

Stochastic Decision Diagrams

John Hooker
Carnegie Mellon University

CPAIOR 2022
Los Angeles, USA

Motivation


- Decision diagrams (DDs) have proved useful for solving **discrete optimization** problems.
 - Especially those with recursive **dynamic programming** (DP) models.
- Yet many (most) DP models are **stochastic**.
 - We therefore generalize DDs to **stochastic DDs** (SDDs) by adding probabilities to arcs.

Motivation

- It's **no big deal** to put probabilities in a DD.
 - So why is this **interesting**?
- One reason:
 - It allows us to **extend DD-based relaxation techniques** to **stochastic** DP models.
 - Relaxation is **essential** to solving hard problems **to optimality**.

Review: Deterministic DDs

- A deterministic (binary) DD is a graphical representation of a **Boolean function**.
 - Often used for logic circuit design, product configuration.
 - Bounded-width DDs can represent **exponentially many** solutions.
 - **Same principle** lies behind “deep learning” and DP!



Review: Deterministic DDs


- A deterministic (binary) DD is a graphical representation of a **Boolean function**.
 - Often used for logic circuit design, product configuration.
 - Bounded-width DDs can represent **exponentially many** solutions.
 - **Same principle** lies behind “deep learning” and DP!
- A **weighted DD** can represent a discrete **optimization problem**.

Hadzic and JH (2006, 2007)

 - For example, the **maximum clique** problem...

Maximum clique

Max clique example



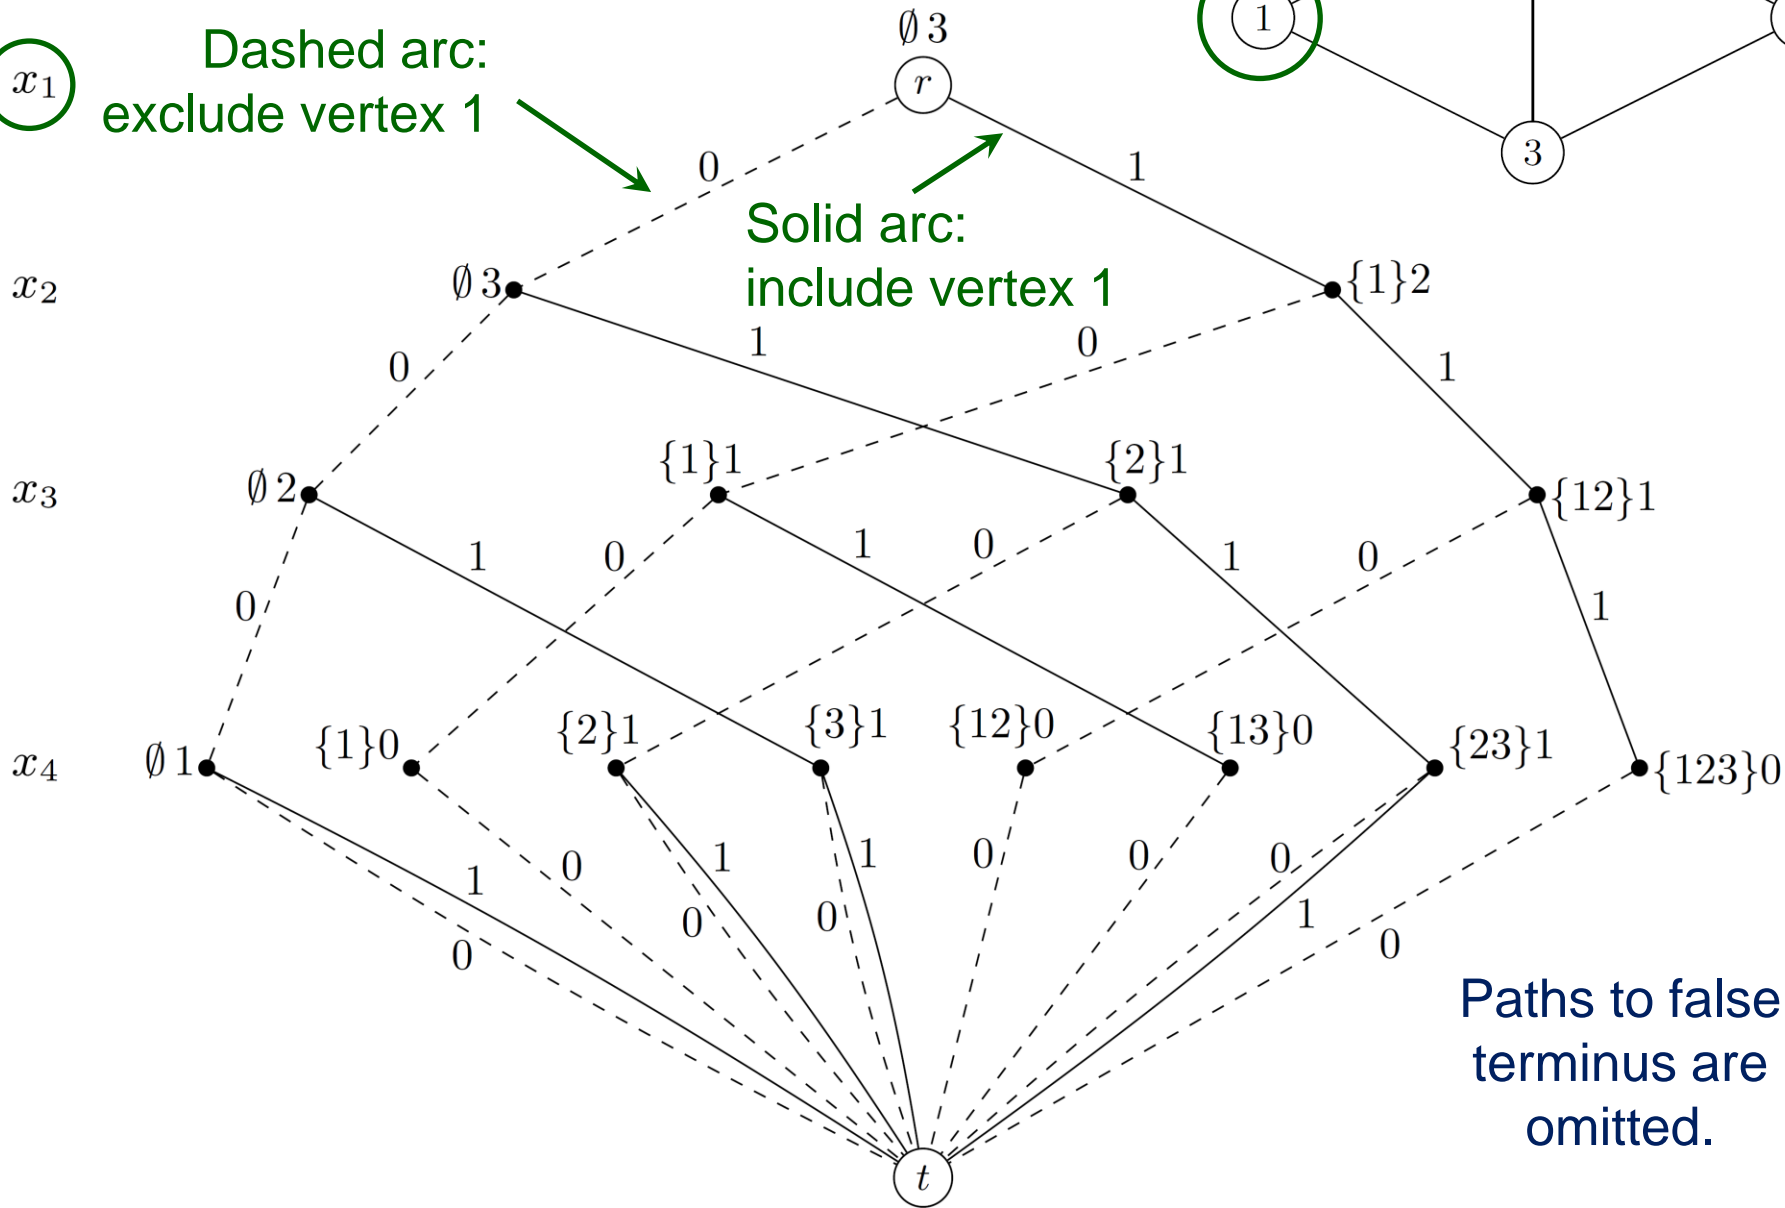

Let binary variable $x_i = 1$
when vertex i is in the clique.

