neq e \cap (V \setminus S)\}$ is symmetric

and submodular, but $h(\boldsymbol{P}) \neq ½\, c(\boldsymbol{P})$, unless k or $\rho\,(H) = 2$

- an edge that intersects $t \geq 2$ parts in $\boldsymbol{P}$ contributes $t \cdot c_e$ to $c(\boldsymbol{P})$

Hypergraph k-cuts (2)

Hypergraph *k*-cut is a submodular *k*-cut problem [Okumoto, Nagamochi & Ibaraki 2012]

- Choose a root $v(e) \in e$ for each edge $e \in E$
- For every $S \subseteq V$ let

$$f_v(S) = \Sigma\{c_e : v(e) \in S \text{ and } e \setminus S \neq \emptyset\}$$

Then f_v is submodular, but not symmetric, and for every k-cut $\boldsymbol{P}$ we have

$$h(\boldsymbol{P}) = f_v(\boldsymbol{P}) = \sum_{i=1..k} f_v(P_i)$$

Therefore, we also have:

G k-cut $\subset$ rank-r H k-cut $\subset$ H k-cut $\subset$ S k-cut

(for every fixed $r \geq 2$)

Unconstrained *k*-cut Problems

k	2	3	4	5	≥ 6	variable
G	√	√	√	√	√	X
rank-r H	√	√	√	√	√	X
H	√	√				X
SS	√	√	√			X
S	√	√				X

- $k = 2$: Symmetric SFM (and more specialized algorithms)
- G k-cuts: [Goldschmidt & Hochbaum, 1994] and many others
- rank-r H k-cuts: [Fukunaga, 2010]
- H 3-cuts: [Xiao, 2010]
- S 3-cuts: [Okumoto, Fukunaga & Nagamochi, 2012]

k-terminal Problems

Given k terminals $t_1, \ldots, t_k$ in V, a ***k*-terminal cut** is a k-cut $\boldsymbol{P}$ such that $t_i \in P_i$ for all $i = 1, \ldots, k$

- also known as ***t-u* cuts** when $k = 2$, $t_1 = t$ and $t_2 = u$

The optimum 2-terminal cut problem solvable in polytime

- as SFM when cost is submodular
- many specialized algorithms
- also size-constrained versions, for any *fixed* minimum part-size vector $\boldsymbol{s}$

Theorem [Dalhaus, Johnson, Papadimitriou, Seymour and Yannakakis, 1994]

The k-terminal cut problem in a graph is NP-hard for every fixed $k \geq 3$

2. Main Solution Approaches

1-a. Fix-a-Part [GH] (Goldschmidt & Hochbaum 1994)

- enumerate a polynomial number of **candidate parts** $C_1, \ldots, C_p$
- for each C_i solve the $(k-1)$-cut problem on $V \setminus C_i$

- preserves the graph, hypergraph (and rank), and submodularity structures
- but does not preserve symmetry

1-b. Divide and Conquer [Kamidoi, Wakabayashi & Yoshida, 2002]:

- enumerate a polynomial number of **candidate divisions** $(V_1, k_1), \ldots, (V_p, k_p)$
- for each (V_i, k_i) solve the k_i-cut problem on V_i and the $(k-k_i)$-cut problem on $V \setminus V_i$

- Leads to a more balanced enumeration tree
  - [Xiao 2008] $p = O(n^{2k-1})$ and each $k_i = \lfloor k/2 \rfloor$ for G k-cut

Divide and Conquer

Which candidate parts?

A. (*S*,*T*)-cut based methods:

Theorem [GH]: For G k-cut where $|V| \geq 2k - 1$, there exists an optimum $\boldsymbol{P}$ in which either (i) the cheapest part P_1 has at most $k - 2$ vertices (2 vertices, when $k = 3$) or (ii) $\delta(P_1)$ is the edge set of a maximal minimum cut separating a set S of $k - 1$ vertices from another set T of k vertices

- Many variations for G k-cuts

Remark: [Queyranne, Aussois 1999] claimed that this Theorem and the resulting polytime algorithm extend to SS k-cut, but

  - only proved for SS 3-cut and SS 4-cut

Divide and Conquer (2)

Which candidate parts?

B. 2-cut ranking methods:

- enumerate $C_1, \ldots, C_p$ so $c(C_1) \leq c(C_2) \leq \ldots \leq c(C_p)$ (excluding complements) until $k \cdot c(C_p) \geq c(\boldsymbol{P'})$ for some k-cut $\boldsymbol{P'}$ [Nagamochi, Katayama & Ibaraki 2000]

How?

