

Structural Routability of n -Pairs Information Networks^{*}

Girish N. Nair

Abstract Information does not generally behave like a conservative fluid flow in communication networks with multiple sources and sinks. However, it is often conceptually and practically useful to be able to associate separate data streams with each source-sink pair, with only routing and no coding performed at the network nodes. This raises the question of whether there is a nontrivial class of network topologies for which achievability is always equivalent to ‘routability’, for any combination of source signals and positive channel capacities. This chapter considers possibly cyclic, directed, errorless networks with n source-sink pairs and mutually independent source signals. The concept of *downward dominance* is introduced and it is shown that, if the network topology is downward dominated, then the achievability of a given combination of source signals and channel capacities implies the existence of a feasible multicommodity flow.

1 Introduction

In an n -pairs or *multiple unicast* communication network, n source signals must be conveyed to their corresponding sinks without exceeding any channel capacities. Until quite recently, the belief was that this was possible iff there existed a *routing* solution, i.e. if every symbol generated by a source could be carried without modification, over channels and through network nodes, until it reached the sink.

Girish N. Nair
Dept. Electrical and Electronic Engineering, University of Melbourne, VIC 3010, Australia, e-mail: gnair@unimelb.edu.au

^{*} Published in *Information and Control in Networks* (eds. G. Como, B. Bernhardsson, A. Rantzer), Lecture Notes in Control and Information Sciences v. 40, Springer, pp. 215–239, 2014. This work was supported by Australian Research Council grant DP110102401. A preliminary version was presented in [14].

At a macroscopic level, this is equivalent to presuming the existence of a feasible *multicommodity flow* [12].

However, in [15, 2], an example was constructed of a 2-pairs communication network that did not admit a routing solution, but became admissible if nodes could perform modulo-2 arithmetic on incoming bits. This counter-intuitive result started the field of *network coding*, in which nodes are permitted to not just route incoming symbols, but also to perform causal functions on them, so as to better exploit the network structure and the available channel capacities.

It is now known that the capacity regions for n -pairs networks are not generally given by feasible multicommodity flows. In [1], n -pairs networks were constructed with coding capacity much larger than the routing capacity. Other related work includes [9], in which a necessary and sufficient condition for broadcasting correlated sources over erroneous channels was found, and [13], in which linear network coding was shown to achieve capacity for a multicast network.

Notwithstanding the power of network codes, routing/multicommodity flow solutions are appealing in several respects. Most obviously they are simpler, because network nodes are not required to perform extra mathematical operations on arriving bits. In addition, because different data streams are not ‘hashed’ together by means of some function, there is arguably less potential for cross-talk between different source-sink pairs, arising for instance from nonidealities during implementation in the physical layer. For similar reasons, routing may be preferred over network coding if security and privacy are important. Furthermore, being able to treat information as a conservative fluid flow could potentially provide a simple basis to analyse communication requirements in areas outside traditional multiterminal information theory, e.g. networked feedback control and multi-agent coordination/consensus problems - see, e.g. [3].

These considerations raise the natural questions of whether there is a general class of network topologies on which achievability is always equivalent to the existence of a feasible multicommodity flow. This chapter aims to answer this questions for possibly cyclic, directed, errorless networks with n source-sink pairs and mutually independent source signals, where the goal is to reconstruct source-signals perfectly at their respective sinks. The structural concept of *downward dominance* (Def. 4.4) is introduced, and the main result (Thm. 4.1) is that if the network topology is downward dominated then the existence of an achievable combination of source signals and channel capacities always implies the existence of a feasible multicommodity flow.

The proof relies on the iterative construction of an *entropically feasible* multicommodity flow (Def. 5.2). As downward dominance inheres solely in the topology of the network, this result suits situations where channels, switches, transceivers and interfaces are expensive to set up and difficult to move, or where channel capacities and source-signal statistics are unknown. On these structures, information can always be treated like a flow of conservative, immiscible fluids.

Downward dominance is a more general condition than the notion of ‘triangularisability’ that was introduced in the conference version [14] of this chapter. While it is not generally easy to verify in arbitrary n -pairs networks, Lemmas 4.1 and 4.2)

give simpler, sufficient conditions for it to hold. Several examples are then provided in Sect. 6 to illustrate the applicability of Theorem 4.1 to various example networks, both cyclic and acyclic, including but not confined to the directed cycles and lines studied in [11, 10].

Although downward dominance is sufficient to guarantee that routing can always achieve the full coding capacity of a network, it is not necessary, and the important question of finding a more general - or even tight - structural condition remains open. In the concluding section, a potential directions for future work are outlined.

1.1 Notation and Basic Terminology

For convenience, the basic notation and terminology used in this chapter are described below.

- The set of nonnegative integers (i.e. whole numbers) is denoted by \mathbb{W} , the set of positive integers (i.e. natural numbers) by \mathbb{N} , and the set of positive reals by $\mathbb{R}_{>0}$.
- A contiguous set $\{i, i+1, \dots, j\}$ of integers is denoted $[i : j]$.
- Other sets are usually written in boldface type.
- Random variables (rv's) are written in upper case and their realisations are indicated in corresponding lower case.
- The set operation $\mathbf{A} \setminus \mathbf{B}$ denotes $\mathbf{A} \cap \mathbf{B}^c$.
- A discrete-time random signal or process $(F(k))_{k=0}^{\infty}$ is denoted F , and the finite sequence $(F(k))_{k=s}^t$ is denoted $F(s : t)$.
- Given a subscripted rv or signal F_j , with j belonging to a countable set \mathbf{J} , $F_{\mathbf{J}}$ denotes the tuple $(F_j)_{j \in \mathbf{J}}$, arranged according to the order on \mathbf{J} .
- The *entropy* of a discrete-valued rv E is denoted $\mathbf{H}[E] \geq 0$, and the conditional entropy of E given another rv F is $\mathbf{H}[E|F] := \mathbf{H}[E, F] - \mathbf{H}[F]$.
- The *mutual information* between rv's E and F is denoted $\mathbf{I}[E; F] := \mathbf{H}[E] - \mathbf{H}[E|F] \geq 0$, and the *conditional mutual information* between rv's E and F given G is denoted $\mathbf{I}[E; F|G] := \mathbf{H}[E|G] - \mathbf{H}[E|F, G]$.
- If E and F are random processes and E is discrete-valued, then the *entropy rates* of E , and the *conditional entropy rate* of E given (past and present) F are respectively defined as

$$\begin{aligned} \mathbf{H}_{\infty}[E] &:= \underline{\lim}_{t \rightarrow \infty} \frac{\mathbf{H}[E(0 : t)]}{t + 1}, \\ \mathbf{H}^{\infty}[E] &:= \overline{\lim}_{t \rightarrow \infty} \frac{\mathbf{H}[E(0 : t)]}{t + 1}, \\ \mathbf{H}_{\infty}[E|F] &:= \underline{\lim}_{t \rightarrow \infty} \frac{\mathbf{H}[E(0 : t)|F(0 : t)]}{t + 1}, \end{aligned}$$

- If E, F and G are random processes, then the *mutual information rates* of E and F , and the *conditional mutual information rate* of E and F given (past and present) G are respectively defined as

$$\begin{aligned}
I^\infty[E;F] &:= \overline{\lim}_{t \rightarrow \infty} \frac{I[E(0:t);F(0:t)]}{t+1}, \\
I_\infty[E;F] &:= \underline{\lim}_{t \rightarrow \infty} \frac{I[E(0:t);F(0:t)]}{t+1}, \\
I_\infty[E;F|G] &:= \underline{\lim}_{t \rightarrow \infty} \frac{I[E(0:t);F(0:t)|G(0:t)]}{t+1}.
\end{aligned}$$

- A *directed graph (digraph)* (\mathbf{V}, \mathbf{A}) consists of a set \mathbf{V} of *vertices*, and a set \mathbf{A} of *arcs* that each represent a directed link between a particular pair of vertices.
- The initial vertex of an arc is called its *tail* and the terminal vertex, its *head*.
- A *walk* in a digraph is an alternating sequence $\omega = (v_1, \alpha_1, v_2, \alpha_2, \dots, \alpha_k, v_{k+1})$, $k \geq 0$, of vertices and arcs, beginning and ending in vertices, s.t. each arc α_l connects the vertex v_l to v_{l+1} . Each vertex v_j and arc α_l in the sequence is said to *be in* the walk; with a minor abuse of notation, this is denoted $v_j \in \omega$.
- A *path* is a walk with no loops, i.e. it passes through no vertex more than once, including the initial one.
- An *undirected path* is an alternating sequence $\omega = (v_1, \alpha_1, v_2, \alpha_2, \dots, \alpha_k, v_{k+1})$, $k \geq 0$, of vertices and arcs, beginning and ending in vertices, s.t. no vertex is repeated and each arc α_l connects the vertex v_l to v_{l+1} , or v_{l+1} to v_l .
- A *cycle* is a walk in which the initial and final vertices are identical, but every other vertex occurs once.
- A *subpath* of a path $(v_1, \alpha_1, v_2, \alpha_2, \dots, \alpha_k, v_{k+1})$ is a segment $(v_l, \alpha_l, v_{l+1}, \dots, v_j)$ of it, where $1 \leq l \leq j \leq k+1$.
- A vertex v is said to be *reachable* from another vertex μ , denoted $\mu \rightsquigarrow v$, if \exists a path leading from μ to v . Equivalently, it is said that μ *can reach* v . The same terminology and notation apply, with analogous meaning, for pairs of arcs as well as mixed pairs of arcs and vertices. E.g. given an arc β , $\mu \rightsquigarrow \beta$ means that there is a path from the vertex μ to the tail of β .
- Similarly, a (vertex or arc) set \mathbf{W} is said to be *reachable* from another set \mathbf{U} , denoted $\mathbf{U} \rightsquigarrow \mathbf{W}$, if there is an element of \mathbf{W} that is reachable from an element of \mathbf{U} ; equivalently, it is said that \mathbf{U} *can reach* \mathbf{W} .
- For any vertex set $\mathbf{U} \subseteq \mathbf{V}$, $\text{ARCS}(\mathbf{U}) \subseteq \mathbf{A}$ is the set of arcs with tails in \mathbf{U} .
- The notation $\text{OUT}(\mathbf{U})$ ($\text{IN}(\mathbf{U})$) represents the set of arcs in \mathbf{A} that have tails (resp. heads) in a vertex set $\mathbf{U} \subseteq \mathbf{V}$ and heads (tails) $\in \mathbf{V} \setminus \mathbf{U}$. If $\text{OUT}(\mathbf{U})$ ($\text{IN}(\mathbf{U})$) consists of a single arc, this arc is denoted $\text{out}(\mathbf{U})$ ($\text{in}(\mathbf{U})$). When \mathbf{U} is a singleton $\{\mu\}$, the braces are omitted.

