

IBM Research Report

Games of Connectivity

Pradeep Dubey
pdubey@notes.cc.sunysb.edu
Center for Game Theory in Economics
Stony Brook University
& Cowles Foundation, Yale University (USA)
and
Rahul Garg
grahul@in.ibm.com
IBM India Research Lab
New Delhi (INDIA)

IBM Research Division,
IBM India Research Lab,
Hauz Khas, New Delhi - 110016. INDIA.
Phone: +91-11-2686-1100, Fax: +91-11-2686-1555

IBM Research Division

Almaden - Austin - Beijing - Delhi - Haifa - T.J. Watson - Tokyo - Zurich

LIMITED DISTRIBUTION NOTICE: This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties). Copies may be requested from IBM T.J. Watson Research Center, Publications, P.O. Box 218, Yorktown Heights, NY 10598 USA (email: reports@us.ibm.com). Some reports are available on the internet at <http://domino.watson.ibm.com/library/CyberDig.nsf/home>

Games of Connectivity

Pradeep Dubey*

Rahul Garg[†]

Abstract

We consider a communications network in which users transmit beneficial information to each other at a cost. We pinpoint conditions under which the induced cooperative game is convex. Our analysis is in a lattice-theoretic framework, which is at once simple and able to encompass a wide variety of seemingly disparate models.

Keywords: information lattice, multicast/unicast transmission, cooperative games, Shapley value, convex games.

JEL Classification: C71, D82, L96.

1 Introduction

The Shapley value [3] constitutes a scheme for the fair division of the benefits in a cooperative game. Unfortunately it is not always “stable” in that some coalitions may have incentive to break away because they can obtain more on their own than what the Shapley value awards them. In other words, the Shapley value can fail to be in the core of the game.

In a seminal paper [4], Shapley identified the class of “convex” games in which the Shapley value is not only in the core, but is the “center of gravity” of the core. Such games are those that exhibit increasing returns to cooperation: the marginal contribution of a player to a coalition goes up as the coalition is enhanced.

In this paper we pinpoint conditions under which certain games of connectivity are convex. Players in our model are located at the vertices of a communications network and can stand to gain a lot by sharing disparate bits of information that they initially hold. Indeed information is more amenable to sharing than standard commodities. Commodities are typically lost to the person who gives them away. Information in contrast has “the quality of mercy”, blessing him that gives and him that takes, since the giver retains all his information even as he sends it out. Nevertheless it is not automatic that all information will be shared. This is because, though costlessly duplicable, information may be costly to transmit (e.g., on account of setup costs of links in the communications network). Any coalition must do a careful cost-benefit analysis, choosing that pattern of information transmission which minimizes the total net benefit to its members.

* (pdubey@notes.cc.sunysb.edu) Center for Game Theory in Economics, Stony Brook University and Cowles Foundation, Yale University (USA).

[†] (grahul@in.ibm.com) IBM India Research Lab, New Delhi (India).

It should be pointed out that our model, though much more general, takes its cue from a multicast transmission game presented in [2]. There, too, the Shapley value was examined but the focus was on its computation. Somewhat surprisingly, it was overlooked that the game is convex. But the convexity of the game is not besides the point. It ensures, as was said, that the Shapley value is not only fair but also stable, further bolstering the plausibility of the Shapley value as a solution concept in such games.

An important feature of our approach is that we formulate information in terms of a *lattice*. This leads to a framework that is at once universal and simple. We can encompass a wide variety of seemingly different models, involving unicast and multicast modes of transmission, setup and variable costs in the communications network, and information that comes in various guises (from finite dimensional vectors, to partitions of a set, to layered encoding). The lattice framework makes for a remarkably transparent analysis in all cases.

The paper is organized as follows. In Section 2 we present some motivating examples, starting with the model in [2]. The abstract lattice-theoretic framework is presented in Section 3. In Section 4 we establish our main result which states that games of connectivity are convex. Section 5 points out a monotonicity property of optimal transmissions. Finally, in Section 6, we show how to fit the examples into our lattice-theoretic framework; and we also examine the tightness of our assumptions and indicate some generalizations of the model.

2 Examples

We present a series of examples of information transmission in a network, all of which yield convex games.

2.1 Multicast transmission

First let us recall the game presented in [2]. There is a finite tree Γ with a sender δ located at its root and a distinct receiver at each leaf (terminal vertex). Any receiver can get information from δ if he is connected to δ using the edges of Γ . The tree Γ is viewed as a digital network which carries a public broadcast by δ , and it is assumed that information flowing into any vertex of the tree can be costlessly duplicated and sent out (multicast) on any subset of the outgoing edges. But the edges of Γ do have setup costs associated to them. Offsetting these costs are benefits $B(\alpha)$ to α when he receives information from δ .

A cooperative game is induced on the player-set N of receivers in a natural manner. Any coalition $S \subset N$ can use an arbitrary subtree Γ' of Γ at the cost $C(\Gamma')$ of all the edges of Γ' . The benefit S derives from Γ' is $B(S, \Gamma') = \sum_{\alpha} B(\alpha)$, where the summation runs over all α in S which are connected to δ via Γ' . Thus the “worth” $w(S)$ of coalition S (i.e., the most S can guarantee to itself) is obtained by maximizing the net benefit $B(S, \Gamma') - C(\Gamma')$ over all possible subtrees Γ' .

