Submodel Decomposition Bound for Influence Diagrams

Junkyu Lee
junkyul@uci.edu
University of California, Irvine
IBM Research AI

Radu Marinescu
radu.marinescu@ie.ibm.com
IBM Research Europe

Rina Dechter
dechter@ics.uci.edu
University of California, Irvine

35th AAAI Conference on Artificial Intelligence

Outline

- Backgrounds
  - Influence Diagrams (IDs) and Limited Memory IDs (LIMIDs)
  - Decomposition of IDs and LIMIDs
  - Bounding Schemes for Maximum Expected Utility (MEU)
- Submodel Decomposition for IDs and LIMIDs
  - Motivation and Contributions
  - A Submodel-Tree Clustering Scheme
  - A Bounding Scheme over Submodel-Tree
- Experiments and Case Study
  - Upper bounds in IDs, LIMIDs, and MDP/POMDP planning domain
- Conclusion and Future Directions

Outline

- Backgrounds
  - Influence Diagrams (IDs) and Limited Memory IDs (LIMIDs)
  - Decomposition of IDs and LIMIDs
  - Bounding Schemes for Maximum Expected Utility (MEU)
- Submodel Decomposition for IDs and LIMIDs
  - Motivation and Contributions
  - A Submodel-Tree Clustering Scheme
  - A Bounding Scheme over Submodel-Tree
- Experiments and Case Study
  - Upper bounds in IDs, LIMIDs, and MDP/POMDP planning domain
- Conclusion and Future Directions

Sequential Decision Making Under Uncertainty

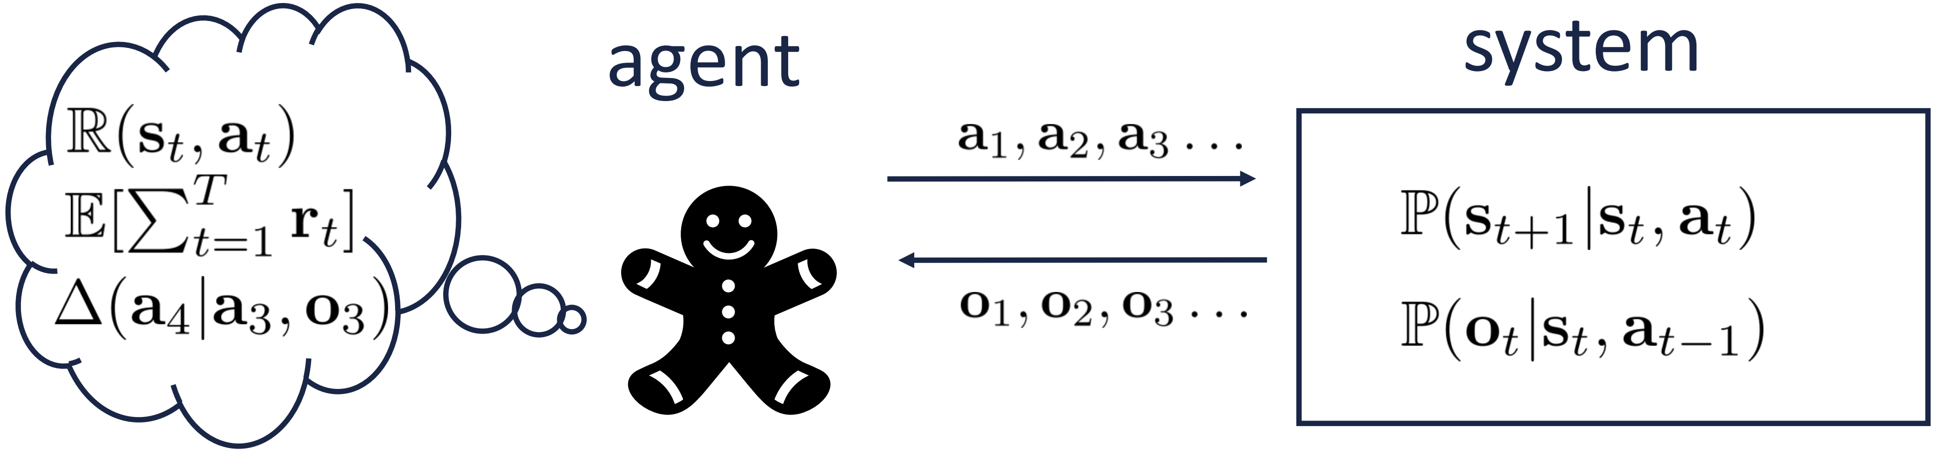

$\mathbb{P}(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ stochastic dynamics over factored state variables

$\mathbb{P}(\mathbf{o}_t | \mathbf{s}_t, \mathbf{a}_{t-1})$ stochastic partial observation

$\Delta(\mathbf{a}_4 | \mathbf{a}_3, \mathbf{o}_3)$ stochastic, non-stationary, limited memory policy

Influence Diagrams [Howard and Matheson, 1984]

$$\mathcal{M} := \langle \mathbf{X}, \mathbf{D}, \mathbf{P}, \mathbf{U}, \mathcal{O} \rangle$$

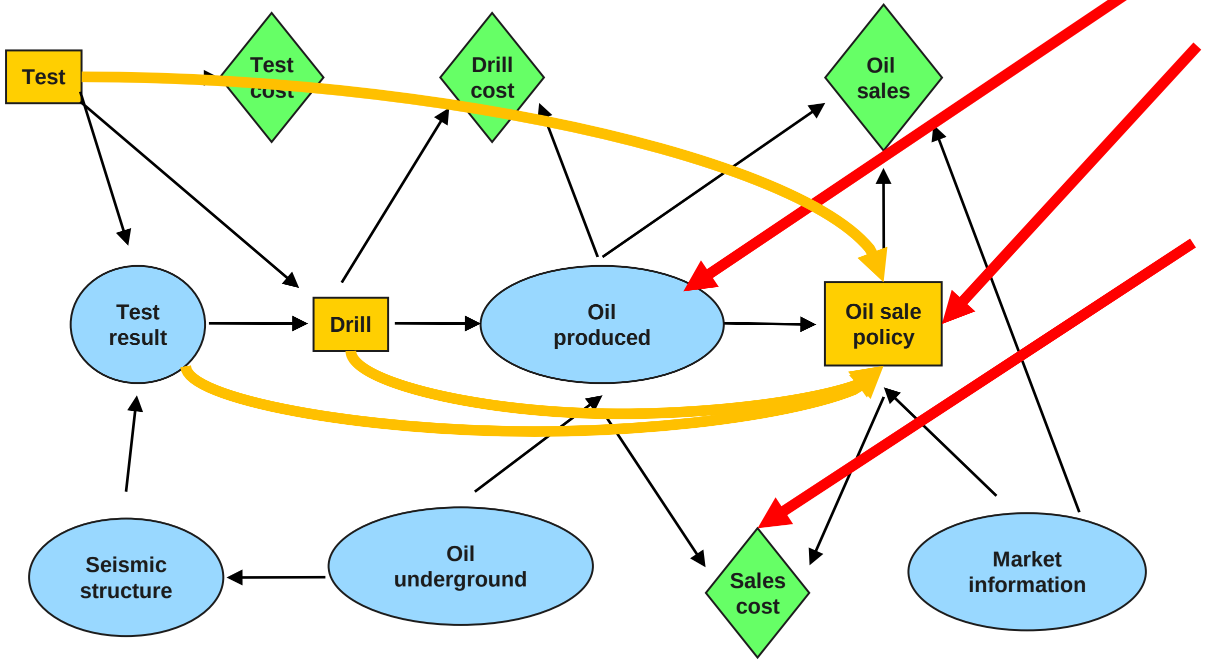

Chance variables $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$

Decision variables $\mathbf{D} = \{D_1, D_2, \dots, D_m\}$

Probability functions $\mathbf{P} = \{P_1, P_2, \dots, P_n\}$

Utility functions $\mathbf{U} = \{U_1, U_2, \dots, U_r\}$

Policy functions $\mathbf{\Delta} = \{\Delta_1, \dots \Delta_m\}$

$$\mathcal{O} = \{\text{pa}(D_1) \prec D_1 \prec \cdots \prec \text{pa}(D_m) \prec D_m\}$$

$$\Delta_i(D_i | \text{hist}(D_i))$$

Maximum expected Utility $\max_{\mathbf{\Delta}} \mathbb{E}_{P(\mathbf{X},\mathbf{D})}[\sum_{U_i \in \mathbf{U}} U_i]$ $P(\mathbf{X}, \mathbf{D}) = \prod_{P_i \in \mathbf{P}} P_i \times \prod_{\Delta_i \in \mathbf{\Delta}} \Delta_i$

Optimal strategy $\mathbf{\Delta}^* = \text{argmax}_{\mathbf{\Delta}} \mathbb{E}[\sum_{U_i \in \mathbf{U}} U_i]$

Limited Memory Influence Diagrams [Lauritzen and Nilsson, 2001]

$$\mathcal{M} := \langle \mathbf{X}, \mathbf{D}, \mathbf{P}, \mathbf{U}, \mathcal{O} \rangle$$

