Markov Chain-based Policies for Multi-stage Stochastic Integer Linear Programming

with an Application to Disaster Relief Logistics

Joint work with

Margarita Castro (PUC-Chile) and Yongjia Song (Clemson)

ICERM

February 2023

Computational Optimization Lab

Motivating App. Hurricane Disaster Relief Planning

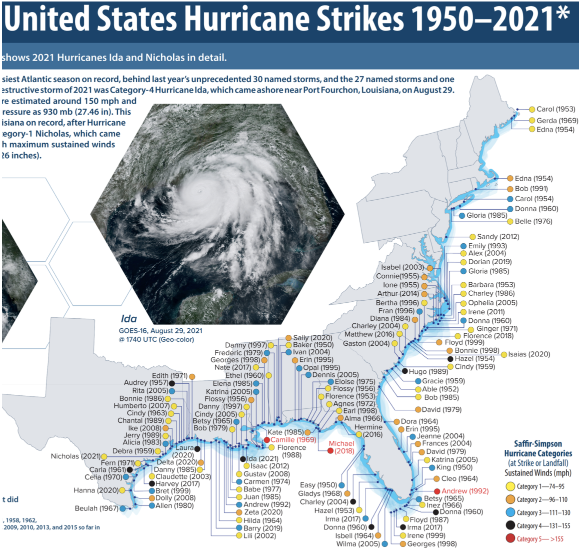

(Source: Wikipedia)

Motivating App. Hurricane Disaster Relief Planning

- Hurricane Katrina (2005):
  - $>$ 1200 fatalities
  - $>$ \$190 billion damage
- Hurricane Ian (2022):
  - $>$ 160 fatalities
  - $>$ \$113 billion damage

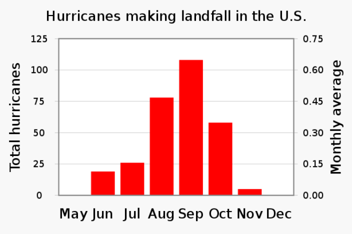

(Source: Wikipedia)

- Hurricanes can be detected a few days before their landfall
- $\approx$ 5 days before landfall, National Hurricane Center provides info about
  - hurricane's predicted trajectory, speed, intensity
  - endangered areas

Computational Optimization Lab

Motivating App. Hurricane Disaster Relief Planning

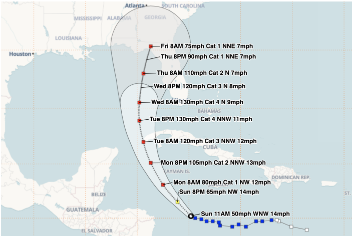

Predictions are leveraged by humanitarian and governmental agencies to prepare and allocate hurricane relief supplies

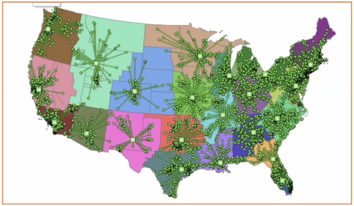

When, where, and how to preposition relief supplies ahead of an impeding hurricane?

How to distribute supplies to affected population in an efficient way?

Computational Optimization Lab

Motivating App. Hurricane Disaster Relief Planning

Contingency modality activation:

- Not all types of operations are fully adaptive
- Some operations are “all or nothing” type decisions, e.g., contingency modality activation for expanding DC capacities

Computational Optimization Lab

Talk is about: Multi-stage SP & A bit of two-stage SP

First need to move some key players to a different stage!

Sequential Decision-making Under Uncertainty

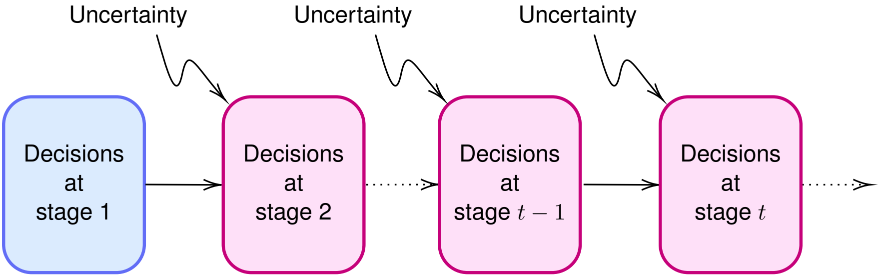

- Uncertainty is gradually observed
- Decisions are dynamically adapted to:
  - Observed uncertainty
  - Previous decisions