The problem addressed in [2] is how to divide the total benefit $w(N)$ among the receivers in N . The Shapley value of the game w is proposed and the focus is on its computation.

Generalizations of the model come immediately to mind. There can be several senders located at different vertices of the tree, each with its own distinctive information to transmit. Moreover

not all senders need be “dummies” as in [2]. Some of them could be bona fide players in the game with the power to withhold their information. One could also imagine them to have different transmission trees, possibly with significant overlap.

In spite of these complications, the game remains convex and so the Shapley value continues to be centrally located in the core (but its computation may no longer be as felicitous as in [2]).

2.2 Unicast transmission

Imagine a set of users connected to each other through a hierarchical network (as in telephony). In particular suppose they are located on the leaves of a tree Γ with other vertices acting as relays. But the communication is private rather than public, and the users transmit information to each other on a one-to-one basis.

The user at leaf α can choose the amount of information $\tau_{\alpha\beta} \in [0, m], m > 0$, to be sent to β . The total benefit derived at β is $\sum_{\alpha} B_{\alpha\beta}(\tau_{\alpha\beta})$, where $B_{\alpha\beta}$ is an arbitrary non-decreasing function. As before, it costs to use the tree. Each edge now has not only a setup cost, but also an arbitrary non-decreasing variable cost for every α -to- β flow on it (additive across flows).

This unicast scenario also gives rise to a cooperative game in an obvious way. Any coalition S chooses $\tau = \{\tau_{\alpha\beta} : \alpha \in S, \beta \in S\}$, and a subforest of Γ to carry τ , so as to maximize the net benefit.

It turns out that this game is also convex.

2.3 Transmission of layered information

We turn to a situation where information is encoded or organized in layers (e.g., as in a video transport system, see [5]). To be precise, suppose layer L_i consists of “information bricks” numbered by integers $m_{i-1} + 1, m_{i-1} + 2, \dots, m_i$. The bricks in $L = \cup_{i=1}^k L_i$ are, however, distributed arbitrarily among the n players located at the vertices of a communication tree Γ , with no duplication. So, denoting by Σ_{α} the set of bricks held at vertex α , we have $\Sigma_{\alpha} \cap \Sigma_{\beta} = \phi$ if $\alpha \neq \beta$. Players wish to receive bricks in order to build a “knowledge pyramid”, but they cannot construct layer L_i unless all previous layers L_1, L_2, \dots, L_{i-1} are in place. Of course, since these bricks are not standard commodities but signify information, no player loses any of his own bricks by sending them to others. The player at vertex α may transmit any subset $Q_e \subset \Sigma_{\alpha}$ on any edge e emanating from α . Then for any edge e' that follows from e , he can send $Q_{e'} \subset Q_e$, and so on. In short he can contemplate multicast transmission on Γ with α as the root.

There is a set-up cost for every edge e as earlier, and additional flow costs $C_{e,\alpha}(x)$ for $x \in \Sigma_{\alpha}$.

Benefits accrue as follows. Denoting by $Q_{\beta\alpha} \subset \Sigma_{\beta}$ the subset of bricks that α receives from β , the benefit to α is $f_{\alpha}(n)$, where

$$n = \max\{j : L_i \subset \Sigma_{\alpha} \cup (\cup_{\beta} Q_{\beta\alpha}) \forall i \leq j\}$$

and $f(n)$ is an arbitrary non-decreasing function.

The idea here, as was said, is that information is organized in pyramidal form. Information of layer L_i is not usable unless all layers L_1, L_2, \dots, L_i are complete.

The cooperative game, arising in this setup, is once again convex.

2.4 Transmission of information partitions

Let $Q = \{1, 2, \dots, k\}$. Let $\{Q_\alpha : \alpha \in V\}$ be a partition of Q . (Here V denotes the set of vertices of Γ .) Further let P_α be a partition of Q_α . The information held at vertex α is the partition $\{P_\alpha, Q \setminus Q_\alpha\}$ of Q . Each vertex can transmit its partition (or any coarsening thereof) to other vertices prior to the production stage. If vertex α winds up with the partition P of Q , it can earn (via production):

$$\begin{aligned} & \max f_\alpha(x_1, x_2, \dots, x_k) \\ & \text{Subject to: } x_i \leq a_i(\alpha) \\ & \quad \quad \quad x_i \geq 0 \\ & \text{and } i \sim_P j \Rightarrow x_i = x_j \end{aligned}$$

where $i \sim_P j$ means that i and j are in the same cell of the partition P . We assume that the production function f_α is supermodular on R_+^k , i.e., (assuming differentiability):

$$\frac{\partial}{\partial x_i} \frac{\partial f_\alpha}{\partial x_j} \geq 0$$

for all i, j and α . In other words the inputs x_1, x_2, \dots, x_k are complementary goods.

When a coalition S forms, its members can transmit information to each other through any subforest of Γ after paying the setup costs.

