Sequential team form and its simplification using graphical models

Aditya Mahajan and Sekhar Tatikonda
Yale University

Allerton, September 30, 2009

Outline

- Sequential team
- Team form
- Simplification of team form
- Representation of team form as a graphical model
- Automated simplification of the graphical model

Multi-agent decentralized systems: a classification

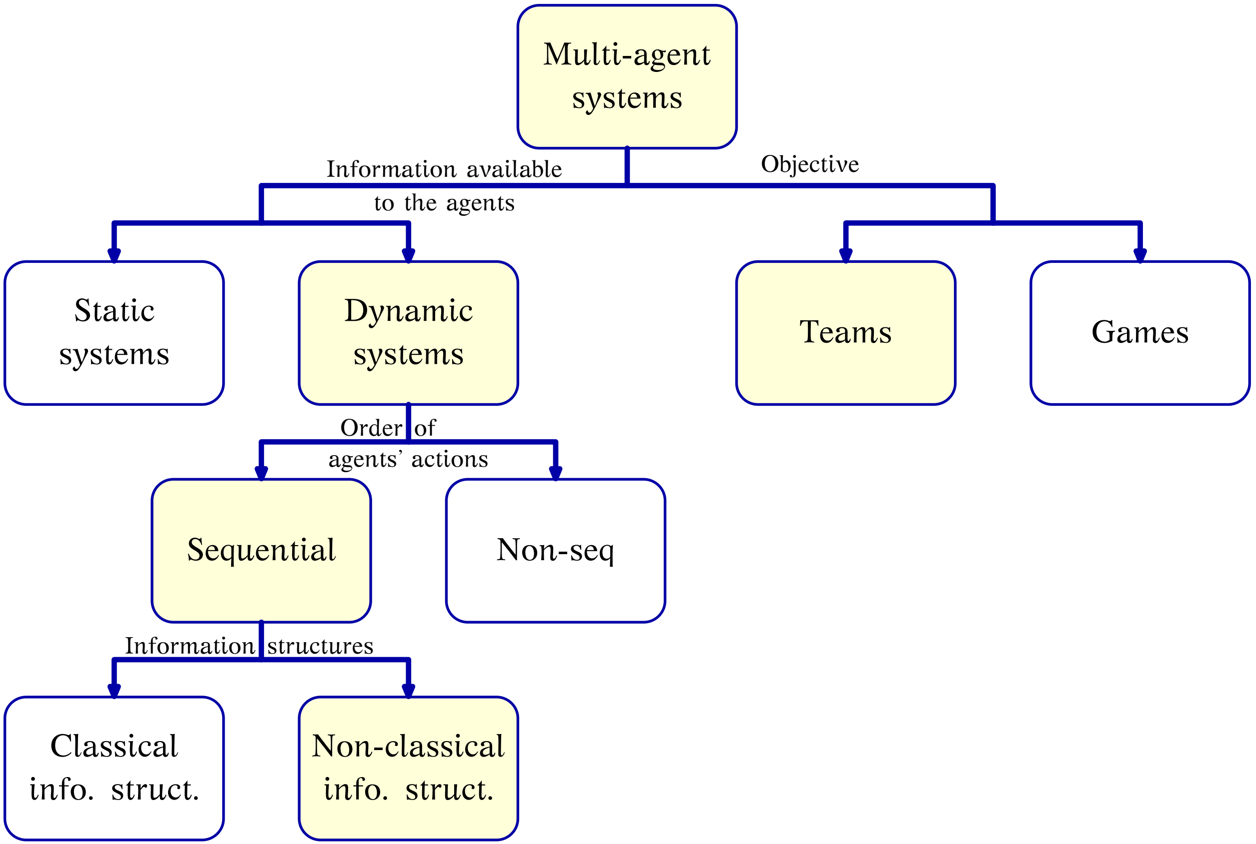

Multi-agent decentralized systems: a classification

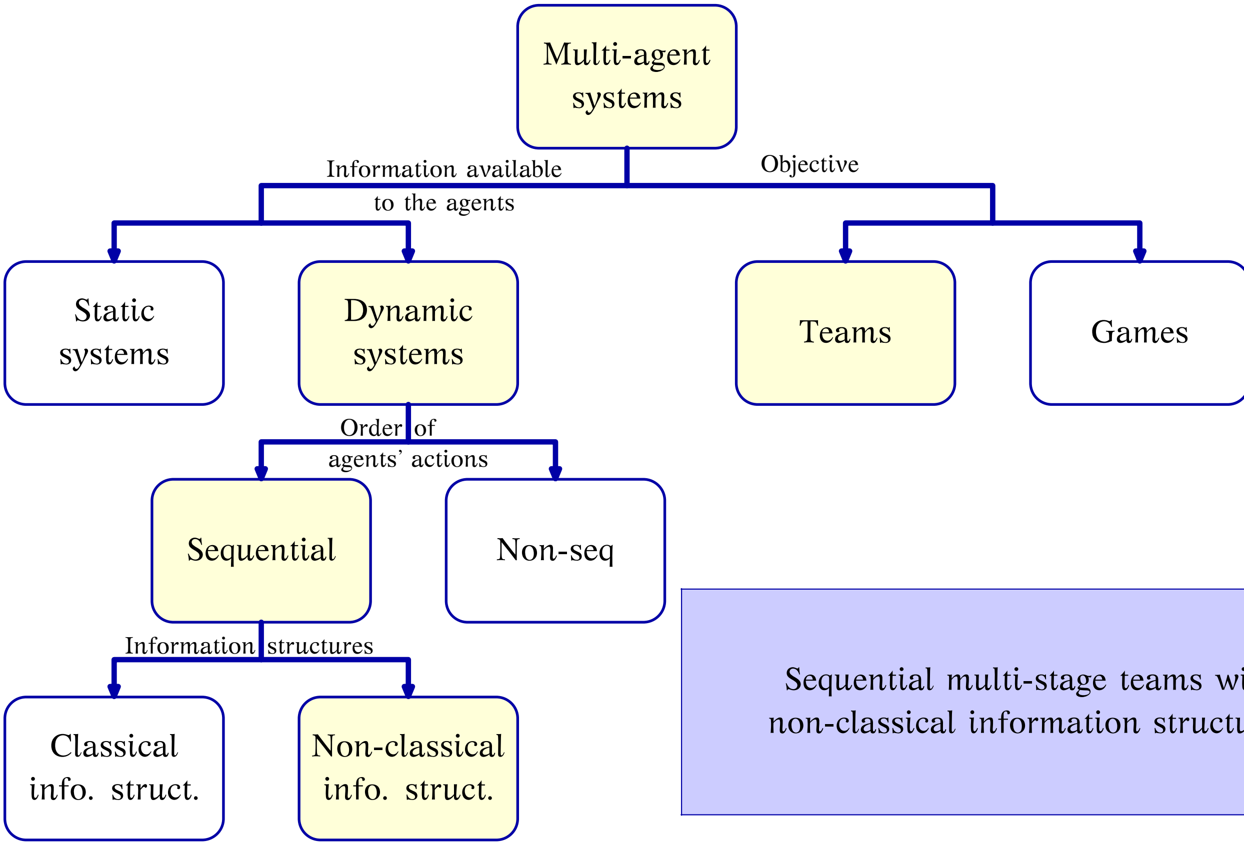

Sequential multi-stage teams with non-classical information structures

Notation

For a set M

- Variables: $X_M = (X_m : m \in M)$.
- Spaces: $\mathcal{X}_M = \prod_{m \in M} \mathcal{X}_m$
- σ-algebras: $\mathfrak{F}_M = \bigotimes_{m \in M} \mathfrak{F}_m$

Model for a sequential team

- A collection of n system variables, $(X_k, k \in N)$ where $N = \{1, \dots, n\}$
- A collection $\{(\mathcal{X}_k, \mathfrak{F}_k)\}_{k \in N}$ of measurable spaces.
- A collection $\{I_k\}_{k \in N}$ of information sets such that $I_k \subset \{1, \dots, k-1\}$.
- A set $A \subset N$ of controllers/agents.
- A set $R \subset N$ of rewards.
- The variables $X_{N \setminus A}$ are chosen by nature according to stochastic kernels $\{p_k\}_{k \in N \setminus A}$ where p_k is a stochastic kernel from $(\mathcal{X}_{I_k}, \mathfrak{F}_{I_k})$ to $(\mathcal{X}_k, \mathfrak{F}_k)$.

Objective

- Choose a strategy $\{g_k\}_{k \in A}$ such that the control law g_k is a measurable function from $(\mathcal{X}_{I_k}, \mathfrak{F}_{I_k})$ to $(\mathcal{X}_k, \mathfrak{F}_k)$.

- Joint measure induced by strategy $\{g_k\}_{k \in N}$

$$P(dX_N) = \bigotimes_{k \in N \setminus A} p_k(dX_k | X_{I_k}) \bigotimes_{k \in A} \delta_{g_k(X_{I_k})}(dX_k)$$

- Choose a strategy to maximize

$$E^{g_A}\left\{ \sum_{i \in R} X_i \right\}$$

This maximum reward is called the value of the team

Generality of the model

This model is a generalization of the model presented in

Hans S. Witsenhausen, Equivalent stochastic control problems, Math. Cont. Sig. Sys.-88

which in turn in equivalent to the intrinsic model presented in

Hans S. Witsenhausen, On information structures, feedback and causality, SICON-71

which is as general as it gets.

Team form

A (sequential) team form is the team problem where the measurable spaces $\{(\mathcal{X}_k, \mathfrak{F}_k)\}_{k \in N}$ and the stochastic kernels $\{p_k\}_{k \in N \setminus A}$ are not pre-specified.

$\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$: system variables, control variables, reward variables, and the information sets are specified.

