


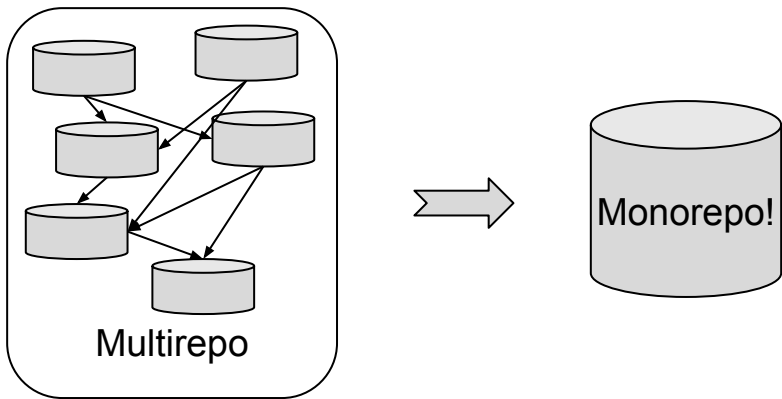
# Keeping Master Green at Scale

**Sundaram Ananthanarayanan**, Masoud Saeida Ardekani, Denis Haenikel,  
Balaji Varadarajan, Simon Soriano, Dhaval Patel, Ali-Reza Adl-Tabatabai  
(<https://eng.uber.com/research/keeping-master-green-at-scale/>)



# Monorepo is popular!

- Single, shared repo hosting companies' software assets



UBER



Google



Advantages of a Monorepo [Ciera et al. @ICSE'18]

- ✓ Simplified Dependency Management
- ✓ Improved Code Visibility

# Always **green** master considered hard

- **Monorepos handle a huge volume of commits every day**
- **Existing CI workflows do not guarantee an always green master**
  - Too hard at scale
- **Submit Queue guarantees an always-green master at scale**

# Outline

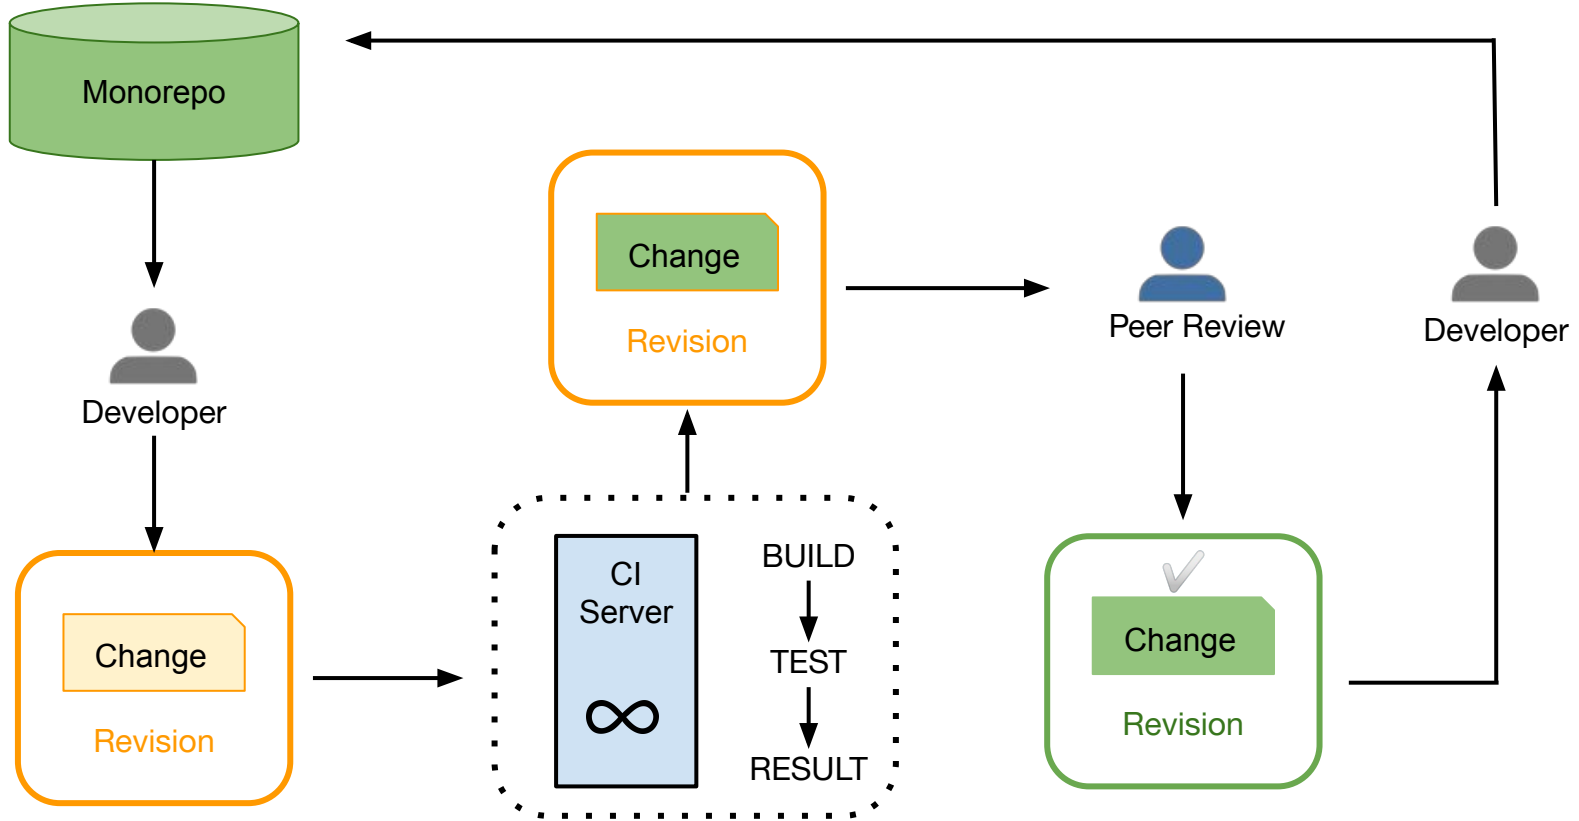
**01** Why **green** master is hard

**02** Probabilistic Speculation


**03** Conflict Analyzer

**04** Evaluation

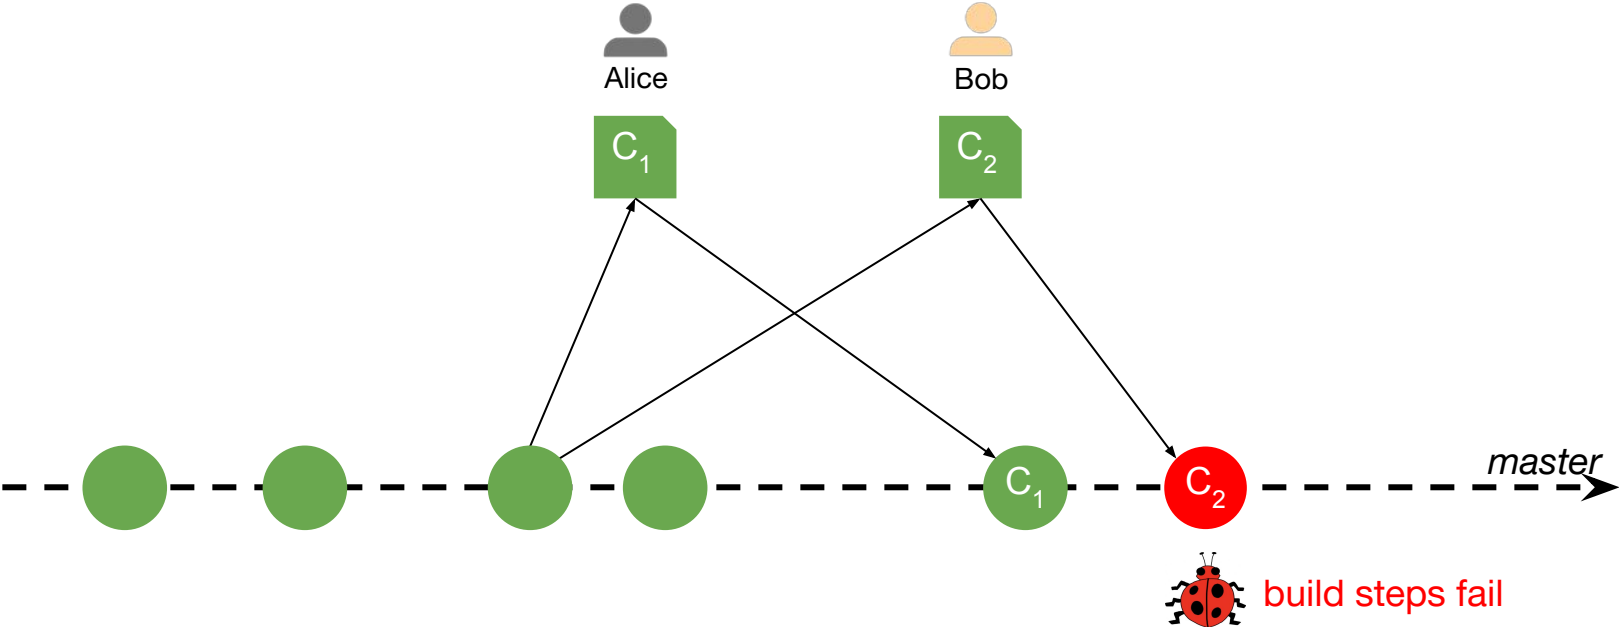
# Lifecycle of a change in monorepo



# Challenge: Concurrent conflicting changes



# Challenge: Concurrent conflicting changes



# Example of a real conflict



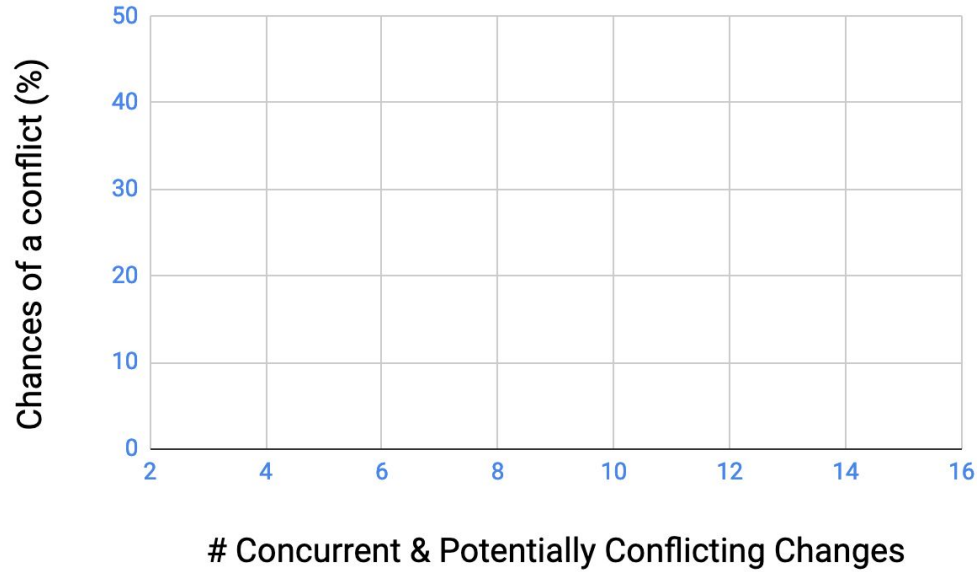
```
@@ -287,6 +288,9 @@ + (void)registerProtocols
    [currentGraph registerProtocol:@protocol(AnalyticsDeveloper)
        withImplementation:analyticsManager];
+
+ DependencyProvider *theProvider = [DependencyProvider theProvider];
+ [theProvider setAnalytics:analyticsManager];
+
```