2 Problem Formulation

A network of unidirectional, point-to-point channels may be modelled using a digraph (\mathbf{V}, \mathbf{A}) , where the vertex set \mathbf{V} represents information sources, sinks, repeaters, routers etc., and the arc set \mathbf{A} indicates the directions of any channels between nodes. As usual with digraphs, it is assumed that no arc leaves and enters the

same vertex, and that at most one arc leads from the first to the second element of any given ordered pair of vertices. In other words, every arc in \mathbf{A} may be uniquely identified with a tuple $(\mu, \nu) \in \mathbf{V}^2$, with $\mu \neq \nu$.² It is also assumed that the digraph is *connected*, i.e. there is an undirected path between any distinct pair of vertices.

In an n -pairs information network, the locations of sources and sinks are respectively represented by disjoint sets $\mathbf{S} = \{\sigma_1, \dots, \sigma_n\}$ and $\mathbf{T} = \{\tau_1, \dots, \tau_n\}$ of distinct vertices in \mathbf{V} , with each source σ_i aiming to communicate to exactly one sink τ_i . It is assumed that $\sigma_i \rightsquigarrow \tau_i$. Let \mathbf{P} denote the sequence $((\sigma_i, \tau_i))_{i=1}^n$ of source-sink pairs, arranged in a specified order. Without loss of generality, it is assumed that every source (sink) has no in-coming (resp. out-going) arcs and exactly one out-going (in-coming) arc.³ The *boundary* $\partial\mathbf{V}$ of the network is the set $\mathbf{S} \cup \mathbf{T}$ of source and sink vertices, and its *interior* is $\text{int}\mathbf{V} := \mathbf{V} \setminus \partial\mathbf{V}$.

Each channel in the network can transfer bits errorlessly up to a maximum average rate, as specified by a positive *arc-capacity* $c_\alpha \in \mathbb{R}_{>0}$. In some situations, it may be natural to assign infinite capacity to certain arcs,⁴ and the set of all such arcs is denoted $\mathbf{A}_\infty \subset \mathbf{A}$. In particular, the arcs leaving sources are by convention assigned infinite capacity. The set of finite-capacity arcs is written $\mathbf{A}_f = \mathbf{A} \setminus \mathbf{A}_\infty$, with associated arc-capacity vector $c := (c_\alpha)_{\alpha \in \mathbf{A}_f} \in \mathbb{R}_{>0}^{|\mathbf{A}_f|}$. The *structure* of the n -pairs information network is defined as the tuple $\Sigma = (\mathbf{V}, \mathbf{A}_f, \mathbf{A}_\infty, \mathbf{P})$.

The communication signals in the network are represented by a vector $S \equiv (S_\alpha)_{\alpha \in \mathbf{A}}$ of discrete-valued random processes called *arc signals*. In particular, the arc signals leaving sources and entering sinks respectively represent the exogeneous inputs to and outputs from the network. For convenience, the input signal $S_{\text{out}(\sigma_i)}$ generated by the i -th source $\sigma_i \in \mathbf{S}$ is called X_i , and the output signal $S_{\text{in}(\tau_i)}$ entering the i -th sink $\tau_i \in \mathbf{T}$ is called Y_i . It is assumed throughout this chapter that the signals X_1, \dots, X_n are mutually independent processes with strictly positive entropy rates $H_\infty[X_i] > 0$.

The arc-signal vector S is assumed to have the following property:

Definition 2.1 (Setwise Causality and Signal Graphs). An arc-signal vector S is called *setwise causal* on a structure $\Sigma = (\mathbf{V}, \mathbf{A}_f, \mathbf{A}_\infty, \mathbf{P})$ if all arc signals leaving vertices in any internal vertex-set $\mathbf{U} \subseteq \text{int}\mathbf{V}$ are causally determined by those entering \mathbf{U} from outside it. That is, $\forall \mathbf{U} \subseteq \text{int}\mathbf{V}, \exists$ an operator $g_{\mathbf{U}}$ s.t.

$$S_{\text{ARCS}(\mathbf{U})}(t) = g_{\mathbf{U}}(t, S_{\text{IN}(\mathbf{U})}(0:t)), \quad \forall t \in \mathbb{W}, \quad (1)$$

where $\text{ARCS}(\mathbf{U}) \subseteq \mathbf{A}$ denotes the set of arcs leaving vertices of \mathbf{U} .

The tuple (Σ, S) is then called a *signal graph*.

◇

² Such digraphs are sometimes called *simple*.

³ If a source or sink were actually connected to multiple nodes in the network, it would be represented in the digraph by an auxiliary vertex connected by an arc (of infinite capacity) with a multiply-connected vertex.

⁴ For instance, when a single network node is represented as two ‘virtual’ vertices connected by an arc of unbounded capacity.

Remark: Setwise causality is a strengthened version of the basic concept of *well-posedness* [16] in feedback control theory. In a well-posed feedback system, the current values of all internal and output signals are uniquely determined by the past and present values of external inputs.⁵ Setwise causality essentially imposes an analogous condition on any subcollection of nodes and associated signals, treated as a system. In acyclic digraphs (i.e. in which every walk is a path), it is equivalent to causality at every internal vertex. However, feedback signals may be present in cyclic digraphs, in which case vertex-wise causality cannot guarantee (1) without further assumptions, e.g. a positive time-delay at every vertex.

In the n -pairs network problem studied here, the objective is for each sink to perfectly reconstruct each source signal, block-by-block, using only causal operations and without exceeding any arc-capacities. This leads to the following definition:

Definition 2.2 (Achievability). Consider an n -pairs information network with structure Σ , source-signal vector X and arc-capacity vector $c \in \mathbb{R}_{>0}^{|\mathbf{A}_f|}$. The tuple (Σ, X, c) is called *achievable* if \exists a setwise-causal arc-signal vector S (Def. 2.1) and a positive integer $m \in \mathbb{N}$ s.t.

$$S_{\text{out}(\sigma_i)} = X_i, \quad \forall i \in [1 : n], \quad (2)$$

$$Y_i(km - 1) = X_i((k - 1)m : km - 1), \quad \forall k \in \mathbb{N}, i \in [1 : n], \quad (3)$$

$$H^\infty[S_\alpha] \leq c_\alpha, \quad \forall \alpha \in \mathbf{A}_f. \quad (4)$$

Such an S is called a *solution* to the n -pairs information network problem (Σ, X, c) . The arc-capacity vector c is called achievable on (Σ, X) and (X, c) is called achievable on Σ .

◇

Remarks: This differs from standard definitions of network coding solutions in several minor respects. For instance, in [2, 7, 11, 5] and most of [10], the inequalities (4) are replaced by bounds either on the cardinalities of channel alphabets, or on block-coding rates over a period of time. In addition, in previous formulations, the sinks typically must reconstruct the source signal either perfectly and instantaneously [7, 5, 10], which corresponds to setting $m = 1$ in (3), or else with arbitrarily small probability of decoding error over blocks of sufficiently large length m [2, 11].

In this work, bounds are imposed directly on entropies, as in sec. VIII of [10], in order to focus on the information-theoretic aspects of the problem. Errorless reconstruction is demanded so as to enable the graphical characterisation of *informational dominance* from [10] to be used with very minor changes. However, perfect reconstruction is not required instantaneously in (3), but only in blocks of length m . This allows a solution S to be interpreted operationally in terms of variable bit-rate codes.⁶

⁵ In the linear, time-invariant context of [16], this is equivalent to the corresponding transfer functions being well-defined and proper.

⁶ In other words, if S solves (Σ, X, c) , then there exist variable bit-rate codes for each arc that yield errorless, block-by-block reconstruction of the source-signals at their sinks, with expected bit-

Finally, it is conjectured that the results in this paper also apply if (3) is relaxed so that Y_i is causally determined by X_i , with $H_\infty[Y_i] > 0$.