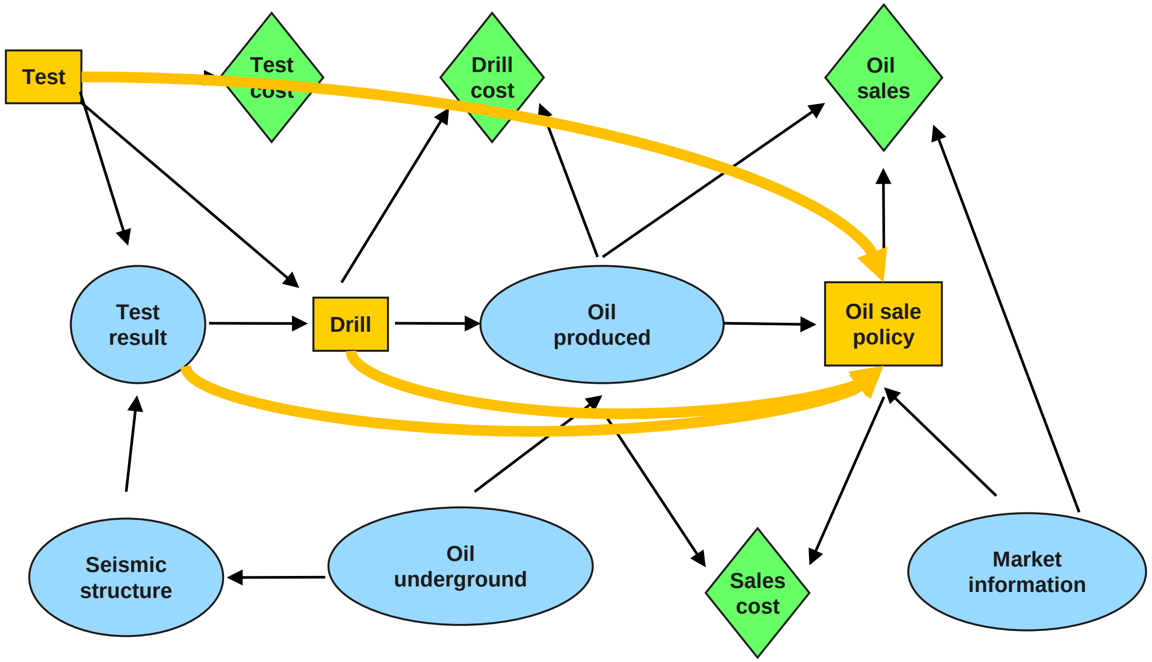

Policy functions $\boldsymbol{\Delta} = \{\Delta_1, \ldots \Delta_m\}$

$$\Delta_i(D_i|\mathrm{pa}(D_i))$$

Graphical Models

$$\mathcal{M} := \langle \mathbf{X}, \mathbf{D}, \mathbf{F} \rangle$$

Variables $\mathbf{X} = \{X_1, X_2, \ldots, X_n\}$

Domains $\mathbf{D} = \{D_1, D_2, \ldots, D_n\}$

Functions $\mathbf{F} = \{F_1, F_2, \ldots, F_r\}$

Global Function $F(\mathbf{X}) = \prod_{F_i \in \mathbf{F}} F_i(\mathbf{X}_{F_i})$

combination operator: $\otimes, \times, +, \bowtie$

Inference Task $Z = \sum_{\mathbf{X}} \prod_{F_i \in \mathbf{F}} F_i(\mathbf{X}_{F_i})$

elimination operator: $\sum$, max, min

Complexity $(\max |D_i|)^{\text{tree-width}}$ [Decther, 1999]

Bayesian Networks [Pearl 1998]

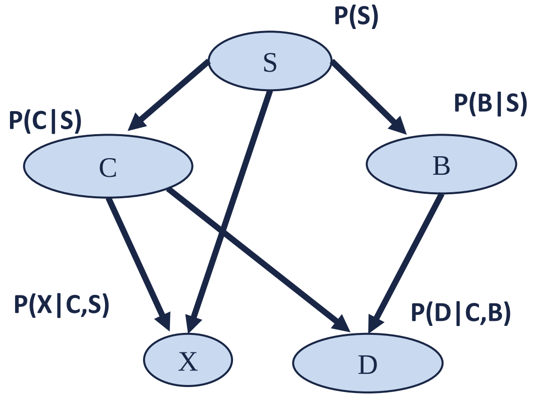

$P(S, C, B, X, D) = P(S)\, P(C|S)\, P(B|S)\, P(X|C,S)\, P(D|C,B)$

Primal Graph

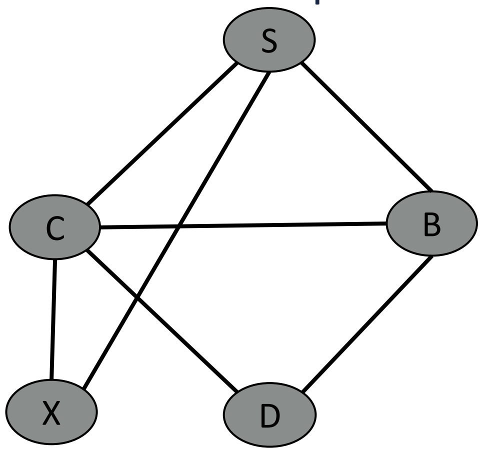

Join-Tree

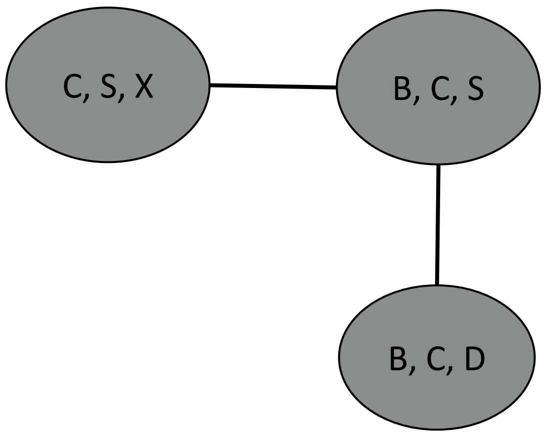

Decomposition of IDs with Perfect Recall

- Constrained Junction-Tree for IDs [Jensen, 1994]
  - Transform influence diagram to primal graph
  - Use restricted elimination order to obtain constrained tree decomposition

- Decomposition of IDs [Nielsen and Jensen1999] [Nielsen, 2001]
  - Identify requisite observation in IDs
  - Extract required subset of variables and functions for each decision variable

- MC-DAG for IDs [Pralet, et. al. 2006]
  - Re-write MEU expression and identify the most relaxed variable elimination order for computing MEU

Decomposition of LIMIDs

- Soluble LIMIDs [Zhang and Poole, 1992] [Lauritzen and Nilsson, 2001]
  - Identify a subclass of LIMIDs that can be solved by variable elimination
  - local search algorithm that improves single policy function at each iteration
- Local Search for LIMIDs [Detwarasiti and Shacter, 2005] [Maua, 2016]
  - Improve multiple policy functions at each iteration
  - Identify relevant subset of nodes for updating multiple policy functions

UCI IBM Research

Upper bounds for MEU in IDs

- IDs with perfect recall
  - Information Relaxation [Nielsen and Hohle,2001] [Yuan, et. al. 2010]
  - Join-Graph Decomposition Bounds [Lee, et. al. 2018]
  - Weighted Mini-bucket Decomposition Bounds [Lee, et. al. 2019]

- LIMIDs [Maua and Cozman, 2016]
  - Theoretical Bounds

- Translating IDs with perfect recall
  - Marginal MAP [Liu and Ihler, 2012] [Maua, 2016]
  - MILP Encodings [Parmentier et. al,2020]

UCI

IBM Research

Outline

- Backgrounds
  - Influence Diagrams (IDs) and Limited Memory IDs (LIMIDs)
  - Decomposition of IDs and LIMIDs
  - Bounding Schemes for Maximum Expected Utility (MEU)

- Submodel Decomposition for IDs and LIMIDs
  - Motivation and Contributions
  - A Submodel-Tree Clustering Scheme
  - A Bounding Scheme over Submodel-Tree

- Experiments and Case Study
  - Upper bounds in IDs, LIMIDs, and MDP/POMDP planning domain

- Conclusion and Future Work

Motivations and Contributions

- Graph-based method for decomposing IDs and LIMIDs
  - Remove some restrictions in earlier approaches
    - 1 decision per 1 time step, regularity condition, perfect recall
  - Extend tree clustering framework for reasoning in graphical models
    - Identify subproblems from graph
    - Extract a cluster tree for exact algorithms
    - Characterize complexity
- Upper-Bounds for MEU in IDs and LIMIDs
  - Don’t inflate problem size by translation
  - Avoid difficult non-convex optimization formulations in earlier works

Partial Evaluation and Local MEU

- (Definition) Local Maximum Conditional Expected Utility

$$\mathrm{LMEU}_{\mathcal{M}(\mathbf{D}',\mathbf{U}')} := \max_{\mathbf{\Delta}'} \mathbb{E}\left[\sum_{U_i \in \mathbf{U}'} U_i | \mathrm{pa}(\mathbf{D}')\right]$$