- Lawler-Murty ranking approach [Hamacher, Picard & Queyranne 1984; Vazirani & Yannakakis 1992]

Until when?

- [Burlet & Goldschmidt 1997] G 3-cut:
  - until $c(C_p) \geq 4/3\ c(C_1)$, so $p = O(n^2)$
- [Nagamochi & Ibaraki 1999] G 3-cut and G 4-cut:
  - until C_p crosses one of $C_1, \ldots, C_{p-1}$, so $p = O(n)$

Main Solution Approaches (2)

2. Contraction methods

- enumerate a polynomial number of **contractible subsets** C_1, …, C_p
- for each C_i solve the k-cut problem on $V|C_i$ where C_i has been contracted to a single node

- preserves **all** graph, hypergraph and rank, symmetry, and submodularity structures
- but requires each $|C_i| \geq 2$ for making progress
  - leads to minimum-size constraints

2-a. Deterministic contraction

- For 2-cuts:
  - [Nagamochi & Ibaraki 1992] G 2-cut, using MA-orderings
  - [Queyranne 1998]: extension to SS 2-cut

Contraction Methods

2-a. **Deterministic Contraction** (continued)

- For 3-cuts:

  Theorem [Xiao 2010] For H 3-cut where $|V| \geq 6$, there exists an optimum $\boldsymbol{P}$ such that either (i) a part of $\boldsymbol{P}$ is a singleton, or (ii) no optimum (2,2)-size 2-cut crosses $\boldsymbol{P}$

  - [Okumoto, Fukunaga & Nagamochi 2012]: extend to S 3-cut
- SS 4-cuts:
  - We prove that a similar result to Xiao's Theorem holds for SS 4-cut using an optimum **(2,2,1)-size 3-cut**

2.b. **Randomized Contraction**

- [Karger 1993, Karger & Stein 1996] G 2-cut:
  - `while` $|V| > 2$ `do`: choose $C_1 = e \in E$ with probability $c_e / c(E)$ and contract e to a single node
  - expand the final 2-node graph back to a 2-cut $\boldsymbol{P}$ in G and output it

  Theorem: The probability of obtaining *any given* optimum 2-cut $\boldsymbol{P}$ is $\Omega(n^{-2(k-1)})$

  $\Rightarrow$ $O(n^{2(k-1)} \log n)$ independent runs yield $\boldsymbol{P}$ with high probability

Contraction Methods (2)

2-b. **Randomized Contraction** (continued)

- [Chekuri & Korula 2010] H k-cuts
- For ***s*-size** H k-cuts:
  - **`while`** $|V| > \sigma = \sum_{i=1..k} s_i$ and $c(E) > 0$ **`do`**:
    - choose $C_1 = e \in E$ with probability $c_e / c(E)$
    - contract e (or $|V| - \sigma - 1$ nodes from e if e is too large)
  - choose uniformly at random an ***s***-size k-cut in the final contracted hypergraph,, expand it and output it

Theorem: If H is connected, rank-r and $|V| \geq \sigma$, where σ and r are fixed, then the probability of outputting any given optimum ***s***-size k-cut $\boldsymbol{P}$ is $\Omega(n^{-r\sigma})$

- $O(n^{r\sigma} \log n)$ independent runs yield $\boldsymbol{P}$ w.h.p.
- Proof of Theorem relies on an **Upper Bound Lemma**:

  If $|V| \geq \sigma$ then

$$OPT \leq \sum_{e \in E} c_e \left(1 - \prod_{j=1}^{|e|} \left(1 - \frac{\sigma_{k-1} - j + 1}{|V| - j + 1}\right)\right) \leq c(E) \left(1 - \left(1 - \frac{\sigma_{k-1}}{|V|}\right)^r\right)$$

Main Solution Approaches (3)

3. Tree Packing methods

- [Thorup 2008] G k-cut:
  - construct a greedy tree packing $\boldsymbol{T}$
    - a sequence $(T_1, \ldots, T_p)$ of spanning trees
    - each T_i is MST w.r.t. edge utilization $L_e^{\boldsymbol{T}-}/c_e$ where $\boldsymbol{T}^- = (T_1, \ldots, T_{i-1})$ and the load $L_e^{\boldsymbol{T}-}$ is the number of trees in $\boldsymbol{T}^-$ containing edge e

