Collaborative Diagnosis of Over-subscribed Temporal Plans

Peng Yu

Model-based Embedded and Robotics Systems Group, CSAIL

October 7th, 2016

The Problem of Over-subscription

- We often ask for more than what we can do, while underestimating the length and uncertainty of different activities.

'I want to have Chinese food then watch a movie tonight.'

'Sorry, you cannot do both because your movie starts at ***7pm****. How about eating after the movie?'*

The Problem of Over-subscription

- Over-subscription is a significant issue in the operations of many transportation and robotics systems.


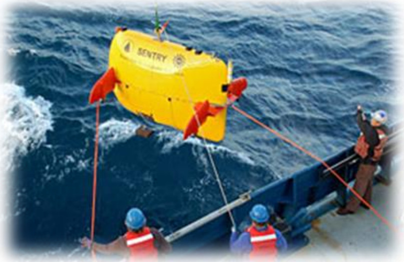


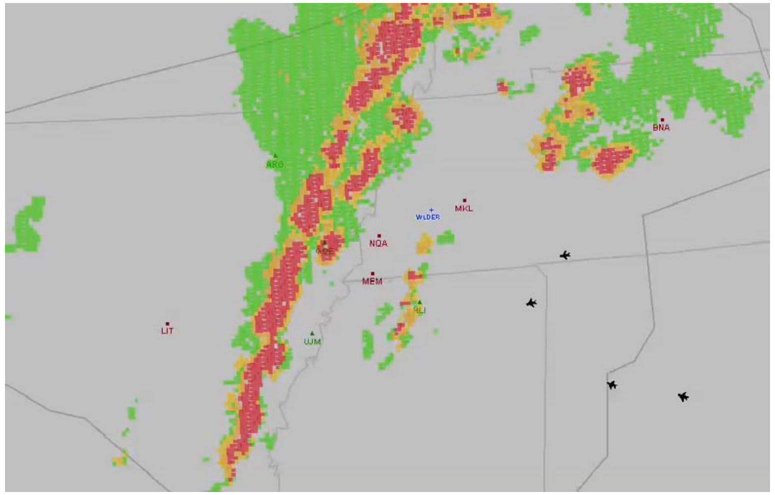

Motivation

- Develop an **advisory system** that helps us resolve over-subscribed plans.

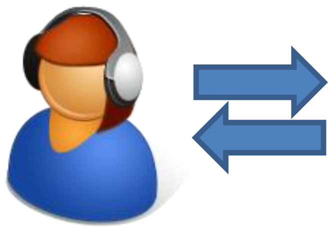

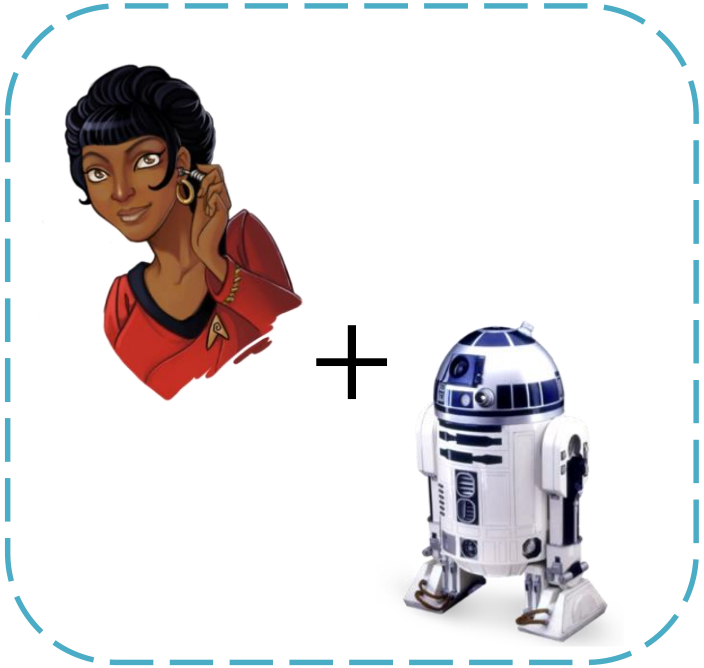

Uhura

Uhura: Plan A Trip to Logan

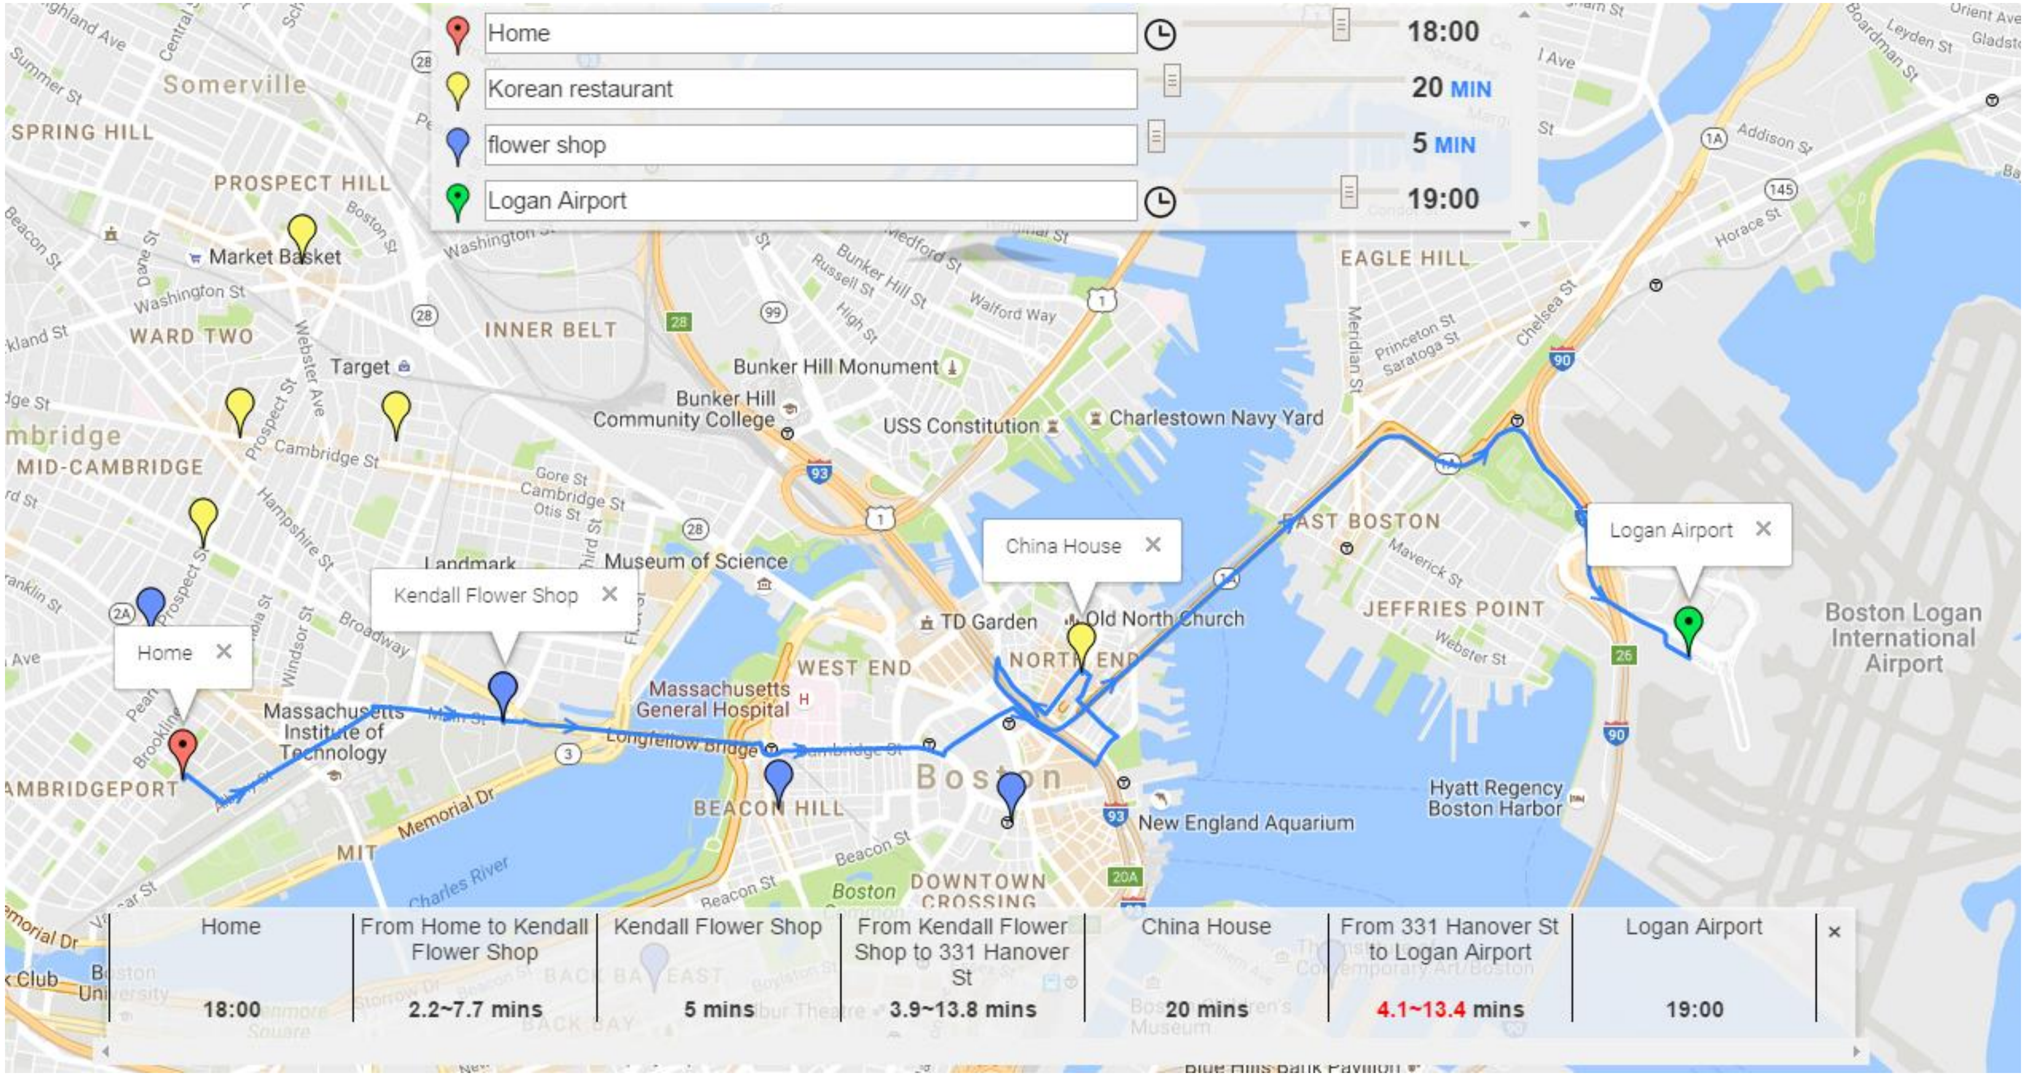

Key Features

- **Continuous relaxation** for temporal constraints.

Can you shift leaving home from 18:00 to 17:54?

- **Risk-bound** relaxation against uncertainty.

There is 29.4% risk that you will not arrive on time, is that ok?

- **Domain relaxation** for unachievable goals.

Can you relax your cuisine requirement from Korean to Chinese?

- **Incorporates** feedbacks from users.

"I cannot leave before 6pm."

"I cannot take more than 1% risk."

Thesis Claim

Resolving over-subscribed plans using a **variety** of **efficient** partial relaxation techniques leads to greater **flexibility** in plan adaptation.

System Architecture

- Uhura: a dialog-based assistant for travel planning problems:

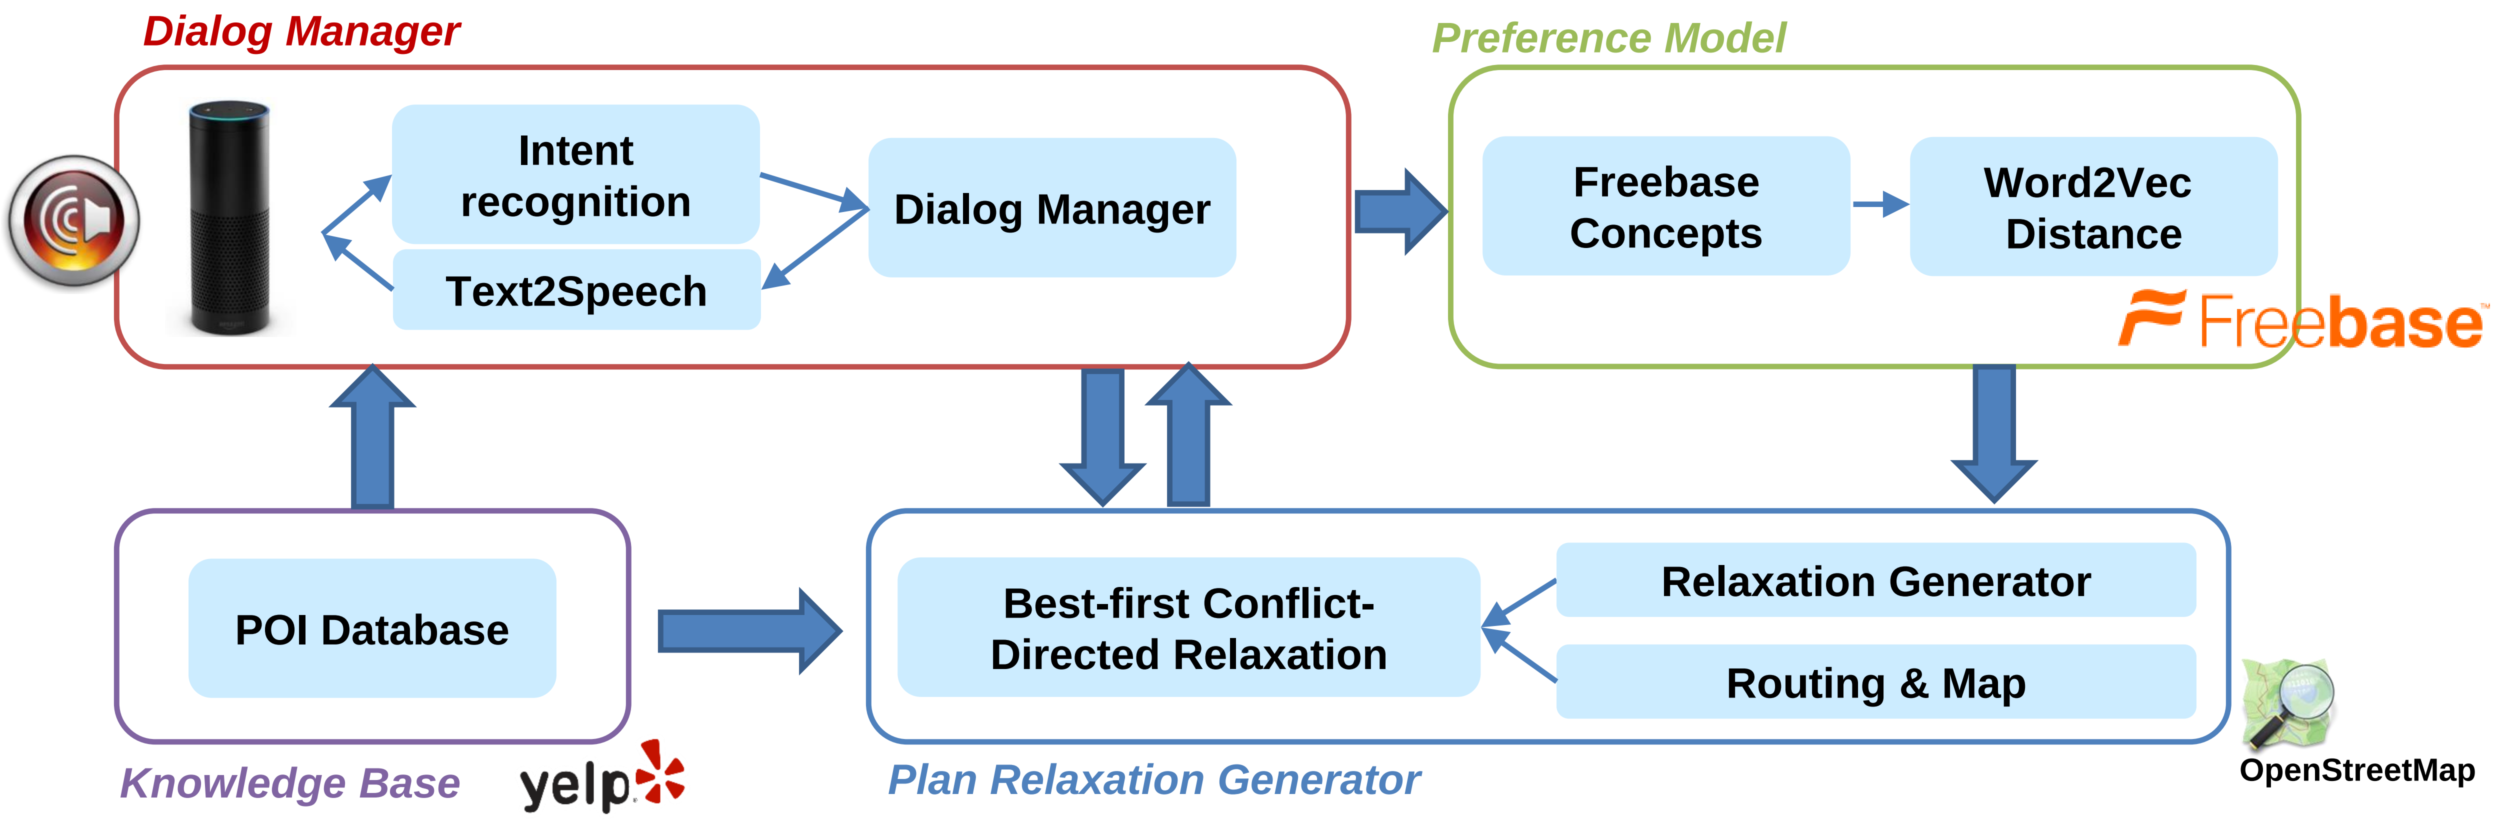

* [Yu, Shen, Yeh and Williams, 2016b]

System Architecture

- This seminar: Best-first Conflict-Directed Relaxation (BCDR).

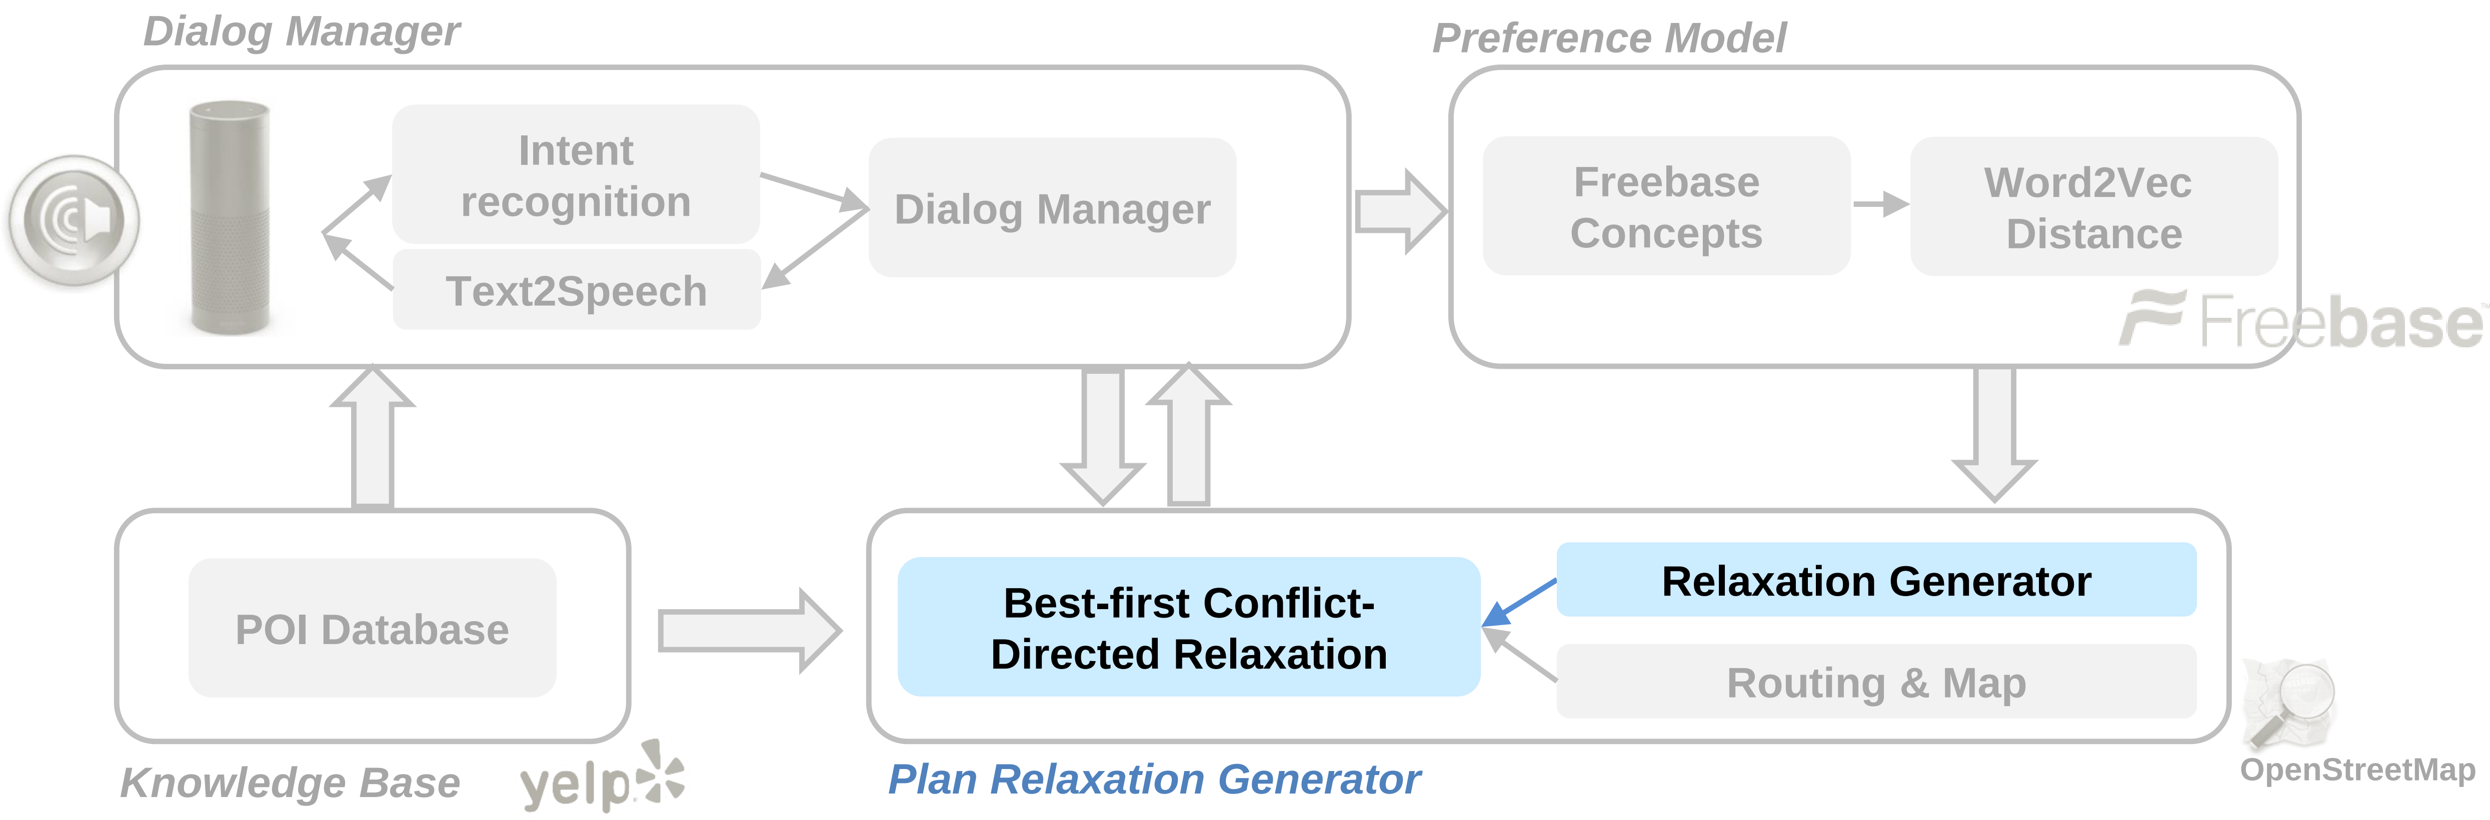

Best-first Conflict-Directed Relaxation (BCDR)

- A collaborative generate-and-test approach:
  - Supports a set of **orthogonal** and **complementary** relaxation techniques.
  - Efficient enumeration of preferred relaxations.

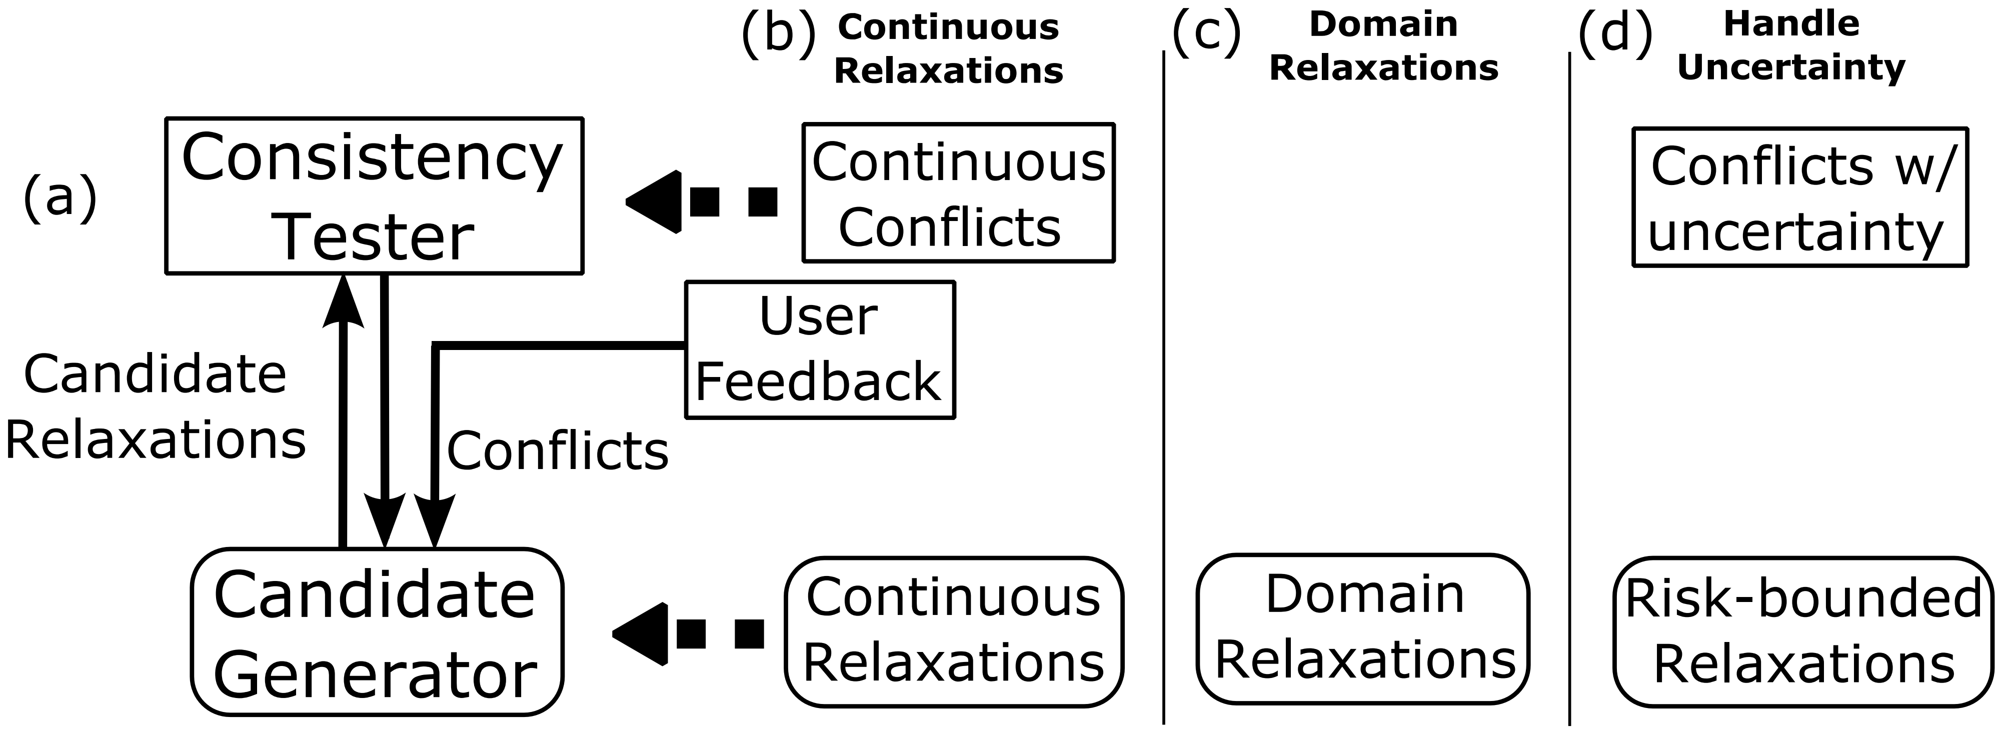

Driven by Multiple Applications

Applications

Research Problems

Continuous *Relaxation*
[Yu, Williams, 2013]

Probal *Schedu*
[Fang, Yu
2014]

tegration with ***ialog Systems***
u, Shen, Yeh, Williams,
16b]

Temporal ***Uncertainty***
[Yu, Fang, Williams, 2014]

Risk-bounded *Relaxations*
[Yu, Fang, Williams, 2015]

Multi-vehicle Coordination *with Temporal Constraints*
[Karpas, Levine, Yu, Williams 2015]

Domain relaxation *using semantically similar alternatives*
[Yu, Shen, Yeh, Williams, 2016a]

This Presentation

- Best-first Conflict-Directed Relaxation.
- Continuous Relaxation for Temporal Constraints.
- Domain Relaxation for Parameterized Variables.
- Risk-bounded Relaxation under Temporal Uncertainty.

A Diagnosis Approach

- We can efficiently enumerate relaxations through conflict-directed diagnosis.

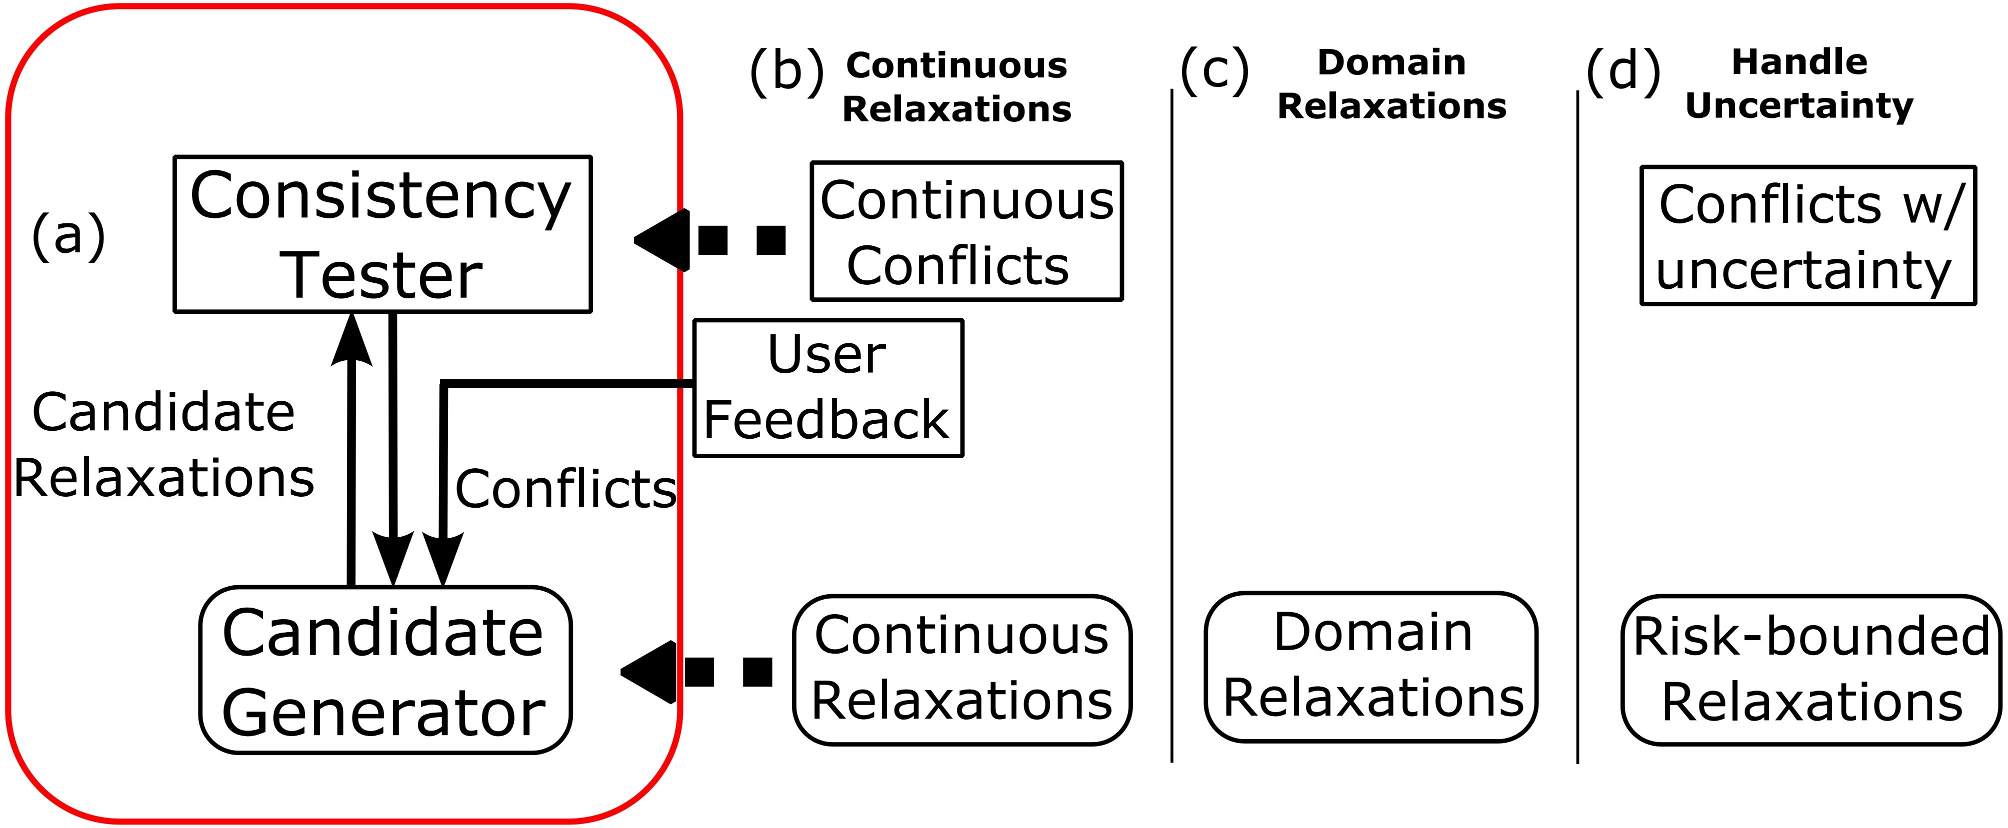

Problem Formulation

- The input is a Temporal Plan.

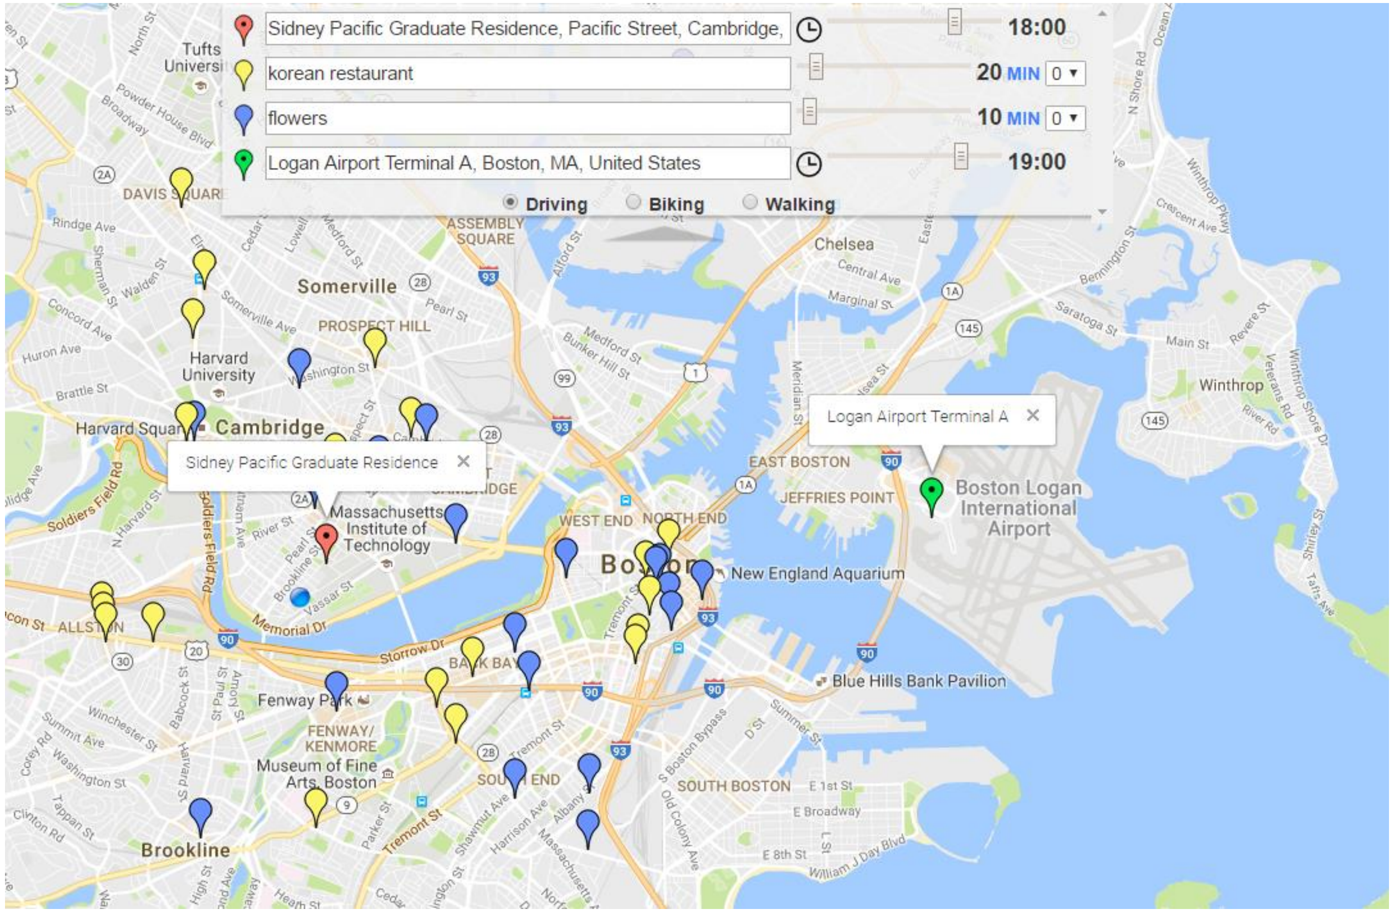

$$\langle P, Q, V, E, L_e, L_p \rangle$$

P, Q: decision variables and assignments.

V, E: guarded events and episodes.

L_e, L_p: guard assignments for episodes and variables.

Problem Formulation

- The input is a Temporal Plan.

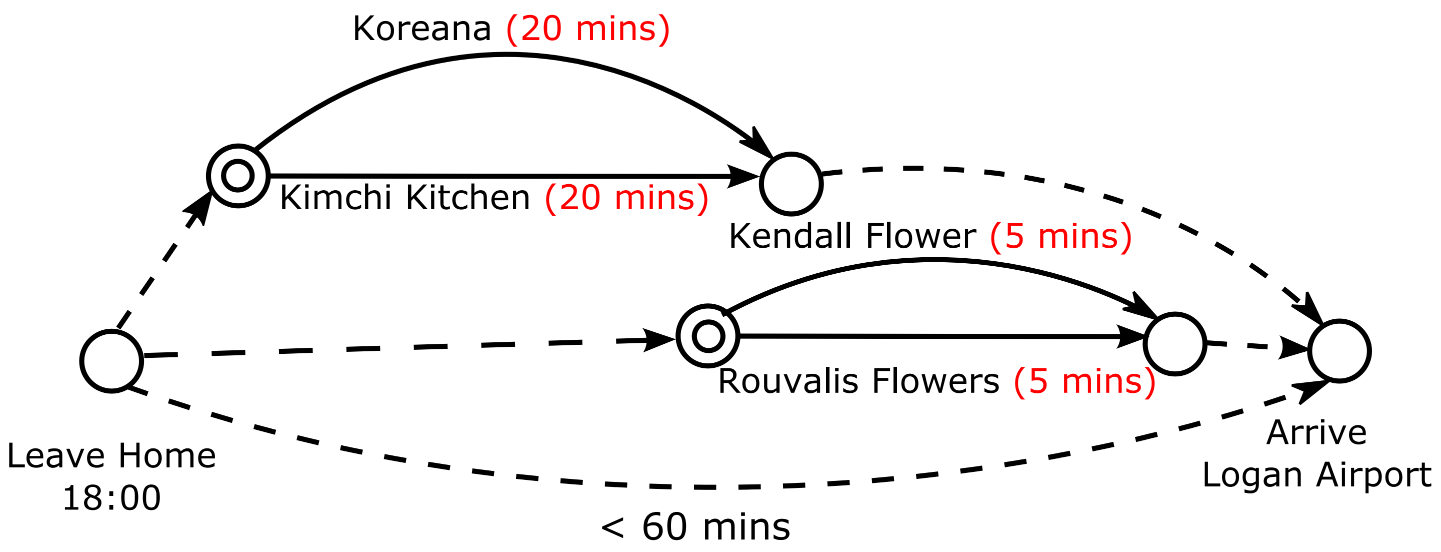

$$\langle P, Q, V, E, L_e, L_p \rangle$$

P, Q: decision variables and assignments.

V, E: guarded events and episodes.

L_e, L_p: guard assignments for episodes and variables.

Output

- The output P is represented by 4-tuple $< A, S, R_e, R_d >$:
  - A feasible set of episodes that supports the activities specified by assignments A, following the sequence assignments S.