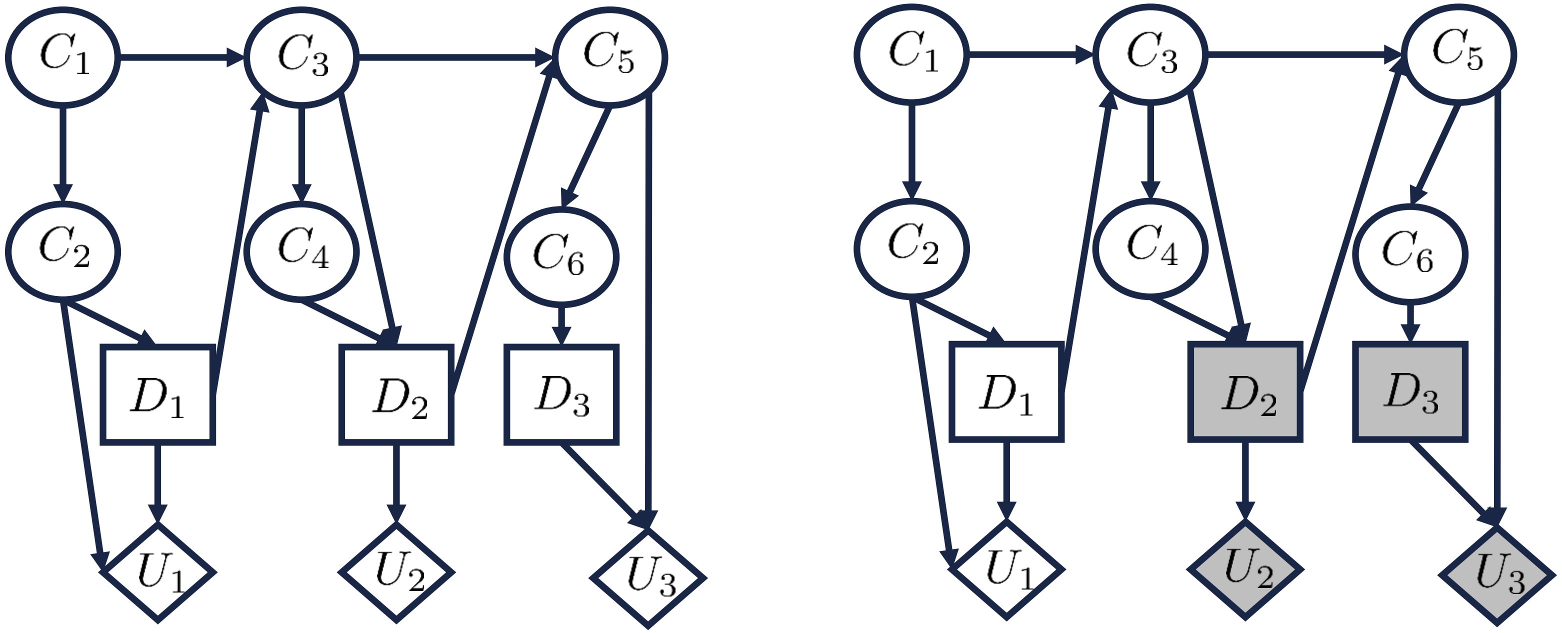

$$\max_{\Delta(D_2|C_3,C_4),\Delta(D_3|C_6)} \sum_{\mathbf{X},\mathbf{D}} \frac{P(\mathbf{X},\mathbf{D})}{P(C_3,C_4,C_6)} \left[U_2 + U_3\right]$$

Submodel

- (Definition) Submodel $\mathcal{M}'(\mathbf{D}', \mathbf{U}')$ is a relevant subset of model $\mathcal{M}$ for computing LMEU on $\mathbf{D}' \subseteq \mathbf{D}, \mathbf{U}' \subseteq \mathbf{U}$

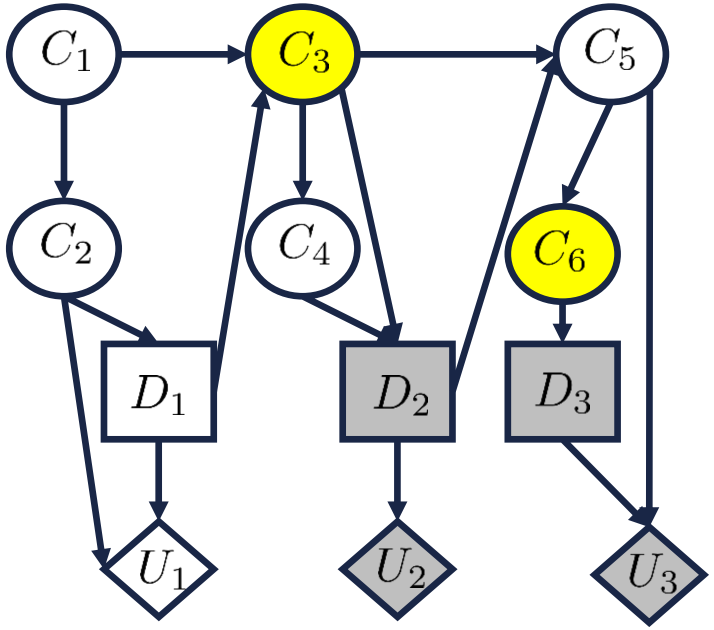

Relevant Observed Variables

$\mathrm{REL}_O(\mathbf{D}', \mathbf{U}')$

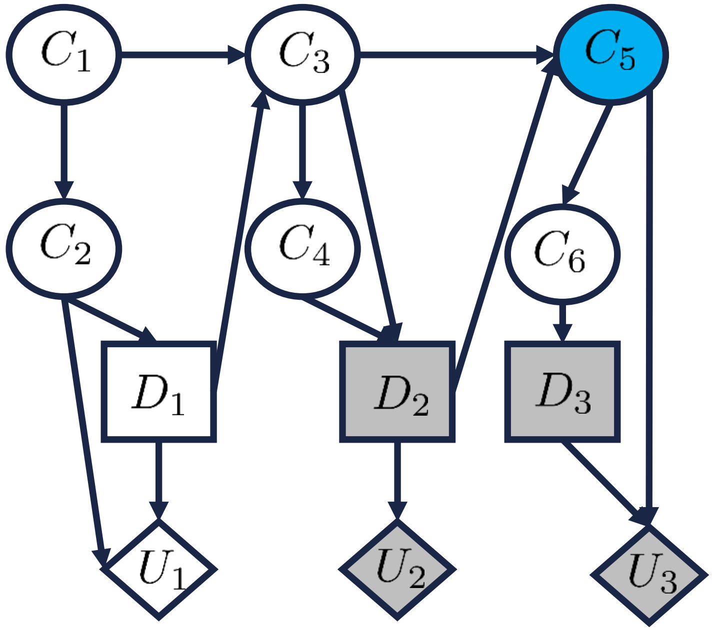

Relevant Hidden Variables

$\mathrm{REL}_H(\mathbf{D}', \mathbf{U}')$

Stable Submodel

- (Definition) Submodel $\mathcal{M}'(\mathbf{D}', \mathbf{U}')$ is stable when there is no decision variables in $\mathrm{REL}_H(\mathbf{D}', \mathbf{U}')$

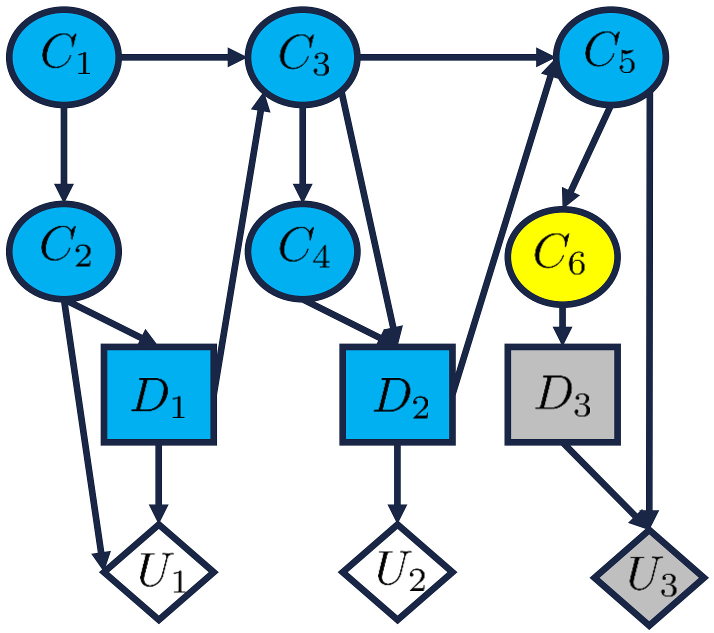

Unstable Submodel $\mathcal{M}'(\{D_3\}, \{U_3\})$

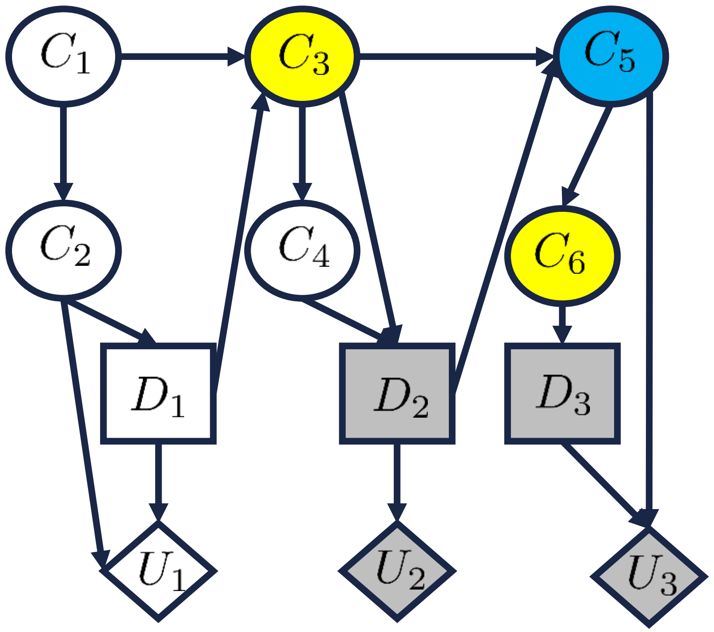

Stable Submodel $\mathcal{M}'(\{D_2, D_3\}, \{U_2, U_3\})$

Graph-based Identification of Submodels

- $\mathrm{REL}_U(\mathbf{D}')$ is descendant utility nodes of decision nodes [Nielsen and Jensen1999]