Equivalence of team forms

Two team forms $\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$ and $\mathcal{T}' = (N', A', R', \{I'_k\}_{k \in N'})$ are equivalent if the following conditions hold:

1. $N = N'$, $A = A'$, and $R = R'$;
2. for all $k \in N \setminus A$, we have $I_k = I'_k$;
3. for any choice of measurable spaces $\{(\mathcal{X}_k, \mathfrak{F}_k)\}_{k \in N}$ and stochastic kernels $\{p_k\}_{k \in N \setminus A}$, the values of the teams corresponding to $\mathcal{T}$ and $\mathcal{T}'$ are the same.

The first two conditions can be verified trivially. There is no easy way to check the last condition.

Simplification of team forms

A team form $\mathcal{T}' = (N', A', R', \{I'_k\}_{k \in N'})$ is a simplification of a team form $\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$ if

$$\mathcal{T}' \text{ is equivalent to } \mathcal{T}$$

and

$$\sum_{k \in A} |I'_k| < \sum_{k \in A} |I_k| \, .$$

$\mathcal{T}'$ is a strict simplification of $\mathcal{T}$ if $\mathcal{T}'$ is equivalent to $\mathcal{T}$, $|I'_k| \leqslant |I_k|$ for $k \in N$, and at least one of these inequalities is strict.

Given a team form, can we simplify it?

Asking for simplification of a team form is same as asking for structural properties that do not depend on the nature of the process (discrete or continuous values), the specific form of probability measure (Gaussian, uniform, binomial , etc.) and the specific properties of cost function (convex, monotone, etc.)

Some Preliminaries

Partial Orders

A strict partial order $\prec$ on a set S is a binary relation that is transitive, irreflexive, and asymmetric. i.e., for a, b, c in S, we have

1. if $a \prec b$ and $b \prec c$, then $a \prec c$ (transitive)
2. $a \not\prec a$ (irreflexive)
3. if $a \prec b$ then $b \not\prec a$ (asymmetric)

The reflexive closure $\preceq$ of a partial order $\prec$ is given by

$$a \preceq b \text{ if and only if } a \prec b \text{ or } a = b$$

Partial Order

Let A be a subset of a partially ordered set $(S, \prec)$. Then, the lower set of A, denoted by $\overleftarrow{A}$ is defined as

$$\overleftarrow{A} \coloneqq \{b \in S : b \preceq a \text{ for some } a \in A\}.$$

By duality, the upper set of A, denoted by $\overrightarrow{A}$ is defined as

$$\overrightarrow{A} \coloneqq \{b \in S : a \preceq b \text{ for some } a \in A\}.$$

Sequential teams and partial orders

- Hans S. Witsenhausen, On information structures, feedback and causality, SICON-71

- Hans S. Witsenhausen, The intrinsic model for discrete stochastic control: Some open problems, LNEMS-75

> A team problem is sequential if and only if
> there is a partial order between the agents

Partial orders can be
represented by directed graphs

So, sequential teams can be
represented as directed graphs

Representing teams using directed graphs

Hans S. Witsenhausen, Separation of estimation and control for discrete time systems, Proc. IEEE-71.

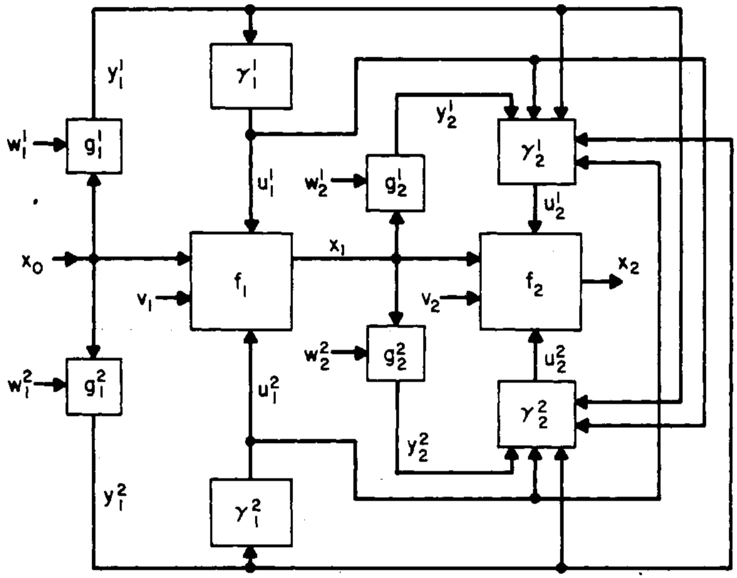

Fig. 1.

Representing teams using directed graphs

Yu-Chi Ho and K'ai-Ching Chu, Team Decision Theory and Information Structures in Optimal Control Problems—Part I, TAC-72.

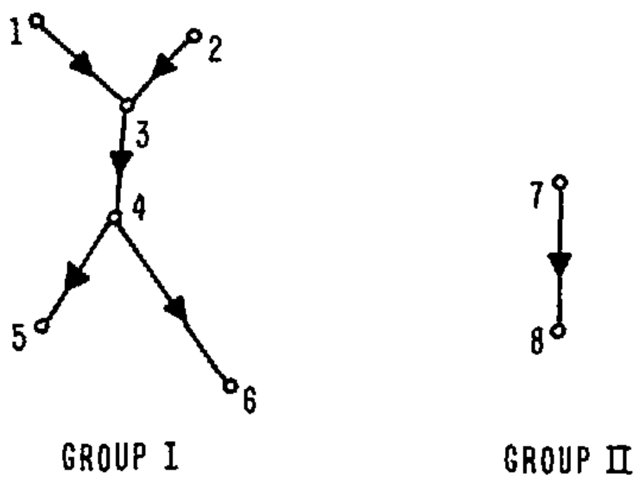

Fig. 3.

Representing teams using directed graphs

Tseneo Yoshikawa, Decomposition of Dynamic Team Decision Problems, TAC-78.

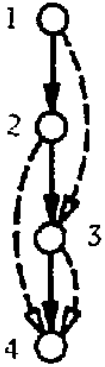

Fig. 1. Precedence diagram.

Representing teams using directed graphs

Steffen L. Lauritzen and Dennis Nilsson, Representing and Solving Decision Problems with Limited Information, Management Science-2001.

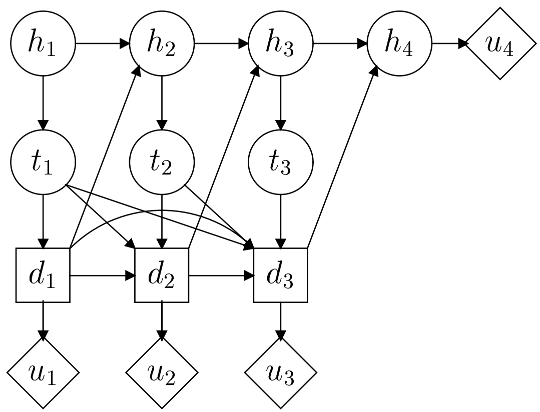

None of these fit our requirements perfectly. So, we use DAFG (Directed Acyclic Factor Graphs)

A graphical model for sequential team forms

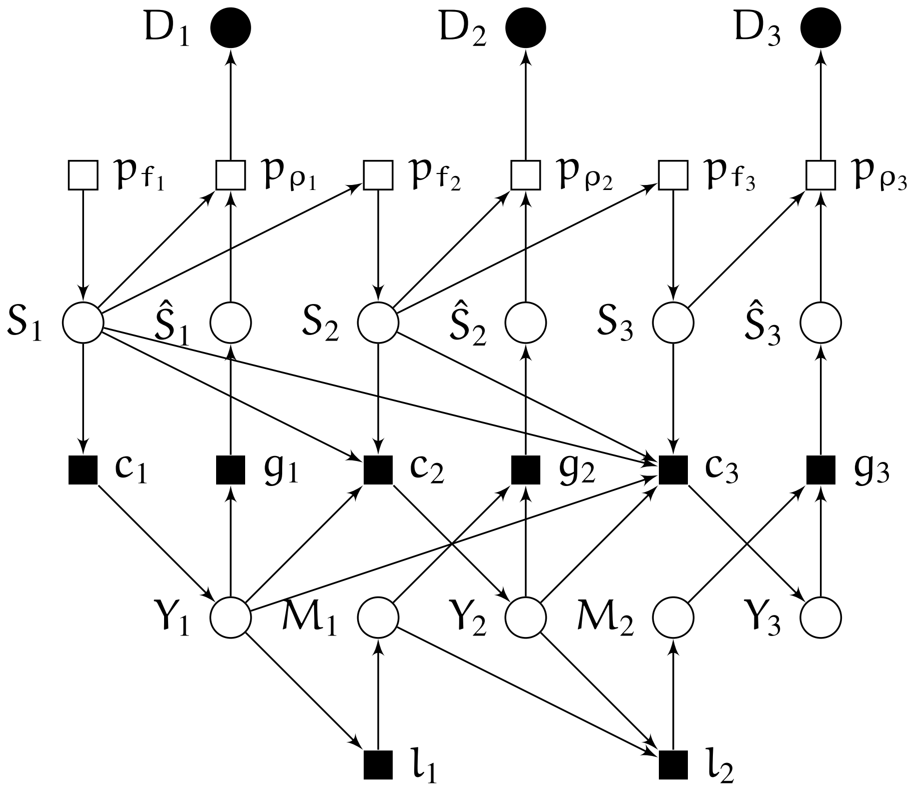

A graphical model for sequential team forms

Directed Acyclic Factor Graph $\mathcal{G} = (V, F, E)$ for $\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$

$$V = N \times \{0\}, \quad F = N \times \{1\}$$

$$E = \{(k^1, k^0) : k \in N\} \cup \{(i^0, k^1) : k \in N, i \in I_k\}$$

- Vertices

  Variable Node $k^0 \equiv$ system variable X_k

  Factor node $k^1 \equiv$ stochastic kernel p_k or control law g_k.