Weighted DD



Dashed arc:
exclude vertex 1

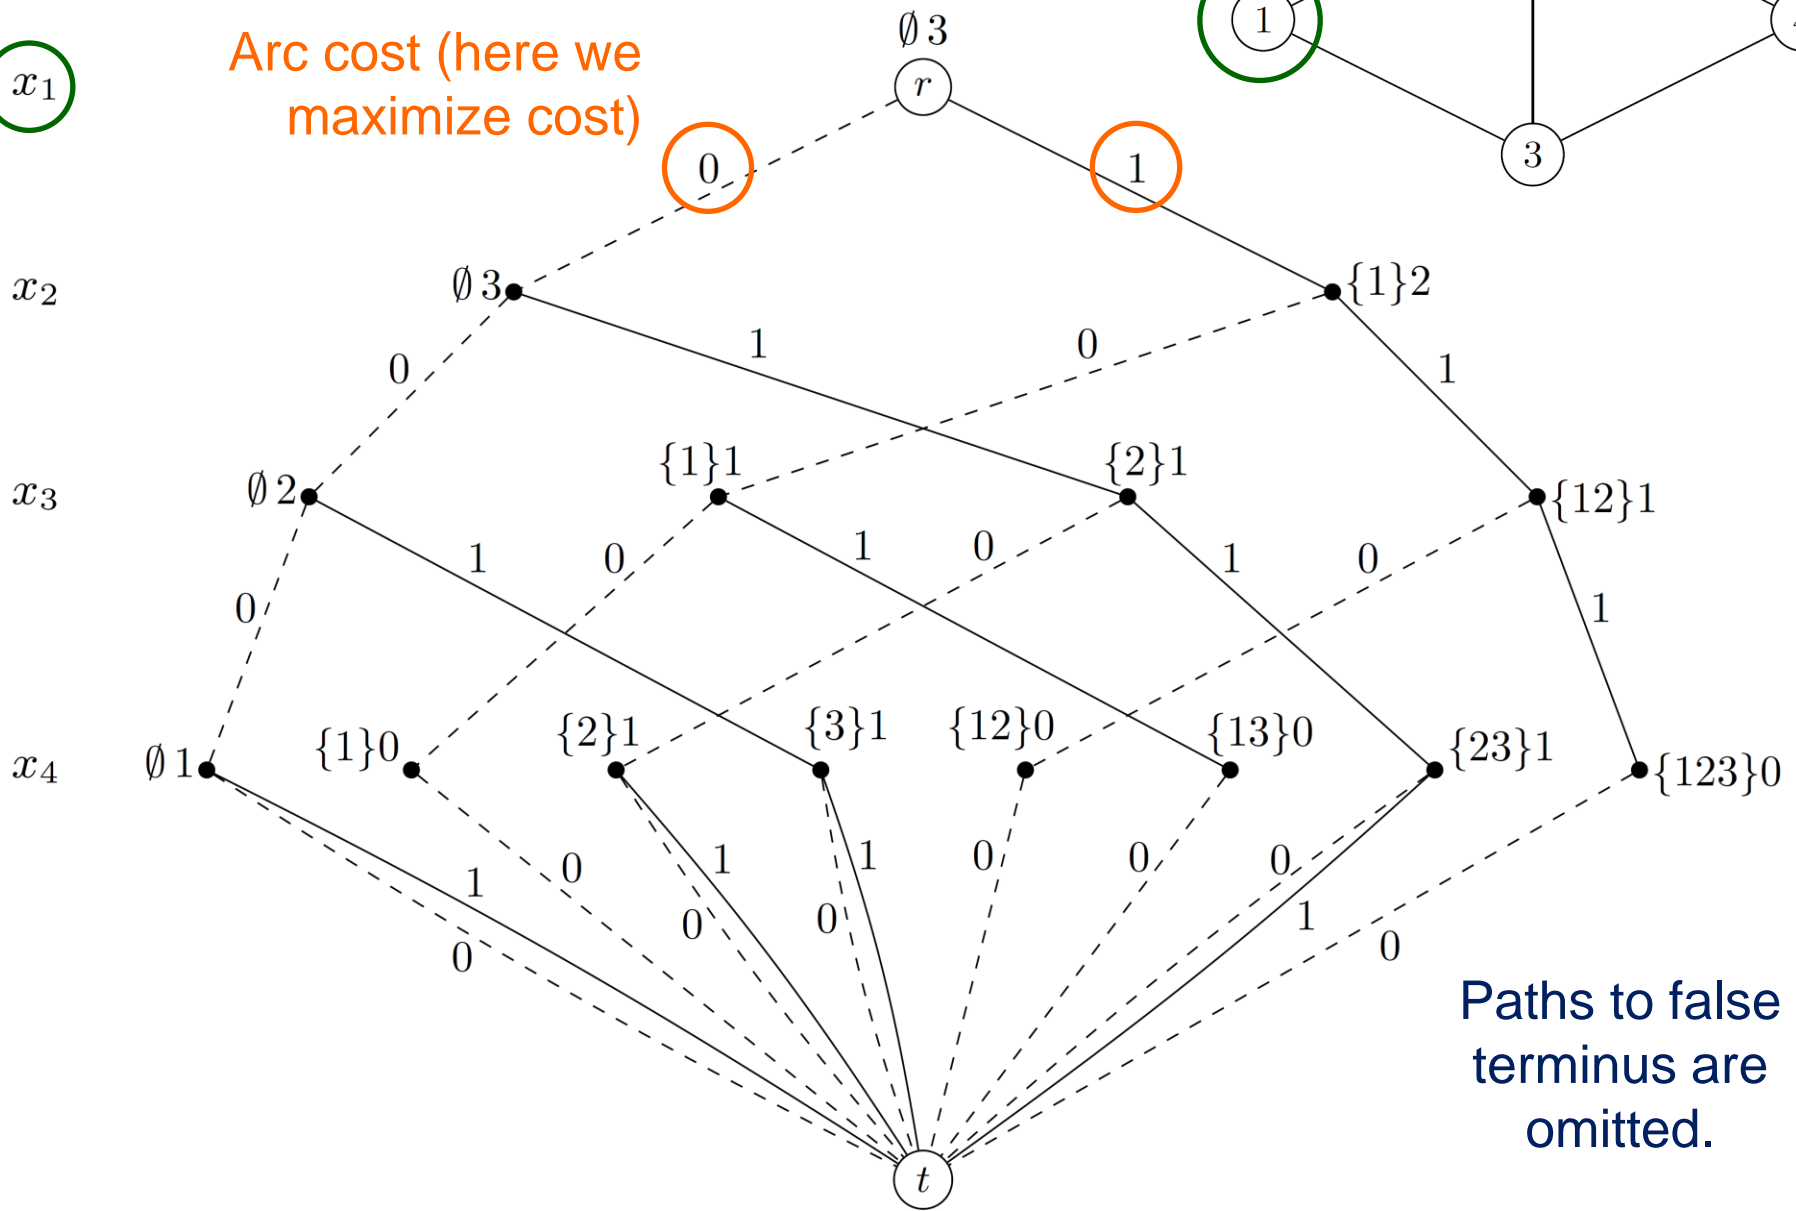

Solid arc:
include vertex 1



Paths to false terminus are omitted.