This, too, induces a game that is convex.

2.5 General network with controlled edges

Let G be an arbitrary undirected graph with edge set E and vertex set V . For each vertex $\alpha \in V$, let $\Gamma(\alpha) \subset G$ be a tree rooted at α on which α is constrained to transmit its information. Further suppose that edges of G are subject to the control of coalitions.

Thus when a coalition S forms, each $\alpha \in S$ has access to only those edges in $\Gamma(\alpha)$ whose controllers are contained in S .

In this setup, players who are neither senders nor receivers of information, may nevertheless have a vital role to play in the game on account of their control of edges (such as cable operators or monopoly network providers).

All of our preceding examples can be embedded in this larger framework. The games induced will still be convex.

We build, in the next section, an abstract lattice-theoretic model of information and its transmission, which unifies all of the above (and more) examples and makes for a particularly transparent analysis.

3 The Abstract Model

3.1 The communications network

Let $G = (V, E)$ be a graph where V is a finite set of vertices and E is a set of undirected edges.

For every $\alpha \in V$ there is a tree $\Gamma(\alpha) \equiv (V(\alpha), E(\alpha)) \subset G$, rooted at α , that can be used by α to transmit its information to other vertices.

3.2 Information

Information is modeled as a lattice \mathcal{L} with \geq denoting the partial order and \vee, \wedge the join and the meet operators¹. We assume that $0 \equiv \wedge \{x : x \in \mathcal{L}\}$ exists in \mathcal{L} and that \wedge distributes over \vee , i.e.,

$$x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$$

for all $x, y, z \in \mathcal{L}$. This property holds in a variety of contexts and is well-known (see [1]).

The canonical examples we have in mind is that \mathcal{L} is the power set of a finite set with \geq corresponding to the set-theoretic notion of \supseteq ; or that \mathcal{L} is the set of all partitions of a finite set with \geq corresponding to refinement; or that \mathcal{L} is a closed interval of the real line with \geq corresponding to the standard order; or that \mathcal{L} is the product lattice of finitely many such lattices. In all of these cases 0 exists in \mathcal{L} and the distributive property holds.

Any vertex $\alpha \in V$ can transmit information from a sub-lattice $\mathcal{L}(\alpha)$ of \mathcal{L} . A key assumption we make is that the information held at different vertices is disjoint, i.e.,

$$x \in \mathcal{L}(\alpha), y \in \mathcal{L}(\beta), \alpha \neq \beta \Rightarrow x \wedge y = 0$$

We also assume that each vertex can opt to send no information, i.e., $0 \in \mathcal{L}(\alpha)$ for all $\alpha \in V$.

3.3 Location of players and public facilities

Let $N = \{1, 2, \dots, n\}$ be the set of players. There is an additional dummy player, labeled $n + 1$, used to model public facilities available to all players in N . Denote $\tilde{N} = N \cup \{n + 1\}$.

Each vertex is occupied by a player² as specified by a *location map*

$$\eta : V \rightarrow \tilde{N}$$

where $\eta(\alpha)$ denotes the player (possibly, dummy) at vertex α . Let $V(S)$ represent the set of all the vertices occupied by players in $S \cup \{n + 1\}$ i.e.,

$$V(S) = \{\alpha \in V : \eta(\alpha) \in S \cup \{n + 1\}\}$$

3.4 Control of edges

Edges are controlled by coalitions of players in accordance with a *control map*

$$\kappa : E \rightarrow 2^N$$

where $\kappa(e)$ denotes the coalition that controls³ the use of edge e . (If $\kappa(e) = \emptyset$, then e is accessible to everyone.)

¹Recall (see e.g. [1]) that for any x and y in \mathcal{L} , there exists a greatest lower bound w.r.t. \geq (denoted $x \wedge y$) and a least upper bound (denoted $x \vee y$).

²The case where several players occupy a vertex is included in our set-up (see remark 3 in Section 6).

³A natural case: if $e = (\alpha, \beta)$, then $\kappa(e) = (\eta(\alpha) \cup \eta(\beta)) \cap N$.

3.5 The transmission of information

Each vertex α can transmit information $x \in \mathcal{L}(\alpha)$ to other vertices on its tree $\Gamma(\alpha) \equiv (V(\alpha), E(\alpha))$. Aggregating across vertices, the total transmission may be viewed as a map $\tau : E \times V \rightarrow \mathcal{L}$ with the interpretation that $\tau(e, \alpha)$ is the information transmitted by the vertex α on the edge e . Some natural conditions must be imposed on this map τ . Any vertex α can send information only out of $\mathcal{L}(\alpha)$ i.e.,

$$\tau(e, \alpha) \in \mathcal{L}(\alpha) \quad (1)$$

for all $\alpha \in V$ and $e \in E(\alpha)$. Moreover, no vertex α can send any (except null) information on edges outside its tree i.e.,

$$\tau(e, \alpha) = 0 \text{ if } e \notin E(\alpha) \quad (2)$$

for all $\alpha \in V$ and $e \in E$. Finally, the join of all the information of α that flows out of a vertex must be no more than the information of α that arrives at it, i.e.,

$$\tau(e, \alpha) \geq \vee \{ \tau(e', \alpha) : e' \in F(e, \alpha) \} \quad (3)$$

for all $\alpha \in V$ and $e \in E(\alpha)$, where $F(e, \alpha)$ denotes the set of immediate offspring edges of e in the tree $\Gamma(\alpha)$.