- Edges

  (k^1, k^0) for each $k \in N$

  (i^0, k^1) for each $k \in N$ and $i \in I_k$

An Example: Real-time communication

Hans S. Witsenhausen, On the structure of real-time source coders, BSTJ-79

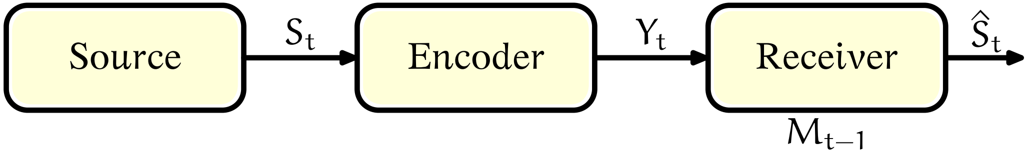

First order Markov source $\{S_t, t = 1, \dots, T\}$.

Real-Time Encoder: $Y_t = c_t(S^t, Y^{t-1})$

Real-Time Finite Memory Decoder: $\begin{array}{|l} \hat{S}_t = g_t(Y_t, M_{t-1}) \\ M_t = l_t(Y_t, M_{t-1}) \end{array}$

Instantaneous distortion $\rho(S_t, \hat{S}_t)$

Objective: minimize $E\left\{ \sum_{t=1}^{T} \rho(S_t, \hat{S}_t) \right\}$

An Example: Real-time communication

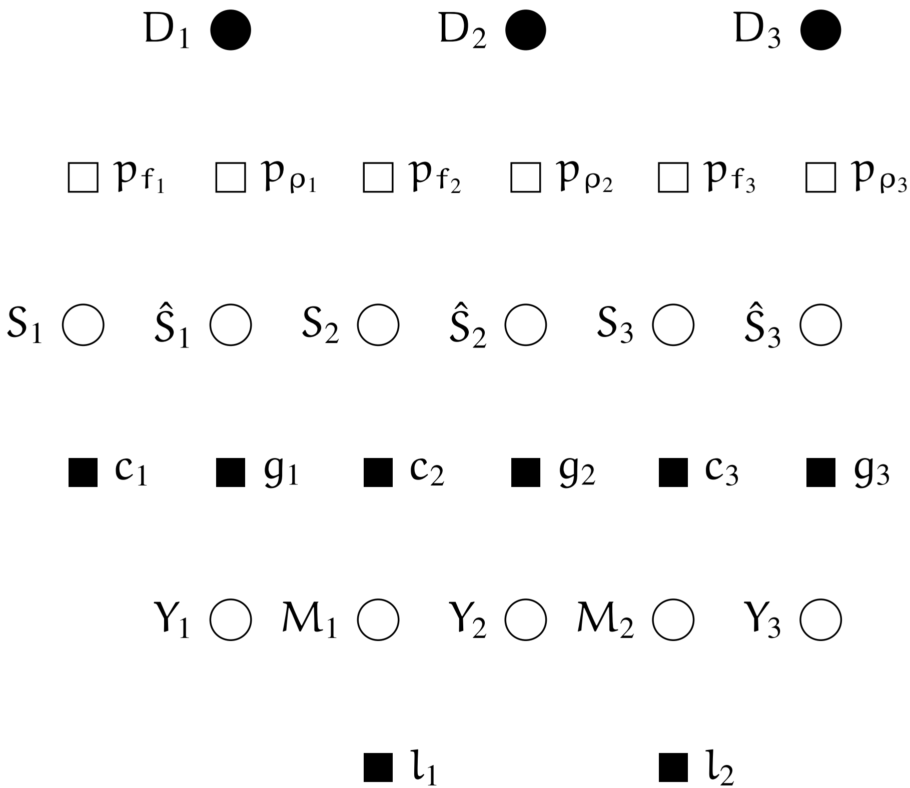

An Example: Real-time communication

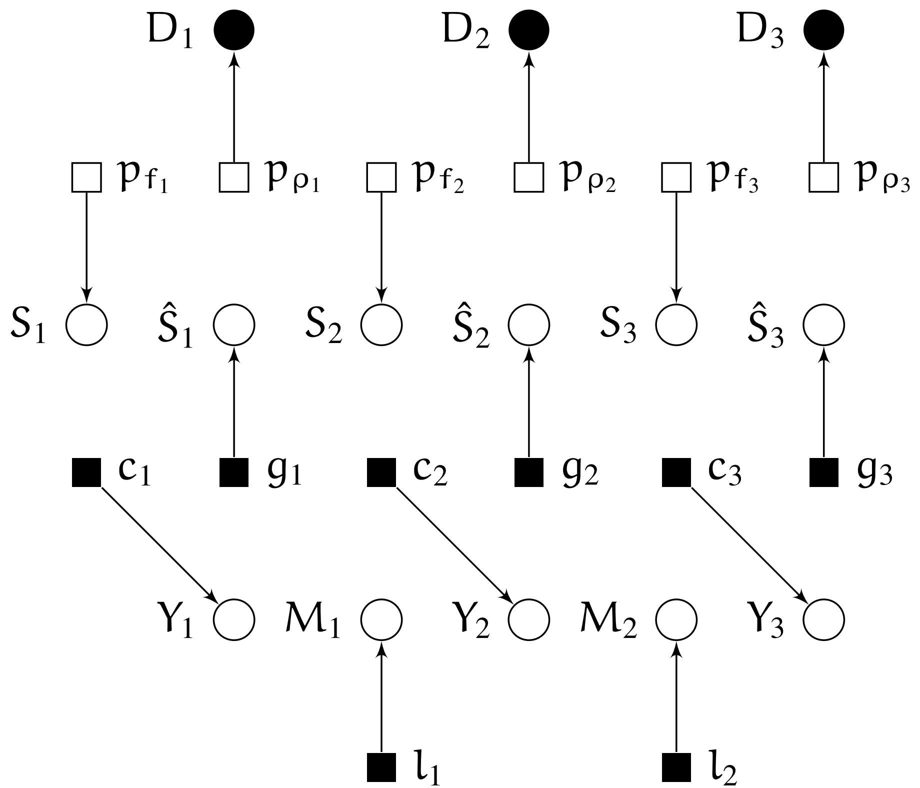

An Example: Real-time communication

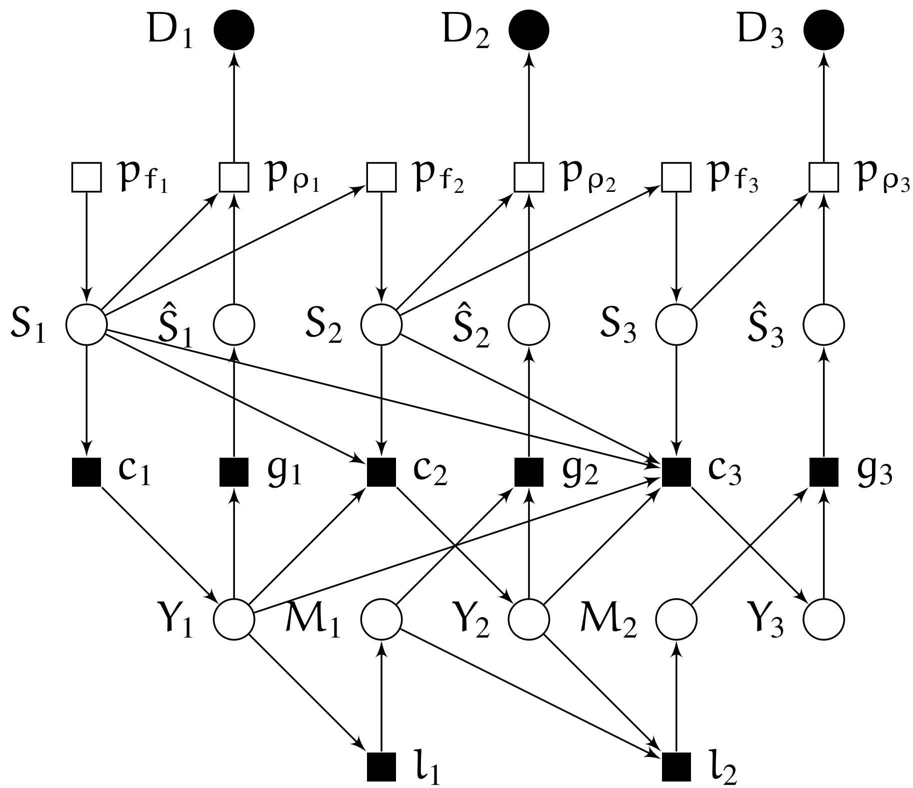

Checking conditional independence

Dan Geiger, Thomas Verma, and Judea Pearl, Identifying independence in Bayesian networks, Networks-90.

> Conditional independence can be efficiently checked on a directed graph.

Given a DAFG $\mathcal{G} = (V, F, E, D)$ and sets $A, B, C \subset V$, X_A is irrelevant to X_B given X_C if X_A is independent to X_B given X_C for all joint measures $P(dX_V)$ that recursively factorize according to $\mathcal{G}$.

Data irrelevant to X_A given X_C is

$$R^-_{\mathcal{G}}(X_A|X_C) = \{k \in C : X_k \text{ is irrelevant to } X_A \text{ given } X_C\}$$

Back to simplification of team forms

Completion of a team

A team form $\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$ is complete if for $k, l \in A$, $k \neq l$, such that $I_k \subset I_l$ we have $X_k \in I_l$. (If l knows the data available to k, then l also knows the action taken by k).