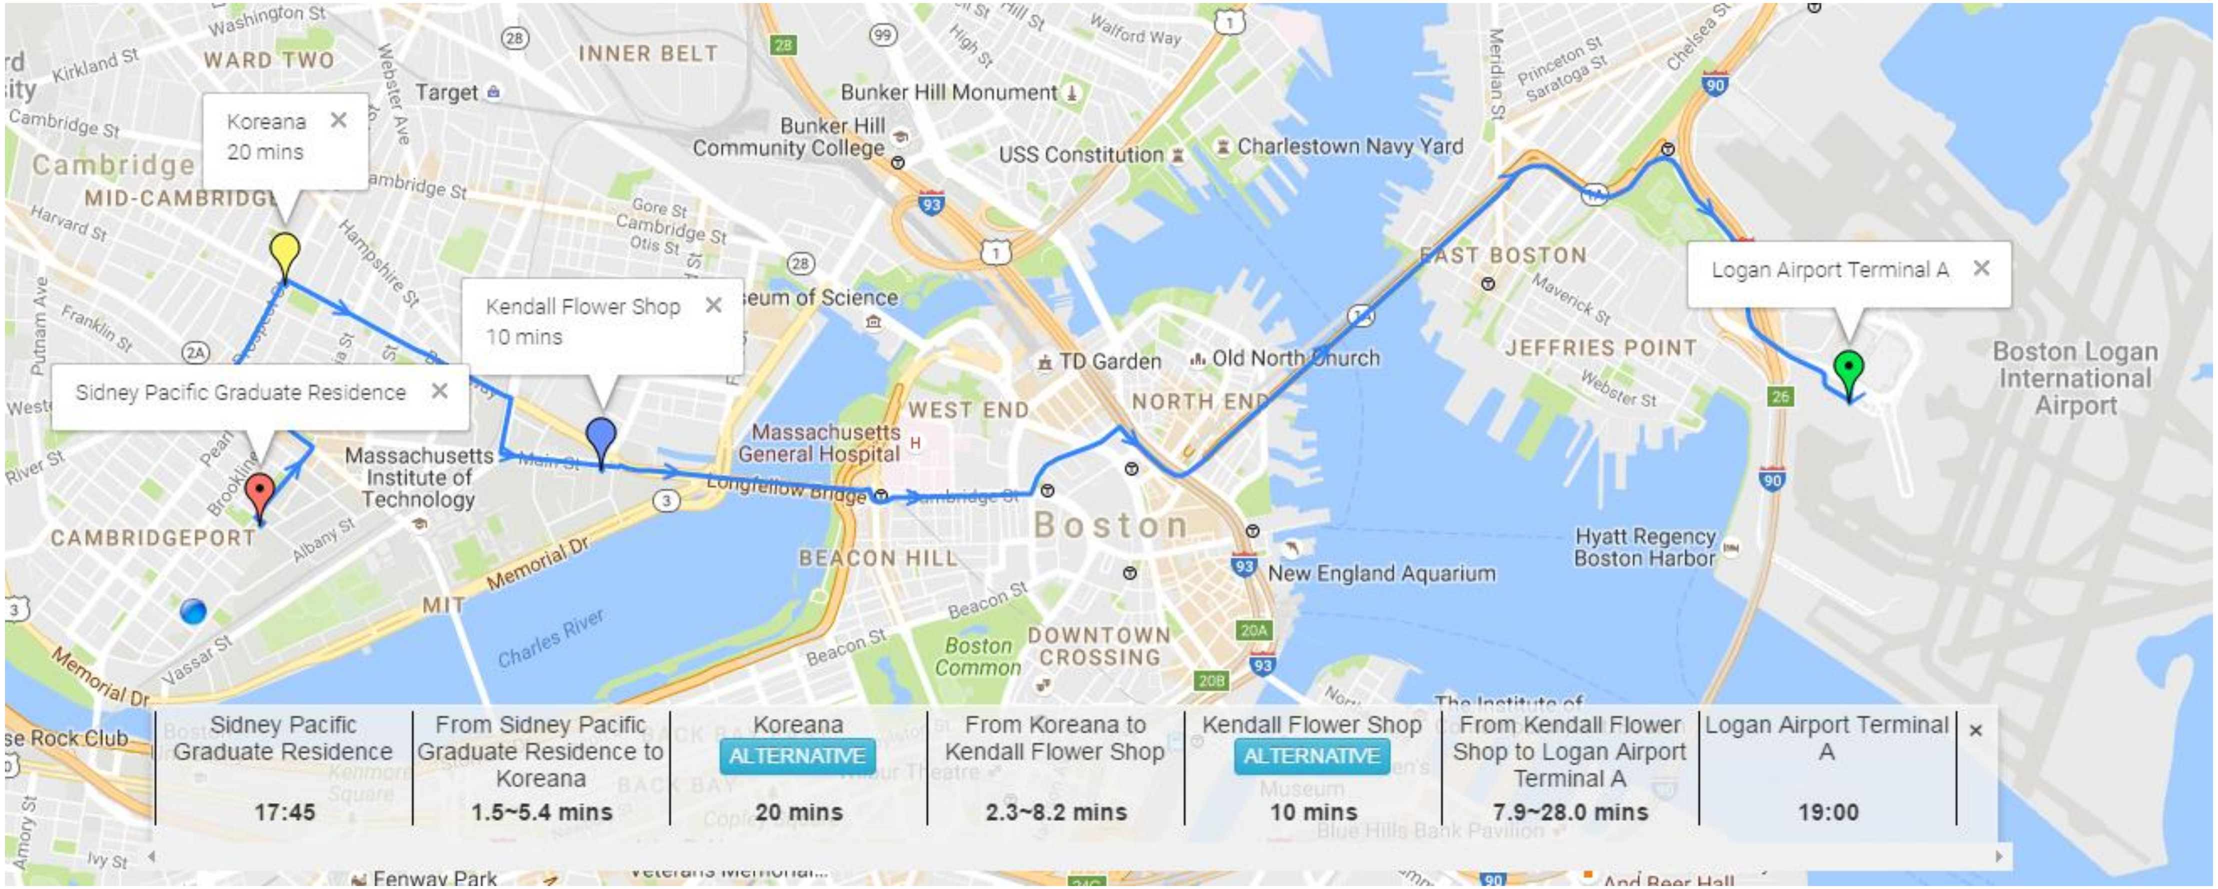

Temporal Relaxation

– R_e is a set of temporal relaxations for the durations of E, which are necessary for making P executable.

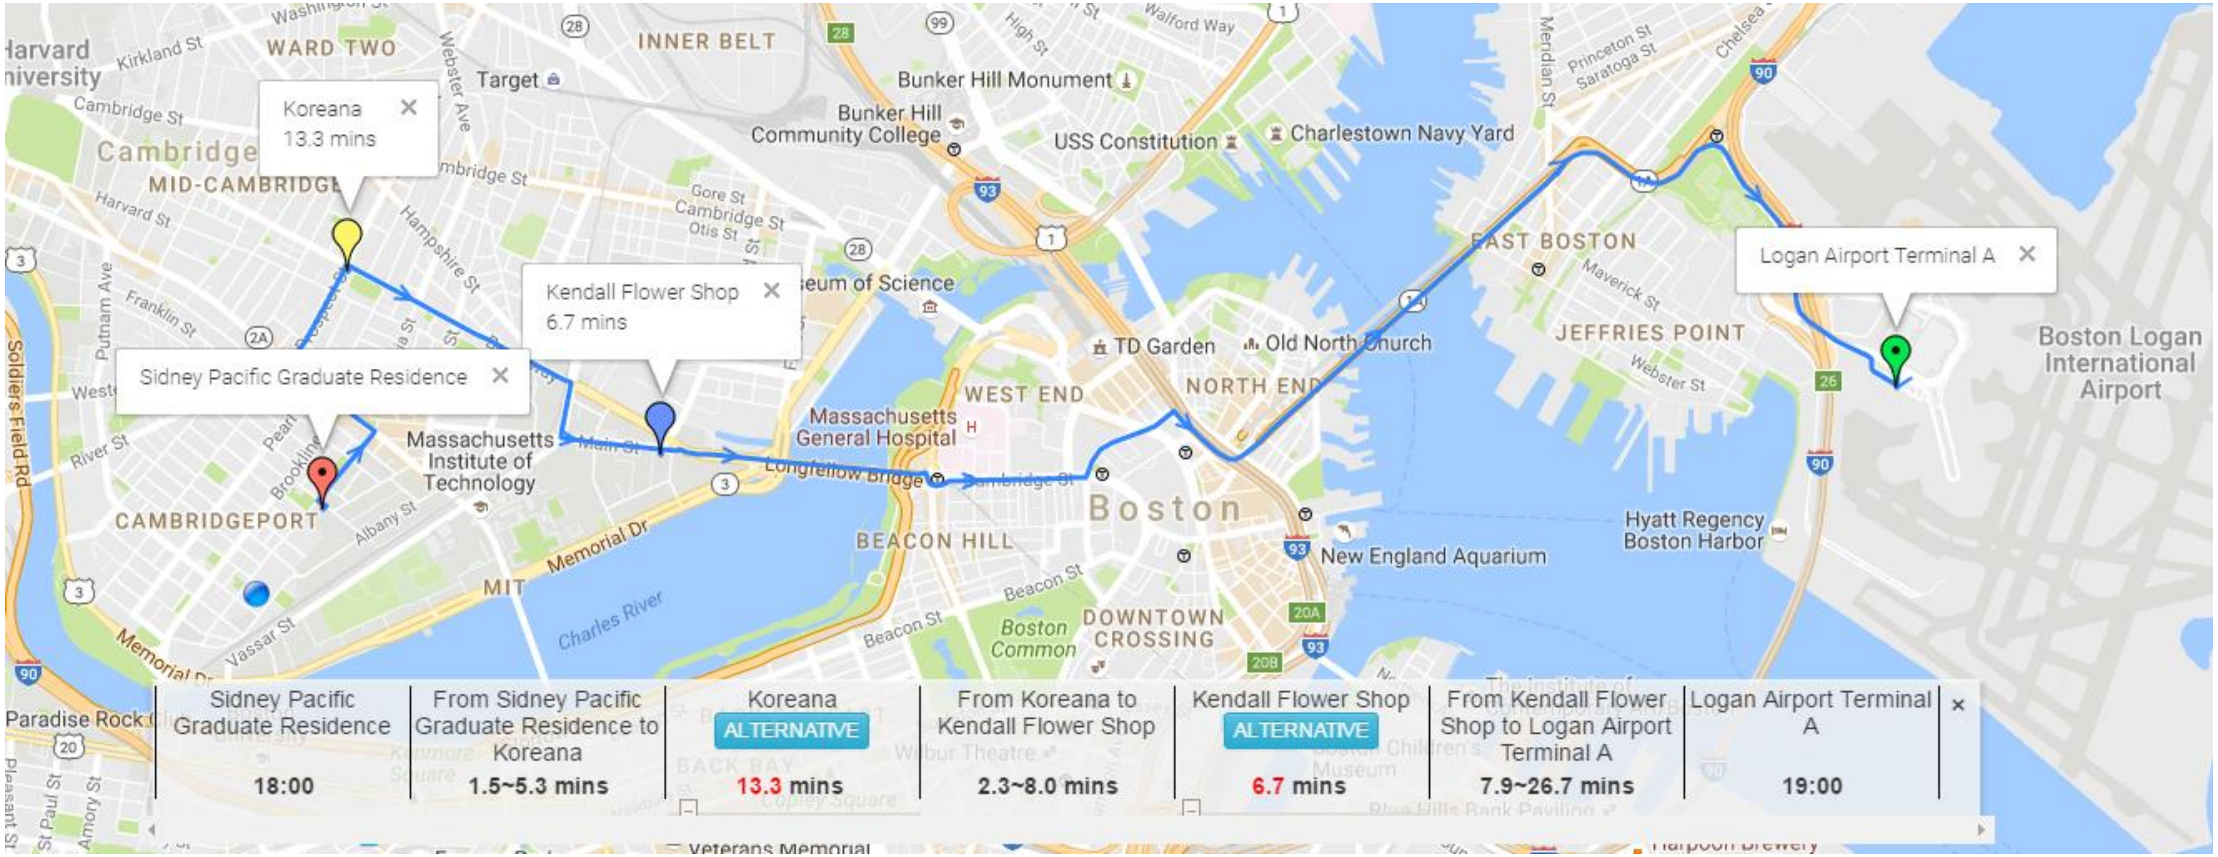

Domain Relaxation

- R_d is a set of domain relaxations to the variables, that allows additional alternative options to be considered in P.

```
rating(r): "4" → "3"
cuisine(r): "Chinese" → "Korean"
location(r): "Cambridge" → "Boston"
```

```
Rating(x)>4 ∧ Cuisine(x)=Korean ∧ Location(x)=Cambridge
```

```
Rating(x)>3
```

{Cho Cho's, Bonchon Chicken}

```
Cuisine(x)=Chinese
```

{Royal East, Shanghai Fresh}

```
Location(x)=Boston
```

{Ginger Exchange, InBoston}

A Diagnosis Approach

- Given an over-subscribed temporal plan, BCDR:
  - **Detects** the conflicting sets of constraints.
  - **Relaxes** constraints to resolve conflicts.
  - **Enumerates** alternative plans and relaxations in best-first order.

- It is like solving a diagnosis problem, but on a broken plan instead of hardware.

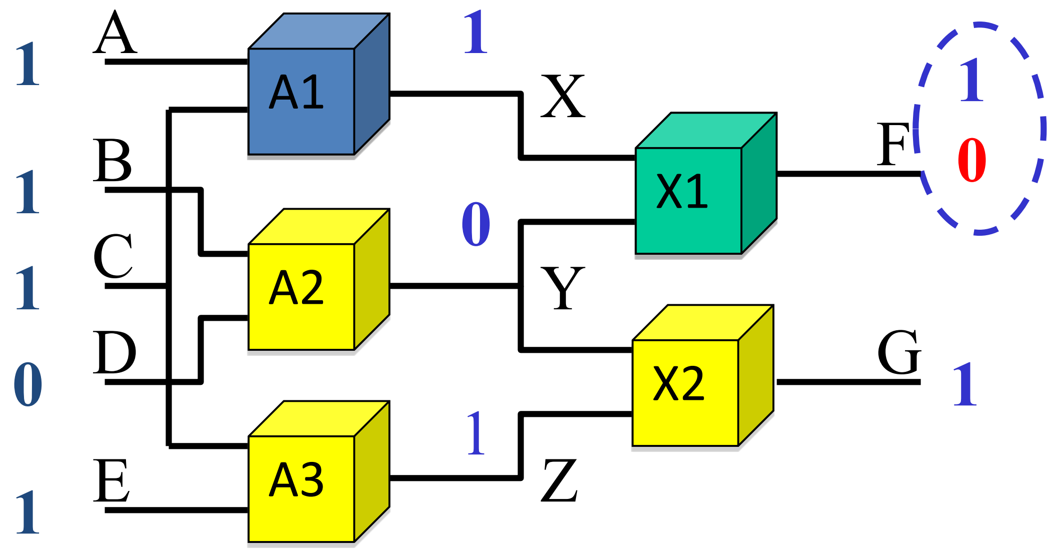

Prior Work on Diagnosis and Temporal Relaxation

- **General Diagnosis Engine** detects likely cause of failures and computes diagnoses*.

- **Conflict-Directed A*** algorithm† prunes search space efficiently using learned conflicts, and enumerates kernels in best-first order.

- Prior work on resolving inconsistent scheduling problems through **constraint suspension**‡.
  - Enumerate minimal suspensions[4].
  - Compute preferred suspensions[5].
  - Efficient domain pruning using learned conflicts[6].

* [de Kleer and Williams, 1989]
† [Williams and Ragno, 2003]
‡ [Davis, 1984]
[4] [Previti and Marques-Silva, 2013]
[5] [Peintner, Moffitt and Pollack, 2005]
[6] [Bailey and Stucky, 2005]

Constraint Suspension

- Resolve over-constrained temporal plan C by **suspending** episodes.
  - Valid Relaxation: $R \subseteq P$ $such\ that\ P \backslash R\ is\ consistent$.
  - Minimal Relaxation: $\forall c \in R, (P \backslash R) \cup \{c\}\ is\ inconsistent$.

Remove arrival time constraint:

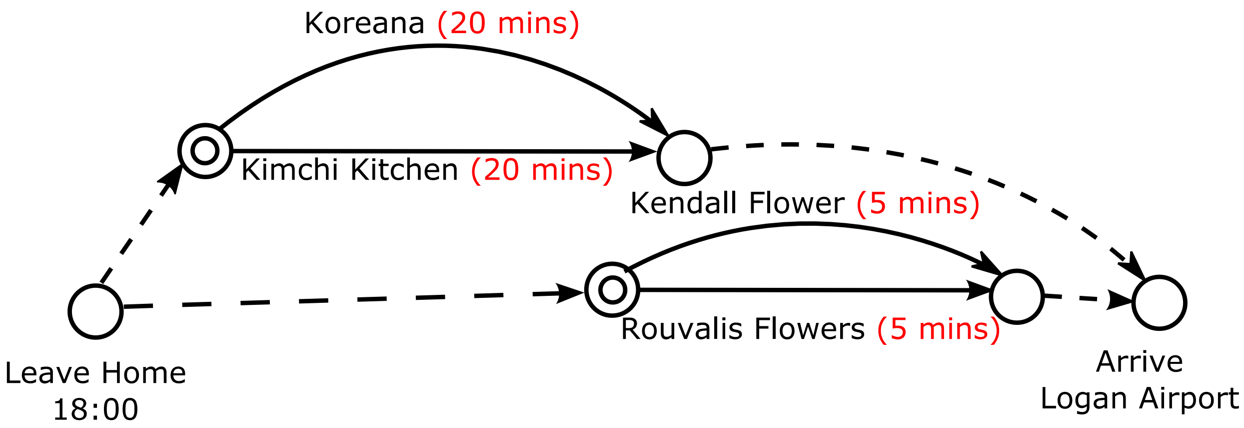

Remove dinner task:

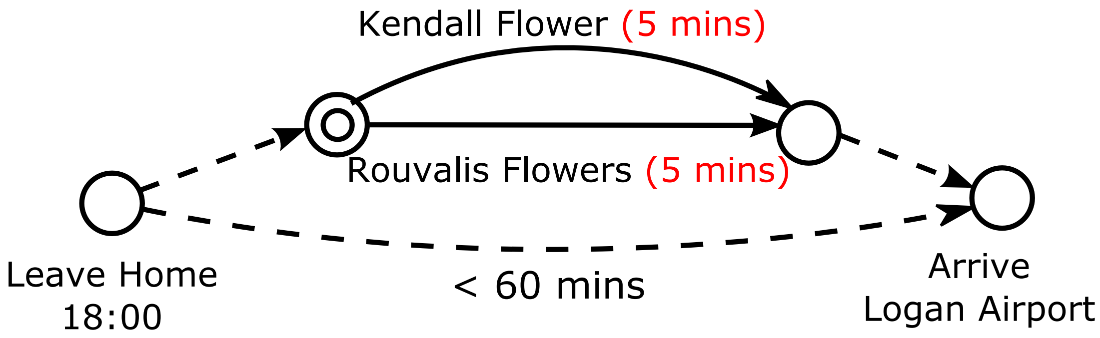

Approach

- BCDR = Conflict-Directed A* + Constraint Suspension.
  - Best-first enumeration of minimal relaxations for over-subscribed plans*.

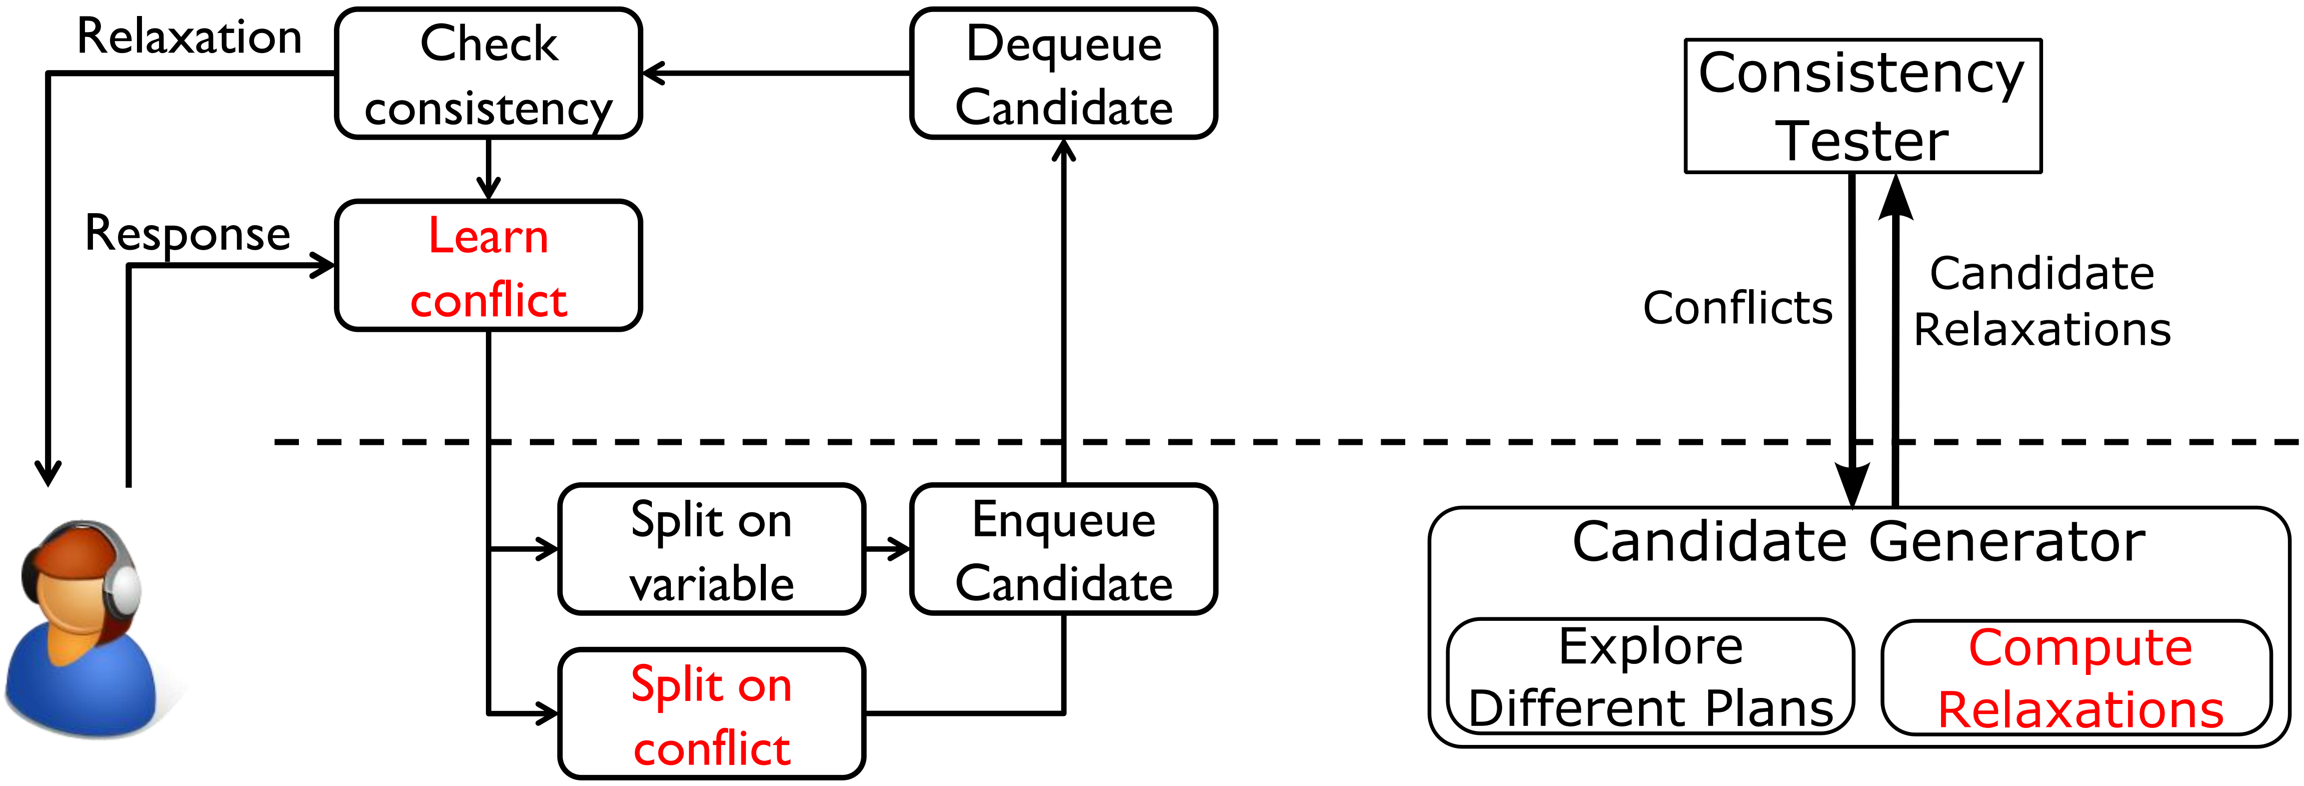

* [Yu, Master Thesis, 2013]

Enumerate Preferred Relaxations

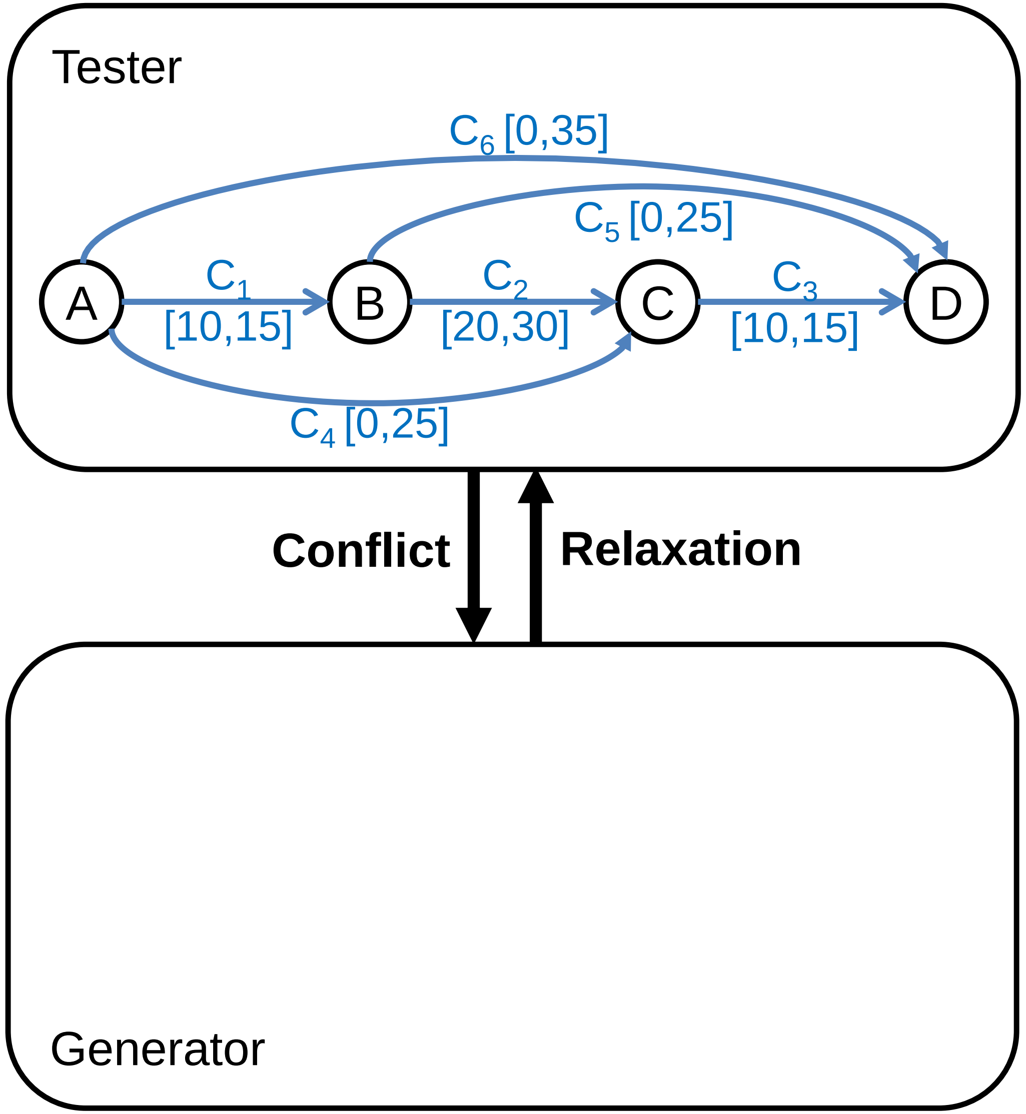

$\{C_2, C_4, C_5, C_6\}$

$\{C_2, C_4, C_5\}$ $\{C_2, C_4, C_6\}$ $\{C_2, C_5, C_6\}$ $\{C_4, C_5, C_6\}$

$\{C_2, C_4\}$ $\{C_2, C_5\}$ $\{C_2, C_6\}$ $\{C_4, C_5\}$ $\{C_4, C_6\}$ $\{C_5, C_6\}$

$\{C_2\}$ $\{C_4\}$ $\{C_5\}$ $\{C_6\}$

$\{\}$

Candidates: [{}]

Enumerate Preferred Relaxations

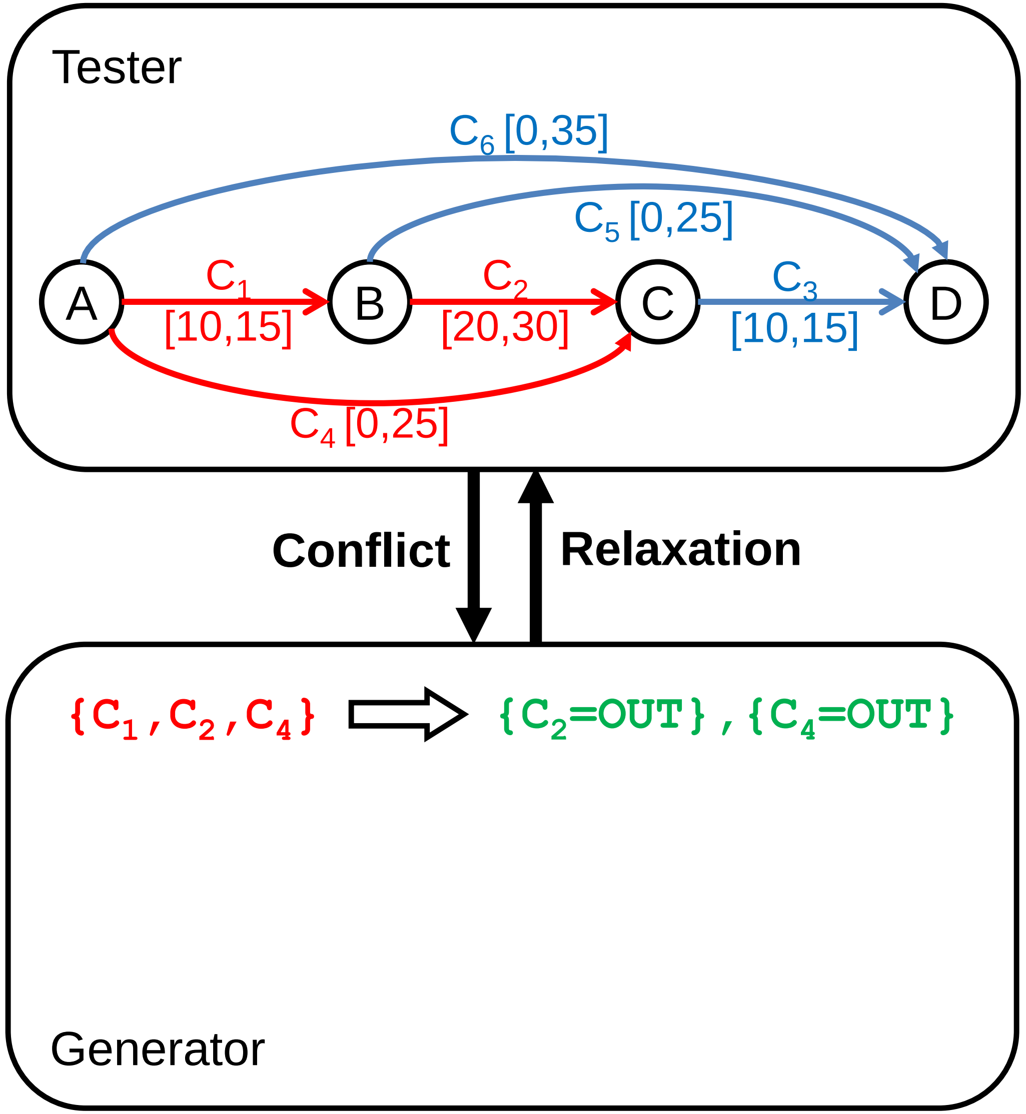

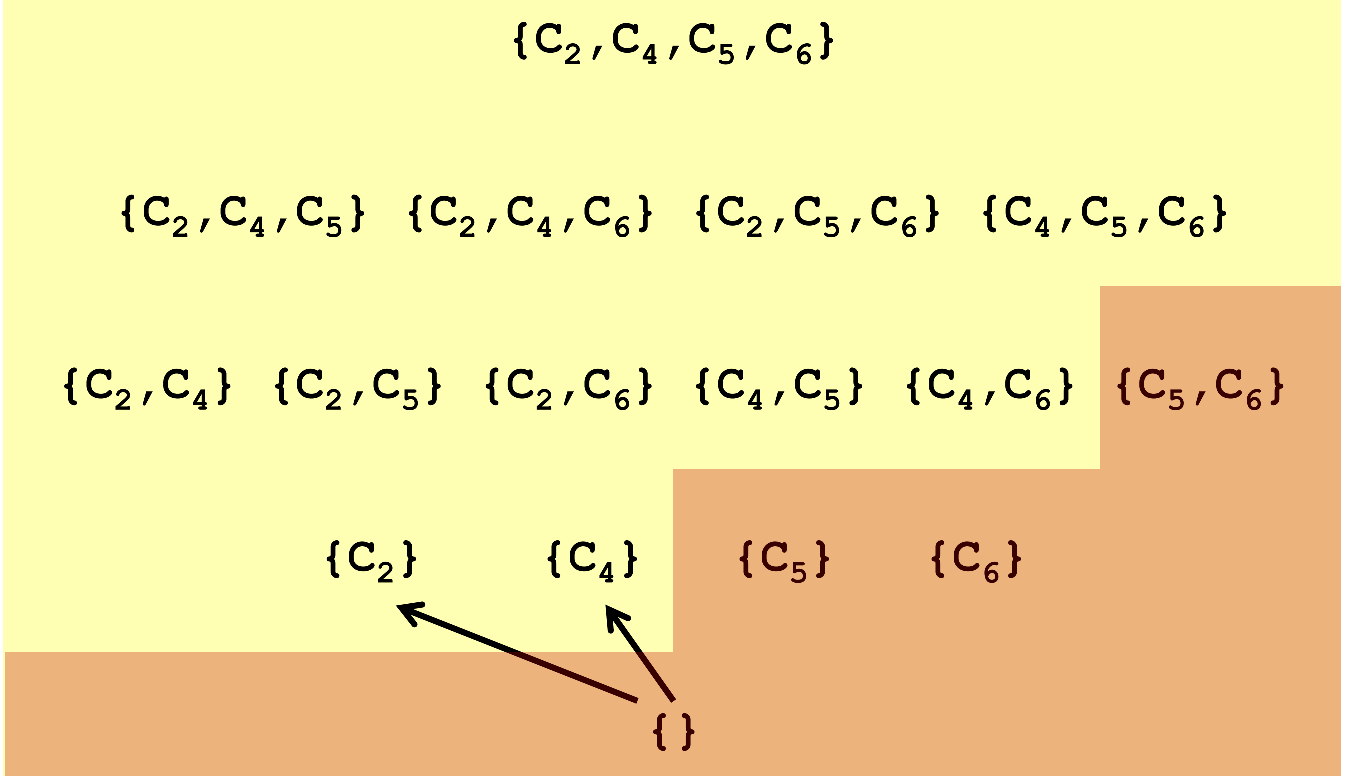

Candidates : [{C_2=OUT}, {C_4=OUT}]

Cost 10 Cost 100

Enumerate Preferred Relaxations

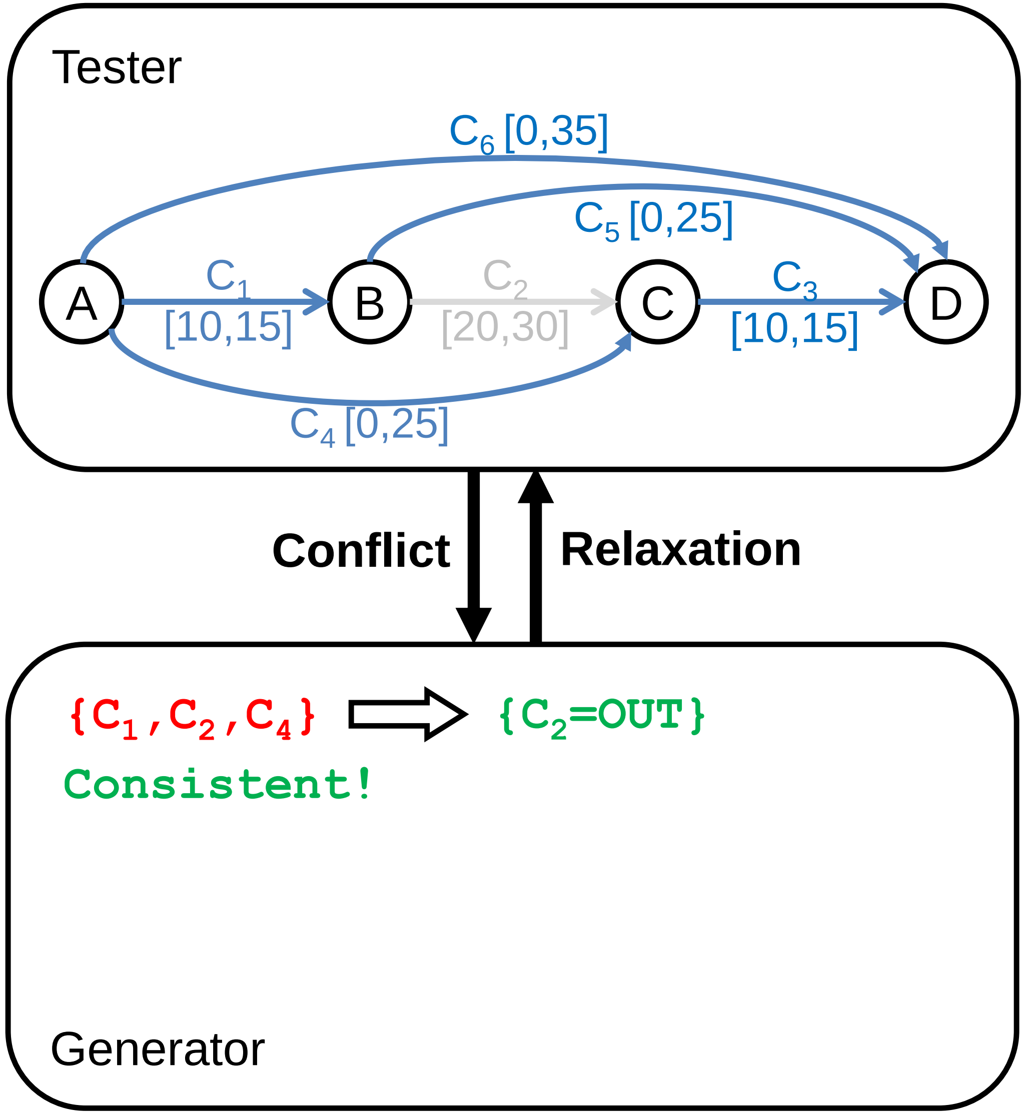

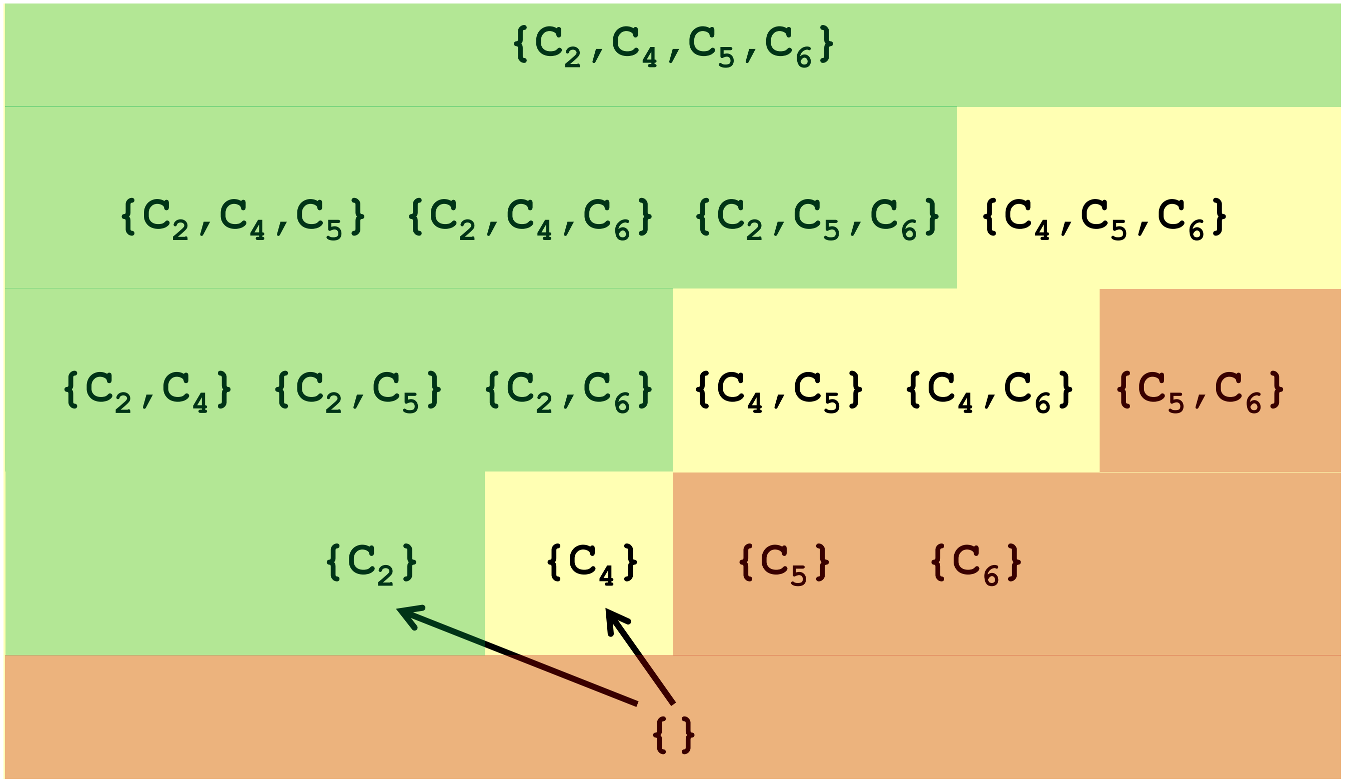

Candidates : [{C_4=OUT}]

