


Good Network Updates for Bad Packets

Arne Ludwig,
Matthias Rost, Damien Foucard, Stefan Schmid




Updates happen


- Network updates happen
 - Changing security policies
- Network updates are challenging
 - Even with global view
- Potential high damage if fail
 - Security policy violation




Example



Example




Example



Waypoint Enforcement (WPE)


Example



- Eventual consistency


Example

Bad packet




- ✓ Eventual consistency
- Transient consistency?

Example




- ✓ Eventual consistency
- Transient consistency?

Example



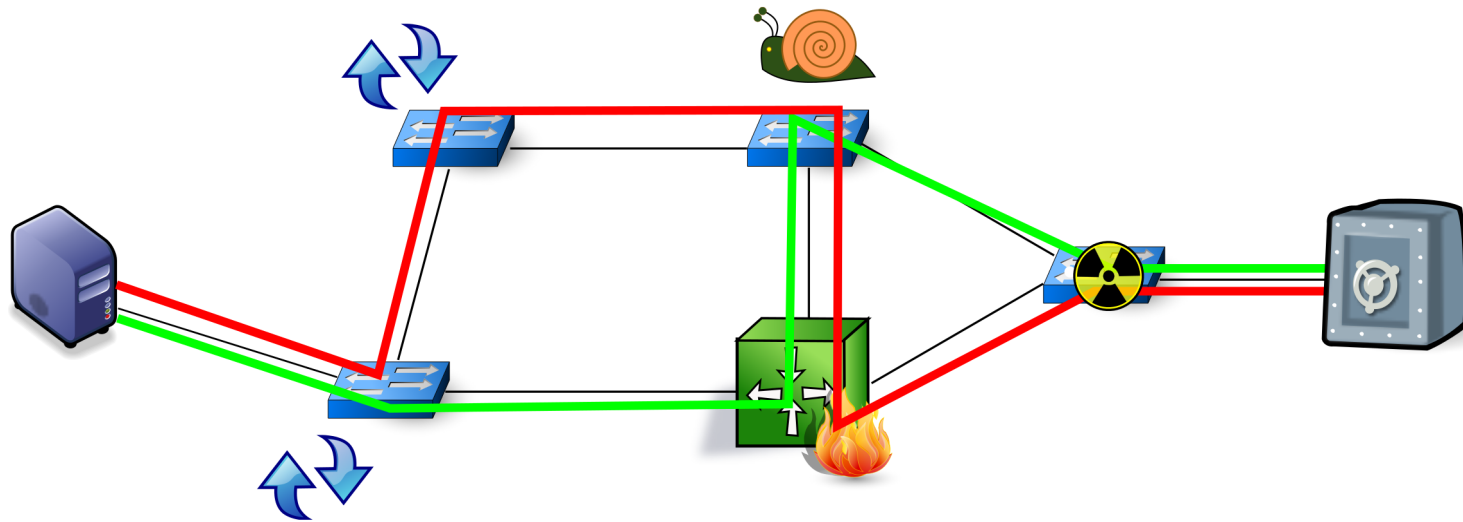
- ✓ Eventual consistency
- Transient consistency?

Example




- ✓ Eventual consistency
- Transient consistency?

Example



- ✓ Eventual consistency
- Transient consistency?

Example




- ✓ Eventual consistency
- × Transient consistency

Outline

- What could possibly go wrong?
- It's not a trivial thing!
- But we present an optimal solution.


Model and a Trivial Compression

| Solid lines = current path




Model and a Trivial Compression


- | Solid lines = current path
- : Dashed lines = new path
- Flow-specific path




Model and a Trivial Compression




- | Solid lines = current path
- ⋮ Dashed lines = new path
- Flow-specific path



Model and a Trivial Compression




| Solid lines = current path
| Dashed lines = new path
| Flow-specific path




- Safe to be updated
- Safe to be left untouched

Consistency Properties

- WPE = every packet traverses the waypoint at least once
- LF = loop freedom



Update all “simultaneously“?



Update all “simultaneously“?

Not possible in practice!

What could possibly go wrong?


Update all “simultaneously“?

Not possible in practice!

What could possibly go wrong?


