2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Network Coding meets Decentralized Control: Capacity-Stabilizabililty Equivalence

Se Yong Park and Anant Sahai {separk,sahai}@eecs.berkeley.edu

***Abstract*— The main difference between centralized and decentralized control is the communication. Controllers in a decentralized system can communicate with each other to achieve their common goal. In this paper, we argue that even linear time-invariant controllers in a decentralized linear system "communicate" via linear network coding to stabilize the plant. To justify this argument, we propose an algorithm to "externalize" the implicit communication between controllers that we believe must be occurring to stabilize the plant. Based on this, we show that the stabilizability condition for decentralized linear systems comes from an underlying communication limit, which can be described by an algebraic mincut-maxflow theorem.**

I. Introduction

This paper is inspired by the similarity between the algebraic characterization of fixed modes [1] and the min-cut bound in information theory [2]. The algebraic condition for λ to be a fixed mode [1, Theorem 4.1] is

$$\min_{V \subseteq \{1,2,\cdots,v\}} \text{rank} \begin{bmatrix} A-\lambda I & B_V \\ C_{V^c} & 0 \end{bmatrix} \geq \dim(A). \quad (1)$$

The information-theoretic min-cut bound [2, Theorem 15.10.1] is

$$\min_{V \subseteq \{1,2,\cdots,v\}} I(X_V; Y_{V^c} | X_{V^c}) \geq \sum_{i \in V, j \in V^c} R_{ij}. \quad (2)$$

We can see that the left-hand sides of both (1) and (2) have a minimization over all subsets. Moreover, in noiseless relay networks the mutual information is essentially equal to the rank of the channel matrix[1] [3, Theorem 4.4]. Therefore, the left-hand sides of (1) and (2) can be considered to be exactly the same. Identifying the right hand sides of (1) and (2) with each other, we can see that the dimension of A corresponds to a rate of information flow. Moreover, fixed modes are closely connected to stabilizability. Thus, we can conjecture that a decentralized system is stabilizable if and only if enough information flow can be supported to stabilize the plant. In this paper, we make this conjecture rigorous.

This work was supported by a Samsung Scholarship and the National Science Foundation (CCF-0729122, CNS-0932410)

The authors are with the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley.

[1]Information is traditionally measured in bits and the rate of bits that a channel can carry is computed by the mutual information $I(X;Y)$. However, in continuous-alphabet channels like the AWGN (additive white Gaussian noise) channel, the mutual information depends crucially on the signal-to-noise ratio and scales as log SNR. It was noticed that when the channel has multiple-inputs and multiple-outputs (MIMO) — like when there are multiple antennas involved in wireless communication — the mutual information increases as the rank of the channel matrix **times** log SNR. This fact inspired the creation of the finite-field noiseless MIMO channel model, within which the mutual information is equal to the rank of the channel matrix multiplied by the log of the field size. Therefore, the rank can be considered another measure for information, as measured in units of dimensions or degrees-of-freedom.

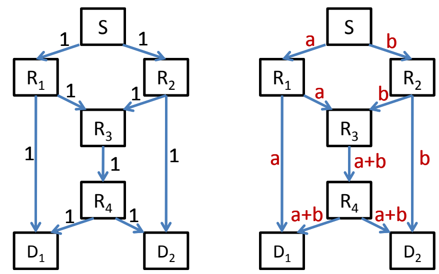

Fig. 1. "Butterfly" example that shows the benefit of network coding in multicast problems.

First, let's review perspectives on information flow in communication networks. Historically, information in a network was believed to behave like a physical commodity. The network was modeled using a graph, and the information flow was thought of as commodities to be transported from the source to the destination by routing them through the nodes. The most important result is the celebrated mincut-maxflow theorem [4], [5], which reveals that the maximum amount of commodity flow through a graph is equal to the minimum cut of the graph. Moreover, this maximum flow is achievable by a routing scheme. This optimality result made researchers stick to routing solutions for decades.

However, in [6] it was found that information flow in networks does not really behave like physical commodities do. The "butterfly" example of Fig. 1 (introduced in [6]) shows this. The source S has two bits of binary data a, b which both destinations D_1, D_2 want to receive. Every edge has unit capacity as we see in the left side of Fig. 1. We can easily check that by routing alone, it is impossible to send a, b to both destinations. The link between R_3 and R_4 is a bottleneck. D_2 wants it used to send the bit sent from S to R_1, and D_1 wants it to send the bit from S to R_2.

Instead of routing, let the relay R_3 mix its incoming data a, b and relay $a+b$ (XOR in the binary field). As shown in the right side of Fig. 1, the destinations D_1 and D_2 receive $a, a+b$ and $b, a+b$ respectively. Thus, both destinations can decode a, b. The paradigm that allows processing or mixing data at the nodes is called *network coding*. The set of problems like Fig. 1 which have one source and multiple destinations that want the same data are called *multicast* problems.