If a team is not complete, it can be completed by sequentially adding "missing links"

Depending on the order in which we proceed, we can end up with different completions. However,

all completions of a team form are equivalent.

Completion of a team form

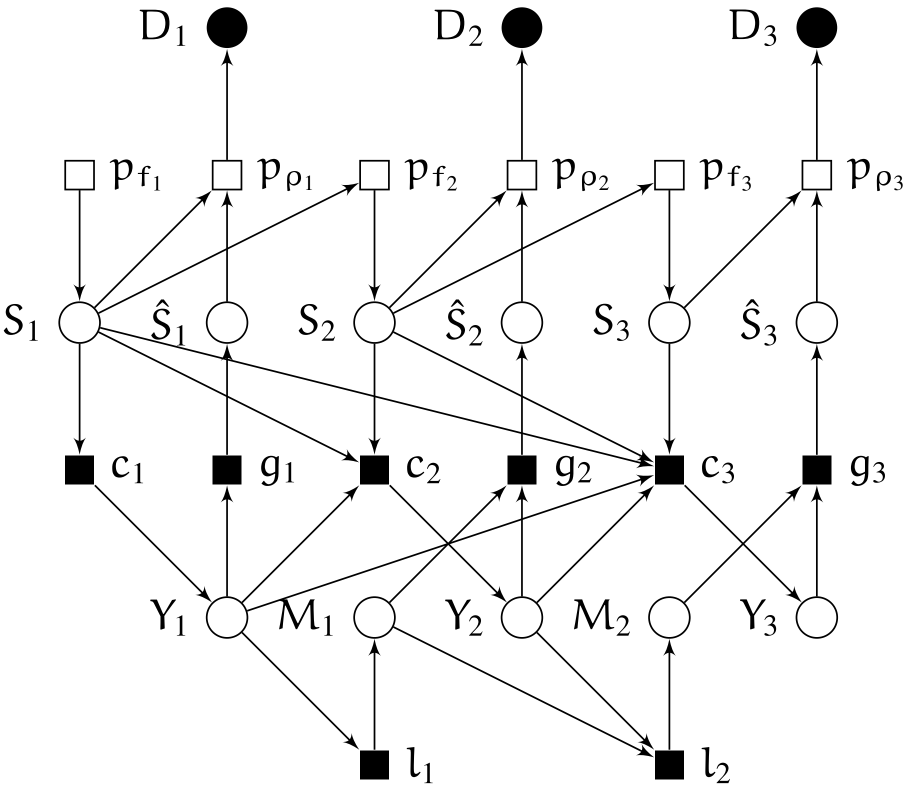

Completion of a team

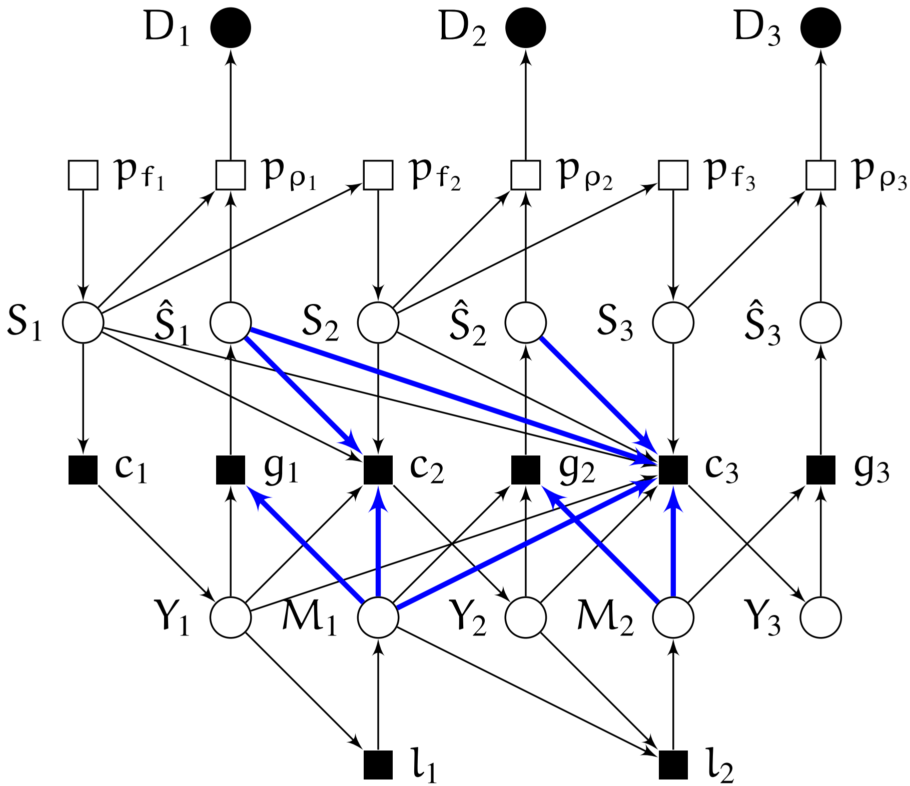

Simplification of team forms

Step 1: Complete the team form.

(Note: All completions of a team form are equivalent to the original)

Removing irrelevant nodes

Recall Given a DAFG $\mathcal{G} = (V, F, E, D)$ and sets $A, B, C \subset V$, X_A is irrelevant to X_B given X_C if X_A is independent to X_B given X_C for all joint measures $P(dX_V)$ that recursively factorize according to $\mathcal{G}$ and

$$R^{-}_{\mathcal{G}}(X_A | X_C) = \{k \in C : X_k \text{ is irrelevant to } X_A \text{ given } X_C\}$$

For any $k \in A$ in a team form $\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$, replacing X_{I_k} by $X_{I_k} \setminus \left(R^{-}_{\mathcal{G}}(X_R \cap \overrightarrow{X_k} \mid X_{I_k}, X_k) \setminus X_k\right)$ does not change the value of the team.