As mentioned in the introduction, it was once thought that a network was achievable⁷ iff it admitted a routing solution. In the present context, this is equivalent to presuming the existence of an (X, c) -feasible multicommodity flow, i.e. of a non-negative tuple $f = (f_{\alpha,j})_{\alpha \in \mathbf{A}, j \in [1:n]} \in \mathbb{R}_{\geq 0}^{|\mathbf{A}|n}$, of bit-rates on each arc associated with every source-sink pair, s.t.

$$\sum_{j=1}^n f_{\alpha,j} \leq c_\alpha, \quad \forall \alpha \in \mathbf{A}_f \quad (\text{capacity bound}); \quad (5)$$

$$f_{\text{in}(\tau_j),j} = f_{\text{out}(\sigma_j),j} = H_\infty[X_j], \quad \forall j \in [1:n] \quad (\text{supply equals demand}), \quad (6)$$

$$\sum_{\alpha \in \text{IN}(v)} f_{\alpha,j} = \sum_{\alpha \in \text{OUT}(v)} f_{\alpha,j} \quad (\text{conservation of flow}), \quad (7)$$

for any $j \in [1:n]$ and $v \in \mathbf{V} \setminus (\{\sigma_j\} \cup \{\tau_j\})$. Via an explicit counter-example, the article [2] showed that this intuitive notion was incorrect, i.e. that although the existence of a feasible multicommodity flow is sufficient for achievability, it is not generally necessary. This laid the foundations for *network coding*, in which nodes are permitted to not just route incoming bits, but also to perform functions on them.

Nonetheless, routing/multicommodity-flow solutions have certain virtues, as discussed in Sect. 1. This chapter poses the question: is there a general class of n -pairs information network structures Σ in which the achievability of (X, c) is equivalent to the existence of an (X, c) -feasible multicommodity flow f (5)–(7)?

Any n -pairs information network structure Σ can support (X, c) -feasible multicommodity flows if the arc-capacities are sufficiently larger than the source entropy rates, provided each sink is reachable from its source. However, there are examples of structures on which an (X, c) -feasible multicommodity flow does not exist if arc-capacities are reduced, even though (X, c) is still achievable (see Sect. 6).

The aim of this chapter is to isolate certain structural properties that ensure routability *over all achievable combinations of (X, c)* . Such properties would inhere solely in Σ , suiting situations in which channels, switches, transceivers and interfaces are expensive to set up and difficult to move, and/or where channel capacities and source-signal statistics are variable or unknown.

rates at worst negligibly larger than arc-capacities. Conversely, if there exists a distributed entropy coding scheme that achieves perfect reconstruction of source-signals at their sinks in blocks of length m , and with expected bit-rates no larger than the arc-capacities, then this yields a solution S as defined above. However, these operational interpretations will not be used in this article.

⁷ ignoring differences in the definition of achievability

3 Preliminary Notions

Before proceeding, several existing graph-theoretic notions are needed. Throughout this section, $\Sigma = (\mathbf{V}, \mathbf{A}, \mathbf{P}) \equiv (\mathbf{V}, \mathbf{A}_f, \mathbf{A}_\infty, \mathbf{P})$ is the structure of an n -pairs information network as described in Sect. 2, and $\Gamma = (\Sigma, S)$ is its setwise-causal signal graph (Def. 2.1), with source- and sink-signal vectors X and Y .

First, some largely familiar concepts are revisited. A path in an n -pairs information network that goes from a source σ_i to its sink τ_i is called an i -path. The set of all i -paths is called an i -bundle, i.e. the set of all acyclic walks via which information can be routed from σ_i to τ_i . Given a set $\mathbf{J} \subseteq [1 : n]$, the set of all i -paths with $i \in \mathbf{J}$ is called a \mathbf{J} -bundle (not the same as the set of $\sigma_{\mathbf{J}} \rightsquigarrow \tau_{\mathbf{J}}$ -paths, which contains it). Let $(\mathbf{V}^{\mathbf{J}}, \mathbf{A}^{\mathbf{J}})$ denote the subgraph formed by all the vertices and arcs in the \mathbf{J} -bundle. In particular, $(\mathbf{V}^i, \mathbf{A}^i)$ is the subgraph formed by the i -bundle. A vertex set $\mathbf{U} \subset \mathbf{V}^i$ such that $\sigma_i \in \mathbf{U}$ and $\tau_i \notin \mathbf{U}$ is called an i -cut.

The following concepts are adapted from [10], with minor changes in terminology.

Definition 3.1 (Indirect i -Walks – Based on [10]). An indirect i -walk (ii -walk) ω is an alternating sequence $(\alpha_1, \beta_1, \dots, \alpha_{j-1}, \beta_{j-1}, \alpha_j)$ of forward- and reverse-oriented paths in the n -pairs structure Σ such that

1. α_1 begins with the i -th source vertex σ_i ;
2. both α_ℓ and β_ℓ end with the same vertex $\mu_\ell, \forall \ell \in [1 : j - 1]$;
3. both β_ℓ and $\alpha_{\ell+1}$ begin from the same source vertex, $\forall \ell \in [1 : j - 1]$;
4. α_j ends with the sink vertex τ_i ; and
5. every arc and vertex in ω can reach τ_i .

An ii -walk ω is said to *bypass* an arc-set \mathbf{C} if no arc in ω lies in \mathbf{C} .

◇

Remarks: Note that the fifth condition above is equivalent to the requirement that each joint vertex μ_ℓ reaches $\tau_i, \forall \ell \in [1 : j - 1]$.

An ii -walk as defined above is, in the terminology of [10], an *indirect walk* from $\text{out}(\sigma_i)$ to $\text{in}(\tau_i)$ in a subgraph $G(\emptyset, i)$. Similarly, an ii -walk that bypasses \mathbf{C} is an indirect walk from $\text{out}(\sigma_i)$ to $\text{in}(\tau_i)$ in a subgraph $G(\mathbf{C}, i)$; if such a bypass exists, then Y_i is not always fully determined by $S_{\mathbf{C}}$, even if all i -paths go through \mathbf{C} . See Fig. 3 and Defs. 10 – 11 in [10].

Indirect i -walks are related to the concept of *fd-separation* [11]. In particular, if $S_{\mathbf{C}}$ fd-separates X_i and Y_i for any setwise causal S (Def. 2.1), then all ii -walk's pass through \mathbf{C} ; that is, an ii -walk that bypasses \mathbf{C} corresponds to an undirected path between X_i and Y_i in a *functional dependence* subgraph $\mathcal{G}_{X_i, S_{\mathbf{C}}, Y_i}$ constructed according to the procedure in [11].

However, the converse is not generally true, i.e. ‘ ii -separation’ is a less stringent requirement. This is because paths connecting X_i and Y_i in $\mathcal{G}_{X_i, S_{\mathbf{C}}, Y_i}$ do not have to satisfy an analogue of the fifth condition, which arises from the requirement that

each sink reproduce its source signal with perfect fidelity. For this to be possible, it turns out that each joint vertex μ_ℓ in an ii -walk must be able to reach τ_i .

Put another way, requiring S_C to fd-separate X_i and Y_i is equivalent to requiring that $a)$ C be an i -cut, and $b)$ for each $j \neq i$, either all $\sigma_j \rightsquigarrow \tau_i$ -paths (if any) bypass C , or all pass through it. Under ii -separation, (a) must still hold, but (b) is relaxed: a source σ_j can have a path π to τ_i that bypasses C as well as another that passes through C , provided that π is not the last leg of an ii -walk that bypasses C .

Definition 3.2 (Structural Dominance – Based on [10]). For any arc-set $B \subseteq A$ in an n -pairs network, $SDOM(B)$ is the smallest arc-set $C \subseteq A$ that satisfies the conditions below:

1. $C \supseteq B$
2. $out(\sigma_i) \in C$ iff $in(\tau_i) \in C$
3. If $\alpha \in A$ is *downstream* from C – i.e. all paths from sources to the tail of α pass through C – then $\alpha \in C$.
4. If all indirect i -walks (Def.3.1) pass through C then $out(\sigma_i), in(\tau_i) \in C$.

The arcs in $SDOM(B)$ are said to be *structurally dominated* by B .

◇

Remarks: Note that $SDOM(B)$ is the smallest such arc-set in the sense of being contained by every $C \subseteq A$ that satisfies criteria 1 – 4.

As noted in [8] (pp. 199–200), $SDOM(B)$ can be constructed by setting $C = B$, letting $T = A \setminus B \neq \emptyset$ be the set of arcs to be tested, and then following this greedy algorithm:

- (i) Pick any arc $\alpha \in T$.
- (ii) If α satisfies any of the conditions 2 – 4 in Def. 3.2, update $C \leftarrow C \cup \{\alpha\}$ and then $T \leftarrow A \setminus C$; else keep C the same and update $T \leftarrow T \setminus \{\alpha\}$.
- (iii) If $T = \emptyset$ then exit; else go to step (i).

The final set C is then $SDOM(B)$. However, the following lemma gives two quicker conditions for guaranteeing that a specific arc lies in $SDOM(B)$.

Lemma 3.1 (Based on [10]).

1. If an arc $\alpha \in A$ is *downstream* from B , then $\alpha \in SDOM(B)$.
2. If all indirect i walks (Def. 3.1) pass through B then $out(\sigma_i), in(\tau_i) \in SDOM(B)$.

Proof. If either of these criteria hold, then the relevant arcs – α , $out(\sigma_i)$, $in(\tau_i)$ – must lie inside any arc-set C that satisfies the 1st to 4th conditions in Def. 3.2. As $SDOM(B)$ is such a set, the lemma follows. □

The significance of structural dominance arises from the following result:

Theorem 3.1 (Informational Dominance – Based on [10]). Consider any arc $\alpha \in A$ and arc-set $B \subseteq A$ in an n -pairs network with structure Σ . If $\alpha \in SDOM(B)$ (Def.

3.2), then for any setwise causal arc-signal vector S (Def. 2.1) and positive integer $m \in \mathbb{N}$ that satisfy (2) and (3), \exists a function γ such that

$$S_\alpha(0 : km) = \gamma(k, S_{\mathbf{B}}(0 : km)), \quad \forall k \in \mathbb{N}. \quad (8)$$

Conversely, if for any setwise causal S that meets (2) – (3) with block-length m there is a function γ ensuring that (8) holds, then $\alpha \in \text{SDOM}(\mathbf{B})$.