Even if physical commodity flows (which we can only route) and information flows (which we can process and mix) are different, the graph-theoretic concepts and insights originally developed for commodity flows continue to be helpful. The main difference is that the amount of flow, which is naturally measured by the number (or weight or volume) of commodities in physical commodity flows, must instead be measured in "dimensions" of the signal for information flows. However, the mincut-maxflow theorem remains the main tool to understand network information flows. For example, in the multicast problem the relevant mincut is the minimum of the mincut to each destination, and the mincut-maxflow theorem still holds [6]. Moreover, this maximum flow is achievable by linear time-invariant network coding [7].

Once information-theorists had the freedom to mix and process signals inside the nodes that they could design, they also started to consider such operations as also existing outside these nodes. The signals from the relay nodes could be broadcast to multiple receiving nodes or superposed with other signals at a receiving node. In fact, such extensions were a natural fit to wireless communication [3]. The operations outside the nodes modeled the communication channels and such wireless channel models had long been valuable even when restricted to be linear time-invariant.

At this point, we can see the similarities between network-coding problems [3] and decentralized-linear-control problems [8]. The channels of the network which we cannot design can be considered as the linear plant. The source, relays and destination nodes which we can design can be considered as decentralized controllers. Just as decentralized controllers process and combine their observations to generate their control inputs, the relay nodes process and combine their incoming signals from the channel to generate their outgoing signals.

Despite these similarities, many differences between the communication and control problems had been preventing a firm connection being made between them. First of all, network-coding information-theorists work in finite fields, whereas control-theorists default to infinite fields like the reals or complex numbers. Moreover, information-theorists tend not to have any explicit state in the system, preferring an input-output perspective. Most importantly, the information-theorists have a clearly specified source and destination, and their goal is to push information from one to the other. Control-theorists tend not to have explicit sources and destinations, and instead there is a dynamic evolution that needs to be controlled or stabilized.

The main goal of this paper is to bridge these differences and make a concrete connection between network coding and decentralized linear control. As shown in Theorems 3, we will prove that if a decentralized linear system is LTI[2]-stabilizable, then there must exist a corresponding implicit information flow sufficient to stabilize the system.

The rest of the paper is organized as follows: In section II, we introduce some preliminary facts about decentralized linear systems and LTI networks. Section III shows a representative example that clearly illustrates the implicit information flows in decentralized systems. Section IV gives the main result of the paper, the capacity-stabilizability equivalence theorem. Because of space limitations here, we defer the technical details and the further discussions to [11].

II. Preliminary

A. Decentralized Linear System

Decentralized linear systems have multiple controllers, each of which has its own observations and its own control inputs. Formally, the decentralized linear system, $\mathcal{L}(A, B_i, C_i)$, is defined as follows:[3]

$$
\begin{aligned}
x[n+1] &= Ax[n] + B_1 u_1[n] + \cdots + B_v u_v[n] \qquad (3)\\
y_1[n] &= C_1 x[n]\\
&\vdots\\
y_v[n] &= C_v x[n]
\end{aligned}
$$

where $A \in \mathbb{C}^{m\times m}, B_i \in \mathbb{C}^{m\times q_i}$ and $C_i \in \mathbb{C}^{r_i\times m}$. Then, an interesting question is under what conditions such systems are stabilizable using only LTI controllers:

Definition 1 (Stabilizability): A decentralized linear system is called LTI-stabilizable if there exist linear time-invariant (LTI) controllers $\mathcal{K}_i$ (possibly with internal memories) that connect y_i to u_i whose resulting closed-loop system has only stable poles.

The stabilizability condition for a decentralized linear system is given in [8] using the concept of fixed modes.

Definition 2: [8, Definition 2] λ is called a fixed mode of $\mathcal{L}(A, B_i, C_i)$ if $\lambda \in \bigcap_{K_i\in\mathbb{C}^{q_i\times r_i}} \sigma(A + \sum_{1\le i\le v} B_i K_i C_i)$ where $\sigma(\cdot)$ is the set of eigenvalues of the matrix.
The intuition behind this definition is that if an eigenvalue is fixed for all choices of the controllers, this eigenvalue is either unobservable or uncontrollable. Thus, if we have unstable fixed modes, we cannot stabilize the plant.

Theorem 1: [8, Theorem 1] $\mathcal{L}(A, B_i, C_i)$ is stabilizable if and only if all of its fixed modes are within the unit circle.
The algebraic characterization of fixed modes (1) is reported in [1]. This condition turns out to be a special case of the mincut-maxflow theorem.