Removing irrelevant nodes

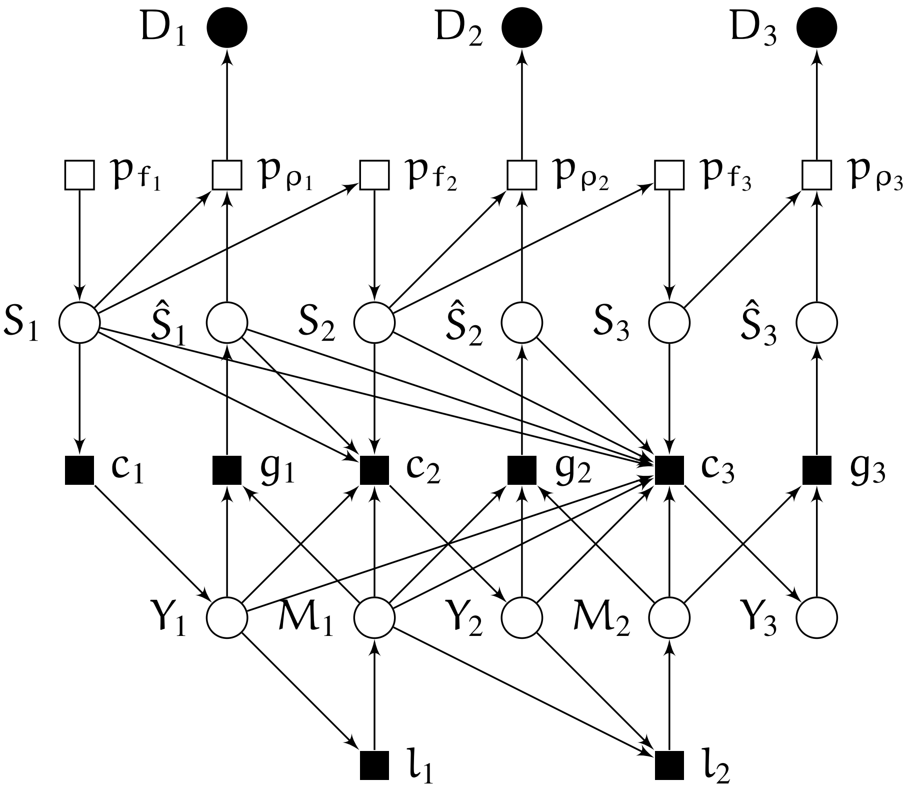

Removing irrelevant nodes

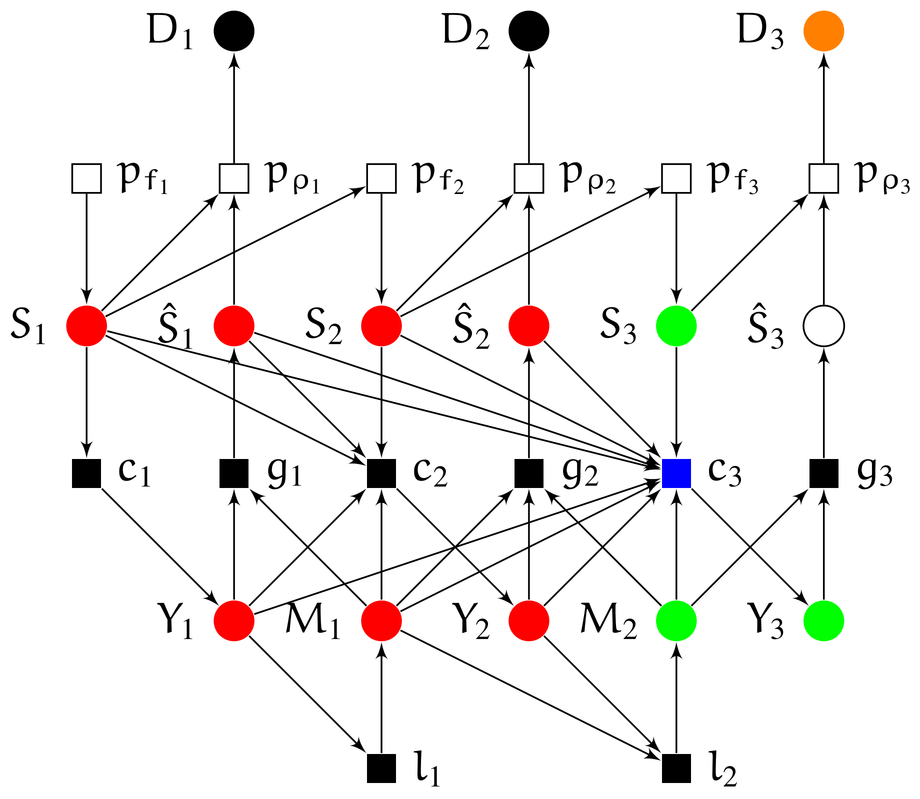

Removing irrelevant nodes

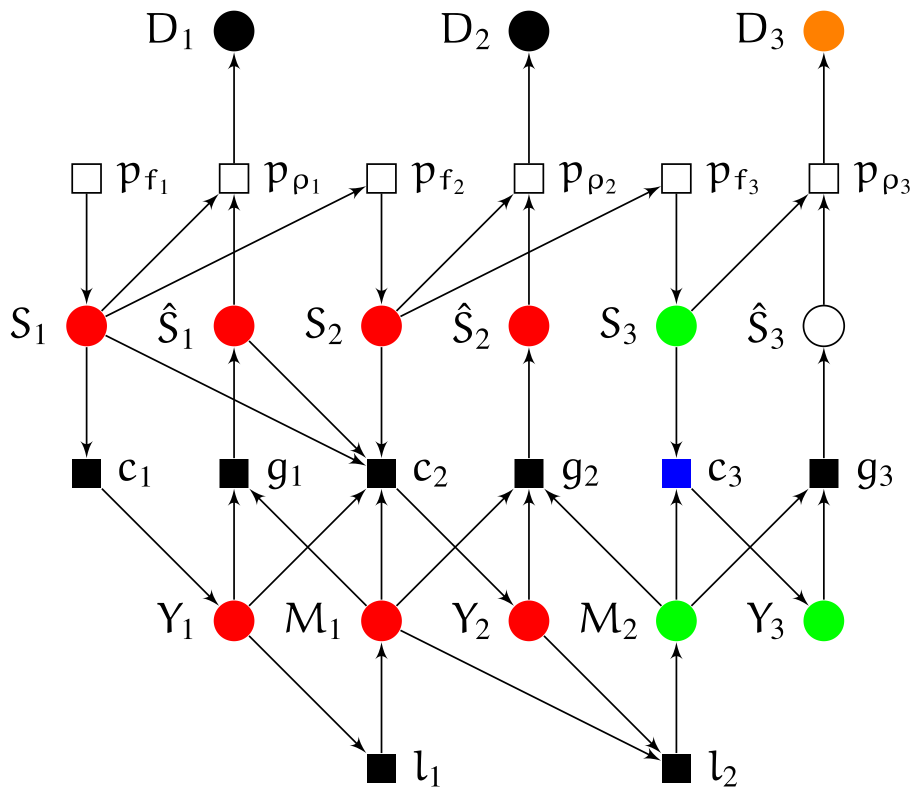

Removing irrelevant nodes

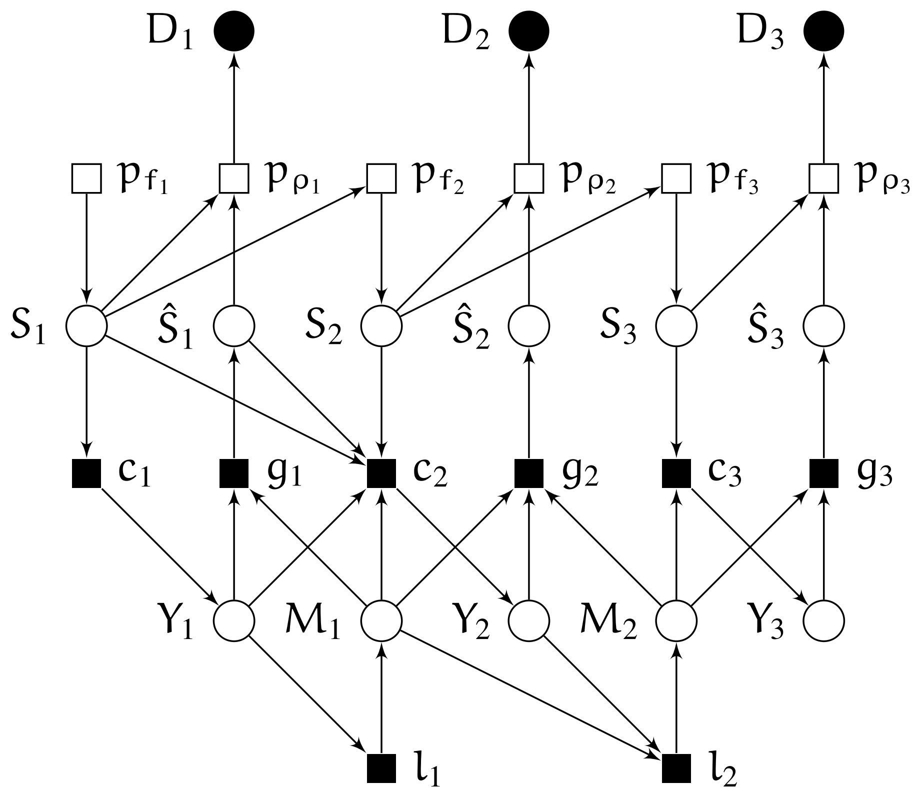

Removing irrelevant nodes

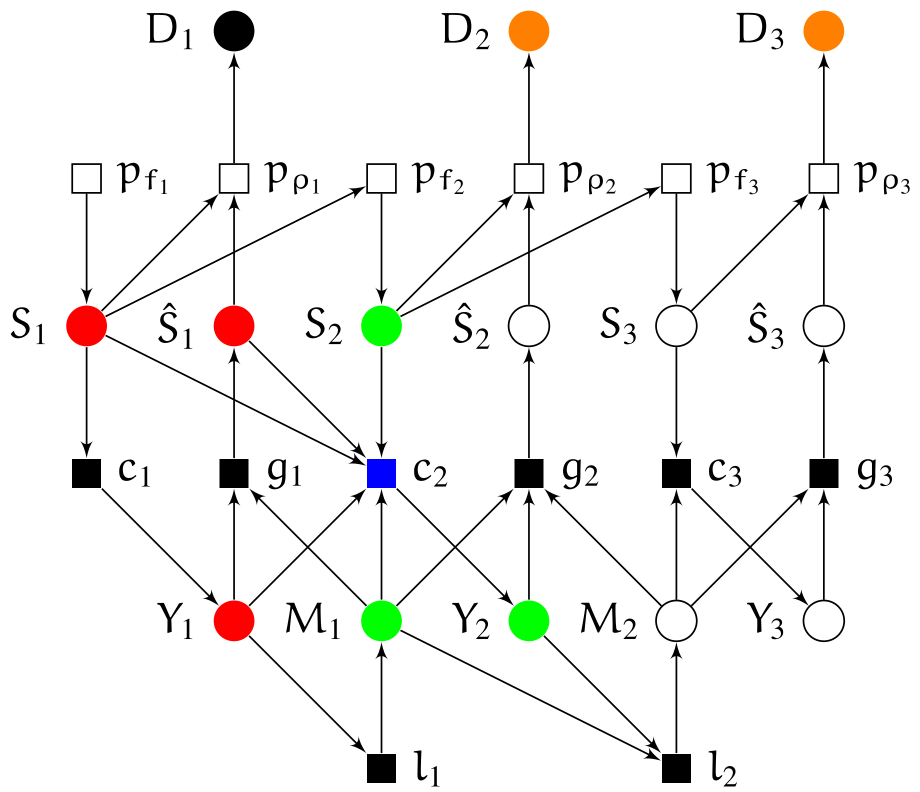

Removing irrelevant nodes

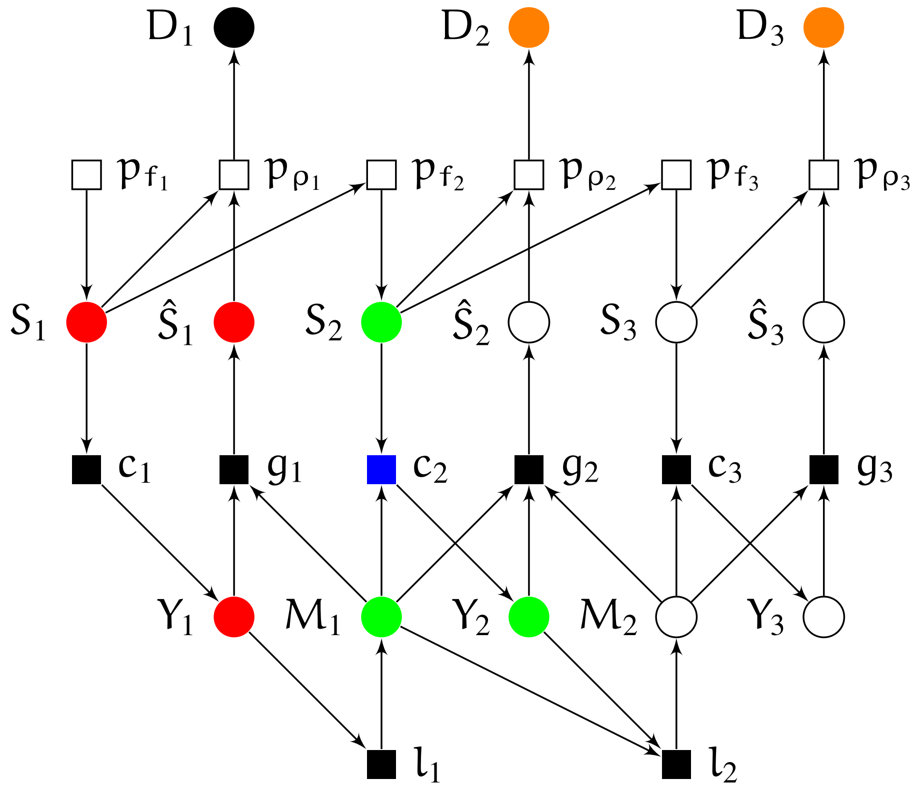

Removing irrelevant nodes

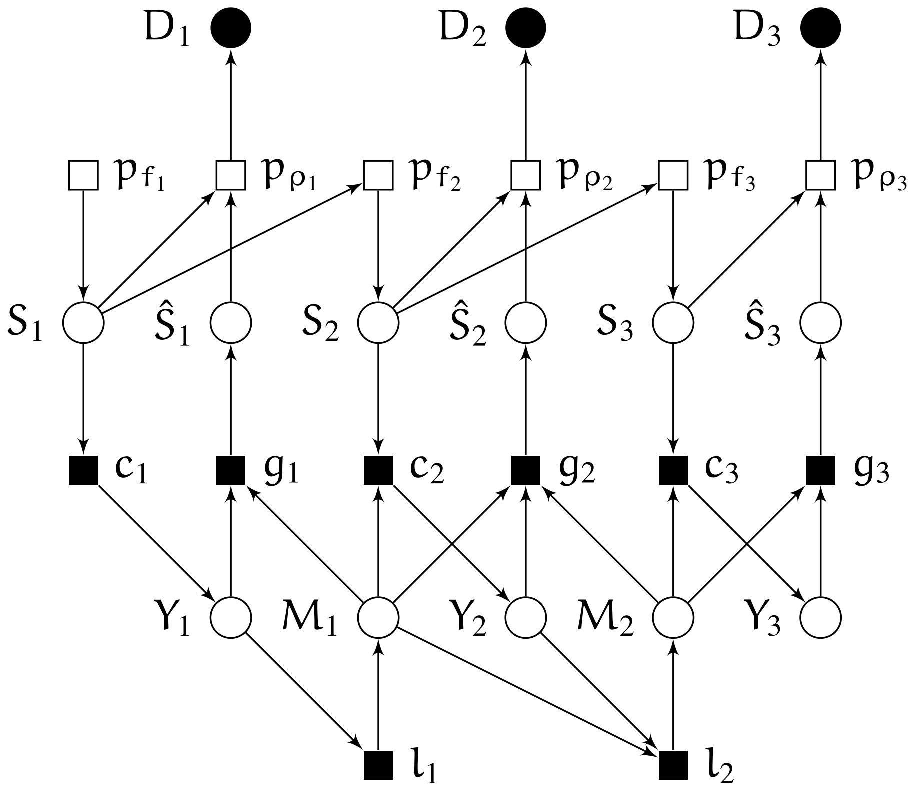

Removing irrelevant nodes

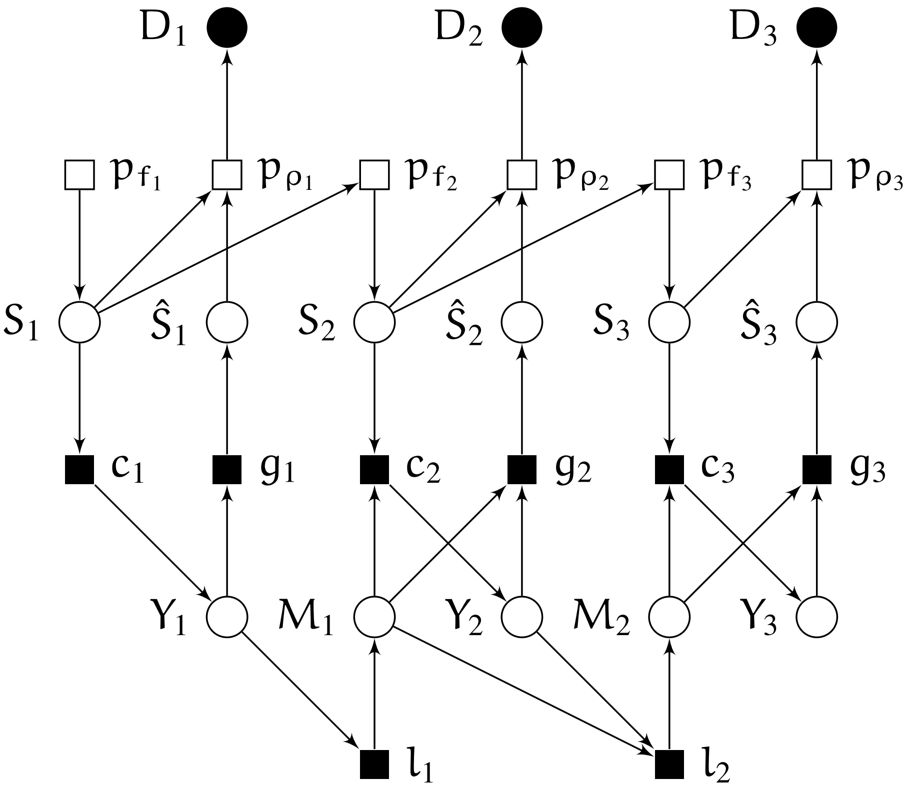

Removing irrelevant nodes

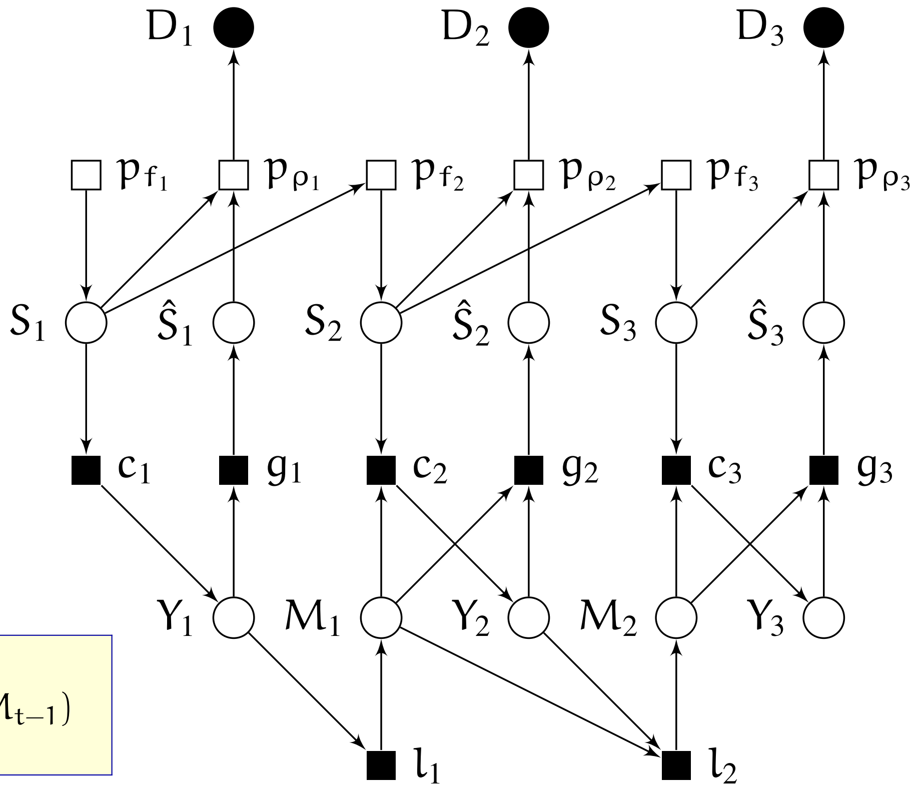