Enumerate Preferred Relaxations

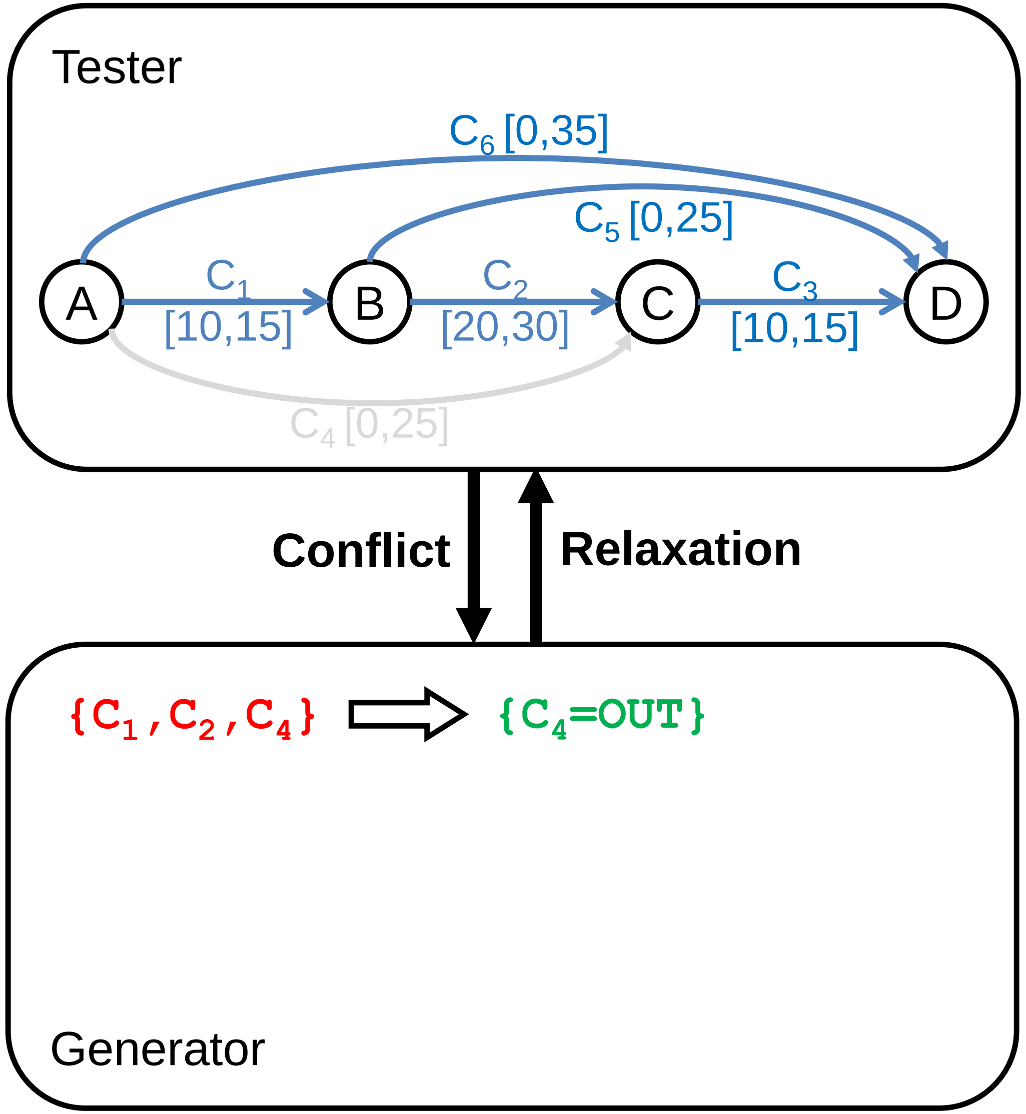

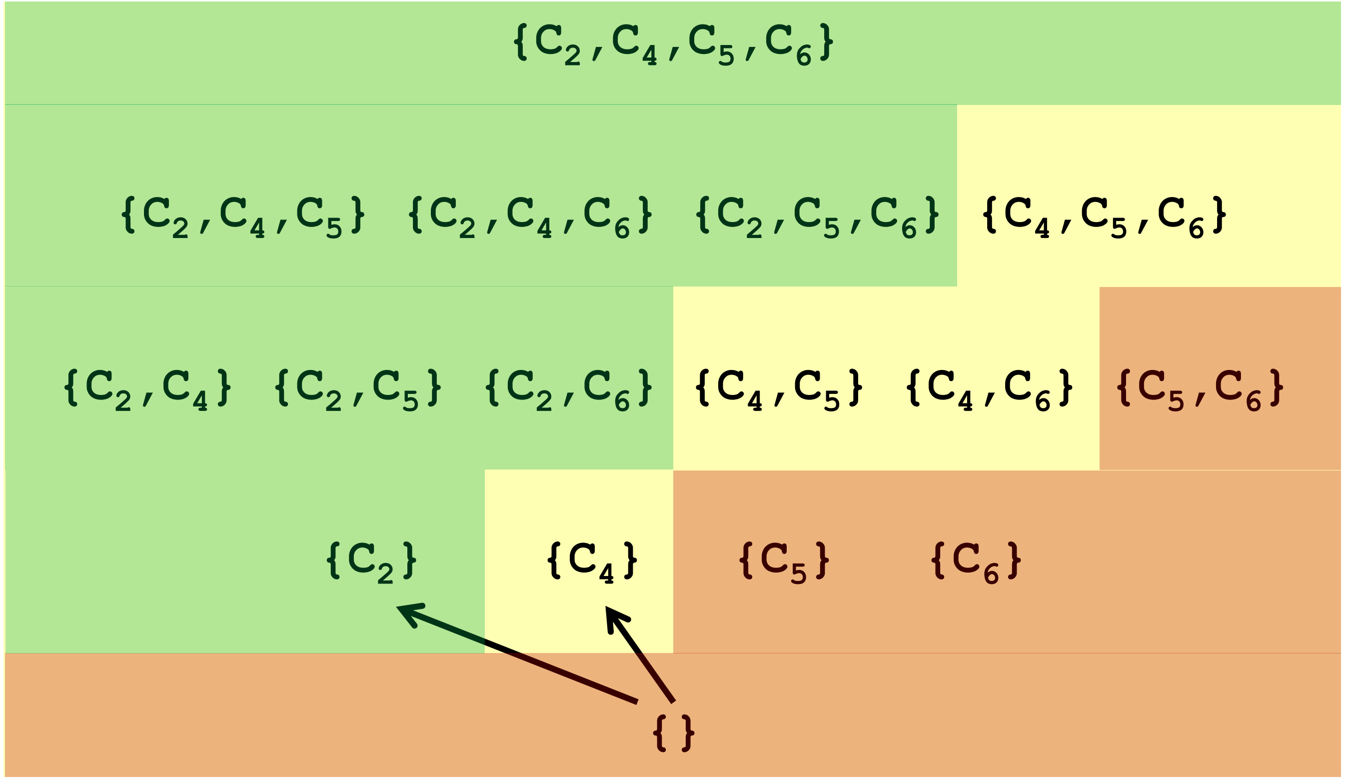

Enumerate Preferred Relaxations

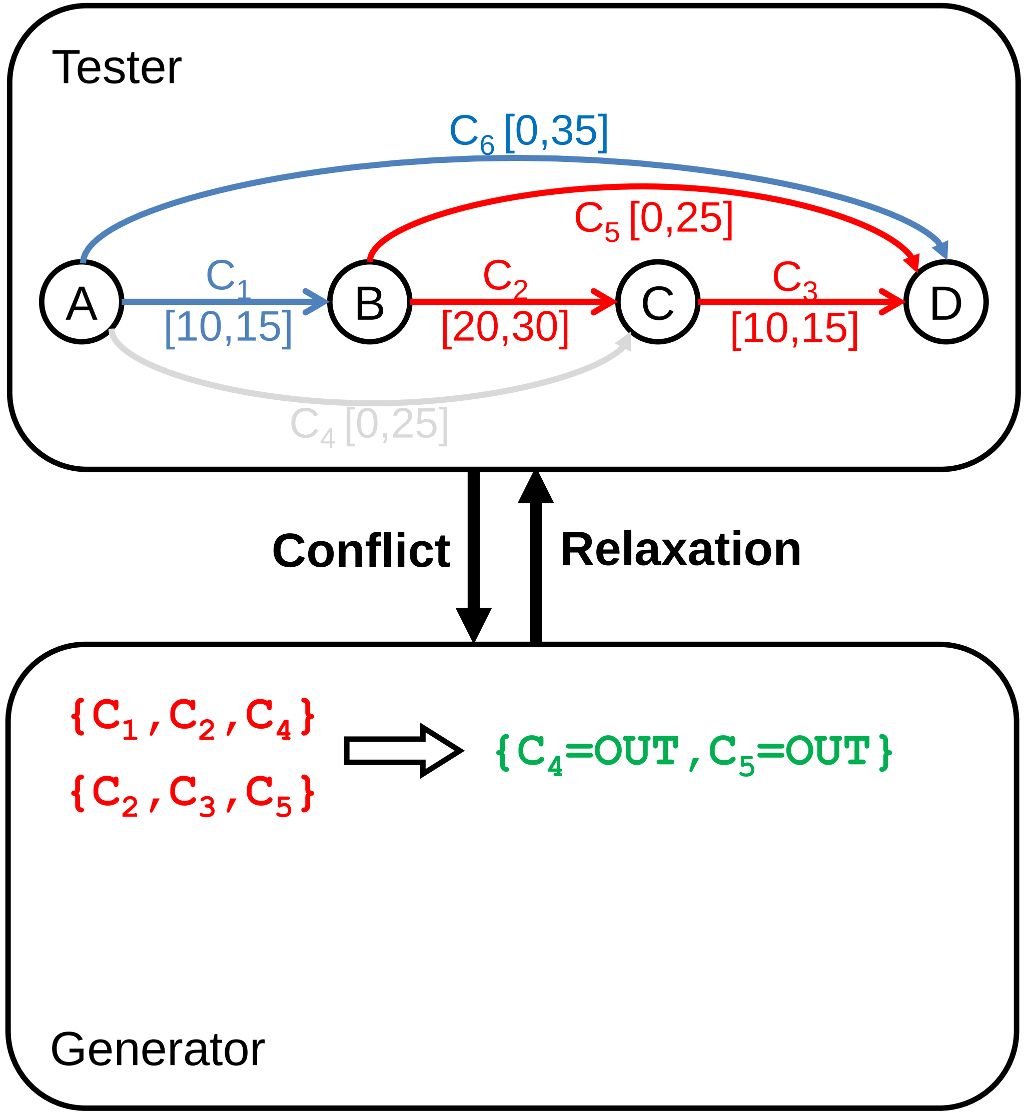

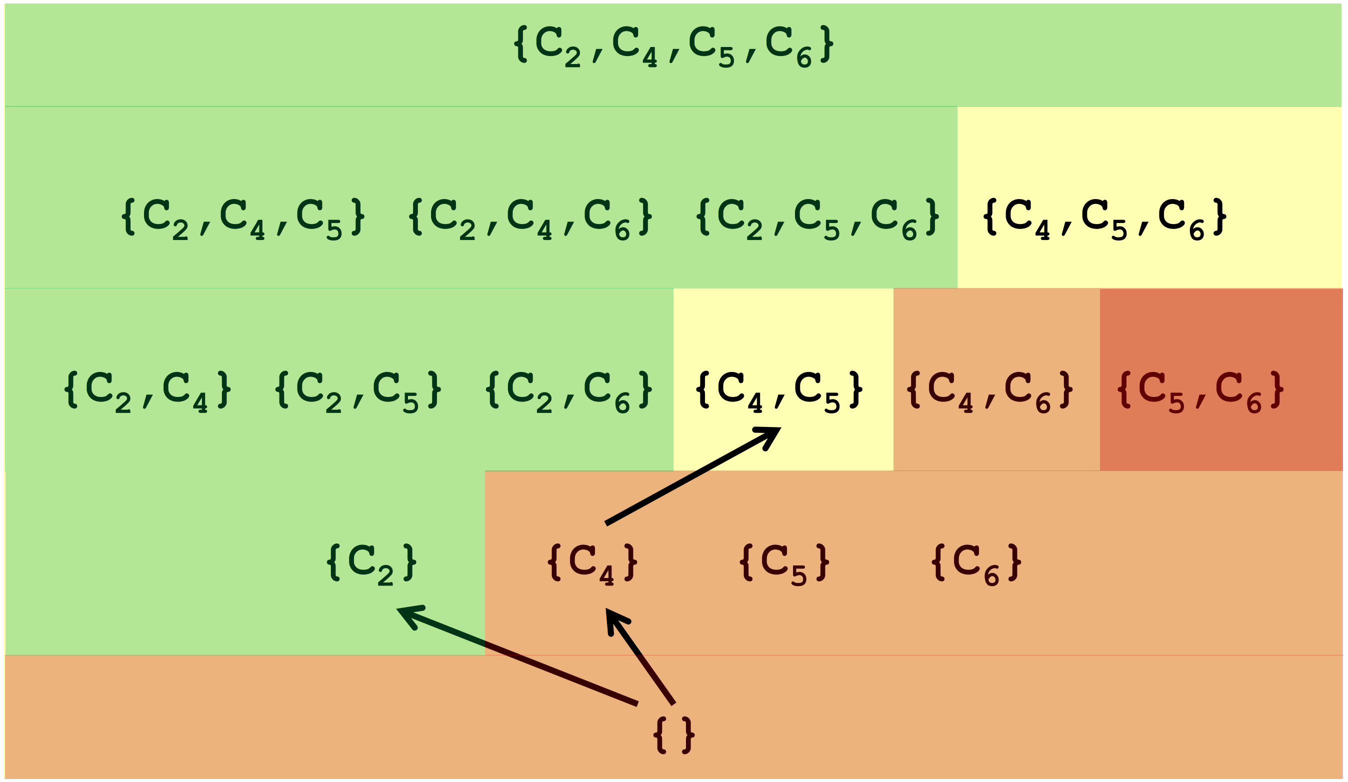

Candidates : [{C_4=OUT, C_5=OUT}]

Enumerate Preferred Relaxations

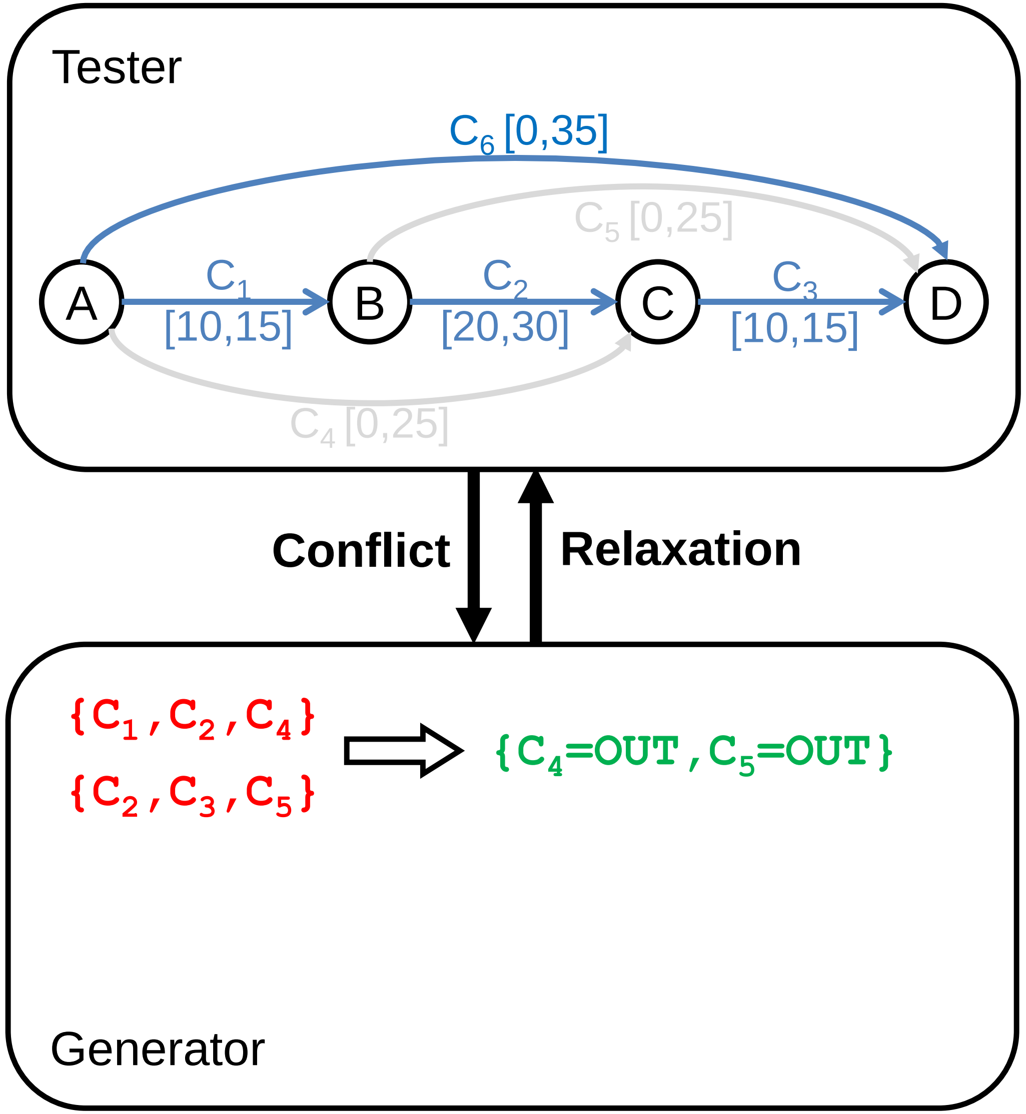

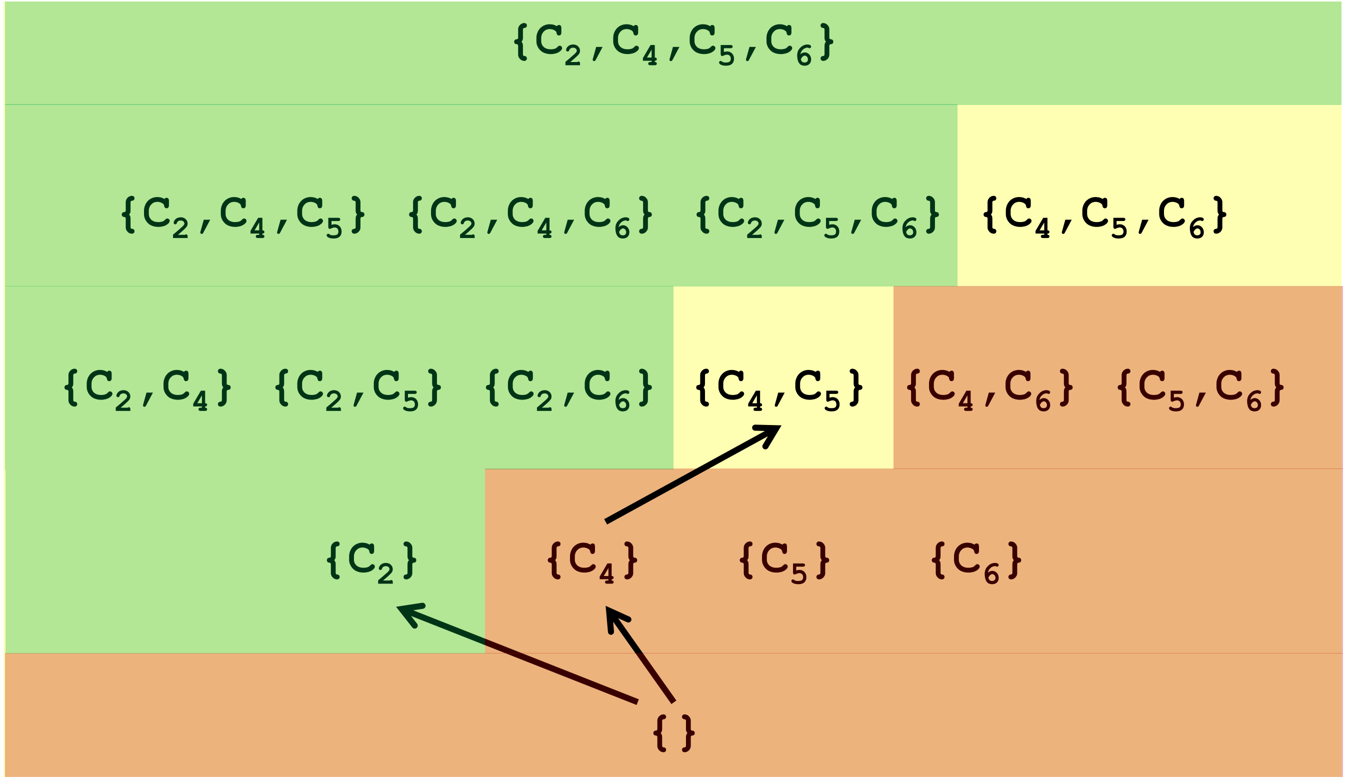

Enumerate Preferred Relaxations

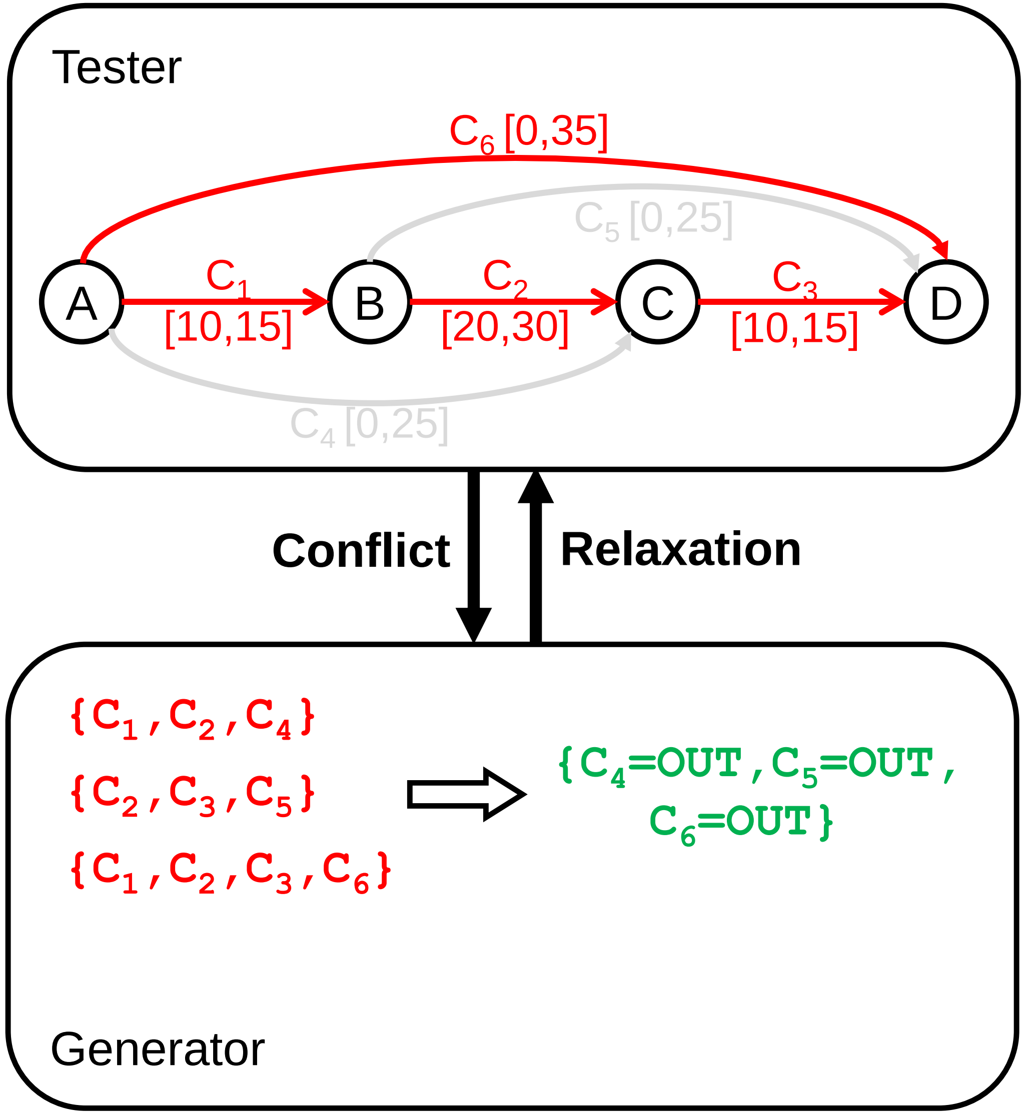

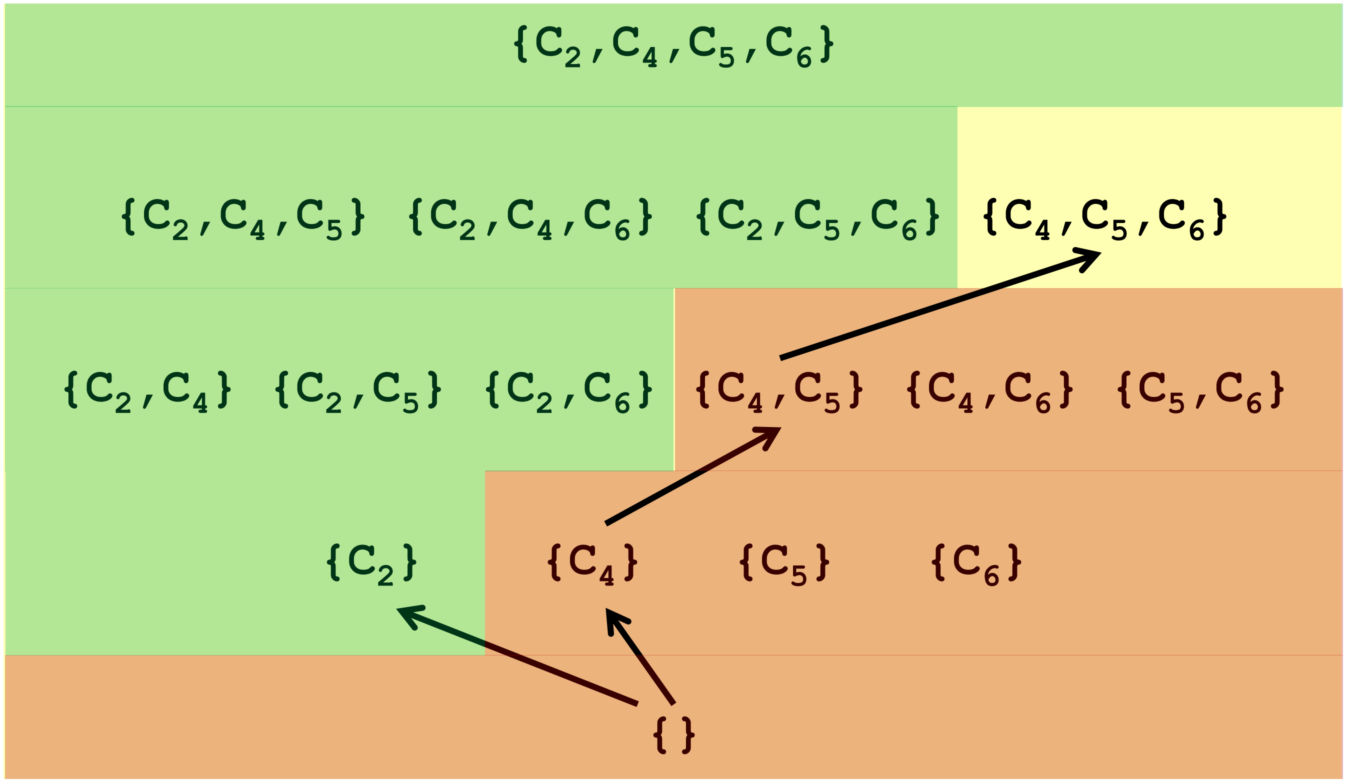

Candidates : [{C_4=OUT, C_5=OUT, C_6=OUT}]

Enumerate Preferred Relaxations

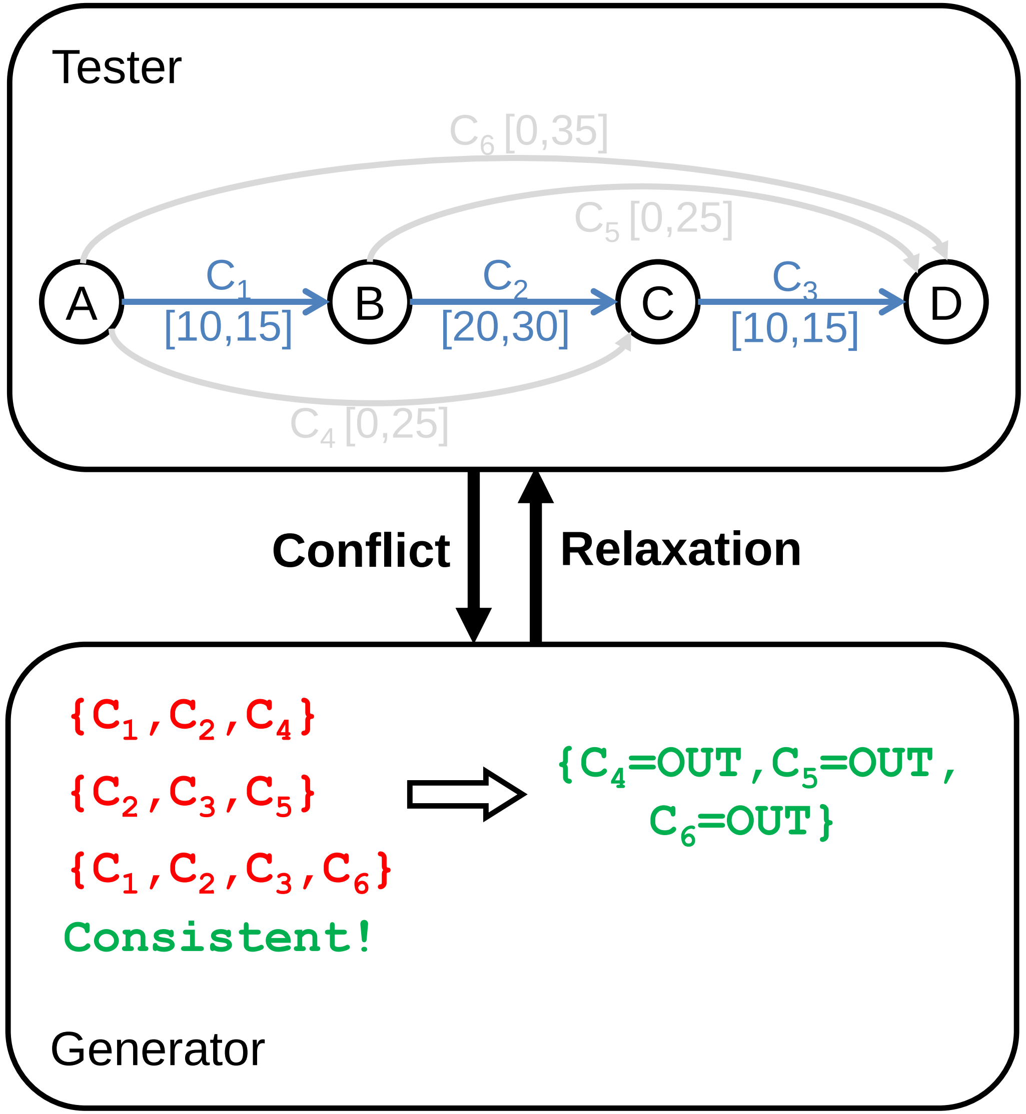

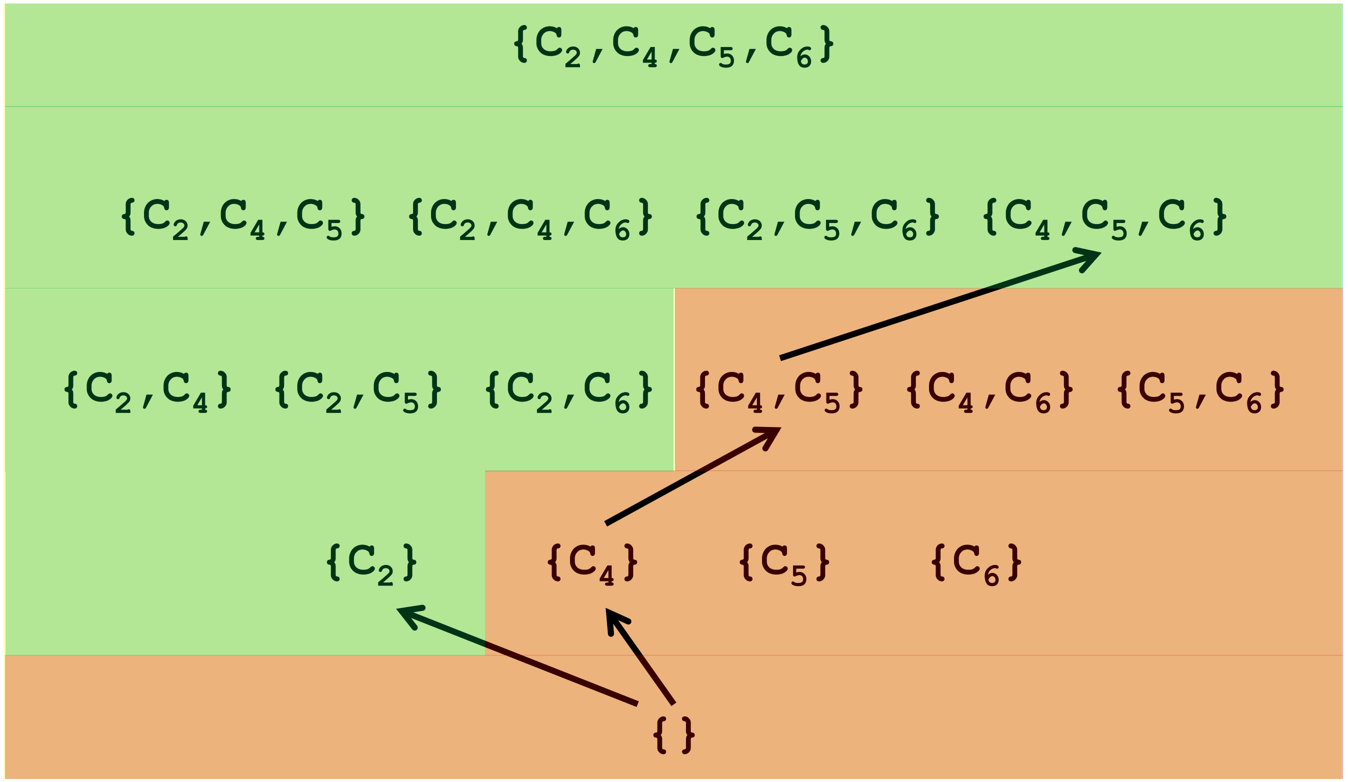

This Presentation

- Best-first Conflict-Directed Relaxation.
- Continuous Relaxation for Temporal Constraints.
- Domain Relaxation for Parameterized Variables.
- Risk-bounded Relaxation under Temporal Uncertainty.

Continuous Relaxation for Temporal Constraints

- We can resolve conflicts through **weakening**, instead of completely suspending temporal constraints.

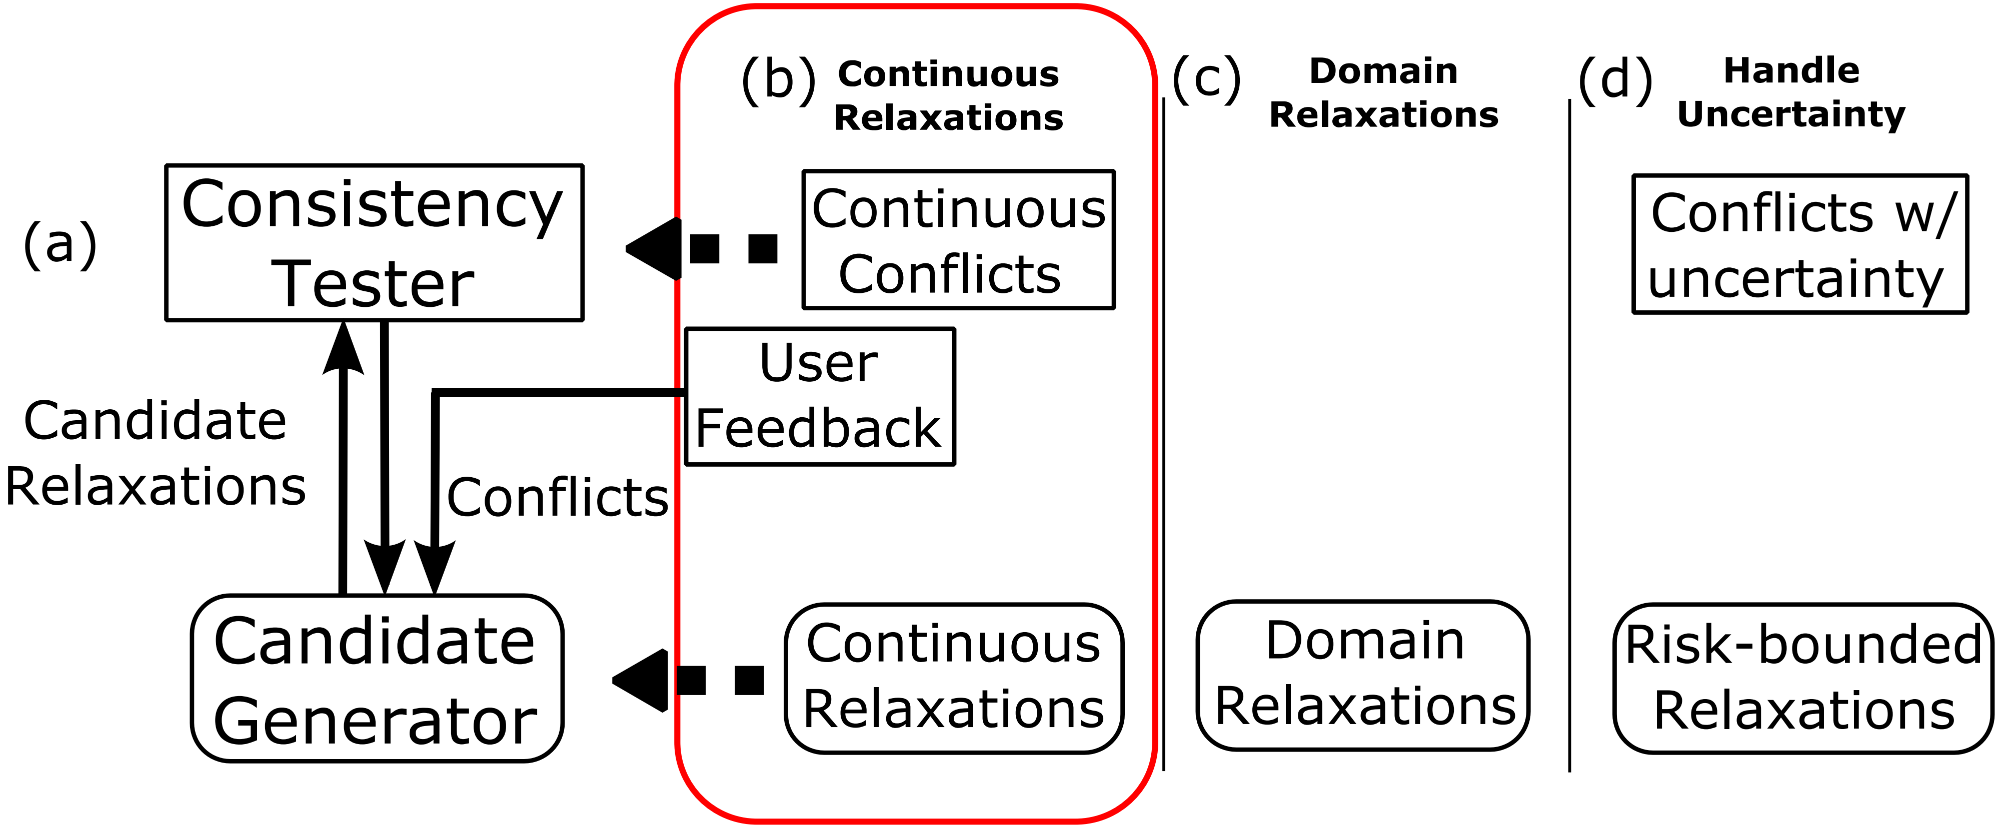

Continuous Relaxation for Temporal Constraints

- Introducing a set of relaxable episodes, $RE \subseteq E$.
  - A continuous relaxation, tr_i, **continuously** weakens the temporal bounds of $e_i \in RE$ from $[LB, UB]$ to $[LB', UB']$, where $LB' \leq LB$ and $UB' \geq UB$.

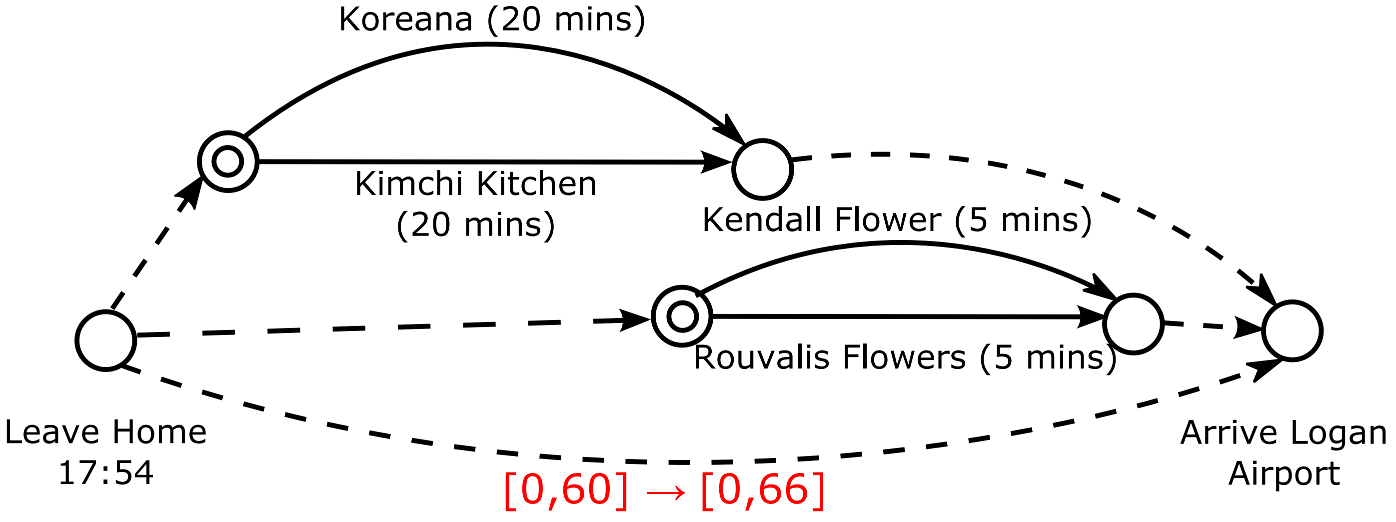

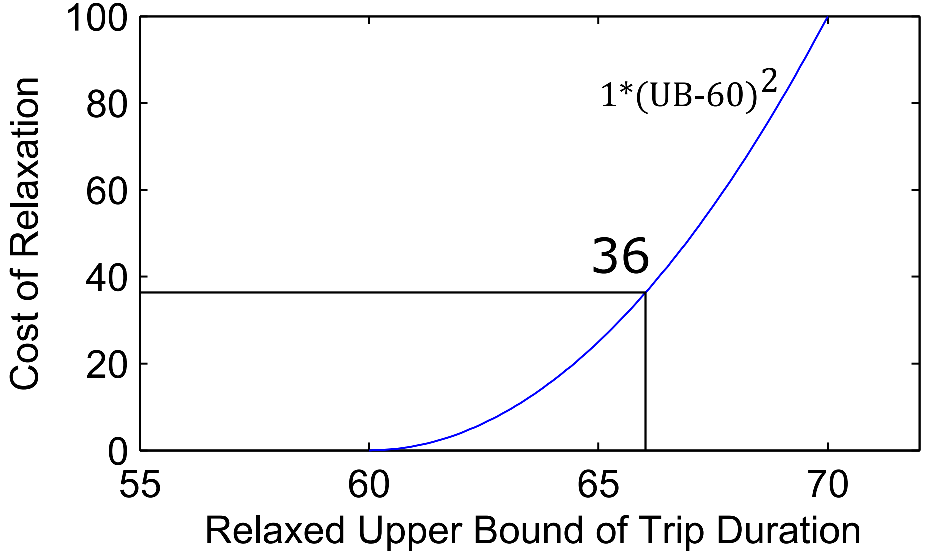

Relaxation cost:

$$tr = TripDuration: [0,60] \rightarrow [0,66]$$
$$f_e(tr) = f_e(66) = 36$$

Continuous Conflicts and Relaxations

- Generalize the discrete conflicts and relaxations, to continuous conflicts and relaxations.
  - a linear expression over constraint bounds that form a negative cycle*.

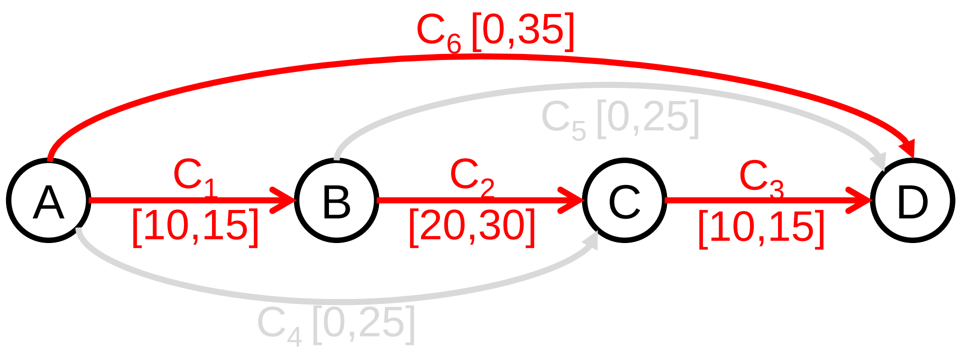

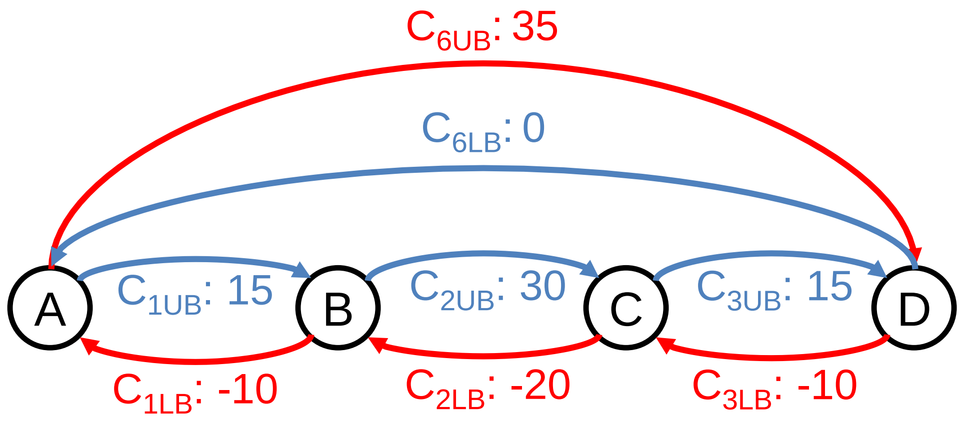

Discrete conflict: $\{C_1, C_2, C_3, C_6\}$

Discrete constituent relaxations:

$\{C_2=\mathrm{OUT}, C_6=\mathrm{OUT}\}$

Continuous conflict: $-C_{1LB}-C_{2LB}-C_{3LB}+C_{6UB} < 0$

Continuous constituent relaxation:

$-C_{1LB}-C_{2LB}-C_{3LB}+C_{6UB} \geq 0$

* [Yu and Williams, 2013]

Preferred Continuous Relaxations

- We define linear/quadratic cost functions over the relaxed bounds.