B. LTI Networks

We now introduce communication-theoretic notions for linear relay networks[4]. In our companion paper [12], a point-to-point LTI network $\mathcal{N}(z)$ is defined as follows: There is

[2]It is in our focus on stabilizability using only linear *time-invariant* control laws that the results in this paper differ from the results in [9], [10] where *time-varying* control laws are permitted. The overall perspectives however are compatible in that we are also interested in cutsets and information flows.

[3]In this paper, we consider discrete-time systems since they are easier to connect to communication theory. We believe that the underlying phenomena discussed here also exist in continuous-time. Furthermore, we assume the matrices here are complex since we will use the Jordan form which can be complex. However, if the system were real we could prove corresponding results restricting the controller design to be real without changing the stabilizability condition.

[4]The LTI networks considered here are essentially the same as the linear deterministic model studied in [3] except that the LTI network restricts the relay design to be linear time-invariant and the underlying field is complex rather than a finite field.

one transmitter node labeled by "tx", one receiver node labeled by "rx", and v relay nodes labeled by numbers between 1 to v. Each node generates inputs to the channels and receives outputs from the channels. The transmitter node has only inputs, the receiver node has only outputs, and the relay nodes have both. Every input is potentially connected to every output by the channel matrix $H_{i,j}(z)$ for $i, j \in \{tx, 1, \cdots, v, rx\}$, including possibly self-loops. Because the channel is LTI, each element of the channel matrix is given in z-transform[5], *i.e.* belongs to $\mathbb{C}[z]$ which is the rational function field on z with coefficients in $\mathbb{C}$. Each relay node is free to connect its channel outputs with its channel inputs by a matrix K_i, which is of the form $\begin{bmatrix} k_{i,1,1} & \cdots & k_{i,1,b} \\ \vdots & \ddots & \vdots \\ k_{i,a,1} & \cdots & k_{i,a,b} \end{bmatrix}$. Each element of $K_1, \cdots, K_v$ is considered as a different variable that takes a value from $\mathbb{C}[z]$ when it is realized[6].

Denote the transfer matrix from the transmitter to the receiver of the LTI network as $G_{tx,rx}(z, K_i)$. The elements of $G_{tx,rx}(z, K_i)$ belong to $\mathbb{C}[z, K_i]$ which is the rational function field on z and the elements of K_i with coefficients in $\mathbb{C}$. It is well-known that the capacity of MIMO (multiple-input multiple-output) channels is closely related to the rank of the channel matrix [13].

Definition 3 (Degree of Freedom Capacity): For a given LTI network $\mathcal{N}(z)$, we say that the degree of freedom (d.o.f.) capacity of the network $\mathcal{N}(z)$ is k if its transfer matrix $G_{tx,rx}(z, K_i)$ is rank k.

Therefore, the object of relay design is to maximize the transfer-matrix rank from the transmitter input to the receiver output. Fortunately, we can easily show that almost all generic choices of K_i from $\mathbb{C}[z]$ achieve the maximum feasible rank of the transfer matrix, which is $\mathrm{rank}(G_{tx,rx}(z, K_i))$. Here, since $\mathrm{rank}(G_{tx,rx}(z, K_i))$ is the maximum d.o.f. that we can communicate through the LTI network, it is called the maxflow or the capacity of the LTI network.

It has to be mentioned that by choosing the relay gain matrices K_i we are allowing the relays to process the signals rather than just routing them. Thus, to achieve the maxflow of the LTI network the relays communicate via network coding.

One key fact about LTI networks is that the well-known mincut-maxflow theorem [5], [4] can be extended to them. As proved in our companion paper [12], the rank of the transfer matrix is equal to the rank of the mincut.

Theorem 2 (Algebraic Mincut-Maxflow Theorem): [12] Consider the LTI network $\mathcal{N}(z)$ defined above. For sets $A, B \subseteq \{tx, 1, \cdots, v, rx\}$, let $H_{A,B}$ be the channel matrix from the inputs generated by the nodes in A to the outputs observed by the nodes in B, then

$$\begin{aligned} &\mathrm{rank}(G_{tx,rx}(z, K_i)) \\ &= \min_{V \subseteq \{tx,1,\cdots,v,rx\}, V \ni tx, V \not\ni rx} \mathrm{rank}(H_{V,V^c}(z)). \end{aligned}$$

[5]We assume that all channels are causal so there is no question regarding the relevant regions of convergence.

[6]In other words, relays are allowed to have memory.

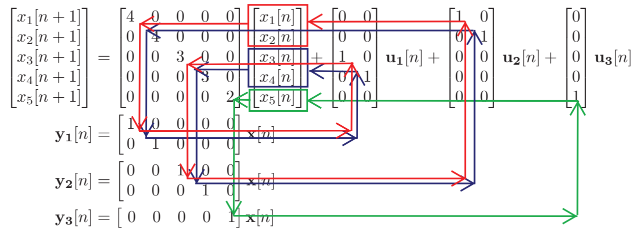

Fig. 2. An example of an implicit information flow in a decentralized linear system.