$$Y_t = c_t(S_t, M_{t-1})$$

Simplification of team forms

Step 1: Complete the team form.

(Note: All completions of a team form are equivalent to the original)

Step 2: At control factor node k, remove incoming edges from nodes irrelevant to $X_R \cap \overrightarrow{X}_k$ given (X_{I_k}, X_k)

(Note: The resultant team form is equivalent to the original)

Does not always work

Another Example: Shared randomness

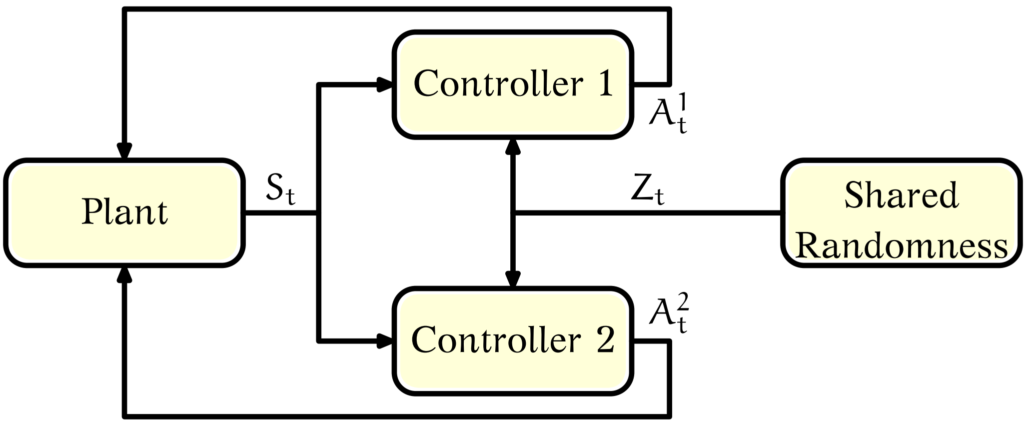

Plant: $S_{t+1} = f_t(S_t, A^1_t, A^2_t, W_t)$

Shared Randomness: $\{Z_t, t = 1, \ldots, T\}$

independent of plant disturbance and observation noise.