Minimize

$$f_c(C_{2LB}) + f_c(C_{4UB}) + f_c(C_{5UB}) + f_c(C_{6UB})$$

Subject to

$$-C_{1LB} - C_{2LB} + C_{4UB} \geq 0$$

$$-C_{2LB} - C_{3LB} + C_{5UB} \geq 0$$

$$-C_{1LB} - C_{2LB} - C_{3LB} + C_{6UB} \geq 0$$

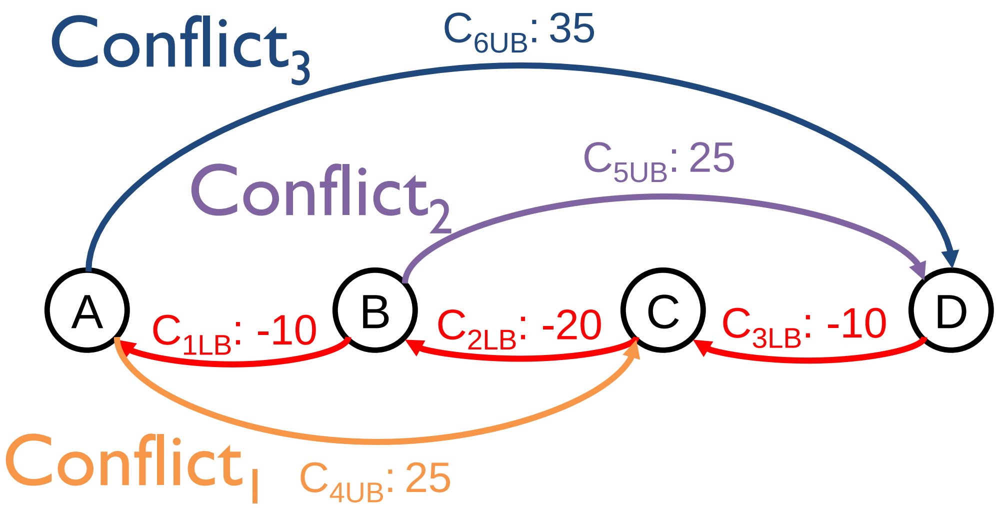

Enumerate Preferred Relaxations

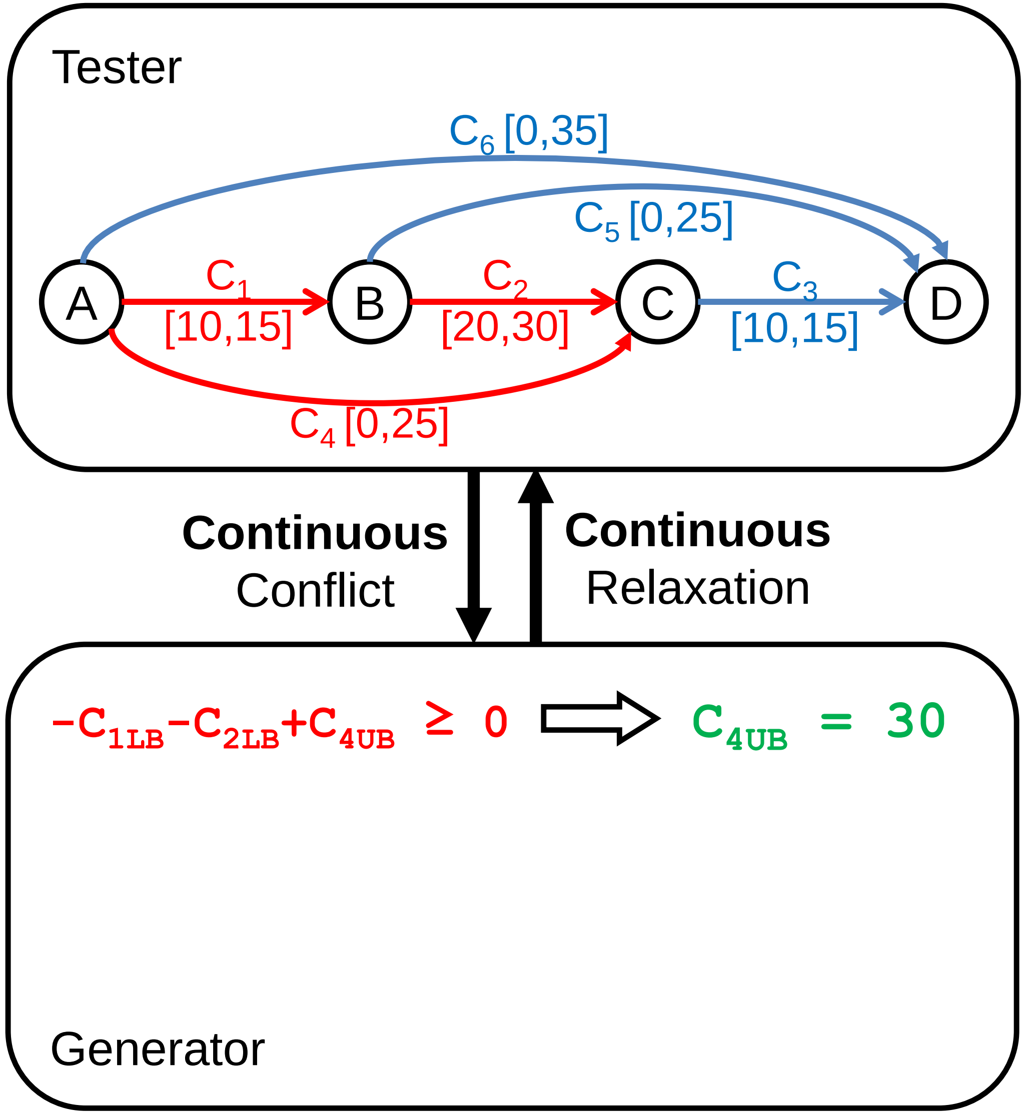

- The **tester** now checks the consistency of the problem, and extract **continuous** conflicts.
- The **generator** computes preferred continuous relaxations subject to **all** known conflicts.

Enumerate Preferred Relaxations

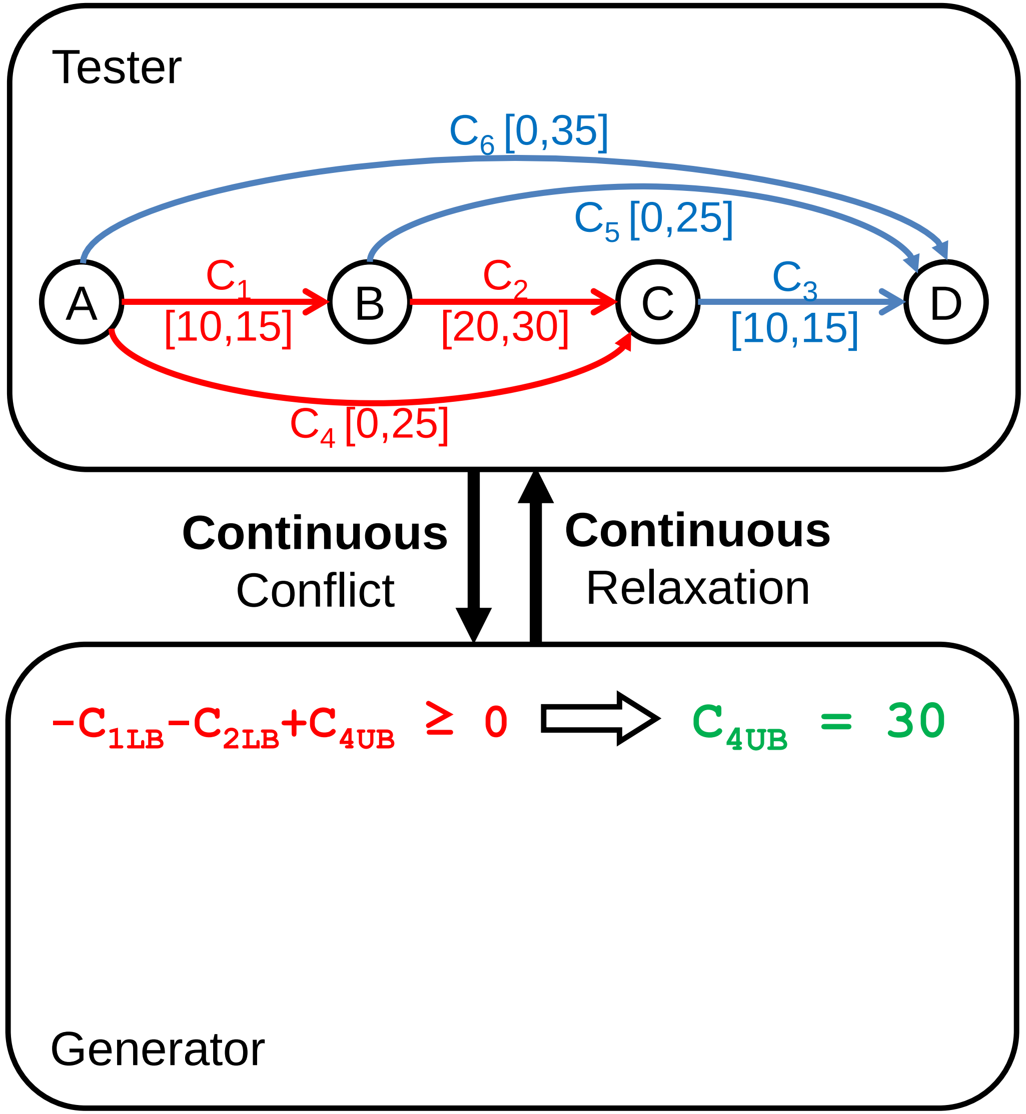

$C_{4UB} = 30$

$\uparrow$

$-C_{1LB}-C_{2LB}+C_{4UB} \geq 0$

$\uparrow$

{ }

Enumerate Preferred Relaxations

Tester

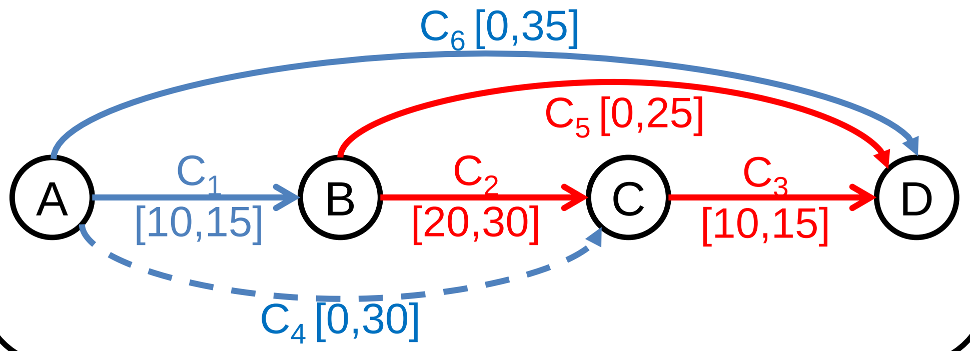

Continuous Conflict

Continuous Relaxation

$-C_{1LB}-C_{2LB}+C_{4UB} \geq 0$

$-C_{2LB}-C_{3LB}+C_{5UB} \geq 0$

$\Rightarrow$ $C_{2LB} = 15$

Generator

$C_{2LB} = 15$

$\uparrow$

$-C_{3LB}-C_{2LB}+C_{5UB} \geq 0$

$\uparrow$

$C_{4UB} = 30$

$\uparrow$

$-C_{1LB}-C_{2LB}+C_{4UB} \geq 0$

$\uparrow$

{ }

Enumerate Preferred Relaxations

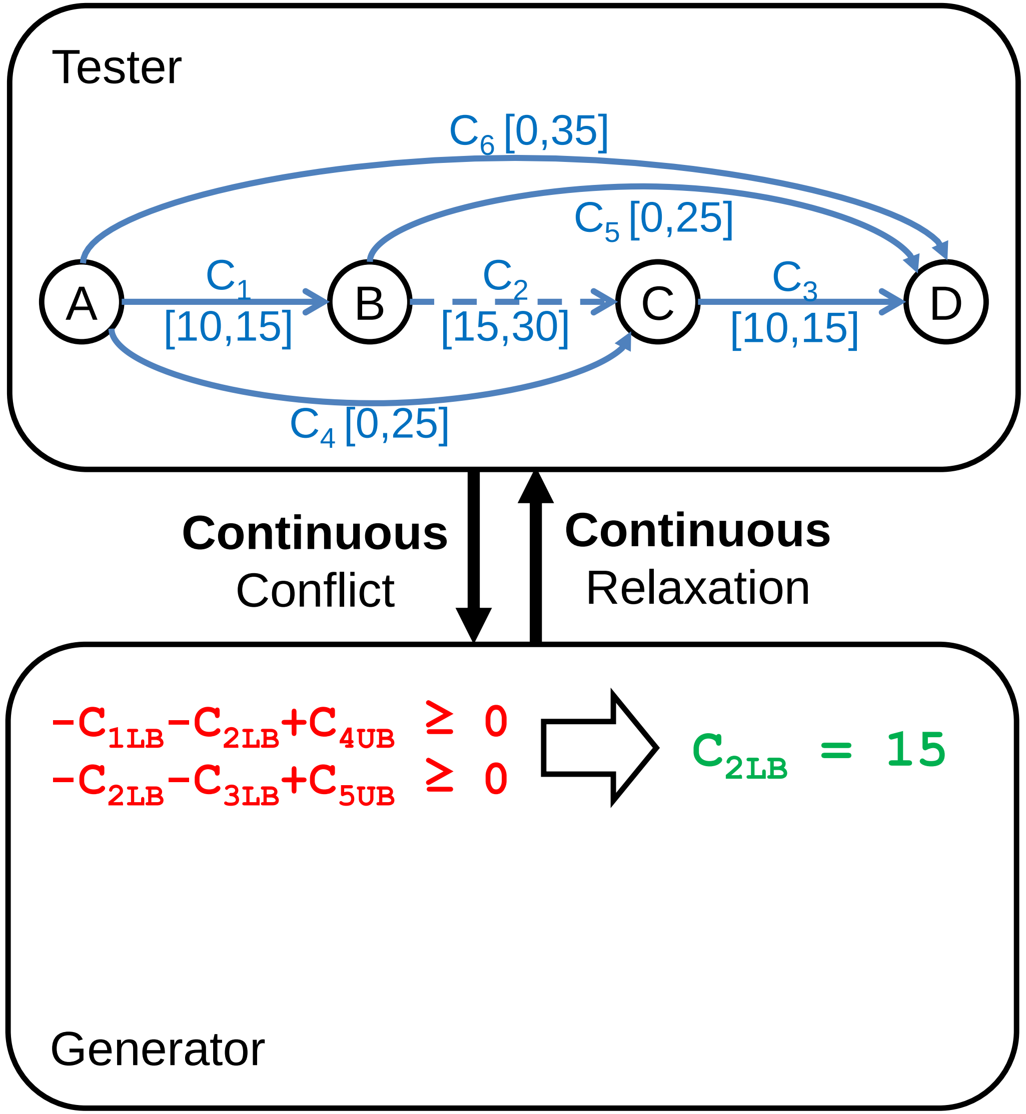

$C_{2LB} = 15$

↑

$-C_{3LB}-C_{2LB}+C_{5UB} \geq 0$

↑

$C_{4UB} = 30$

↑

$-C_{1LB}-C_{2LB}+C_{4UB} \geq 0$

↑

{ }

Enumerate Preferred Relaxations

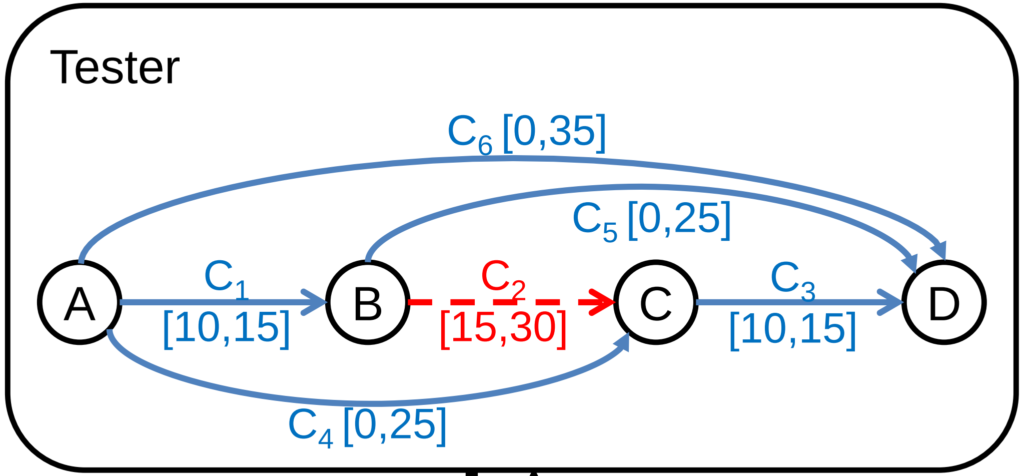

Continuous Conflict

Continuous Relaxation

$$-C_{1LB} - C_{2LB} + C_{4UB} \geq 0$$
$$-C_{2LB} - C_{3LB} + C_{5UB} \geq 0$$
$$C_{2LB} \geq 20$$

$$C_{4UB} = 30$$
$$C_{5UB} = 30$$

User: "The lower bound of C2 has to be at least 20."

Generator

$$C_{4UB} = 30,\ C_{5UB} = 30$$
$$\uparrow$$
$$C_{2LB} \geq 20 \quad \longleftarrow \quad \textbf{User}$$
$$\uparrow$$
$$C_{2LB} = 15$$
$$\uparrow$$
$$-C_{3LB} - C_{2LB} + C_{5UB} \geq 0$$
$$\uparrow$$
$$C_{4UB} = 30$$
$$\uparrow$$
$$-C_{1LB} - C_{2LB} + C_{4UB} \geq 0$$
$$\uparrow$$
$$\{\}$$

Enumerate Preferred Relaxations

Tester

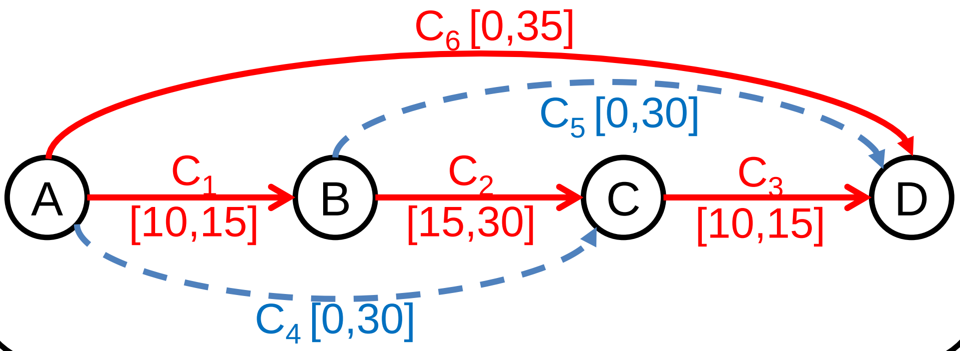

Continuous Conflict

Continuous Relaxation

$-C_{1LB}-C_{2LB}+C_{4UB} \geq 0$

$-C_{2LB}-C_{3LB}+C_{5UB} \geq 0$

$C_{2LB} \geq 20$

$-C_{1LB}-C_{2LB}-C_{3LB}+C_{6UB} \geq 0$

$\Rightarrow$

$C_{4UB} = 30$

$C_{5UB} = 30$

$C_{6UB} = 40$

Generator

$C_{4UB} = 30,\ C_{5UB} = 30,\ C_{6UB} = 40$

$\uparrow$

$-C_{1LB}-C_{2LB}-C_{3LB}+C_{6UB} \geq 0$

$\uparrow$

$C_{4UB} = 30,\ C_{5UB} = 30$

$\uparrow$

$C_{2LB} \geq 20 \leftarrow$ **User**

$\uparrow$

$C_{2LB} = 15$

$\uparrow$

$-C_{3LB}-C_{2LB}+C_{5UB} \geq 0$

$\uparrow$

$C_{4UB} = 30$

$\uparrow$

$-C_{1LB}-C_{2LB}+C_{4UB} \geq 0$

$\uparrow$

{ }

Demonstration

This Presentation

- Best-first Conflict-Directed Relaxation.
- Continuous Relaxation for Temporal Constraints.
- Domain Relaxation for Parameterized Variable.
- Risk-bounded Relaxation under Temporal Uncertainty.

Domain Relaxations for Variable Descriptions.

- We can resolve conflicts by gradually widening the set of options allowed in the plan.

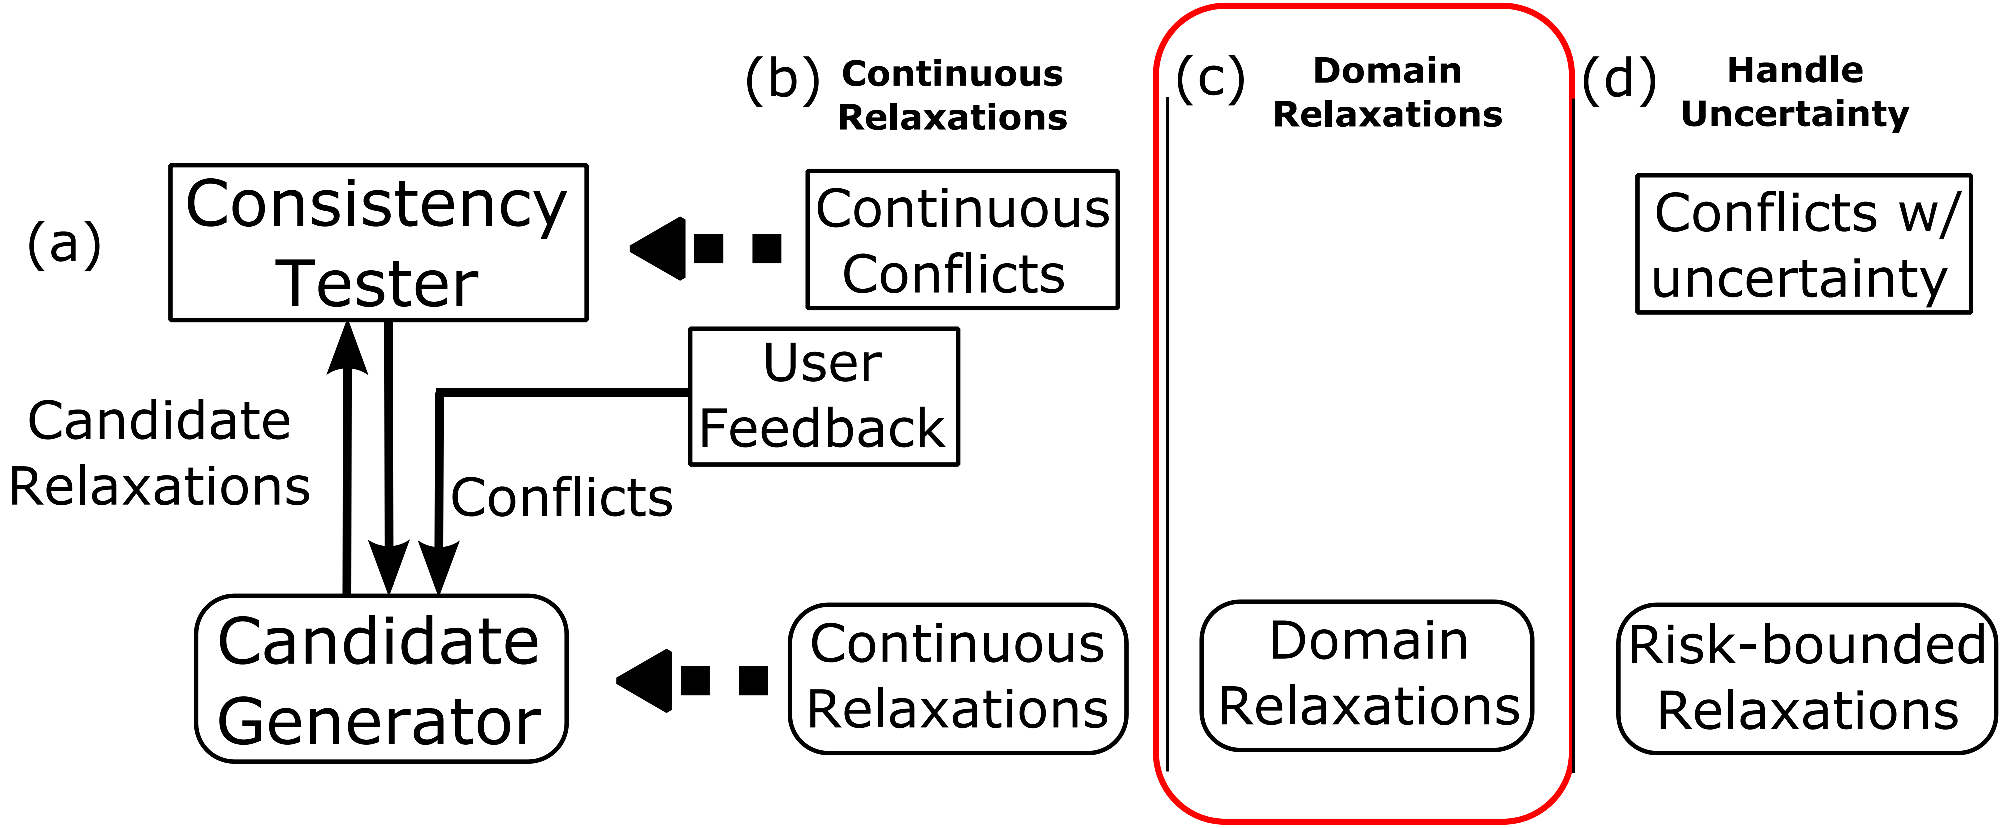

Domain Constraints

- The variable for each activity in the plan is associated with a set of domain constraints, such as:

$\forall r$:

```
cuisine(r)  = Korean
location(r) = Cambridge
Rating(r) ≥ 4
```

POI Database

r = {Koreana, Kimchi Kitchen}

  - The domain of a variable is defined by these constraints.

Resolving Conflicts with Domain Relaxations

- Weakens the domain constraints of variables, which allows additional options to be considered to resolve conflicts.

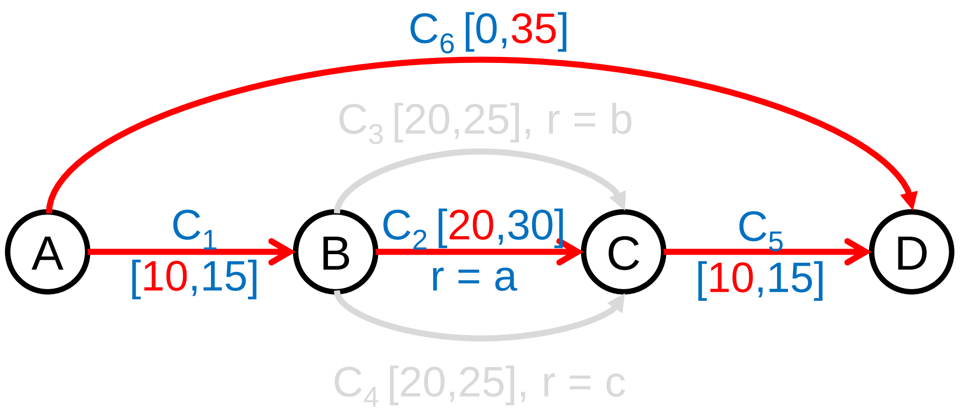

$\forall r$, Rating(r) $\geq$ 4
$\wedge$ Cuisine(r)=Korean
$\wedge$ Location(r)=Cambridge

Discrete constituent relaxations:

{r=b, r=c}

Continuous constituent relaxation:

$-C_{1LB}-C_{2LB}-C_{5LB}+C_{6UB} \geq 0$

Constituent domain relaxations:

{Rating(x)<4,
Cuisine(x)≠Korean,
Location(x)≠Cambridge}

‘Continuous’ Domain Relaxation

- Prior work on domain relaxation suspends constraints to enable more options*.

Rating(r)=4
∧ Cuisine(r)=Korean
∧ Location(r)=Cambridge

Rating(r) ≥ 4
∧ Location(r)
= Cambridge

Cuisine(r)=Korean
∧ Rating(r) ≥ 4

- Instead, we took the ***continuous***† approach to weaken domain constraints until a consistent option becomes available.

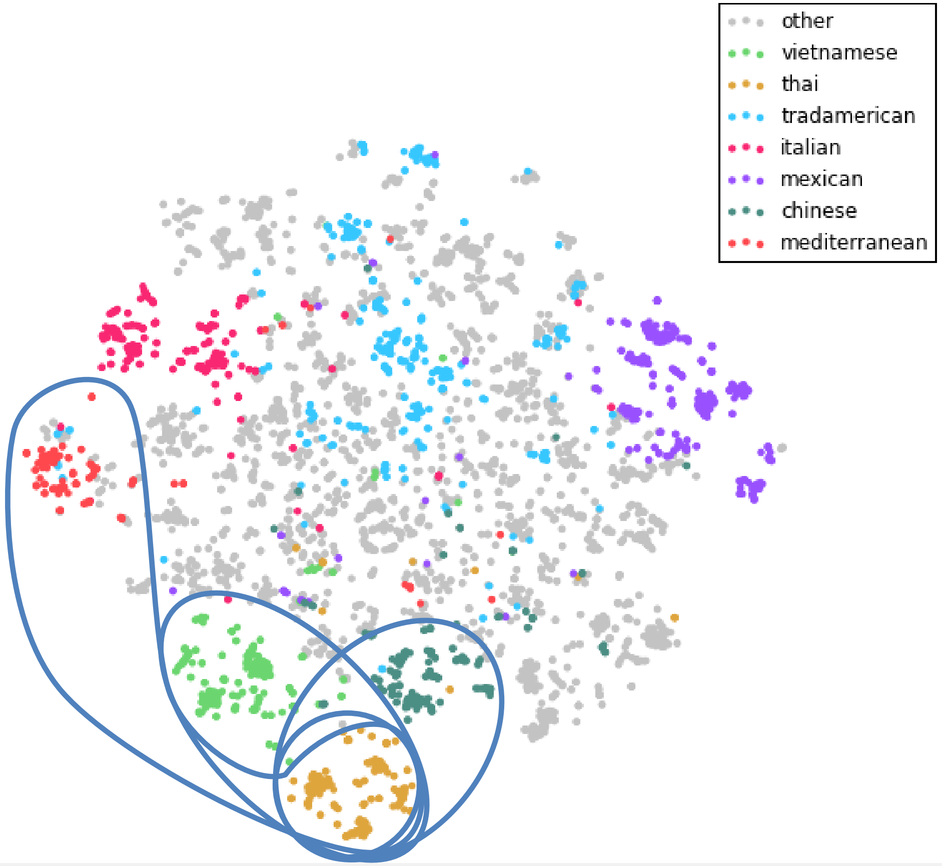

* [Thompson, Goker and Langley, 2004]
† [Raiman, 2015]

Similarity Measurement

- We use a phrase similarity model*, generated by the word2vec package over Freebase concepts†.
  - Supports a total ordering between alternatives.
  - Works across multiple domains.

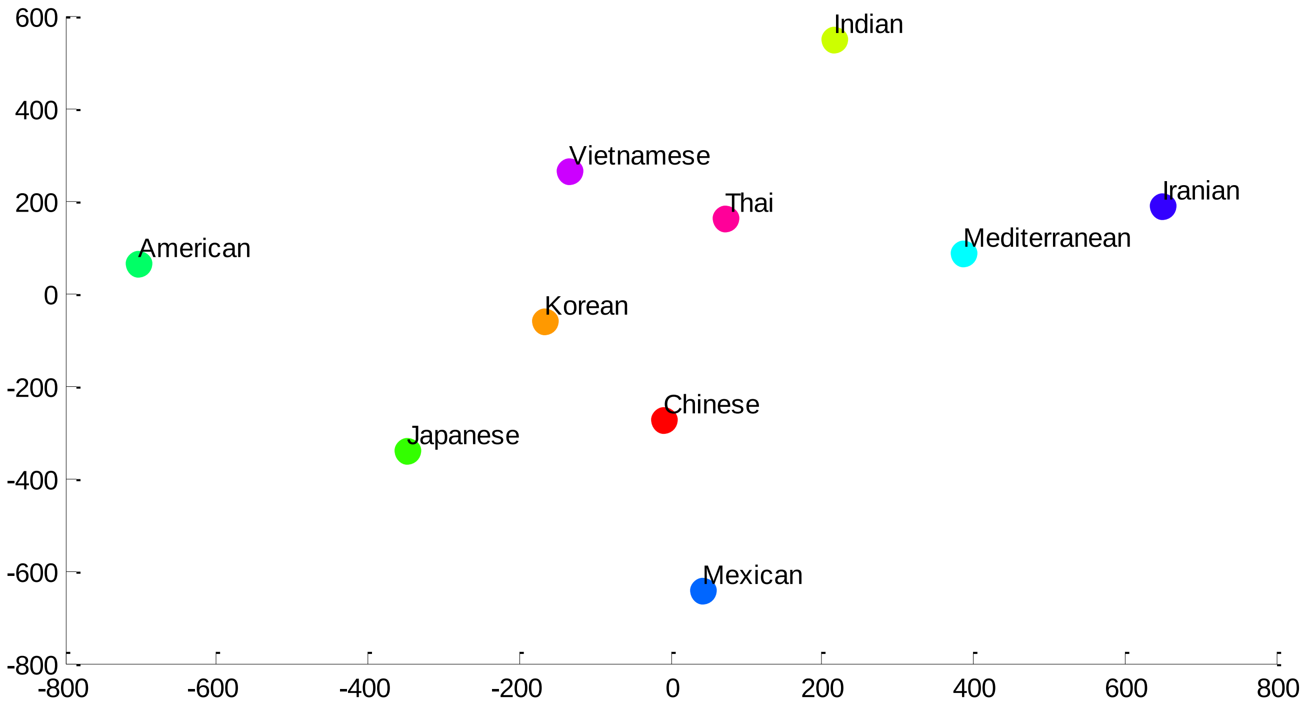

$$s = \cos(\theta) = \frac{\vec{A} \cdot \vec{B}}{\|\vec{A}\|\|\vec{B}\|}$$

Cuisine KOREAN `m.048vr`	→	CHINESE `m.01xw9` 0.7140	→	JAPANESE `m.042ck` 0.6626	→	MEXICAN `m.051zk` 0.4380	→	AMERICAN `m.01z1jf2` 0.3367

* [Mikolov, Sutskever, Chen, Corrado and Dean, 2013]
† [word2vec, Freebase Skipgram Vectors, 2013]

Enumerating Temporal & Domain Relaxations

- BCDR simultaneously computes discrete, temporal and domain relaxations
  - by resolving conflicts using all three types of constituent relaxations*.

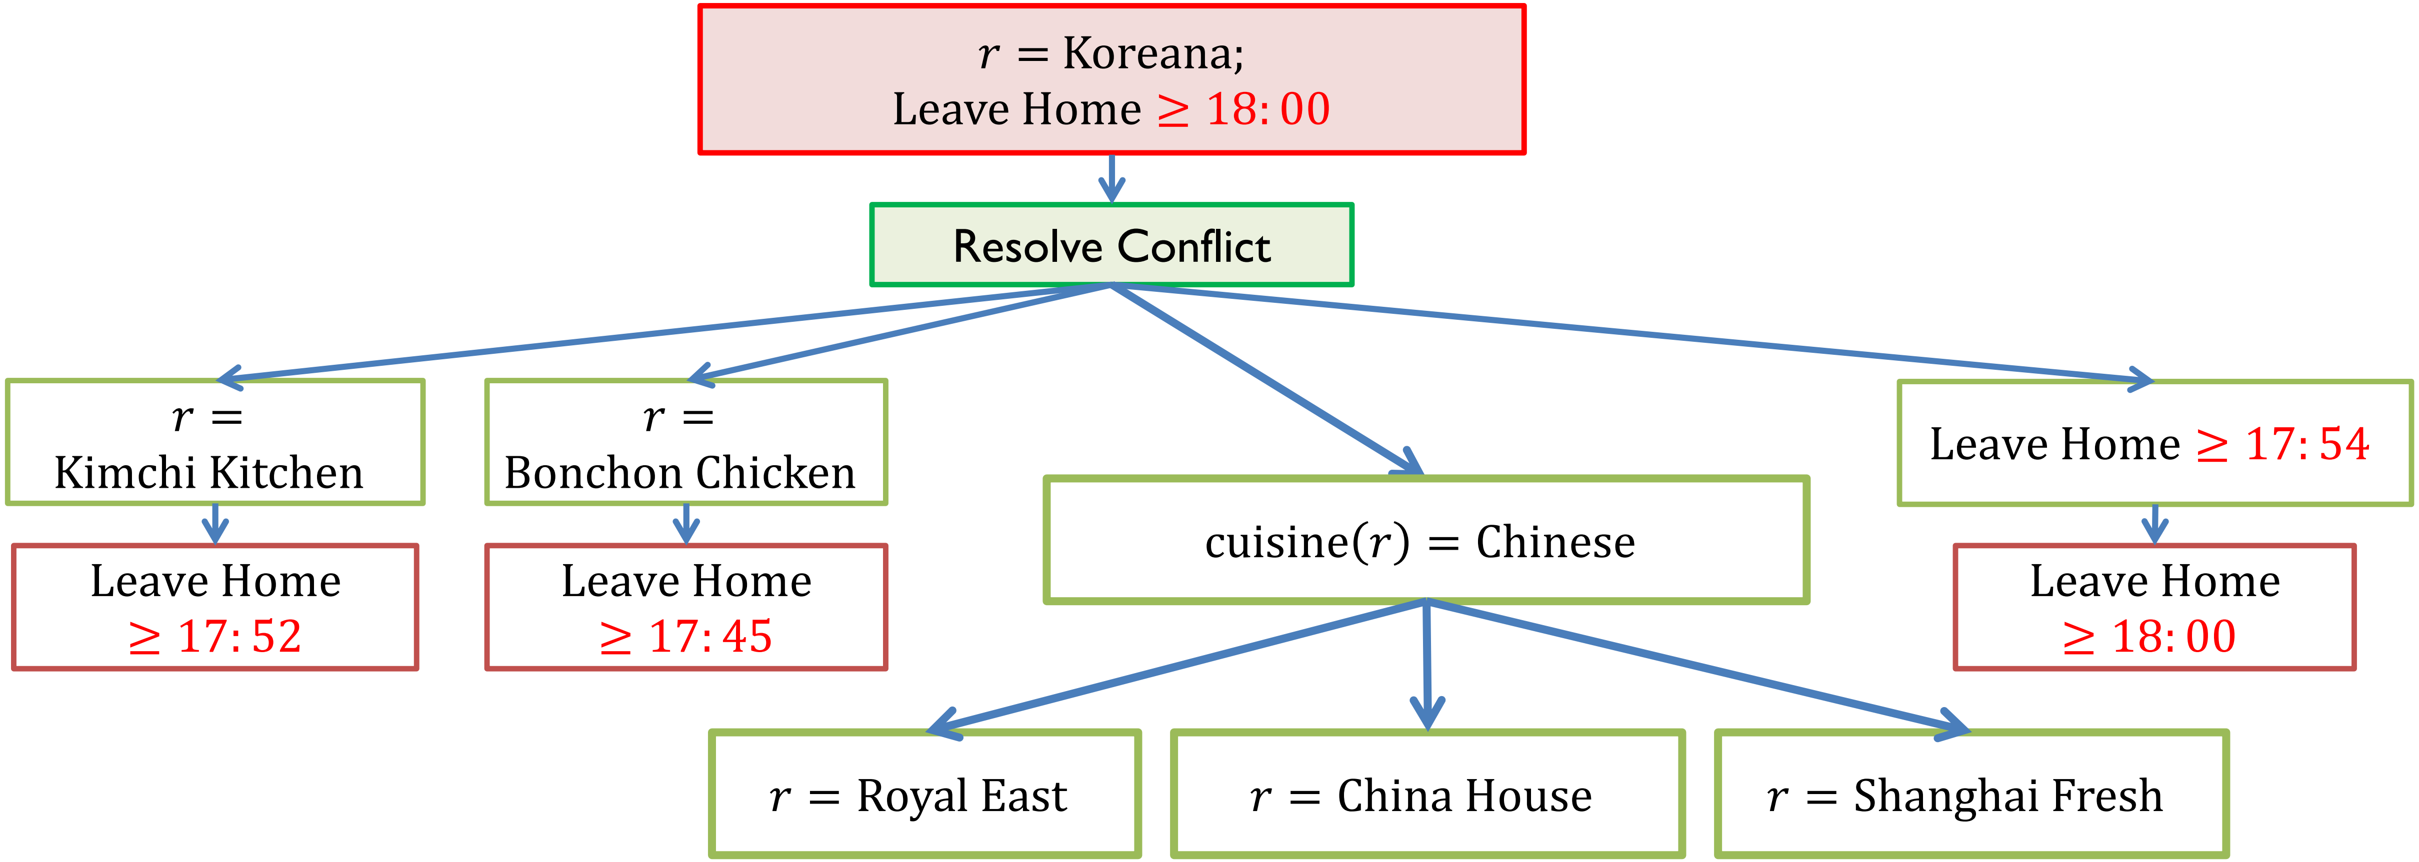

* [Yu, Shen, Yeh and Williams, 2016a]

User Survey

- We invited 9 participants to join a survey for Uhura, by using the travel advisor to manage their day-to-day tasks.