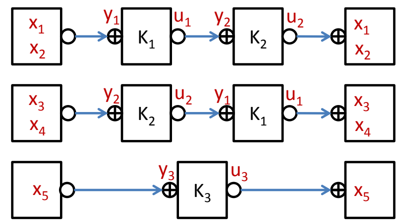

Fig. 3. Conceptual representations of the information flows within the example of Fig. 2

In this paper, we will also consider LTI networks at a specific generalized frequency, $z = \lambda$.To indicate that the LTI network is considered at the generalized frequency $z = \lambda$, we write the network as $\mathcal{N}(\lambda)$. Then, the capacity definition is naturally generalized to $\mathcal{N}(\lambda)$.

Definition 4: For a given LTI network $\mathcal{N}(z)$, we say that the degree of freedom (d.o.f.) capacity of the network $\mathcal{N}(z)$ is k at frequency $z = \lambda$ if its transfer matrix $G_{tx,rx}(\lambda, K_i)$ is rank k.

Here we can see that the transfer matrix only makes sense at $z = \lambda$ when it does not have any pole at λ. Thus, we assume that $H_{i,j}$ has no pole at $z = \lambda$ and also restrict the relay design K_i to those rational functions on z that do not have a pole at λ. Then, the algebraic mincut-maxflow theorem also holds for $\mathcal{N}(\lambda)$ as before.

Corollary 1: Given the LTI network $\mathcal{N}(\lambda)$ above,

$$\begin{aligned} &\mathrm{rank}(G_{tx,rx}(\lambda, K_i)) \\ &= \min_{V \subseteq \{tx,1,\cdots,v,rx\}, V \ni tx, V \not\ni rx} \mathrm{rank}(H_{V,V^c}(\lambda)). \end{aligned}$$

Proof: Easily follows from Theorem 2. See [11] for the details. ■

In summary, the stabilizability of decentralized linear systems is related to fixed modes, and the capacity of networks is related to the rank of transfer matrices.

III. Example: Information Flow in a Decentralized Linear System

Before we discuss a general algorithm to externalize the implicit communication between controllers, it will be helpful to see the information flows that we want to capture in an illustrative example. By now, we have mounting evidence[7] that in linear systems, the states themselves are the sources and, at the same time, the destinations of information flows. Consider a linear plant controlled by one controller. The states of the system will be excited by the disturbance, *i.e.* the states are generating uncertainties. Then, the states will be observed by the controller, *i.e.* the uncertain information flows from the state to the controller. Finally, the controller will compensate for the disturbance, *i.e.* the information flows back to the states.

When there is more than one controller, the situation becomes more complicated since the controllers can implicitly communicate with each other through the plant [20], [21]. The example shown in Fig. 2 (adapted from [22]) illustrates this phenomenon. As we can see, the states $x_1[n]$ and $x_2[n]$ are associated with the eigenvalue 4. However, the controller $\mathcal{K}_1$ can only observe $x_1[n], x_2[n]$, the controller $\mathcal{K}_2$ can only control $x_1[n], x_2[n]$, and the controller $\mathcal{K}_3$ can neither observe nor control $x_1[n], x_2[n]$. Therefore, to stabilize $x_1[n], x_2[n]$ the controller $\mathcal{K}_1$ intuitively has to relay its observations to controller $\mathcal{K}_2$ through the implicit channel provided by the states $x_3[n], x_4[n]$.

The red arrow of Fig. 2 shows the information flow to stabilize $x_1[n], x_2[n]$. First, $x_1[n], x_2[n]$ are observed by $\mathcal{K}_1$ through $\mathbf{y_1}[n]$. Then, $\mathcal{K}_1$ relays its observations to $\mathcal{K}_2$ by $\mathbf{u_1}[n]$ through the channel $x_3[n], x_4[n]$. $\mathcal{K}_2$ receives the relayed signals through $\mathbf{y_2}[n]$, and finally controls the states by $\mathbf{u_2}[n]$. Thus, we expect that the implicit information flow to stabilize $x_1[n], x_2[n]$ should be roughly representable as the first LTI network of Fig. 3. We can see the same kind of information flow to stabilize the states $x_3[n], x_4[n]$ as indicated by the blue arrow. Meanwhile the state $x_5[n]$ can be stabilized by the controller $\mathcal{K}_3$ as indicated by the green arrow. Conceptually, these information flows can be represented as the second and third LTI networks of Fig. 3.