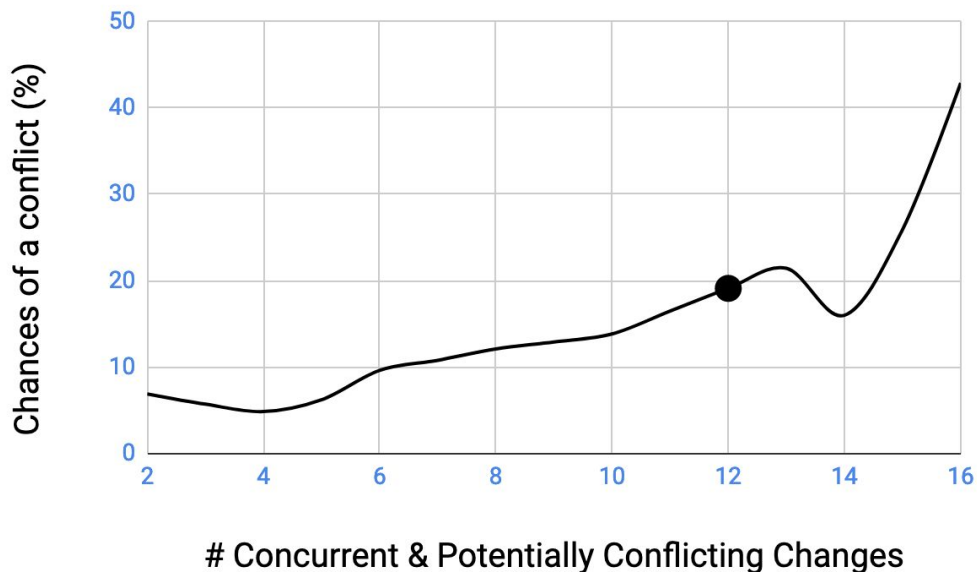
```
@@ -377,10 +378,14 @@ + (void)registerProtocols
+
+ #pragma mark - Networking
+ [currentGraph registerProtocol:@protocol(NetworkingConfiguration)
-         withFactory:^id {
-             return [[NetworkingConfiguration alloc] init];
-         }];
+         withImplementation:networkingConfiguration];
+
+ HTTPProvider *theProvider = [HTTPProvider theProvider];
+ [theProvider setLogger:logger];
```



# How often conflicts happen?



# How often conflicts happen?



**Observation:** Chances of a conflict  $\uparrow$  from 5% to 40% as #. of concurrent & potentially conflicting changes  $\uparrow$

# Drawbacks of a **red** master

## Delayed rollouts



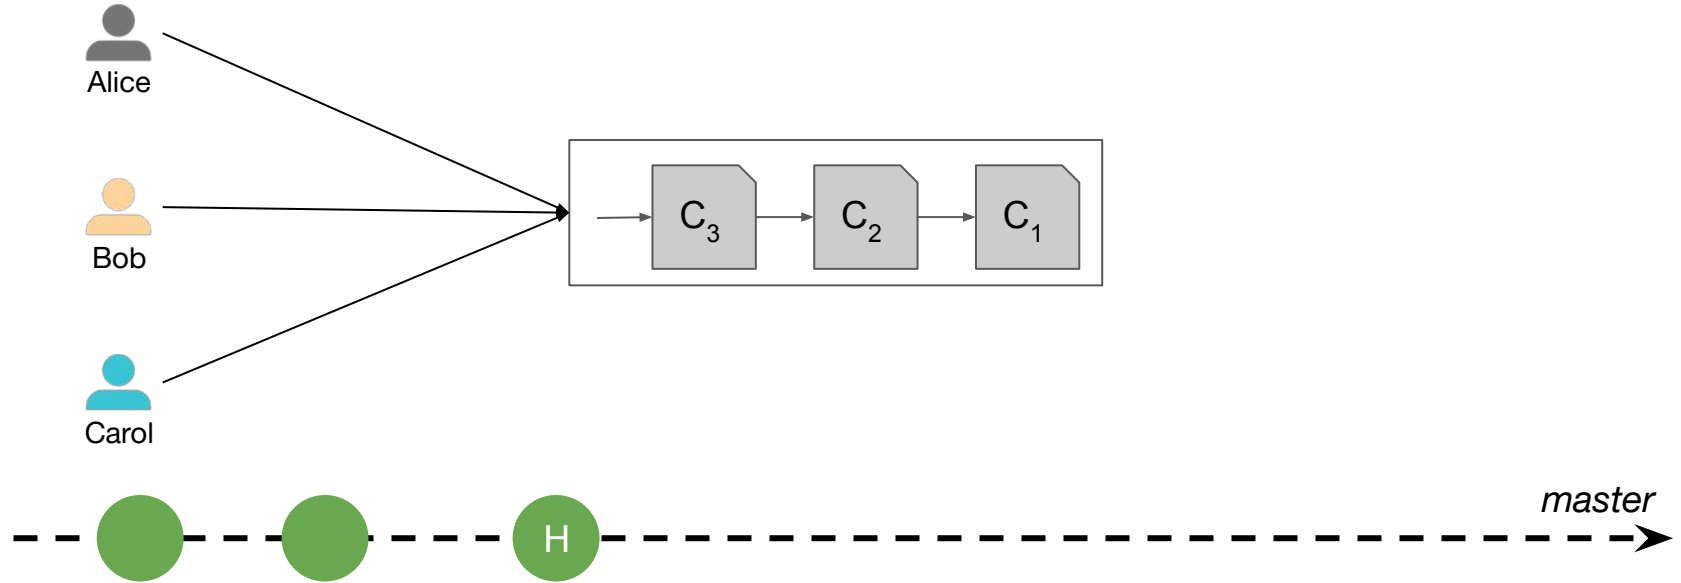
## Hampered Productivity



## Complex rollbacks

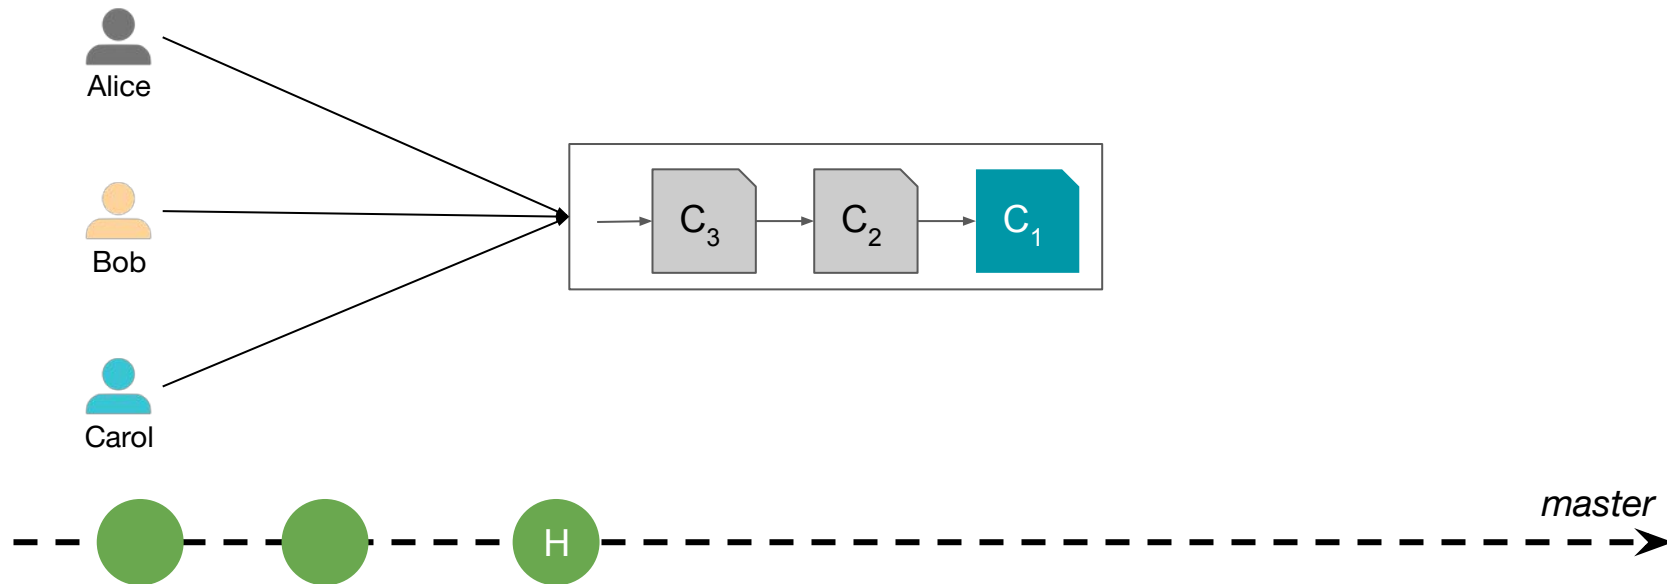


# Keeping master green: Queue



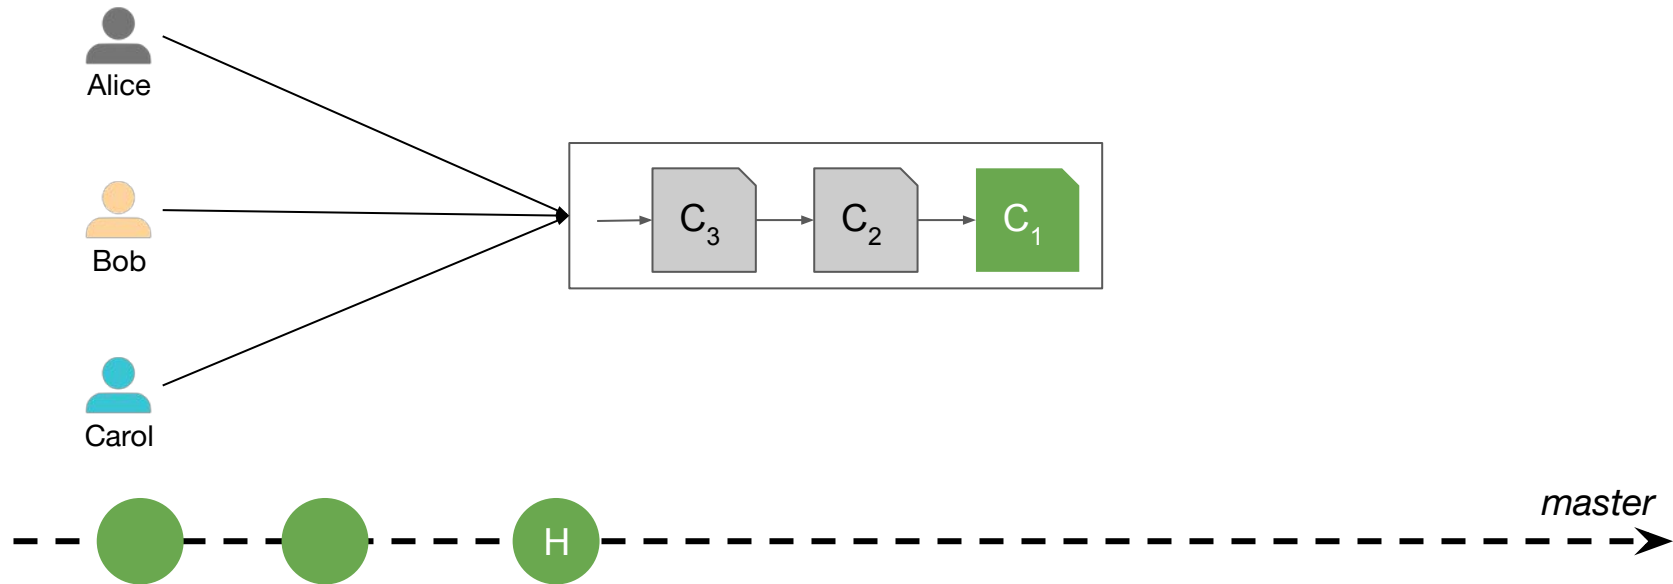
Alice, Bob, Carol enqueue changes they want to commit