Weighted DD

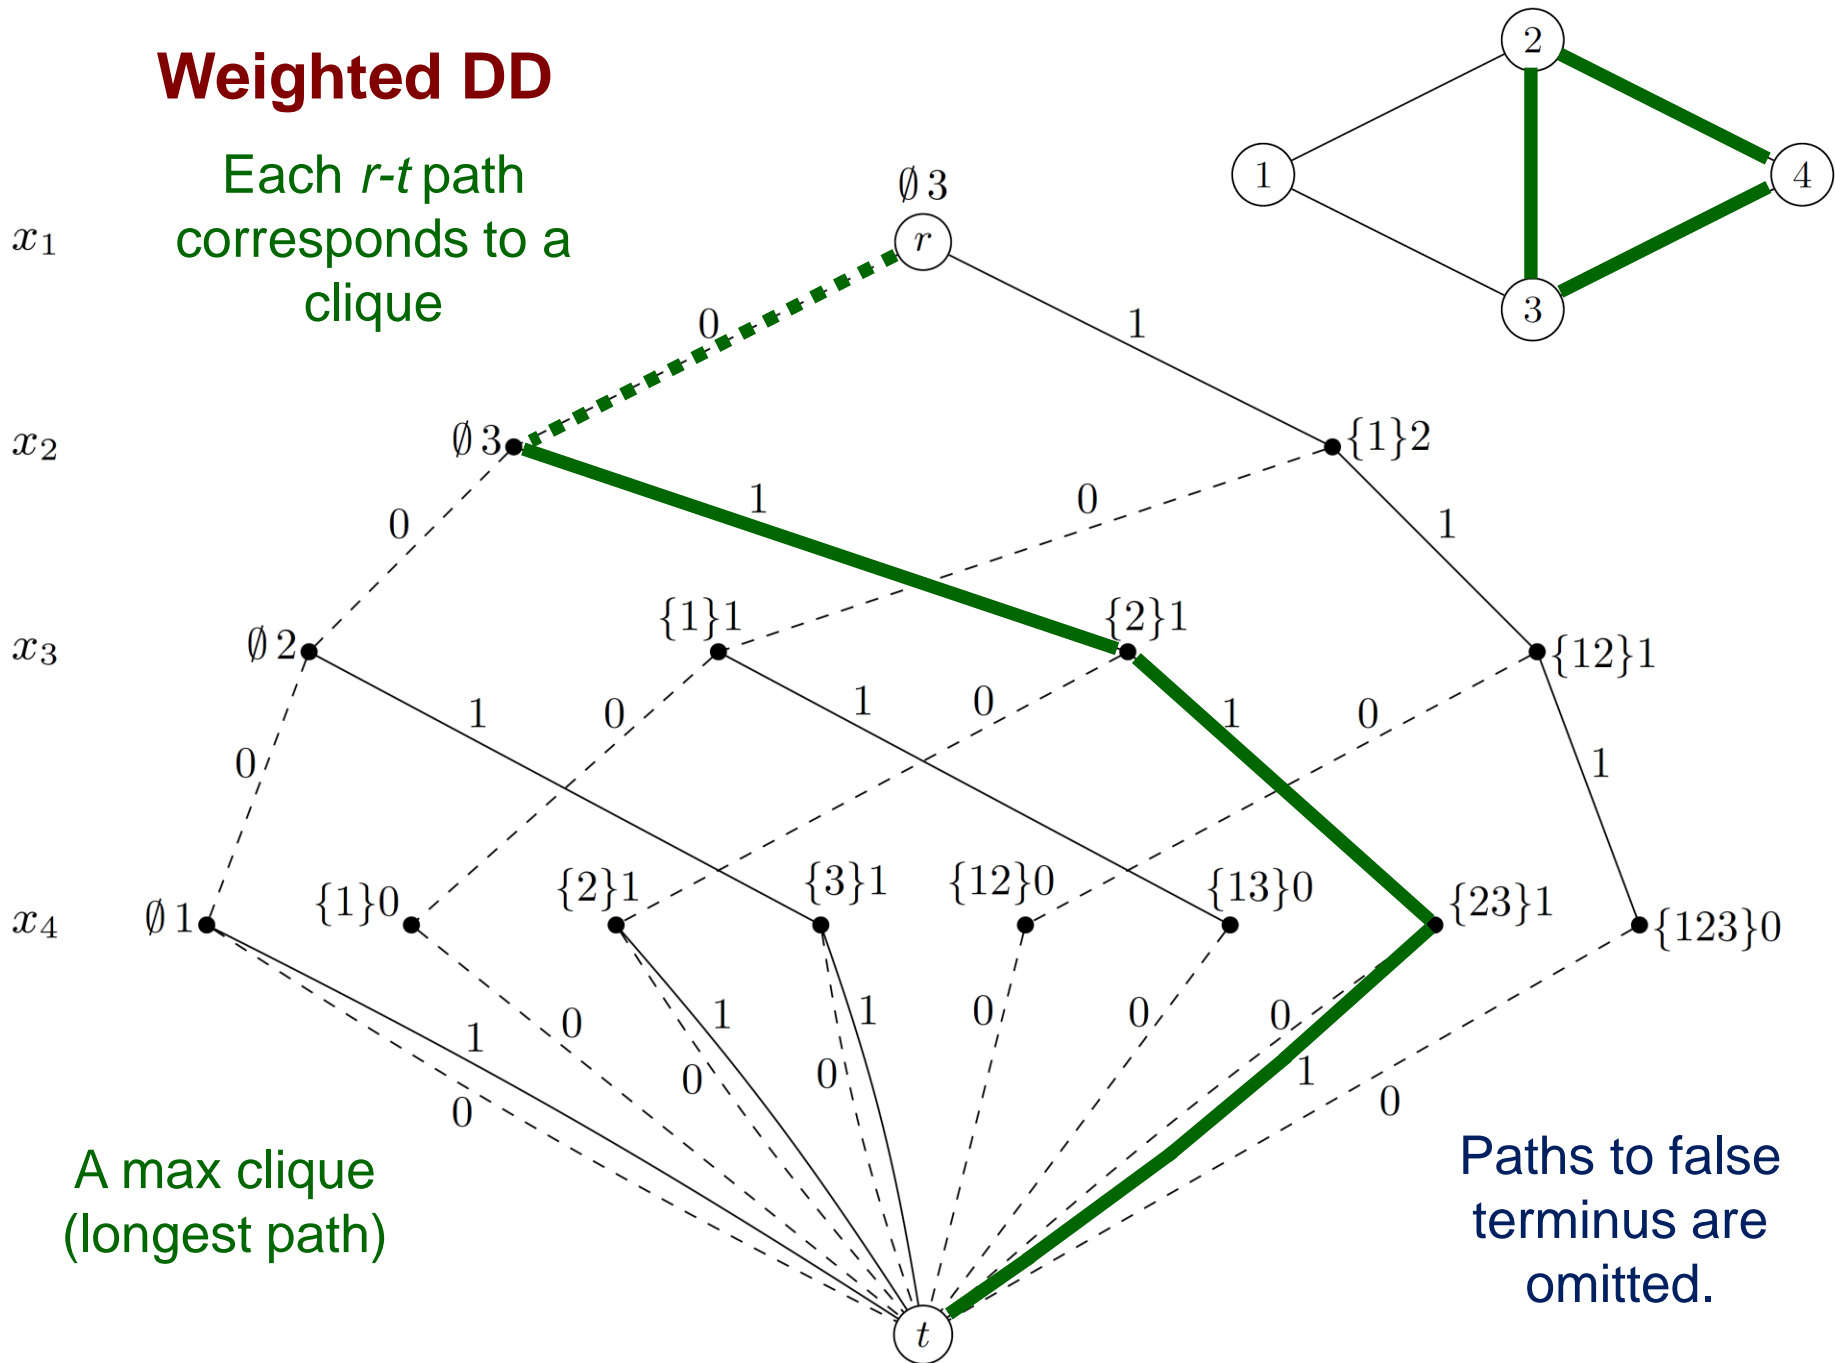
Arc cost (here we maximize cost)



Paths to false terminus are omitted.

Weighted DD

Each r - t path corresponds to a clique




A max clique (longest path)

Paths to false terminus are omitted.

Weighted DD

Each r - t path corresponds to a clique



Another max clique

Paths to false terminus are omitted.

Dynamic Programming

- The **state transition graph** of a dynamic programming (DP) problem can be interpreted as a DD.
 - By associating states with nodes of the DD.
 - This opens the door to using DD **relaxation** techniques to obtain bounds for DPs.

Andersen, Hadzic, JH, Tiedemann (2007)
Bergman, Cire, van Hoeve, JH (2013)

- ...and to solving the DPs by **branch and bound**.

Bergman, Cire, van Hoeve, JH (2014)

- For example, the **maximum clique** problem...

Deterministic DDs

Max clique DP model

State variable $S_i = \{\text{vertices selected so far in stage } i\}$.

The recursion is

$$h_i(S_i) = \max \left\{ \underset{\substack{\uparrow \\ x_j = 0}}{h_{i+1}(S_i)}, \quad 1 + \underset{\substack{\uparrow \\ x_j = 1}}{h_{i+1}(S_i \cup \{i\})} \right\}, \quad i = n, \dots, 1$$

cost to go

(max number of vertices that can be added to the clique, given that the vertices in S_i have been added so far)

Deterministic DDs

Max clique DP model

State variable $S_i = \{\text{vertices selected so far in stage } i\}$.


The recursion is

$$h_i(S_i) = \max \left\{ \underset{\substack{\uparrow \\ x_j = 0}}{h_{i+1}(S_i)}, \quad 1 + \underset{\substack{\uparrow \\ x_j = 1}}{h_{i+1}(S_i \cup \{i\})} \right\}, \quad i = n, \dots, 1$$

In general,

$$h_i(\mathbf{S}_i) = \min_{\substack{\uparrow \\ \text{control}}} \left\{ \underset{\substack{\uparrow \\ \text{immediate} \\ \text{cost}}}{c_i(\mathbf{S}_i, x_i)} + \underset{\substack{\uparrow \\ \text{cost to go}}}{h_{i+1}(\underset{\substack{\uparrow \\ \text{state transition function}}}{\phi_i(\mathbf{S}_i, x_i)})} \right\}, \quad i = n, \dots, 1$$

Weighted DD



Deterministic MDDs

Job sequencing example

Let control x_i be i th job in sequence.


Time window $[r_i, d_i]$ and processing time t_i for each job i .

Minimize total tardiness.

Weighted DD

State = (S_i, f_i) , where

$S_i = \{\text{jobs sequenced so far}\}$, $f_i = \text{finish time of previous job}$



Deterministic DDs

Job sequencing DP model

State is (S_i, f_i) .

The recursion is

$$h_i(S_i, f_i) = \min_{x_i \notin S_i} \left\{ c_i(S_i, f_i) + h_{i+1}(\phi_i((S_i, f_i), x_i)) \right\}$$

where

$$c_i((S_i, f_i), x_i) = \max \{ 0, \max\{r_{x_i}, f_i\} + t_{x_i} - d_{x_i} \}$$

$$\phi_i((S_i, f_i), x_i) = (S_i \cup \{x_i\}, \max\{r_{x_i}, f_i\} + t_{x_i})$$

Dynamic Programming

- Note: a state transition graph is a **different concept** than a DD.
 - A DD does not need **state information**.

Dynamic Programming


- Note: a state transition graph is a **different concept** than a DD.
 - A DD does not need **state information**.
- The DD perspective yields advantages:
 - A DD can be often be **reduced** by identifying isomorphic portions of the DD that are associated with different states. Bryant (1986 etc.)
 - This occasionally results in **radical** simplification (e.g., inventory management). JH (2013)
 - DP can benefit from **relaxation techniques** that have been developed for DDs...