Multi-stage stochastic programs

Finite-horizon sequential decision-making problems under uncertainty

- $T \geq 2$ decision stages
- Stochastic process: $\{\boldsymbol{\xi}_t\}_{t \in [T]}$ $\qquad [T] = \{1, \ldots, T\}$
- History: $\boldsymbol{\xi}^t := (\xi_1, \ldots, \boldsymbol{\xi}_t)$
- Dynamics:

$\boxed{t = 1}$ $\cdots$ $\boxed{\text{Stage } t-1 \text{ decisions}}$ $\xrightarrow[\boldsymbol{\xi}_t]{\text{observe}}$ $\boxed{\text{Stage } t \text{ decisions}}$ $\cdots$ $\boxed{t = T}$

y_1, x_1 $\qquad$ y_{t-1}, x_{t-1} $(\boldsymbol{\xi}^{t-1})$ $\dashrightarrow$ y_t, x_t $(\boldsymbol{\xi}^t)$

- Decision variables: (nonanticipative)
  - State variables: $x_t(\boldsymbol{\xi}^t)$
  - Recourse (stage) variables: $y_t(\boldsymbol{\xi}^t)$
- For convenience: $\xi_1 = 1$

Our Setup

Commonly used scenario-tree based approach

Usually an exponentially large tree

Uncertainty model:

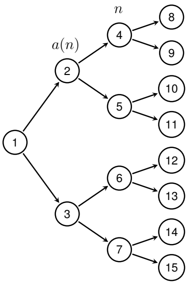

- Governed by a Markov chain (MC)

MSILP model:

$$\min \sum_{n \in \mathcal{N}} p_n f_n(x_n, z_n, y_n)$$

s.t. $\forall n \in \mathcal{N}$:

$$(x_n, z_n, y_n) \in \mathcal{X}_n(x_{a(n)}, z_{a(n)})$$

$y_n \in \mathbb{R}^m \rightarrow$ cont. local variables

$x_n \in \mathbb{R}^r \rightarrow$ cont. state variables

$z_n \in \mathbb{Z}^\ell \rightarrow$ int. state variables

- Linear objective and constraints

Motivating Example

Hurricane Disaster Relief Planning

- Produce & distribute resources from distribution centers to shelters
- **Decision stages:** When hurricane originates to its landing
- **Uncertainty:**
  - Demand is a function of the hurricane's state
  - Evolution of the hurricane is modeled by a MC
  - MC states: Intensity + location
- **Decisions:**
  - Local variables: Production, distribution, unsatisfied demand
  - Continuous state variables: Inventory and capacity
  - Integer state variables: Contingency modality activation (to increase DC capacities)

Computational Optimization Lab

Existing Methodologies: Purely continuous case

- General case: Nested Benders decomposition [Birge, 1985]

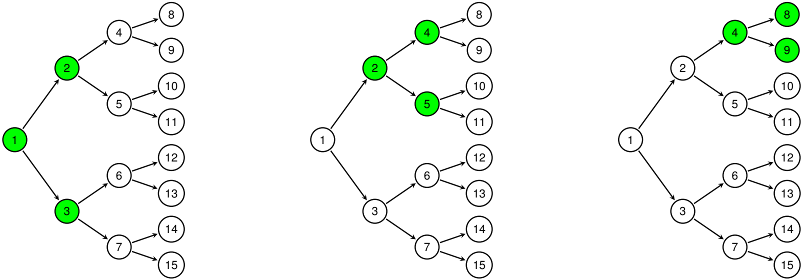

Via Benders cuts, approximate the expected cost-to-go functions:

$$Q_n(x_{a(n)}) = \min_{(x_n, y_n)} f_n(x_n, y_n) + \sum_{m \in \mathcal{C}(n)} \bar{p}_{nm} Q_m(x_n)$$

Existing Methodologies: Purely continuous case

- Stage-wise independent case: SDDP [Pereira and Pinto, 1991]
  - Each stage has its own independent set of realizations
  - Can recombine the scenario tree:

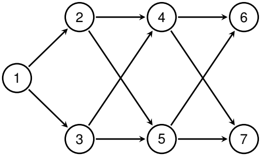

  - One expected cost-to-go function per stage instead!
  - Many fewer nodes!