# Keeping master green: Queue



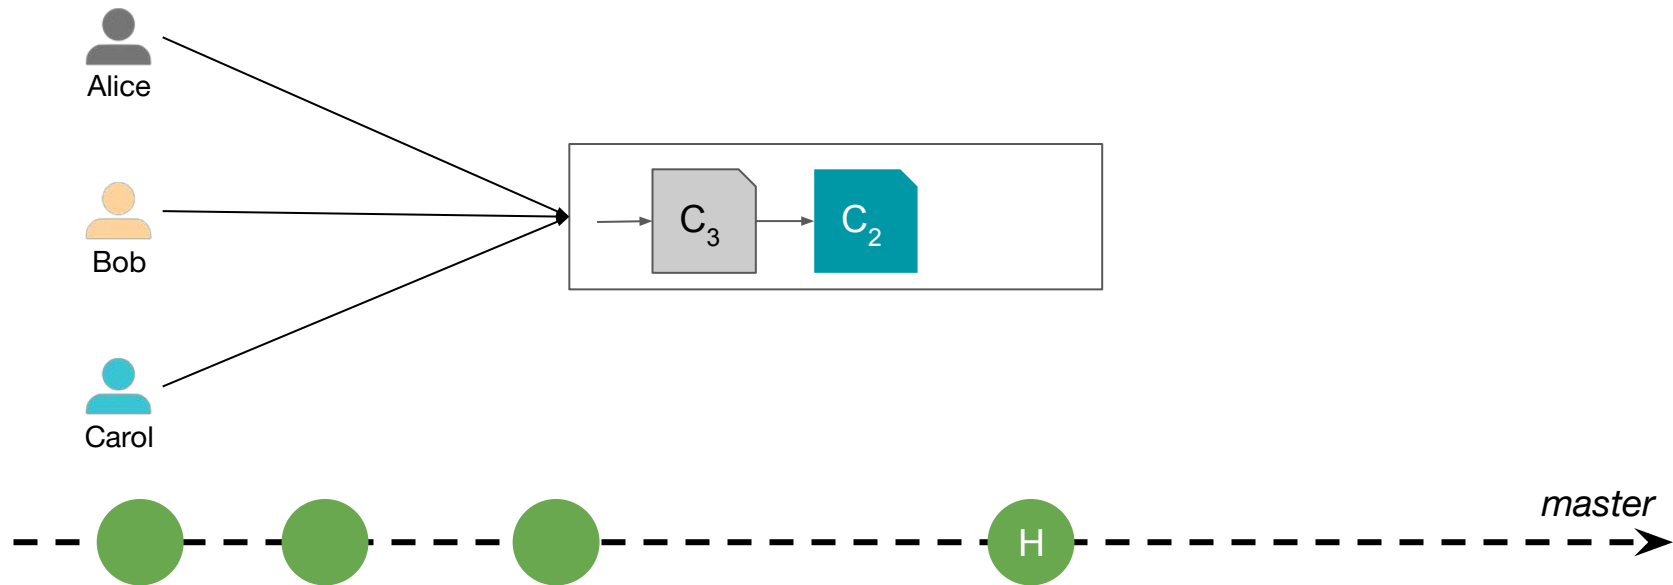
C<sub>1</sub> is built and tested against mainline head (H).

# Keeping master green: Queue

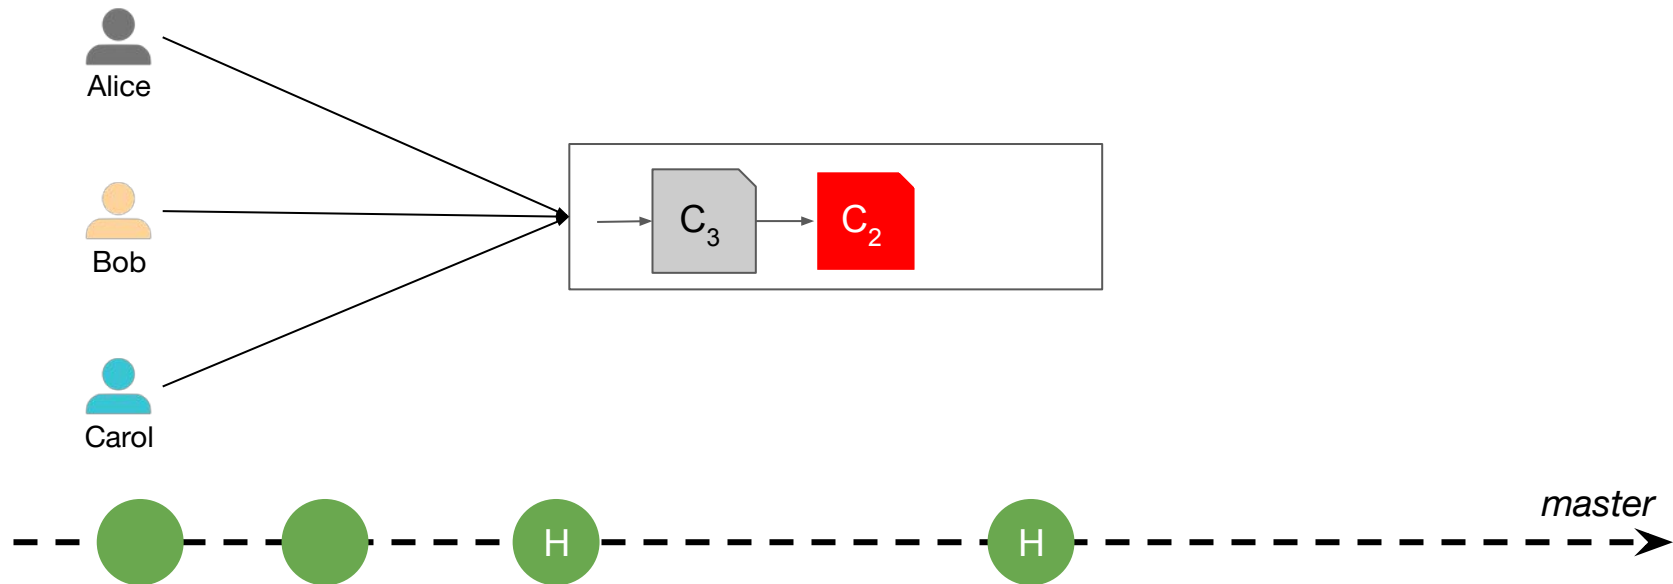


Build steps for  $H \oplus C_1$  succeed.

# Keeping master green: Queue



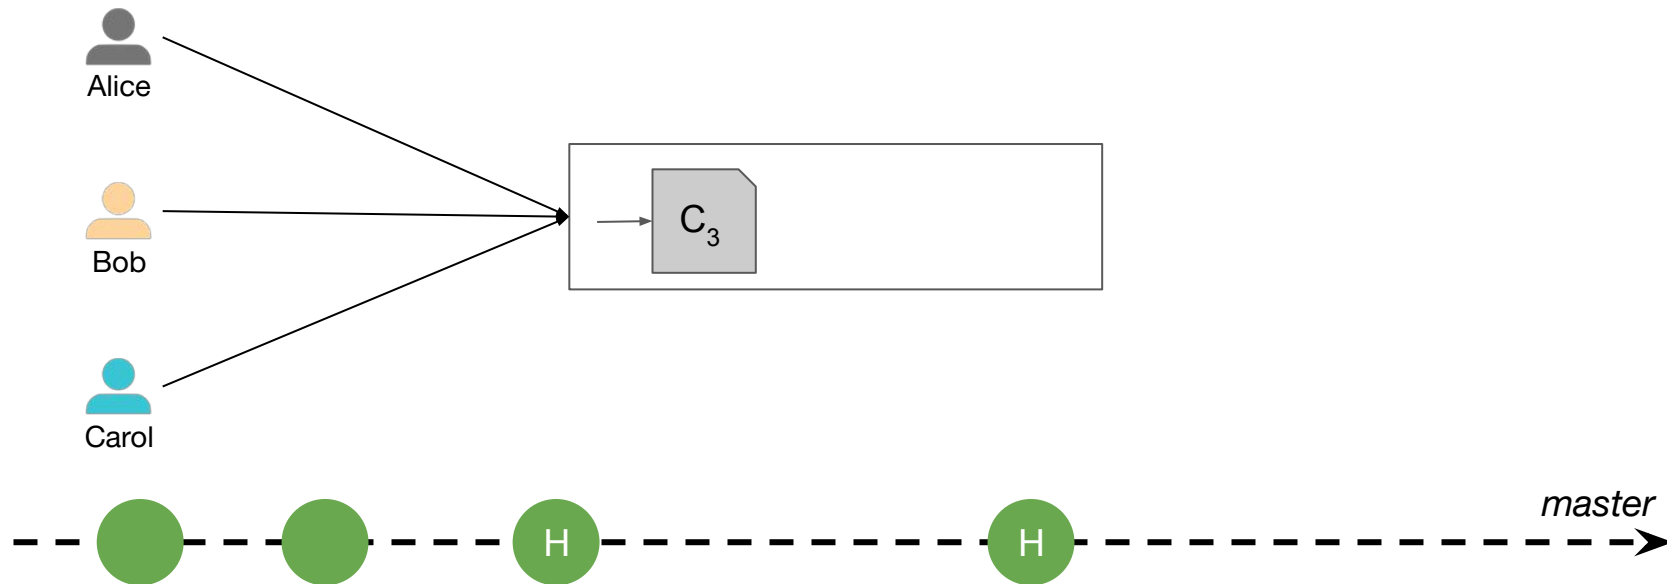
# Keeping master green: Queue



Build steps for  $H \oplus C_2$  fails and  $C_2$  is rejected.