Relaxed DDs

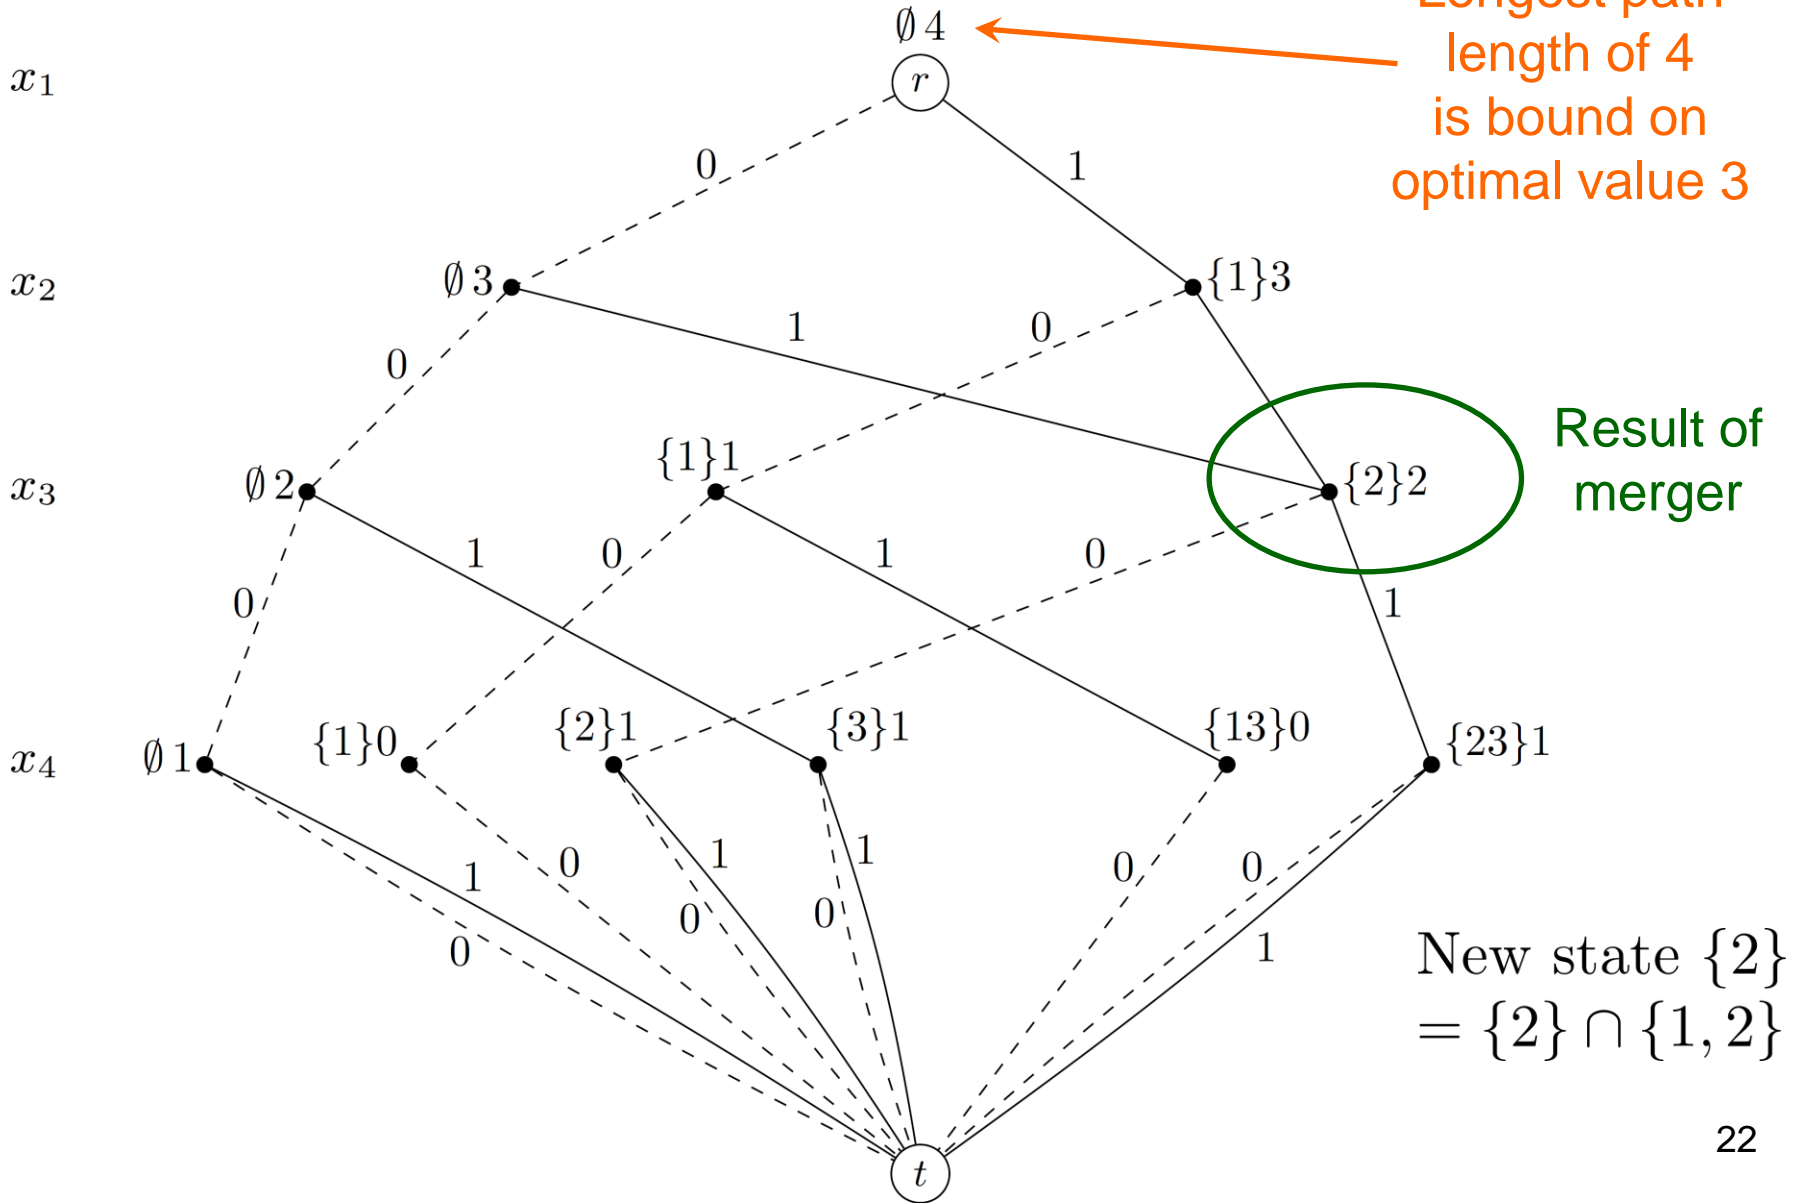
- A **relaxed** DD represents a superset of feasible solutions.
 - Can provide a **bound** on the optimal value.
- Created during top-down compilation of the DD.
 - By **node merger** or **node splitting**.
 - We focus on **node merger**.
 - Mergers result in **smaller DD** but weaker bound.
 - Can obtain bound of **any desired quality** by controlling width of relaxed DD.

Andersen, Hadzic, JH, Tiedemann (2007)
Bergman, Cire, van Hoeve, JH (2013)


Weighted DD for max clique



Relaxed DD for max clique




Weighted DD for job sequencing



State merger rule: $(S_i, f_i) \oplus (S'_i, f'_i) = ((S_i \cap S'_i), \min\{f_i, f'_i\})$

Relaxed DD for job sequencing



$$\text{New state } (\{2\}, 5) = (\{1, 2\} \cap \{2, 3\}, \min\{6, 5\})$$

Traditional state space relaxation

- Requires creation of alternate (smaller) state space for every problem.

Christofides, Mingozzi, Toth (1981)
Baldacci, Mingozzi, Roberti (2012)

- General practice is to use **approximate DP** instead.

Powell (2011)

Traditional state space relaxation

- Advantages of DD-based relaxation.
 - Uses **same state variables** as original problem.
 - This allows DD-based **branch-and-bound** method to solve problem. Bergman, Cire, van Hoesve, JH (2014)
 - Relaxation constructed **dynamically**.
 - Can be **tightened** by filtering, Lagrangian relaxation. Bergman, Cire, van Hoesve (2015); JH (2017, 2019)
 - Can be sized to provide bound of **any desired quality**.

Stochastic DDs

- A stochastic decision diagram (SDD) has **probabilistic** transitions to the next layer.
 - A control can have several possible **outcomes**, each with a known probability.
 - The outcome determines which arc is followed.


Stochastic DDs

- A stochastic decision diagram (SDD) has **probabilistic** transitions to the next layer.
 - A control can have several possible **outcomes**, each with a known probability.
 - The outcome determines which arc is followed.
- A solution is now a **policy**.
 - The control in a given layer depends on the **state** (node).
 - We **learn** from previous decisions
 - Original occurrence of learning in optimization?