Control Station 1: $A^1_t = g^1_t(S^t, A^{1,t-1}, Z^t)$ Control Station 2: $A^2_t = g^2_t(S^t, A^{2,t-1}, Z^t)$

Instantaneous cost: $\rho_t(S_t, A^1_t, A^2_t)$

Another Example: Shared randomness

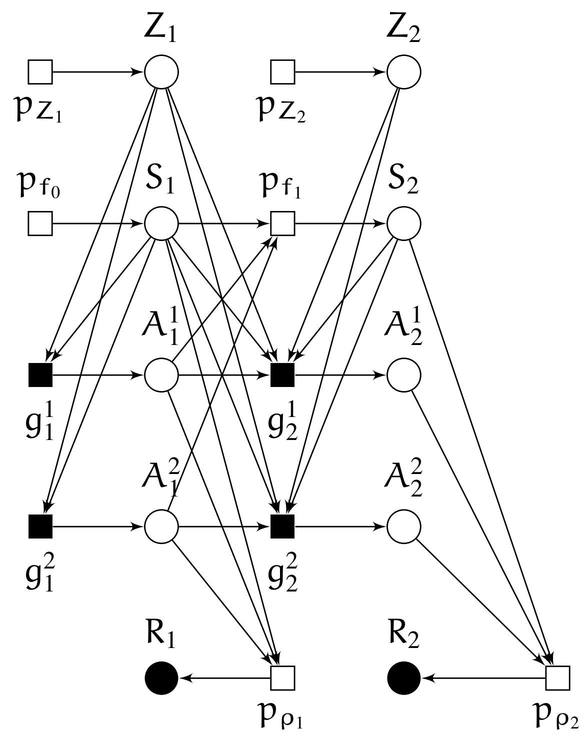

Another Example: Shared randomness (Step 1)

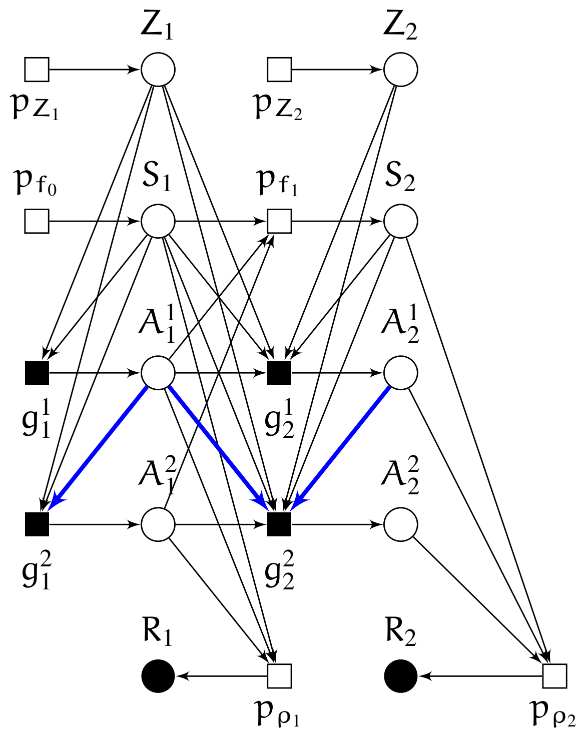

Coordinator for a subset of agents

For $a, b \in A$, consider a coordinator that observes $X_C \coloneqq X_{I_a} \cap X_{I_b}$ and chooses partial functions $\hat{g}_a : X_{I_a \setminus C} \to X_a$ and $\hat{g}_b : X_{I_b \setminus C} \to X_b$.

Agent a and b simply carry out the computations prescribed by $\hat{g}_a$ and $\hat{g}_b$

Remove irrelevant incoming edges at the coordinator!

Equivalently, at agents a and b, remove edges from nodes that are irrelevant to $X_R \cap \overrightarrow{X}_{\{a,b\}}$ given $(X_C, X_{\{a,b\}})$.

Coordinator for a subset of agents

For any $B \subset A$ in a team form $\mathcal{T} = (N, A, R, \{I_k\}_{k \in N})$

and any $b \in B$, let $X_C = \bigcap_{b \in B} X_{I_b}$. Then, replacing

X_{I_b} by $X_{I_b} \setminus \left(R^{-}_{\mathcal{G}}(X_R \cap \overrightarrow{X}_B \mid X_C, X_B) \setminus X_B\right)$

does not change the value of the team

Simplification of team forms

Step 1: Complete the team form.

(Note: All completions of a team form are equivalent to the original)

Step 2: At control factor node k, remove incoming edges from nodes irrelevant to $X_R \cap \overrightarrow{X}_k$ given (X_{I_k}, X_k)

(Note: The resultant team form is equivalent to the original)

Step 3: At all nodes of any subset B of A, remove incoming edges from nodes irrelevant to $X_R \cap \overrightarrow{X}_B$ given $(\bigcup_{b \in B} X_{I_b}, X_B)$.

(Note: The resultant team form is equivalent to the original. Furthermore, this computation can be carried out efficiently on a lattice of shared information.)

Another Example: Shared randomness (Step 3)

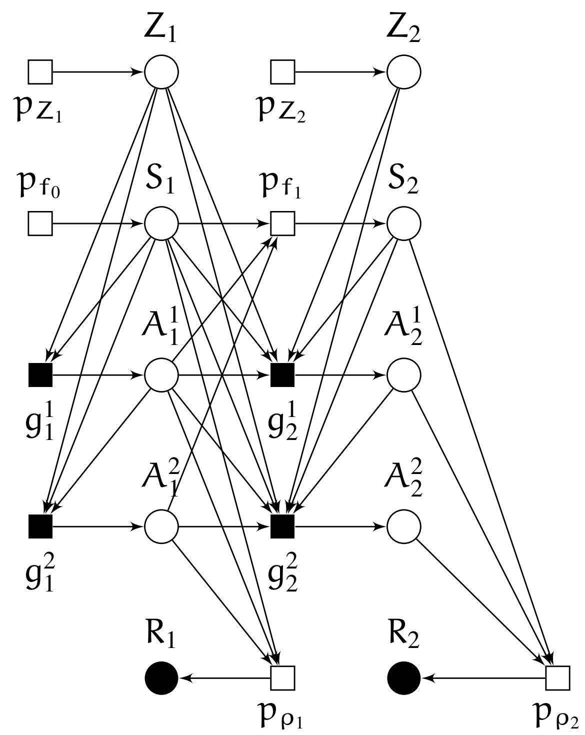

Another Example: Shared randomness (Step 3)

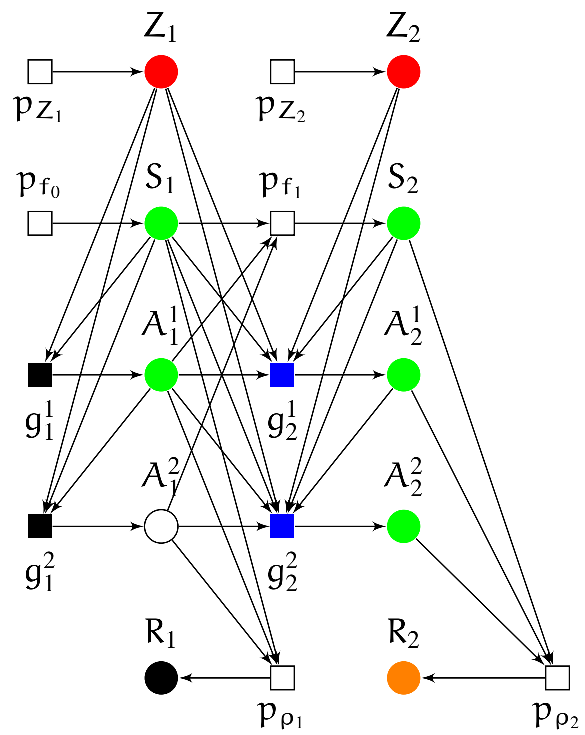

Another Example: Shared randomness (Step 3)

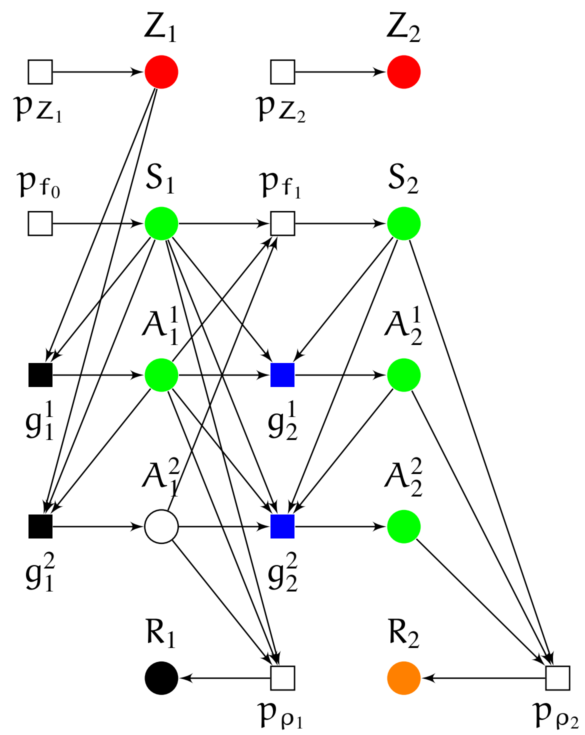

Another Example: Shared randomness (Step 3)