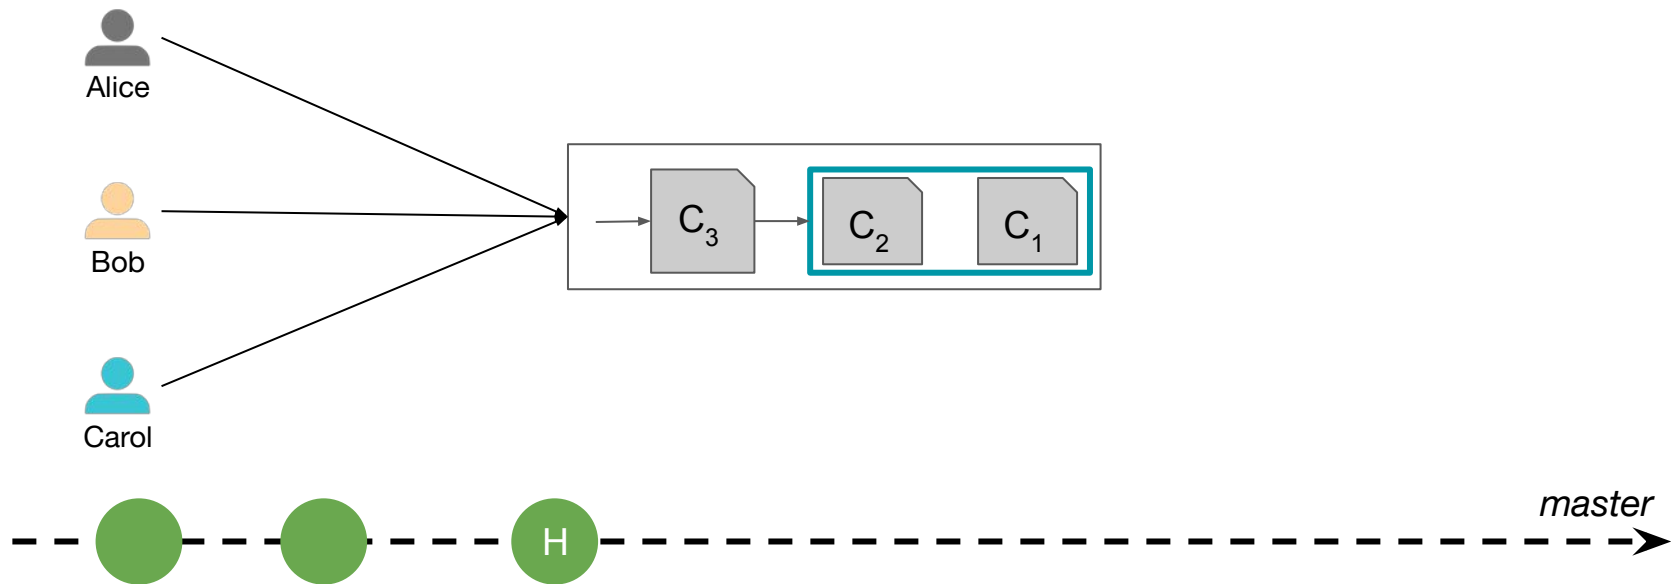


# Keeping master green: Queue



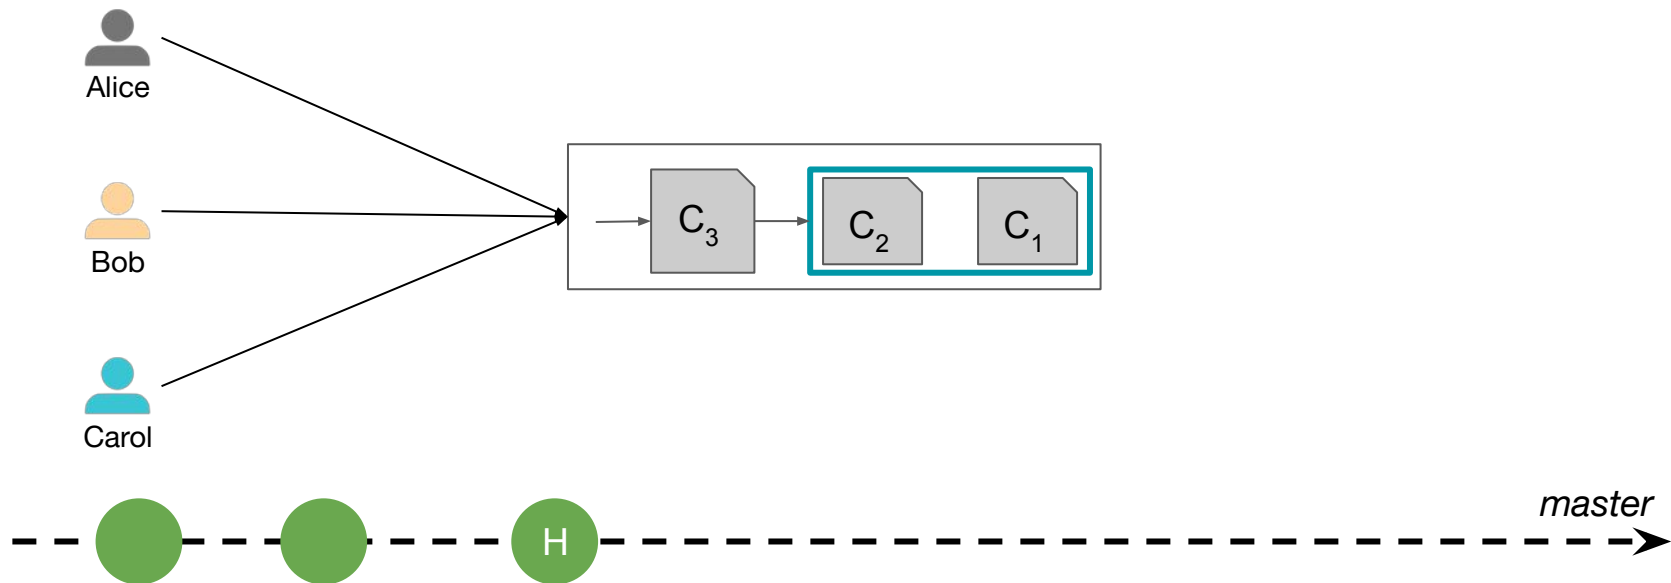
- ✓ Guarantees an always green master by serializing changes
- ✗ Does not scale to 1000s of changes/day

# Keeping master green: Batching changes



C<sub>1</sub> and C<sub>2</sub> are batched and build steps are run.

# Keeping master green: Batching changes



- ✓ Improves the throughput if batches succeed more often than not
- ✗ Testing batches masks intermediate changes that fail
- ✗ Batches will fail often as the size of the batch increases

**What happens when batches fail?**

# Keeping master green: Goals

## Guarantee serializability

- Illusion of a single queue when committing changes
- Git only offers serializability of patches

## Provide reasonable SLAs


- Overheads should be short enough for developers to trade speed for correctness!

**Challenge:** how to do this at scale? (1000s of commits/day)

# Submit Queue: Overview


## Speculation Engine

- Speculates on success/failure of changes
- Builds speculation graph



## Conflict Analyzer


- Determines independent changes
- Constructs conflict graph



## Planner Engine


- Selects most valuable builds from speculation engine
- Execute builds and commit changes