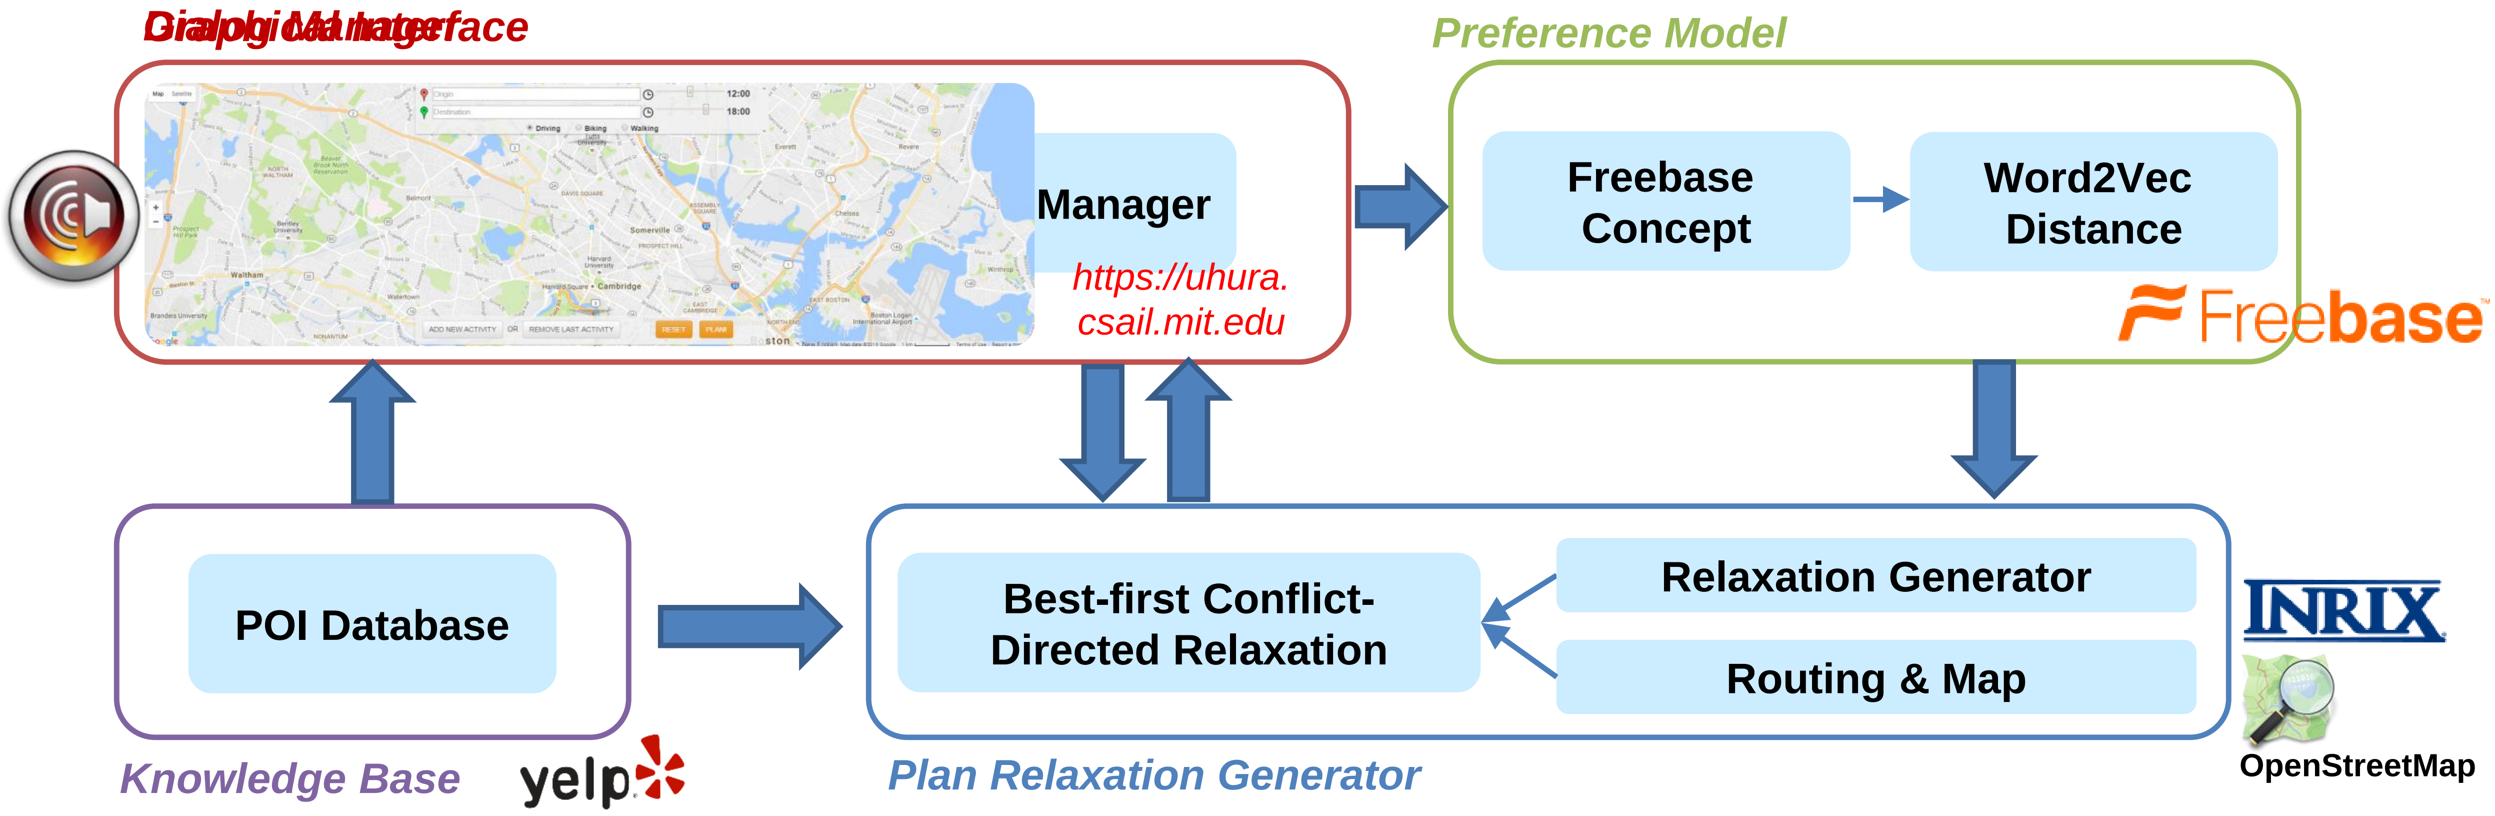

Results for Domain Relaxation

- Uhura found acceptable solutions in **52 out of 54** scenarios.
- With temporal relaxation only approach, solutions found in **43** scenarios.

Session	Temporal Relaxations	Domain Relaxations
1	2.0 (2.6)	2.1 (2.7)
2	1.3 (2.9)	3.0 (3.3)
3	2.9 (3.0)	3.1 (2.8)
4	0.3 (0.7)	1.7 (3.4)
5	1.9 (2.6)	1.7 (3.0)
6	0.6 (1.1)	0.0 (0.0)

This Presentation

- Best-first Conflict-Directed Relaxation.
- Continuous Relaxation for Temporal Constraints.
- Domain Relaxation for Parameterized Variables.
- Risk-bounded Relaxation under Temporal Uncertainty.

Uncertainty in Deep-sea Explorations

- Temporal uncertainty is commonly encountered while operating Remotely Operated Vehicles (ROVs) and AUVs.

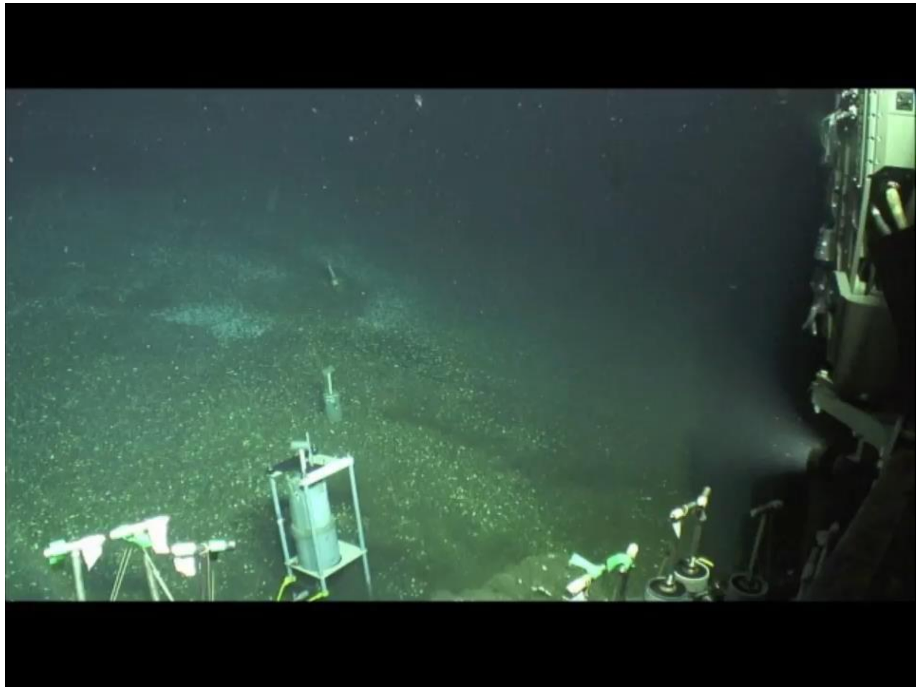

Uncertain traversal and task durations*

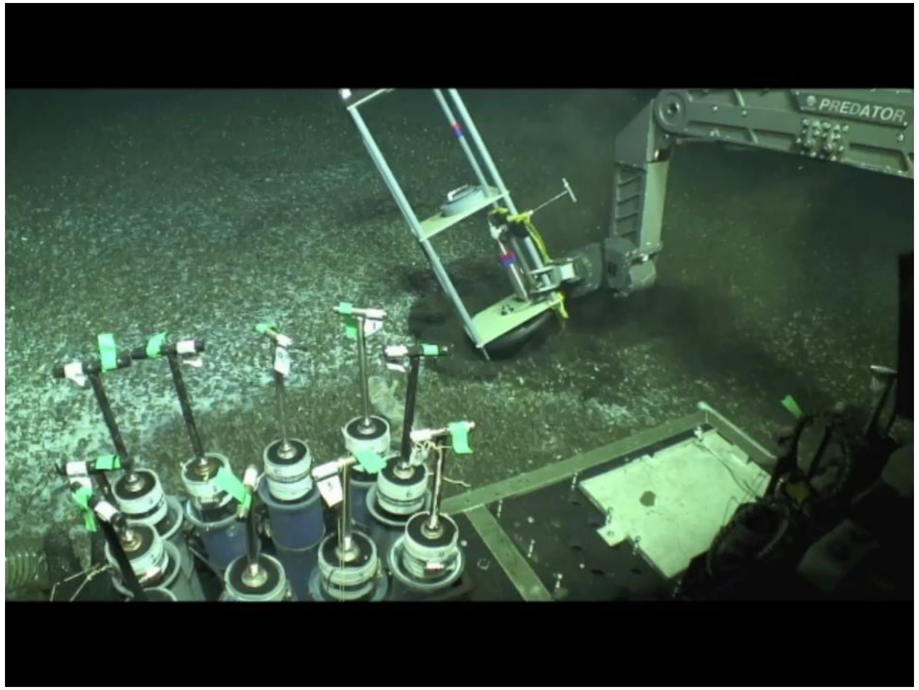

Uncertainty in operator performance*

*Video clips from Jason's cameras, WHOI SEEPS 13 cruise, 09/2013

Risk-bounded Relaxation under Temporal Uncertainty

- We can resolve conflicts by only handling a limited set of outcomes from uncontrolled variables.
  - Instead of widening the domain of controlled ones.

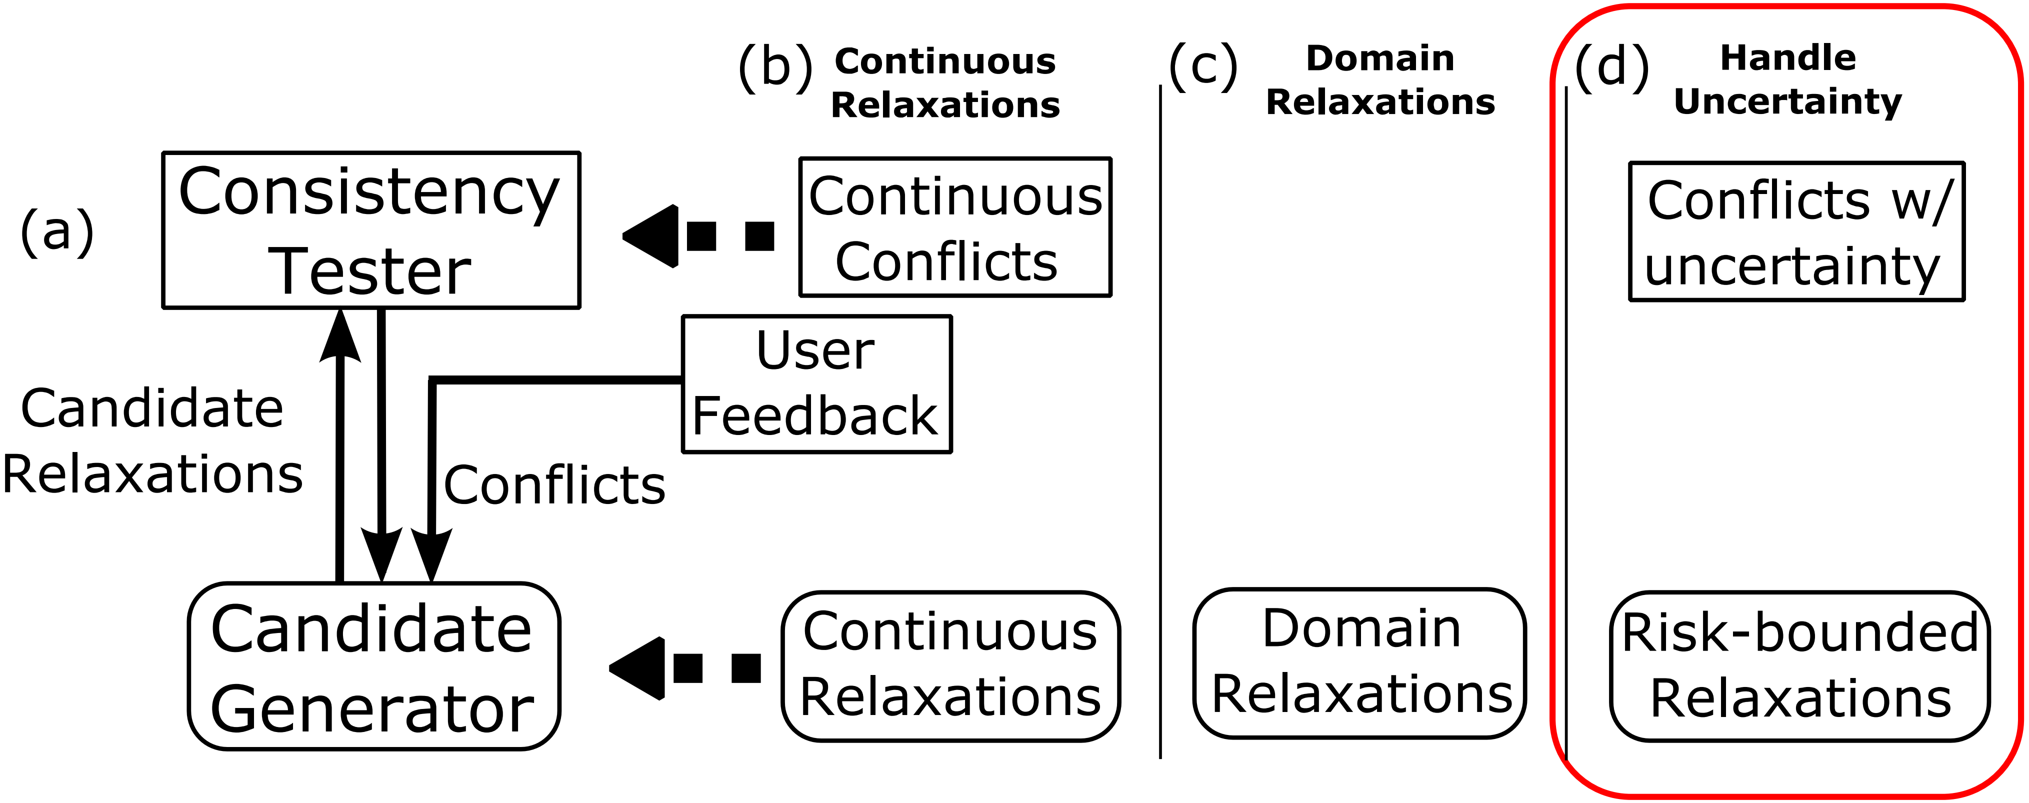

Modeling Uncertain Durations

- An uncertain duration $e_i \in E_u$ is modeled using a random variable ω.
  - ω may be described using a probabilistic distribution, which encodes the likelihood of different outcomes*.

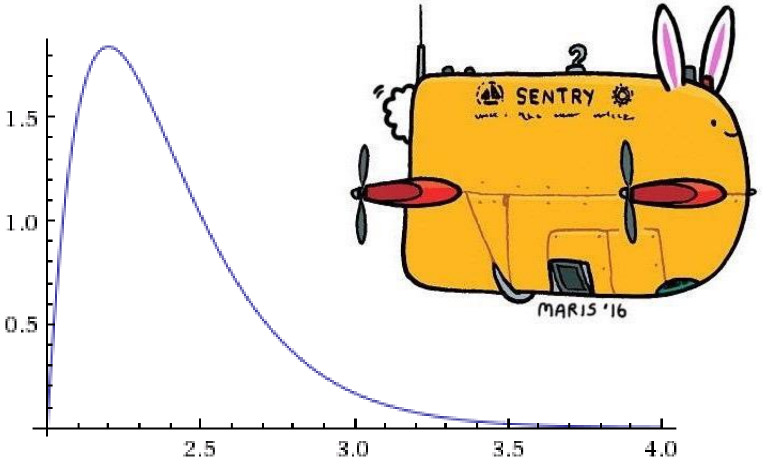

* Probabilistic Simple Temporal Problems, [Tsamardinos, 2002]

Encoding Risk Bound

- Introducing Δ_t, the required bound on the risk of failure.
  - Given a solution, the chance of violating any temporal constraints in the plan must be less than Δ_t.

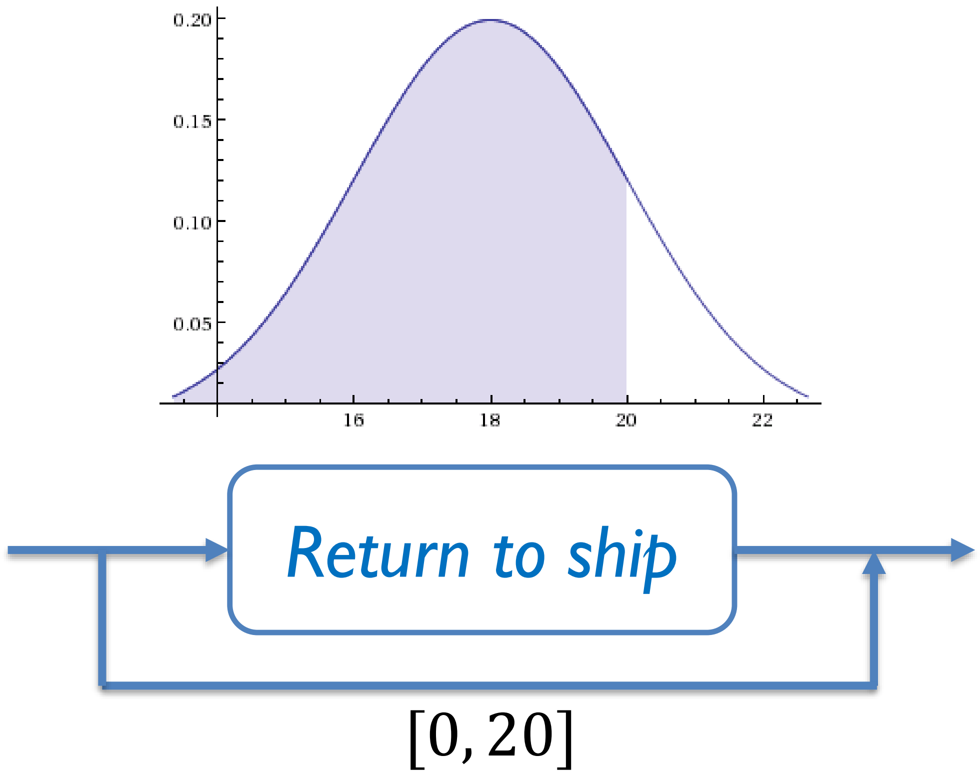

$$\Delta_t = 20\%$$

Risk-Bounded Relaxations

- Over-subscription means no solution exists that provides the required guarantee of success.

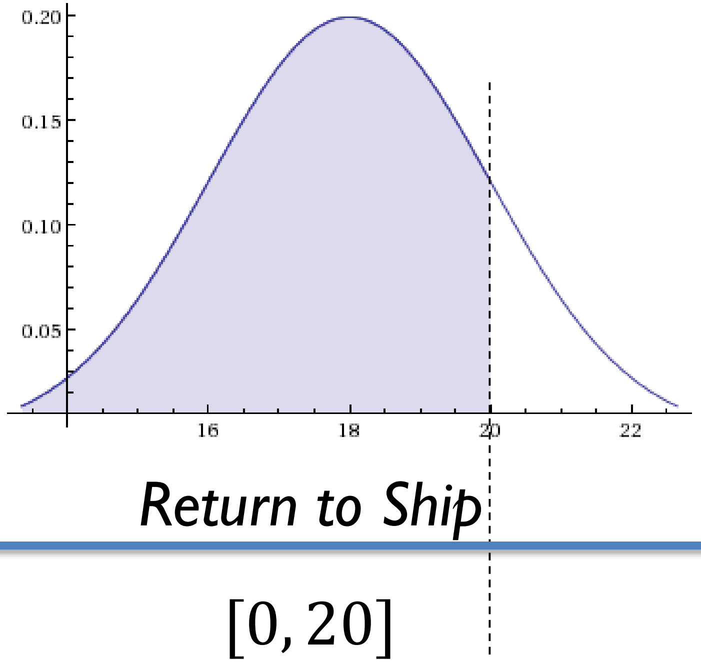

Return to Ship

$[0, 20]$

$\Delta_t = 10\%$

$[0, 20] \rightarrow [0, 20.57]$

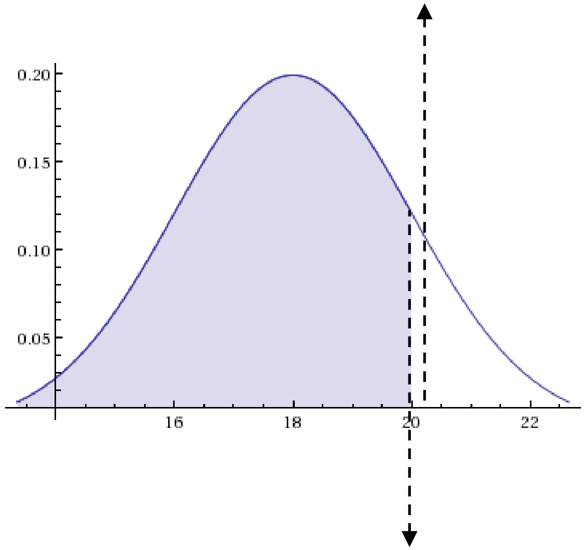

$\Delta'_t = 15.9\%$

Change of Workflow

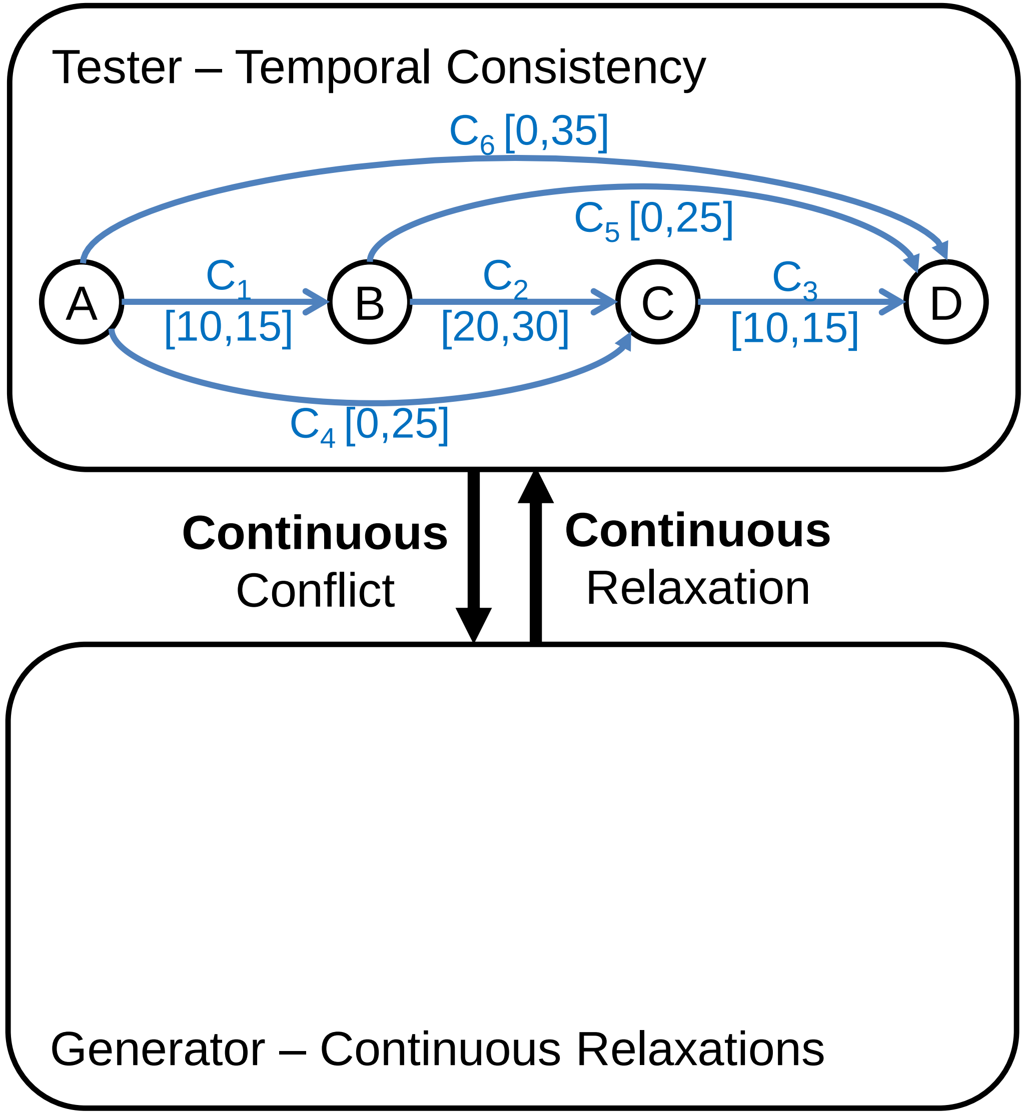

Change of Workflow

Tester – Temporal Controllability

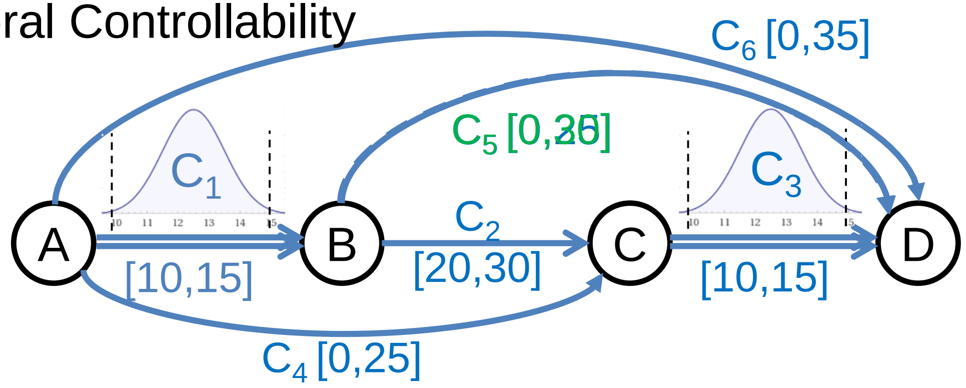

Continuous Conflict w/ **Uncertain** Duration

Relaxations **and Risk Allocation**

$$-C_{2LB} - C_{3UB} + C_{5UB} \geq 0$$

$$Risk(C_1) + Risk(C_3) \leq \Delta_t$$

$$C_{5UB} = 30$$

$$C_{1LB} = 10,\ C_{1UB} = 15$$
$$C_{3LB} = 10,\ C_{3UB} = 15$$

$$\Delta_t = 0.05$$

Generator – Risk-bounded Relaxations

Learning Conflicts with Uncertain Durations

- During the reduction for controllability checking*, **record the 'contribution'** of each constraint and duration to the distance graph†.

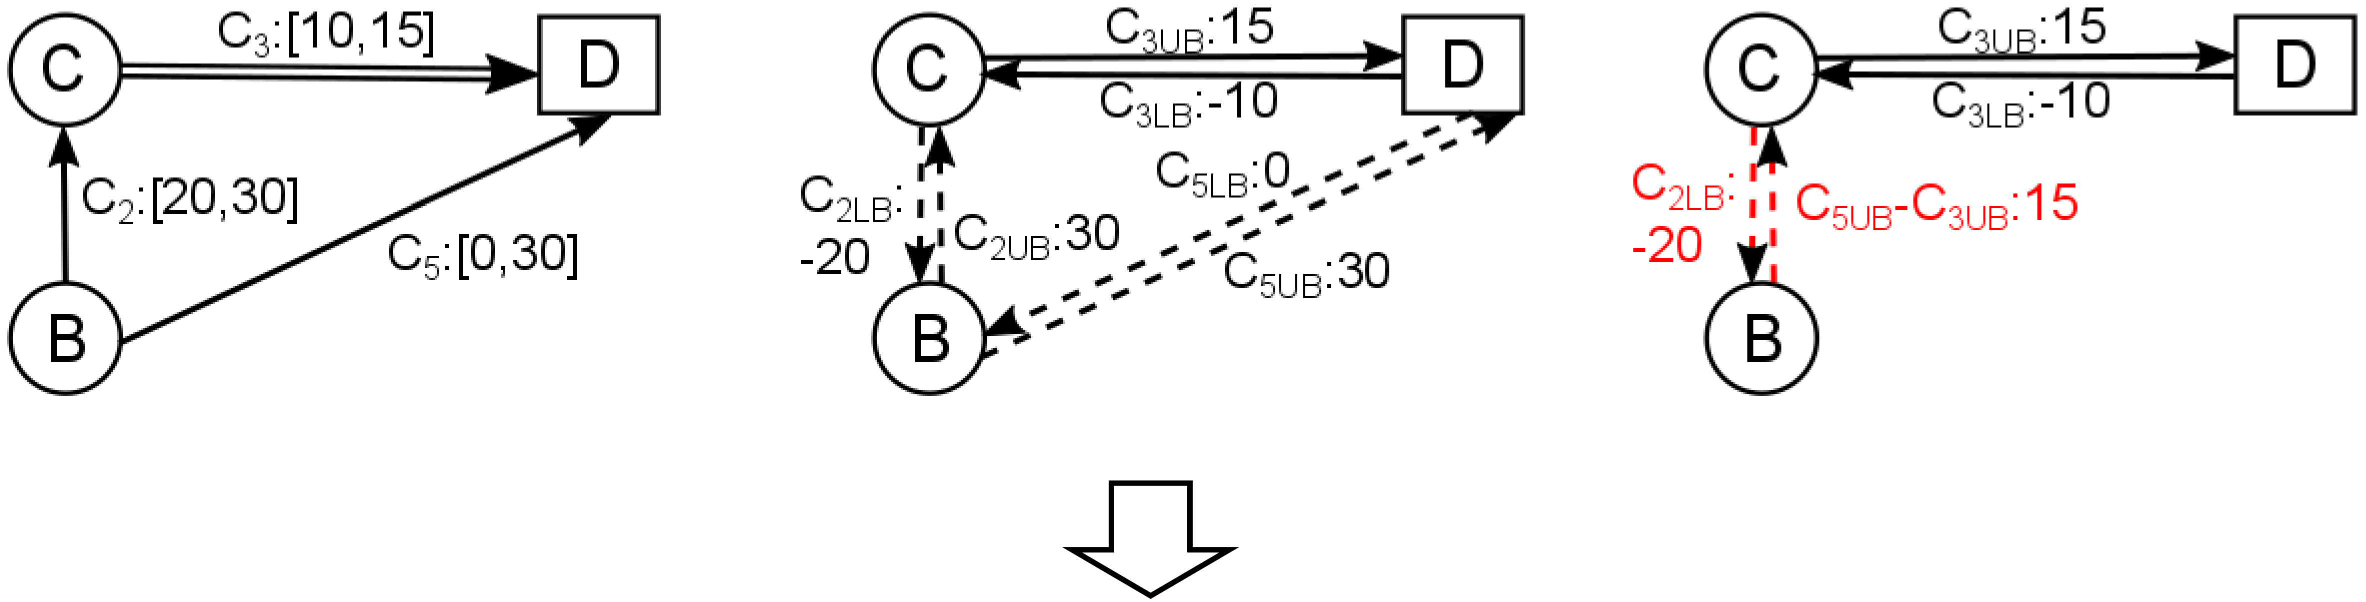

$$-C_{2LB}-C_{3UB}+C_{5UB} \geq 0$$

* [Vidal and Fargier, 1999]

†[Yu, Fang and Williams, 2014], [Karpas, Levine, Yu and Williams, 2015]

Computing Risk-bound Relaxations

Minimize $\sum_i f_{tr}(tr_i) + f_{\Delta_t}(\Delta_t')$

- Minimize the **cost** of temporal (tr_i) and risk bound ($\Delta_t{}'$) relaxations†.

Subject to $conflict_1 \geq 0$
$conflict_2 \geq 0$
… …

$$\sum_j (1 - \int_{lb_j}^{ub_j} p(\omega_j)) \leq \Delta_t{}'$$

- Respect all continuous **constituent** relaxations.
- Bound the **risk allocated** over all uncertain durations (ω_j)*.

* [Fang, Yu and Williams, 2014], [Wang and Williams, 2015]
† [Yu, Fang and Williams, 2015]

Empirical Evaluations

- In addition to travel planning, BCDR has been applied to problems in the following domains:
  - managing autonomous underwater vehicle (AUV) operations;
  - vehicle dispatching for transit routes;
  - evaluating the robustness of job-shop schedules.

- Expriment domains:
  - AUV missions – **computational cost** for handling uncertainty.
  - Transit route management - **performance against** state-of-the-art solvers.

Scheduling AUV Missions

- Scheduling missions for AUV *Sentry*.
  - Randomly generated relay missions of multiple dives and vehicles.
  - Compute continuous relaxations for over-subscribed requirements, subject to uncertain traversal and task durations.

Example AUV Relay Problems

Empirical Results - Runtime

- We evaluate the number of instances solved in 30sec by:
  - BCDR (consistency).
  - BCDR-C(SC) (risk-bounded with strong controllability).
  - BCDR-C(DC) (risk-bounded with dynamic controllability).

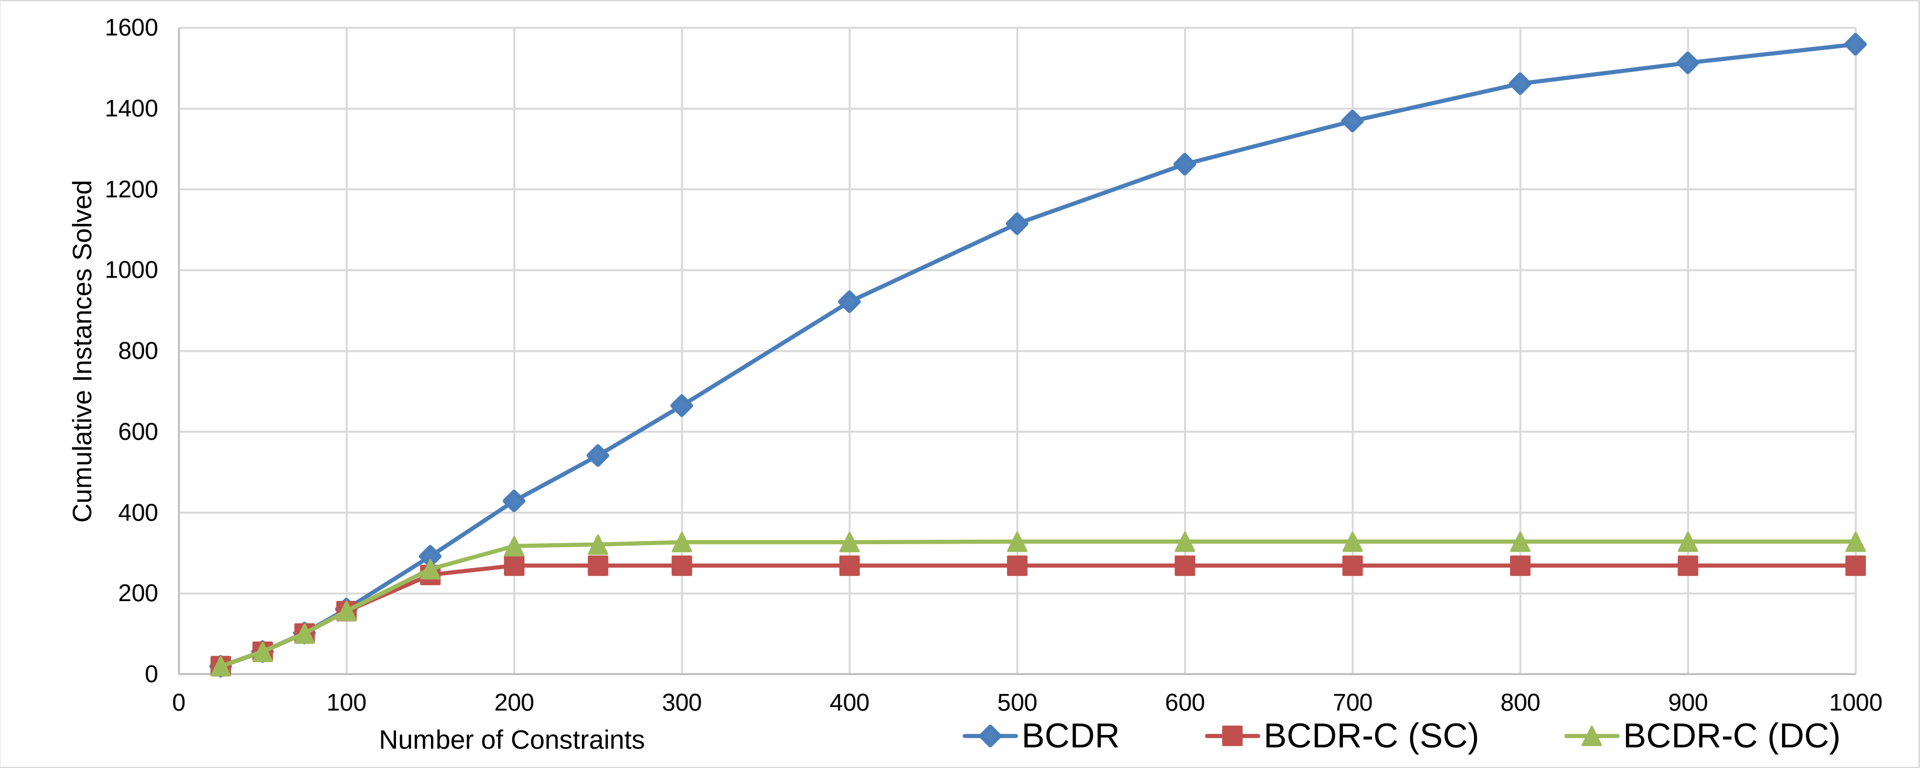

Empirical Results - Conflicts

- The numbers of conflicts resolved by the three algorithms.

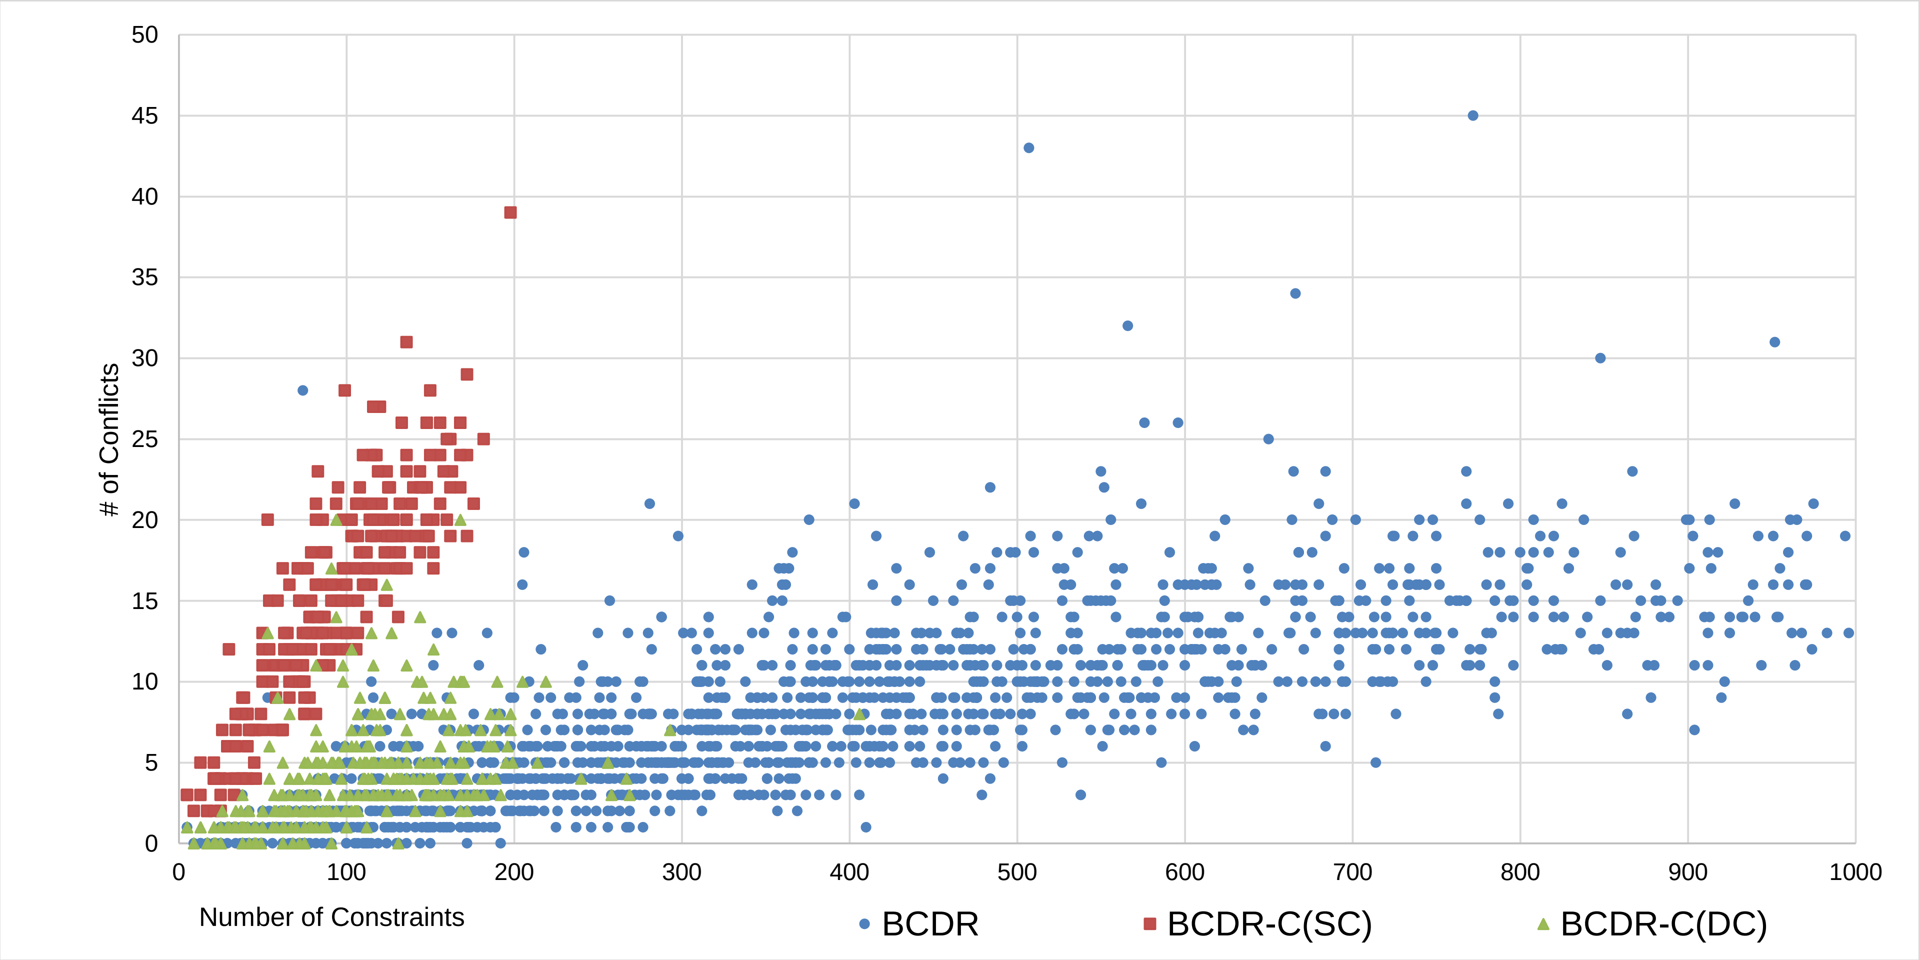

Handling Uncertainty is Important

- We ran simulations to check the robustness of relaxation results.
  - For each solvable case, we generate 100 scenarios by randomly sampling the outcomes of uncertain durations.