Remarks: The property specified in (8) is a version of the concept of *informational dominance* introduced in []; the important of this result lies in giving this functional concept a purely structural characterisation. The proof follows similar lines as that of Theorem 10 in [10] and is omitted. Minor differences are that m is not constrained to be 1 here, and that cyclic networks are handled using the notion of setwise causality (Def. 2.1), rather than by introducing channel delays and then ‘unwrapping’ the network over time to yield an infinite directed acyclic graph.

4 Main Result

The main result of this paper is presented in this section. In order to do so, several nonstandard graph-theoretic notions are needed. Throughout this section, $\Sigma = (\mathbf{V}, \mathbf{A}, \mathbf{P}) \equiv (\mathbf{V}, \mathbf{A}_f, \mathbf{A}_s, \mathbf{P})$ is the structure of an n -pairs information network as described in Sect. 2, and $\Gamma = (\Sigma, S)$ is its setwise-causal signal graph (Def. 2.1), with source- and sink-signal vectors X and Y .

Definition 4.1 (J-Disjointness). Given an index set $\mathbf{J} \subseteq [1 : n]$, an arc set $\mathbf{B} \subseteq \mathbf{A}$ is **J-disjoint** if each path in the **J**-bundle passes through at most one arc in **B**.

If $\mathbf{J} = \{i\}$ for some $i \in [1 : n]$, then **B** is called *i-disjoint*.

◇

Remarks: It is easy to see that empty and singleton arc-sets are automatically **J**-disjoint, that every $\mathbf{B} \subseteq \mathbf{A}$ is \emptyset -disjoint, and that every subset of a **J**-disjoint set inherits its **J**-disjointness. With a little effort, it can also be shown that **J**-disjoint arc-sets satisfy the ‘augmentation’ property. Thus **J**-disjoint sets form a *finite matroid* on **A**.

Structural dominance (Def. 3.2) and $[1 : i]$ -disjointness are next used to define nested families of arc-sets with certain structural properties. These properties are needed later to inductively extract *entropically feasible* multicommodity flows (Def. 5.2). First, for any arc-set $\mathbf{E} \subseteq \mathbf{A}$ and $h \in [1 : n]$ define the source-augmented set

$$\mathbf{E}_{*h} := \mathbf{E} \cup \text{OUT} \left(\sigma_{\mathbf{J}^{h-1} \cup [h+1:n]} \right), \quad \mathbf{J}^{h-1} \equiv \{j \in [1 : h-1] : \mathbf{E} \cap \mathbf{A}^j = \emptyset\}. \quad (9)$$

That is, **E** is augmented by those source-arcs that either have indices greater than h or that have indices less than h but no source-sink paths going through **E**.

Definition 4.2 (i -Downward Dominated Sets). For each $i \in [1 : n]$, the family \mathcal{D}_i consists of all arc sets $\mathbf{E} \subseteq \mathbf{A}$ such that

1. \mathbf{E} is $[1 : i]$ -disjoint (Def. 4.1), and
2. for each $h \in [1 : i]$, either the h -bundle does not touch \mathbf{E} , i.e. $\mathbf{E} \cap \mathbf{A}^h = \emptyset$, or else the source-augmented arc-set \mathbf{E}_{*h} (9) structurally dominates the source-arc out(σ_h) (Def. 3.2).

Every member-set of \mathcal{D}_i is called *i -downward dominated*.

◇

Remark: Clearly, every \mathcal{D}_i -set is also in \mathcal{D}_{i-1} .

The next concept describes a class of i -cuts that have a special structure:

Definition 4.3 (Viable i -Cuts). Given an index $i \in [1 : n]$, an i -cut $\mathbf{U} \subset \mathbf{V}^i$ is called *viable* under the following conditions:

1. Every arc leaving \mathbf{U} in the i -bundle is finite-capacity, i.e. $\text{OUT}(\mathbf{U}) \cap \mathbf{A}^i \subseteq \mathbf{A}_f$.
2. There is an i -path that leaves \mathbf{U} without re-entering.
3. Each arc in $\text{OUT}(\mathbf{U}) \cap \mathbf{A}^i$ lies in an i -path that either exits \mathbf{U} without re-entering or else lies in the $[1 : i - 1]$ -bundle.
4. Every vertex $v \in \mathbf{U}$ lies on an undirected path π from σ_i to v such that
 - a. all vertices before v on π are in \mathbf{U} , and
 - b. every reverse-oriented arc in π (i.e. pointing from v to σ_i) lies on an i -path that does not re-enter \mathbf{U} .

◇

Remark: Viable i -cuts correspond to possible *min-cuts* in a residual capacitated digraph that is used to prove the main result of this chapter (Thm. 4.1). Further investigation of these min-cuts may yield other structural properties to add to the list above; however, this is left for future work.

Definition 4.4 (Downward Dominance). A structure Σ is called *downward dominated* if for each $i \in [2 : n]$ and viable i -cut \mathbf{U} (Def. 4.3), the set $\mathbf{O}^i = \text{OUT}(\mathbf{U}) \cap \mathbf{A}^i$ of outgoing arcs in the i -bundle satisfies the following two conditions:

1. $\mathbf{O}^i \in \mathcal{D}_{i-1}$ (Def. 4.2), and
2. the source-augmented arc-set \mathbf{O}_{*i}^i (9) structurally dominates the source-arc out(σ_i) (Def. 3.2).

◇

Remarks: Note that 1-pair structures are automatically downward dominated, since the conditions above become empty.

A sequence of simpler and progressively more restrictive sufficient conditions for downward dominance can be found by exploiting Lemma 3.1:

Lemma 4.1 (Simpler Condition 1). *Suppose that for each $i \in [2 : n]$ and viable i -cut $\mathbf{U} \subset \mathbf{V}^i$ (Def. 4.3), the arc-set $\mathbf{O}^i = \text{OUT}(\mathbf{U}) \cap \mathbf{A}^i \subseteq \mathbf{A}_f$ satisfies the following conditions:*

1. \mathbf{O}^i is $[1 : i - 1]$ -disjoint (Def. 4.1), and
2. for each index $h \in [1 : i]$ for which there is a h -path that passes through \mathbf{O}^i , i.e. $\mathbf{O}^i \cap \mathbf{A}^h \neq \emptyset$, all indirect h -walks (Def. 3.1) pass through the source-augmented arc-set \mathbf{O}_{*h}^i (9).

Then Σ is downward dominated (Def. 4.4).

Proof. Follows immediately from applying Lemma 3.1 to Defs. 4.2 and 3.2. \square

Lemma 4.2 (Simpler Condition 2). *Suppose that for each $i \in [2 : n]$ and viable i -cut $\mathbf{U} \subset \mathbf{V}^i$ (Def. 4.3), the arc-set $\mathbf{O}^i = \text{OUT}(\mathbf{U}) \cap \mathbf{A}^i \subseteq \mathbf{A}_f$ satisfies the following conditions:*

1. \mathbf{O}^i is $[1 : i - 1]$ -disjoint (Def. 4.1).
2. For every $h \in [1 : i]$ and $s \in [1 : h]$ such that $\mathbf{O}^i \cap \mathbf{A}^h$ and $\mathbf{O}^i \cap \mathbf{A}^s \neq \emptyset$, all paths from σ_s to τ_h pass through \mathbf{O}^i .

Then Σ is downward dominated (Def. 4.4).

Proof. Let $\mathbf{O}^i \cap \mathbf{A}^h \neq \emptyset$ for some $h \in [1 : i]$. It is asserted that all indirect h -walks (Def. 3.1) must pass through \mathbf{O}_{*h}^i .

To see this, suppose in contradiction that there is an indirect h -walk ω that does not pass through $\mathbf{O}_{*h}^i \equiv \mathbf{O}^i \cup \text{OUT}(\sigma_{\mathbf{J}^{h-1} \cup [h+1:n]})$, where $\mathbf{J}^{h-1} \equiv \{j \in [1 : h-1] : \mathbf{O}^i \cap \mathbf{A}^j = \emptyset\}$. Let σ_s be the last source vertex in ω , and let π be the subpath from σ_s to τ_h . Clearly, $s \notin \mathbf{J}^{h-1} \cup [h+1 : n]$. In addition, $s \neq h$, since otherwise ω reduces to a path from σ_h to τ_h , which by the second condition above must pass through $\mathbf{O}^i \subseteq \mathbf{O}_{*h}^i$.

Thus $s \in [1 : h-1] \setminus \mathbf{J}^{h-1}$, i.e. $\mathbf{O}^i \cap \mathbf{A}^s \neq \emptyset$. By the second condition above, all $\sigma_s \rightsquigarrow \tau_h$ -paths must then pass through \mathbf{O}^i . As π is such a path, the indirect h -walk ω , of which it is a part, passes through $\mathbf{O}^i \subseteq \mathbf{O}_{*h}^i$, yielding a contradiction.

The result then follows from Lemma 4.1. \square

Lemma 4.3 (Simpler Condition 3). *Suppose that for each $i \in [2 : n]$ and every viable i -cut $\mathbf{U} \subset \mathbf{V}^i$ (Def. 4.3), there is exactly arc in $\mathbf{O}^i = \text{OUT}(\mathbf{U}) \cap \mathbf{A}^i \subseteq \mathbf{A}_f$.*

Furthermore, suppose that for each $h \in [1 : i]$ and $s \in [1 : h]$ such that $\mathbf{O}^i \cap \mathbf{A}^h$ and $\mathbf{O}^i \cap \mathbf{A}^s \neq \emptyset$, all paths from σ_s to τ_h pass through \mathbf{O}_i , or none of them do.

Then Σ is downward dominated (Def. 4.4).