# Speculation Tree



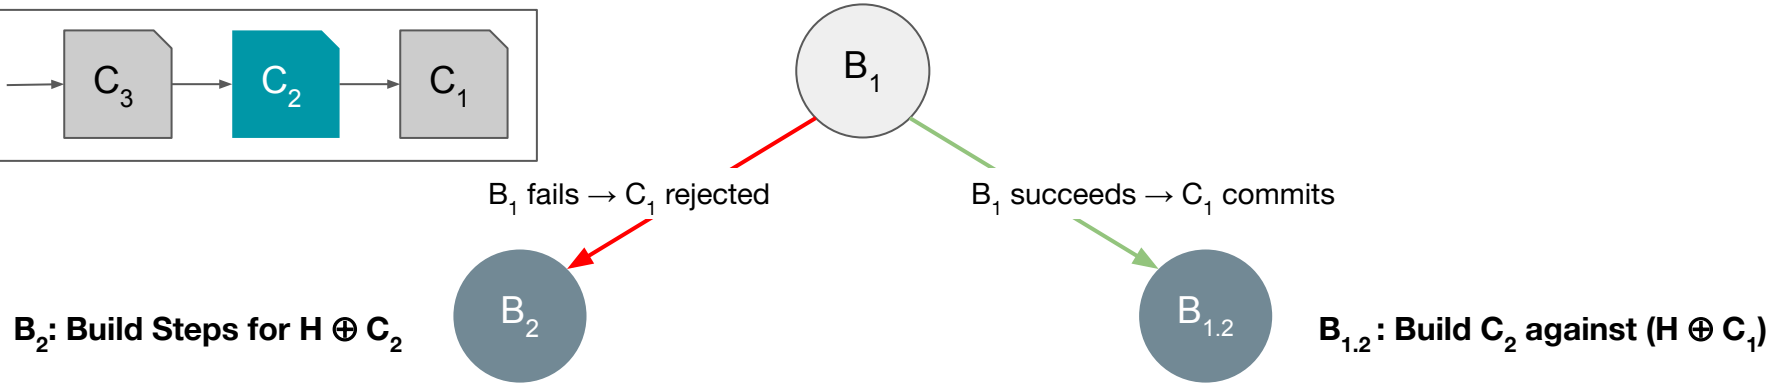

$C_1, C_2, C_3$  - pending changes

# Speculation Tree



$B_1$ : Build Steps for  $H \oplus C_1$

# Speculation Tree

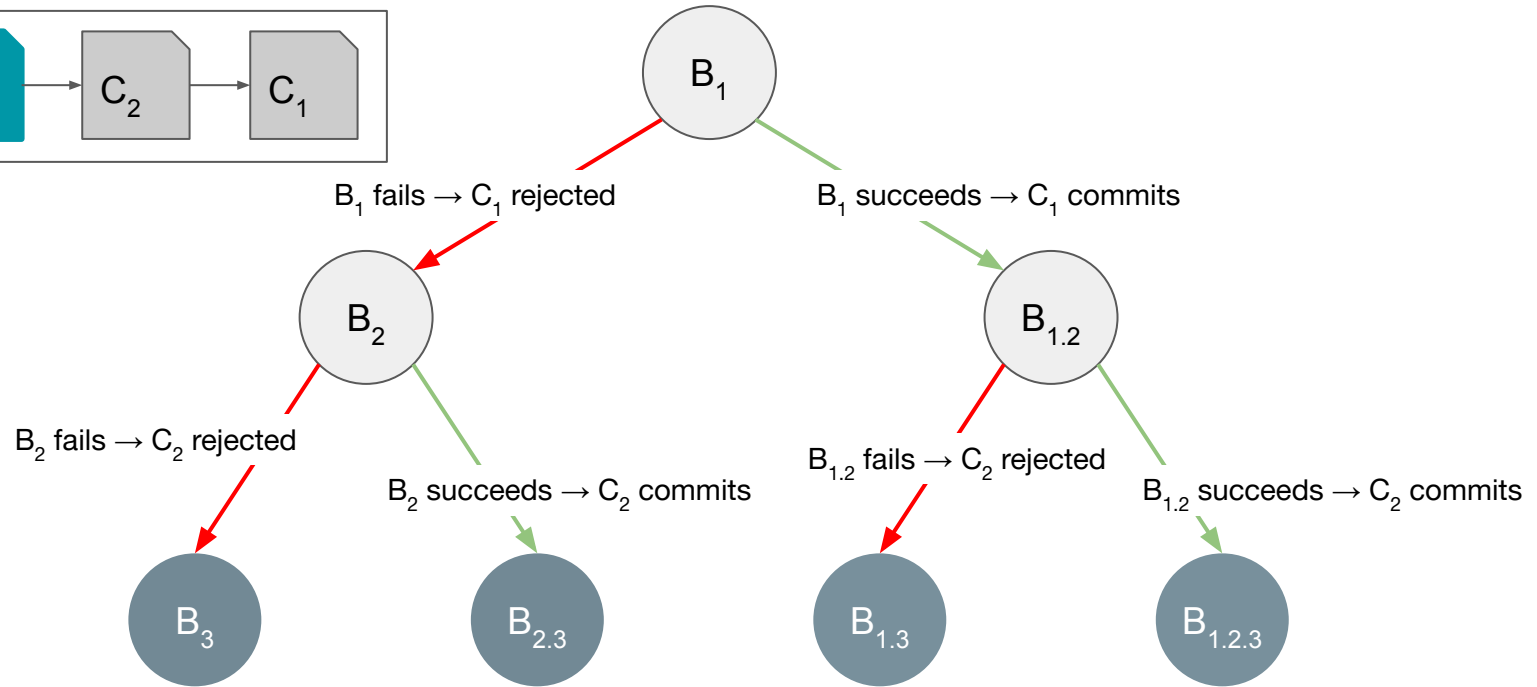



U B E R

1. Precompute the outcome of committing  $C_2$  under different realities
2. Commit or reject  $C_2$  based on the outcome of  $B_1$  and one of  $\{B_2, B_{1.2}\}$



# Speculation Tree




**Challenge:** Which builds to run?

# Approach #1: Speculate Them All

Speculate on all possible outcomes equally


- Selects builds in a breadth-first order



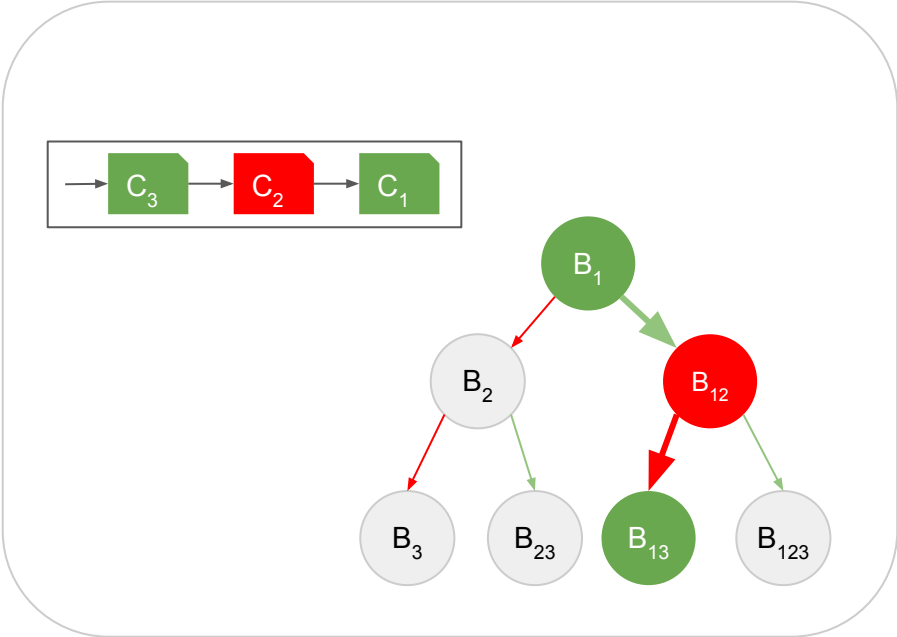
**Does not scale for 1000s of changes/day**

- Need to run  $2^n$  builds in parallel to commit ' $n$ ' changes

**Leads to substantial waste of resources**



# Speculate Them All: Resource Wastage




# Speculate Them All: Observation

If we select and execute builds whose *outcomes* are **most likely to be needed**, then we require only  **$n$**  (out of  $2^n$ ) **builds**.

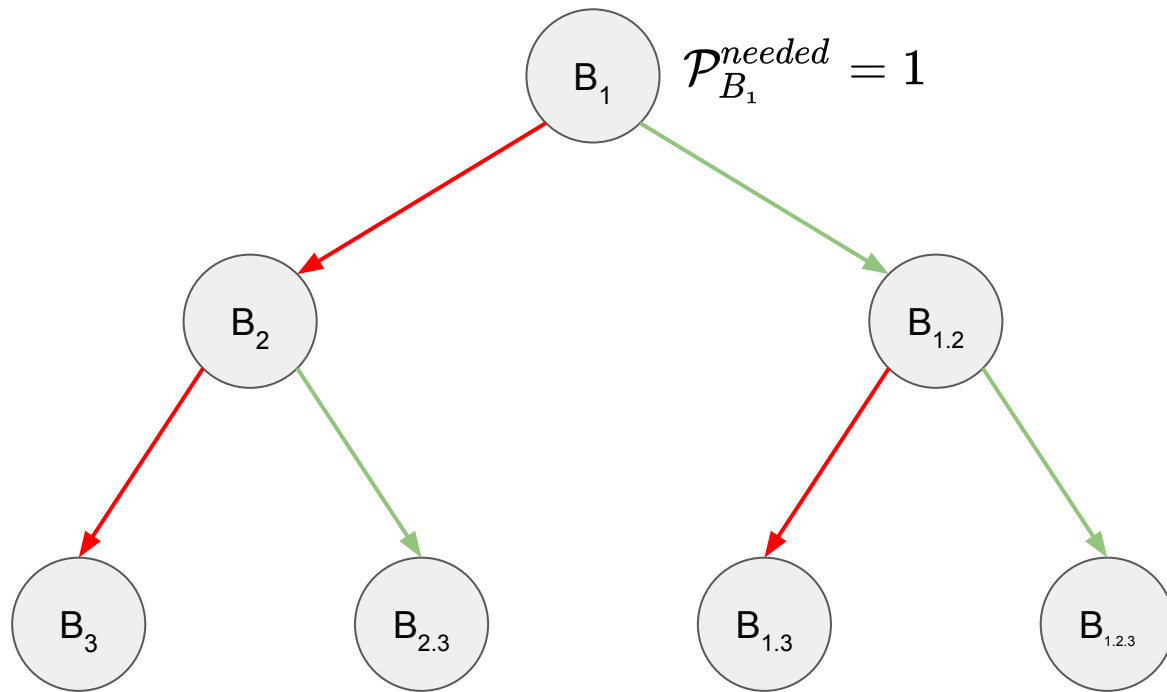
**Challenge:** Which ' $n$ ' builds are likely to be needed?

# Probabilistic Speculation




$\mathcal{P}_{B_c}^{needed}$  represents the prob. the result of the build  $B_c$  is used to make to commit/reject C.

# Probabilistic Speculation



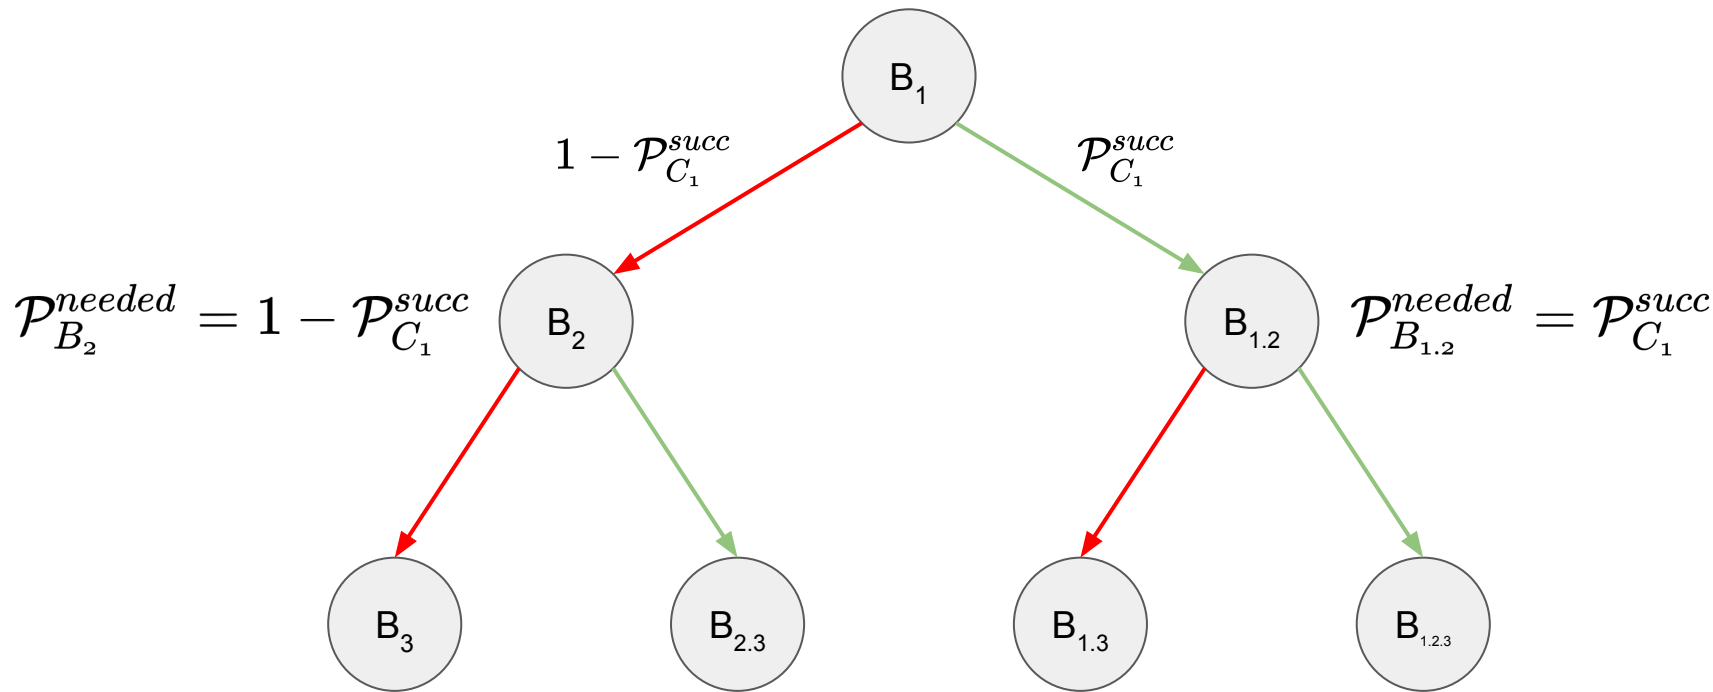
Root  $B_1$  is always needed as is used to determine if  $C_1$  can be committed