Let \mathcal{T} denote the set of all possible transmissions, i.e.,

$$\mathcal{T} = \{ \tau : E \times V \rightarrow \mathcal{L} : \tau \text{ satisfies (1), (2) and (3)} \}$$

The set \mathcal{T} itself forms a lattice under the natural definitions: $\tau \geq \tau'$ if $\tau(e, \alpha) \geq \tau'(e, \alpha)$ for all e, α ; $(\tau \vee \tau')(e, \alpha) = \tau(e, \alpha) \vee \tau'(e, \alpha)$ for all e, α ; $(\tau \wedge \tau')(e, \alpha) = \tau(e, \alpha) \wedge \tau'(e, \alpha)$ for all e, α .

For any coalition $S \subset N$, define the subset $\mathcal{T}(S) \subset \mathcal{T}$ of *transmissions feasible for S* as follows:

$$\mathcal{T}(S) = \{ \tau \in \mathcal{T} : \text{for any } e \text{ and } \alpha, \tau(e, \alpha) > 0 \Rightarrow \kappa(e) \subset S \text{ and } \alpha \in S \cup \{n+1\} \}$$

In other words, only members of S or public vertices can transmit information in $\mathcal{T}(S)$; and only the edges under the control of S may be used.

3.6 The reception of information

A transmission $\tau \in \mathcal{T}$ induces a reception $\sigma(\tau, \alpha) \in \mathcal{L}$ at every vertex $\alpha \in V$ as follows:

$$\sigma(\tau, \alpha) = (x^*(\alpha)) \vee (\vee \{ \tau(e(\beta, \alpha), \beta) : \beta \in V \setminus \{ \alpha \} \text{ and } \alpha \in \Gamma(\beta) \})$$

where $e(\beta, \alpha)$ is the edge coming into α from β in $\Gamma(\beta)$ and $x^*(\alpha) \equiv \vee \{ x : x \in \mathcal{L}(\alpha) \}$.

Here $x^*(\alpha)$ represents the maximum information in $\mathcal{L}(\alpha)$. Since α can costlessly receive its own information, and since information is valuable, we suppose that α always “sends” $x^*(\alpha)$ to itself. The total reception at α is obtained by joining $x^*(\alpha)$ with the bits of information $\tau(e(\beta, \alpha), \beta)$ sent to α by other vertices β ⁴.

⁴In our current setup of “passive” communication, we could drop the constant term $x^*(\alpha)$. But later we shall extend our model to encompass interactive communication wherein α 's own information may get enhanced on account of information received from other vertices, and then the inclusion of the term is vital.

3.7 The cost of a transmission

The cost of transmitting information (originating at different vertices) on any edge is given by⁵ $c_e : \mathcal{L}^V \rightarrow R_+$, where $c_e((x(\alpha))_{\alpha \in V}) \equiv$ the cost of the flow $(x(\alpha))_{\alpha \in V}$ on e . We postulate that c_e is *submodular* on \mathcal{L}^V , i.e.,

$$c_e(x \vee y) + c_e(x \wedge y) \leq c_e(x) + c_e(y)$$

for all $e \in E$ and $x, y \in \mathcal{L}^V$. Such costs can arise in several ways. For instance, suppose there is a set-up cost $f(e)$ for e , and a further set-up cost $f(e, \alpha)$ for every vertex α that uses e , i.e.,

$$c_e((x(\alpha))_{\alpha \in V}) = \begin{cases} 0, & \text{if } x(\alpha) = 0 \text{ for all } \alpha \\ f(e) + \sum_{x(\alpha) > 0} f(e, \alpha), & \text{otherwise} \end{cases}$$

It is evident that this cost function is submodular, and that it remains so if we add variable costs $\sum_{\alpha \in V} g_\alpha(x(\alpha))$ provided each $g_\alpha : \mathcal{L} \rightarrow R_+$ is itself submodular (i.e., evinces economy of scale).

The *cost of transmission* $\tau \in \mathcal{T}$ is the sum of the costs incurred on all the edges, i.e.,

$$C(\tau) = \sum_{e \in E} c_e((\tau(e, \alpha))_{\alpha \in V})$$

It is easy to verify that C is submodular on \mathcal{T} , i.e.,

$$C(\tau) + C(\tau') \geq C(\tau \vee \tau') + C(\tau \wedge \tau') \quad (4)$$

3.8 The benefit from a transmission

For every vertex $\beta \in V$, there is a benefit function $B_\beta : \mathcal{L} \rightarrow R_+$, where $B_\beta(x)$ represents the benefit to β from receiving information $x \in \mathcal{L}$. We assume that B_β is supermodular and non-decreasing for all $\beta \in V$ i.e.,