Here, we can notice some interesting points. First, we are dividing the states according to their associated eigenvalues. In this example, the states are first divided into three sets $\{x_1[n], x_2[n]\}$, $\{x_3[n], x_4[n]\}$ and $\{x_5[n]\}$, and the information flows for these sets are considered separately. Moreover, in each information flow the states associated with the same eigenvalue are considered as both sources and destinations of the information. The remaining states are considered as the channels that are available to implicitly carry this information flow. The controllers themselves are considered as relays.

We can also see the connection between stabilizability and capacity. The eigenvalue 4 has two associated states, $x_1[n]$ and $x_2[n]$. Thus, we can think that the source has 2 d.o.f. to transmit. This information can be successfully transferred since the channel provided by the states $x_3[n]$ and $x_4[n]$ has d.o.f. capacity 2, and so the eigenvalue 4 is not a fixed mode. However, if we remove the state $x_4[n]$ from the system, the implicit channel's d.o.f. capacity becomes 1. Thus, a source with 2 d.o.f. cannot be transferred, and the eigenvalue 4 becomes a fixed mode.

[7]We return to this point in the conclusion, but the evidence here has largely come from contexts in which the communication is explicitly present. On one side, papers like [14], [15], [16] construct feedback communication systems that use unstable states to encode the desired messages. This provides strong evidence for the states acting as information sources. On the other side, papers like [17], [18], [16] talk about networked control systems in which the communication demands on the network come from the states. These argue persuasively for the states in a control system as being destinations of information flows since control and estimation are intimately linked together. The "thermodynamic" perspective on the Kalman filter presented in [19] suggests strongly that such information flows exist even when there is no explicit communication going on.

IV. Externalization of Implicit Communication

In this section, we discuss how to externalize the implicit communication in decentralized linear systems. The main idea can be considered as the reverse of the algebraic approach to network coding. In [7], Koetter and Medard considered network coding as an algebraic problem. In other words, they found that what is important about networks (graphical objects) in network coding is their transfer functions (algebraic objects). What we do is the opposite. First, we will find transfer functions which are closely connected to fixed modes. Then, we will find the LTI networks whose transfer functions these are.

A. Jordan-Form Externalization

It turns out that what is important in externalization is the right choice of transfer function. In section III we saw that the source and the destination of the information flows are the states. However, while fixed modes come from eigenvalues, there is no correspondence between eigenvalues and the original states in a linear system for a general matrix A. To find a natural correspondence between eigenvalues and states, we convert A into Jordan normal form. By a similarity transform an arbitrary linear system $\mathcal{L}(A, B_i, C_i)$ can be converted to an equivalent linear system $\mathcal{L}(A', B_i', C_i')$ with the matrix A' in Jordan form [23].

As we discussed in Section III, to check if an eigenvalue λ is a fixed mode, it is enough to examine the transfer matrix from the states associated with Jordan blocks corresponding to the eigenvalue λ to themselves. For that purpose, we introduce an auxiliary input and output to the closed loop system.

To understand the core ideas, we first consider a diagonal A matrix, *i.e.* $A = \begin{bmatrix} \lambda I_{m_\lambda \times m_\lambda} & 0 \\ 0 & A' \end{bmatrix}$ where $I_{m_\lambda \times m_\lambda}$ is a $m_\lambda \times m_\lambda$ identity matrix, and A' is a diagonal matrix whose diagonal elements are not equal to λ. Because the matrix is diagonal, each Jordan block is just a 1×1 matrix and so m_λ can be thought of as the number of Jordan blocks associated with λ. We will introduce auxiliary input $u_\lambda[n]$ and output $y_\lambda[n]$ that control and observe the states corresponding to the eigenvalue λ. For this, we define B_λ and C_λ as follows:

$$C_\lambda = \begin{bmatrix} I_{m_\lambda \times m_\lambda} & 0 \end{bmatrix}, B_\lambda = \begin{bmatrix} I_{m_\lambda \times m_\lambda} & 0 \end{bmatrix}^T. \qquad (4)$$

Then, the closed loop system is given as

$$x[n+1] = (A + \sum_{1 \le i \le v} B_i K_i C_i) x[n] + B_\lambda u_\lambda[n]$$
$$y_\lambda[n] = C_\lambda x[n].$$