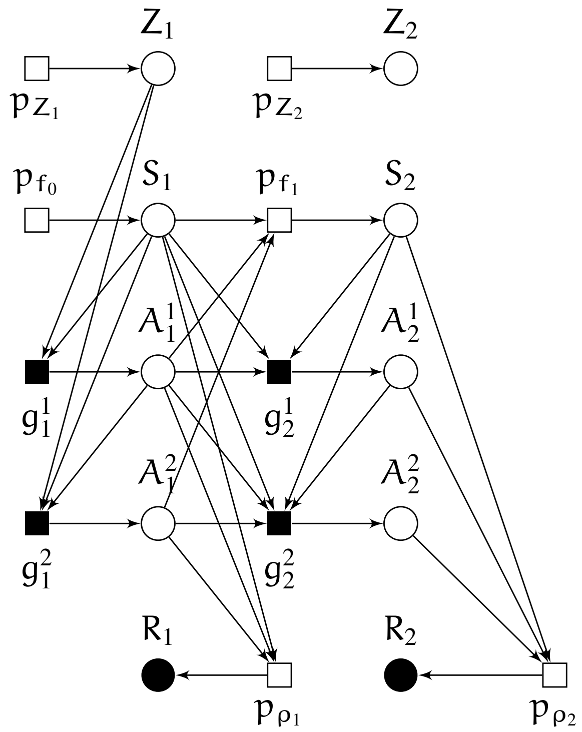

Another Example: Shared randomness (Step 3)

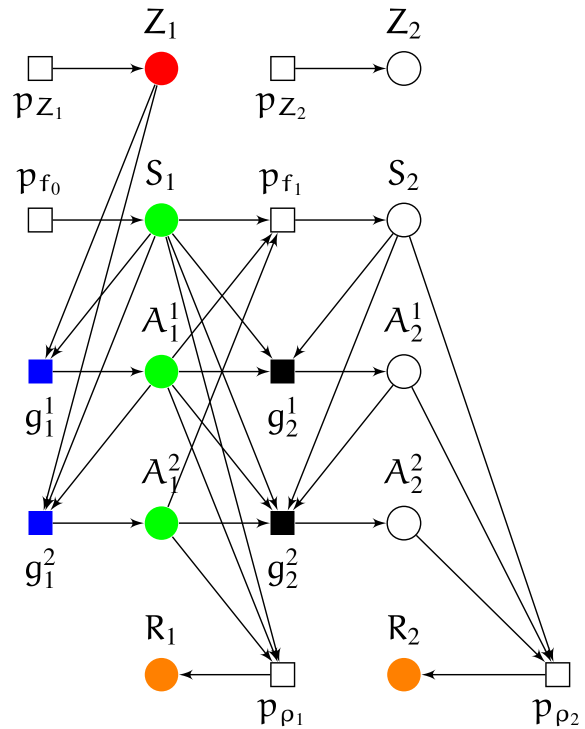

Another Example: Shared randomness (Step 3)

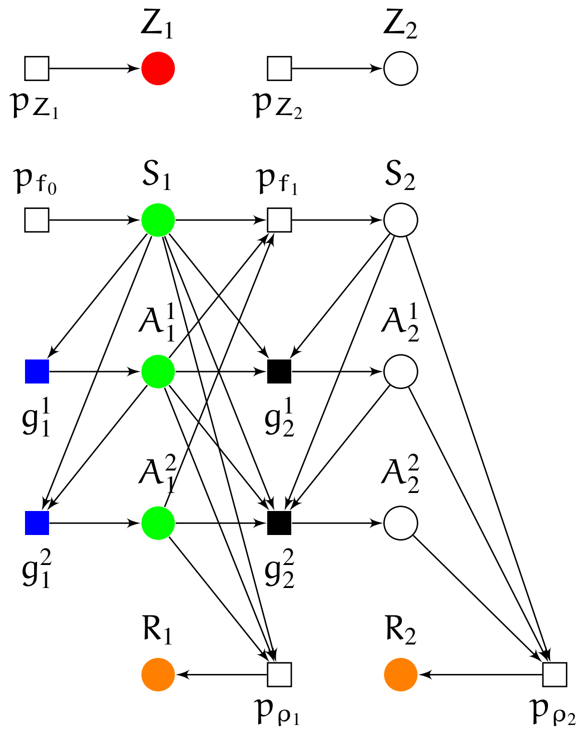

Another Example: Shared randomness (Step 3)

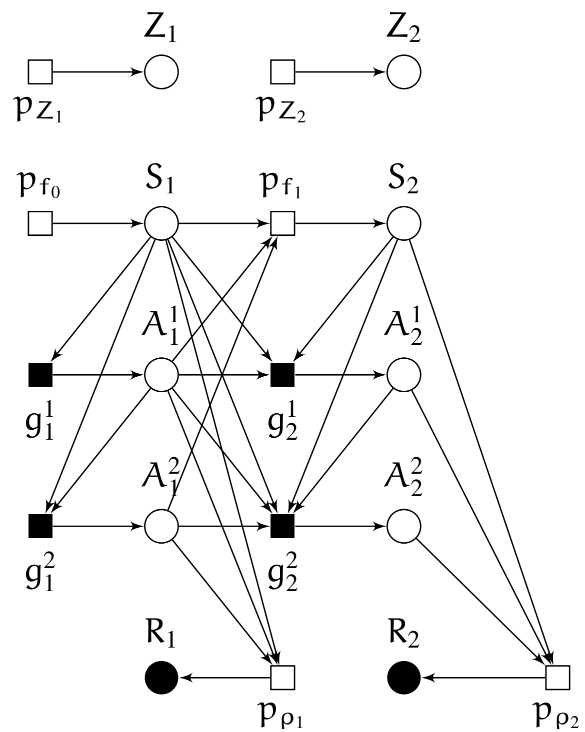

Another Example: Shared randomness (Step 2)

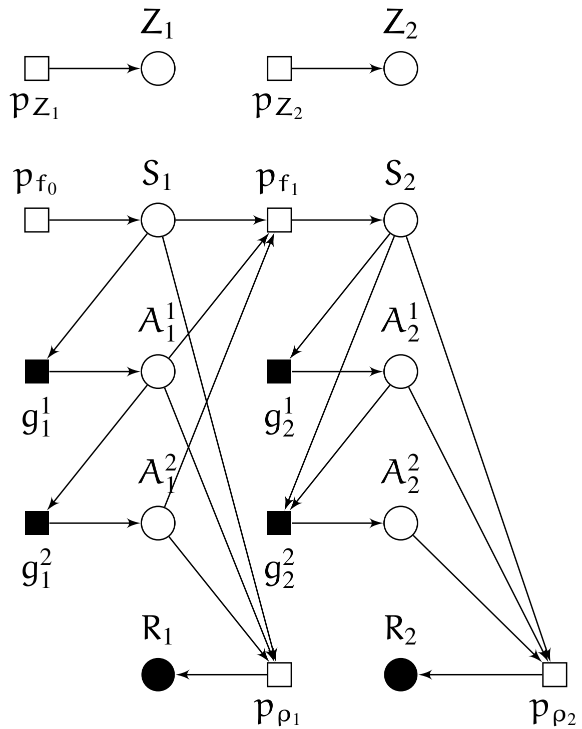

Another Example: Shared randomness (Step 1)

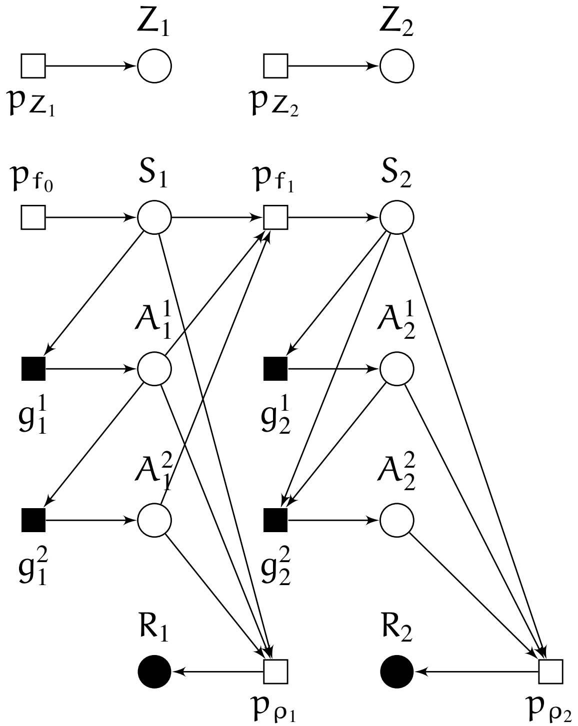

$$A_t^1 = g_t^1(S_t)$$

$$A_t^2 = g_t^2(S_t, A_t^1)$$

Simplification of team forms

Step 1: Complete the team form.

(Note: All completions of a team form are equivalent to the original)

Step 2: At control factor node k, remove incoming edges from nodes irrelevant to $X_R \cap \overrightarrow{X}_k$ given (X_{I_k}, X_k)

(Note: The resultant team form is equivalent to the original)

Step 3: At all nodes of any subset B of A, remove incoming edges from nodes irrelevant to $X_R \cap \overrightarrow{X}_B$ given $(\bigcup_{b \in B} X_{I_b}, X_B)$.

(Note: The resultant team form is equivalent to the original. Furthermore, this computation can be carried out efficiently on a lattice of shared information.)

Conclusion

Team form for sequential teams, equivalence and simplification of team forms.

Representing a team form as a DAFG

Carrying out the simplification of the team form on the DAFG. This process can be automated.

Future Directions

Sequential decomposition of a team form on a DAFG (The sequential decomposition of Witsenhausen's standard form can be carried out efficiently on a DAFG).

Adding belief states / information states (need to study conditional independence properties and define an appropriate notion of simplification)

Thank you