$$B_\beta(x \vee y) + B_\beta(x \wedge y) \geq B_\beta(x) + B_\beta(y)$$

and

$$x \geq y \Rightarrow B_\beta(x) \geq B_\beta(y)$$

The benefit to a coalition $S \subset N$ from transmission $\tau \in \mathcal{T}$ is given by

$$B(S, \tau) = \sum_{\beta \in V(S)} B_\beta(\sigma(\tau, \beta))$$

4 The Connectivity Game

We consider the coalitional game that arises from the communications network. A non-empty coalition $S \subset N$ can choose any $\tau \in \mathcal{T}(S)$ to transmit information between its members or to receive information from public vertices. The coalition obtains total benefit $B(S, \tau)$ but at a cost $C(\tau)$. The maximum net benefit that S can guarantee is therefore given by

$$w(S) = \max_{\tau \in \mathcal{T}(S)} B(S, \tau) - C(\tau)$$

⁵Note that \mathcal{L}^V is a finite product of \mathcal{L} with itself (V times) and is a product lattice.

(with $w(\phi)$ understood to be 0). We call w the *connectivity game*.

Recall that a game $u : 2^n \rightarrow R$ is called *convex* [4] if u is supermodular on the lattice 2^N , i.e.,

$$w(S \cup T) + w(S \cap T) \geq w(S) + w(T)$$

for all $S \subset N$ and $T \subset N$. (An equivalent condition is $w(T \cup \{i\}) - w(T) \geq w(S \cup \{i\}) - w(S)$ whenever $S \subset T \subset N$ and $i \notin T$.) Our main result is:

Theorem 1 *The connectivity game w is convex.*

For the proof we first establish some lemmas.

Lemma 1 *Let $S \subset N$, $T \subset N$, $\tau \in \mathcal{T}(S)$ and $\tau' \in \mathcal{T}(T)$. Then $\tau \vee \tau' \in \mathcal{T}(S \cup T)$ and $\tau \wedge \tau' \in \mathcal{T}(S \cap T)$.*

Proof: Since τ and τ' are in \mathcal{T} , $\tau(e, \alpha)$ and $\tau'(e, \alpha)$ are in $\mathcal{L}(\alpha)$. Since $\mathcal{L}(\alpha)$ is a lattice, $(\tau \vee \tau')(e, \alpha) \equiv \tau(e, \alpha) \vee \tau'(e, \alpha) \in \mathcal{L}(\alpha)$ and $(\tau \wedge \tau')(e, \alpha) \equiv \tau(e, \alpha) \wedge \tau'(e, \alpha) \in \mathcal{L}(\alpha)$.

Next, if $e \notin E(\alpha)$, then $\tau(e, \alpha) = \tau'(e, \alpha) = 0$ and therefore $(\tau \vee \tau')(e, \alpha) = 0$ and $(\tau \wedge \tau')(e, \alpha) = 0$ as well.

Finally, since

$$\tau(e, \alpha) \geq \vee \{\tau(e', \alpha) : e' \in F(e, \alpha)\} \tag{5}$$

$$\tau'(e, \alpha) \geq \vee \{\tau'(e', \alpha) : e' \in F(e, \alpha)\} \tag{6}$$

we have

$$\begin{aligned} (\tau \vee \tau')(e, \alpha) &= \tau(e, \alpha) \vee \tau'(e, \alpha) \\ &\geq \vee \{(\tau \vee \tau')(e', \alpha) : e' \in F(e, \alpha)\} \end{aligned}$$

from the associativity of \vee and the fact that $x \geq x'$ and $y \geq y'$ implies $x \vee y \geq x' \vee y'$. This shows that $\tau \vee \tau' \in \mathcal{T}$. Also from (5) and (6)

$$\begin{aligned} \tau(e, \alpha) \wedge \tau'(e, \alpha) &\geq (\vee \{\tau(e', \alpha) : e' \in F(e, \alpha)\}) \wedge (\vee \{\tau'(e', \alpha) : e' \in F(e, \alpha)\}) \\ &\geq \vee \{(\tau(e', \alpha) \wedge \tau'(e', \alpha)) : e' \in F(e, \alpha)\} \end{aligned}$$

The first inequality follows from the fact that $x \geq x'$ and $y \geq y'$ implies $x \wedge y \geq x' \wedge y'$; the second from the fact $(x \vee y) \wedge z \geq (x \wedge z) \vee (y \wedge z)$ and the commutativity and associativity of \vee, \wedge . This proves that $\tau \wedge \tau' \in \mathcal{T}$.