Stochastic DDs

Max clique example

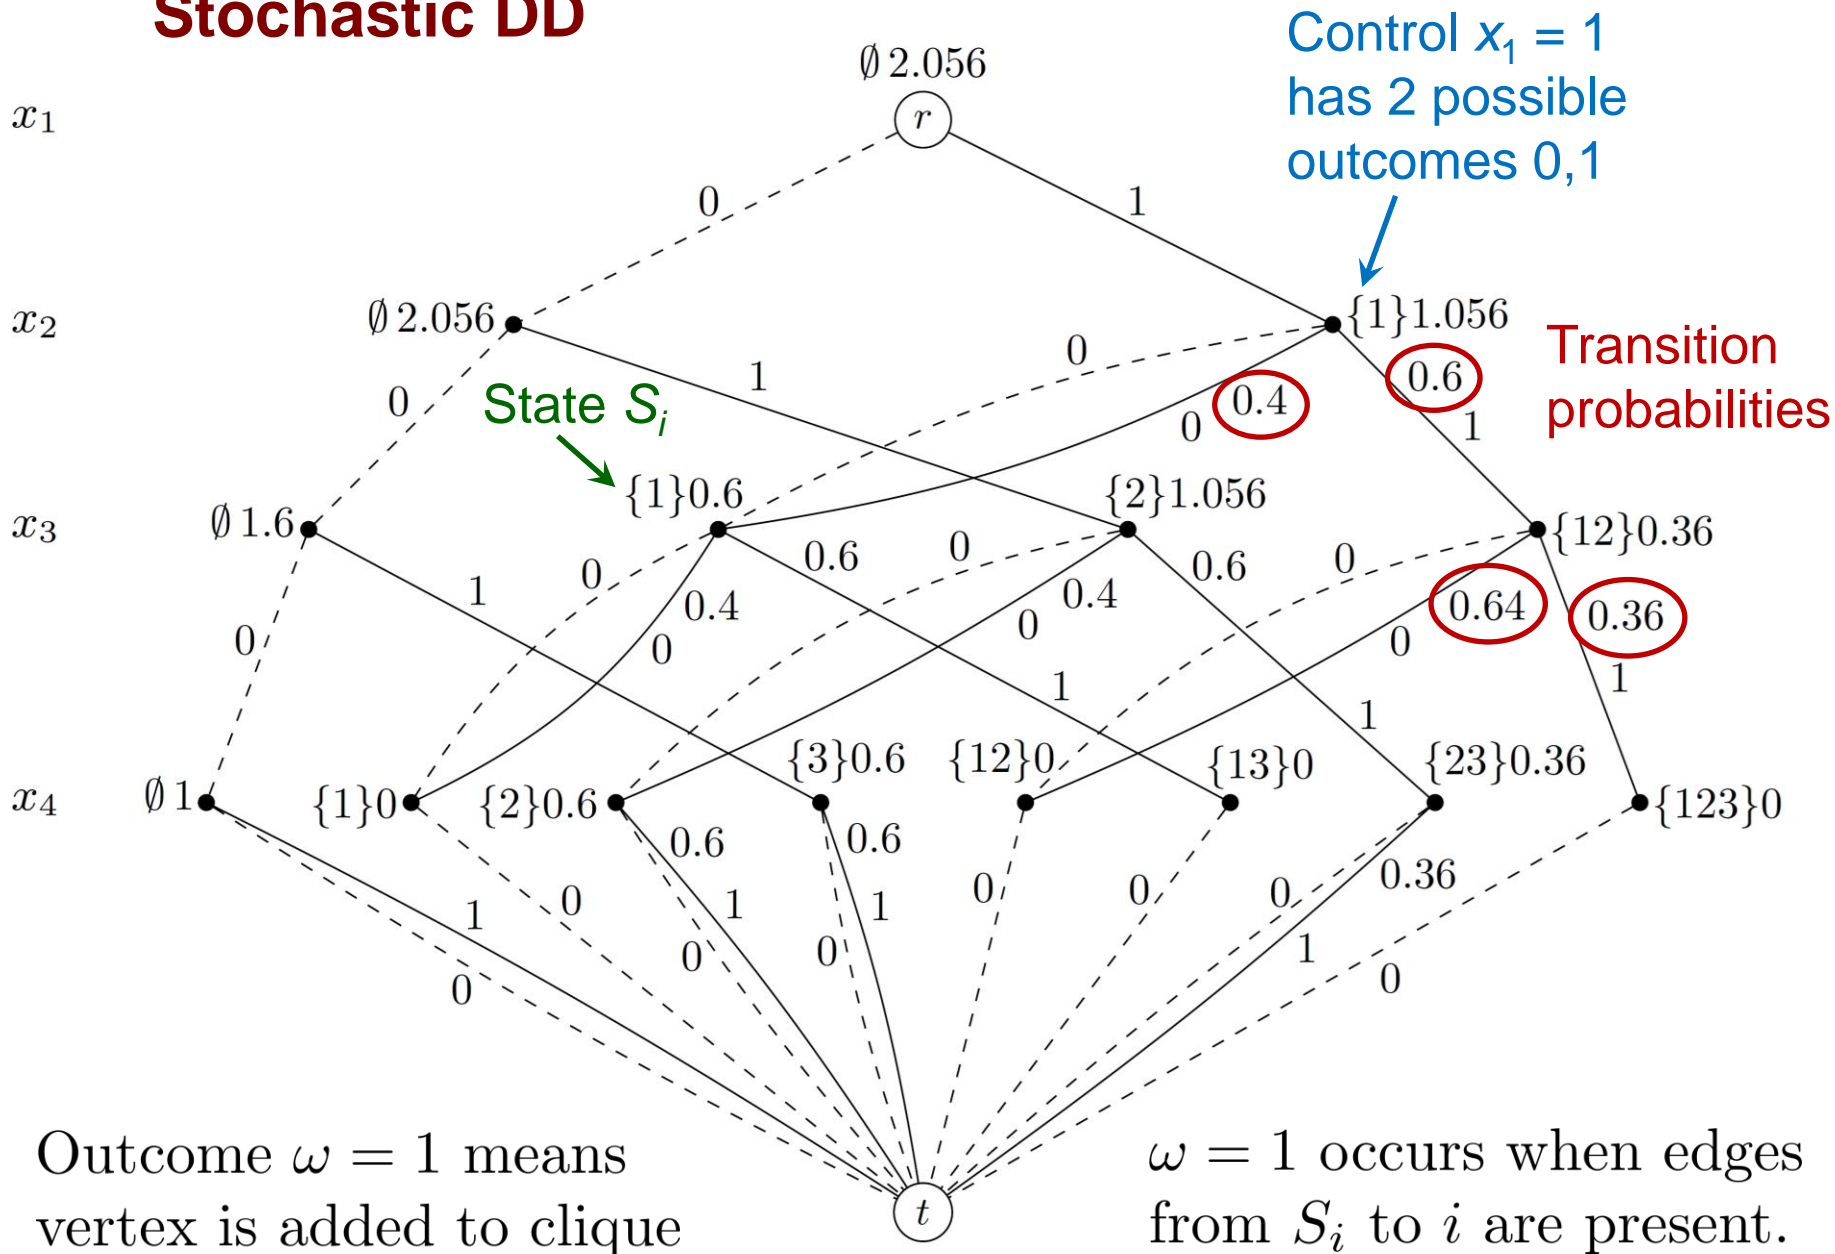
Defined on a
random graph



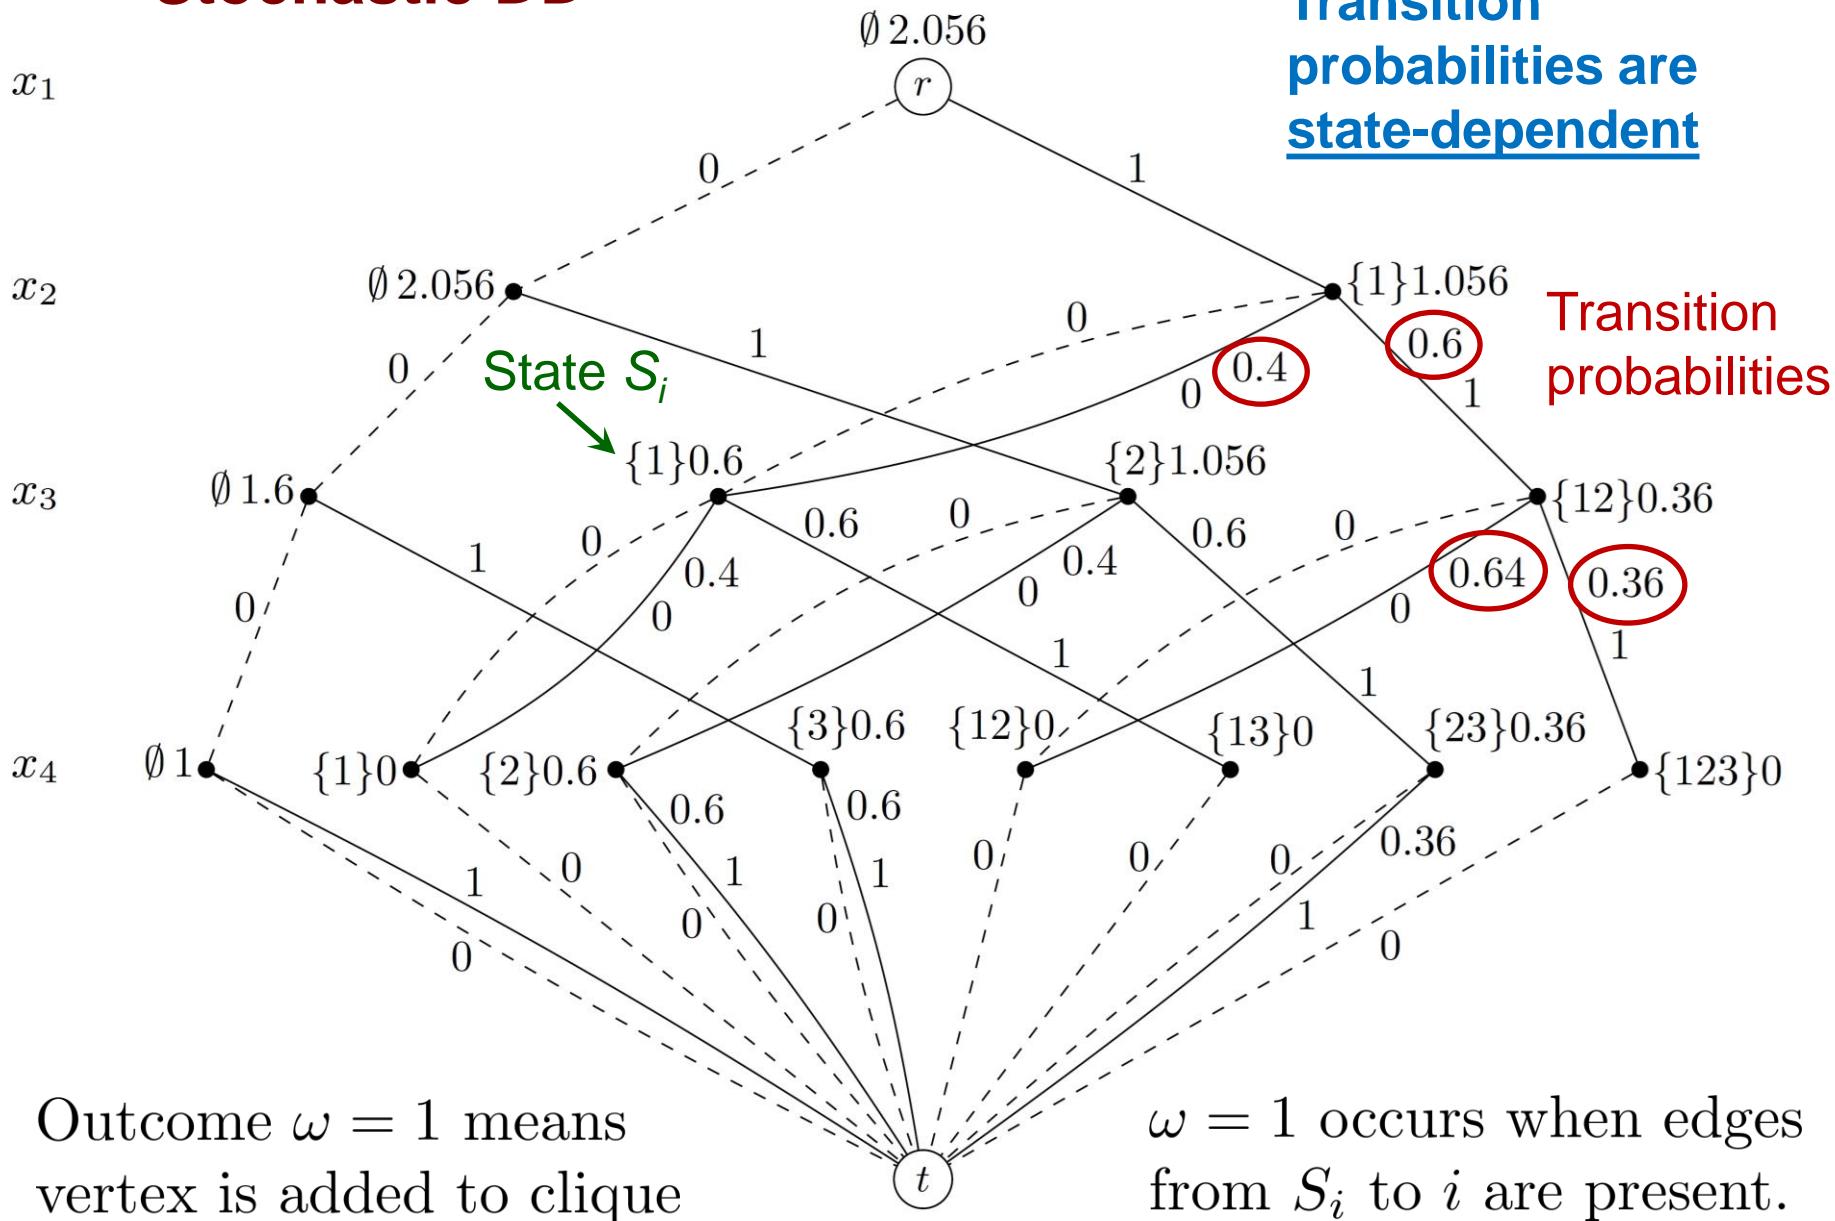
Each arc has probability 0.6

Objective is to maximize expected clique size.

Stochastic DD



Stochastic DD



Stochastic DDs

Max clique DP model

The recursion is

$$h_i(S_i) = \max \left\{ \underset{\substack{\uparrow \\ x_j = 0}}{h_{i+1}(S_i)}, (1 - p(S_i)) \underset{\substack{\uparrow \\ x_j = 1}}{h_{i+1}(S_i)} + p(S_i) h_{i+1}(S_i \cup \{i\}) \right\}$$

where $p(S_i) =$ probability that vertex i can be added to clique

Stochastic DDs

Max clique DP model

The recursion is

$$h_i(\mathcal{S}_i) = \max \left\{ \underset{\substack{\uparrow \\ x_j = 0}}{h_{i+1}(\mathcal{S}_i)}, (1 - p(\mathcal{S}_i)) \underset{\substack{\uparrow \\ x_j = 1}}{h_{i+1}(\mathcal{S}_i)} + p(\mathcal{S}_i) h_{i+1}(\mathcal{S}_i \cup \{i\}) \right\}$$