Update times can vary significantly
(up to 10x higher than median
[Dionysus – SIGCOMM'14])

Update all “simultaneously“?




Update all “simultaneously“?

- Not waypoint enforced!




Delay s_1 ?




Delay s_1 ?


- Not loop free!




Update possible?



Update possible?




Update possible?




Update possible?

- Consistent transient states!



Rounds

- Round = set of parallel updates
- $R_1 = \{s_2\}, R_2 = \{s_3\}, R_3 = \{s_1\}$




→ Minimize number of rounds / communication overhead

Greedy Update Fails

- Greedy approach may:
 - take up to $\Omega(n)$ times more rounds
 - fail to find solution


See paper!



Greedy Update Fails


- Greedy approach may:
 - take up to $\Omega(n)$ times more rounds
 - fail to find solution

See paper!




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)
2. Switches $< WP$ (new), $< WP$ (old)




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)
2. Switches $< WP$ (new), $< WP$ (old)




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)
2. Switches $< WP$ (new), $< WP$ (old)
3. Remaining switches




WPE - Update Algorithm

1. Switches $< WP$ (new), $> WP$ (old)
2. Switches $< WP$ (new), $< WP$ (old)
3. Remaining switches




Constant in 3 rounds, but not LF!

LF and WPE Conflict




LF and WPE Conflict

- s_1, s_2 violate WPE; s_3, s_4 violate LF



Mixed Integer Program

Minimize
Rounds



$$\min R \quad (\text{Obj})$$

$$R \geq r \cdot x_v^r \quad r \in \mathcal{R}, v \in V \quad (1)$$

$$1 = \sum_{r \in \mathcal{R}} x_v^r \quad v \in V \quad (2)$$

$$y_{u,v}^r = 1 - \sum_{r' \leq r} x_u^{r'} \quad r \in \mathcal{R}, (u, v) \in E_{\pi_1} \quad (3)$$

$$y_{u,v}^r = \sum_{r' \leq r} x_u^{r'} \quad r \in \mathcal{R}, (u, v) \in E_{\pi_2} \quad (4)$$

$$a_s^r = 1 \quad r \in \mathcal{R} \quad (5)$$

$$a_v^r \geq a_u^r + y_{u,v}^{r-1} - 1 \quad r \in \mathcal{R}, (u, v) \in E \quad (6)$$

$$a_v^r \geq a_u^r + y_{u,v}^r - 1 \quad r \in \mathcal{R}, (u, v) \in E \quad (7)$$

$$y_{u,v}^{r-1 \vee r} \geq a_u^r + y_{u,v}^{r-1} - 1 \quad r \in \mathcal{R}, (u, v) \in E \quad (8)$$

$$y_{u,v}^{r-1 \vee r} \geq a_u^r + y_{u,v}^r - 1 \quad r \in \mathcal{R}, (u, v) \in E \quad (9)$$

$$y_{u,v}^{r-1 \vee r} \leq \frac{l_v^r - l_u^r - 1}{|V| - 1} + 1 \quad r \in \mathcal{R}, (u, v) \in E \quad (10)$$


$$\bar{a}_s^r = 1 \quad r \in \mathcal{R} \quad (11)$$

$$\bar{a}_v^r \geq \bar{a}_u^r + y_{u,v}^{r-1} - 1 \quad r \in \mathcal{R}, (u, v) \in E_{\overline{\text{WP}}} \quad (12)$$

$$\bar{a}_v^r \geq \bar{a}_u^r + y_{u,v}^r - 1 \quad r \in \mathcal{R}, (u, v) \in E_{\overline{\text{WP}}} \quad (13)$$