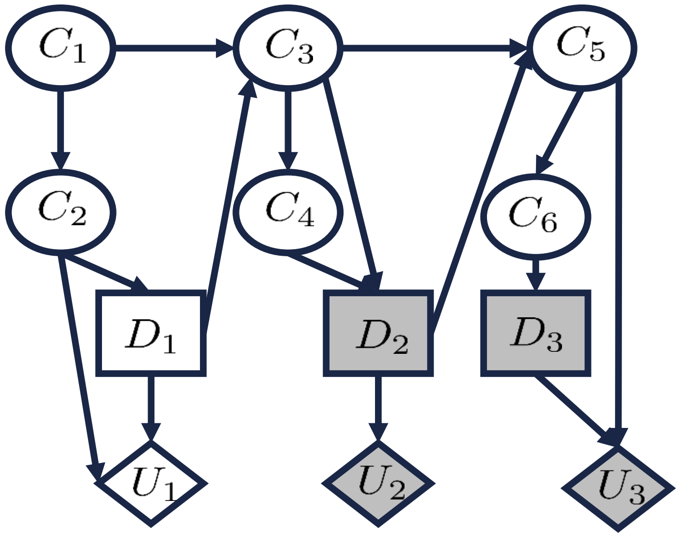

Graph-based Identification of Submodels

- $\mathrm{REL}_O(\mathbf{D}', \mathbf{U}')$ is the backdoor* set between D' and U'

(Backdoor) [Pearl 2009]
a set Z satisfies the backdoor criterion relative to (X, Y)
(1) None of the nodes in Z is a descendant of X
(2) Z blocks every path between X and Y that contain arrow into X

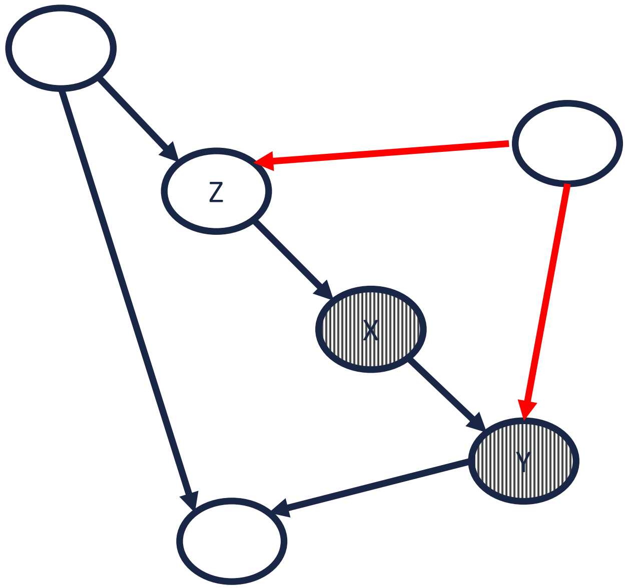

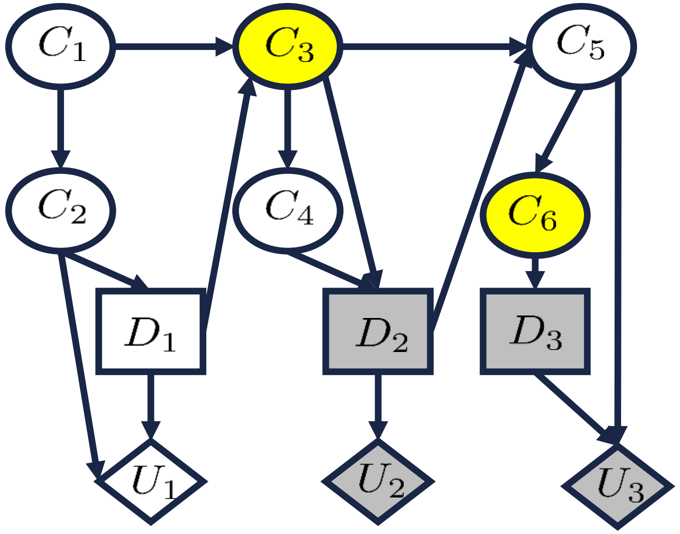

{C3, C6} is a backdoor set relative to ({D2,D3}, {U2, U3})

Removing C3 opens a backdoor path by
C1->C3->D2 ->C3->C5->U3

Graph-based Identification of Submodels

- $\mathrm{REL}_H(\mathbf{D}', \mathbf{U}')$ is the union of all frontdoor* set between pa(D') and ch(U')

(Frontdoor) [Pearl 2009]
a set Z satisfies the frontdoor criterion relative to (X, Y)
(1) Z intercept all directed paths from X to Y
(2) There is no backdoor path from X to Z
(3) All backdoor paths from Z to Y are blocked by X

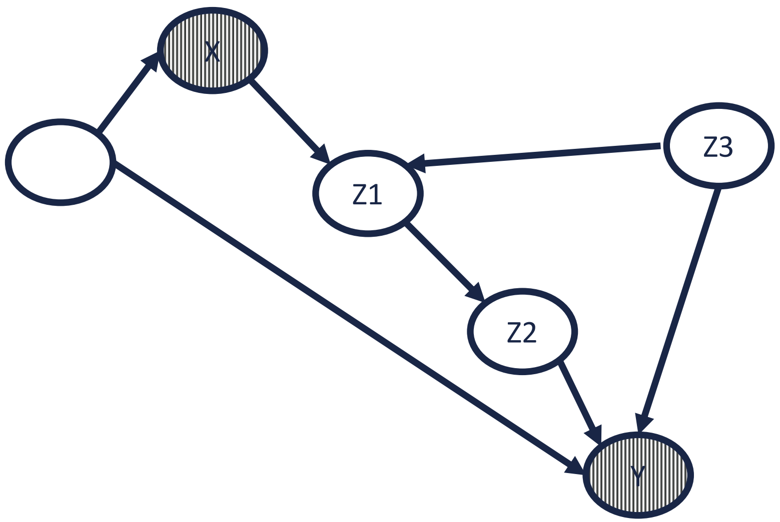

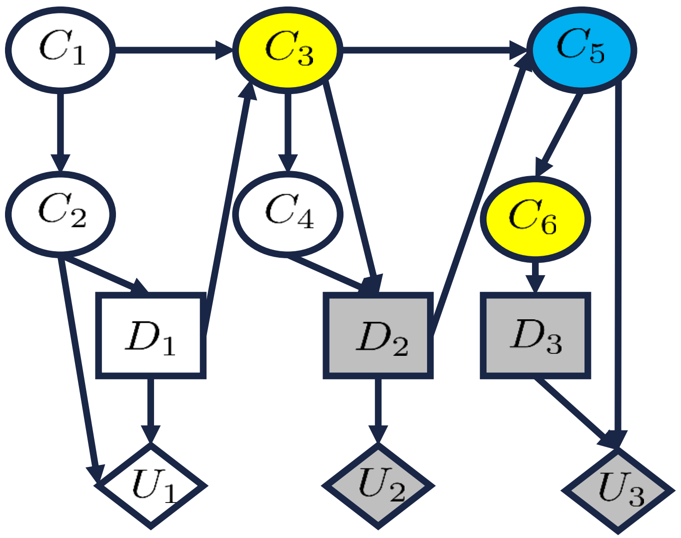

C1, C2, D1, and C4 don't belong to any frontdoor set

Submodel–Tree Clustering

- Process decision nodes in reverse topological order

Partial decision order $\mathcal{O}_D = \{D_1 \prec D_2 \prec D_3\}$

Process decision variables in the order of D3, D2, and D1

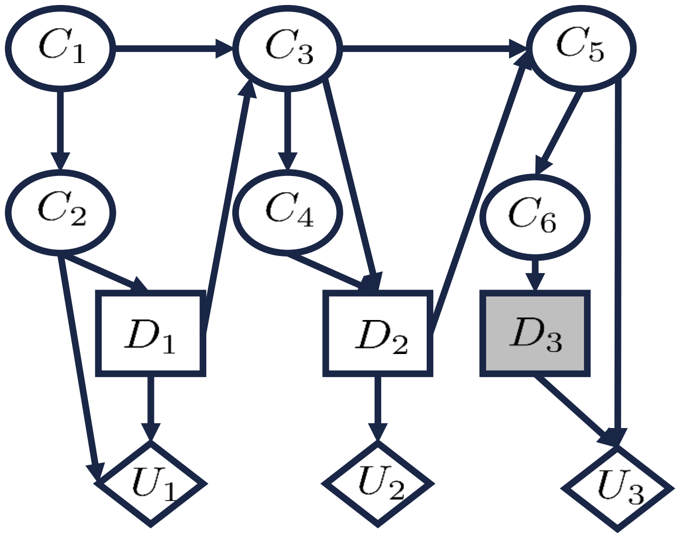

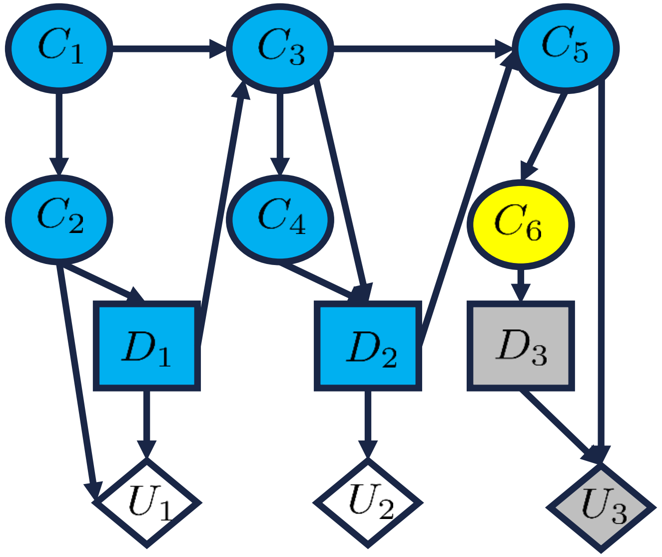

Submodel $\mathcal{M}'(\{D_3\}, \{U_3\})$ is unstable

Submodel–Tree Clustering

- Find a stable submodel

Next combine two submodels $\mathcal{M}'(\{D_3\},\{U_3\}) \otimes \mathcal{M}'(\{D_2\},\{U_2,U_3\})$
and Try $\mathcal{M}'(\{D_2,D_3\},\{U_2,U_3\})$