# Probabilistic Speculation



$\mathcal{P}_{C_1}^{succ}$  represents the prob. that change  $C_1$  succeeds individually

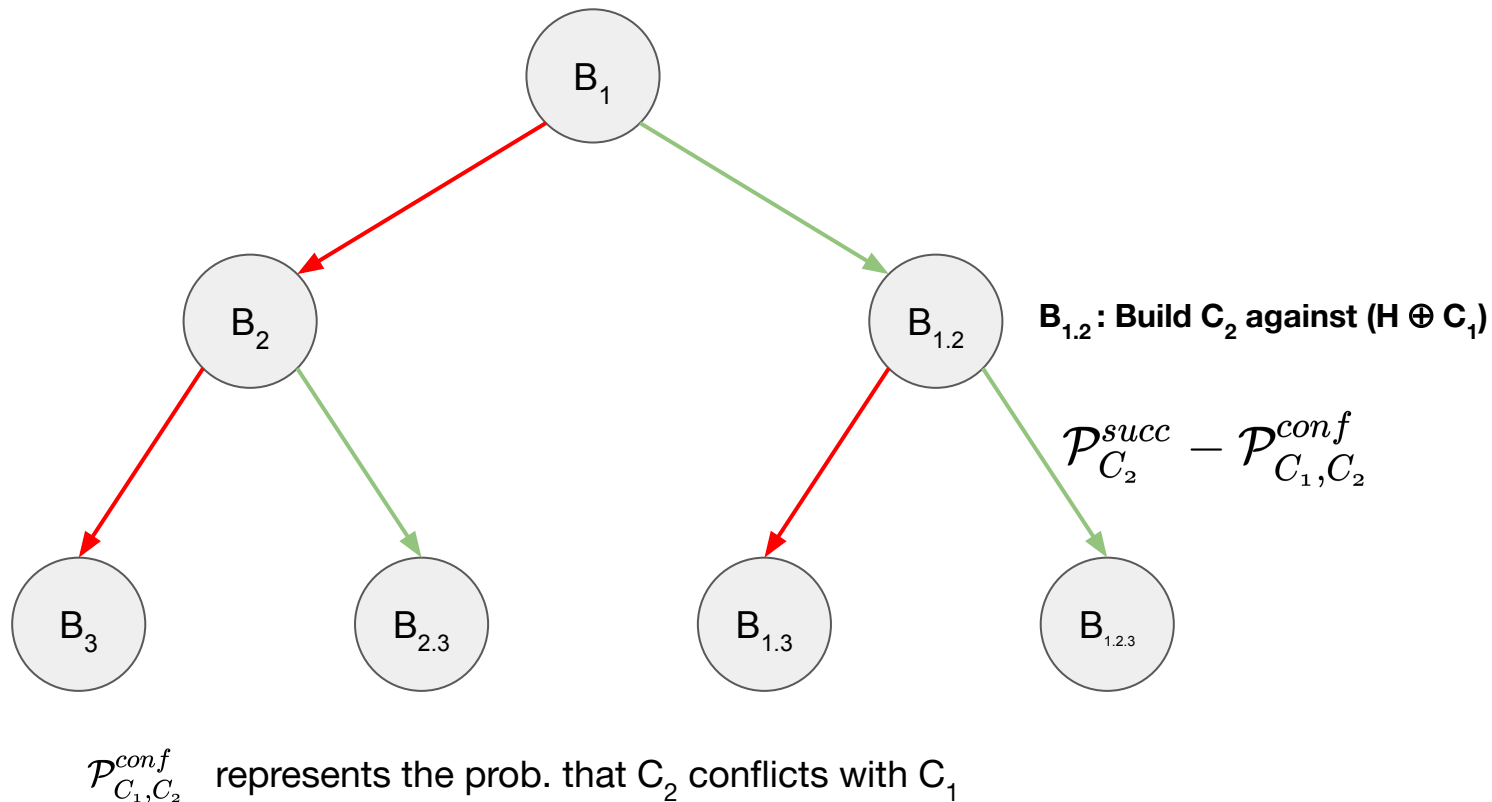
# Probabilistic Speculation




$\mathcal{P}_{C_1}^{succ}$  represents the prob. that change  $C_1$  succeeds individually



# Probabilistic Speculation



# Probabilistic Speculation




$$\mathcal{P}_{B_{1.2.3}}^{needed} = \mathcal{P}_{C_1}^{succ} \cdot (\mathcal{P}_{C_2}^{succ} - \mathcal{P}_{C_1, C_2}^{conf})$$

# Probabilistic Speculation: Summary

Choose **most valuable builds** by determining

- Probability of success of a change  $\mathcal{P}_{C_i}^{succ}$
- Probability of a conflict bet. changes  $\mathcal{P}_{C_i, C_j}^{conf}$



# Evaluating $\mathcal{P}_{C_i}^{succ}$ and $\mathcal{P}_{C_i, C_j}^{conf}$

- **Logistic regression to train prediction models**
  - Feature set includes 100+ hand-picked features
  - Prediction accuracy of 97%

## Change

- # affected targets
- # git commits
- # files changed
- status of pre-submit checks

## Developer

- developer name
- employment proficiencies

## Speculation

- dynamic features to re-adjust weights based on initial predictions
- # speculations succeeded
- # speculations failed

# Features for Training ML Models

## Change

- # affected targets
- # git commits
- # files changed
- status of pre-submit checks

## Revision

- revision is a container for changes
- # changes submitted
- revert and test plans
- # Submit attempts made

## Developer

- developer name
- employment proficiencies


## Speculation

- dynamic features to re-adjust weights based on initial predictions
- # speculations succeeded
- # speculations failed

# Conflict Analyzer

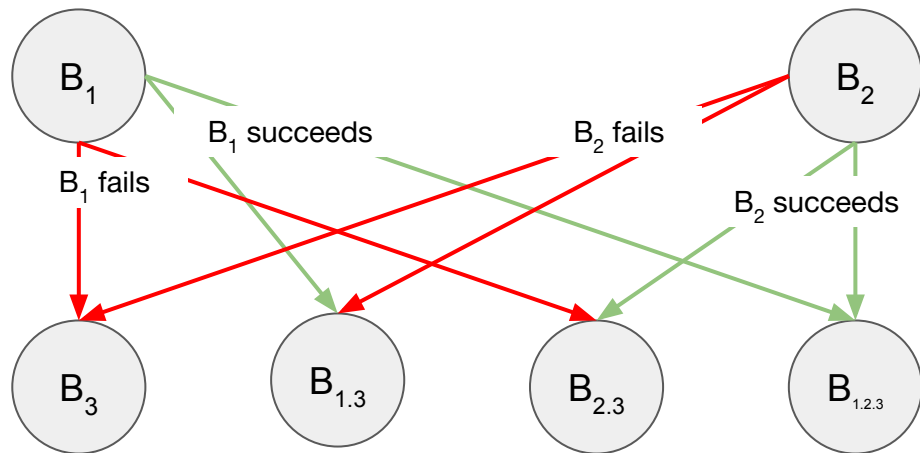

- **So far, we assumed all changes potentially conflict with each other**
  - Cannot commit in parallel
- **What if changes can be proved to be independent?**
  - Commit changes in parallel
  - Trim speculation space
- **We use Conflict Analyzer to find independent changes**

# Conflict Analyzer: Commit Changes in Parallel



Conflict graph for changes  $C_1$ ,  $C_2$ ,  $C_3$  where  $C_1$  and  $C_2$  are independent and conflict with  $C_3$ .


# Conflict Analyzer: Commit Changes in Parallel



**Insight:** Changes  $C_1$  and  $C_2$  can be committed in parallel.

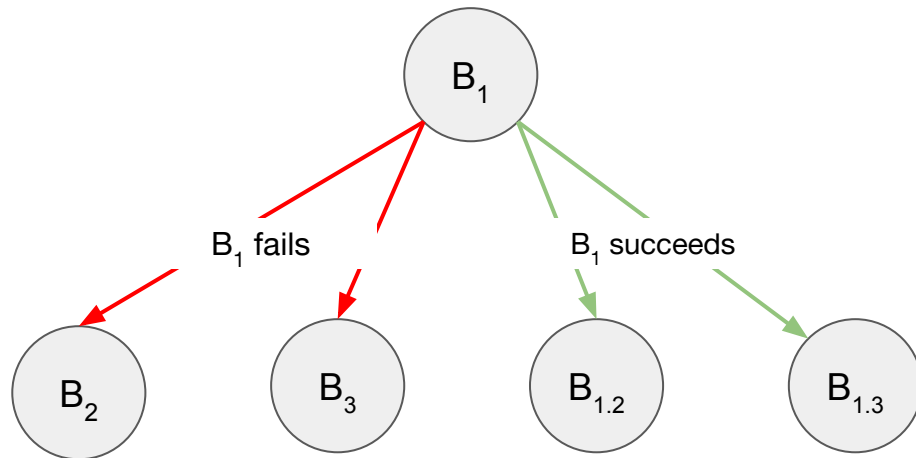



# Conflict Analyzer: Trim Speculation Space



Conflict graph for  $C_1, C_2, C_3$  where  $C_1$  conflicts with independent changes  $C_2$  and  $C_3$ .

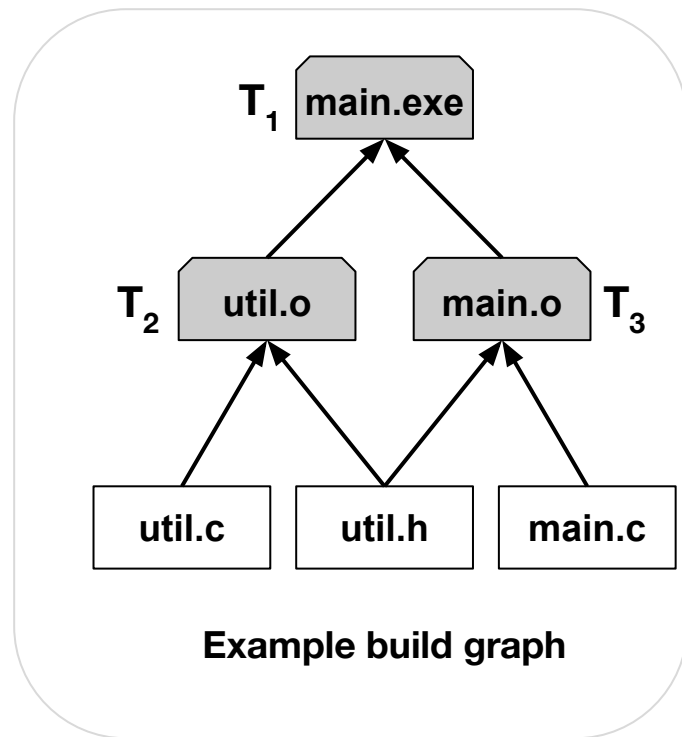
# Conflict Analyzer: Trim Speculation Space