$$\bar{a}_t^r = 0 \quad r \in \mathcal{R} \quad (14)$$

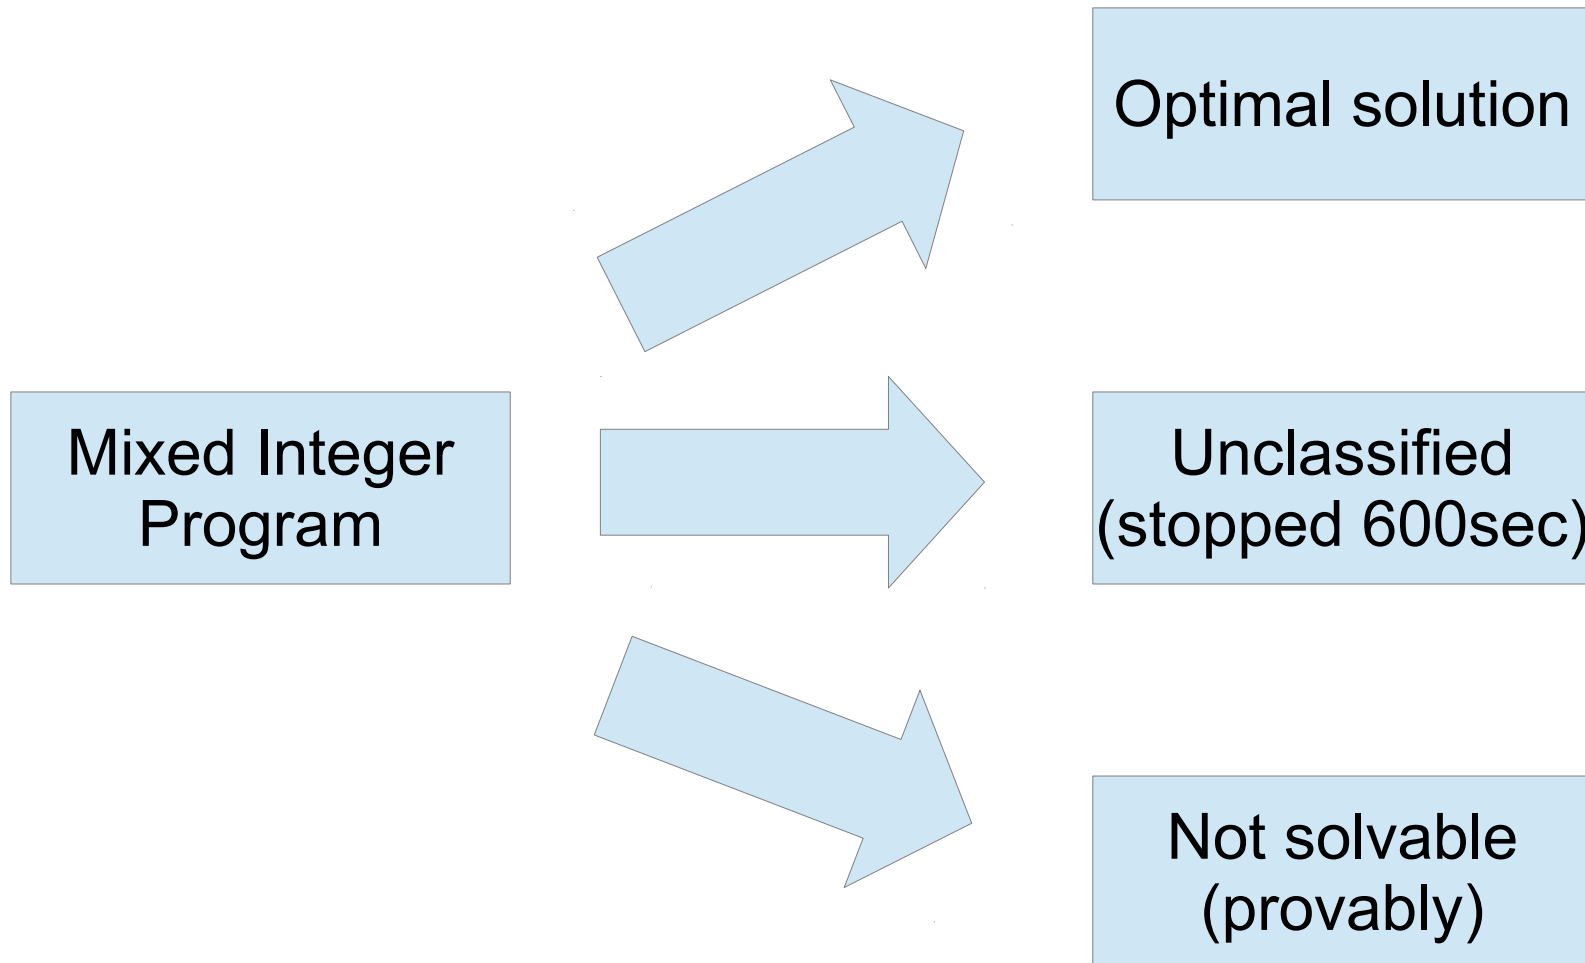