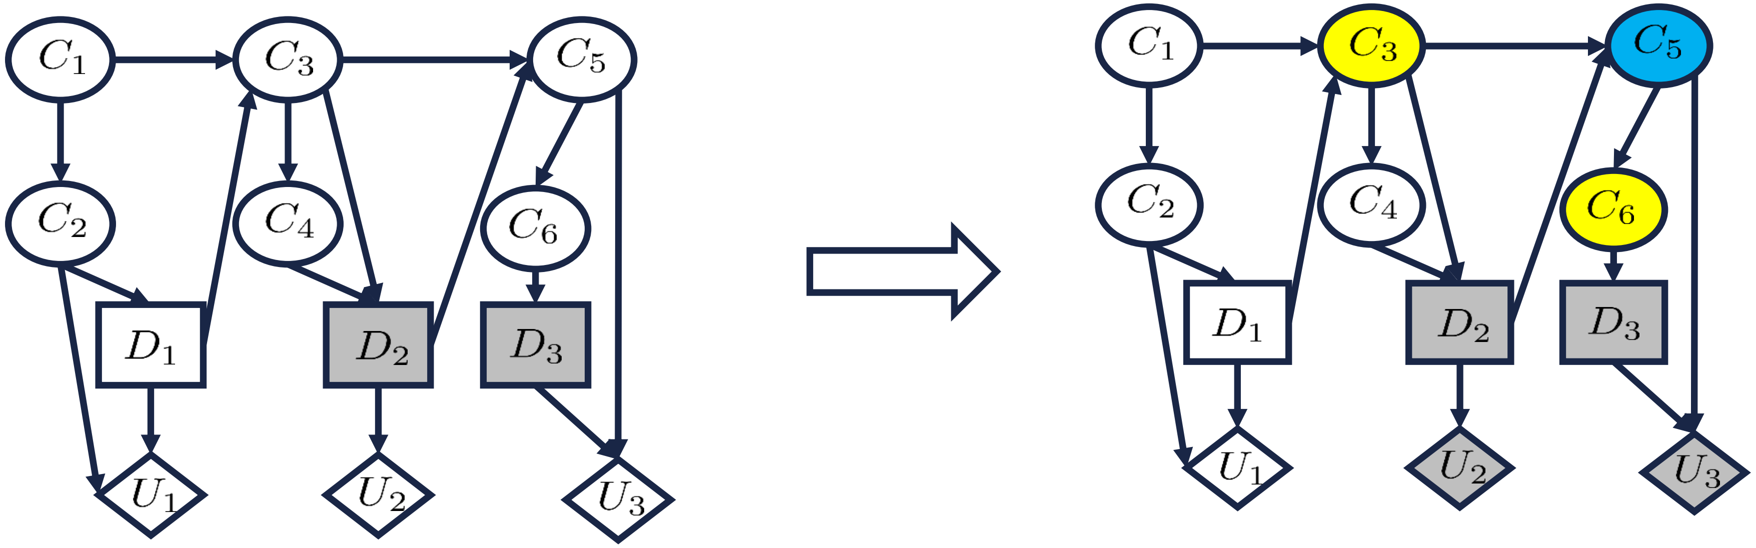

$\mathcal{M}'(\{D_2,D_3\},\{U_2,U_3\})$ is stable

UCI

IBM Research

Submodel–Tree Clustering

- Eliminate submodel $\mathcal{M}'(\{D_2, D_3\}, \{U_2, U_3\})$ from IDs $\mathcal{M}$

Remove D2, D3, U2, U3 and Add V(C3)