  Theorem: If $p = 48 \cdot |E| \cdot k \ \ln(4 \cdot |V| \cdot |E| \cdot k)$ then every optimum k-cut is crossed at most $k - 2$ times by some tree in $\boldsymbol{T}$

- [Fukunaga 2010]: extension to rank-r H k-cut, using hypertree packings

3. Why Does Submodularity Help?

1. **Uncrossing**: let f be a submodular set function
   a. **Uncrossing two sets**: The submodular inequality may be interpreted as an *uncrossing property*:

      if $f(A) \leq f(A \cap B)$ and $f(B) \leq f(A \cup B)$ then

      $f(A) = f(A \cap B)$ and $f(B) = f(A \cup B)$

   b. **Uncrossing a set and disjoint subsets**:

      if $B_1, \ldots, B_p$ are pairwise disjoint then

$$f(A) + \sum_{i=1}^{p} f(B_i) \geq f\left(A \cup \bigcup_{i=1}^{p} B_i\right) + \sum_{i=1}^{p} f(A \cap B_i)$$

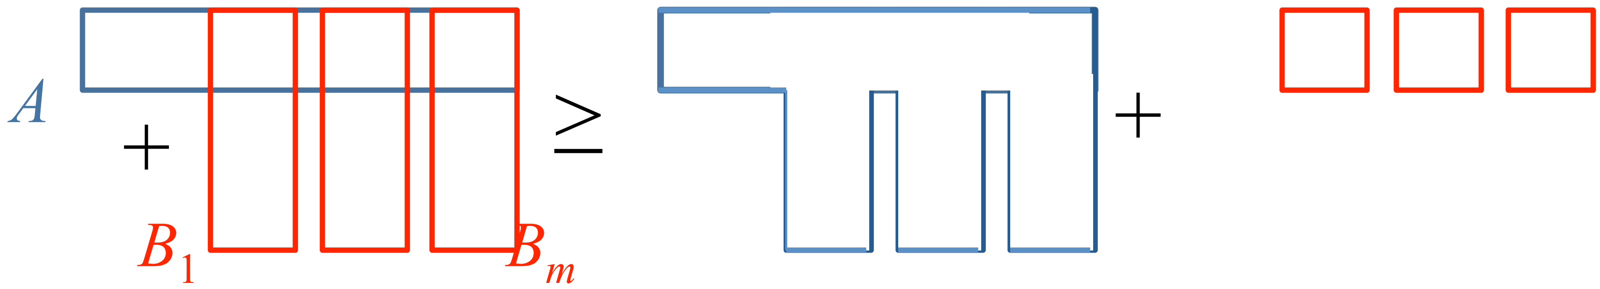

Uncrossing Partitions

c. Uncrossing two partitions

Extending [Xiao 2010] and [Okumoto, Fukunaga & Nagamochi 2012]

Let $\boldsymbol{X} = \{X_1, \ldots, X_h\}$ be an h-cut and $\boldsymbol{Y} = \{Y_1, \ldots, Y_k\}$ a k-cut, and let $Z_{ij} = X_i \cap Y_j$. Then

$$\boldsymbol{X}' = (Z_{11} \cup \bigcup_{j=2}^{k} Y_j, Z_{21}, \ldots, Z_{h1}) \text{ and } \boldsymbol{Y}' = (Z_{11} \cup \bigcup_{i=2}^{h} X_i, Y_{12}, \ldots, Y_{1k})$$

satisfy $f(\boldsymbol{X}) + f(\boldsymbol{Y}) \geq f(\boldsymbol{X}') + f(\boldsymbol{Y}')$