LF




WPE



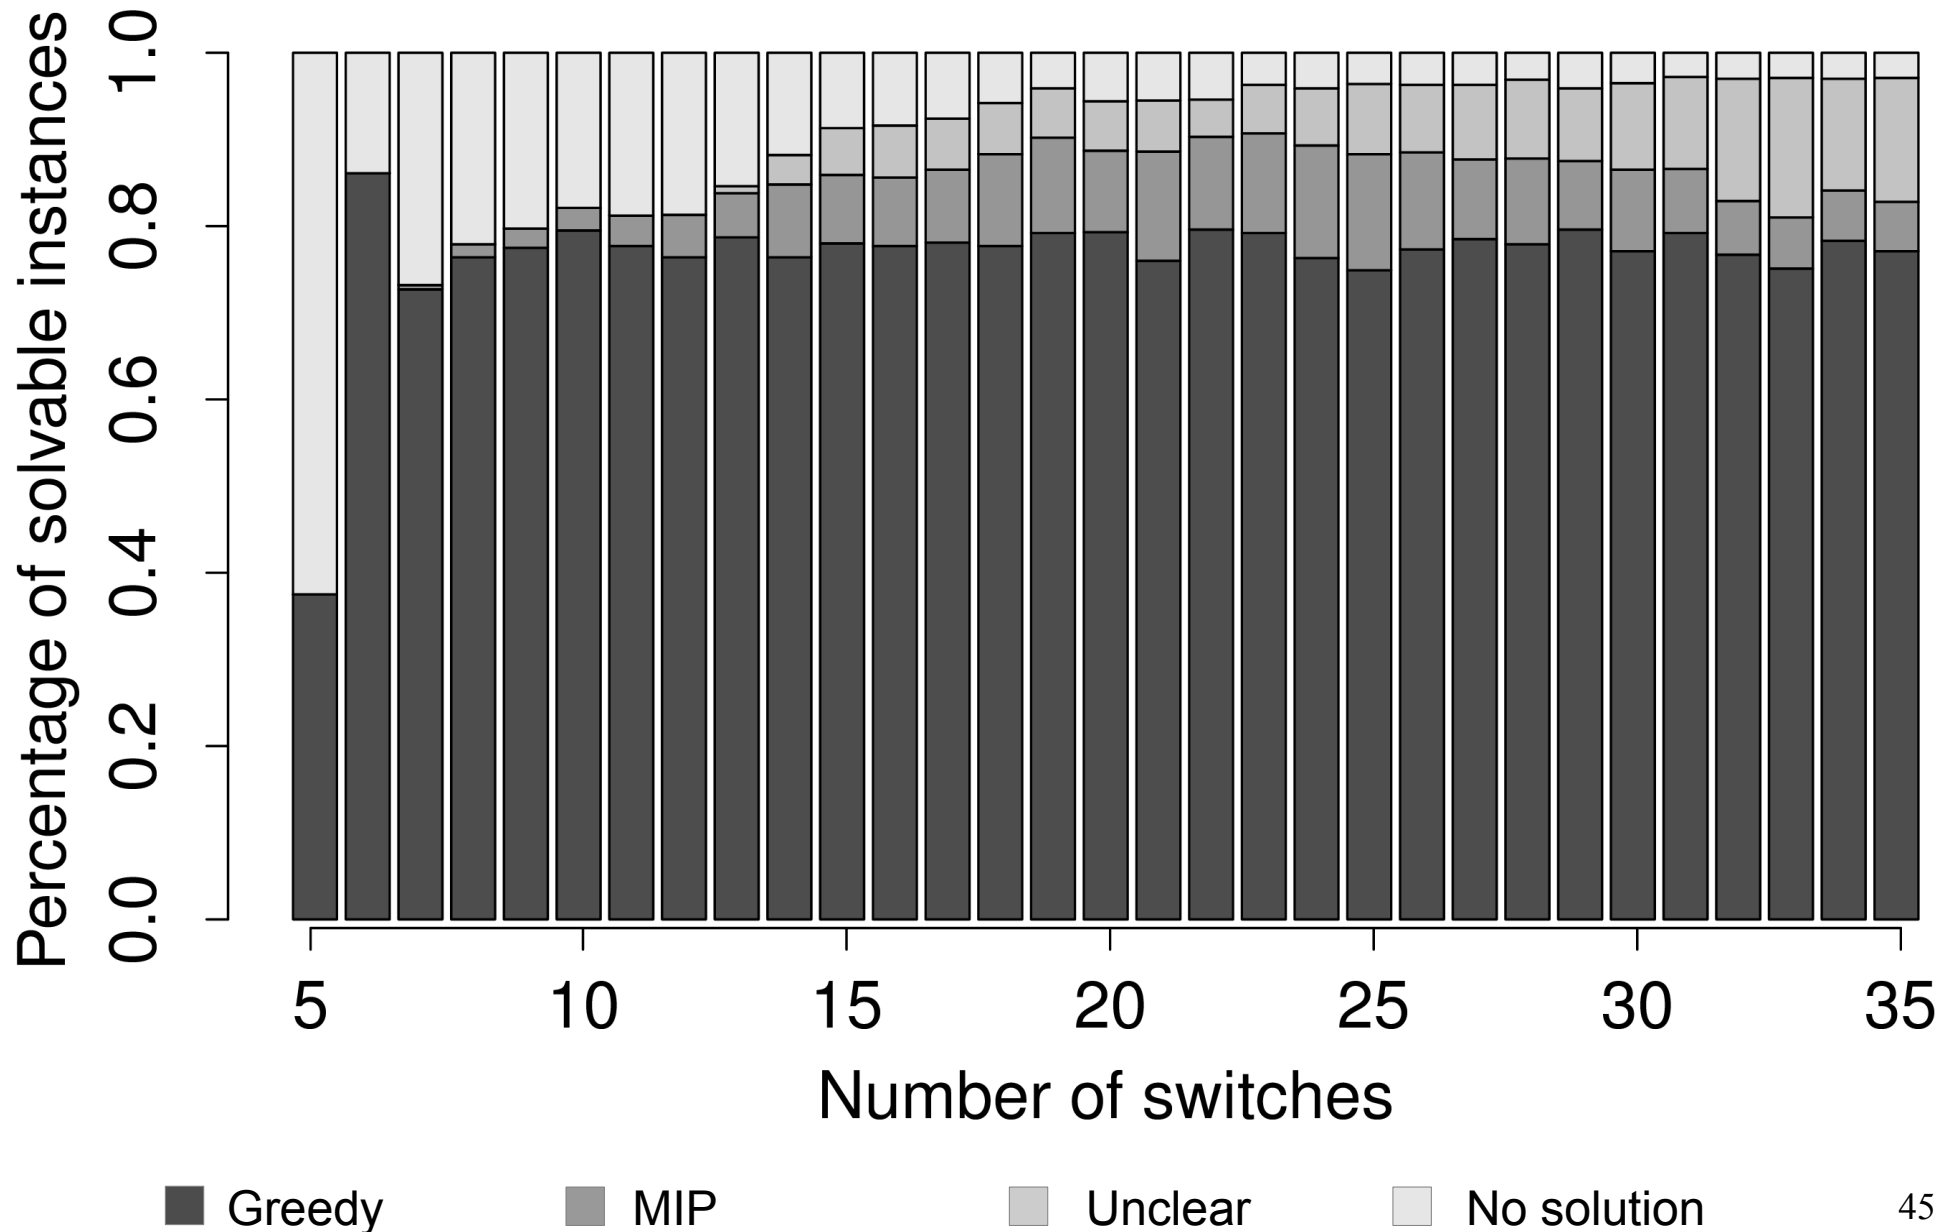
Mixed Integer Program




Solvability Analysis