**Insight:** Because  $C_3$  does not speculate on  $C_2$ , # of possible builds for  $C_3$  reduces to 2.

# Conflict Analyzer: Detecting conflicts at scale

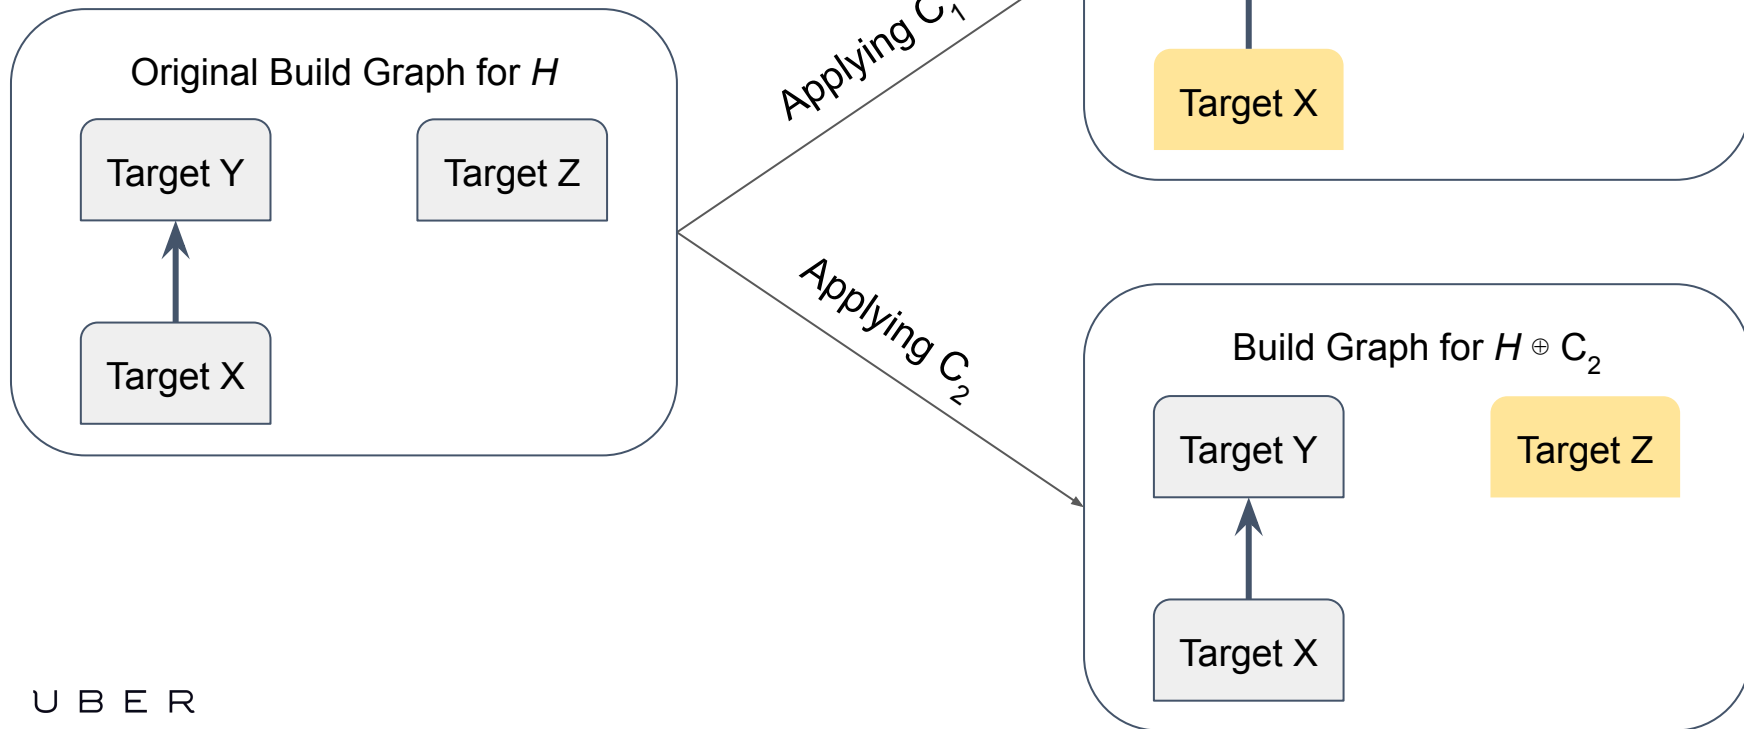
- **Build system** to detect if changes are independent
- Code partitioned into smaller entities called *targets*
- Every change affects a set of *targets*



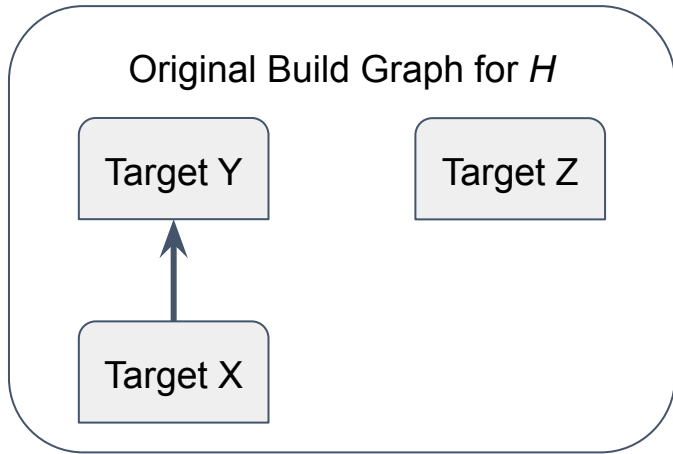
# Detecting Conflicts: Intuition

Two changes are independent if they affect a disjoint set of targets.

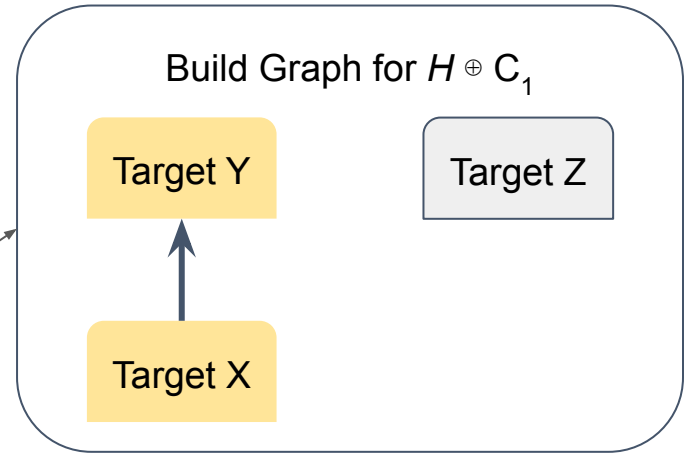
# Detecting Conflicts: Example



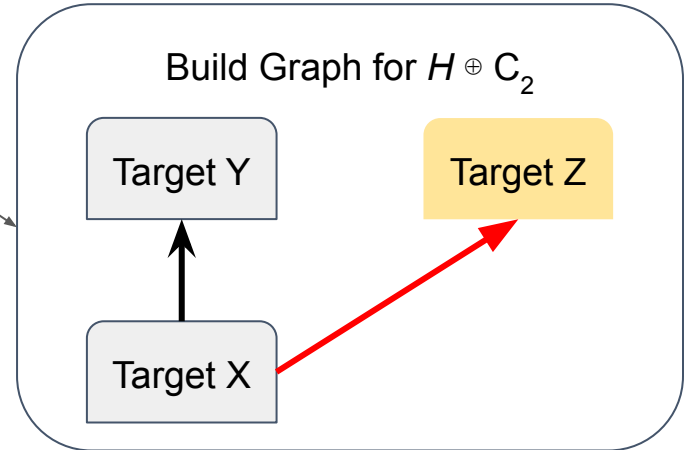
# Detecting Conflicts: Puzzle



Applying  $C_1$

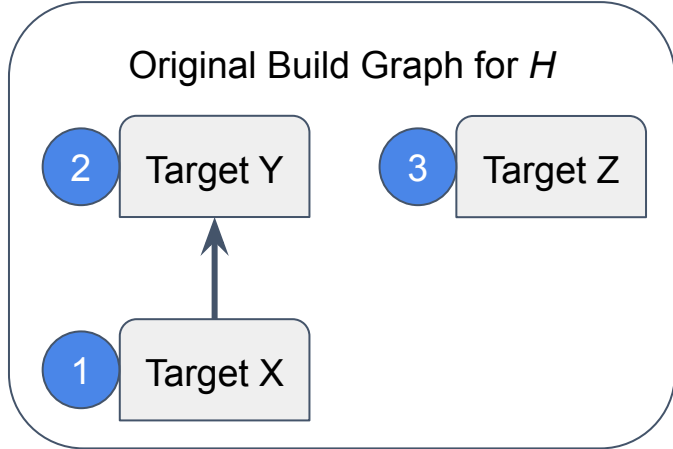


Applying  $C_2$

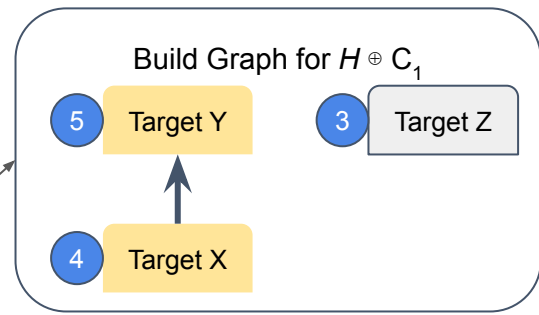


- $C_1$  and  $C_2$  are conflicting
- But, the intersection of affected targets is empty!

# Detecting Conflicts: Composition

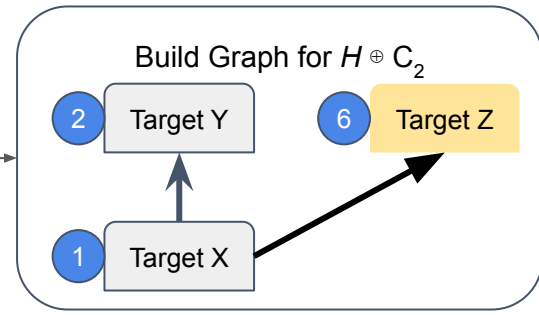


Applying  $C_1$



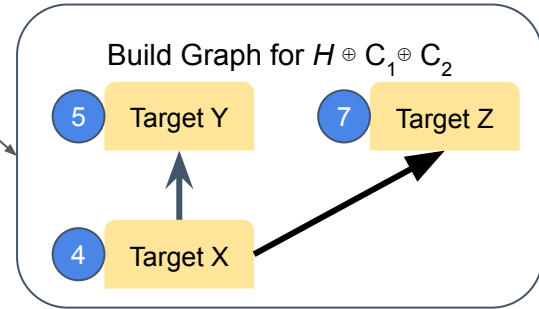
$\{(x, 4), (y, 5)\}$

Applying  $C_2$



$\{(z, 6)\}$

Applying  $C_1 \oplus C_2$



$\{(x, 4), (y, 5), (z, 7)\}$

$\{(x, 4), (y, 5)\} \cup \{(z, 6)\} \neq \{(x, 4), (y, 5), (z, 7)\}$   
Thus,  $C_1$  and  $C_2$  are conflicting!