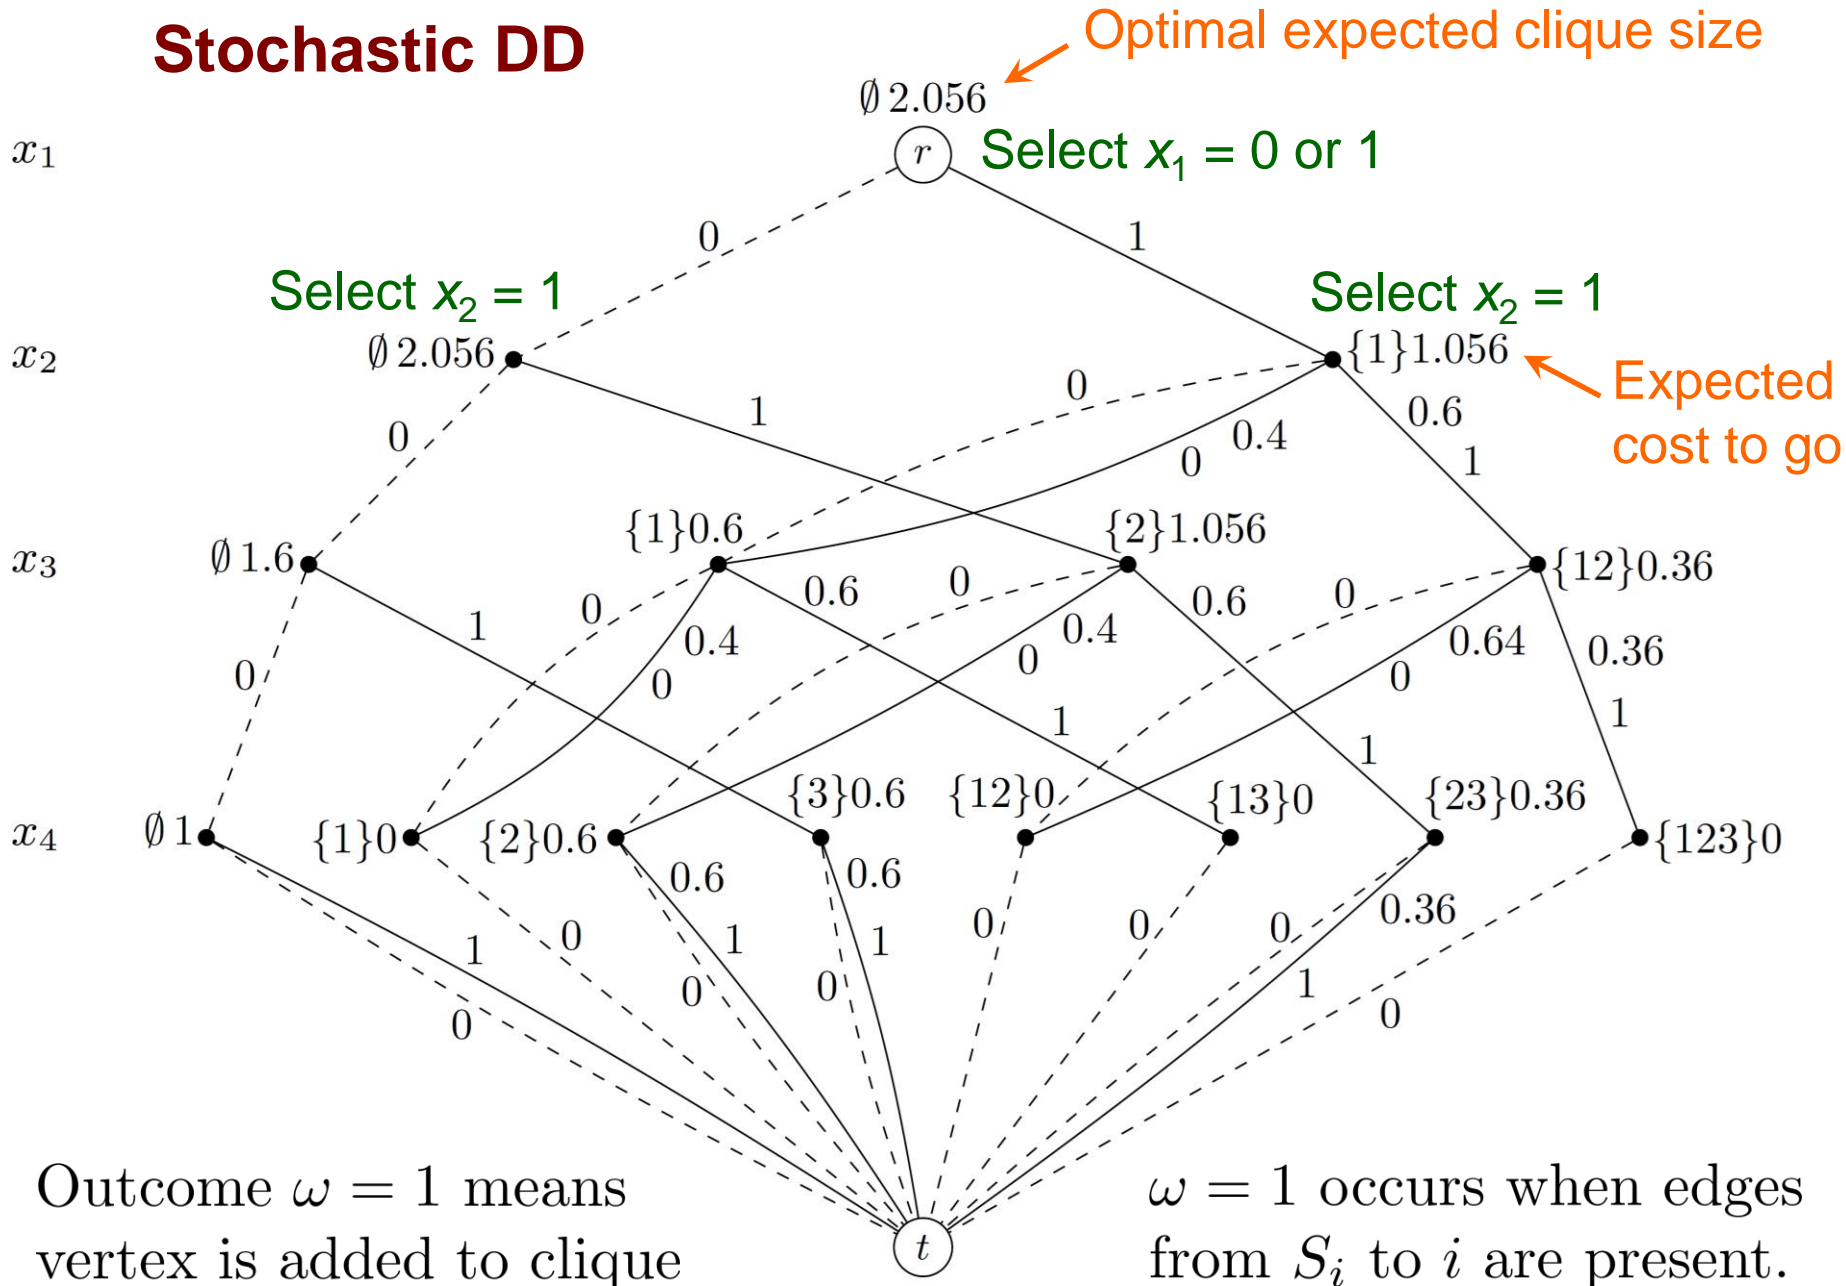
where $p(\mathcal{S}_i)$ = probability that vertex i can be added to clique

In general,

$$h_i(\mathcal{S}_i) = \min_{x_i} \left\{ \sum_{\omega} p_{i\omega}(\mathcal{S}_i, x_i) [c_{i\omega}(\mathcal{S}_i, x_i) + h_{i+1}(\phi_{i\omega}(\mathcal{S}_i, x_i))] \right\}$$

where $p_{i\omega}(\mathcal{S}_i, x_i)$ = prob. of outcome ω given control x_i in state \mathcal{S}_i
and similarly for $c_{i\omega}(\mathcal{S}_i, x_i)$ and $\phi_{i\omega}(\mathcal{S}_i, x_i)$

Stochastic DD



Relaxed SDDs


- A relaxed SDD is one that provides a **valid bound** on optimal expected cost.
 - **Unclear** how to define relaxation in terms of **individual solutions**.
 - ...since solutions are **policies** defined on the **entire SDD**, and a relaxed SDD may have very different structure.

Stochastic diagram \bar{D} relaxes diagram D when \bar{D} and D have the same variables, controls, and possible outcomes, and when optimal cost of $\bar{D} \leq$ optimal cost of D .

Relaxed SDDs


- We will relax SDDs by node merger.
 - We will also provide **sufficient conditions** under which a given merger operation yields a valid relaxed SDD.
 - Conditions must account for **policy-based** solutions rather than simple control sequences.
 - Examples...

Stochastic DD for max clique



Relaxed DD for max clique

Expected longest path length of 2.4736 is bound on optimal value 2.056



Result of merger

We will prove that this is a valid relaxed SDD

New state $\{2\} = \{2\} \cap \{1, 2\}$

Relaxed SDDs

Stochastic job sequencing DP model

Stochastic element is processing time (state independent).

Let $t_{j\omega}$ = job j processing time in outcome ω .

Let $p_{j\omega}$ probability of outcome ω for job j .

Transition probabilities are state-independent

The recursion is

$$h_i(S_i, f_i) = \min_{x_i \notin S_i} \left\{ \sum_{\omega} p_{i\omega}(S_i, x_i) [c_{i\omega}(S_i, f_i) + h_{i+1}(\phi_{i\omega}((S_i, f_i), x_i))] \right\}$$

where

$$c_{i\omega}((S_i, f_i), x_i) = \max \{ 0, \max\{r_{x_i}, f_i\} + t_{x_i\omega} - d_{x_i} \}$$

$$\phi_{i\omega}((S_i, f_i), x_i) = (S_i \cup \{x_i\}, \max\{r_{x_i}, f_i\} + t_{x_i\omega})$$

We can use the same node merger operation as before.

Node Merger in SDDs

We need a concept of **one state relaxing another**.

Relaxation must have the property that state $\bar{\mathbf{S}}_i$ relaxes state \mathbf{S}_i only if

$$p_{i\omega}(\bar{\mathbf{S}}_i, x_i) c_{i\omega}(\bar{\mathbf{S}}_i, x_i) \leq p_{i\omega}(\mathbf{S}_i, x_i) c_{i\omega}(\mathbf{S}_i, x_i)$$

for any control x_i and any outcome ω .

Node Merger in SDDs

We need a concept of **one state relaxing another**.

Relaxation must have the property that state \bar{S}_i relaxes state S_i only if

$$p_{i\omega}(\bar{S}_i, x_i) c_{i\omega}(\bar{S}_i, x_i) \leq p_{i\omega}(S_i, x_i) c_{i\omega}(S_i, x_i)$$

for any control x_i and any outcome ω .

Max clique problem: \bar{S}_i relaxes S_i when $\bar{S}_i \subseteq S_i$.

Job sequencing problem: (\bar{S}_i, \bar{f}_i) relaxes (S_i, f_i) when $\bar{S}_i \subseteq S_i$ and $\bar{f}_i \leq f_i$.

These definitions satisfy the property.

Node Merger in SDDs

Jointly sufficient conditions under which node merger yields a relaxed SDD:

(C1) State $\mathcal{S}_i \oplus \mathcal{S}'_i$ relaxes both \mathcal{S}_i and \mathcal{S}'_i .

(C2) If state $\bar{\mathcal{S}}_i$ relaxes state \mathcal{S}_i , then $\phi_{i\omega}(\bar{\mathcal{S}}_i, x_i)$ relaxes $\phi_{i\omega}(\mathcal{S}_i, x_i)$ for any ω, x_i .

Note: (C1) and (C2) are sufficient for deterministic DDs

Node Merger in SDDs

Jointly sufficient conditions under which node merger yields a relaxed SDD:

(C1) State $\mathcal{S}_i \oplus \mathcal{S}'_i$ relaxes both \mathcal{S}_i and \mathcal{S}'_i .