Proof. Observe that \mathbf{O}^i consists of a single arc α . Thus the first condition of Lemma 4.2 is trivially satisfied. To show that its second condition is also met, suppose that $\mathbf{O}^i \cap \mathbf{A}^h, \mathbf{O}^i \cap \mathbf{A}^s \neq \emptyset$ for some $h \in [1 : i], s \in [1 : h]$. Thus α is on both an s -path and a h -path. Let π_1 be the subpath of the s -path from σ_s to the tail of α and π_2 , the subpath of the h -path from the head of α to τ_h . Then the concatenation $\pi_1 \alpha \pi_2$ is a $\sigma_s \rightsquigarrow \tau_h$ -path that passes through \mathbf{O}^i . By the all-or-nothing condition above, all $\sigma_s \rightsquigarrow \tau_h$ -paths then pass through \mathbf{O}^i . The result then follows from Lemma 4.2. \square

The main result of this chapter can now be stated:

Theorem 4.1 (Downward Dominance \Rightarrow Structural Routability). *If there is an ordering of the source-sink pairs in an n -pairs network so that the structure Σ is downward dominated (Def. 4.4), then the achievability of (X, c) (Def. 2.2) implies the existence of an (X, c) -feasible multicommodity flow (5)–(7).*

Conversely, if X is stationary and there exists an (X, c) -feasible multicommodity flow with (5) holding in strict form, then (Σ, X, c) is achievable. ∇

Remarks: This result defines a non-trivial class of directed network structures for which achievability is essentially equivalent to the existence of a feasible multicommodity flow. On these structures, information can indeed be treated like an incompressible, immiscible fluid flow.

The proof of Thm. 4.1 is given in the next section. In Sect. 6, several network examples are discussed to illustrate the applicability of Thm. 4.1.

5 Proof of Theorem 4.1

In both the proofs of necessity and sufficiency, use will be made of the fact that $\forall i \in [1 : n]$, any single-commodity flow q from σ_i to τ_i in the structure Σ can be decomposed into a superposition of i -path flows and cycle flows (see e.g. [4], Thm. 3.3.1). That is, if $\pi_{1,i}, \dots, \pi_{p,i}$ are the distinct i -paths and $\gamma_1, \dots, \gamma_g$, the distinct cycles, then \exists numbers $u_{1,i}, \dots, u_{p,i} \geq 0$ and $w_{1,i}, \dots, w_{g,i} \geq 0$ s.t.

$$q\alpha = \sum_{1 \leq k \leq p_i: \pi_{k,i} \ni \alpha} u_{k,i} + \sum_{1 \leq l \leq g: \gamma_l \ni \alpha} w_{l,i}. \quad (10)$$

If $w_{l,i} = 0$ for all $l \in [1 : g]$, then the flow q is called *acyclic*.

The proof of sufficiency in Sect. 5.2 is relatively straightforward. Given an (X, c) -feasible multicommodity flow f (5)–(7) on Σ , the decomposition (10) is used directly to devise a routing solution S .

The proof of necessity in Sect. 5.1 is more difficult and involves induction, using the following building blocks.

Definition 5.1 (J-Flow). Given an index set $\mathbf{J} \subseteq [1 : n]$, a nonnegative tuple $f = (f_{\alpha,j})_{\alpha \in \mathbf{A}, j \in \mathbf{J}} \in \mathbb{R}_{\geq 0}^{|\mathbf{A}| |\mathbf{J}|}$ is called a **J-flow** on the structure Σ if $\forall j \in \mathbf{J}$ and $v \in \mathbf{V} \setminus \{\sigma_j\} \cup \{\tau_j\}$,

$$\sum_{\alpha \in \text{IN}(v)} f_{\alpha,j} = \sum_{\alpha \in \text{OUT}(v)} f_{\alpha,j} \quad (j\text{-flow conservation}), \quad (11)$$

As a convention, the \emptyset -flow is defined as the empty sequence $()$.

\diamond

Remark: A \mathbf{J} -flow is a (possibly infeasible) multicommodity flow with source-sink pairs (σ_j, τ_j) , $j \in \mathbf{J}$. If each j -flow $f_{\mathbf{A},j}$ is acyclic, $\forall j \in \mathbf{J}$, then f is called an *acyclic \mathbf{J} -flow*.

The next concept is central to the proof of necessity. It defines a class of feasible $[1 : i]$ -flows that obey certain information-theoretic bounds when only the signals X_j , $j \in [1 : i]$, need to be communicated.

Definition 5.2 (Entropic Feasibility). Given $i \in [1 : n]$ and a solution S to (Σ, X, c) (Def. 2.2), a $[1 : i]$ -flow $f \in \mathbb{R}_{\geq 0}^{|\mathbf{A}|^i}$ (Def. 5.1) is called *entropically feasible* if it satisfies the following conditions:

i) On every arc $\alpha \in \mathbf{A}_f$,

$$\sum_{j=1}^i f_{j,\alpha} \leq c_\alpha. \quad (12)$$

ii) On any i -downward dominated arc set \mathbf{B} (Def. 4.2),

$$\sum_{\alpha \in \mathbf{B}, j \in [1:i]} f_{\alpha,j} \leq H^\infty \left[S_{\mathbf{B}} | X_{\mathbf{J}^i \cup [i+1:n]} \right], \quad (13)$$

where $\mathbf{J}^i = \{j \in [1 : i] : \mathbf{B} \cap \mathbf{A}^j = \emptyset\}$.

iii) On arcs entering sinks and leaving sources,

$$f_{\text{in}(\tau_j),j} = f_{\text{out}(\sigma_j),j} = H_\infty[X_j], \quad \forall j \in [1 : i]. \quad (14)$$

◇

Remarks: Note that the \emptyset -flow is entropically feasible, since the condition (14) disappears and (13) is trivially satisfied due to a zero left-hand side (LHS).

The proof of necessity in the next section proceeds by constructing an entropically feasible $[1 : n]$ -flow on (Σ, c, S) , which automatically gives the desired (X, c) -feasible multicommodity flow (5)–(7).

5.1 Necessity Proof for Theorem 4.1

Let the arc-signal vector S be a solution (Def. 2.2) to the n -pairs information network problem (Σ, X, c) . An entropically feasible $[1 : n]$ -flow (Def. 5.2) f^n will be constructed, using upward induction.

Let Σ be downward dominated (Def. 4.4) and suppose that $f^{i-1} = (f_{\alpha,j})_{\alpha \in \mathbf{A}, j \in [1:i-1]} \in \mathbb{R}_{\geq 0}^{|\mathbf{A}|^{(i-1)}}$ is an entropically feasible, acyclic $[1 : i-1]$ -flow for some $i \in [1 : n]$, noting that the \emptyset -flow f^0 is entropically feasible. An i -flow $(f_{\alpha,i})_{\alpha \in \mathbf{A}} \in \mathbb{R}_{\geq 0}^{|\mathbf{A}|}$ will be constructed in such a way that $f^i \in \mathbb{R}_{\geq 0}^{|\mathbf{A}|^i}$ will be an entropically feasible, acyclic $[1 : i]$ -flow.

On any arc $\alpha \in \mathbf{A}$, let

$$r_\alpha := \begin{cases} c_\alpha - \sum_{j=1}^{i-1} f_{\alpha,j} & \text{if } \alpha \in \mathbf{A}_f \\ \infty & \text{if } \alpha \in \mathbf{A}_\infty \equiv \mathbf{A} \setminus \mathbf{A}_f \end{cases} \quad (15)$$

be the residual capacity after subtracting the relevant components of f^{i-1} . Note that $r_\alpha \stackrel{(12)}{\geq} 0$ since f^{i-1} is an entropically feasible $[1 : i-1]$ -flow. The next step is to find an acyclic i -flow (Def. 5.1) $q \in \mathbb{R}_{\geq 0}^{|\mathbf{A}|}$ from $\sigma_i \rightsquigarrow \tau_i$ that is *a*) \leq the residual capacity on each arc, and *b*) $\geq H_\infty[X_i]$ on the arc entering τ_i . There are two mutually exclusive cases to consider.

5.1.1 1st Case: \exists an i -Path with No Finite-Capacity Arcs

Denote this i -path by π_e , noting that $r_\alpha = \infty, \forall \alpha \in \pi_e$ by the 2nd line of (15). Set the i -path flows as

$$u_k = \begin{cases} H_\infty[X_i] & \text{if } k = e \\ 0 & \text{otherwise} \end{cases}, \forall k \in [1 : p], \quad (16)$$

and the cycle flows equal to zero in the decomposition (10) (dropping the i -subscripts), so that

$$q_\alpha \stackrel{(10)}{=} \sum_{1 \leq k \leq p: \pi_k \ni \alpha} u_k, \forall \alpha \in \mathbf{A}. \quad (17)$$

Evidently q is acyclic and meets the residual capacity constraint on all arcs in \mathbf{A} . Furthermore, since every i -path passes through the single arc entering τ_i ,

$$q_{\text{in}(\tau_i)} \stackrel{(17)}{=} \sum_{1 \leq k \leq p} u_k \stackrel{(16)}{=} u_e = H_\infty[X_i], \quad (18)$$

satisfying the conditional information constraint.

