### Combinatorial Constructions in Testing Concurrent Programs

#### **Rupak Majumdar**

Max Planck Institute for Software Systems (MPI-SWS), Germany

Joint Work with Filip Niksic and Dmitry Chistikov

### Despite Many Formal Approaches...

- ... practitioners test their code
- ... by providing random inputs

And despite our best judgment, ...

... testing is *surprisingly effective in finding bugs* 

In this talk, we explore this unexpected effectiveness

#### Example: Jepsen http://jepsen.io

 A framework for black-box testing of distributed systems by randomly inserting network partition faults



 CAP Theorem: No system has consistency, availability, and partition tolerance

etcd Postgres Redis Riak MongoDB Cassandra Chronos Kafka RabbitMQ Consul Elasticsearch Aerospike Zookeeper

### Tests and Coverage



**Coverage Goals** 

Tests

Covering Family of Tests = Set of tests covering all goals "Small" covering family = More efficient testing

### Random Testing

Pick tests at random, independently





**Coverage Goals** 

Tests

Question: Will random testing find a small covering family?

### Random Testing

Pick tests at random, independently



Suppose  $Pr[Test \bullet covers goal \bullet] \ge p$ Goal: Characterize covering families w.r.t. p and |G|

### Probabilistic Method

Let G be the set of goals and  $Pr[Random test covers a goal] \ge p$ 

**Theorem.** There exists a covering family of size  $p^{-1} \log|G|$ .

Proof.

 $Pr[Random test does not cover goal g] \le 1 - p$ 

 $\Pr[K \text{ ind tests do not cover goal } g] \le (1 - p)^{K}$ 

 $\Pr[K \text{ ind tests is not a covering family}] \le |G| (1 - p)^{K}$ 

For  $K = p^{-1} \log |G|$ , this probability is < 1

Then: There must exist K tests that is a covering family!

Existence, not constructive!

### Probabilistic Method

Let G be the set of goals and  $Pr[Random test covers a goal] \ge p$ 

**Theorem.** There exists a covering family of size  $p^{-1} \log|G|$ .

By repeated sampling, random testing overwhelmingly likely to find a covering family!

### Tests and Coverage



1. What are the coverage goals?

**Coverage Goals** 

2. What are tests? What is the Probability space?

3. Can we bound Pr[Test covers a goal]?



- Start with some combinatorial puzzles
- Connect these puzzles to testing
  - 1. Come up with coverage goals
  - 2. Show how the puzzles relate to coverage goals
  - 3. Bound the probabilities



Puzzle I: Ninjas Training



In a dojo in Kaiserslautern, **n** ninjas are in training:



Training is **complete** if for every pair of ninjas, there is a round where they are in opposing teams

Training is **complete** if for every pair of ninjas there is a round where they are in opposing teams

How many rounds make the training complete?

Naïve: O(n<sup>2</sup>) rounds

Can you do it in log n rounds?

Now **n** ninjas are practicing in **k**-way fights:



Training is **complete** if for every choice of **k** ninjas there is a round where they are each in a different team

How many rounds make the training complete?









Training is **complete** if for every choice of **k** ninjas there is a round where they are each in a different team

How many rounds make the training complete?

Naïve: O(n<sup>k</sup>)

Can you do it in  $O(k^{k+1} (k!)^{-1} \log n)$  rounds?

Exponentially better in n, when k is constant

Puzzle II: Ninjas Eating

### Hungry Ninjas

After training, the ninjas retire to a bucolic Biergarten...



A banquet is **complete** if for every pair of ninjas (**i**, **j**), there's a course that is served to **i** before **j** and one that is served to **j** before **i** 

How many courses make a banquet complete?

### A Complete Banquet

**Two** courses suffice:



### **3-Complete Banquets**



#### A banquet is **3-complete** if for every **triple** of ninjas (**i**, **j**, **k**), there's a course served in the order **i<j<k**

How many courses make a banquet 3-complete?

There is a 3-complete banquet with n<sup>3</sup> courses Can you do it in O(exp(d) log n) rounds?

### Masters at the Banquet

#### Ninjas, of course, form a hierarchy A **master** is always served **before** their **student**



### Masters at the Banquet

Again, two courses suffice for 2-completeness:



### Ninjas at the Banquet

admissible

A banquet is **3-complete** if for every triplet (**i**, **j**, **k**), there's a course served to ninja **i** before **j**, and **j** before **k** 

Naive approach with **n**<sup>3</sup> courses:

Pick a course for each  ${}^{n}C_{3}$ . 3! orders

Can be done with **O(log n)** courses!

From Ninjas to Testing...

#### From Training Ninjas to Distributed Systems with Partition Faults





ninjas weapons rounds complete training nodes in a network blocks in a partition tests = partitions **splitting family of partitions**  Coverage Goal: Splitting Families

Given n nodes and  $k \leq n$ ,

a partition of nodes  $P = \{B_1, ..., B_k\}$  splits a set of nodes  $S = \{x_1, ..., x_k\}$  if  $x_1 \in B_1, ..., x_k \in B_k$ .