To check that $\tau \vee \tau' \in \mathcal{T}(S \cup T)$, observe that, for any e and α

$$\begin{aligned} &\tau(e, \alpha) \vee \tau'(e, \alpha) > 0 \\ \Rightarrow &\tau(e, \alpha) > 0 \text{ or } \tau'(e, \alpha) > 0 \\ \Rightarrow &\kappa(e) \subset S \text{ or } \kappa(e) \subset T \\ \Rightarrow &\kappa(e) \subset S \cup T \end{aligned}$$

To check that $\tau \wedge \tau' \in \mathcal{T}(S \cap T)$, observe that

$$\begin{aligned} & \tau(e, \alpha) \wedge \tau'(e, \alpha) > 0 \\ \Rightarrow & \tau(e, \alpha) > 0 \text{ and } \tau'(e, \alpha) > 0 \\ \Rightarrow & \kappa(e) \subset S \text{ and } \kappa(e) \subset T \\ \Rightarrow & \kappa(e) \subset S \cap T \end{aligned}$$

□

Lemma 2 For $S \subset N$, $T \subset N$, $\tau \in \mathcal{T}$ and $\tau' \in \mathcal{T}$,

$$B(S, \tau_1) + B(S, \tau_2) \leq B(S \cup T, \tau_1 \vee \tau_2) + B(S \cap T, \tau_1 \wedge \tau_2)$$

Proof: From the definition of σ and the associativity of \vee it is immediate that

$$\sigma(\tau \vee \tau', \alpha) = \sigma(\tau, \alpha) \vee \sigma(\tau', \alpha) \quad (7)$$

The analogous result holds for \wedge only when the lattice \mathcal{L} is distributive and the sub-lattices $\mathcal{L}(\alpha)$ are disjoint across $\alpha \in V$, as we now check.

Let α and β be two distinct vertices. Denote by $\rho(\tau, \alpha, \beta)$ the information that α receives from β in the transmission τ , i.e.,

$$\rho(\tau, \alpha, \beta) = \begin{cases} \tau(e(\beta, \alpha), \beta), & \text{if } \alpha \in \Gamma(\beta) \\ 0, & \text{otherwise} \end{cases}$$

where, recall $e(\beta, \alpha)$ is the edge coming into α from β in the tree $\Gamma(\beta)$. Then,

$$\sigma(\tau, \alpha) = x^*(\alpha) \vee (\vee \{\rho(\tau, \alpha, \beta) : \beta \in V \setminus \{\alpha\}\})$$

So,

$$\sigma(\tau, \alpha) \wedge \sigma(\tau', \alpha) = (x^*(\alpha) \vee (\vee \{\rho(\tau, \alpha, \beta) : \beta \in V \setminus \{\alpha\}\})) \wedge (x^*(\alpha) \vee (\vee \{\rho(\tau', \alpha, \beta) : \beta \in V \setminus \{\alpha\}\}))$$

By the distributivity of \wedge over \vee , and the commutativity and associativity of \wedge and \vee , the right hand side of the above equation simplifies to

$$x^*(\alpha) \vee (\vee \{\rho(\tau, \alpha, \beta) \wedge \rho(\tau', \alpha, \beta') : \beta \in V \setminus \{\alpha\}, \beta' \in V \setminus \{\alpha\}\})$$

Since the sub-lattices $\mathcal{L}(\beta)$ and $\mathcal{L}(\beta')$ are disjoint when $\beta \neq \beta'$ all the cross-terms in the above expression disappear, reducing it to

$$x^*(\alpha) \vee (\vee \{\rho(\tau, \alpha, \beta) \wedge \rho(\tau', \alpha, \beta) : \beta \in V \setminus \{\alpha\}\})$$

which obviously equals

$$x^*(\alpha) \vee (\vee \{\rho(\tau \wedge \tau', \alpha, \beta) : \beta \in V \setminus \{\alpha\}\})$$

proving that

$$\sigma(\tau \wedge \tau', \alpha) = \sigma(\tau, \alpha) \wedge \sigma(\tau', \alpha) \quad (8)$$

From the definition of the benefit function B ,

$$B(S, \tau) + B(T, \tau') = \sum_{\beta \in V(S)} B_\beta(\sigma(\tau, \beta)) + \sum_{\beta \in V(T)} B_\beta(\sigma(\tau', \beta))$$

By rearranging terms we get

$$\begin{aligned} B(S, \tau) + B(T, \tau') &= \sum_{\beta \in V(S) \setminus V(T)} B_\beta(\sigma(\tau, \beta)) + \sum_{\beta \in V(T) \setminus V(S)} B_\beta(\sigma(\tau', \beta)) \\ &+ \sum_{\beta \in V(S) \cap V(T)} (B_\beta(\sigma(\tau, \beta)) + B_\beta(\sigma(\tau', \beta))) \end{aligned} \quad (9)$$

From (7), (8) and the supermodularity of B_β we have

$$\begin{aligned} B_\beta(\sigma(\tau, \beta)) + B_\beta(\sigma(\tau', \beta)) &\leq B_\beta(\sigma(\tau, \beta) \vee \sigma(\tau', \beta)) + B_\beta(\sigma(\tau, \beta) \wedge \sigma(\tau', \beta)) \\ &= B_\beta(\sigma(\tau \vee \tau')) + B_\beta(\sigma(\tau \wedge \tau')) \end{aligned}$$