# Detecting Conflicts: Summary

- **Intersection Approach**

✗ Does not detect all kinds of conflicts

- **Union Approach**

✗ Determining conflicts for  $n$  changes requires  $n^2$  build graphs!

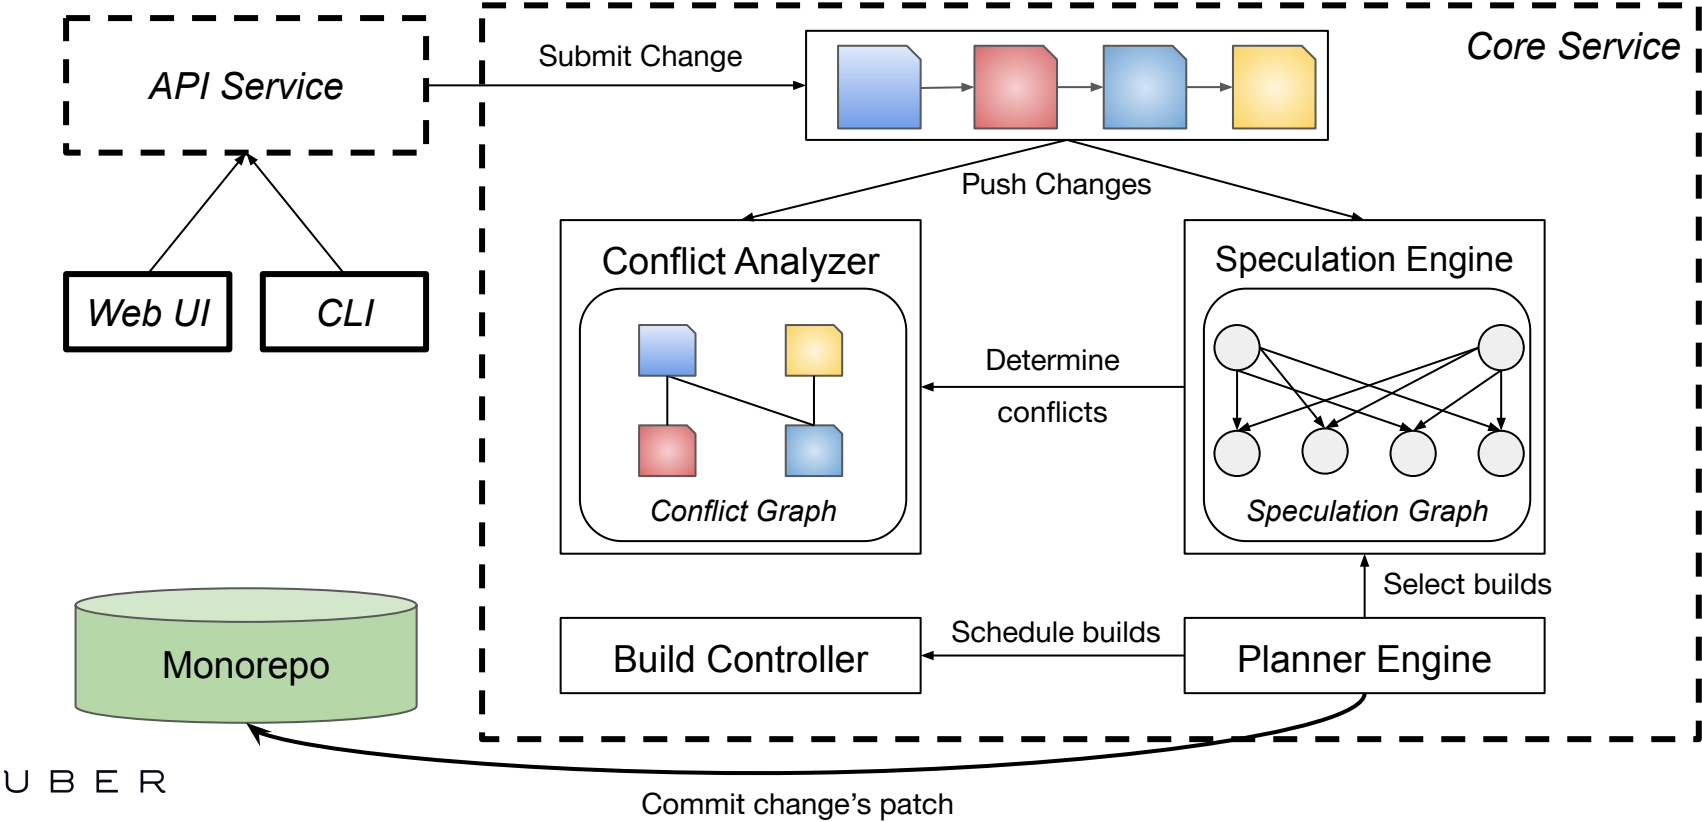
- **Hybrid Approach**

✓ Only 7.9% of changes cause a change to the build graph

- Union Graph Approach (details in paper)



# Submit Queue: Architecture Overview



# Evaluation


## Questions

- How does Submit Queue performance compare against other strategies?
  - Queue, Speculate-all, Optimistic
- What is the impact of conflict analyzer?

## Setup

- Implemented an Oracle that predicts outcome of a change correctly
  - All results normalized against the Oracle
- Ingested real changes into our system at different rates

# Evaluation: Submit Queue Performance



# Evaluation: Submit Queue Performance (P50)

**Speculate-all**

#Changes / Hour	100	200	300	400	500
500	11.21	10.05	9.44	9.19	9.04
400	11.82	10.69	9.80	9.75	9.42
300	13.08	11.87	11.00	10.74	10.58
200	15.30	14.04	13.14	12.90	12.72
100	7.41	6.63	6.46	6.44	6.24

Speculate-all suffers up to **15x** slowdown compared to the Oracle.

# Evaluation: Submit Queue Performance (P50)

Speculate-all

500	11.21	10.05	9.44	9.19	9.04
400	11.82	10.69	9.80	9.75	9.42
300	13.08	11.87	11.00	10.74	10.58
200	15.30	14.04	13.14	12.90	12.72
100	7.41	6.63	6.46	6.44	6.24
	100	200	300	400	500

#Changes / Hour

#Workers

Optimistic speculation

500	8.54	8.72	8.62	8.57	8.77
400	8.75	8.70	8.67	8.74	8.69
300	7.33	7.63	7.64	7.56	7.65
200	9.60	9.62	9.62	9.64	9.64
100	7.46	7.46	7.44	7.44	7.44
	100	200	300	400	500

#Changes / Hour

#Workers

Optimistic speculation performs better than speculate-all esp. under contention.

# Evaluation: Submit Queue Performance (P50)

**Speculate-all**

500	11.21	10.05	9.44	9.19	9.04
400	11.82	10.69	9.80	9.75	9.42
300	13.08	11.87	11.00	10.74	10.58
200	15.30	14.04	13.14	12.90	12.72
100	7.41	6.63	6.46	6.44	6.24
	100	200	300	400	500

#Changes / Hour

#Workers

**Optimistic speculation**

500	8.54	8.72	8.62	8.57	8.77
400	8.75	8.70	8.67	8.74	8.69
300	7.33	7.63	7.64	7.56	7.65
200	9.60	9.62	9.62	9.64	9.64
100	7.46	7.46	7.44	7.44	7.44
	100	200	300	400	500

#Changes / Hour

#Workers

**Submit Queue**

500	2.56	1.77	1.49	1.38	1.26
400	2.57	1.87	1.59	1.47	1.42
300	2.52	1.87	1.44	1.31	1.28
200	2.98	2.04	1.92	1.72	1.54
100	1.83	1.00	1.02	1.00	1.00
	100	200	300	400	500

#Changes / Hour

#Workers

Submit Queue has the best performance among all the approaches.

# Evaluation: Submit Queue Performance (P50)

**Submit Queue**

500	2.56	1.77	1.49	1.38	1.26
400	2.57	1.87	1.59	1.47	1.42
300	2.52	1.87	1.44	1.31	1.28
200	2.98	2.04	1.92	1.72	1.54
100	1.83	1.00	1.02	1.00	1.00
	100	200	300	400	500

**#Workers**

Performance matches Oracle's performance under low contention

# Evaluation: Submit Queue Performance (P99)

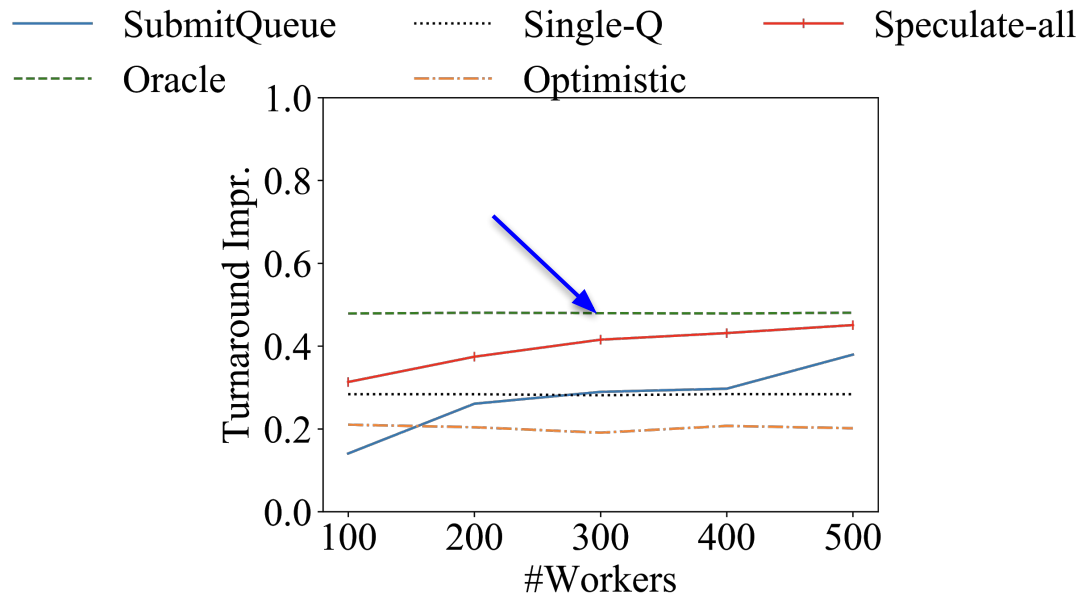
**Submit Queue**

500	2.87	1.84	1.52	1.39	1.21
400	3.25	2.00	1.62	1.51	1.35
300	3.55	2.45	1.83	1.54	1.53
200	3.95	2.57	1.95	1.63	1.51
100	2.19	1.65	1.46	1.38	1.25

P99 turnaround time is only 4x worse under *extreme contention*.  
We don't operate there typically in production.



# Evaluation: Impact of Conflict Analyzer



P95 Turnaround Time Impr. for 500 changes/hour

- Oracle's turnaround time improves by up to 50% with conflict analyzer.
- Benefit for SQ and Speculate-all steadily converges towards Oracle.

# Submit Queue guarantees always-green master

- **Probabilistic speculation** powered by *logistic regression* to select builds that are likely to succeed, and execute them in parallel
- **Conflict analyzer** to commit independent changes in parallel, and trim the speculation space.
- **Evaluated** Submit Queue in production deployment

# Thank you!