A set of partitions F is a **k-splitting family** if for every k-subset of nodes there is a partition in F that splits it.

### From Eating Ninjas to Testing Concurrent Systems





ninjas hierarchy courses d-complete banquet events partial order on events tests = schedules/linearization **d-hitting family of schedules** 

### Coverage Goal: Hitting Families

Given a poset of events, a schedule **hits** a d-tuple of events  $(e_1,...,e_d)$  if it executes the events in the order  $e_1 < ... < e_d$ .

Given a poset of events, a **family of schedules F** is **d-hitting** if for every admissible d-tuple of events there is a schedule in **F** that hits it.

#### From Jepsen https://jepsen.io Why k-Splitting?

- Chronos: A distributed fault-tolerant job scheduler
- Works in conjunction with Mesos and Zookeeper
- Three special nodes: Chronos leader, Mesos leader, Zookeeper leader







# Why d-Hitting?

Many bugs in asynchronous programs involve small number of events —**bug depth d** 

[Lu et al. ASPLOS '08] [Burckhardt et al. ASPLOS '10] [Jensen et al. OOPSLA '15] [Petrov et al. 2012]



Combinatorics: Bounding the Probabilities



- Start with some combinatorial puzzles
- Connect these puzzles to testing
  - 1. Come up with coverage goals
  - 2. Show how the puzzles relate to coverage goals
  - 3. Bound the probabilities

### Small k-Splitting Families

- Fix a k-element set S
- What is the probability a random partition splits S?
- A splitting partition uniquely corresponds to a map  $U \setminus S \rightarrow S$ . There are  $k^{(n-k)}$  such maps
- So probability =  $k^{(n-k)} / \{n P k\}$

Stirling number of 2<sup>nd</sup> kind = Number of partitions of n elements into k parts

### Small k-Splitting Families

Probability that a fixed set is split = k<sup>(n-k)</sup> / {n P k}

 $\geq k^{-k} \cdot k!$ 

- Can we get rid of n?
- Yes: k<sup>n</sup> ≥ k! {n P k}



### Small k-Splitting Families

- Pr[Random partition splits S]  $\ge k!/k^k$
- From our general theorem: There is a k-splitting family of size (k<sup>k</sup> /k!) k log n
- Turns out: uniformly sampling k-partitions is hard
  - But sampling balanced partitions is sufficient
  - The combinatorial arguments are harder

$$k^n \binom{n}{k} \le n^k \binom{n}{k}$$

Problem 11957 AMM 2017



Probabilistic argument: Fix a tuple of d ninjas

What is the probability that a random schedule covers it? 1/d! What is the probability k schedules don't cover it? (1 - 1/d!)<sup>k</sup>

What is the probability k schedules are not a hitting family?

At most (n C 3) . (1 - 1/d!)<sup>k</sup>

Pick  $k > d! d \log n$ 



- Start with some combinatorial puzzles
- Connect these puzzles to testing
  - 1. Come up with coverage goals
  - 2. Show how the puzzles relate to coverage goals
  - 3. Bound the probabilities

### Combinatorics I: Splitting Families & Perfect Hashing

- Andrew Yao. Should Tables Be Sorted? 1981
- Uses k-splitting families to construct "perfect hash functions": for every k-subset S of an n-element domain there is a hash function that is 1-to-1 on S
- Uses perfect hash functions to construct hash tables with constant lookup
- Follow-up work by Fredman, Komlós, Szemerédi

Combinatorics II: Hitting Families & Order Dimension of Posets

- Dushnik and Miller. Partially Ordered Sets. 1941
- Define and study order dimension of a poset:
  Number of linearizations whose intersection is the poset
- Hitting families generalize order dimension
  order dimension = size of the smallest 2-hitting family
- d-completeness: Not studied in the p.o. literature!

### **Explicit Constructions**

Probabilistic method shows existence:

- k-splitting families of size: (k<sup>k+1</sup> / k!) log n
- d-hitting families of size: d d! log n

But may not be optimal:

- 2-splitting family of size log n exists; prob. method gives 2 log n
- d-hitting families for trees of size h of size O(exp(d)-h<sup>d-1</sup>)

OTOH, matching explicit constructions are open:

k-splitting families of size: 4<sup>sqr(k)</sup> (log<sub>2</sub> n)<sup>k-1</sup> by Yao







1. Splitting nodes in network partitions

**Coverage Goals** 

Tests

3. Pr[Test covers a goal]  $\ge \exp(k) \log n$ 



#### **Coverage Goals**

Tests

3. Pr[Test covers a goal]  $\ge \exp(d) \log n$ 





#### Watch out for Filip Niksic's PhD Thesis!

Icons made by Freepik at www.flaticon.com