	BCDR	BCDR-C (SC)	BCDR-C (DC)
# All Test Problems	2400	2400	2400
# Solved in 30sec	1589	269	328
# Simulations	158900	26900	32800
# Consistent	2703	26900	32800
Ratio	1.701%	100%	100%

Coordinating Vehicles in Transit Routes

- Dynamic dispatching and holding strategies for Red Line trains.
  - To maintain the planned headways and prevent ‘train bunching’*.

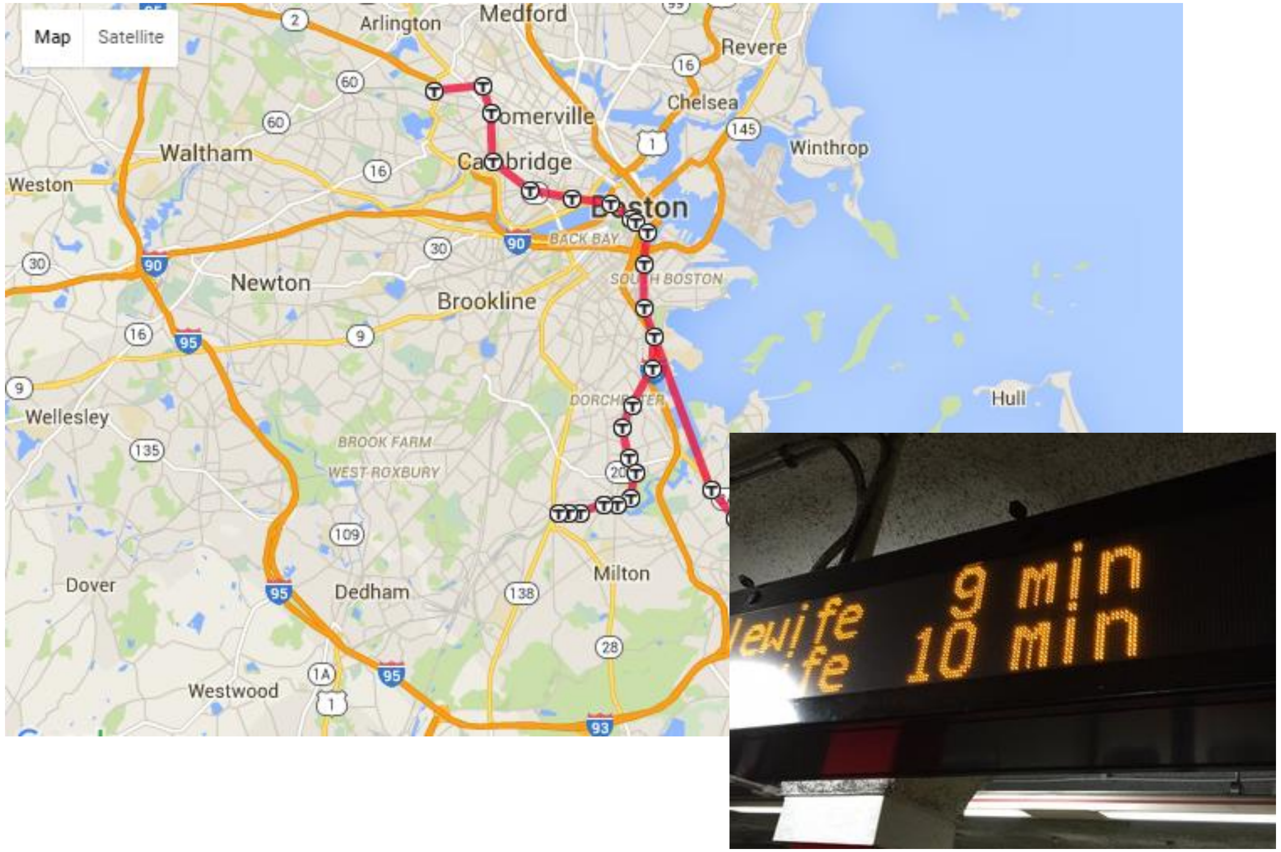

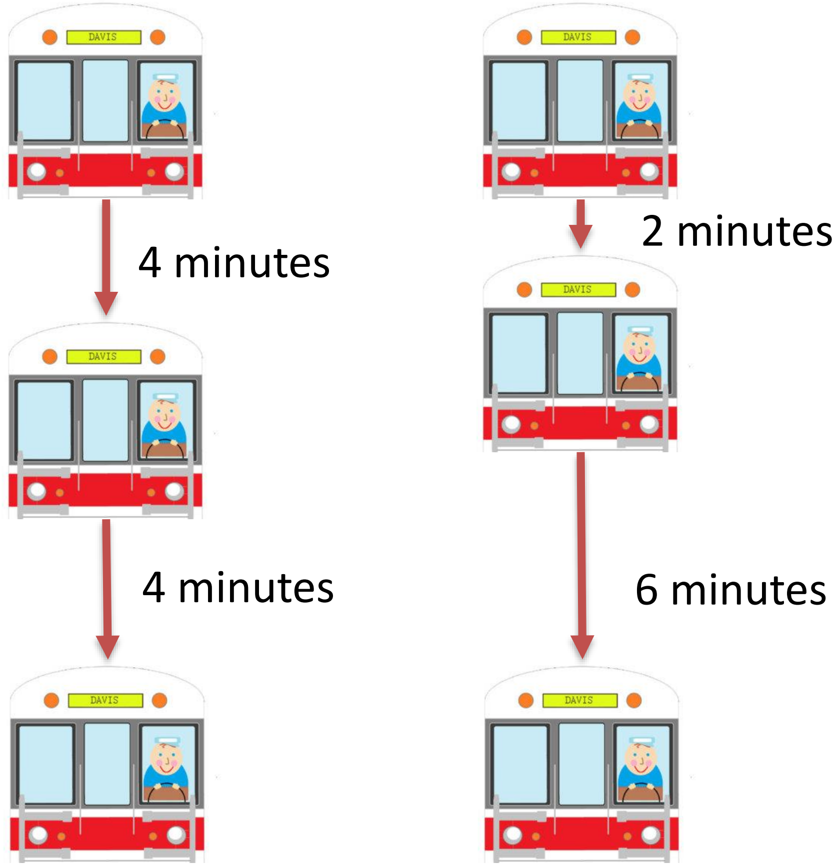

* [Newell and Potts, 1964]

Dynamic Dispatching and Holding Strategies

- Generating a dynamically controllable plan through:
  - adjusting the departure times from origin station.
  - adding additional hold times at various stations.

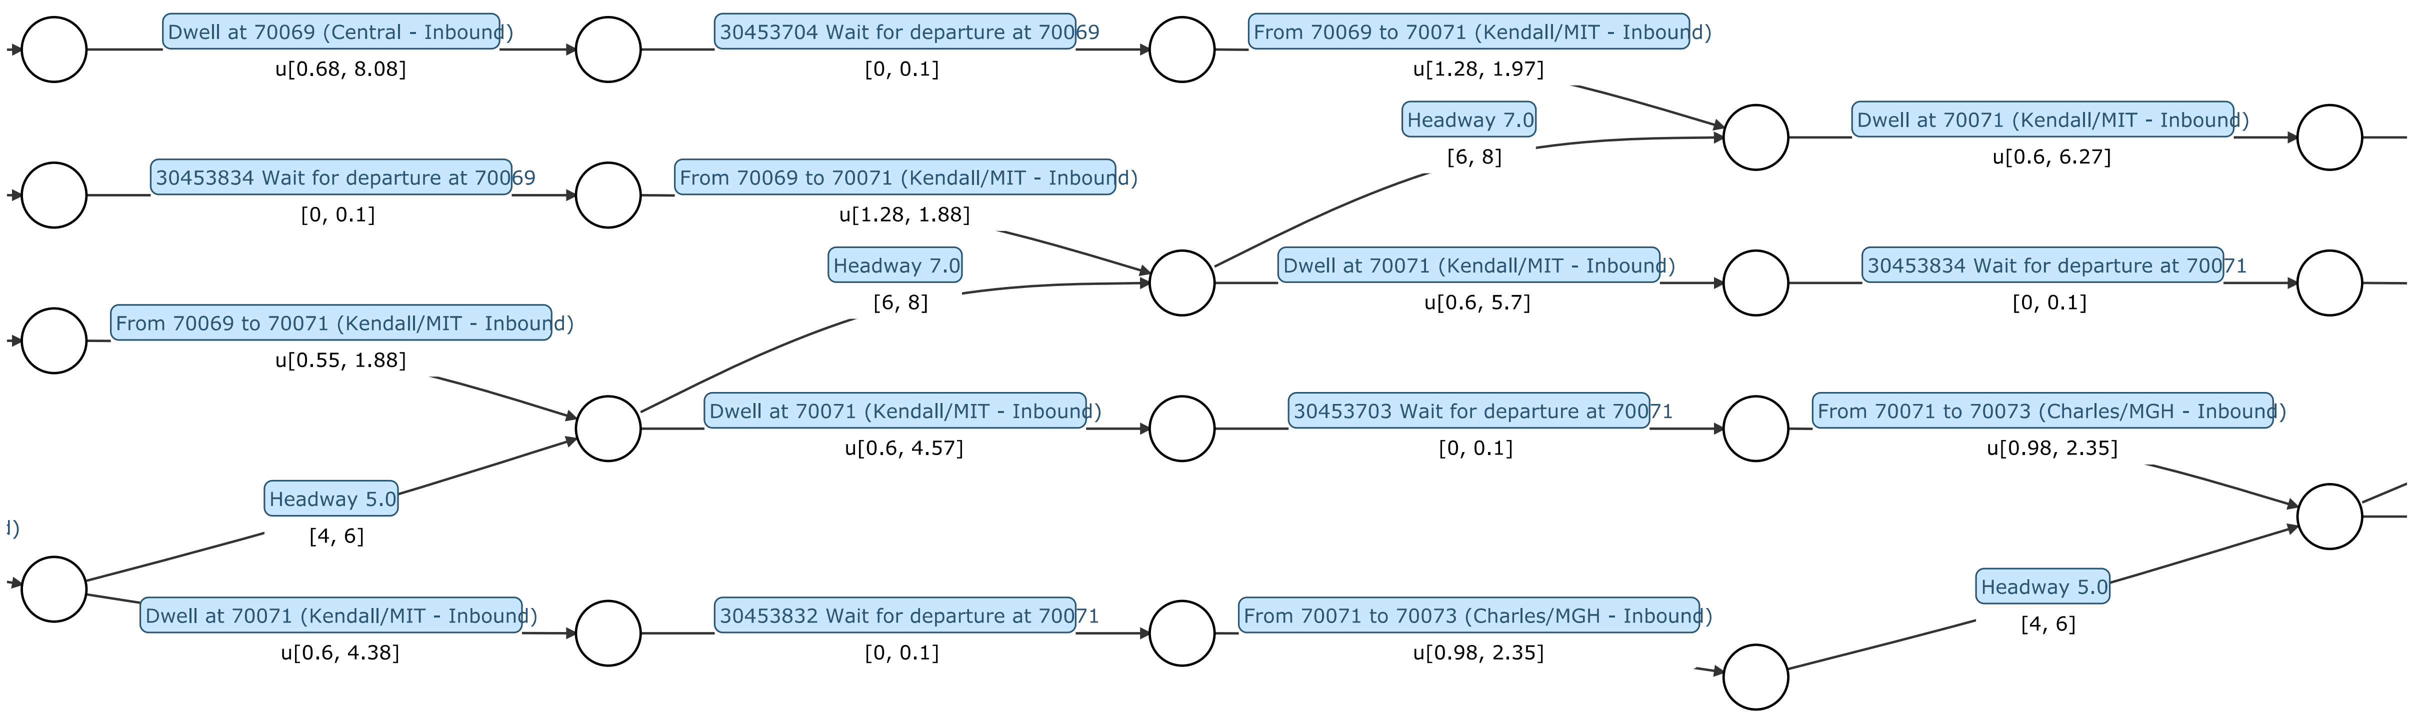

Empirical Results - Runtime

- Compare the performance of BCDR-U(DC) and GuRoBi (with a MIP encoding*).
  - Solve subsets of the original problem and record runtime (timeout: 600sec).

# of Trains

BCDR

7	183	x	x	x	x	x	x	x
6	87.1	x	x	x	x	x	x	x
5	27	315	x	x	x	x	x	x
4	7.16	120	370	x	x	x	x	x
3	1.94	18.4	57.3	162	424	x	x	x
2	0.63	3.22	9.79	19.9	51.4	80.9	201	346
	2	**3**	**4**	**5**	**6**	**7**	**8**	**9**

MIP+GuRoBi

7	x	x	x	x	x	x	x	x
6	x	x	x	x	x	x	x	x
5	x	x	x	x	x	x	x	x
4	x	x	x	x	x	x	x	x
3	30.4	x	x	x	x	x	x	x
2	2.7	12.6	76.6	340	x	x	x	x
	2	**3**	**4**	**5**	**6**	**7**	**8**	**9**

# of Stops

* [Wah and Xin, 2006]

Empirical Results - Conflicts

- Number of conflicts detected by BCDR.

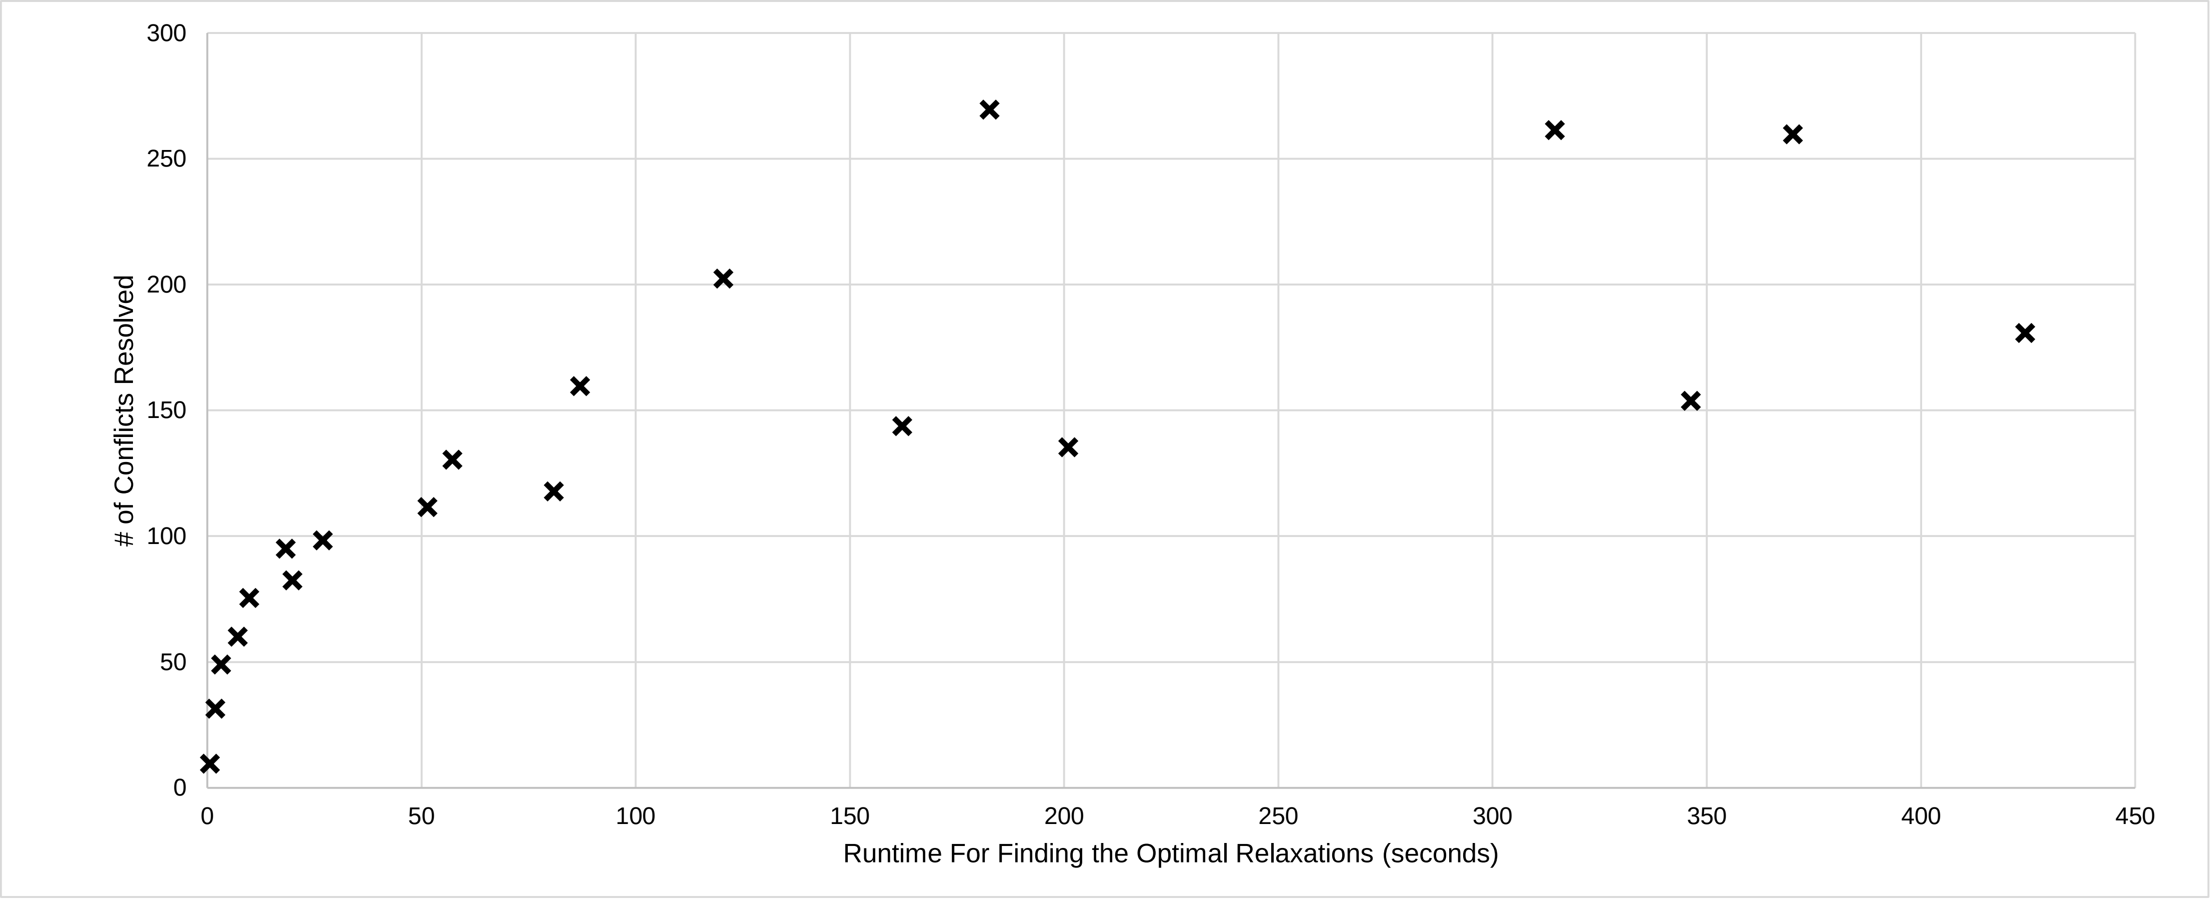

Contributions

Resolving over-subscribed plans using a **variety** of **efficient** partial relaxation techniques leads to greater **flexibility** in plan adaptation.

(b) **Continuous Relaxations**

(c) **Domain Relaxations**

(d) **Handle Uncertainty**

(a) Consistency Tester ◀ ■ ■ Continuous Conflicts

Conflicts w/ uncertainty

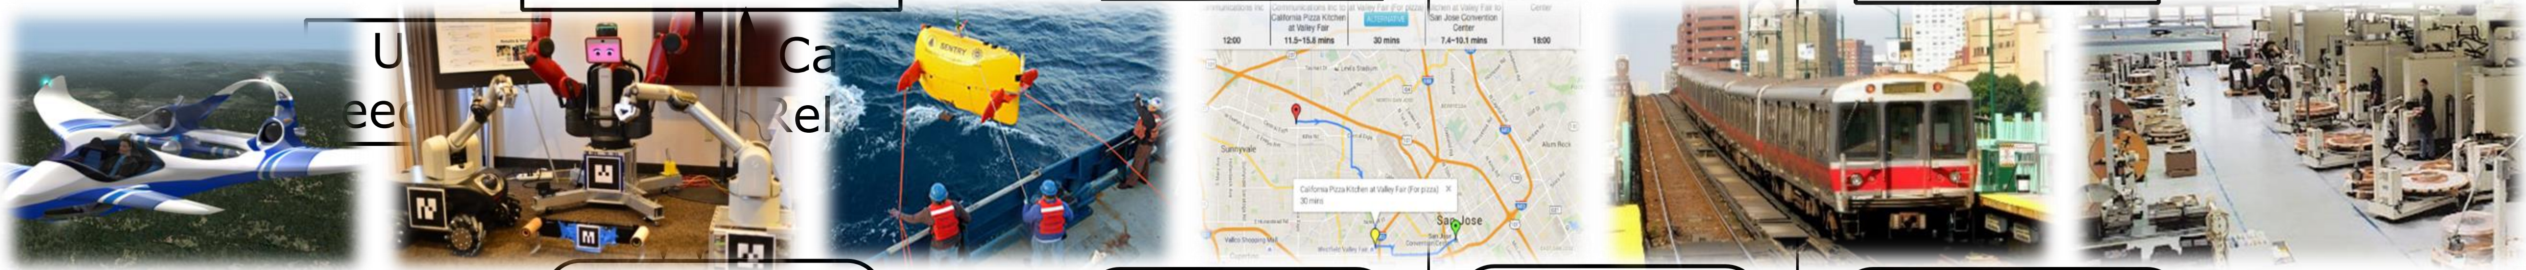

Candidate Generator ◀ ■ ■ Continuous Relaxations

Domain Relaxations

Risk-bounded Relaxations

Future work - Applications

- Driving coach for improving safety.

A relaxation approach for ***generative***, and contingent planning.

Visit the grocery store if you want to make roast turkey tonight.

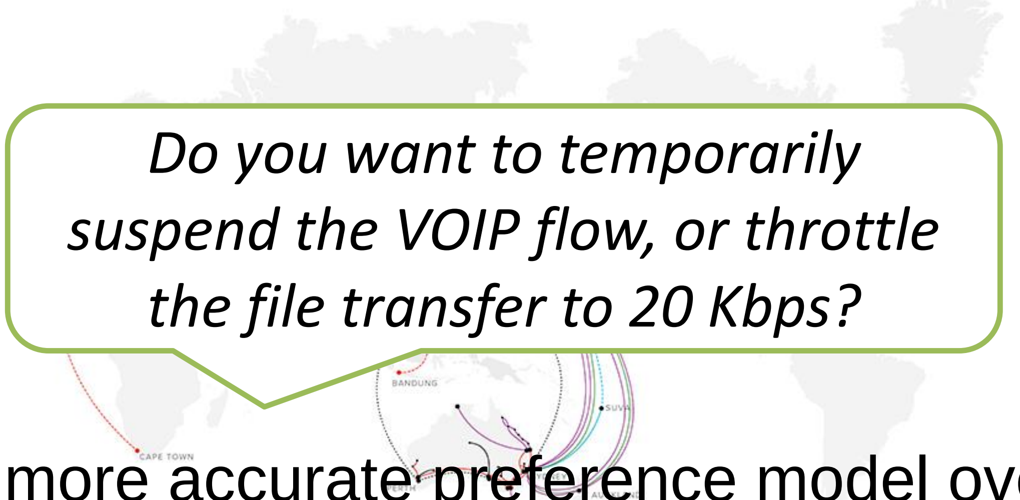

Do you want to temporarily suspend the VOIP flow, or throttle the file transfer to 20 Kbps?

A more accurate preference model over ***discrete*** and ***continuous*** choices.

- Configuration advisor for reliable network communication.

A more efficient approach for generating ***good*** relaxations for large-scale plans.

- Transit network management assistant.

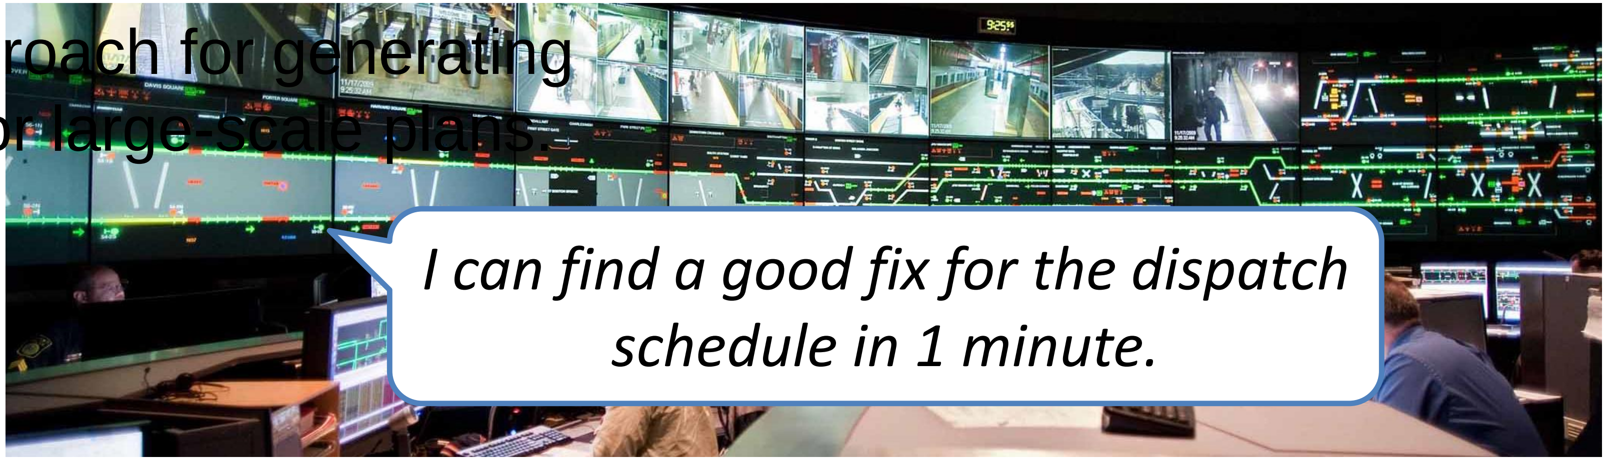

I can find a good fix for the dispatch schedule in 1 minute.

Acknowledgements

- Professor Brian Williams and all MERS members:
  - David, Eric, Erez, Christian, Simon, Andreas, Steve, Andrew, Szymon, Jonathan, Enrique, Tiago, Ben, James, Pedro, Dan, Ameya, Spencer, Shannon, Hiro, Wesley, Larry, Sean, Bobby, Yuening, Nikhil, Matt, Jingkai, Cyrus, Sang and Askhan.

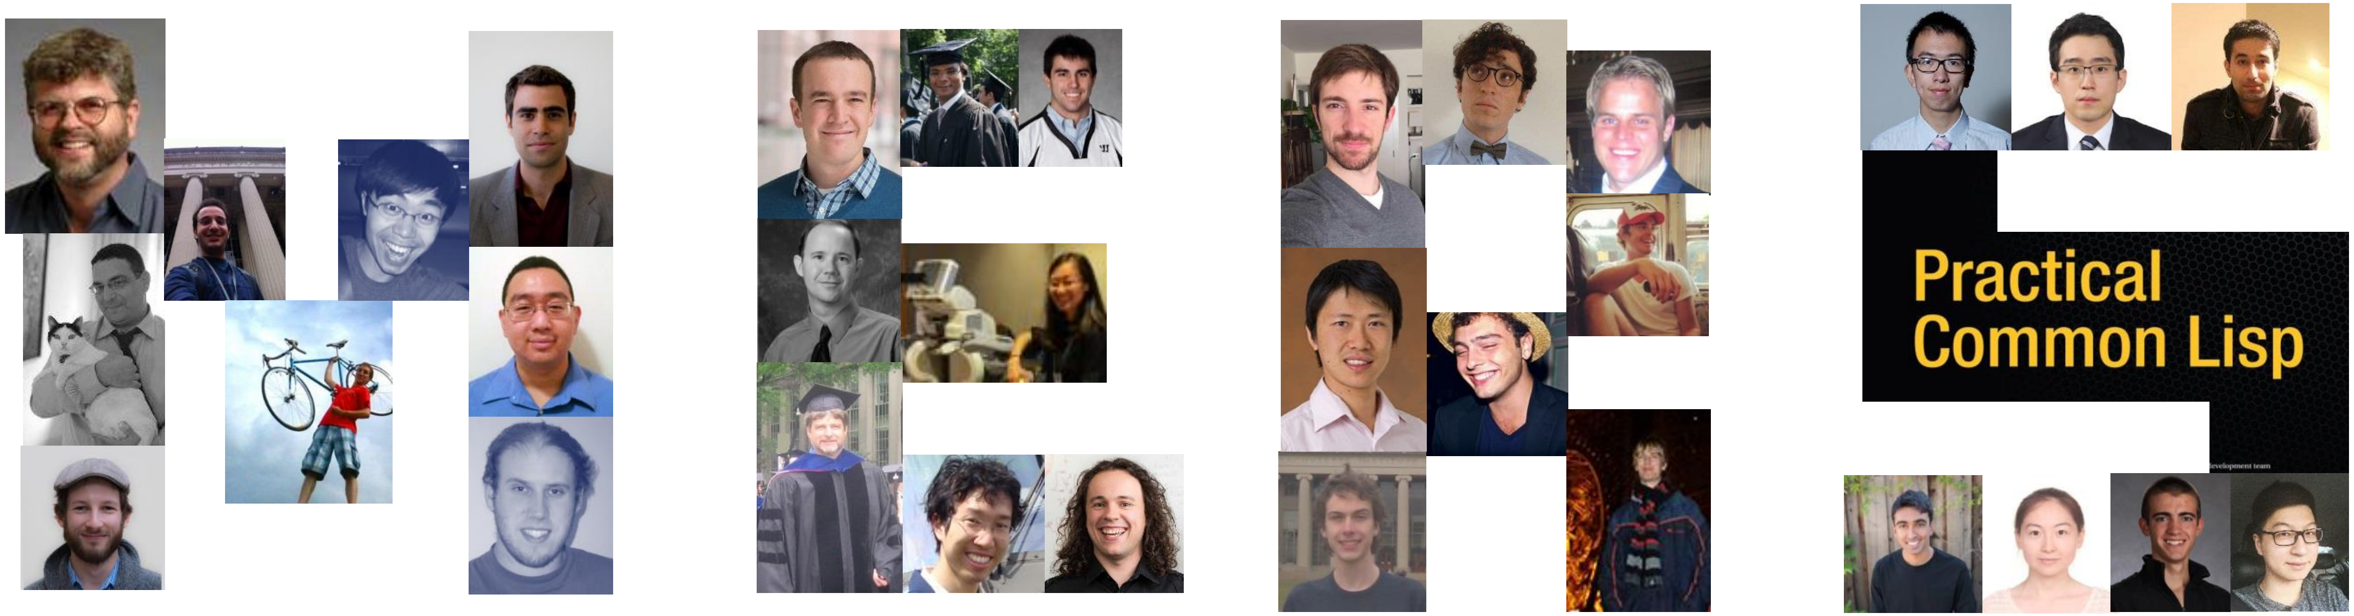

Acknowledgements

- Committee members and readers:
  - Professor Leslie Kaelbling, Professor Randall Davis, Dr. Jiaying Shen and Dr. Christian Muise.
- Collaborators at Boeing Research & Technology, Deep Submergence Lab at WHOI, Nuance NL/AI Lab and NICTA:
  - Scott Smith, Ron Provine, Dr. Rich Camilli, Dr. Peter Z. Yeh, Professor Patrik Haslum and Jing Cui.
- My parents Zhihe Yu and Jie Cui, and my wife, Jennifer.

Contributions

Resolving over-subscribed plans using a **variety** of **efficient** partial relaxation techniques leads to greater **flexibility** in plan adaptation.




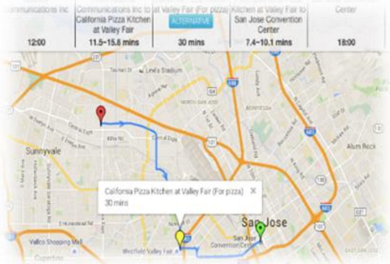

Publications

Publications

- [Yu and Williams, 2013] Peng Yu and Brian Williams, Continuously Relaxing Over-constrained Conditional Temporal Problems through Generalized Conflict Learning and Resolution, *Proceedings of the Twenty- third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013.

- [Yu, Fang and Williams, 2013] Peng Yu and Cheng Fang and Brian Williams, Resolving Uncontrollable Conditional Temporal Problems using Continuous Relaxations, *Proceedings of the Twenty-fourth International Conference on Automated Planning and Scheduling (ICAPS-14)*, Portsmouth, 2014 (Honorable Mention for Best Paper Award).

- [Karpas, Levine, Yu and Williams, 2015] Erez Karpas and Steve Levine and Peng Yu and Brian Williams, Robust Execution of Plans for Human-Robot Teams, *Proceedings of the Twenty-fifth International Conference on Automated Planning and Scheduling (ICAPS-2015),* Jerusalem, 2015.

Publications

- [Fang, Yu and Williams, 2014] Cheng Fang and Peng Yu and Brian Williams, Chance-constrained Probabilistic Simple Temporal Problems, *In Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, Quebec City, 2014.

- [Yu, Fang and Williams, 2015] Peng Yu and Cheng Fang and Brian Williams, Resolving Over-constrained Probabilistic Temporal Problems through Chance Constraint Relaxation, *In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, Austin, 2015.

- [Cui, Yu, Fang, Haslum and Williams, 2015] Jing Cui and Peng Yu and Cheng Fang and Patrik Haslum and Brian Williams, Optimising Bounds in Simple Temporal Networks with Uncertainty under Dynamic Controllability Constraints, *Proceedings of the Twenty-fifth International Conference on Automated Planning and Scheduling (ICAPS-2015),* Jerusalem, 2015.

Publications

- [Yu, Shen, Yeh and Williams, 2016a] Peng Yu and Jiaying Shen and Peter Yeh and Brian Williams, Towards Personal Assistants that can Help Users Plan, In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-2016),* New York, 2016.

- [Yu, Shen, Yeh and Williams, 2016b] Peng Yu and Jiaying Shen and Peter Yeh and Brian Williams, In *Proceedings of the Sixteenth International Conference on Intelligent Virtual Agents (IVA-2016)*, Los Angeles, 2016.

References

References

- [de Kleer and Williams, 1989] Johan De Kleer and Brian Williams, Diagnosing Multiple Faults, *Artificial intelligence* 32.1 (1987): 97-130.
- [Williams and Ragno, 2003] Brian Williams and Robert Ragno, Conflict-directed A* and Its Role in Model-based Embedded Systems, *Journal of Discrete Applied Mathematics*, 155(12):1562–1595, 2002.
- [Previti and Marques-Silva, 2013] Previti, A., & Marques-Silva, J, Partial MUS enumeration, In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI-2013)*, 818–825.
- [Peintner, Moffitt and Pollack, 2005] Peintner, B., Moffitt, M. D., and Pollack, M. E., Augmenting disjunctive temporal problems with finite-domain constraints. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence (AAAI-2005)*.
- [Bailey and Stucky, 2005] Bailey, J., & Stuckey, P.J, Discovery of minimal unsatisfiable subsets of constraints using hitting set dualization. In *Proceedings of the 7th International Symposium on Practical Aspects of Declarative Languages (PADL'05) of LNCS*, Vol. 3350, 174–186.

References

- [Yu and Williams, 2013] Peng Yu and Brian Williams, Continuously Relaxing Over-constrained Conditional Temporal Problems through Generalized Conflict Learning and Resolution, *Proceedings of the Twenty- third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013.

- [Vidal and Fargier, 1999] T. Vidal and H. Fargier. Handling contingency in temporal constraint networks: from consistency to controllabilities. *Journal of Experimental & Theoretical Artificial Intelligence*, 11:23–45, 1999.

- [Yu, Fang and Williams, 2013] Peng Yu and Cheng Fang and Brian Williams, Resolving Uncontrollable Conditional Temporal Problems using Continuous Relaxations, *Proceedings of the Twenty-fourth International Conference on Automated Planning and Scheduling (ICAPS-14)*, Portsmouth, 2014 (Honorable Mention for Best Paper Award).

- [Karpas, Levine, Yu and Williams, 2015] Erez Karpas and Steve Levine and Peng Yu and Brian Williams, Robust Execution of Plans for Human-Robot Teams, *Proceedings of the Twenty-fifth International Conference on Automated Planning and Scheduling (ICAPS-2015),* Jerusalem, 2015.

References

- [Morris and Muscettola, 2005] Morris, P. H., and Muscettola, N., Temporal dynamic controllability revisited. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-2005),* 1193–1198.

- [Morris, 2006] Morris, P., A structural characterization of temporal dynamic controllability. In *Proceedings of the 12th International Conference on Principles and Practice of Constraint Programming (CP-2006)*, 375–389.

- [Fang, Yu and Williams, 2014] Cheng Fang and Peng Yu and Brian Williams, Chance-constrained Probabilistic Simple Temporal Problems*, In Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, Quebec City, 2014.

- [Wang and Williams, 2015] Andrew Wang and Brian Williams, Chance-constrained Scheduling via Conflict-directed Risk Allocation*, In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, Austin, 2015.

References

- [Yu, Fang and Williams, 2015] Peng Yu and Cheng Fang and Brian Williams, Resolving Over-constrained Probabilistic Temporal Problems through Chance Constraint Relaxation*, In Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, Austin, 2015.

- [Newell and Potts, 1964] G.F. Newell and R.B. Potts, “Maintaining a Bus Schedule”, In *Proceedings of the Second Conference of the Australian Road Research Board 2*, 388–393 (1964).

- [Wah and Xin, 2006] Wah, B. W., & Xin, D., Optimization of bounds in temporal flexible planning with dynamic controllability. *IEEE International Conference on Tools with Artificial Intelligence*, 40–48.

- [Crawford, 1996] J. Crawford, An Approach to Resource Constrained Project Scheduling. In *Proceedings of the 1996 Artificial Intelligence and Manufacturing Research Planning Workshop*, 1996.

References

- [Banerjee and Haslum, 2011] Banerjee, D., and Haslum, P.. Partial-order supportlink scheduling. In *Proceedings of the 21st International Conference on Automated Planning and Scheduling (ICAPS),* 307–310.

- [Cui, Yu, Fang, Haslum and Williams, 2015] Jing Cui and Peng Yu and Cheng Fang and Patrik Haslum and Brian Williams, Optimising Bounds in Simple Temporal Networks with Uncertainty under Dynamic Controllability Constraints, *Proceedings of the Twenty-fifth International Conference on Automated Planning and Scheduling (ICAPS-2015),* Jerusalem, 2015.

- [Yu, Master Thesis, 2013] Peng Yu, Continuous Relaxation of Over-constrained Temporal Plans, *Master Thesis*, Massachusetts Institute of Technology, February 2013.

- [Thompson, Goker and Langley, 2004] Thompson, C. A., Göker, M. H., & Langley, P., A Personolized System for conversational Recommendations. *Journal of Artificial Intelligence Research*, 21, 1-36.

References

- [Raiman, 2015] Jonathan Raiman, Building Blocks for the Mind, Master Thesis, Massachusetts Institute of Technology, 2015.

- [Mikolov, Sutskever, Chen, Corrado and Dean, 2013] Mikolov, T., and J. Dean, Distributed representations of words and phrases and their compositionality, *Advances in neural information processing systems* (2013).

- [word2vec, Freebase Skipgram Vectors, 2013] https://code.google.com/archive/p/word2vec/.

- [Yu, Shen, Yeh and Williams, 2016a] Peng Yu and Jiaying Shen and Peter Yeh and Brian Williams, Towards Personal Assistants that can Help Users Plan, In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-2016),* New York, 2016.

References

- [Yu, Shen, Yeh and Williams, 2016b] Peng Yu and Jiaying Shen and Peter Yeh and Brian Williams, In *Proceedings of the Sixteenth International Conference on Intelligent Virtual Agents (IVA-2016)*, Los Angeles, 2016.

Additional Experiment Results

Resource-constrained Project Scheduling

- A variation of the Job-shop scheduling problems with resource constraints*.

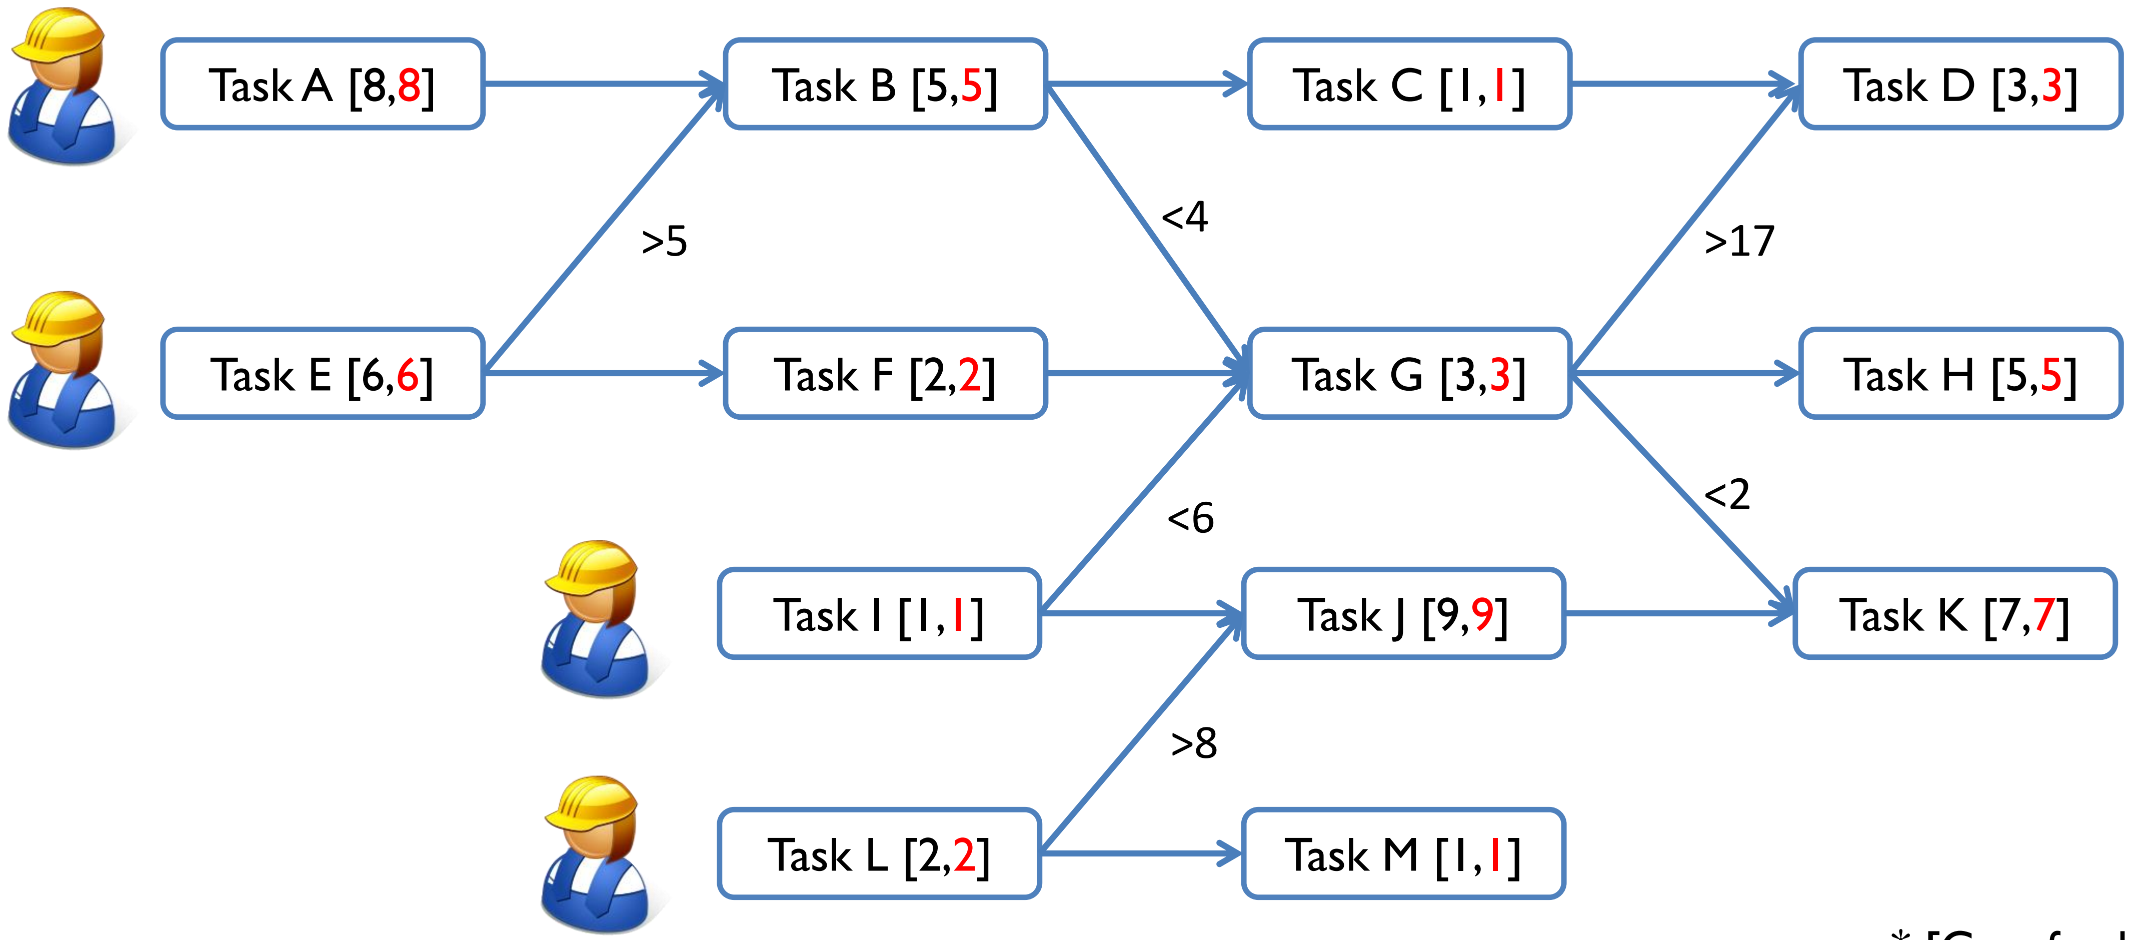

* [Crawford, 1996]

Resource-constrained Project Scheduling

- Given a partial order schedule*, computes the maximum amounts of uncertainty that can be built into it†.