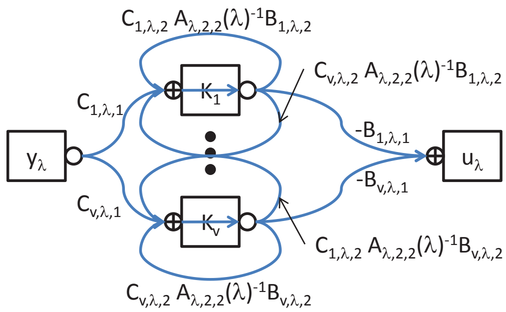

Fig. 4. The graphical representation of $\mathcal{N}_{jd,\lambda}(\lambda)$

Let's set

$$(zI - A) = \begin{bmatrix} A_{\lambda,1,1}(z) & A_{\lambda,1,2}(z) \\ A_{\lambda,2,1}(z) & A_{\lambda,2,2}(z) \end{bmatrix}$$

$$C_i = \begin{bmatrix} C_{i,\lambda,1} & C_{i,\lambda,2} \end{bmatrix}, B_i = \begin{bmatrix} B_{i,\lambda,1} \\ B_{i,\lambda,2} \end{bmatrix}$$

where $A_{\lambda,1,1}(z)$ is a $m_\lambda \times m_\lambda$ matrix, $B_{i,\lambda,1}$ is a $m_\lambda \times q_i$ matrix, $C_{i,\lambda,1}$ is a $r_i \times m_\lambda$ matrix, and the others are the proper implied dimensions. Here, by construction, we can see $A_{\lambda,1,1}(\lambda) = 0$, $A_{\lambda,1,2}(\lambda) = 0$, $A_{\lambda,2,1}(\lambda) = 0$, and $A_{\lambda,2,2}(\lambda)$ is invertible.

Then, we can see that the transfer function from $u_\lambda(z)$ to $y_\lambda(z)$ is given as follows:

$$y_\lambda(z) = \begin{bmatrix} I & 0 \end{bmatrix} \left(\begin{bmatrix} A_{\lambda,1,1}(z) & A_{\lambda,1,2}(z) \\ A_{\lambda,2,1}(z) & A_{\lambda,2,2}(z) \end{bmatrix} - \sum_{1 \le i \le v} \begin{bmatrix} B_{i,\lambda,1} K_i C_{i,\lambda,1} & B_{i,\lambda,1} K_i C_{i,\lambda,2} \\ B_{i,\lambda,2} K_i C_{i,\lambda,1} & B_{i,\lambda,2} K_i C_{i,\lambda,2} \end{bmatrix} \right)^{-1} \begin{bmatrix} I \\ 0 \end{bmatrix} u_\lambda(z)$$

Since the fixed modes show up as poles in the transfer function, checking whether λ is a fixed mode involves checking whether the transfer function from $u_\lambda[n]$ to $y_\lambda[n]$ has a fixed pole. However, checking poles is mathematically troublesome since it results in division by zero. Thus, instead we inspect the zeros of the formal transfer function from $y_\lambda[n]$ to $u_\lambda[n]$. As shown in [11], the formal transfer function from $y_\lambda(z)$ to $u_\lambda(z)$, $G_{jd,\lambda}(z, K_i)$, is given as

$$\begin{aligned} G_{jd,\lambda}(z, K_i) &= (A_{\lambda,1,1}(z) - \sum_{1 \le i \le v} B_{i,\lambda,1} K_i C_{i,\lambda,1}) \\ &+ (A_{\lambda,1,2}(z) - \sum_{1 \le i \le v} B_{i,\lambda,1} K_i C_{i,\lambda,2}) \\ &\cdot (I - (I - A_{\lambda,2,2}(z) + \sum_{1 \le i \le v} B_{i,\lambda,2} K_i C_{i,\lambda,2}))^{-1} \\ &\cdot (-A_{\lambda,2,1}(z) + \sum_{1 \le i \le v} B_{i,\lambda,2} K_i C_{i,\lambda,1}). \end{aligned}$$

Furthermore, in [11] it is proved that the LTI network $\mathcal{N}_{jd,\lambda}(\lambda)$ shown in Fig. 4 has the transfer function $G_{jd,\lambda}(\lambda, K_i)$. Here, the equations on the edges of $\mathcal{N}_{jd,\lambda}(\lambda)$ represent the channel matrices of LTI network, *i.e.*

$$\begin{aligned} &H_{tx,rx}(\lambda) = 0, \\ &H_{tx,i}(\lambda) = C_{i,\lambda,1}, \\ &H_{i,rx}(\lambda) = -B_{i,\lambda,1}, \\ &H_{i,j}(\lambda) = C_{j,\lambda,2} A_{\lambda,2,2}(\lambda)^{-1} B_{i,\lambda,2}. \end{aligned}$$

Now, we are ready to state a capacity-stabilizability equivalence theorem.