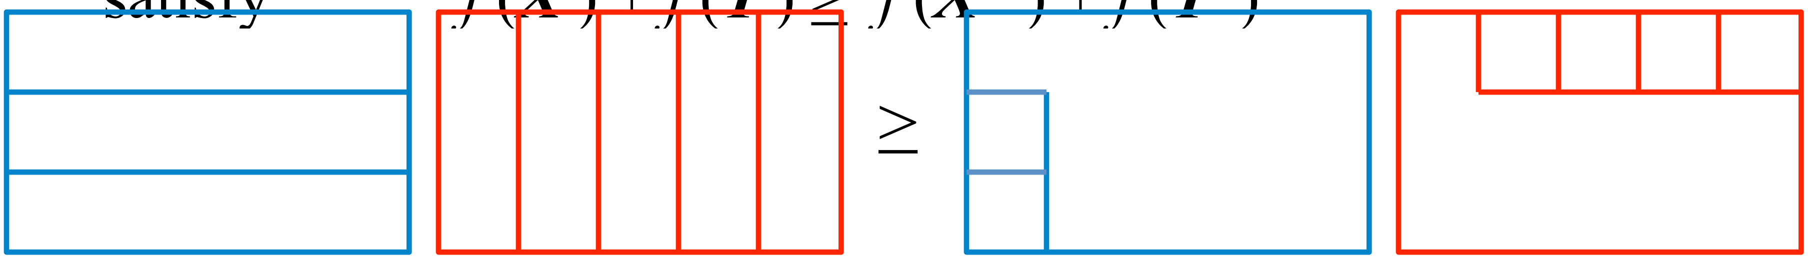

(only useful when all these parts are nonempty)

Uncrossing Partitions (2)

We also use a **Nested Extension** whereby, the last part created is repeatedly subdivided along alternating directions:

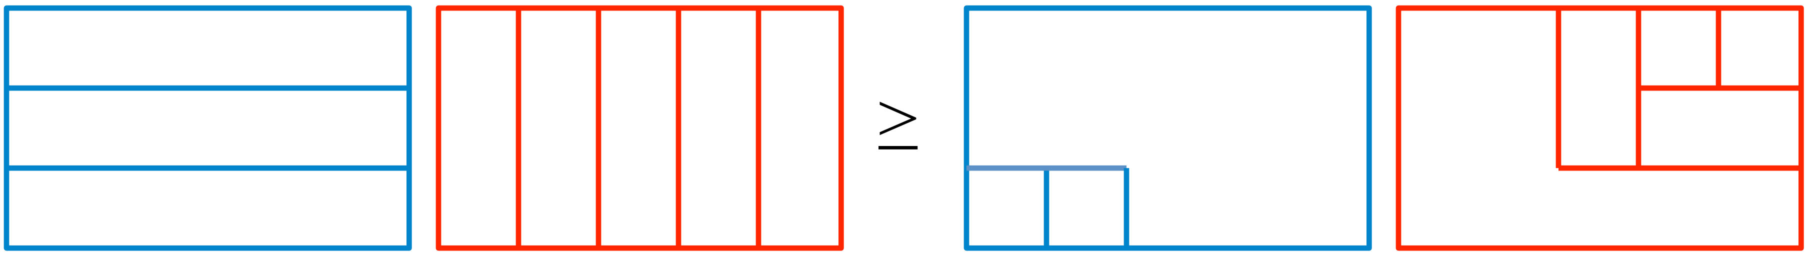

and related constructions when certain Z_{ij}'s are empty:

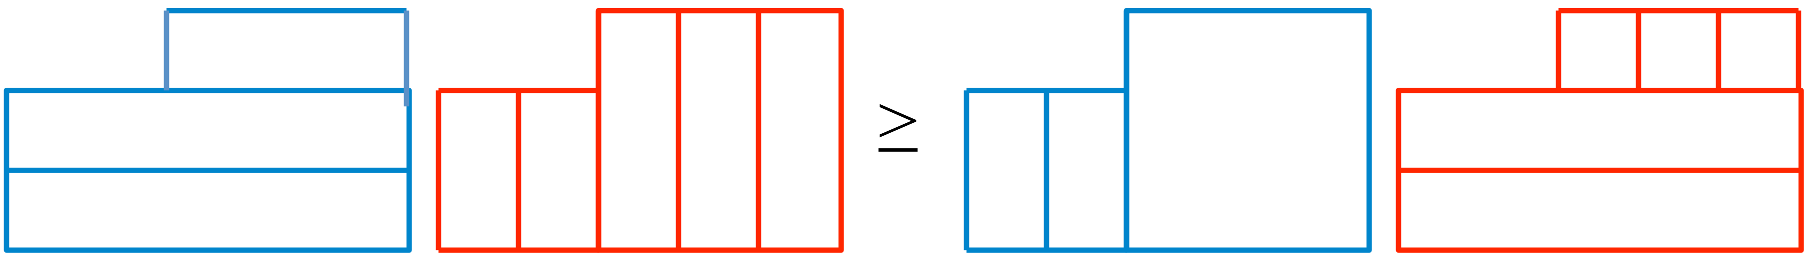

Why Does Submodularity Help? (2)

We use these partition uncrossing inequalities to prove the following reduction:

S 4-cut $\leftarrow$ S (2,2,1) 3-cut

(and also S 4-cut $\leftarrow$ S (2,2,2) 3-cut)

2. When f is **symmetric** and submodular, it is also **posimodular**, i.e., for all A, B,

$$f(A \setminus B) + f(B \setminus A) \le f(A) + f(B)$$

Constrained 3-cuts

111 ⟷ $\quad$ t,u,v 111

211

221

222 ⟷

1 terminal $\quad$ 2 terminals

- (1,1,1) 3-cut is just 3-cut: solved in polytime when f is submodular [Okumoto & al. 2012]
- t,u,v (1,1,1) 3-cut is the 3-terminal cut problem: NP-hard even for graphs [Dalhaus & al. 1994]

⟷ trivial equivalence (t must be in one part)

Constrained 3-cuts (2)

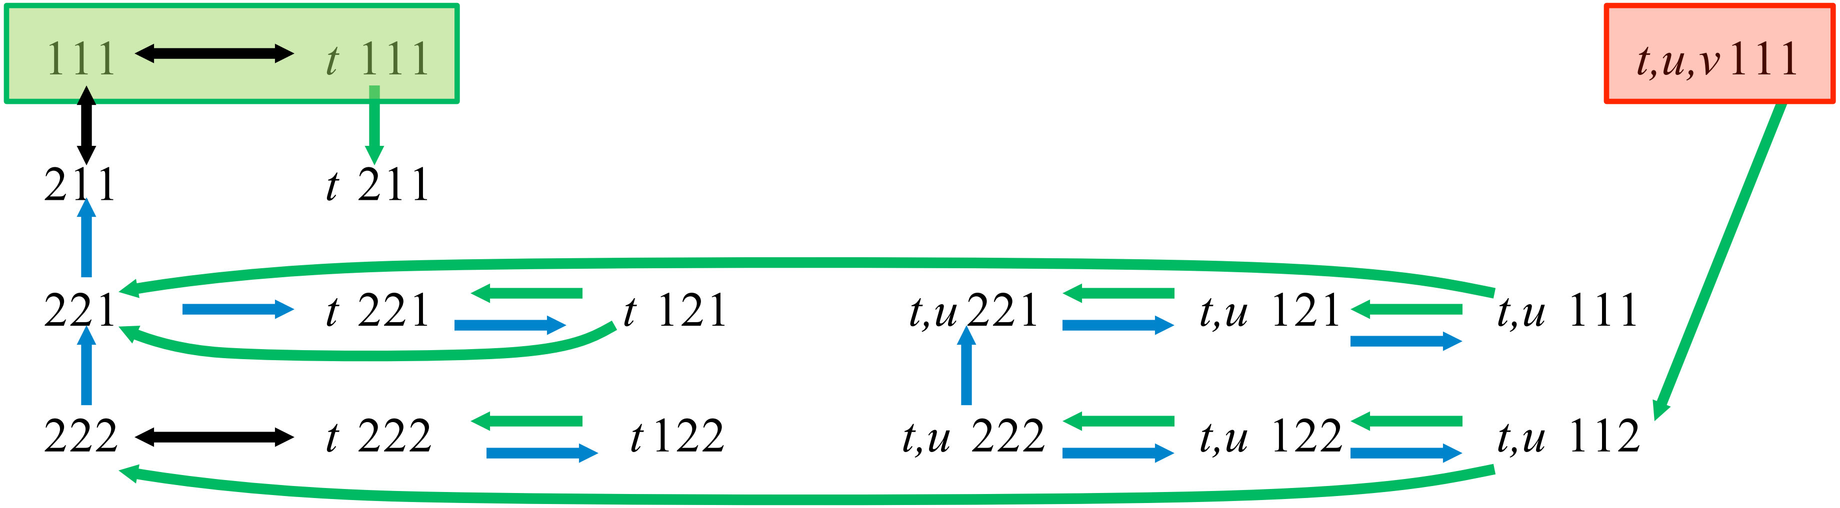

- another trivial reduction, when $n \geq 4$
- reduction by contracting a polynomial number of pairs of vertices
- reduction by enumerating singletons (fix-a- part) and solving 2-cut problems

Constrained 3-cuts (3)