Computational Optimization Lab

Shabbir's SDDP Illustration

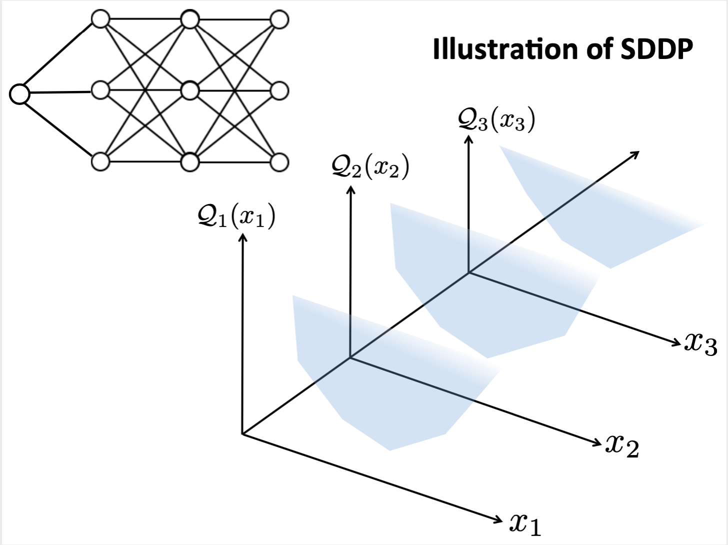

Shabbir's SDDP Illustration

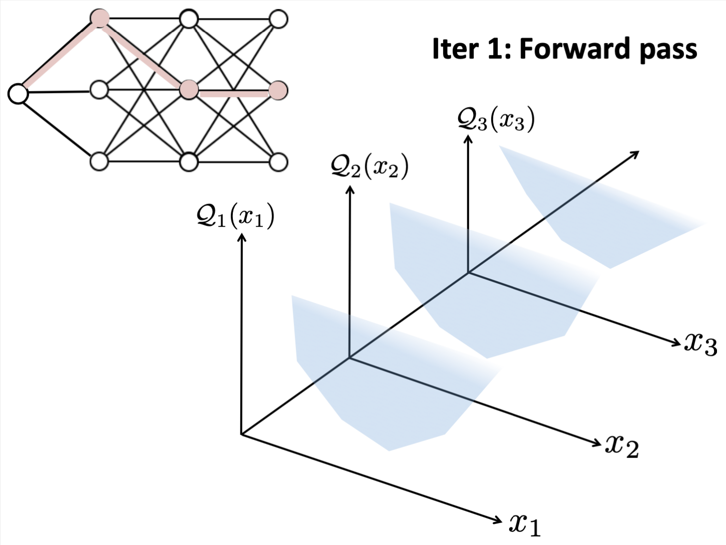

Shabbir's SDDP Illustration

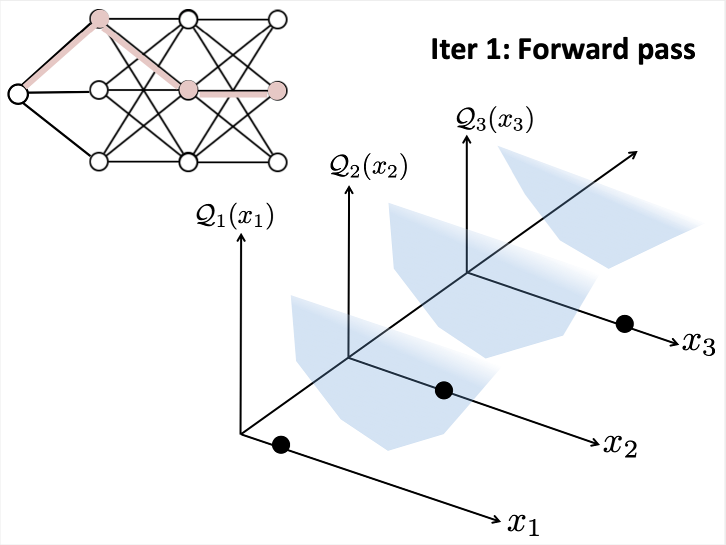

Shabbir's SDDP Illustration

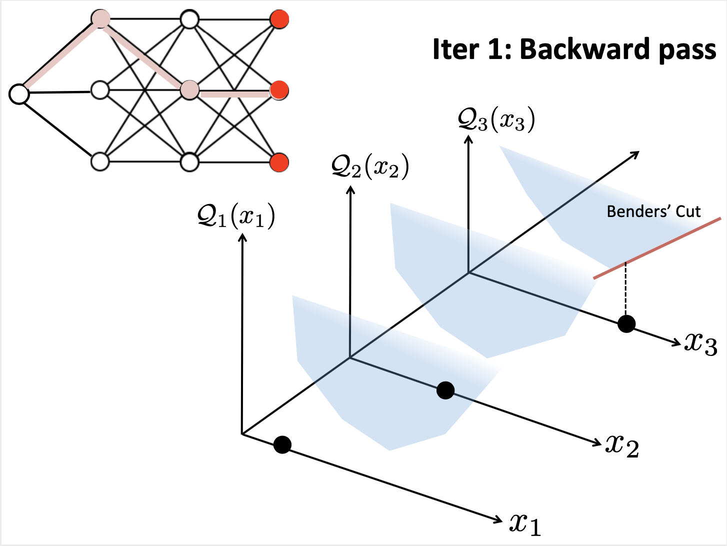

Shabbir's SDDP Illustration

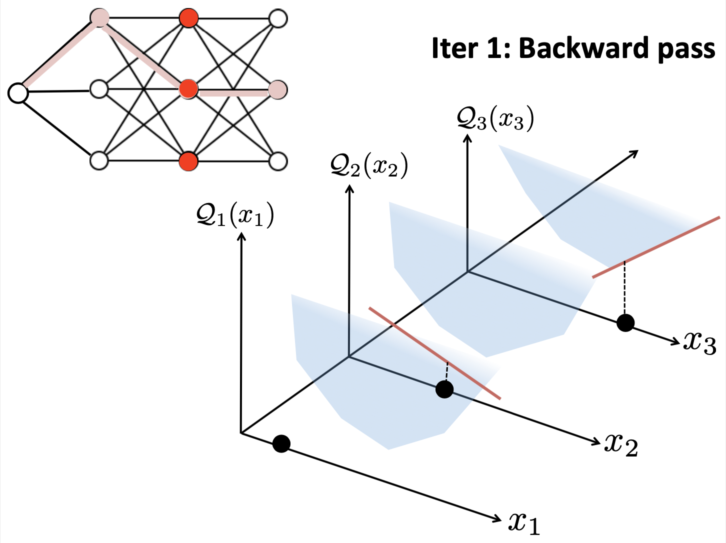

Shabbir's SDDP Illustration

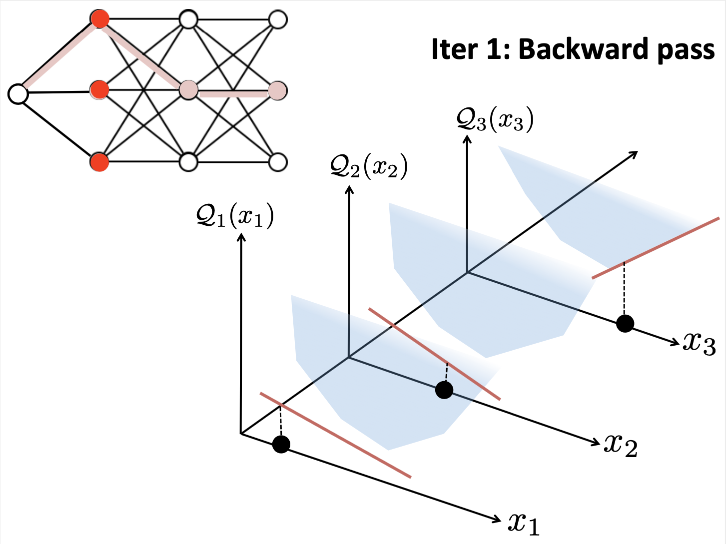

Existing Methodologies

- Purely continuous:
  - General: Nested Benders [Birge, 1985]
  - Stage-wise independence: SDDP [Pereira and Pinto, 1991]
- Pure binary state variables: SDDiP [Zou et al., 2019]
  $\rightarrow$ Lagrangian cuts tight at binary points
- General integer state variables: Binarization + SDDiP
  $\rightarrow$ Large # of binary state variables
- Lipschitz continuous exp. cost-to-go functions: [Ahmed et al., 2020]
  $\rightarrow$ Nonlinear cuts + augmented Lagrangian
- General nonconvex mixed-integer nonlinear:
  - SDDP with generalized conjugacy cuts [Zhang and Sun, 2019]
    $\rightarrow$ Approximate *regularized* exp. cost-to-go functions
  - Nonconvex nested Benders [Füllner and Rebennack, 2022]
    $\rightarrow$ Extends binarization and regularization procedures
    $\rightarrow$ Successful implementation for deterministic multi-stage