Theorem 3: (Capacity-Stabilizability Equivalence) Given the above definitions, the following statements are equivalent.
(1) λ is a fixed mode of the decentralized linear system $\mathcal{L}(A, B_i, C_i)$
(2) $\text{rank}(G_{jd,\lambda}(\lambda, K_i)) < m_\lambda$
(3) (transfer matrix rank of LTI network $\mathcal{N}_{jd,\lambda}(\lambda)) < m_\lambda$
(4) (mincut rank of the LTI network $\mathcal{N}_{jd,\lambda}(\lambda)) < m_\lambda$
(5) $\min_{V \subseteq \{1,\cdots,v\}} \text{rank} \begin{bmatrix} 0 & -B_{V,\lambda,1} \\ C_{V^c,\lambda,1} & C_{V^c,\lambda,2} A_{\lambda,2,2}(\lambda)^{-1} B_{V,\lambda,2} \end{bmatrix} < m_\lambda$

Proof: See [11]. ■

This theorem can be generalized to arbitrary Jordan forms A by introducing auxiliary inputs and outputs from the states associated with λ to themselves. In fact, we can further reduce the dimension of auxiliary input and output. To decide whether λ is a fixed mode, it is enough to examine the transfer matrix from the right-bottom elements of the Jordan blocks corresponding to the eigenvalue λ to their left-top elements. See [11] for the details.

Remark 1: $y_\lambda(z)$ is the signal assigned to the transmitter of $\mathcal{N}_{jd,\lambda}(z)$, and $u_\lambda(z)$ is the signal assigned to the receiver of $\mathcal{N}_{cn}(z)$. Thus, the LTI network connects the states $x[n]$ to themselves, which complies with our discussion of section III.

Remark 2: The statement (1) of the theorem is directly connected to stabilizability by Theorem 1, and statement (3) of the theorem is the d.o.f. capacity of the network at the frequency $z = \lambda$. Thus, this theorem reveals a fundamental equivalence between stabilizability and capacity.

Remark 3: In fact, different externalization is possible depending on the choice of auxiliary input and output. This externalization is minimal in the sense that the dimensions of the transmitter input signal and the receiver output signals are minimal, and the direct link between the transmitter and the receiver does not exist. The canonical externalization corresponding to the maximal is shown in [11] which always result in a simple topology network.

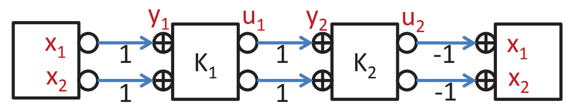

Fig. 5. Jordan-form externalization of the system of Fig. 2 for $\lambda = 4$

The LTI network of Fig. 5 shows the Jordan-form externalization of the Fig. 2 example for $\lambda = 4$. We can easily see that the LTI network of Fig. 5 agrees with the first LTI network of Fig. 3. The information generated at $x_1[n], x_2[n]$ is first observed by the controller $\mathcal{K}_1$, then relayed to the controller $\mathcal{K}_2$, and finally returned to $x_1[n], x_2[n]$. Here, the controller $\mathcal{K}_3$ is correctly omitted since it does not affect the transfer function of the relevant LTI network.

V. Conclusion

In this paper, LTI-stabilizability of decentralized linear systems is found to be equivalent to having sufficient capacity in the relevant LTI networks. We gave an algorithm to make explicit the communication networks that represent the implicit communication required to stabilize the plant. The stabilizability condition of decentralized systems is then easily interpreted using mincut conditions on the corresponding networks. Each eigenvalue is viewed separately, and the number of Jordan blocks corresponding to that eigenvalue corresponds to the number of degrees-of-freedom of implicit communication required to stabilize that eigenvalue. In [11], we further show that the algebraic condition for fixed modes that was reported in [1] and had, in our opinion, remained mysterious for 30 years turns out to be a special case of the algebraic mincut-maxflow theorem. This also confirms that LTI controllers in decentralized control systems implicitly communicate via linear network coding. The connection to network coding becomes even more clear when we consider stabilization problems with an explicit communication network [11].

Taking a step back, the general idea of implicit communication (signaling) between decentralized controllers and information flow in decentralized systems has been recognized since Witsenhausen's counterexample [20]. However, in Witsenhausen's counterexample the need for communication between controllers is justified by the suboptimality of linear controllers, *i.e.* if the decentralized controllers want to communicate with each other for efficient control of the system, they would do so using nonlinear controllers for signaling [24], [25], [21]. However, we showed here that even if we restrict controllers to be linear time-invariant, the controllers still can "communicate" via linear network coding. To an extent, this paper does for implicit communication what [26], [27] did vis-a-vis [16], [28] for explicit communication — it finds a way to discuss the issue within a linear framework. In fact, the existence of implicit communication between linear controllers in decentralized systems has been conjectured for a long time [22], [29], [30], [9], [10]. In a sense, we hope that this paper clarifies these discussions.

Finally, even if this paper focuses on bringing a network-coding information-theoretic perspective to decentralized control, we can also think in the reverse direction. For example, the proof of the algebraic mincut-maxflow theorem [12] exploits a system-theoretic perspective on network coding.