5.1.2 2nd Case: Every i -Path Has One or More Finite-Capacity Arcs

Observe first that for any arc set $\mathbf{B} \subseteq \mathbf{A}$,

$$\begin{aligned} \sum_{\beta \in \mathbf{B}} c_\beta &\stackrel{(4)}{\geq} \sum_{\beta \in \mathbf{B}} \mathbf{H}^\infty[S_\beta] \equiv \sum_{\beta \in \mathbf{B}} \overline{\lim}_{t \rightarrow \infty} \frac{\mathbf{H}[S_\beta(0:t)]}{t+1} \\ &\geq \overline{\lim}_{t \rightarrow \infty} \frac{1}{t+1} \sum_{\beta \in \mathbf{B}} \mathbf{H}[S_\beta(0:t)] \geq \overline{\lim}_{t \rightarrow \infty} \frac{\mathbf{H}[\mathbf{S}_\mathbf{B}(0:t)]}{t+1} \equiv \mathbf{H}^\infty[\mathbf{S}_\mathbf{B}] \end{aligned} \quad (19)$$

$$\geq \mathbf{H}^\infty \left[\mathbf{S}_\mathbf{B} | X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] \quad (20)$$

$$= \overline{\lim}_{t \rightarrow \infty} \left(\frac{\mathbf{H} \left[\mathbf{S}_\mathbf{B}(0:t) | X_{\mathbf{J}^{i-1} \cup [i+1:n]}(0:t) \right] - \mathbf{H} \left[\mathbf{S}_\mathbf{B}(0:t) | X_{\mathbf{J}^{i-1} \cup [i:n]}(0:t) \right]}{t+1} \right.$$

$$\left. + \frac{\mathbf{H} \left[\mathbf{S}_\mathbf{B}(0:t) | X_{\mathbf{J}^{i-1} \cup [i:n]}(0:t) \right]}{t+1} \right)$$

$$= \overline{\lim}_{t \rightarrow \infty} \left(\frac{\mathbf{I} \left[\mathbf{S}_\mathbf{B}(0:t); X_i(0:t) | X_{\mathbf{J}^{i-1} \cup [i+1:n]}(0:t) \right]}{t+1} \right.$$

$$\left. + \frac{\mathbf{H} \left[\mathbf{S}_\mathbf{B}(0:t) | X_{\mathbf{J}^{i-1} \cup [i:n]}(0:t) \right]}{t+1} \right)$$

$$\geq \mathbf{I}_\infty \left[X_i; \mathbf{S}_\mathbf{B} | X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] + \mathbf{H}^\infty \left[\mathbf{S}_\mathbf{B} | X_{\mathbf{J}^{i-1} \cup [i:n]} \right]$$

$$= \mathbf{I}_\infty \left[X_i; \mathbf{S}_\mathbf{B}, X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] + \mathbf{H}^\infty \left[\mathbf{S}_\mathbf{B} | X_{\mathbf{J}^{i-1} \cup [i:n]} \right], \quad (21)$$

where (19) is due to the subadditivity of joint entropy, (20) holds because conditioning cannot increase entropy, and (21) arises from the mutual independence of X_1, \dots, X_n .

Now, consider the residual capacitated digraph $(\mathbf{V}^i, \mathbf{A}^i, r_{\mathbf{A}^i})$ formed by the i -bundle.⁸ Let q be an acyclic maximal flow on it under the constraints

$$0 \leq q_\alpha \leq r_\alpha, \quad \forall \alpha \in \mathbf{A}^i. \quad (22)$$

By the *Min-Cut Max-Flow Theorem* (see e.g. [4], Thm. 3.5.3) \exists an i -cut $\mathbf{U} \subset \mathbf{V}^i$, consisting of every vertex $v \in \mathbf{V}^i$ for which \exists an undirected path π in $(\mathbf{V}^i, \mathbf{A}^i)$ from σ_i to v s.t.

- (*Forward Slack*) every forward-oriented arc α in π (i.e. pointing from σ_i to v) has $q_\alpha < r_\alpha$, and
- (*Backward Flow*) every backward-oriented arc α in π (pointing from v to σ_i) has $q_\alpha > 0$.

As a consequence of this,

⁸ Here, arcs are permitted to have $r_\alpha = 0$.

$$q_\alpha = r_\alpha, \quad \forall \alpha \in \mathbf{O}^i := \text{OUT}(\mathbf{U}) \cap \mathbf{A}^i, \quad (23)$$

$$q_\alpha = 0, \quad \forall \alpha \in \mathbf{I}^i := \text{IN}(\mathbf{U}) \cap \mathbf{A}^i. \quad (24)$$

Note also that since the cyclic flow components w_1, \dots, w_j in (10) are zero,

$$q_\alpha = \sum_{1 \leq k \leq p: \pi_k \ni \alpha} u_k, \quad \forall \alpha \in \mathbf{A}^i. \quad (25)$$

The i -cut \mathbf{U} evidently depends on the residual capacity vector r . However, the following *purely structural* statements may be made about it:

1. Every arc in \mathbf{O}^i lies in \mathbf{A}_f , i.e. is finite-capacity. Otherwise $q_\alpha \stackrel{(23)}{=} r_\alpha \stackrel{(15)}{=} \infty$, implying by (25) that $u_k = \infty$ on some i -path π_k , which is impossible since every i -path in this case travels over at least one finite-capacity arc.
2. Every arc $\alpha \in \mathbf{O}^i$ is in an i -path that exits \mathbf{U} without re-entering, or else α is in the $[1 : i - 1]$ -bundle. To see this, suppose that every i -path π_k passing through α re-enters \mathbf{U} . Evidently, it must then pass through some arc $\beta \in \mathbf{I}^i$. By (24) $q_\beta = 0$, implying by virtue of (25) and nonnegativity that $u_k = 0$. From (23) and (25), this implies that $r_\alpha = 0$. As $c_\alpha > 0$, it must then hold that $f_{\alpha,j} > 0$ for some $j \in [1 : i - 1]$. As the j -flow $(f_{\alpha,j})_{\alpha \in \mathbf{A}}$ is acyclic by construction, α must then lie on a j -path, by (10).
3. There must be an i -path that leaves \mathbf{U} without re-entering. To see this, suppose in contradiction that every i -path re-enters \mathbf{U} . By the preceding argument, all i -paths must then have associated acyclic flow components $u_k = 0$. Pick any i -path and let v be the last vertex in \mathbf{U} that it traverses before leaving \mathbf{U} without further re-entry. Let ω denote its subpath from $v \rightsquigarrow \tau_i$. By the definition of \mathbf{U} , there is an undirected path π from σ_i to v such that all forward-oriented arcs in it are slack and all backward-oriented arcs carry strictly positive q -flow. Note also that all vertices before v in π must also lie in \mathbf{U} , by construction. From (25), any backward arc in π would have to carry an i -path flow component $u_k > 0$, which would be a contradiction. Consequently, all the arcs in π must be forward-oriented, i.e. π is a directed path in \mathbf{U} from $\sigma_i \rightsquigarrow v$. The concatenation of π with ω then yields an i -path that leaves \mathbf{U} exactly once, a contradiction.
4. Finally, by construction of \mathbf{U} , every vertex v in it must lie on an undirected path π from σ_i to v such that
 - a. every vertex before v in π is also in \mathbf{U} (since the subpath from σ_i to v automatically satisfies the defining forward-slack and backward-flow properties), and
 - b. every reverse-oriented arc in π lies on an i -path that does not re-enter \mathbf{U} (since such arcs must by definition carry positive q -flow, and i -paths that re-enter \mathbf{U} carry zero q -flow).

In other words, \mathbf{U} is a *viable i -cut* (Def. 4.3). By downward dominance (Def. 4.4), $\mathbf{O}^i \cup \text{OUT}(\sigma_{j-1 \cup [i+1:n]})$ structurally dominates $\text{out}(\sigma_i)$ (Def. 3.2), and \mathbf{O}^i is a \mathcal{D}_{i-1} -set (Def. 4.2). Using i -flow conservation,

$$\begin{aligned}
q_{\text{in}(\tau_i)} &= \sum_{\beta \in \mathbf{O}^i} q_\beta - \sum_{\alpha \in \mathbf{I}^i} q_\alpha \stackrel{(23),(24)}{=} \sum_{\beta \in \mathbf{O}^i} r_\beta \\
&\stackrel{(15)}{=} \sum_{\beta \in \mathbf{O}^i} c_\beta - \sum_{\beta \in \mathbf{O}^i, j \in [1:i-1]} f_{\beta,j} \\
&\stackrel{(21)}{\geq} \mathbf{I}_\infty \left[X_i; S_{\mathbf{O}^i}, X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] + \mathbf{H}^\infty \left[S_{\mathbf{O}^i} | X_{\mathbf{J}^{i-1} \cup [i:n]} \right] - \sum_{\beta \in \mathbf{B}, j \in [1:i-1]} f_{j,\beta} \\
&\stackrel{(13)}{\geq} \mathbf{I}_\infty \left[X_i; S_{\mathbf{O}^i}, X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right]. \tag{26}
\end{aligned}$$

As $\text{out}(\sigma_i) \in \text{SDOM} \left(\mathbf{O}^i \cup \text{OUT}(\sigma_{\mathbf{J}^{i-1} \cup [i+1:n]}) \right)$, it follows that $X_i(0 : km - 1)$ is a function of $S_{\mathbf{O}^i}(0 : km - 1)$ and $X_{\mathbf{J}^{i-1} \cup [i+1:n]}(0 : km - 1)$. Consequently, $\forall k \in \mathbb{N}$,

$$\mathbf{I} \left[X_i(0 : km - 1); S_{\mathbf{O}^i}(0 : km - 1), X_{\mathbf{J}^{i-1} \cup [i+1:n]}(0 : km - 1) \right] = \mathbf{H} [X_i(0 : km - 1)].$$

As entropy and mutual information are monotonic, a sandwich argument with $k \rightarrow \infty$ then yields that the RHS of (26) is just $\mathbf{H}_\infty[X_i]$, so that

$$q_{\text{in}(\tau_i)} \geq \mathbf{I}_\infty \left[X_i; S_{\mathbf{O}^i}, X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] = \mathbf{H}_\infty[X_i], \tag{27}$$

as desired.