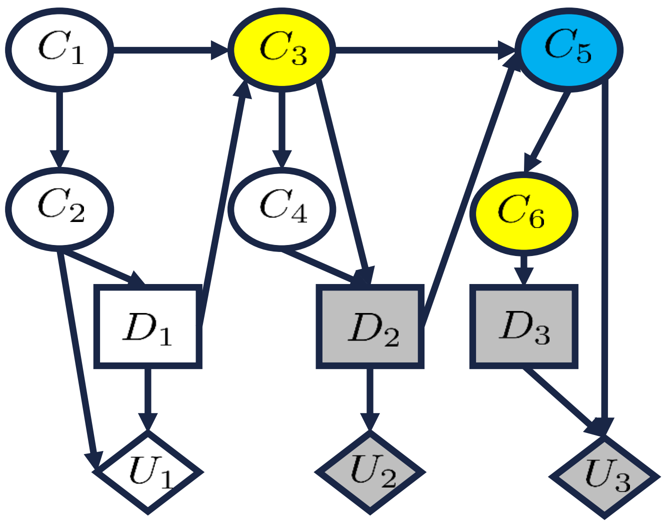

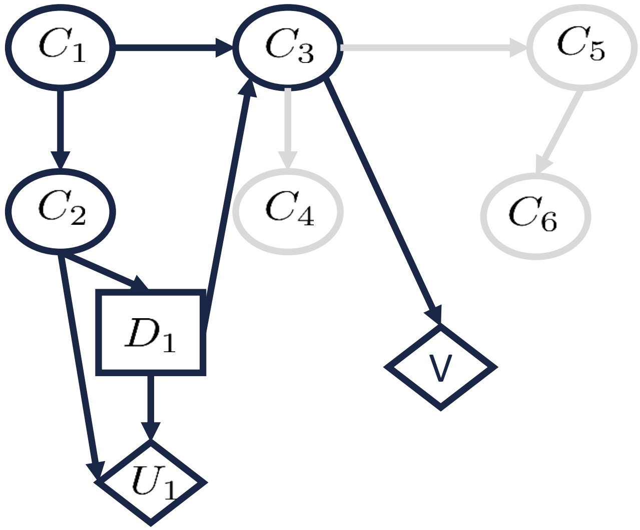

Remove barren chance nodes C4, C5, C6

Submodel–Tree Clustering

- Identify the next submodel and find a submodel-tree

$\mathcal{M}'(\{D_1\},\{U_1,V\})$ $\mathcal{M}'(\{D_2,D_3\},\{U_2,U_3\})$

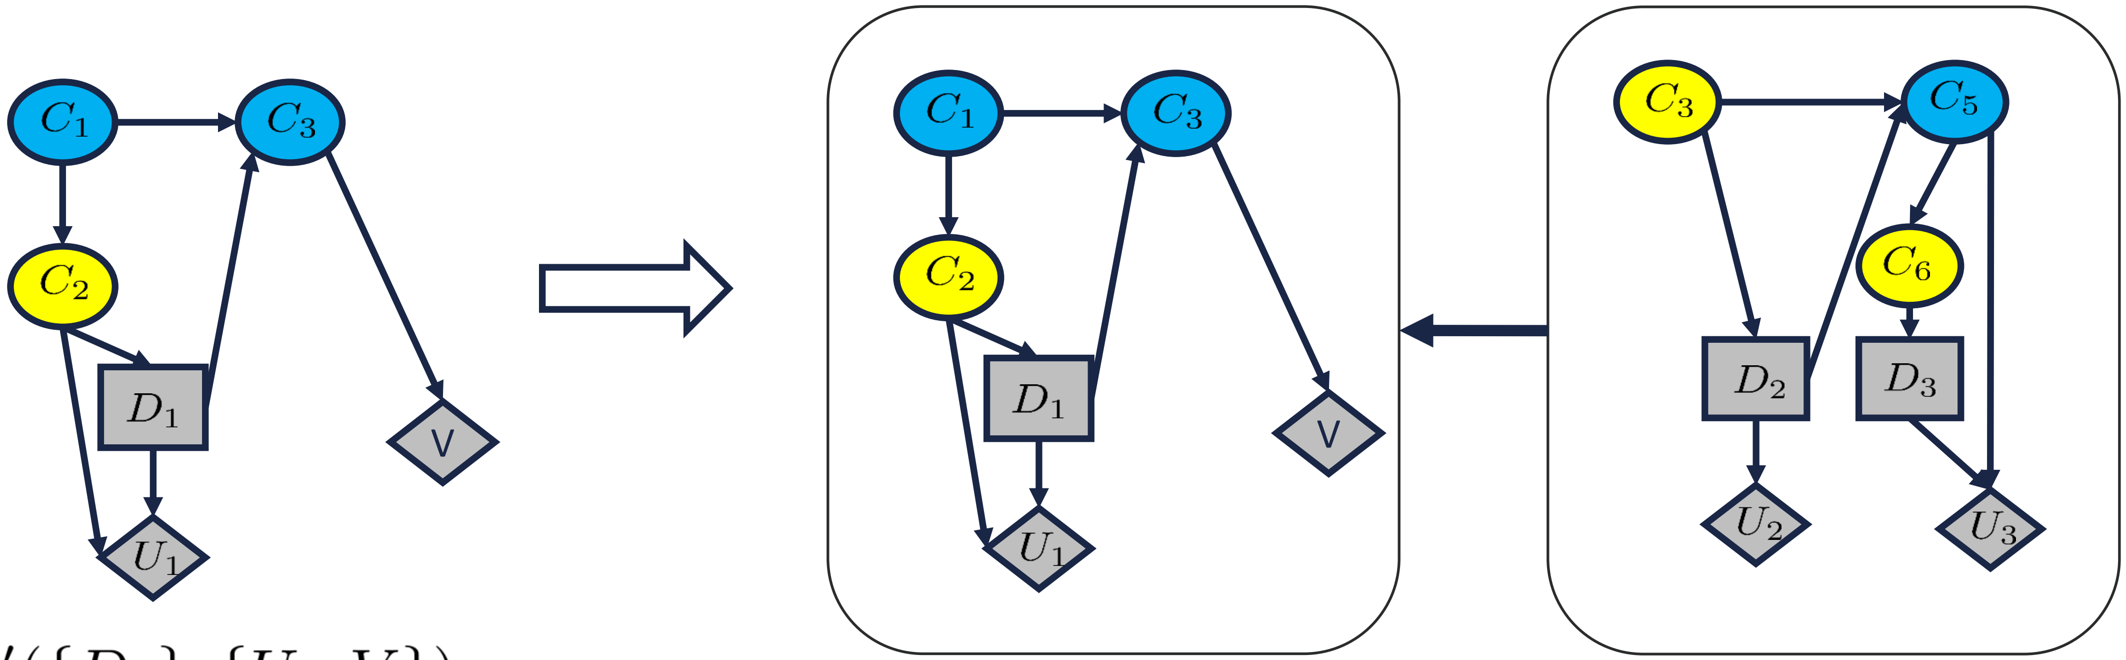

$\mathcal{M}'(\{D_1\},\{U_1,V\})$

Submodel Cluster is a single-stage ID

Submodel Cluster Propagates Conditional MEU

Valuation Algebra over Stable Submodels

- Given an ID $\mathcal{M}$, $\Upsilon_{\mathcal{M}} = \langle \mathbb{M}_{\mathcal{O}_D}, \mathbb{D}_{\mathcal{O}_D}, \otimes, \Downarrow \rangle$ is a valuation algebra

[Shenoy 1997] [Kohlas and Shenoy, 2000]

$\mathcal{O}_D$	Partial decision order read from ID
$\mathbf{M}_{\mathcal{O}_D}$	A set of stable submodels in $\mathcal{M}$ subject to $\mathcal{O}_D$
$\otimes$	Combination operator for a submodel
$\Downarrow$	Projection operator for a submodel
$\mathbb{M}_{\mathcal{O}_D}$	A closure of $\mathbf{M}_{\mathcal{O}_D}$ under the combination
$\mathrm{dom}(\mathcal{M}')$	Domain of a submodel (all variables in $\mathcal{M}'$)
$\mathbb{D}_{\mathcal{O}_D}$	A set of domains of submodels in $\mathbb{M}_{\mathcal{O}_D}$

Valuation Algebra over Stable Submodels

- Given an ID $\mathcal{M}$, $\Upsilon_{\mathcal{M}} = \langle \mathbb{M}_{\mathcal{O}_D}, \mathbb{D}_{\mathcal{O}_D}, \otimes, \Downarrow \rangle$ is a valuation algebra

[Kohlas and Shenoy, 2000]

Semi-group of submodels: $\mathbb{M}_{\mathcal{O}_D}$ is a semi-group with the combination operation

Domain of combination: $\mathrm{dom}(\mathcal{M}'_1 \otimes \mathcal{M}'_2) = \mathrm{dom}(\mathcal{M}'_1) \cup \mathrm{dom}(\mathcal{M}'_2)$

Marginalization: $\Downarrow_{\mathbf{X}} \mathcal{M}' = \Downarrow_{\mathbf{X} \cap \mathrm{dom}(\mathcal{M}')} \mathcal{M}'$

$\mathrm{dom}(\Downarrow_{\mathbf{X}} \mathcal{M}') = \mathbf{X} \cap \mathrm{dom}(\mathcal{M}')$

$\Downarrow_{\mathrm{dom}(\mathcal{M}')} \mathcal{M}' = \mathcal{M}'$

Transitivity of marginalization: $\Downarrow_{\mathbf{X}} (\Downarrow_{\mathbf{Y}} \mathcal{M}') = \Downarrow_{\mathbf{X} \cap \mathbf{Y}} \mathcal{M}'$

Distributivity of marginalization over combination: $\Downarrow_{\mathbf{X}} (\mathcal{M}'_1 \otimes \mathcal{M}'_2) = \mathcal{M}'_1 \otimes (\Downarrow_{\mathbf{X}} \mathcal{M}'_2)$

Neutral elements: $\mathcal{M}'_{\mathbf{0}(\mathbf{X})} \otimes \mathcal{M}'_{\mathbf{0}(\mathbf{Y})} = \mathcal{M}'_{\mathbf{0}(\mathbf{X} \cup \mathbf{Y})}$

- Valuation algebra satisfies axioms of local computation [Shenoy 1997]

Submodel–Tree Decomposition

- Given an ID $\mathcal{M}$, and the set of stable submodels $\mathbf{M}_{\mathcal{O}_D}$ relative to $\mathcal{O}_D$, submodel-tree decomposition is a tuple $\mathcal{T}_{\mathrm{ST}} := \langle T(\mathcal{C}, \mathcal{S}), \chi, \psi \rangle$

$T(\mathcal{C}, \mathcal{S})$ Tree of submodel cluster nodes $\mathcal{C}$ and separator edges $\mathcal{S}$

$\chi : \mathcal{C} \rightarrow 2^{\mathrm{dom}(\mathcal{M})}$ Label a cluster with a subset of variables in $\mathcal{M}$

$\psi : \mathcal{C} \rightarrow 2^{\mathbf{M}_{\mathcal{O}_D}}$ Label a cluster with a subset of submodels in $\mathbf{M}_{\mathcal{O}_D}$

Tree-decomposition satisfies running intersection property

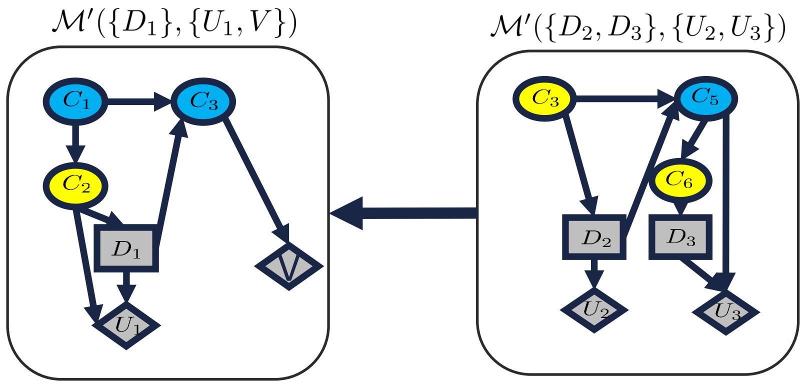

Submodel–Tree Decomposition

- Minimal submodel-tree decomposition

A submodel-tree decomposition is minimal
if submodels assigned at each cluster is not a combination of two stable submodels

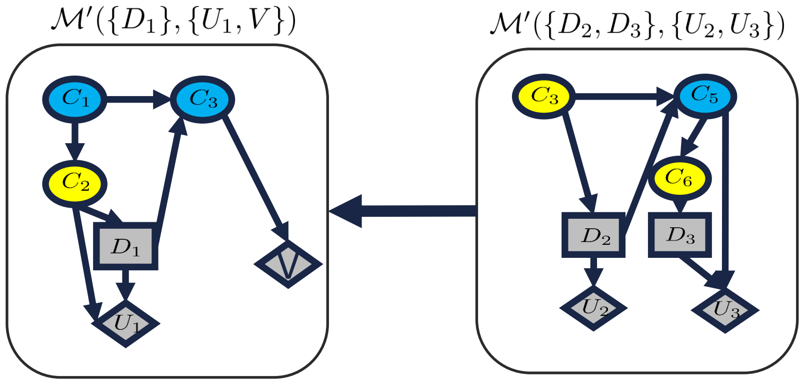

- Given an ID, minimal submodel-tree decomposition is unique.
- For IDs with perfect recall, the minimal submodel-tree is equivalent to MC-DAG
- For IDs with perfect recall, each submodel cluster is one time-step ID
- For IDs without perfect recall, each submodel cluster defines the scope of exhaustive search