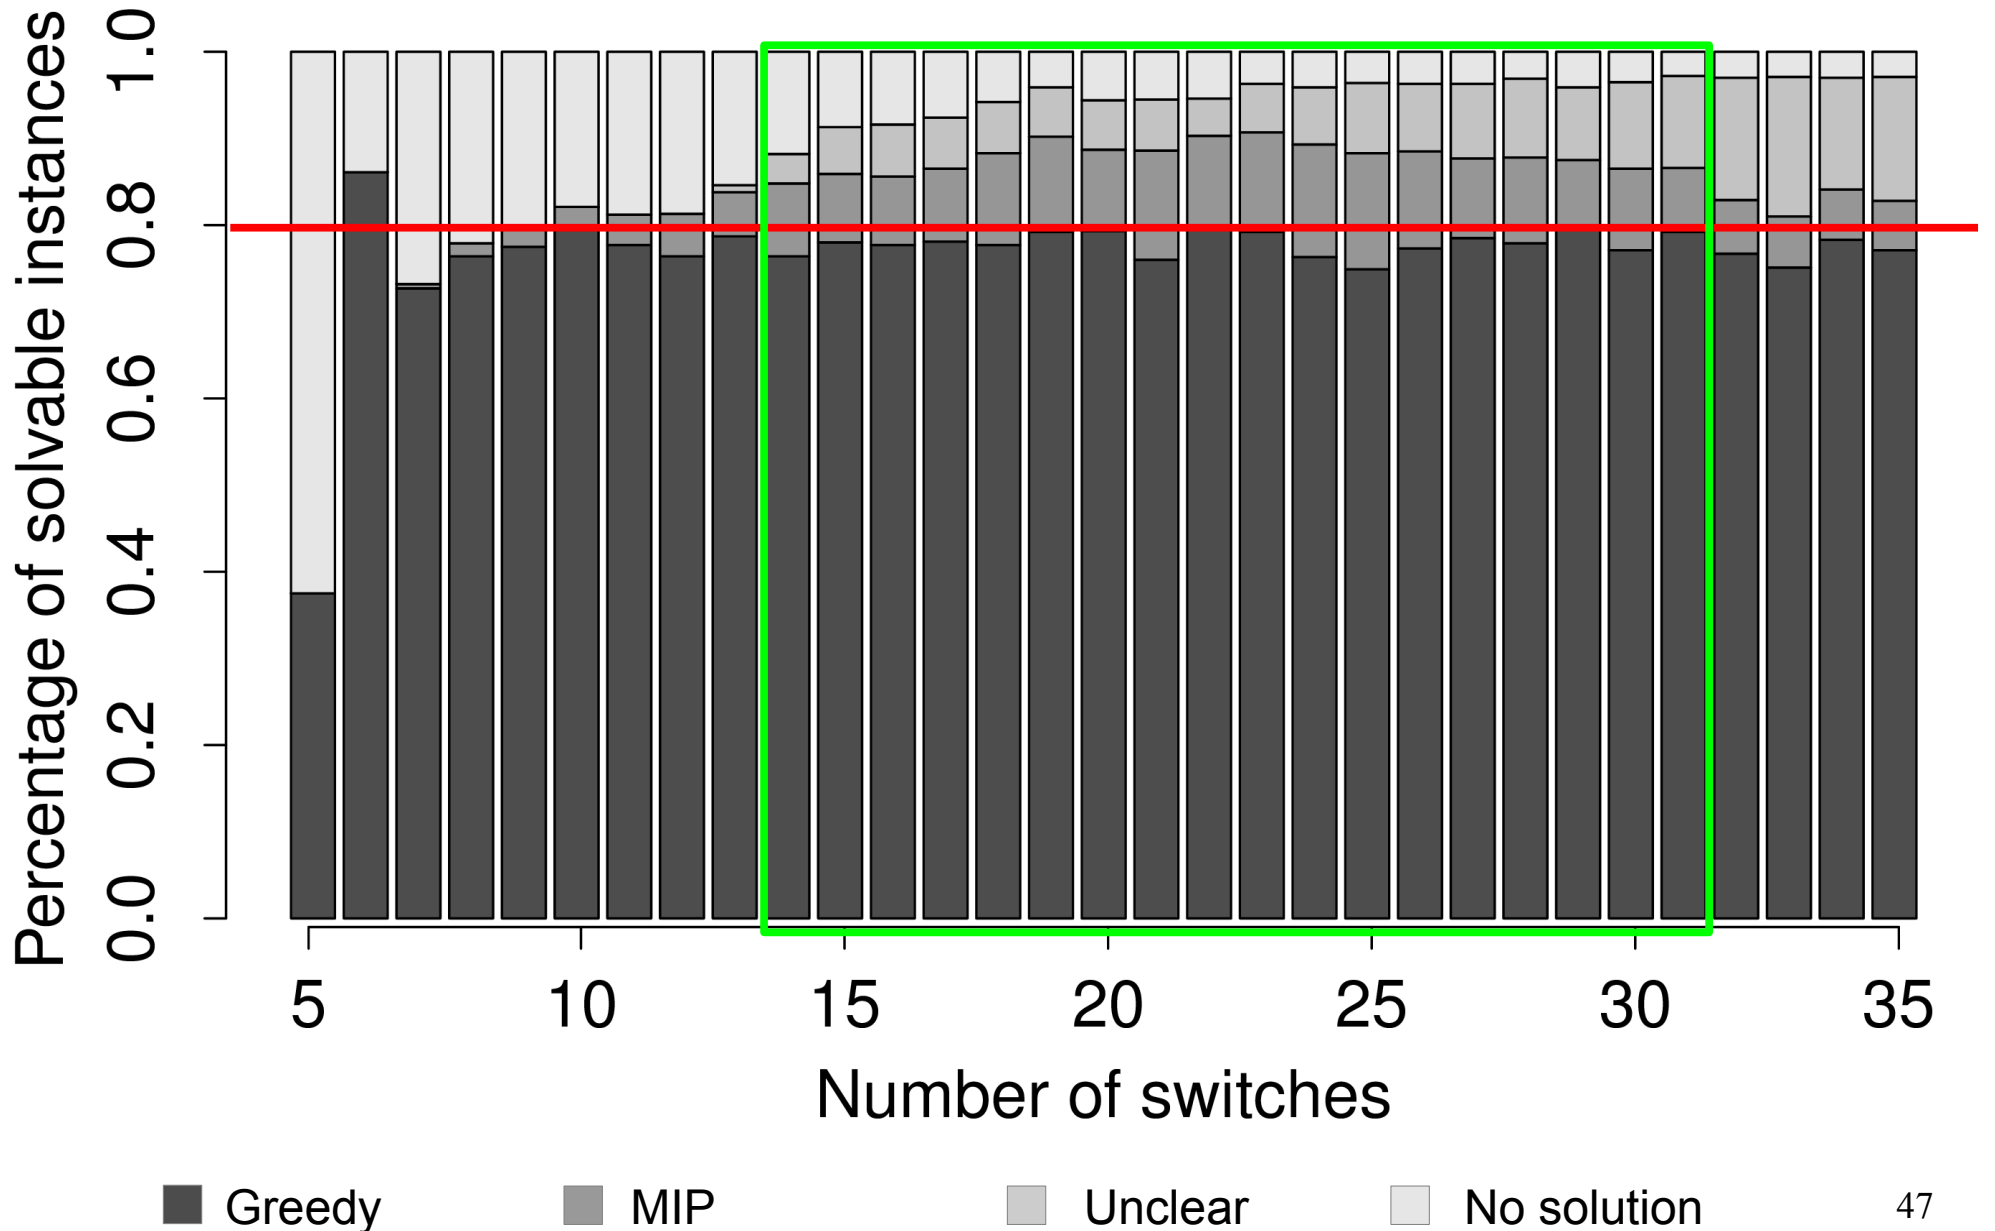
Solvability Analysis



Solvability Analysis



Solvability Analysis



Conclusion

- Transient consistency is not easy to guarantee
- LF and WPE might even conflict
- Greedy can fail to find consistent updates

Dynamic WPE + LF updates are hard to find!