Computational Optimization Lab

Our Idea

Challenge: Approximating nonconvex expected cost-to-go functions **(due to integer state variables)**

- Existing works: Develop exact lower-bounding techniques for the nonconvex expected cost-to-go functions
- Our work: Relocate all integer state variables to the first stage
  $\Rightarrow$ the resulting expected cost-to-go functions are **convex**
  $\Rightarrow$ can be approximated (exactly) by a decomposition scheme (e.g., nested Benders or SDDP)

Partially Extended Formulation

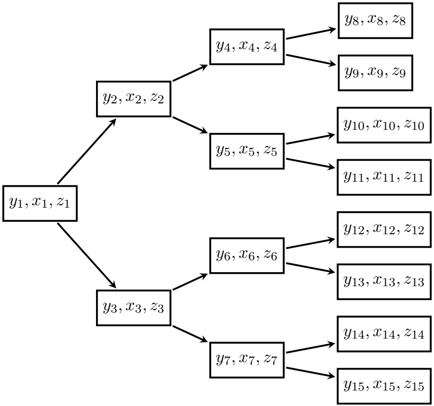

Computational Optimization Lab

Partially Extended Formulation

Our idea: Relocate all integer state variables to the first stage

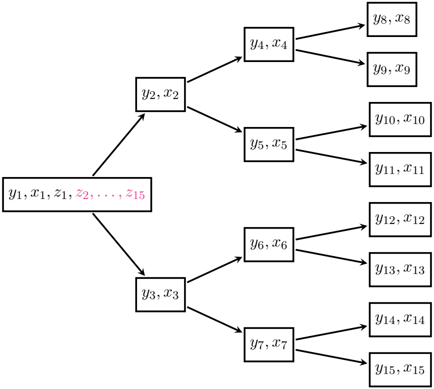

$$Q_n^{\mathtt{Ref}}(x_{a(n)}, z) = \min_{(x_n, y_n)} f_n(x_n, y_n, z_n) + \sum_{m \in \mathcal{C}(n)} \bar{p}_{nm} Q_m^{\mathtt{Ref}}(x_n, z)$$

Too many first-stage (integer) variables!

Our Contributions

- Goal: Leverage the structure of the underlying stochastic process to obtain high-quality policies
- Proposal:
  - Aggregation framework

    $\rightarrow$ Impose additional structure to the **integer state** variables based on the stochastic process (e.g., Markov Chain)
- Methodology:
  - Branch-and-cut algorithm integrated with SDDP

    $\rightarrow$ Exact and approximation methods
  - MC-based two-stage linear decision rules

    $\rightarrow$ Approximation method
- Application:
  - Hurricane disaster relief logistics planning

Computational Optimization Lab

Aggregation Framework

Idea: Simply enforce $z_n = z_{n'}$ for some pairs of nodes based on MC

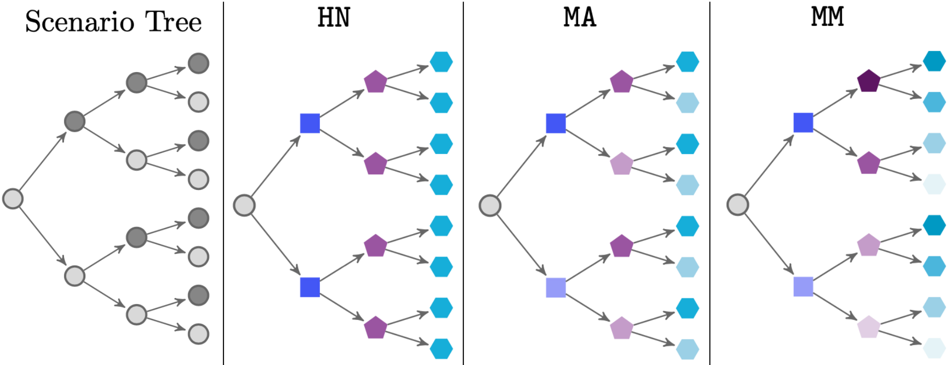

	HN	MA	MM
	Here-and-now (current stage)	Markov-based (current MC state)	Previous and current MC state