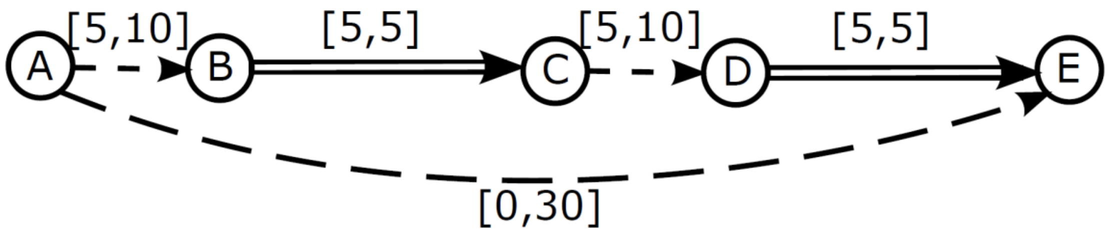

$$max \min_{ub_i'}(ub_i' - lb_i)$$

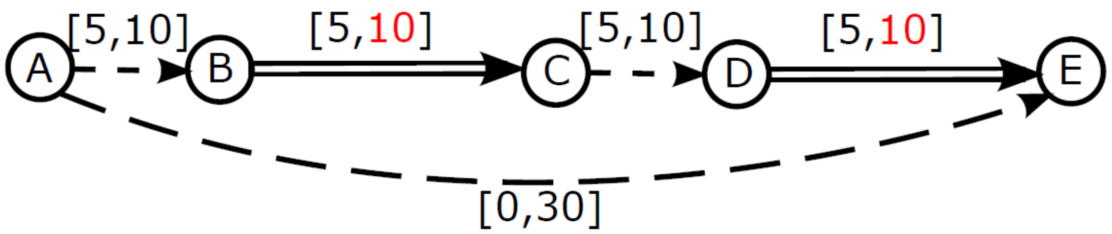

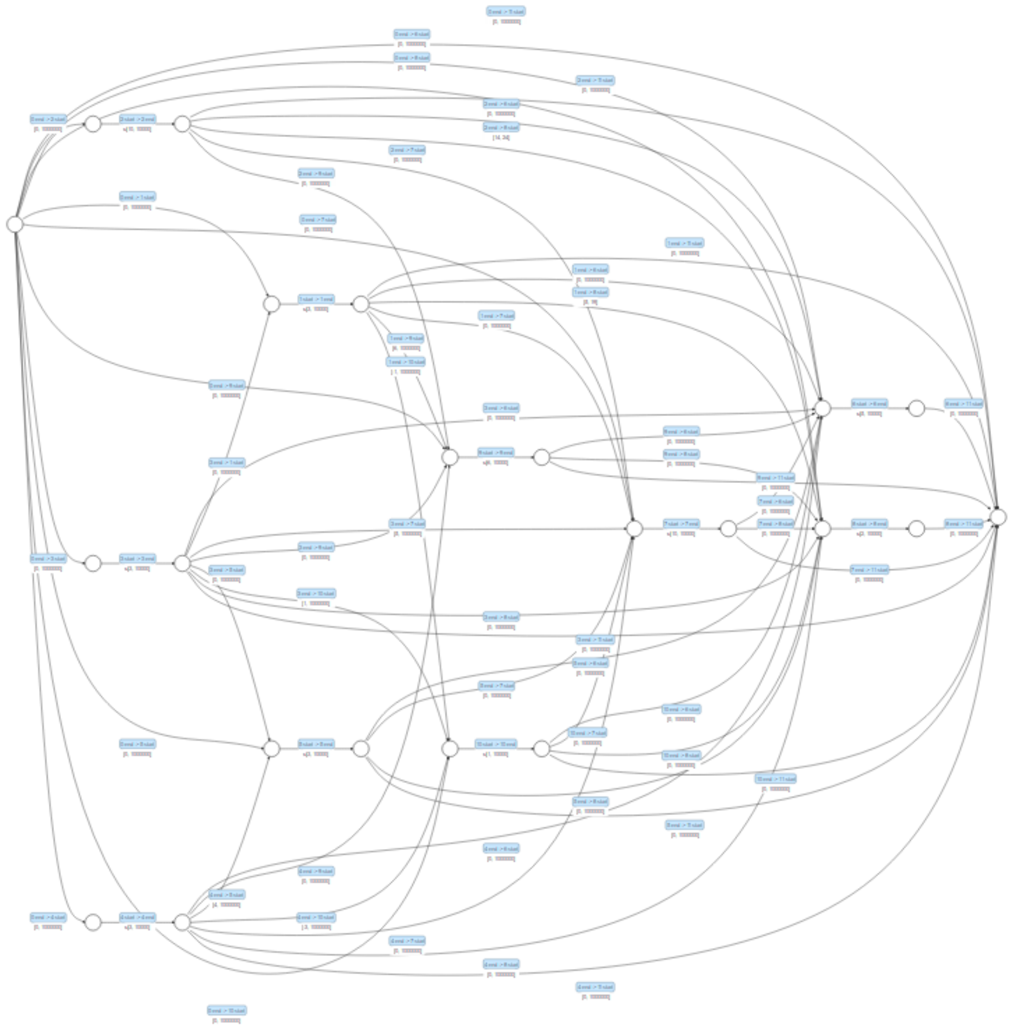

* [Banerjee and Haslum, 2011]

† [Cui, Yu, Fang, Haslum and Williams, 2015]

Empirical Results - Runtime

- Comparing the performance of BCDR-U(DC) and GuRoBi.
  - Solve all 326 instances in the J10 dataset. Each test repeated for 5 times.

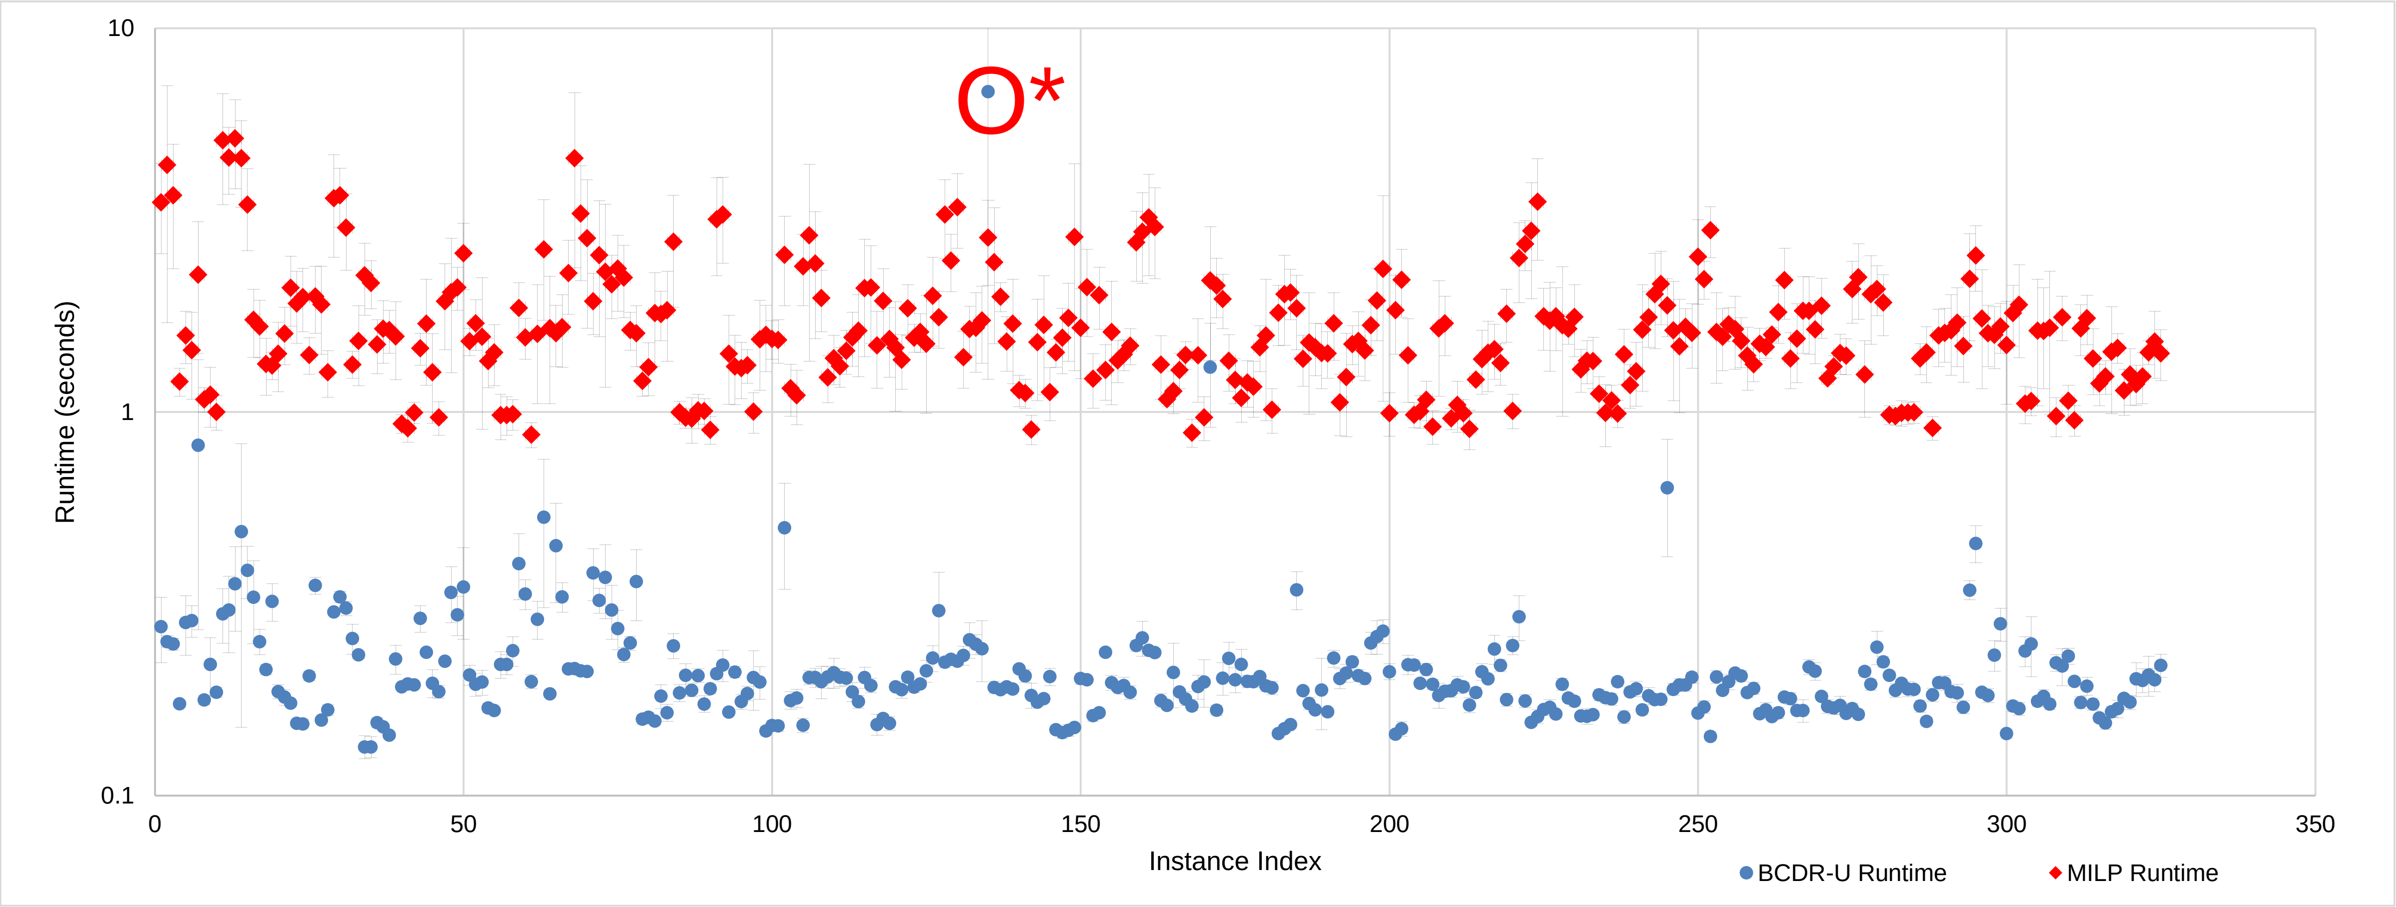

* [Yu, Williams, Fang, Cui and Haslum, submitted to JAIR]

Empirical Results

- We evaluate the number of instances solved in 30sec by:
  - BCDR, BCDR-U (SC), BCDR-U (DC), BCDR-C (SC), BCDR-C (DC).

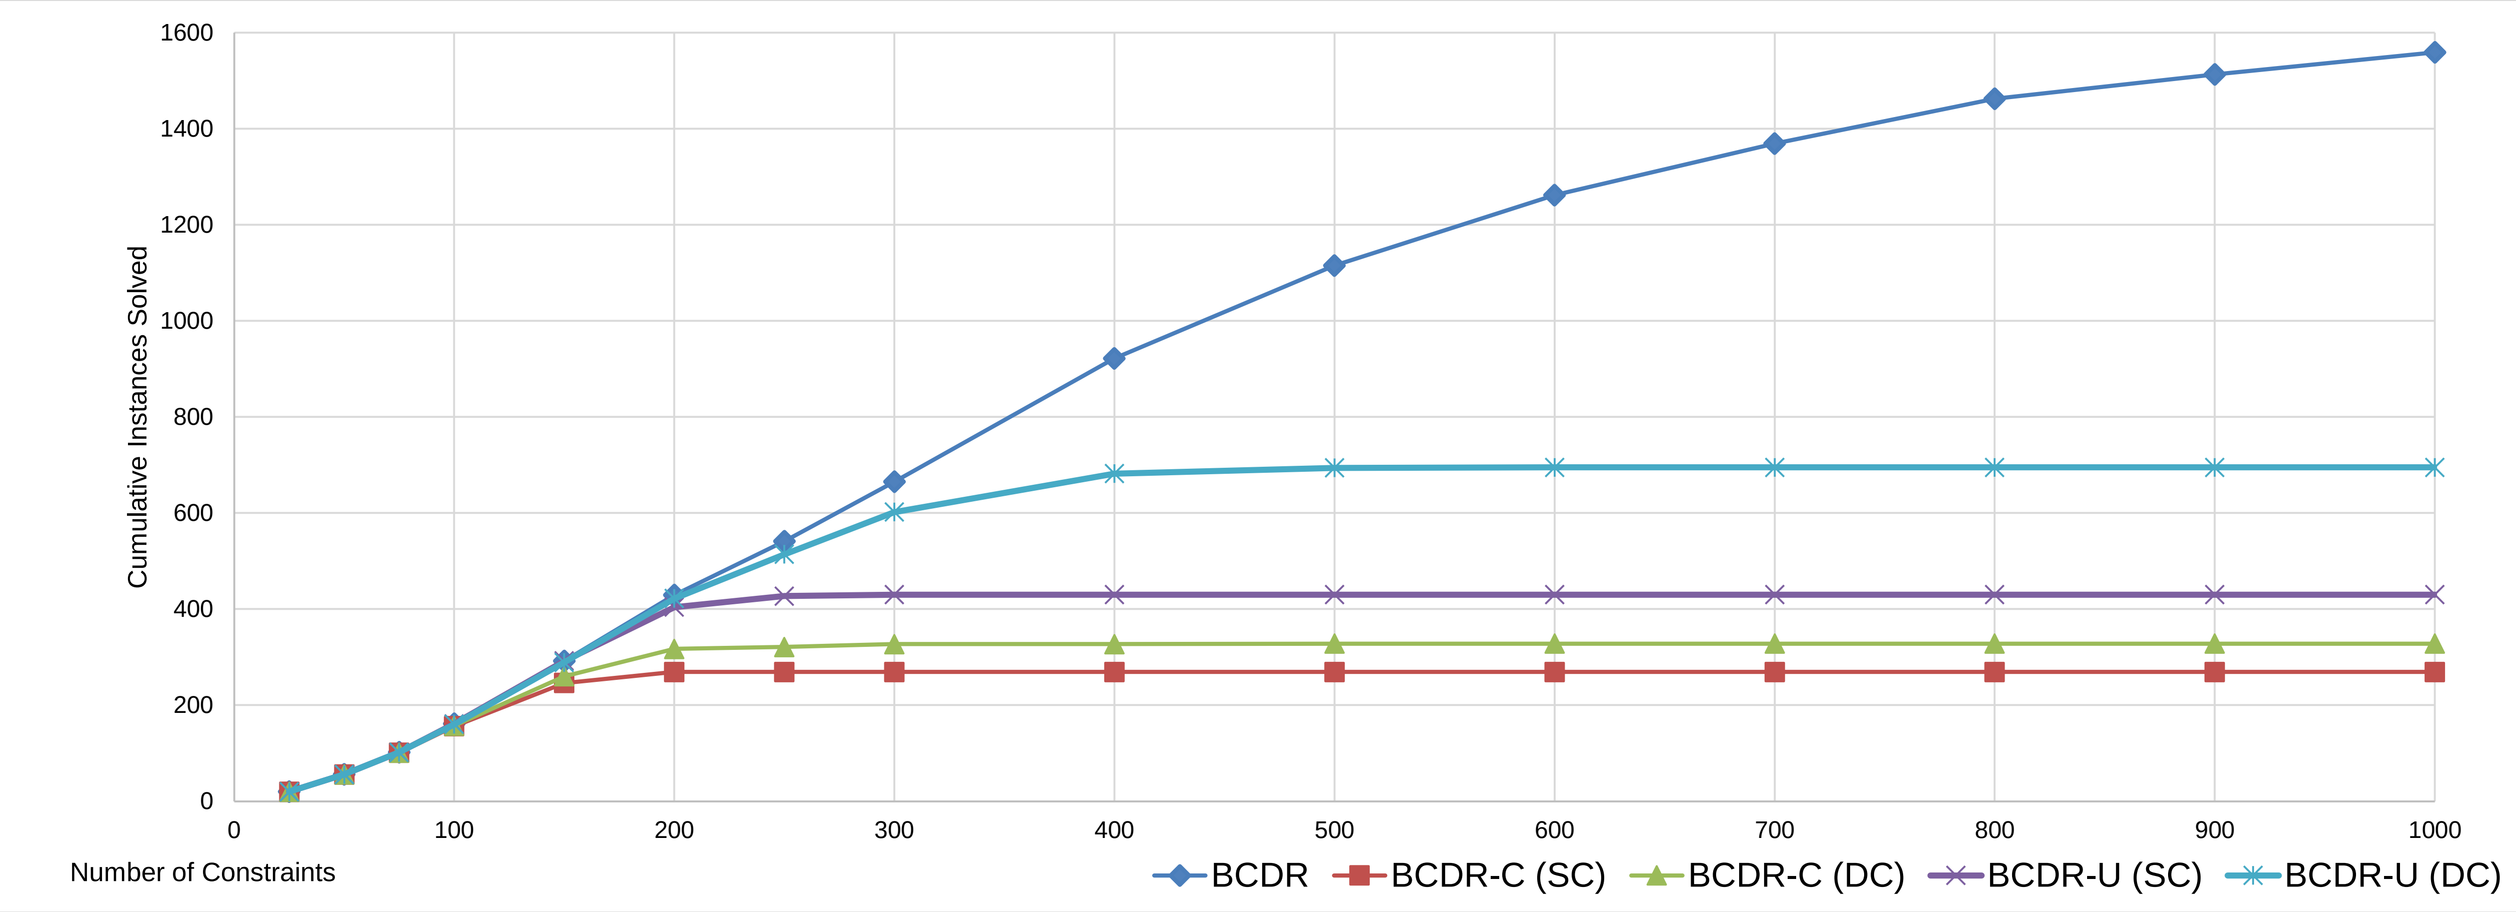

Empirical Results - Conflicts

- The numbers of conflicts resolved by the three algorithms.

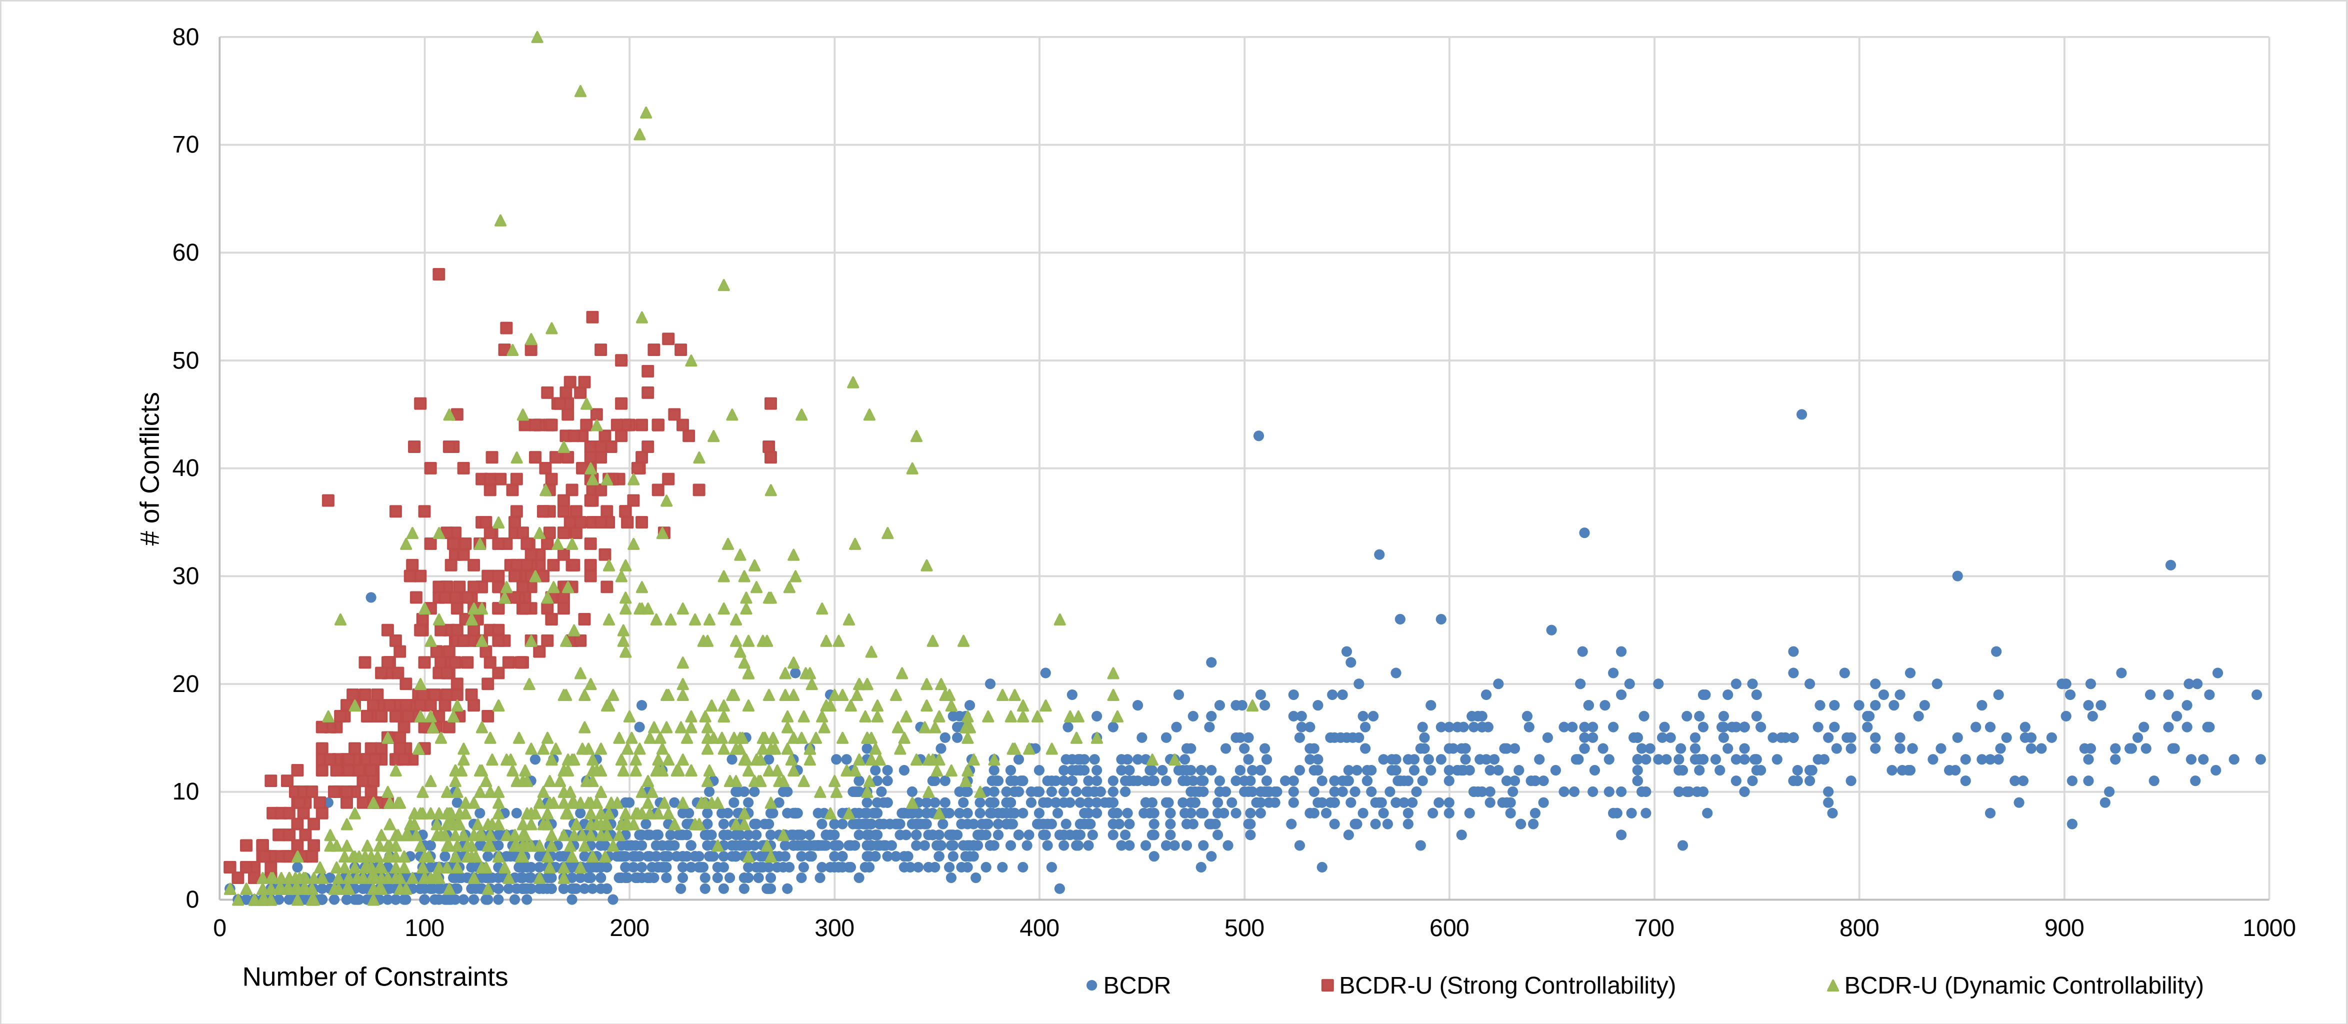

Empirical Results – BCDR vs DAA

- Runtime for computing **the preferred** vs **all** minimal relaxations (on AUV test cases).

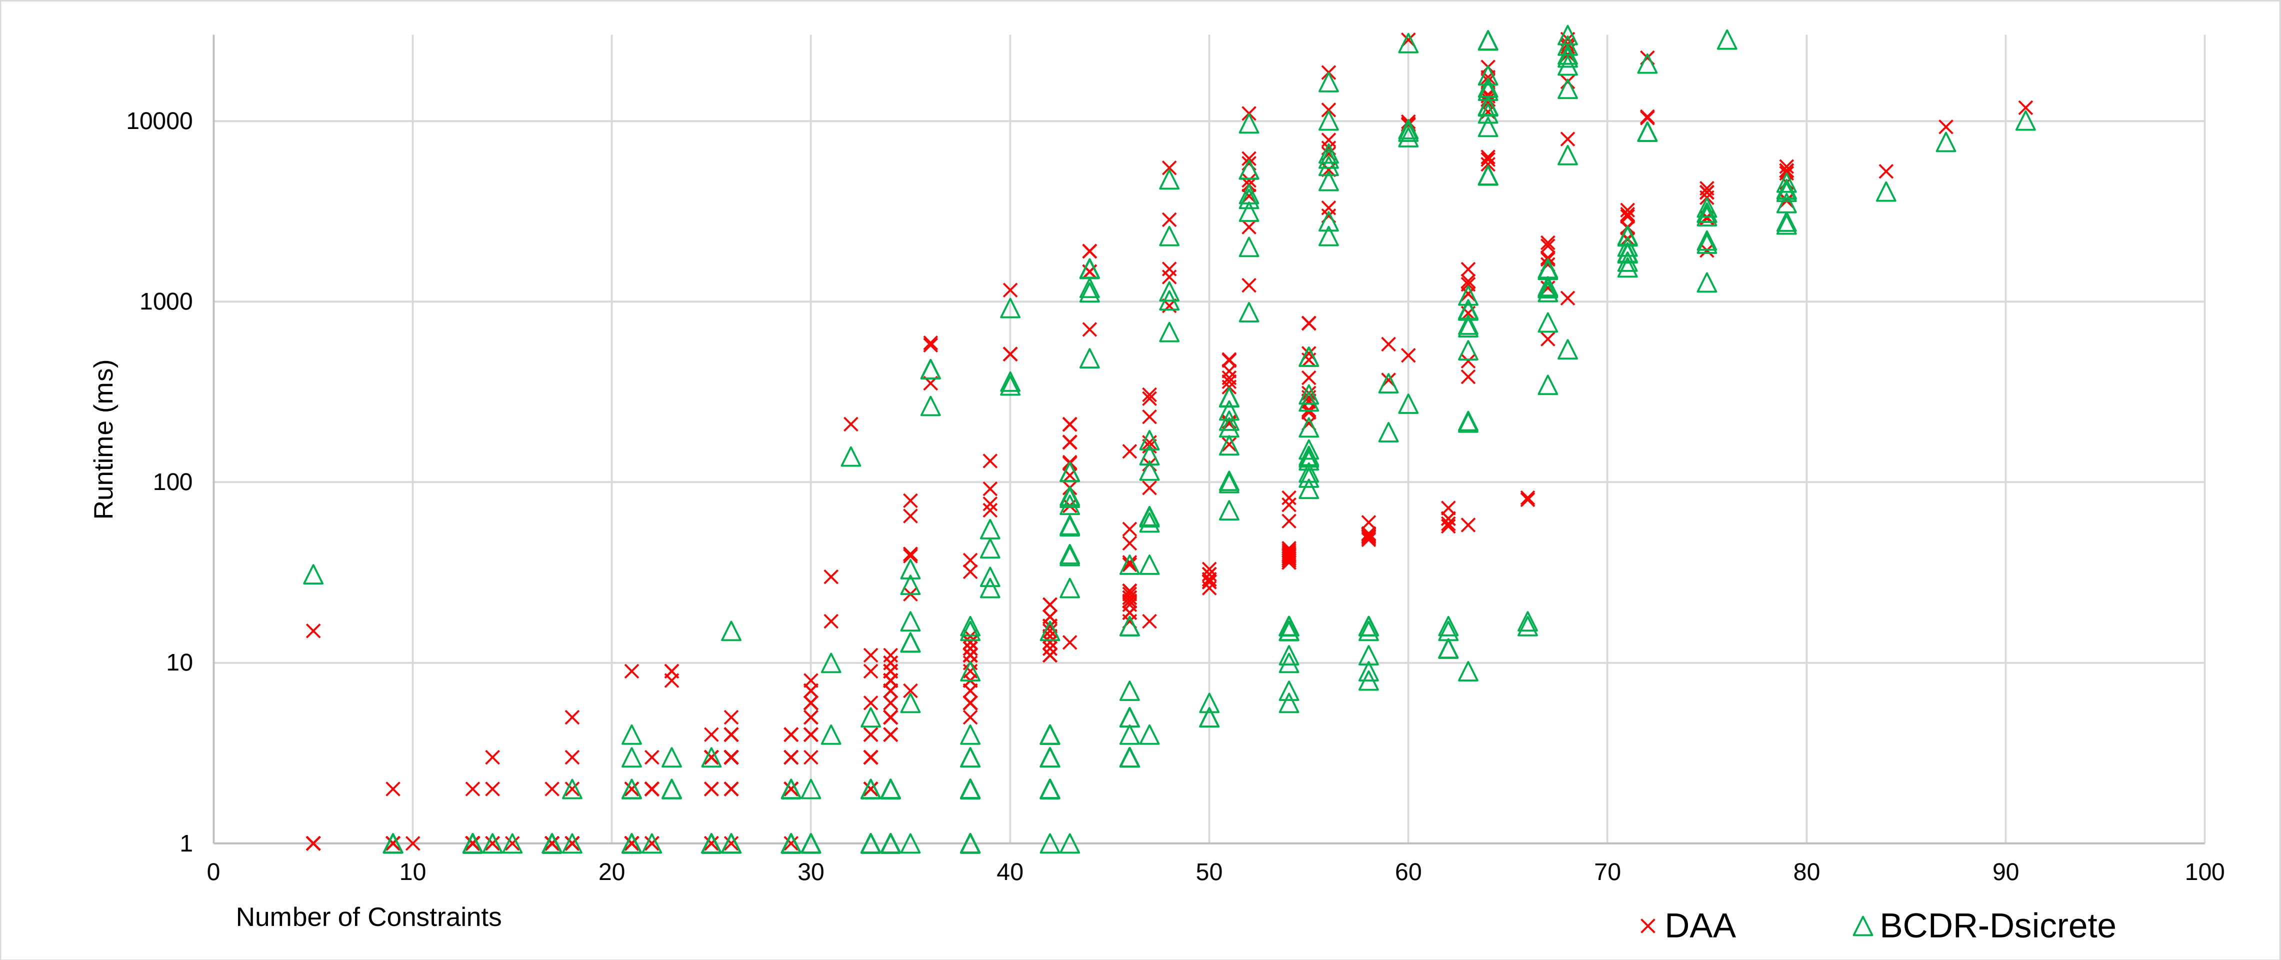

Empirical Results – Discrete vs Continuous

- Runtime for computing minimal discrete and preferred continuous relaxations (on AUV test cases).

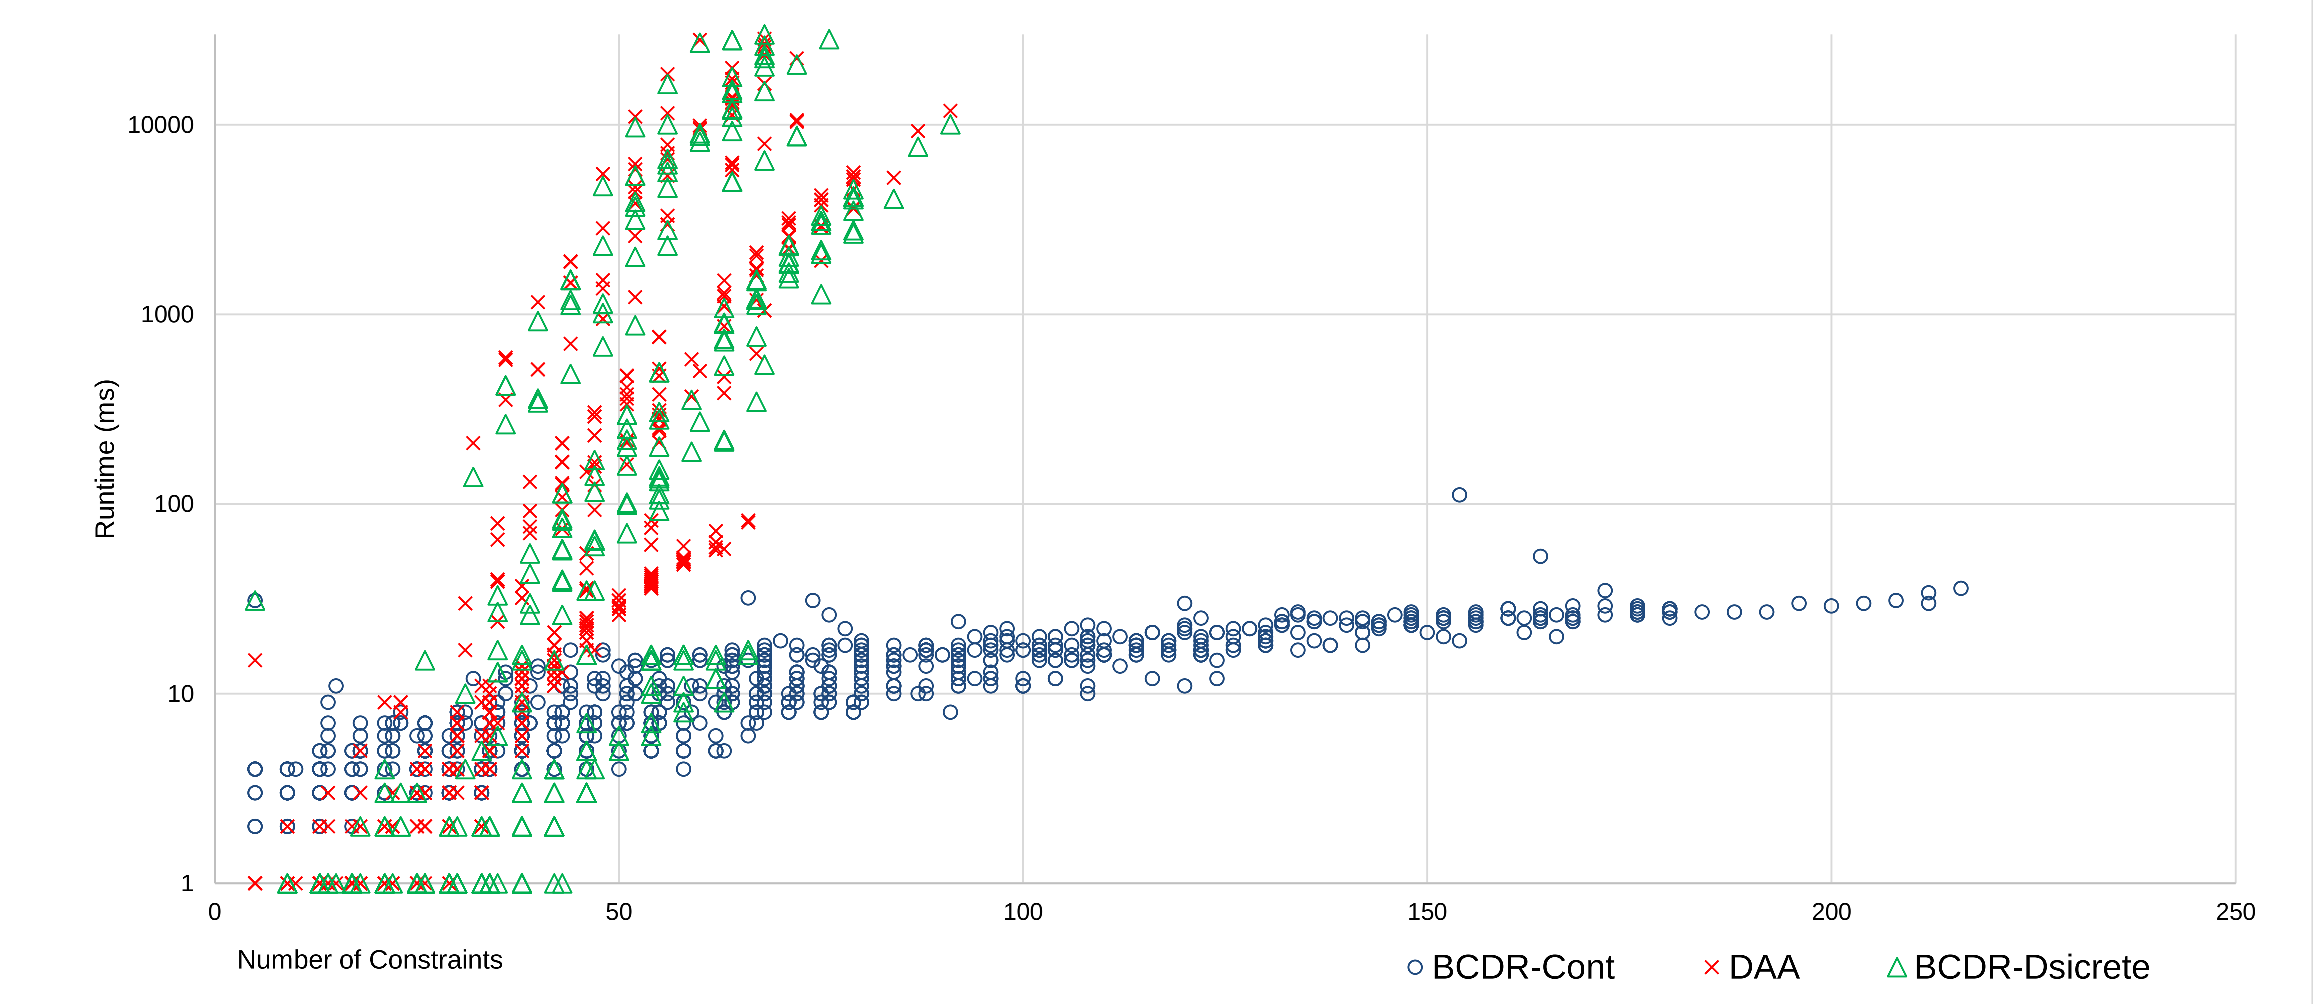

Empirical Results – Discrete vs Continuous

- Runtime for computing preferred continuous relaxations using BCDR and LPSolve (on AUV test cases).