5.1.3 Construction of f^i in Both Cases

For both cases above, let

$$f_{\alpha,i} := \underbrace{\frac{\mathbf{H}_\infty[X_i]}{q_{\text{in}(\tau_i)}}}_{=:v} q_\alpha \equiv v q_\alpha, \quad \forall \alpha \in \mathbf{A}, \tag{28}$$

where $v \in (0, 1]$ by (27). Clearly, $f_{\mathbf{A},i}$ is still an acyclic i -flow since it just a scaled version of q . Furthermore,

$$\sum_{j=1}^i f_{\alpha,j} \stackrel{(28)}{=} v q_\alpha + \sum_{j=1}^{i-1} f_{\alpha,j} \stackrel{(27)}{\leq} q_\alpha + \sum_{j=1}^{i-1} f_{\alpha,j} \stackrel{(15)}{\leq} c_\alpha.$$

The next step is to verify that $f^i = f_{\mathbf{A} \times [1:i]}$ satisfies the remaining conditions (13)–(14) for an entropically feasible $[1 : i]$ -flow. First (13) is checked. Let \mathbf{E} be any arc-set in \mathcal{D}_i (Def. 4.2). If $\mathbf{E} \cap \mathbf{A}^i = \emptyset$, then

$$\begin{aligned} \sum_{\eta \in \mathbf{E}, j \in [1:i]} f_{\eta,j} &= \sum_{\eta \in \mathbf{E}, j \in [1:i-1]} f_{\eta,j} \\ &\stackrel{(13)}{\leq} \mathbf{H}^\infty \left[\mathcal{S}_{\mathbf{E}} | X_{\mathbf{J}^{i-1} \cup [i:n]} \right] = \mathbf{H}^\infty \left[\mathcal{S}_{\mathbf{E}} | X_{\mathbf{J}^i \cup [i+1:n]} \right] \end{aligned}$$

since $\mathbf{E} \in \mathcal{D}_{i-1}$ automatically, and where the last equality follows because $\mathbf{J}^{i-1} \cup \{i\} = \mathbf{J}^i$. Else if $\mathbf{E} \cap \mathbf{A}^i \neq \emptyset$, write

$$\sum_{\eta \in \mathbf{E}, j \in [1:i]} f_{\eta,j} = \sum_{\eta \in \mathbf{E}} f_{\eta,i} + \sum_{\eta \in \mathbf{E}, j \in [1:i-1]} f_{\eta,j} \quad (29)$$

and bound each sum on the RHS as follows. First, note that since $\mathbf{E} \in \mathcal{D}_{i-1}$,

$$\sum_{\eta \in \mathbf{E}, j \in [1:i-1]} f_{\eta,j} \stackrel{(13)}{\leq} \mathbf{H}^\infty \left[\mathcal{S}_{\mathbf{E}} | X_{\mathbf{J}^{i-1} \cup [i:n]} \right]. \quad (30)$$

Then write

$$\begin{aligned} \sum_{\eta \in \mathbf{E}} f_{\eta,i} &\stackrel{(28)}{=} \sum_{\eta \in \mathbf{E}} v q_{\eta} \stackrel{(25),(17)}{=} v \sum_{\eta \in \mathbf{E}} \left(\sum_{1 \leq k \leq p: \pi_k \ni \eta} u_k \right) \\ &= v \sum_{1 \leq k \leq p} u_k \left(\sum_{\eta \in \mathbf{E}: \eta \in \pi_k} 1 \right) \\ &\leq v \sum_{1 \leq k \leq p} u_k \equiv v q_{\text{in}(\tau_i)} \stackrel{(28)}{=} \mathbf{H}_\infty[X_i], \end{aligned} \quad (31)$$

where the inequality arises because the i -path flows $u_1, \dots, u_p \geq 0$ and each i -path π_k transits over at most one arc in \mathbf{E} . As $\text{out}(\sigma_i) \in \text{SDOM} \left(\mathbf{E} \cup \text{OUT}(\sigma_{\mathbf{J}^{i-1} \cup [i+1:n]}) \right)$, the same arguments that lead to the equality in (27) show that $\mathbf{H}_\infty[X_i] = \mathbf{I}_\infty \left[X_i; \mathcal{S}_{\mathbf{E}}, X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right]$. Substituting this into (31) and then combining with (29) and (30) yields

$$\begin{aligned} \sum_{\eta \in \mathbf{E}, j \in [1:i]} f_{\eta,j} &\leq \mathbf{I}_\infty \left[X_i; \mathcal{S}_{\mathbf{E}}, X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] + \mathbf{H}^\infty \left[\mathcal{S}_{\mathbf{E}} | X_{\mathbf{J}^{i-1} \cup [i:n]} \right] \\ &\stackrel{(20),(21)}{\leq} \mathbf{H}^\infty \left[\mathcal{S}_{\mathbf{E}} | X_{\mathbf{J}^{i-1} \cup [i+1:n]} \right] = \mathbf{H}^\infty \left[\mathcal{S}_{\mathbf{E}} | X_{\mathbf{J}^i \cup [i+1:n]} \right], \end{aligned} \quad (32)$$

since $\mathbf{J}^i = \mathbf{J}^{i-1}$ in this case. This confirms that f^i satisfies (13). As f^{i-1} is an entropically feasible $[1 : i-1]$ -flow, (14) is satisfied $\forall j \in [1 : i-1]$. Using flow conservation,

$$f_{\text{out}(\sigma_i),i} = f_{\text{in}(\tau_i),i} \stackrel{(28)}{=} v q_{\text{in}(\tau_i)} = \mathbf{H}_\infty[X_i],$$

verifying (14) when $j = i$. Thus f^i is an entropically feasible $[1 : i]$ -flow.

By induction, f^n is an entropically feasible $[1 : n]$ -flow, giving the desired (X, c) -feasible multicommodity flow (5)–(7).

5.2 Sufficiency of Multicommodity Flows

The converse part of Thm. 4.1 is easier to establish, since it is not difficult to see that the existence of a feasible multicommodity flow implies achievability. Thus only the key steps are provided below.

Suppose f is an (X, c) -feasible multicommodity flow (5)–(7) on an n -pair network structure Σ , with X stationary, and further suppose that (5) is satisfied strictly. In the decomposition (10) for each i -flow $f_{\mathbf{A},i}$, no cycle flow can enter any sink, since it has no departing arcs. Consequently, the cycle flows may be taken to be zero in (10) without violating (5)–(7), yielding

$$f_{\alpha,i} = \sum_{1 \leq k \leq p_i: \pi_{k,i} \ni \alpha} u_{k,i}, \quad (33)$$

where $\pi_{1,i}, \dots, \pi_{p_i,i}$ are the i -paths and $u_{1,i}, \dots, u_{p_i,i} \geq 0$, the i -path flows. In particular,

$$\mathbf{H}_{\infty}[X_i] \stackrel{(6)}{=} f_{\text{out}(\sigma_i)} = f_{\text{in}(\tau_i)} = \sum_{k=1}^{p_i} u_{k,i}. \quad (34)$$

For an arbitrary $\varepsilon > 0$, divide the time axis \mathbb{W} into epochs of sufficiently long duration $m \in \mathbb{N}$ such that $\forall j \in \mathbb{N}, i \in [1 : n]$,

$$\frac{\mathbf{H}[X_i((j-1)m : jm-1)]}{m} = \frac{\mathbf{H}[X_i(0 : m-1)]}{m} \leq \mathbf{H}_{\infty}[X_i] + \varepsilon, \quad (35)$$

where the first equality arises from stationarity. Next use Huffman coding [6] to losslessly encode each source-block $X_i((j-1)m : jm-1)$, $j \in \mathbb{N}$, into binary code-words $Z_{i,j}$ of variable length $L_{i,j}$, where

$$\mathbf{E}[L_{i,j}] \leq \mathbf{H}[X_i(0 : m-1)] + 1 \stackrel{(35)}{\leq} m\mathbf{H}_{\infty}[X_i] + m\varepsilon + 1. \quad (36)$$

Then partition the bits of $Z_{i,j}$ into p consecutive sub-blocks $Z_{i,j,k}$, $k \in [1 : p_i]$, of length $L_{i,j,k} := \left\lceil \frac{u_{k,i}}{\mathbf{H}_{\infty}[X_i]} L_{i,j} \right\rceil$. This is always possible since $\sum_{k=1}^{p_i} L_{i,j,k} \stackrel{(34)}{\geq} L_{i,j}$, padding the last sub-blocks with zeros if necessary.