(C2) If state $\bar{\mathcal{S}}_i$ relaxes state \mathcal{S}_i , then $\phi_{i\omega}(\bar{\mathcal{S}}_i, x_i)$ relaxes $\phi_{i\omega}(\mathcal{S}_i, x_i)$ for any ω, x_i .

(C3) If state $\bar{\mathcal{S}}_i$ relaxes state \mathcal{S}_i , then given any control x_i and any set of numbers $\{\eta_\omega \mid \text{all } \omega\}$, there is a control \bar{x}_i such that

$$\sum_{\omega} p_{i\omega}(\bar{\mathcal{S}}_i, \bar{x}_i) (c_{i\omega}(\bar{\mathcal{S}}_i, \bar{x}_i) + \eta_\omega) \leq \sum_{\omega} p_{i\omega}(\mathcal{S}_i, x_i) (c_{i\omega}(\mathcal{S}_i, x_i) + \eta_\omega)$$

Note: (C1) and (C2) are sufficient for deterministic DDs

Node Merger in SDDs

- Key to proofs: work with **fully articulated SDDs**.
 - All states are represented, even those that are reached with zero probability.
 - Node merger becomes rearrangement of probabilities.

Lemma. If condition (C2) is satisfied, and $\bar{\mathbf{S}}_i$ relaxes \mathbf{S}_i , then cost to go of $\bar{\mathbf{S}}_i \leq$ cost to go of \mathbf{S}_i .

Proof by backward induction on layers.

Node Merger in SDDs

Theorem. If **(C1)-(C3)** are satisfied, then node merger yields a relaxed SDD.

Proof by forward induction on layers of partially compiled SDDs.

Node Merger in SDDs

Theorem. If **(C1)-(C3)** are satisfied, then node merger yields a relaxed SDD.

Proof by forward induction on layers of partially compiled SDDs.

Corollary. The **max clique** state merger operation yields a relaxed SDD.

The operation satisfies (C1)–(C3), but the proof is nontrivial due to the strength of condition (C3).

Node Merger in SDDs

Theorem. If probabilities are state-independent, then a merger operation that satisfies (C1) and (C2) alone yields a relaxed SDD.

Corollary. The **job sequencing** merger operation yields a relaxed SDD.

Probabilities are state-independent, and it is easy to show the merger operation satisfies (C1)–(C2).

Computational Experiments

- **Major barrier to computational testing:**
 - We **don't know** optimal solutions of nontrivial stochastic DDs.
 - We need **optimal** (or very good) solutions to judge the quality of bounds from DDs.

Computational Experiments

- **Approach 1:** Use deterministic **max clique** instances in DIMACS library...
 - ...and **add edge** probabilities.
- **Why?**
 - Relaxed DDs provide good bounds for the **deterministic** problem.
 - Better than **full cutting plane resources** of MIP.

Bergman, Cire, van Hoesve, JH (2013)

Computational Experiments

- **Approach 1:** Use deterministic **max clique** instances in DIMACS library...
 - ...and **add edge** probabilities.
- **Why?**
 - **Simple model** even with **state-dependent** transition probabilities.
 - Relaxed DDs provide good bounds for the **deterministic** problem (better than MIP).

Bergman, Cire, van Hoeve, JH (2013)

Computational Experiments

- **Approach 1:** Use deterministic **max clique** instances in DIMACS library...
 - ...and **add edge** probabilities.
- **Why?**
 - **Simple model** even with **state-dependent** transition probabilities.
 - Relaxed DDs provide good bounds for the **deterministic** problem (better than MIP).

Bergman, Cire, van Hoesve, JH (2013)

- **However...**
 - Can't compare with **optimal** solutions
 - 2 exceptions, 1 of which required **24 hours** to solve.

Computational Experiments

- **Approach 2: Random instances.**
 - **Sized** to be tractable and nontrivial.
 - Exact solutions found with **complete SDDs**, since state space enumeration is the only available method.

Computational Experiments


- **Merger heuristic...**
 - Standard method: Merge **less attractive** nodes first
 - ...as measured by path lengths to root node in deterministic problem.
- **Control size** of relaxed SDD by limiting **width**.
 - Width = max number of nodes in a layer.

Random instances


Solved to optimality

Bound vs time

Density 0.6




Density 0.7




Random instances

Solved to optimality

Density 0.4



Density 0.5




DIMACS instances

Instance	Vertices	Density	Instance	Vertices	Density
brock200_1	200	0.7417	hamming6-2	64	0.8906
cfat200-1	200	0.0767	johnson8-4-4	70	0.7571
cfat500-1	500	0.0357	keller4	171	0.6453
c125.9	125	0.8913	p_hat300-1	300	0.2430
DSJC500_5	500	0.5010	san200_0.7_1	200	0.6965
gen200_p0.9_44	200	0.8955	sanr_0.7	200	0.6934

2 DIMACS instances

Solved to optimality

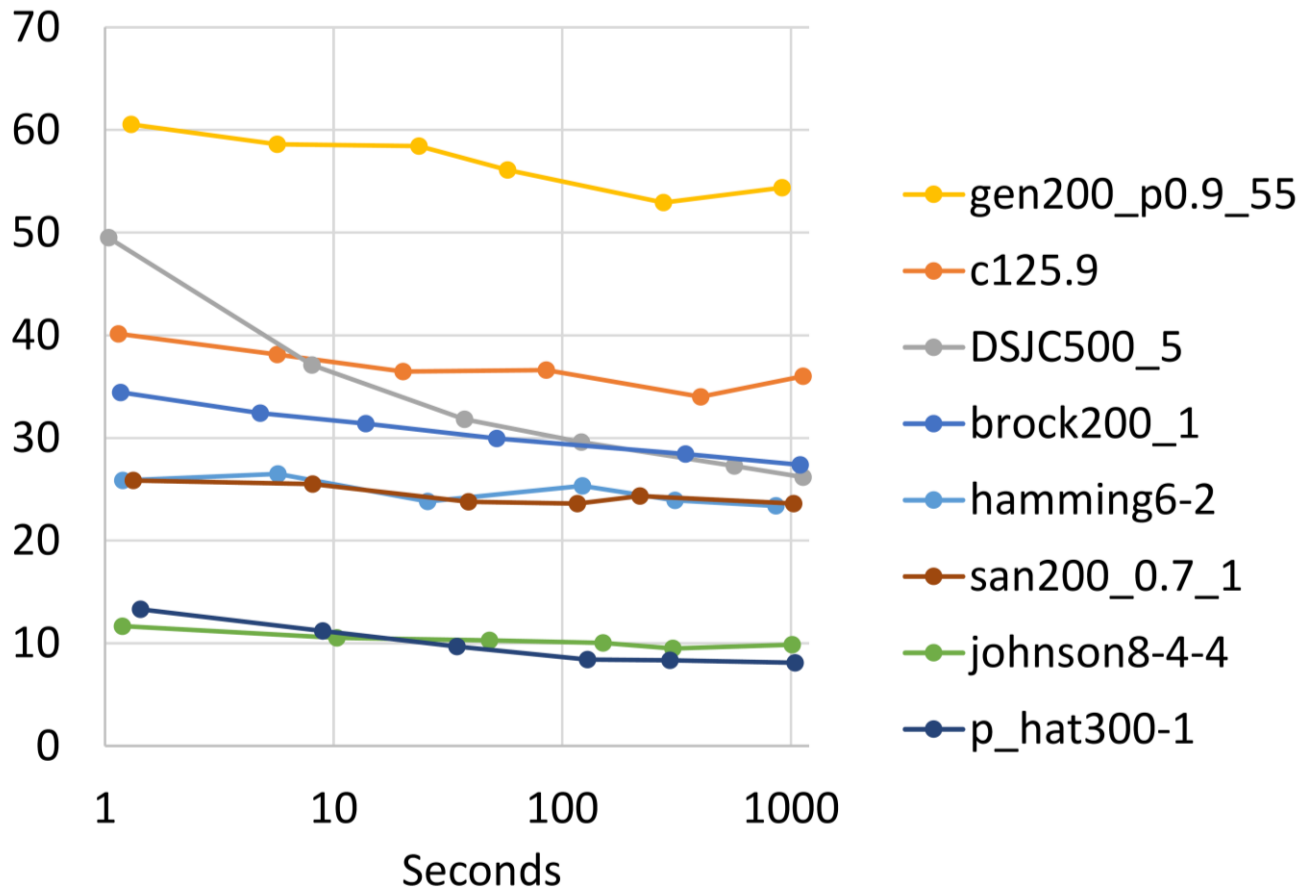
Bound vs time



DIMACS instances

Not solved to optimality

Bound vs time



Computational Experiments

- Bound quality degrades **gradually** with reduction in SDD width/time investment.
 - Even reduction down to a **few seconds**.
 - Indicates that SDDs can provide **useful bounds** for DP models.
- Roughly **logarithmic** relationship.
 - In most cases.
 - May allow estimate of how bound will **improve** with greater time investment.

Research Issues

- Use SDD bounds to solve moderate-sized problems by **branch and bound**.
 - Based on previous experience with deterministic problems.
- Use relaxed SDDs to compute bounds for **approximate DP**.
 - Find solution with traditional approximate DP, which **estimates** costs to go.
 - Use relaxed SDDs to compute **bounds** on costs to go, using same controls as in approximate DP solution.