Therefore (9) becomes

$$\begin{aligned} B(S, \tau) + B(T, \tau') &\leq \sum_{\beta \in V(S) \setminus V(T)} B_\beta(\sigma(\tau, \beta)) + \sum_{\beta \in V(T) \setminus V(S)} B_\beta(\sigma(\tau', \beta)) \\ &+ \sum_{\beta \in V(S) \cap V(T)} (B_\beta(\sigma(\tau \vee \tau', \beta)) + B_\beta(\sigma(\tau \wedge \tau', \beta))) \\ &\leq \sum_{\beta \in V(S) \setminus V(T)} B_\beta(\sigma(\tau \vee \tau'), \beta) + \sum_{\beta \in V(T) \setminus V(S)} B_\beta(\sigma(\tau \vee \tau', \beta)) \\ &+ \sum_{\beta \in V(S) \cap V(T)} B_\beta(\sigma(\tau \vee \tau', \beta)) + \sum_{\beta \in V(S) \cap V(T)} B_\beta(\sigma(\tau \wedge \tau', \beta)) \\ &= B(S \cup T, \tau \vee \tau') + B(S \cap T, \tau \wedge \tau') \end{aligned}$$

(The last inequality follows from the fact that B_β is a non-decreasing function on \mathcal{L} for all $\beta \in V$).

□

Proof of Theorem 1: Let S and T be arbitrary coalitions of N . Let τ_1^*, τ_2^* be optimal transmissions for coalitions S, T respectively, i.e.,

$$v(S) = B(S, \tau_1^*) - C(\tau_1^*)$$

$$v(T) = B(T, \tau_2^*) - C(\tau_2^*).$$

From Lemma 2 and the fact that C is submodular (see (4)), we have

$$\begin{aligned} v(S) + v(T) &\leq B(S \cup T, \tau_1^* \vee \tau_2^*) - C(\tau_1^* \vee \tau_2^*) \\ &+ B(S \cap T, \tau_1^* \wedge \tau_2^*) - C(\tau_1^* \wedge \tau_2^*) \end{aligned} \quad (10)$$

Since τ_1^* is an optimal transmission for coalition S , $\tau_1^* \in \mathcal{T}(S)$. Similarly $\tau_2^* \in \mathcal{T}(T)$. By Lemma 1, $\tau_1^* \vee \tau_2^* \in \mathcal{T}(S \cup T)$ and $\tau_1^* \wedge \tau_2^* \in \mathcal{T}(S \cap T)$. But then,

$$w(S \cup T) \geq B(S \cup T, \tau_1^* \vee \tau_2^*) - C(\tau_1^* \vee \tau_2^*) \quad (11)$$

$$w(S \cap T) \geq B(S \cap T, \tau_1^* \wedge \tau_2^*) - C(\tau_1^* \wedge \tau_2^*) \quad (12)$$

Inequalities (10), (11) and (12) give

$$w(S) + w(T) \leq w(S \cup T) + w(S \cap T)$$

showing that the game w is convex.

5 The Growing Transmissions Property

It is worth noting that optimal transmissions grow with the coalitions in the sense made precise by Theorem 2 below.

Theorem 2 *Let $S \subset T \subset N$ and let $\tau_1 \in \mathcal{T}(S)$ be an optimal transmission for S . Then there exists an optimal transmission $\tau \in \mathcal{T}(T)$ for T such that $\tau \geq \tau_1$.*

Proof: Let τ_2 be an optimal transmission of T . Denote $\tau' \equiv \tau_1 \wedge \tau_2$ and $\tau \equiv \tau_1 \vee \tau_2$. By Lemma 1 and the fact that $S \subset T$, we have $\tau' \in \mathcal{T}(S)$ and $\tau \in \mathcal{T}(T)$.

The optimality of τ_1 for S implies

$$B(S, \tau_1) - B(S, \tau') \geq C(\tau_1) - C(\tau')$$

By the submodularity of C we have

$$C(\tau_1) - C(\tau') \geq C(\tau) - C(\tau_2)$$

From Lemma 2 we also have

$$B(T, \tau) - B(T, \tau_2) \geq B(S, \tau_1) - B(S, \tau')$$

The above three inequalities imply

$$\begin{aligned} B(T, \tau) - B(T, \tau_2) &\geq C(\tau) - C(\tau_2) \\ \Rightarrow B(T, \tau) - C(\tau) &\geq B(T, \tau_2) - C(\tau_2) \end{aligned}$$

Since τ_2 is an optimal transmission for T , the above inequality shows that τ is also optimal for T . But $\tau \equiv \tau_1 \vee \tau_2 \geq \tau_1$, proving the theorem. \square

6 Remarks

Remark 1 (Embedding the examples) We briefly indicate how to fit our examples (from Section 2) into the abstract model.

For Section 2.1, take $\Gamma(\alpha) = \Gamma$ rooted at α , $\kappa(e) = \phi$ for all e , $\mathcal{L}(\delta) = \{0, 1\}$, $\mathcal{L}(\alpha) = \{0\}$ for all $\alpha \neq \delta$, $\mathcal{L} =$ the cross product of all these lattices, $B_\delta = 0$, $B_\alpha(0) = 0$ and $B_\alpha(1) = B(\alpha)$ for all $\alpha \neq \delta$. Finally the cost of an edge is its setup cost if there is a non-zero transmission on it and zero otherwise.