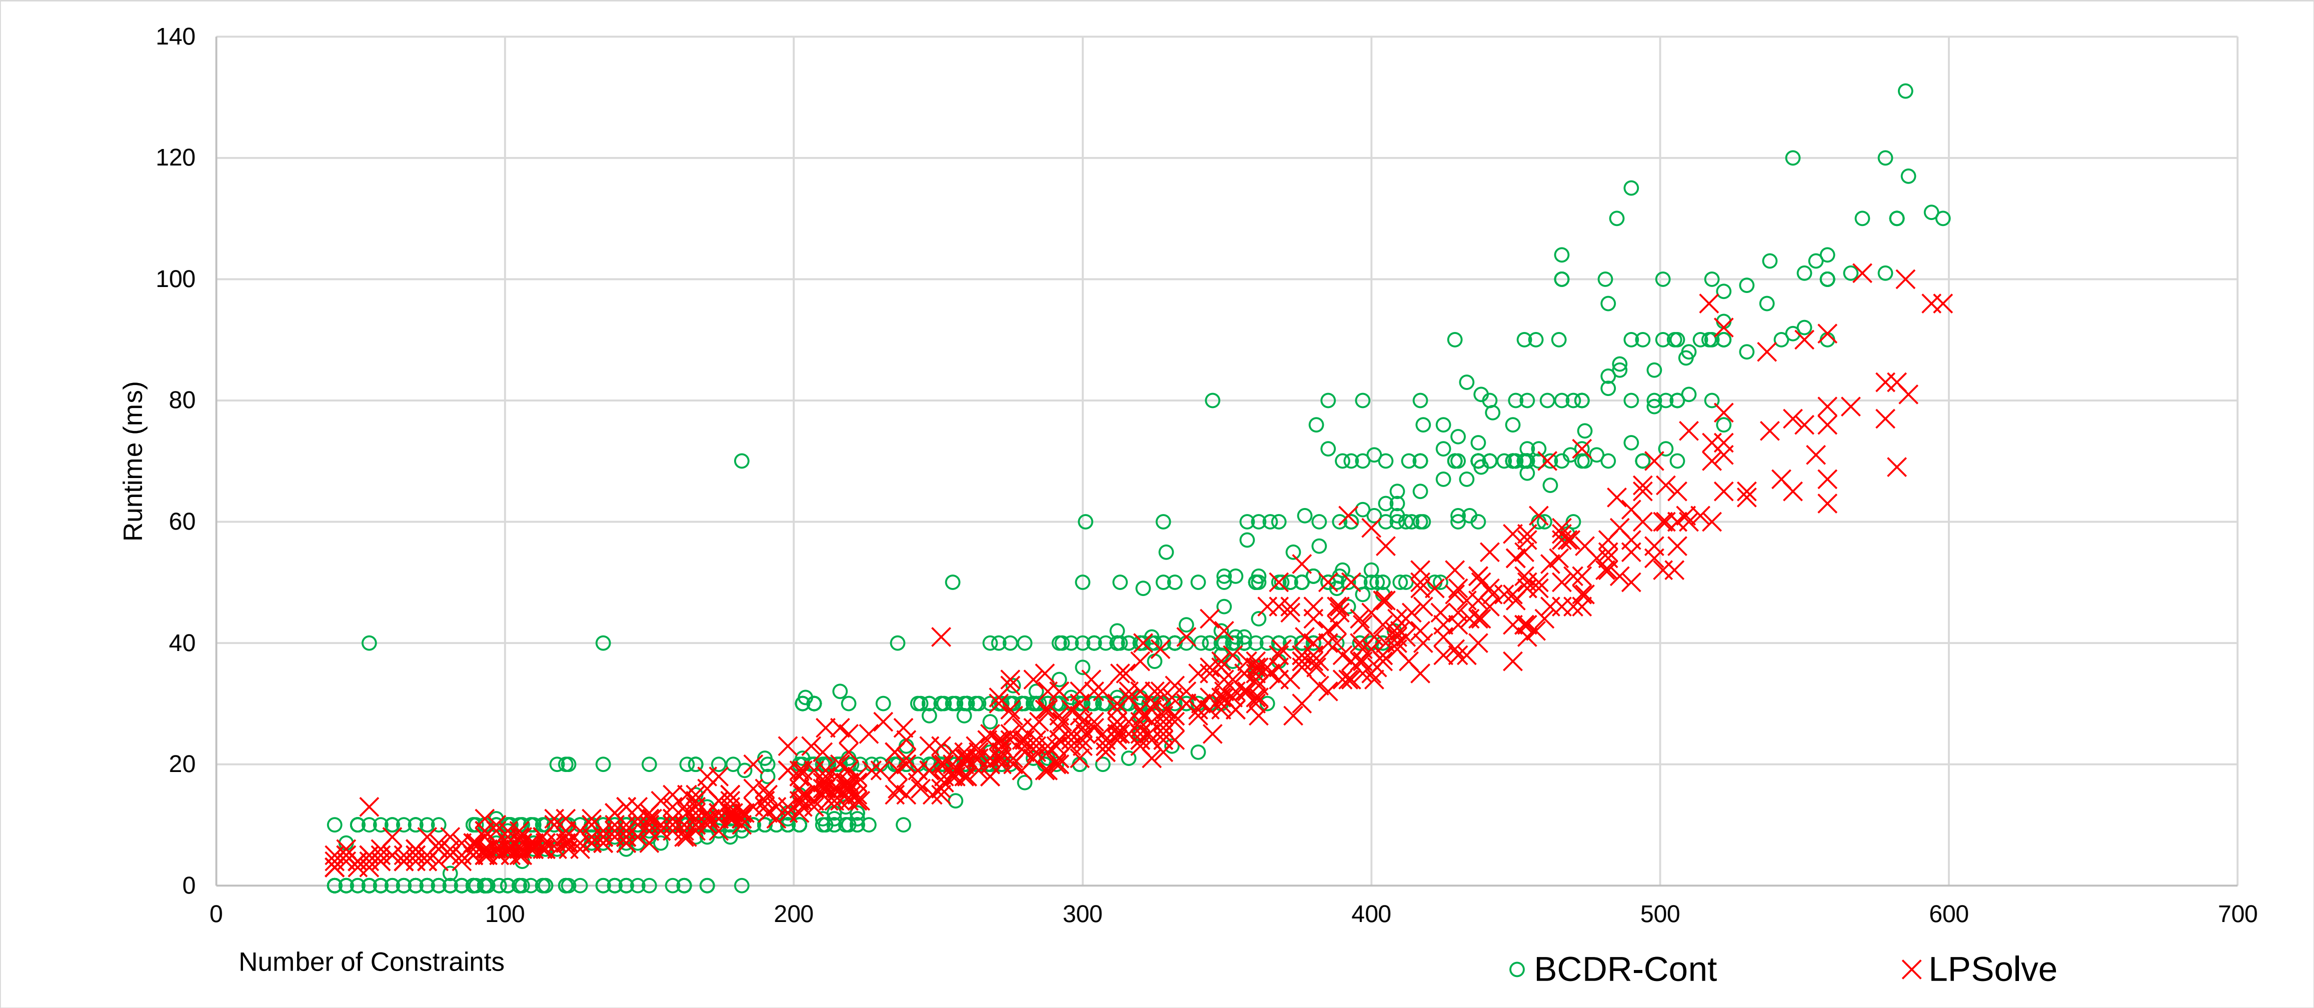

Example AUV Relay Problems

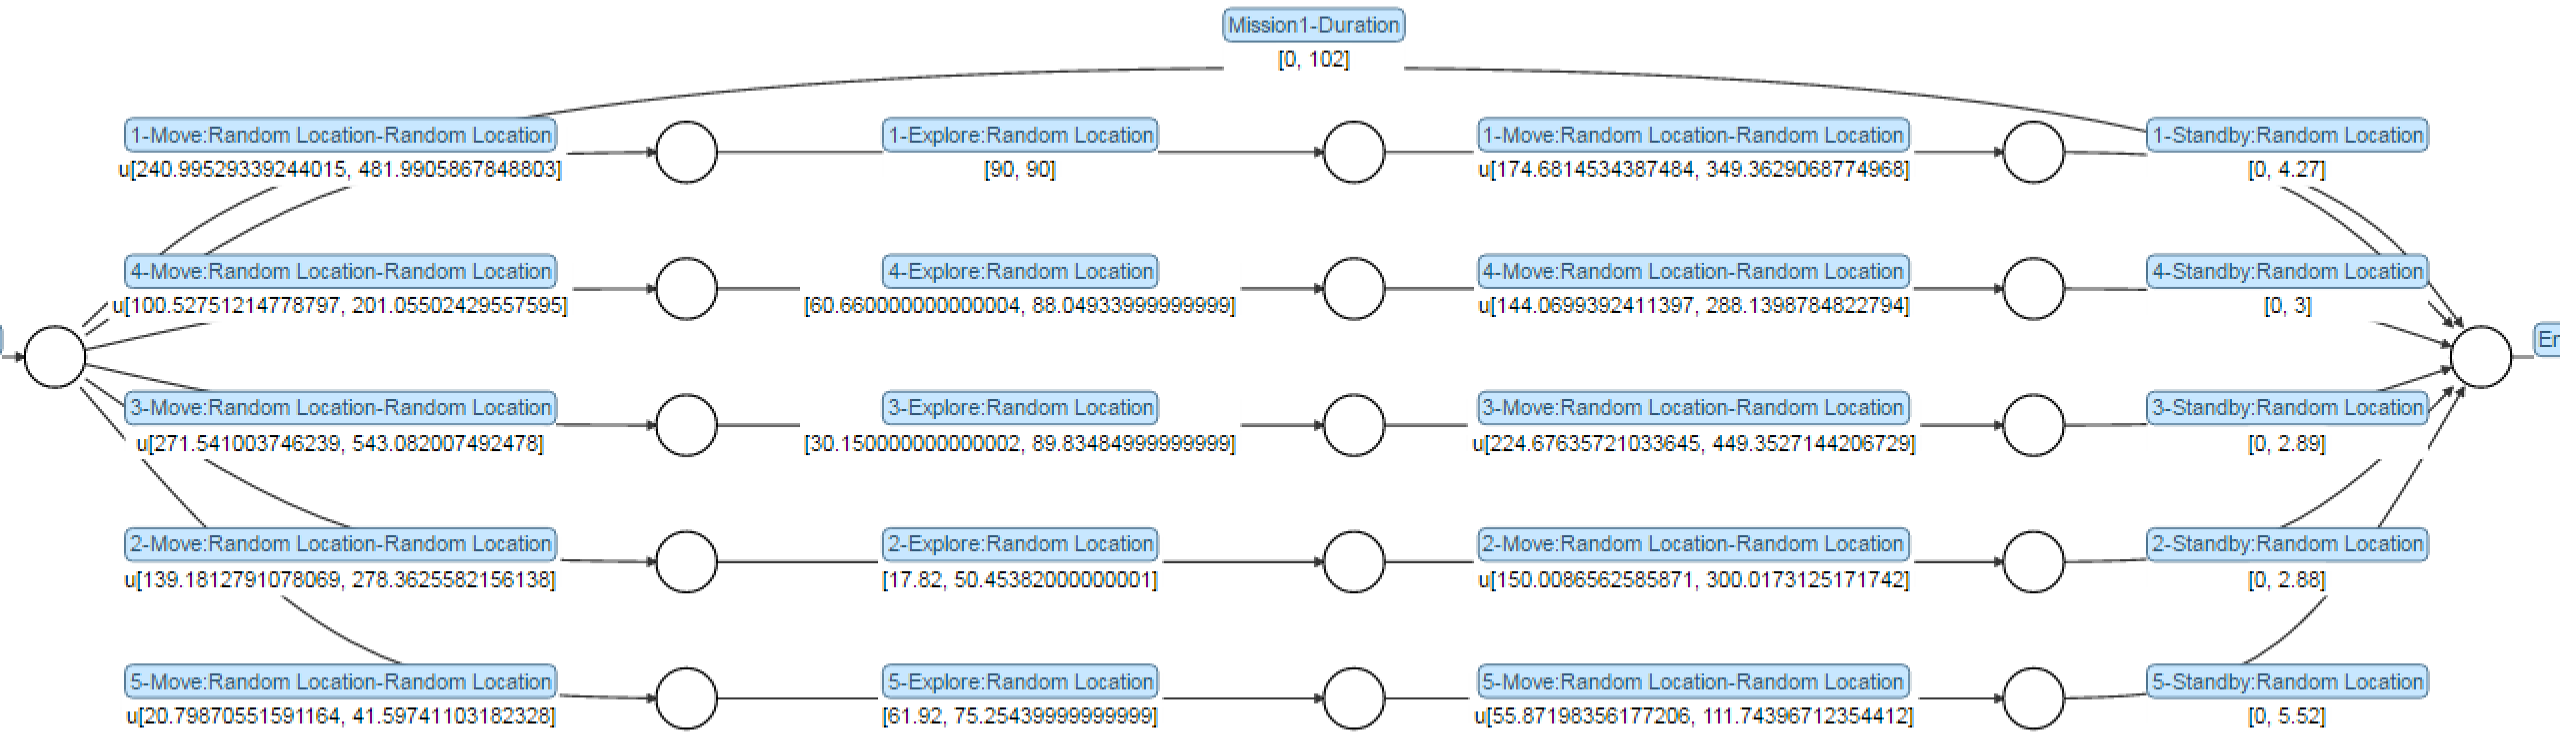

Dynamic Dispatching and Holding Strategies

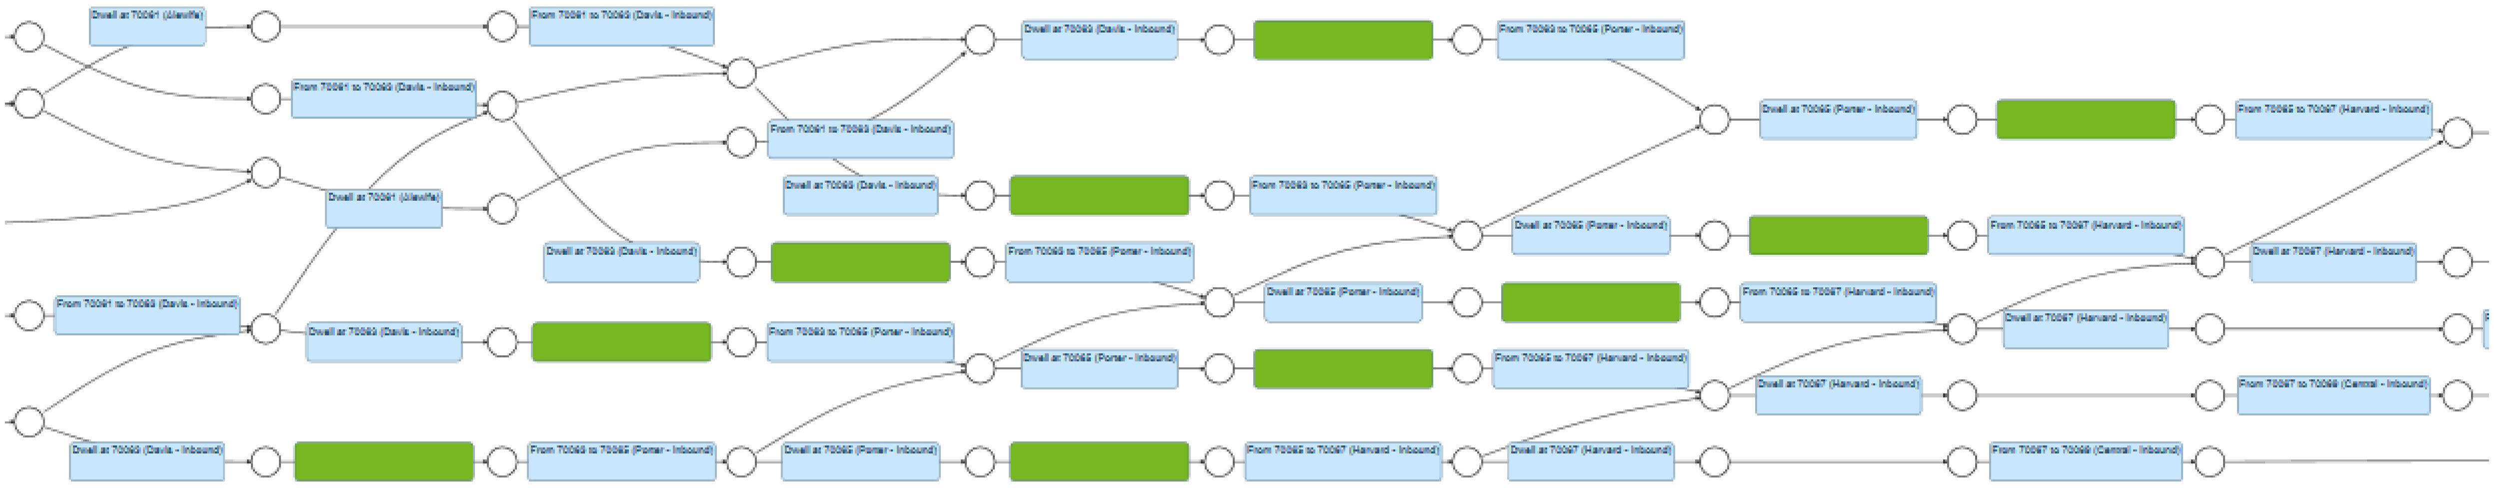

Results for Domain Relaxation

- Uhura found satisfying solutions for users in **52 out of 54** sessions.
- With temporal relaxation only approach, solutions found in **43** sessions.

Session	Quality Score	Relaxations (Temporal / Domain)
1	3.3 (1.4)	2.0 (2.6) / 2.1 (2.7)
2	2.4 (1.5)	1.3 (2.9) / 3.0 (3.3)
3	2.7 (1.5)	2.9 (3.0) / 3.1 (2.8)
4	3.7 (1.6)	0.3 (0.7) / 1.7 (3.4)
5	3.2 (1.4)	1.9 (2.6) / 1.7 (3.0)
6	3.3 (1.5)	0.6 (1.1) / 0.0 (0.0)

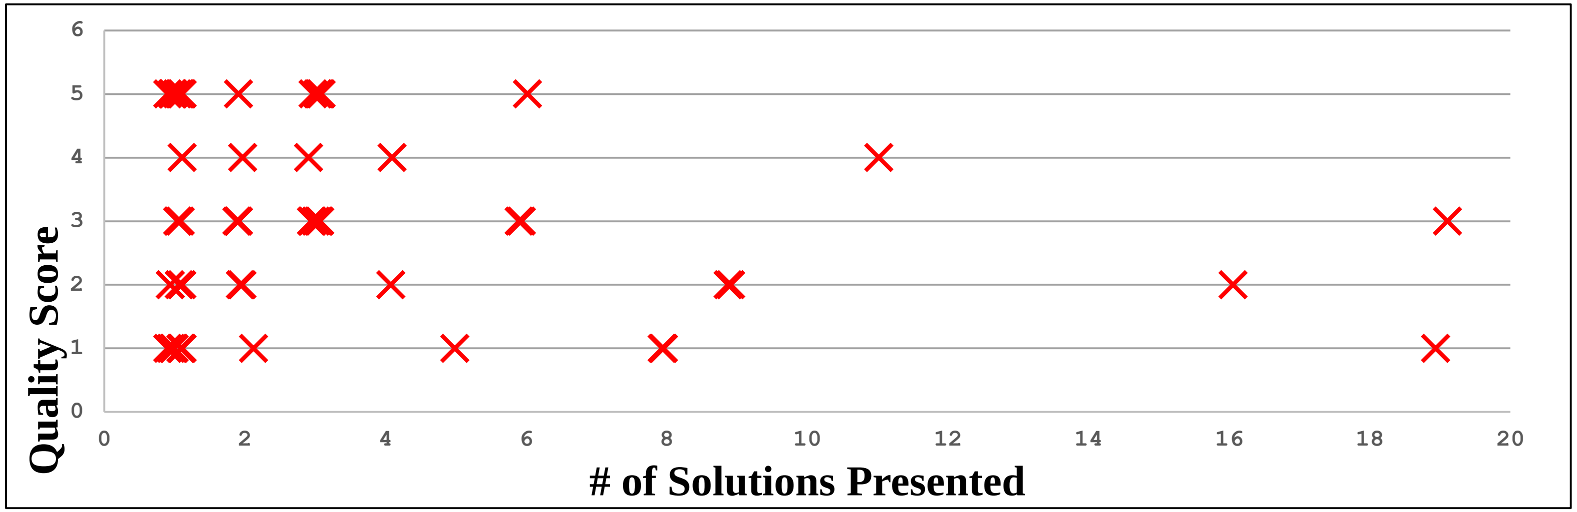

More Details on Conflict Learning and Resolution

Conflicts from Dynamical Controllability Checking

- **Record contributing** constraints during the iterative reduction procedure*.

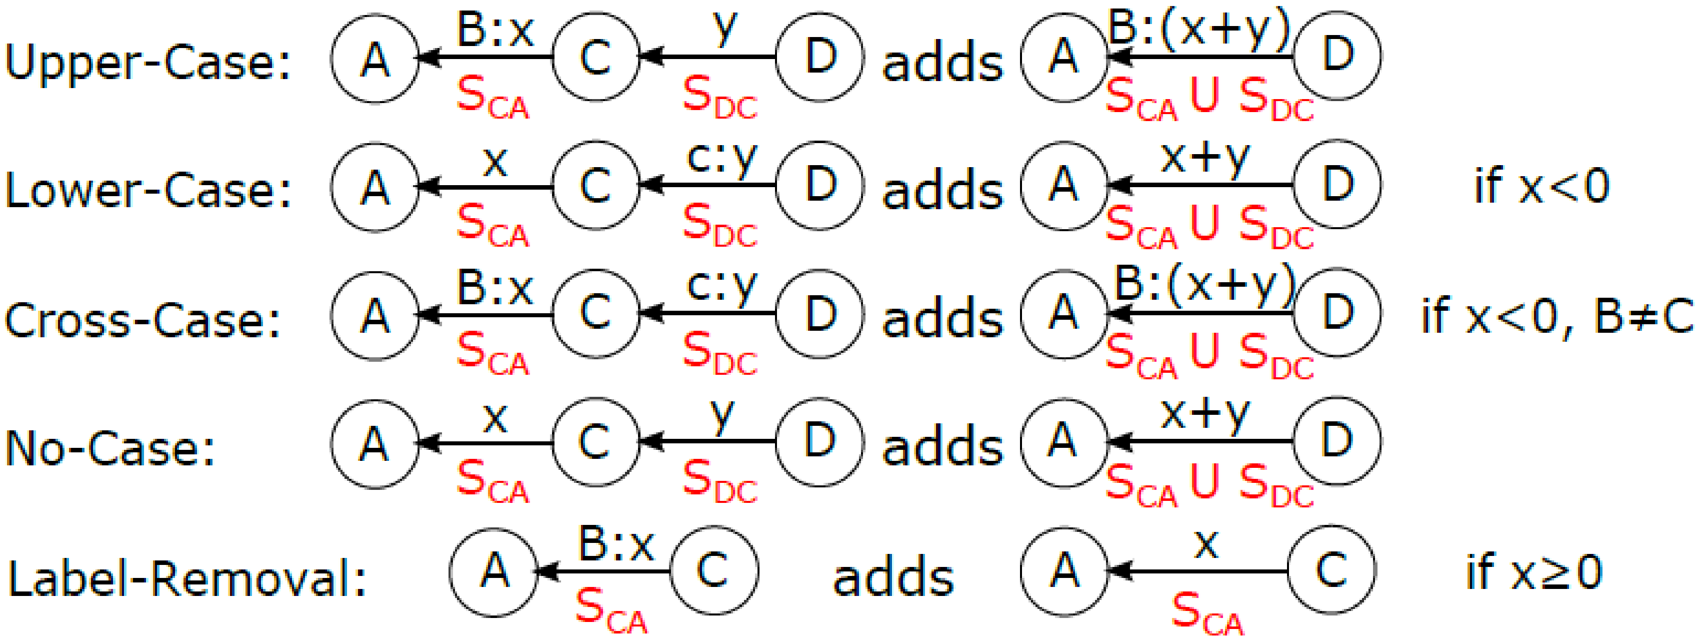

* [Morris and Muscettola, 2005]

- In addition to the negative cycle itself, we can also resolve a DC conflict by **disabling** reductions† that produce edges in the cycle*.

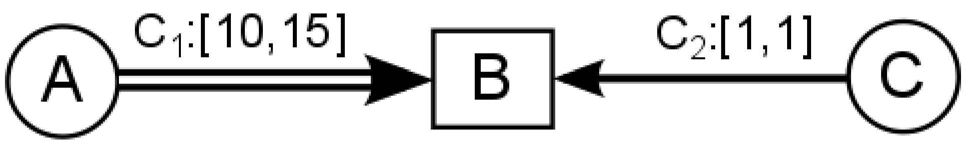

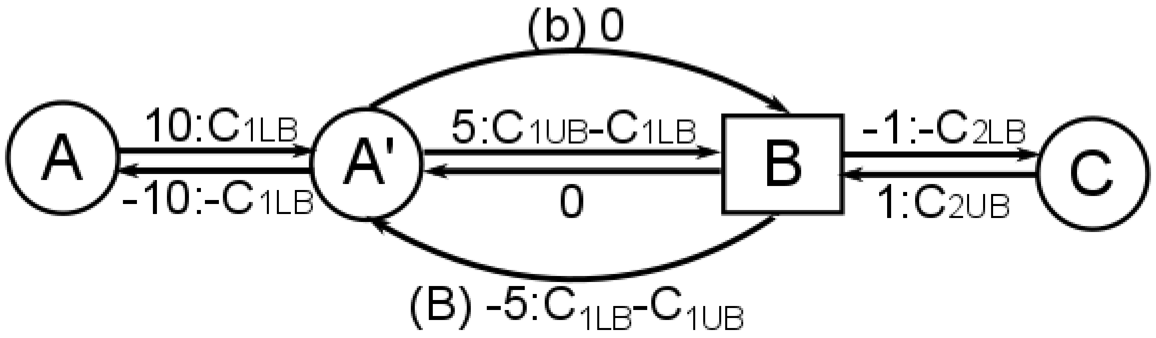

$$C_{2UB}+C_{1LB}-C_{1UB}-C_{2LB} \geq 0$$

or

$$-C_{2LB} \geq 0$$

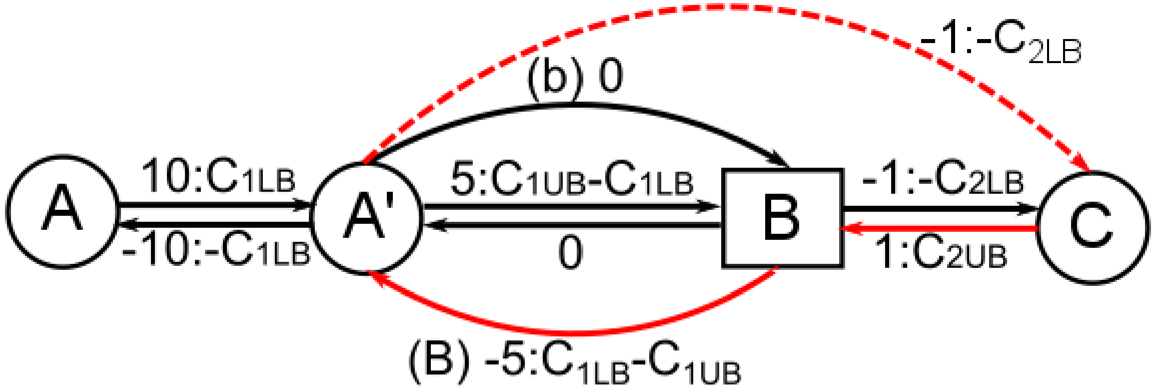

* [Morris, 2006]

† [Yu, Fang and Williams, 2014]

Retrieving Candidates for Domain Relaxations

- We query the knowledge base for additional candidates using the original & weakened domain constraints.

KOREAN RESTAURANT: $\forall$r
(restaurant r) $\wedge$
(cuisine c) $\wedge$
(servesCuisine r c) $\wedge$
(surface c '**KOREAN**')

Relaxing Cuisine Constraint

CHINESE RESTAURANT: $\forall$r
(surface c '**CHINESE**')

KNOWLEDGE BASE

{Koreana,
Kimchi Kitchen,
… …}

{Mary Chung,
Royal East,
… …}

Relaxing Multiple Domain Constraints

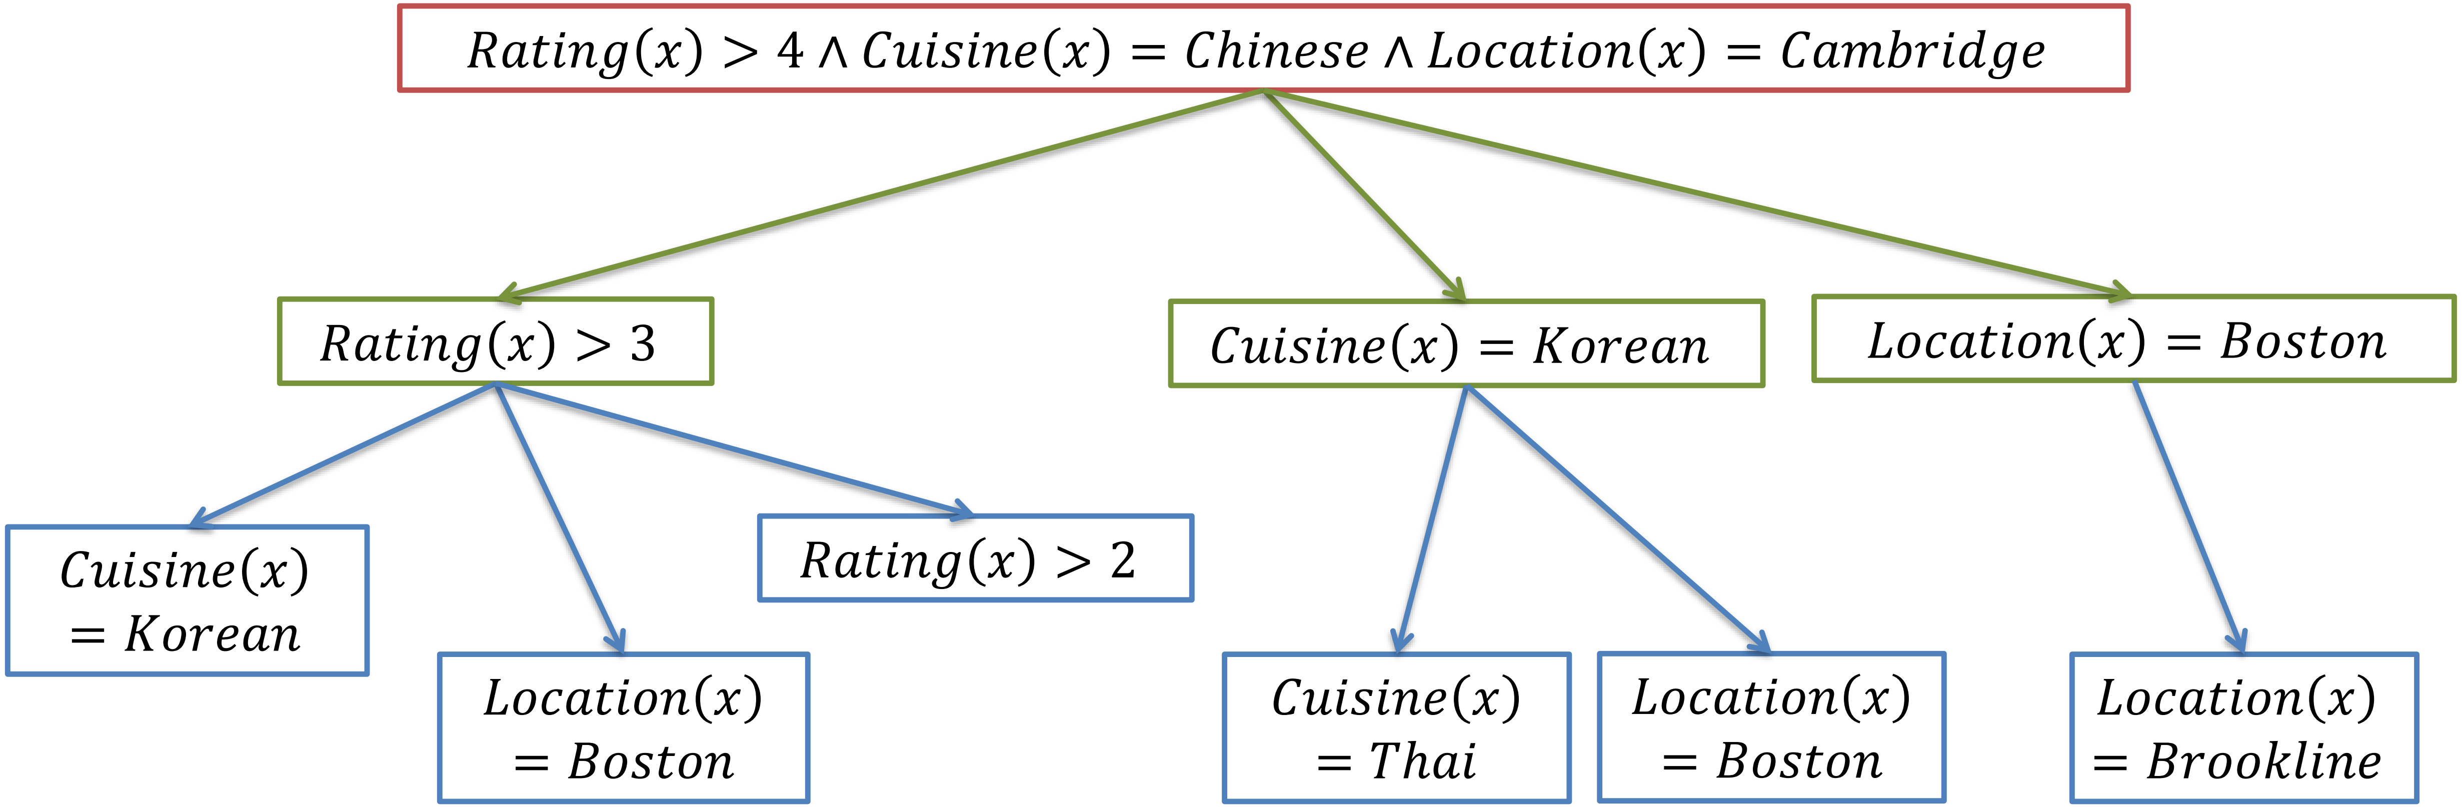

Computing Semantic Distance

- Using a 1000D word vector model of Freebase concepts, trained on news articles, to support the distance calculation between semantic constraints.

Chinese food
/m/01xw9

Cosine Distance

Concept	Distance
Korean food /m/048vr	0.7140
Thai food /m/07hxn	0.6945
Mexican food /m/051zk	0.5169
American food /m/01z1jf2	0.3183

Modeling Uncertain Durations

- An uncertain duration $e_i \in E_u$ can be described using a random variable ω, which encodes the possible outcomes for the duration.
  - ω may be described using a set bound [LB,UB].

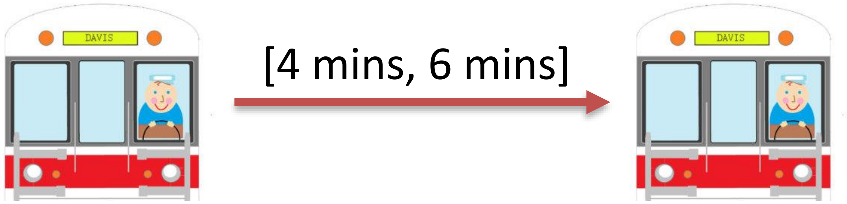

  - or a probabilistic distribution, which provides more details about the likelihood of different outcomes.

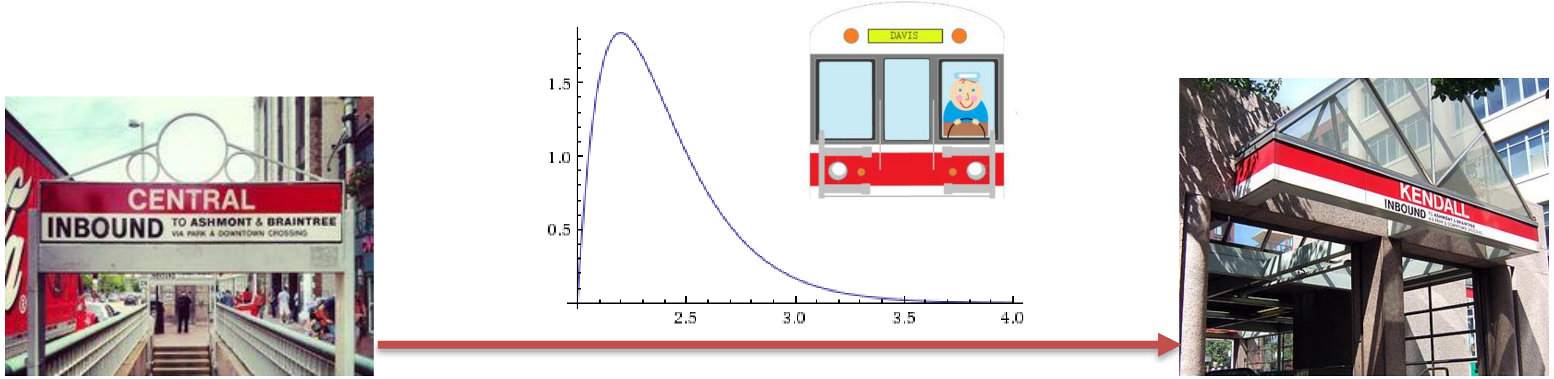

Reward and Cost Functions

- Each variable assignment is mapped to a positive reward by function f_p.
- Each constraint relaxation is mapped to a positive cost by function f_e.

Grocery	Shaws	40
	Trader Joes	100
Dinner	Panda Express	60
	Hong Kong Eatery	80

Assignment reward:

$$\alpha = \{Grocery = \text{Trader Joes}, Lunch = \text{Panda Express}\}$$

$$f_p(\alpha) = 100 + 60 = 160$$

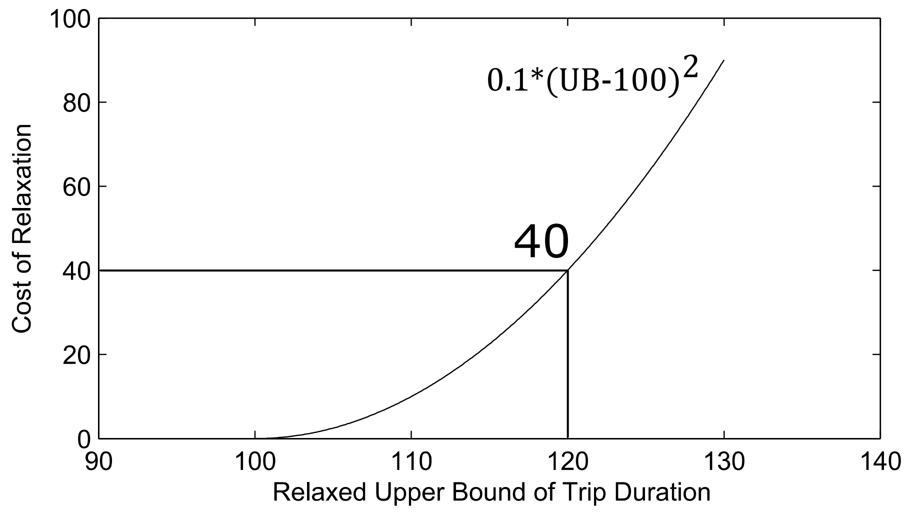

Relaxation cost:

$$tr = TripDuration: [0,100] \rightarrow [0,120]$$

$$f_e(tr) = f_e(120) = 40$$

Details about Uhura’s Architecture

Functional Architecture

- Capture users' goals and requirements.
- Present consistent plans and negotiate different relaxations.

- Compare and enumerate alternative choices.
- Prioritize different relaxations.

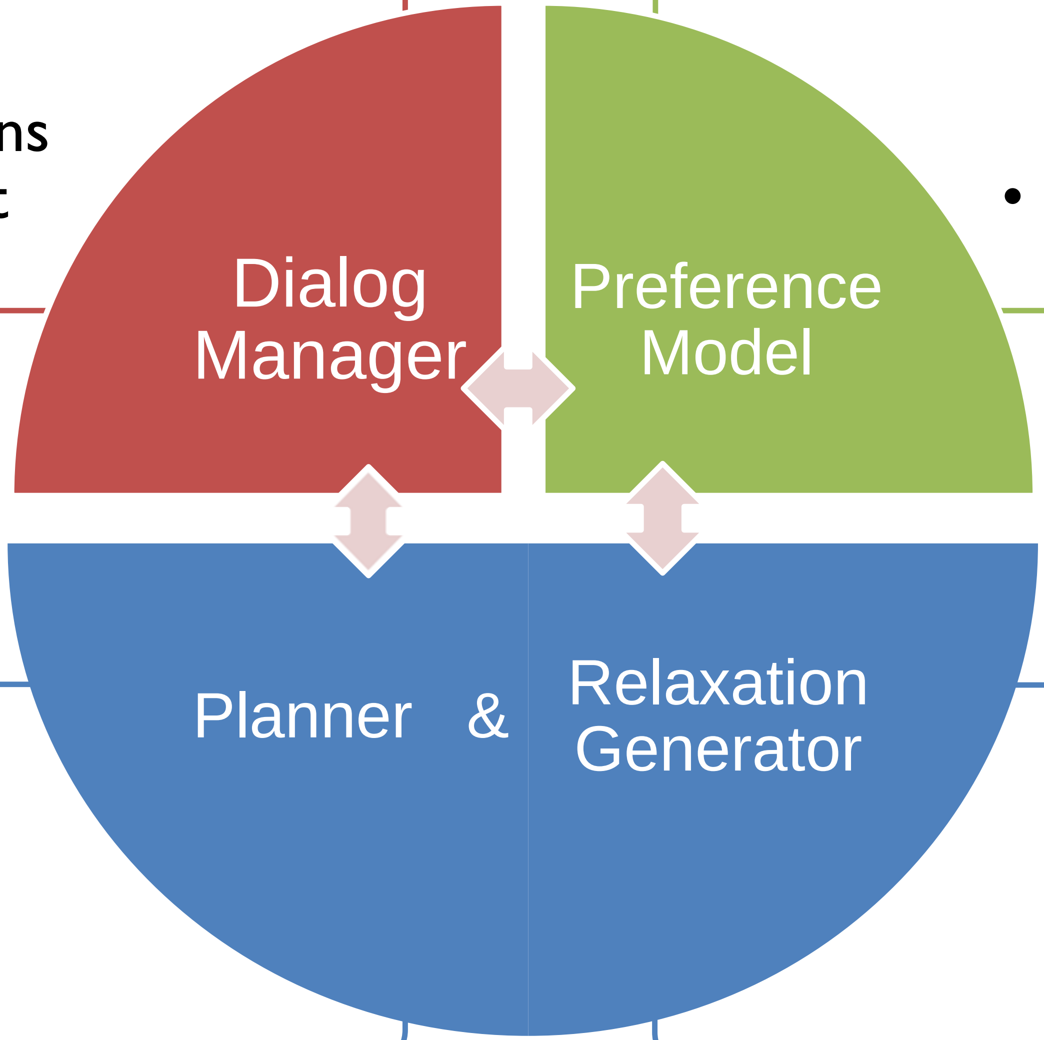

- Generate consistent plans.
- Extract conflicts between choices and constraints.

- Compute preferred temporal and domain relaxations to resolve conflicts.

Summary – MERS Only

‘Pengo 9000’
The R2D2 for over-subscribed planning problems

‘Pengo 900’
Resolving temporally infeasible plans through relaxations