	Scenario Tree	HN	MA	MM
# $z^{\mathtt{A}}$'s:	15	4	7	11

Branch-and-cut + SDDP

Decomposition for the aggregated model:

$$y_1, x_1, z^{\mathrm{A}} \underset{\text{Benders cuts}}{\overset{\text{Candidate solution}}{\rightleftarrows}} \{(y_n, x_n)\}_{n \in \mathcal{N} \setminus \{1\}}$$

Exact algorithm:

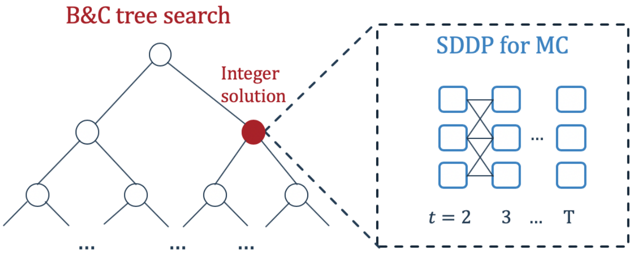

SDDP can be expensive ⇒ lighter version to get LBs

Two-stage Linear Decision Rules

[B. and Luedtke, 2022]

Stage 1	$\cdots$	Stage $t-1$	Stage t	$\cdots$	Stage T
$y_1, x_1, z^{\mathtt{A}}$		$y_{t-1}(\boldsymbol{\xi}^{t-1})$	$y_t(\boldsymbol{\xi}^t)$		$y_T(\xi^T)$
		$x_{t-1}(\boldsymbol{\xi}^{t-1})$	$x_t(\boldsymbol{\xi}^t)$		$x_T(\xi^T)$

$$x_t(\boldsymbol{\xi}^t) = \mu_t^\top \boldsymbol{\xi}^t$$

Stage 1	Stage 2
$y_1, x_1, z^{\mathtt{A}}$,	$\{y_t(\boldsymbol{\xi}^t)\}_{t\in[2,T]}$
$\{\mu_t\}_{t\in[T]}$	

Computational Optimization Lab

Our Scenario-tree Version

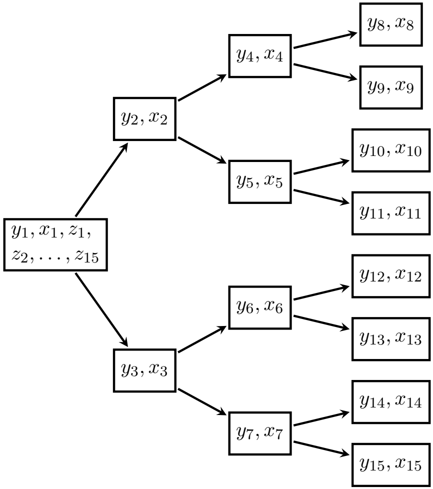

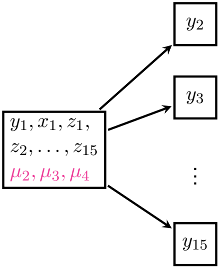

2S-LDR Alternatives

Random variable realizations: $\xi_n^t = \{\xi_{1,\mathrm{r}}, \ldots, \xi_{t-1,a(n)}, \xi_{t,n}\}$

- Stage-history LDR:

$$x_n = \mu_t^\top \xi_n^t$$

- Stage-based LDR:

$$x_n = \mu_t^\top \xi_{t,n}$$

- MC-based LDR:

$$x_n = \mu_{t,m(n)}^\top \xi_{t,n}$$

Computational Optimization Lab

Solving the 2S-LDR Model

Decomposition:

$$y_1, x_1, z^{\mathrm{A}}, \mu^{\mathtt{LDR}} \underset{\text{Benders cuts}}{\overset{\text{Candidate solution}}{\rightleftarrows}} \{(y_n)\}_{n\in\mathcal{N}\setminus\{1\}}$$

Algorithm:

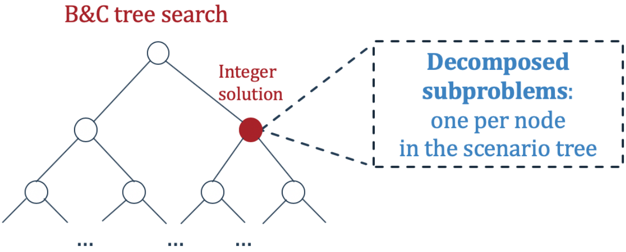

Yields feasible solutions, thus UBs

Hurricane Disaster Relief Planning

with contingency modality

Hurricane Disaster Relief Planning

- Produce & distribute resources from distribution centers to shelters
- **Decision stages:** When hurricane originates to its landing
- **Objective:** Minimize cost
  - Transportation, production, and inventory
  - Unsatisfied demand
- **Decisions:**
  - Local variables: Production, distribution, unsatisfied demand
  - Continuous state variables: Inventory and capacity
  - Integer state variables: Contingency modality activation
    - Choose only one modality
    - Ones active, stays active

Uncertainty Model

Demand is a function of the hurricane's MC state

MC model for the hurricane

- Region represented by a grid
- States: intensity + location
- Cone-shape movement until landing
- MC for intensity [Pacheco and Batta, 2016]

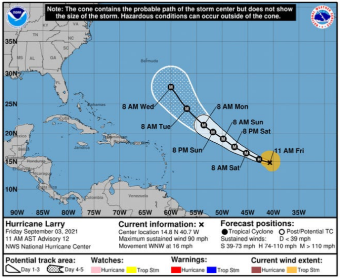

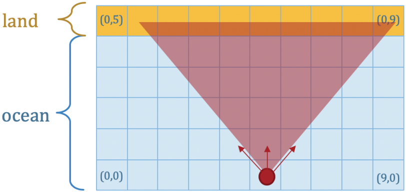

Initial state
Location: (7,0)
Intensity: 4

Computational Optimization Lab

Proposed Restrictions

For contingency modalities

(P.S. $m_t(n) = (m_t^{\mathtt{loc}}, m_t^{\mathtt{int}})$)

- **HN:** Stage-based

$$z^{\mathtt{A}}_{t}$$

- **MA:** MC-based

$$z^{\mathtt{A}}_{t,m_t(n)}$$

- **MM:** Double MC-based

$$z^{\mathtt{A}}_{t,m_t(n),m_{t-1}(a(n))}$$

- **PM:** MC + Intensity

$$z^{\mathtt{A}}_{t,m_t(n),m^{\mathtt{int}}_{t-1}(a(n))}$$

- **FH:** No restriction ($z^{\mathtt{A}} = z$)

For inventories

- **LDR-T:** Stage-based

$$x_{nj} = \sum_{i \in \mathcal{I}} \mu_{t(n)ji} d_{ni}$$

- **LDR-TH:** Stage + history

$$x_{nj} = \sum_{n' \in \mathcal{P}(n)} \sum_{i \in \mathcal{I}} \mu_{t(n')tji} d_{n'i}$$

- **LDR-M:** MC-based

$$x_{nj} = \sum_{i \in \mathcal{I}} \mu_{tm_t(n)ji} d_{ni}$$

Computational Optimization Lab

Experimental Setup

Experiments:

- CPLEX 20.1 + callback
- Single thread
- Time limit: 6 hours

Methods:

- Extensive model
- B&C + SDDP
- 2S-LDR

Instances:

- Medium size: 4x5 grid and 5 stages
- Large size: 5x6 grid and 6 stages
- 6 intensity levels
- Initial capacity: 20%, 30% of max demand
- Modality options:
  - Setting 1: 10%, 20%, 30%, 40%
  - Setting 2: 15%, 30%, 45%, 60%
- 10 instances per configuration

Computational Optimization Lab

Computational Experiments: Main Findings

- **Value of MC-based Policies:**
  - Inclusion of MC info from the previous stage $\rightarrow$ significant improvement (PM and MM policies)
    (Closes $>$ %50 gap between HN and FH policies)
  - MC intensity info of the previous stage captures the most needed (PM policy)
- **Exact Methods:**
  - Poor SDDP performance due to large # of subproblems
  - Extensive model cannot solve larger instances
- **Approximation Methods:**
  - Lower bounding technique via integrated B&C + SDDP offers strong bounds
  - 2S-LDR generates high-quality feasible solutions with reasonable computational time

Computational Optimization Lab

Summary

- Aggregation framework for MSILP with mixed-integer state variables
- Several policies based on the stochastic process (Markov chain)

- B&C framework integrated with SDDP
- MC-based 2S-LDR

- Hurricane disaster relief planning application
- Empirical results showing trade-offs

(Preprint of the paper is available online)

Computational Optimization Lab