For Section 2.2, let $\mathcal{L}(\alpha) = [0, m]^V$, each of whose elements specifies the information sent by α to all the other vertices. The lattice operations \vee and \wedge are obtained by taking component-wise maximum and minimum. \mathcal{L} as usual is the cross product of all the $\mathcal{L}(\alpha)$. The cost functions are obvious. The rest of the construction is as before. (Notice that despite the fact that the components of the benefit and cost functions have no convexity or concavity assumptions on them, the benefit/cost functions are supermodular/submodular in our lattice framework. This follows from the fact that they are additive over their components and that super or sub-modularity is no constraint on a function of one variable.)

For the example in Section 2.3, take $\mathcal{L}(\alpha)$ to be the totally ordered set $\{0\} \cup \Sigma_\alpha$, and \mathcal{L} to be the cross product. We leave it to the reader to verify that the benefit function is supermodular.

Finally, for the example in Section 2.4, take $\mathcal{L}(\alpha)$ to be the lattice of all partitions of Q which are coarser than $\{P_\alpha, Q \setminus Q_\alpha\}$. The supermodularity of the benefit functions follows from that of $f_\alpha, \alpha \in V$.

Remark 2 (Acyclicity) Cycles in the transmissions network $\Gamma(\alpha)$ can cause our result to breakdown. Consider the network in Figure 1 in which players 1, 2, 3, 4, each have access to the whole graph, with costs as shown and with $\epsilon < 1$.




Figure 1: Cycles in the communications network

Further suppose that 1, 2, 3 each derive benefit $B > 2(1 + \epsilon)$ from being connected to 4. Then it is clear that

$$\begin{aligned} w(2, 4) &= B - 2 \\ w(2, 3, 4) &= 2B - 2(1 + \epsilon) \\ w(1, 2, 4) &= 2B - 2(1 + \epsilon) \\ w(1, 2, 3, 4) &= 3B - 3(1 + \epsilon) \end{aligned}$$

Hence the marginal contribution of player 3 to the coalition $\{2, 4\}$ is $w(2, 3, 4) - w(2, 4) = B - 2\epsilon$. This is greater his marginal contribution $w(1, 2, 3, 4) - w(1, 2, 4) = B - 1 - \epsilon$ to the coalition $\{1, 2, 4\}$, violating the “snowball” effect which is equivalent to the convexity of w .




Figure 2: Modeling multiple players at a vertex

Remark 3 (Multiple players at a vertex) Our model allows for many players to be located at the same vertex α . Indeed, by creating a new vertex for each player present at α , and joining these with zero-cost edges to α , we create an expanded graph which fits our model (see Figure 2).

Remark 4 (Control of vertices) Our model also permits coalitions to control vertices by the graph expansion shown in Figure 3. Every edge incident at α is intercepted with a zero-cost edge controlled by the coalition controlling α .




Figure 3: Modeling control of vertices

Remark 5 (Veto players) A more general control of edges by veto players renders our results invalid. Consider a player set $\{1, 2, 3\}$ and suppose that there is common tree available to everyone, which consists of just one zero-cost edge connecting player 1 to a public vertex. The edge can be sanctioned by player 1 (the veto player), in conjunction with any player in $\{2, 3\}$. The only benefit B is obtained by player 1 when he gets connected to the public vertex. In this game $w(1) = 0$ and $w(1, 2) = w(1, 3) = w(1, 2, 3) = B$. Hence $w(1, 2) - w(1) = B > 0 = w(1, 2, 3) - w(1, 3)$, showing that w is not convex.

Remark 6 (Dropping distributivity) In the special case where \mathcal{L} is the cross product of the lattices $\mathcal{L}(\alpha)$ over $\alpha \in V$, our results hold without postulating that \wedge distributes over \vee . In this case, the analogue of (7) for \wedge holds trivially (and this was the only step that required distributivity).

But in general distributivity is indispensable.

Remark 7 (Enhancement of information) So far we have taken information to be fixed apriori. But it could well happen that the information of an agent gets enhanced by virtue of the information he receives from others. He can turn around and send his enhanced information back to them, enhancing theirs', and so on. Even in this setting, under suitable hypotheses, the induced cooperative game is well-defined (i.e., the enhancement sequence converges) and is convex, as we shall show in a sequel paper.

References

- [1] Garrett Birkhoff. Lattice theory. *Encyclopaedia Britannica, Macropaedia, Volume 1, Encyclopaedia Britannica Inc., Helen Hemingway Benton, London*, pages 519–524, 1977.
- [2] Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. In *ACM Symposium on Theory of Computing (STOC'00)*, Portland, Oregon, USA, May 2004.
- [3] Lloyd Shapley. A value for n-person games. *Contributions to the Theory of Games, edited by H. W. Kuhn and A. W. Tucker. Annals of Mathematics Studies*, 2(28), 1953.
- [4] Lloyd Shapley. Cores of convex games. *International Journal of Game Theory*, 1:11–26, 1971.
- [5] Y. Wang and Q. Zhu. Error control and concealment for video communications: A review. *Proceedings of IEEE, special issue on Multimedia Signal Processing*, pages 974–997, May 1998.