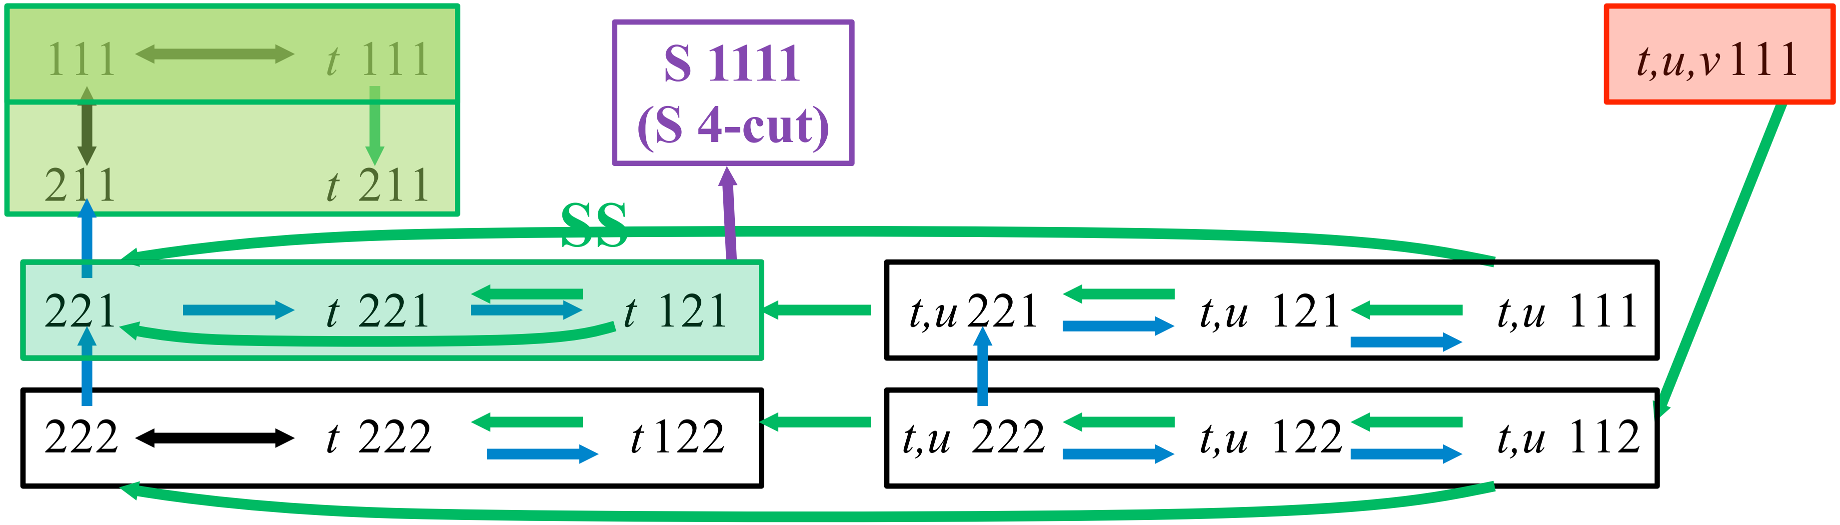

- 6 classes of polynomially equivalent problems
- reduction S 4-cut ← S (2,2,1) 3-cut
- (2,2,1) 3-cut solved in polytime when *f* is **symmetric** and submodular

4. Solved and Unsolved Problems

k	2	3	4	5	≥ 6	variable
G	√	√	√	√	√	X
rank-r H	√	√	√	√	√	X
H	√	√				X
SS	√	√	√			X
S	√	√				X

SS 4-cut solved in polytime by 3 different approaches:

1. (S,T)-cut based fix-a-part, extending Goldschmidt & Hochbaum
2. 2-cut ranking based fix-a-part, from [Nagamochi & Ibaraki 2000]
3. Contraction, based on finding an optimum (2,2,1) 3-cut

In all cases we reduce to *nonsymmetric* S 3-cuts [Okumoto & al., 2012]

Some Open Questions

- Minimal unsolved problems:
  - **Hypergraph 4-cut** (without bounded rank assumption)
  - **Symmetric submodular 5-cut**
- Related questions:
  - **Part-size constrained S** (and **SS**) **3-cut**
  - **2-terminal *k*-way cut**
  - Does **Randomized Contraction** work for (***s***-size) **H *k*-cut** *without* the bounded rank assumption?
  - *Deterministic* polytime algorithm for ***s*-size H *k*-cut**?
  - Can Graph (or Bounded-Rank Hypergraph) *k*-cuts be **ranked** in polytime?
    - the Lawler-Murty approach requires solving NP-hard *k*-terminal problems