References

[1] B. Anderson and D. Clements, "Algebraic characterization of fixed modes in decentralized control," *Automatica*, vol. 17, no. 5, pp. 703–712, Sep. 1981.

[2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, 2006.

[3] S. Avestimehr, S. Diggavi, and D. Tse, "Wireless network information flow: a deterministic approach," *IEEE Transactions on Information Theory*, to appear in 2011.

[4] P. Elias, A. Feinstein, and C. Shannon, "A note on the maximum flow through a network," *IRE Transactions on Information Theory*, vol. 2, no. 4, pp. 117–119, Dec. 1956.

[5] L. R. Ford and D. R. Fulkerson, "Maximal flow through a network," *Canadian Journal of Mathematics*, vol. 8, p. 399404, 1956.

[6] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung, "Network information flow," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.

[7] R. Koetter and M. Medard, "An algebraic approach to network coding," *IEEE ACM Transations On Networking*, vol. 11, pp. 782–795, 2003.

[8] S. H. Wang and E. Davison, "On the stabilization of decentralized control systems," *IEEE Trans. Automat. Contr.*, vol. 18, no. 5, pp. 473–478, Oct. 1973.

[9] S. Yuksel, "Decentralized computation and communication in stabilization of distributed control systems," *Information Theory and Applications Workshop*, 2009.

[10] C. Zhang and G. Dullerud, "Decentalized control with communication bandwidth constraints," *American Control Conference*, Jul. 2007.

[11] S. Y. Park and A. Sahai, "Network coding meets decentralized control: Capacity-stabilizability equivalence," Tech. Rep., 2011. [Online]. Available: eecs.berkeley.edu/∼separk

[12] ——, "Algebraic mincut-maxflow theorem," Submitted to ISIT 2011.

[13] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[14] J. P. M. Schalkwijk and T. Kailath, "A coding scheme for additive noise channels with feedback – I: No bandwidth constraint," *IEEE Trans. Inform. Theory*, vol. 12, no. 2, pp. 172–182, Apr. 1966.

[15] N. Elia, "When Bode meets Shannon: Control-oriented feedback communication schemes," *IEEE Trans. Automat. Contr.*, vol. 49, no. 9, pp. 1477–1488, Sep. 2004.

[16] A. Sahai and S. Mitter, "The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link - part I: Scalar systems," *IEEE Trans. Inform. Theory*, vol. 52, no. 8, pp. 3369–3395, Aug. 2006.

[17] J. Baillieul, "Feedback designs for controlling device arrays with communication channel bandwidth constraints," *In Lecture Notes of the Fourth ARO Workshop on Smart Structures*, 1999.

[18] S. Tatikonda and S. K. Mitter, "Control under communication constraints," *IEEE Trans. Automat. Contr.*, vol. 49, no. 7, pp. 1056–1068, Jul. 2004.

[19] S. K. Mitter and N. J. Newton, "Information flow and entropy production in the Kalman-Bucy filter," *Journal of Statistical Physics*, vol. 118, no. 1, pp. 145–175, Jan. 2005.

[20] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147, Jan. 1968.

[21] P. Grover, A. Sahai, and S. Y. Park, "The finite-dimensional witsenhausen counterexample," *IEEE Trans. Automat. Contr.*, Submitted 2010.

[22] B. Anderson and J. Moore, "Time-varying feedback laws for decentralized control," *IEEE Trans. Automat. Contr.*, vol. 26, no. 5, pp. 1133–1139, Oct. 1981.

[23] C.-T. Chen, *Linear System Theory and Design*, 3rd ed. New York, NY: Oxford University Press, 1999.

[24] H. S. Witsenhausen, "Separation of estimation and control for discrete time systems," *Proceedings of the IEEE*, vol. 59, no. 11, pp. 1557–1566, Nov. 1971.

[25] Y.-C. Ho, M. P. Kastner, and E. Wong, "Teams, signaling, and information theory," *IEEE Trans. Automat. Contr.*, vol. 23, no. 2, Apr. 1978.

[26] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. Jordan, and S. Sastry, "Kalman filtering with intermittent observations," *IEEE Trans. Automat. Contr.*, vol. 49, no. 9, pp. 1453–1464, Sep. 2004.

[27] N. Elia, "Remote stabilization over fading channels," *Systems and Control Letters*, vol. 54, no. 3, pp. 239–249, Mar. 2005.

[28] A. Sahai, "Anytime information theory," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2001.

[29] J. P. Corfmat and A. S. Morse, "Decentralized control of linear multivariable systems," *Automatica*, vol. 12, Sep. 1976.

[30] B. Anderson, "Transfer function matrix description of decentralized fixed modes," *IEEE Trans. Automat. Contr.*, vol. 6, Dec. 1982.