Message Passing over a Submodel-Tree

$\mathcal{M}'(\{D_1\},\{U_1,V\})$

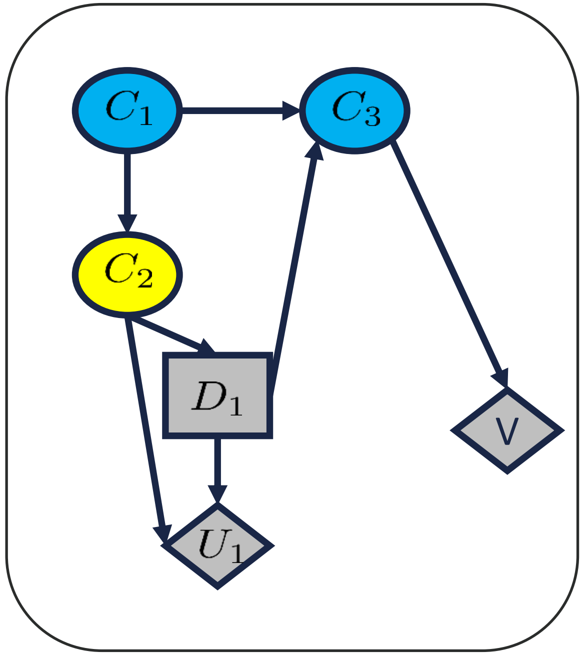

$V(C_3) = \max_{\Delta(D_2,D_3|C_3,C_6)} \mathbb{E}[U_2+U_3|C_3]$

$P(C_3) = \sum_{C_1,C_2,D_1} P(C_1,C_2,C_3,D_1)$

$\mathcal{M}'(\{D_2,D_3\},\{U_2,U_3\})$

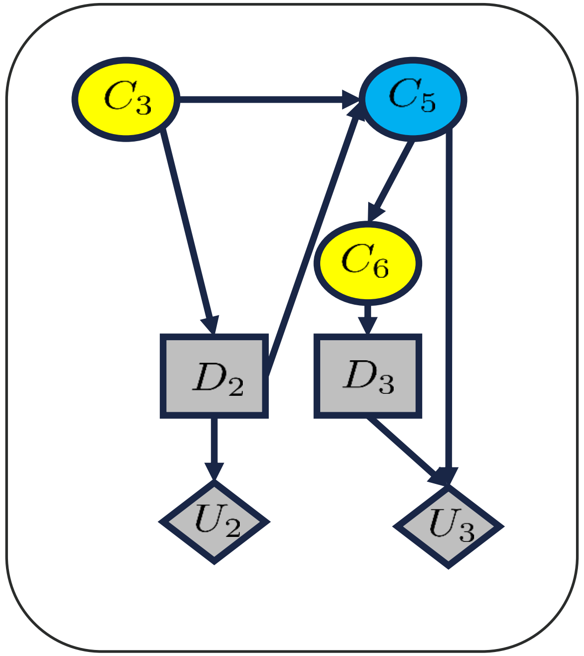

- Each submodel can be solved by any exact algorithm for propagating messages

Message Passing over a Submodel-Tree

Algorithm 1 Hierarchical Message Passing over $\mathcal{T}_{\mathrm{ST}}$

Require: $\mathcal{T}_{\mathrm{ST}} := \langle T(\mathcal{C}, \mathcal{S}), \chi, \psi \rangle$
Ensure: Submodel-tree augmented with messages sent out from clusters, MEU
MEU $\leftarrow 0$
for each cluster C from the leaves to the root in T **do**
  Pull messages from incoming edges
  Update submodel $\mathcal{M}^C$ at C with the received messages
  $V^C \leftarrow Eval(\mathcal{M}^C)$
  if C is the root node or V^C is a constant **then**
    MEU $\leftarrow$ MEU $+ V^C$
  else
    Push V^C to the outgoing edge
return MEU

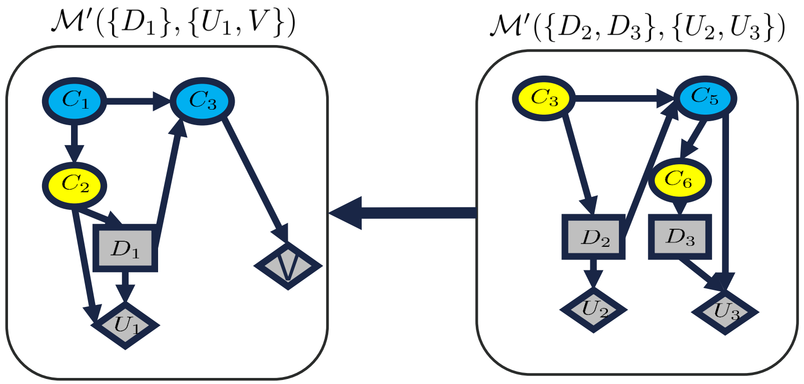

The time and space complexity for solving IDs over the submodel-tree decomposition is exponential in submodel-tree width $w_s : \max_{C \in \mathcal{C}} w_c(C)$, where $w_c(C)$ is the constrained tree-width of the submodel at C

UCI

IBM Research

Bounding MEU of Each Submodel

- Exponentiated Utility Bounds for MEU

For each submodel cluster, we can apply Jensen's inequality to bound MEU

$$\max_{\Delta} \mathbb{E}\Big[\sum_{U_i \in \mathbf{U}} U_i(\mathbf{X}_{U_i})\Big] \leq \max_{\Delta} \log \mathbb{E}\Big[e^{\sum_{U_i \in \mathbf{U}} U_i(\mathbf{X}_{U_i})}\Big]$$

$$= \log \max_{\Delta} \mathbb{E}\Big[e^{\sum_{U_i \in \mathbf{U}} U_i(\mathbf{X}_{U_i})}\Big]$$

$$= \log \max_{\Delta} \mathbb{E}\Big[\prod e^{U_i(\mathbf{X}_{U_i})}\Big]$$

LSH: MEU expression with additive utility function

RHS: Upper bound of MEU with log-partition function with exponentiated utility functions

- Use "any" upper bounding scheme for MMAP on RHS

UCI IBM Research

Outline

- Backgrounds
  - Influence Diagrams (IDs) and Limited Memory IDs (LIMIDs)
  - Decomposition of IDs and LIMIDs
  - Bounding Schemes for Maximum Expected Utility (MEU)
- Submodel Decomposition for IDs and LIMIDs
  - Motivation and Contributions
  - A Submodel-Tree Clustering Scheme
  - A Bounding Scheme over Submodel-Tree
- Experiments and Case Study
  - Upper bounds in IDs, LIMIDs, and MDP/POMDP planning domain
- Conclusion and Future Directions

UCI IBM Research

Benchmark Domains

Finite Horizon MDP

	n	c	d	f	p	u	k	s	w
Min	25	20	3	30	20	10	2	4	5
Average	105.7	99.6	6.1	134.1	99.6	34.5	3.1	7.1	25.5
Max	170	160	10	240	160	80	5	9	43

Finite Horizon POMDP

	n	c	d	f	p	u	k	s	w
Min	15	12	3	18	12	6	2	3	10
Average	55.9	52.4	3.5	73.5	52.4	21.1	2.4	5.5	28
Max	96	92	5	140	92	48	3	9	46

Random Influence Diagrams

	n	c	d	f	p	u	k	s	w
Min	22	20	2	22	20	2	2	3	6
Average	56	47	9	56	47	9	2	3	17
Max	91	70	21	91	70	21	2	3	34

IDs converted from BN

	n	c	d	f	p	u	k	s	w
Min	54	48	3	54	48	3	2	6	12
Average	84	77	7	84	77	7	2	8	21
Max	115	109	12	115	109	12	2	10	42

Experiments: Synthetic IDs

Domain	n	w_c	w_s	ST-GDD(i=1)	ST-GDD(i=5)	ST-WMB(i=10)	JGDID(i=1)	WMBEID(i=10)
ID-BN	84.6	30.2	21.8	0.19	0.15	**0.13**	0.33	0.74
IDBN14w57d12	115	57	42	103.89	96.24	95.37	1420	2.2E+4
FH-MDP	105.7	25.5	25.4	**0.06**	0.07	0.18	0.16	0.44
mdp9-32-3-8-3	99	43	43	18.92	19.71	25.31	23.09	111.81
FH-POMDP	55.9	28.1	28.1	0.31	0.22	**0.06**	0.56	0.72
pomdp8-14-9-3-12-14	96	47	46	73.53	76.37	67.18	5.E+08	5.E+09
RAND	56.2	20.5	17.9	**0.22**	0.24	0.24	0.23	0.46
rand-c70d21o1	84	32	34	1309.89	1791.93	1752.47	1743.6	2.E+04

- ST-GDD: submodel-tree decomposition + GDD for MMAP [ping et al 2015]
- ST-WMB: submodel-tree decomposition + WMBMMM for MMMAP [marinescu et al 2014]
- JGDID: constrained-join graph + GDD for IDs [Lee et al 2018]
- WMBMEID: constrained mini-bucket tree + WMB/GDD for IDs [Lee et al 2019]