Transmit and route each sub-block $Z_{i,j,k}$ along the k -th i -path $\pi_{k,i}$. On every arc $\alpha \in \mathbf{A}$ apart from those leaving sources, let the arc-signal be $S_{\alpha}(t) = 0$ when $t \pmod{m} \neq m-1$ and by $S_{\alpha}(t) = (Z_{i,j,k})_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha}$ when $t \equiv jm-1$, $j \in \mathbb{N}$.⁹ The arc signals leaving sources are set to the respective source signals to satisfy (2). Clearly S is setwise causal (Def. 2.1), since every arc-signal is constructed by routing blocks along acyclic paths. In addition, $\forall \alpha \in \mathbf{A}_f$,

⁹ If an arc is not on any i -path, then its arc signal may be taken to be 0.

$$\begin{aligned} \mathbf{H}^\infty[S_\alpha] &= \frac{1}{m} \mathbf{H} \left[(Z_{i,j,k})_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha} \right] \\ &\leq \sum_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha} \frac{\mathbf{H}[Z_{i,j,k}]}{m} \end{aligned} \quad (37)$$

$$\leq \sum_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha} \frac{\mathbf{E}[L_{i,j,k}]}{m} \quad (38)$$

$$\leq \sum_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha} \frac{1}{m} + \frac{u_{k,i}}{m \mathbf{H}_\infty[X_i]} \mathbf{E}[L_{i,j}]$$

$$\leq \mathbf{O}(1/m) + \sum_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha} u_{k,i} \frac{m \mathbf{H}_\infty[X_i] + m\varepsilon + 1}{m \mathbf{H}_\infty[X_i]}$$

$$= \sum_{i \in [1:n], k \in [1:p_i]: \pi_{k,i} \ni \alpha} u_{k,i} + \mathbf{O}(\varepsilon) + \mathbf{O}(1/m)$$

$$\stackrel{(33)}{=} \sum_{i \in [1:n]} f_{\alpha,i} + \mathbf{O}(\varepsilon) + \mathbf{O}(1/m)$$

$$= f_\alpha + \mathbf{O}(\varepsilon) + \mathbf{O}(1/m) \leq c_\alpha$$

for ε sufficiently small and m sufficiently large. In the above, the bound (37) is due to the subadditivity of entropy, and (38) is due to the fact that the expected number of bits needed to uniquely specify the value of a random variable is never less than its entropy. Furthermore,

$$Y_i(jm - 1) = (Z_{i,j,k})_{k=1}^{p_i} \equiv Z_{i,j} \equiv X_i((j-1)m : jm - 1).$$

Consequently, S is a solution to the n -pairs information network problem (Σ, X, c) , establishing achievability (Def. 2.2).

6 Examples

In this section, several examples are given to illustrate the applicability of Thm. 4.1. However, to begin with a well-known counterexample is discussed.

To avoid cluttering the Figures in this section, arcs leading out of sources and into sinks are not explicitly depicted.

6.1 Butterfly Network

The first example, a 2-pairs butterfly network, is adapted from [15, 11] and depicted in Fig. 1. For this network it is well-known that routing does not achieve coding capacity, and it is a useful exercise to verify that it is not downward dominated.




Fig. 1 Butterfly network

Consider the viable 2-cut having the set $\mathbf{O}^2 = \{\alpha\}$ of outgoing arcs in the 2-bundle. Clearly, both β, γ are downstream of $\mathbf{O}^2 = \mathbf{O}_{*2}^2$, so $\mathbf{C} = \{\alpha, \beta, \gamma\} \subseteq \text{SDOM}(\mathbf{O}^2)$. No other arcs are downstream of \mathbf{C} . Furthermore, the indirect 2-walk concisely represented by $(\varphi, \delta, \epsilon)$ does not pass through α , and neither does the indirect 1-walk (δ, φ, χ) . Thus \mathbf{C} is the smallest set satisfying all the conditions of Def. 3.2), i.e. $\mathbf{C} = \text{SDOM}(\mathbf{O}^2)$. As \mathbf{C} does not include any source or sink arcs, this network is not downward dominated (Def. 4.4) and Thm. 4.1 does not apply.

6.2 Examples that Satisfy Lemma 4.3

Any network where there is at most one (directed) path from any vertex to any other automatically satisfies the conditions of Lemma 4.3, and is therefore downward dominant and structurally routable (Thm. 4.1). This includes in the first instance both directed lines and directed cycles, agreeing with results in [11, 10]. It also covers more complicated structures, for instance directed trees (Fig. 2), and directed cycles arranged in a line or tree structure via one or more gateway nodes (Fig. 3). In all these networks, routing achieves coding capacity regardless of where sources and sinks are placed.

In networks where there are vertex pairs with two or more connecting paths, downward dominance will still hold by virtue of Lemma 4.3) if there is at most one path between each pair of source and sink vertices, or at least from each σ_s to each τ_h , where $1 \leq s \leq h \leq n$. Examples include directed versions of the undirected Okamura-Seymour network (Fig. 4).




Fig. 2 A directed tree. Sources and sinks may be attached to any of the nodes depicted.




Fig. 3 A tree of directed cycles. Sources and sinks may be attached to any of the nodes depicted.

6.3 Examples that Satisfy Lemma 4.2

Now consider the acyclic 2-pairs network in Fig. 5. Observe that there is one 1-path, concisely represented by the arc-sequence $\beta\varepsilon$, but two 2-paths, $\alpha\beta$ and γ . Hence Lemma 4.3 cannot be applied. Neither would it become applicable if the indices 1




Fig. 4 A directed version of the Okamura-Seymour Network. Only one path exists from any source to any sink.

and 2 were relabelled 2' and 1' respectively. To see this, consider the viable 2'-cut with $\mathbf{O}^{2'} = \{\beta\}$. Clearly $\mathbf{A}^{2'} \cap \mathbf{O}^{2'} \neq \emptyset$ and $\mathbf{A}^{1'} \cap \mathbf{O}^{2'} \neq \emptyset$, since the 1'-path $\alpha\beta$ and 2'-path $\beta\epsilon$ both pass through $\mathbf{O}^{2'}$. However, the path γ from σ'_1 to τ'_1 does not.

In this instance, Lemma 4.2 can be applied. The possible viable 2-cuts have sets \mathbf{O}^2 of outgoing arcs in the 2-bundle equal to either $\{\alpha, \gamma\}$ or $\{\beta, \gamma\}$. In the first case, \mathbf{O}^2 has no intersection with any arcs in the 1-bundle, and all 2-paths obviously pass through it. In the second case, all paths from σ_s to σ_h , $1 \leq s \leq h \leq 2$, pass through \mathbf{O}^2 . This the requirements of the Lemma are met and the network is downward dominant.

Another example of the use of Lemma 4.2) is the the cyclic 2-pairs network of Fig. 6. Observe that there is one 1-path, $\epsilon\beta$ and two 2-paths, φ and $\beta\gamma$. The possible viable 2-cuts have sets \mathbf{O}^2 of outgoing arcs in the 2-bundle equal to either $\{\varphi, \beta\}$ or $\{\varphi, \gamma\}$. In the second case, \mathbf{O}^2 has no intersection with any arc in the 1-bundle, and all 2-paths obviously pass through it. In the first case, \mathbf{O}^2 intersects all 2-paths and a 1-path, $\epsilon\beta$, and it can be seen that all from σ_s to σ_h , $1 \leq s \leq h \leq 2$, pass through \mathbf{O}^2 . This the requirements of the Lemma are met and the network is downward dominant.

7 Conclusion

This chapter examined the routability of possibly cyclic n -pairs information networks from a structural perspective. The concept of downward dominance was in-




Fig. 5 An acyclic network covered by Lemma 4.2.




Fig. 6 A cyclic network covered by Lemma 4.2.

roduced, and it was shown that for networks with downward dominated structures, routability and achievability are equivalent, i.e. a given combination of source signals, demand rates and channel capacities is achievable iff the network supports a feasible multicommodity flow.

Downward dominance is a conservative structural condition, and future work will focus on trying to relax it. The inductive nature of the proof of necessity here requires it directly, so any generalisation may need a very different analysis technique.

Acknowledgements The author acknowledges discussions on the decentralised control version of this problem with Prof. Rob Evans at the University of Melbourne. He is also indebted to the anonymous ISIT11 reviewer who pointed out an error in the original version of (14).

References

1. Adler, M., Harvey, N.J.A., Jain, K., Kleinberg, R., Lehman, A.R.: On the capacity of information networks. In: Proc. 17th Annual ACM-SIAM Symp. Discrete Algorithms, pp. 241–50. Miami, USA (2006)
2. Ahlswede, R., Cai, N., Li, S.W.R., Yeung, R.W.: Network information flow. *IEEE Trans. Info. The.* **46**(4), 1204–16 (2000)
3. Antsaklis, P., Baillieul, J. (eds.): Special Issue on the Technology of Networked Control Systems, in Proc. IEEE, vol. 95. IEEE (2007)
4. Bang-Jensen, J., Gutin, G.: *Digraphs Theory, Algorithms and Applications*. Springer (2007)
5. Cannons, J., Dougherty, R., Freiling, C., Zeger, K.: Network routing capacity. *IEEE Trans. Info. The.* **52**(3), 777–788 (2006)
6. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. Wiley (1991)
7. Dougherty, R., Freiling, C., Zeger, K.: Insufficiency of linear coding in network information flow. *IEEE Trans. Info. The.* **51**(8), 2745–59 (2005)
8. Fragouli, C., Soljanin, E.: *Network Coding Applications*. Now Publishers (2008)
9. Han, T.S.: Multicasting multiple correlated sources to multiple sinks over a noisy channel network. *IEEE Trans. Info. The.* **57**(1), 4–13 (2011)
10. Harvey, N.J.A., Kleinberg, R., Lehman, A.R.: On the capacity of information networks. *IEEE Trans. Info. The.* **52**(6), 2345–64 (2006)
11. Kramer, G., Savari, S.A.: Edge-cut bounds on network coding rates. *Jour. Network and Syst. Management* **14**(1), 49–67 (2006)
12. Leighton, T., Rao, S.: Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Jour. ACM* **46**(6), 787–832 (1999)
13. Li, S.Y.R., Yeung, R.W., Cai, N.: Linear network coding. *IEEE Trans. Info. The.* **49**(2), 371–81 (2003)
14. Nair, G.N.: When is n -pairs information a multicommodity flow? In: *IEEE Int. Symp. Info. Theory Proc.*, pp. 169–73. St. Petersburg, Russia (2011)
15. Yeung, R.W.: Multilevel diversity coding with distortion. *IEEE Trans. Info. The.* **41**(2), 412–22 (1995)
16. Zhou, K., Doyle, J., Glover, K.: *Robust and Optimal Control*. Prentice Hall, USA (1996)