- Evaluation: average of the gap $\frac{U - U_{\min}}{U}$

Experiments: Synthetic LIMIDs

	w_s	|C|	STWMBMM.10.	STWMBMM.20	kpu-UB
ID_from_BN_0_w28d6_limid	26	1	44.30	43.86	9.08E+06
ID_from_BN_0_w29d6_limid	22	2	47.47	44.79	9.08E+06
ID_from_BN_0_w32d11_limid	31	2	83.81	81.74	2.74E+07
ID_from_BN_0_w33d11_limid	34	2	87.14	86.60	2.74E+07
ID_from_BN_14w42d6_limid	33	1	45.25	43.61	1.22E+08
ID_from_BN_14w57d12_limid	46	2	94.45	91.87	4.70E+08
ID_from_BN_78_w18d3_limid	15	1	23.54	22.80	4.61E+08
ID_from_BN_78_w19d3_limid	17	2	24.07	23.28	4.61E+08
ID_from_BN_78_w23d6_limid	17	1	46.97	45.21	1.83E+09
ID_from_BN_78_w24d6_limid	25	1	48.72	46.08	1.83E+09
mdp1-4_2_2_5_limid	5	6	3.74	3.74	4.42E+05
mdp10-32_3_8_4_limid	42	10	37.96	33.39	3.74E+09
mdp2-8_3_4_5_limid	10	7	12.20	12.20	7.45E+07
mdp3-10_3_5_10_limid	14	14	43.40	31.08	2.60E+09
mdp4-10_3_5_10_limid	14	11	39.11	32.73	2.60E+09
mdp5-16_3_8_10_limid	21	16	79.42	48.98	2.32E+10
mdp6-20_5_5_5_limid	29	6	31.56	21.13	4.58E+15
mdp7-28_3_6_5_limid	41	11	44.01	34.90	4.66E+11
mdp8-28_3_6_4_limid	41	9	32.42	27.77	2.98E+11
mdp9-32_3_8_3_limid	45	11	25.23	22.29	1.68E+11

	w_s	|C|	STWMBMM.10.	STWMBMM.20	kpu-UB
pomdp1-4_2_2_2_3_limid	11	1	4.34	4.12	1.79E+05
pomdp10-12_7_3_8_4_limid	40	9	29.54	24.17	5.06E+07
pomdp2-2_2_2_2_3_limid	10	1	4.51	4.51	5.55E+04
pomdp3-4_4_2_2_3_limid	16	1	4.76	3.62	1.97E+05
pomdp4-4_4_2_2_5_limid	27	2	7.69	5.74	5.39E+05
pomdp5-6_4_3_5_3_limid	18	3	11.23	9.19	2.72E+07
pomdp6-12_6_2_6_3_limid	29	5	19.69	14.07	1.19E+08
pomdp7-20_10_2_10_3_limid	46	8	37.30	30.45	1.22E+08
pomdp8-14_9_3_12_4_limid	47	12	46.90	39.49	4.28E+08
pomdp9-14_8_3_10_4_limid	44	8	41.70	32.90	4.27E+08
rand-c20d2o1-01_limid	6	1	170.63	170.63	3.76E+04
rand-c30d3o1-01_limid	11	1	272.77	272.71	8.27E+04
rand-c30d6o1-01_limid	13	2	537.87	537.75	2.15E+05
rand-c30d9o1-01_limid	16	2	735.74	733.01	4.11E+05
rand-c50d10o1-01_limid	23	2	863.38	863.11	5.91E+05
rand-c50d15o1-03_limid	27	2	1243.53	1241.16	1.13E+06
rand-c50d5o1-01_limid	17	1	429.09	428.37	2.25E+05
rand-c70d14o1-01_limid	31	2	1194.34	1193.46	1.15E+06
rand-c70d21o1-01_limid	41	2	1735.30	1731.70	2.21E+06
rand-c70d7o1-01_limid	22	1	657.22	655.98	4.38E+05

- |C|: number of clusters in submodel tree
- Kpu-UB: Analytical bound by [Maua and Cozman 2016]

Experiments: SysAdmin MDP/POMDP [Guestrin, et. al 2003]

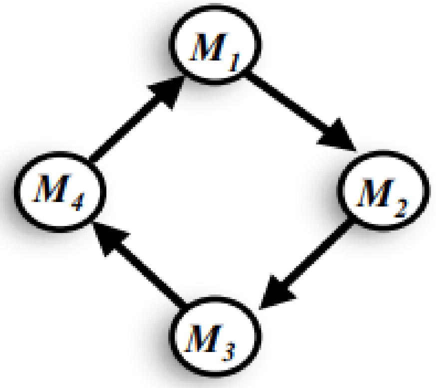

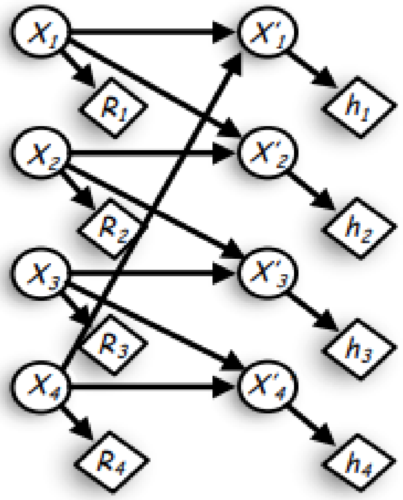

$P(X_i' = t \mid X_i, X_{i-1}, A)$:

	Action is reboot: machine i	Action is reboot: other machine
$X_{i-1} = f \wedge X_i = f$	1	0.0238
$X_{i-1} = f \wedge X_i = t$	1	0.475
$X_{i-1} = t \wedge X_i = f$	1	0.0475
$X_{i-1} = t \wedge X_i = t$	1	0.95

- Evaluation
  - UB: WMBMM-EXP (i=20)
  - LB: Online planner to obtain lower bounds

$$\text{gap} = 1 - \frac{\text{LB}}{\text{UB}}$$

SysAdmin MDP

Instance	c	d	p	u	k	s	w	utime (sec)	ub wmbmm	ltime (sec)	lb sogbofa	gap ($\frac{ub-lb}{ub}$)
mdp9-s50-t3	227	150	227	300	3	8	107	217	161.292	240	142.931	11%
mdp9-s50-t4	286	200	286	400	3	8	108	241	218.210	320	184.250	16%
mdp9-s50-t5	345	250	345	500	3	8	109	298	274.907	400	222.925	19%
mdp9-s50-t6	404	300	404	600	3	8	109	353	329.248	480	261.319	21%
mdp9-s50-t7	463	350	463	700	3	8	109	49406	383.290	560	296.269	23%
mdp9-s50-t8	522	400	522	800	3	8	109	530	438.786	640	328.550	25%
mdp9-s50-t9	581	450	581	900	3	8	108	370	496.466	720	355.263	28%
mdp9-s50-t10	640	500	640	1000	3	8	109	129	547.757	800	385.263	30%
mdp10-s50-t3	218	150	218	300	3	11	112	149	162.368	240	142.731	12%
mdp10-s50-t4	274	200	274	400	3	11	112	184	217.515	320	183.650	16%
mdp10-s50-t5	330	250	330	500	3	11	113	257	273.332	400	221.450	19%
mdp10-s50-t6	386	300	386	600	3	11	113	19741	327.268	480	257.394	21%
mdp10-s50-t7	442	350	442	700	3	11	113	28013	383.312	560	291.988	24%
mdp10-s50-t8	498	400	498	800	3	11	113	41748	439.826	640	316.600	28%
mdp10-s50-t9	554	450	554	900	3	11	113	34739	494.662	720	345.844	30%
mdp10-s50-t10	610	500	610	1000	3	11	113	58270	549.867	800	364.569	34%

SysAdmin POMDP

Instance	c	d	p	u	k	s	w	utime (sec)	ub wmbmm	ltime (sec)	lb snap	gap $(\frac{ub-lb}{ub})$
pomdp9-s50-t3	371	150	371	300	3	10	300	7660	177.597	240	141.575	20%
pomdp9-s50-t4	478	200	478	400	3	10	400	20747	244.415	320	182.775	25%
pomdp9-s50-t5	585	250	585	500	3	10	501	27369	315.351	400	223.000	29%
pomdp9-s50-t6	692	300	692	600	3	10	600	69085	384.860	480	257.350	33%
pomdp9-s50-t7	799	350	799	700	3	10	701	107660	454.910	560	290.425	36%
pomdp9-s50-t8	906	400	906	800	3	10	800	224727	528.582	640	323.575	39%
pomdp9-s50-t9	1013	450	1013	900	3	10	900	207883	599.923	720	348.325	42%
pomdp9-s50-t10	1120	500	1120	1000	3	10	1000	231511	668.883	800	372.750	44%
pomdp10-s50-t3	371	150	371	300	3	10	300	6249	180.257	240	141.950	21%
pomdp10-s50-t4	478	200	478	400	3	10	400	14333	253.043	320	184.925	27%
pomdp10-s50-t5	585	250	585	500	3	10	501	36078	331.108	400	221.775	33%
pomdp10-s50-t6	692	300	692	600	3	10	600	66848	409.456	480	258.525	37%
pomdp10-s50-t7	799	350	799	700	3	10	701	121312	480.291	560	290.300	40%
pomdp10-s50-t8	906	400	906	800	3	10	800	116597	563.324	640	321.425	43%
pomdp10-s50-t9	1013	450	1013	900	3	10	900	290003	633.134	720	346.500	45%
pomdp10-s50-t10	1120	500	1120	1000	3	10	1000	244446	707.226	800	375.900	47%

UCI IBM Research

Conclusion and Future Directions

- Extend Tree-Clustering Framework in PGM for IDs and LIMIDs
  - Graph-based tree-clustering procedure for IDs and LIMIDs
  - Hierarchical message passing algorithm for exact inference

- Simple and Scalable Bounding Scheme for IDs
  - Exponentiating utility functions and reuse decomposition bounds for MMAP

- Future Directions
  - Guide heuristic search for finding MEU in IDs and LIMIDs
  - Extend relaxation schemes in PGM to submodel-tree decomposition
  - Submodel-tree clustering framework for multi-agent IDs