Cascading Failures, Rare Events, and Heavy Tails

Bert Zwart

CWI and TU/e

March 5, 2019

Joint work with Tommaso Nesti, Fiona Sloothaak, Sem Borst

CWI
Centrum Wiskunde & Informatica

Blackouts in power grids

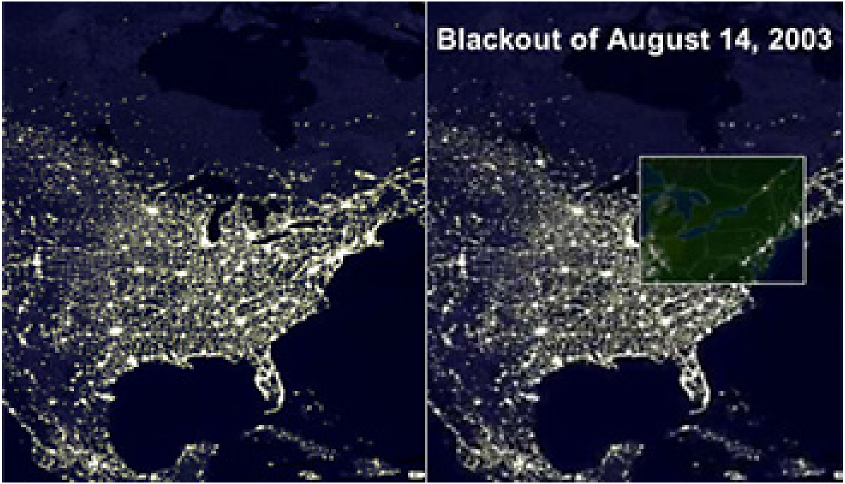

Blackouts in the past fifty years

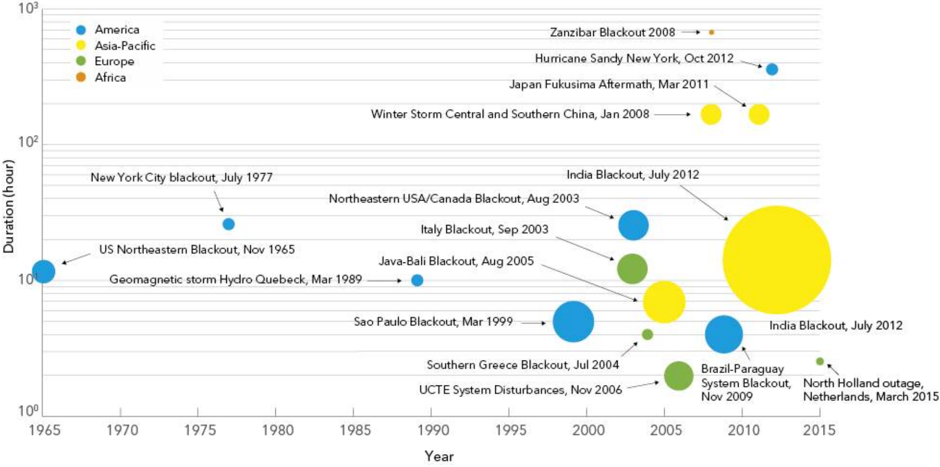

(source: dnv-gl)

“Near-Blackouts” in German HV grid

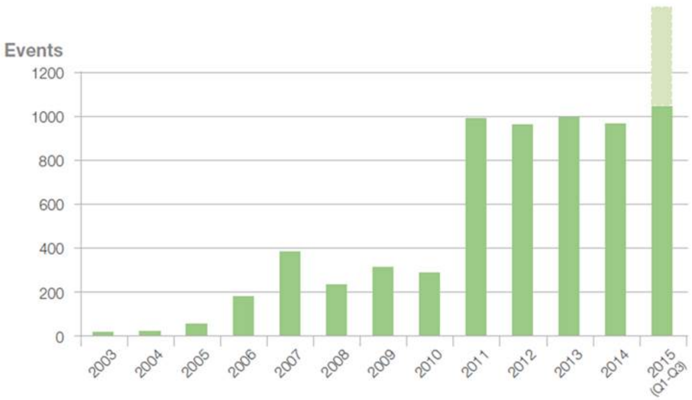

Figure 4: Incidents with counter-measures in the control area of TenneT Germany
(without voltage/reactive power problems; numbers for 2015 including first to third quarter Q1-Q3)

Violations of $N-1$ safety criterion.

Unrest in South Australia (2016 - 2017)

- Rolling blackouts during heat wave
- Renewable energy (wrongfully?) blamed
- Problems mitigated by 100 MW Tesla battery

Successful squirrel attacks in 2016

(NL: no squirrels, one pigeon in Tilburg, one marten in Aalten)

1/23/2017 CyberSquirrel1.com

Cyber Squirrel 1

Disrupting at the highest levels, its #CyberWar4Ever!

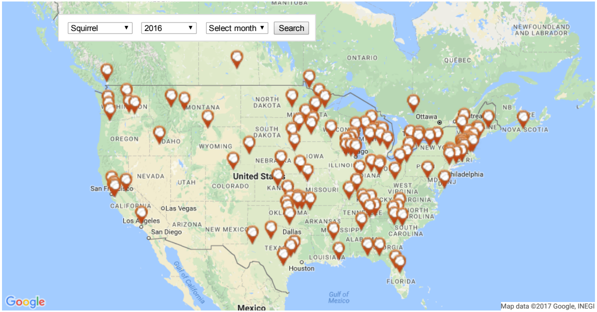

TOTAL SUCCESSFUL CYBER WAR OPS AS OF 2017.01.08 - 1748

Agent	Success
Squirrel	879
Bird	434
Snake	83
Raccoon	72

ABOUT THIS MAP

This map lists all unclassified Cyber Squirrel Operations that have been released to the public that we have been able to confirm. There are many more executed ops than displayed on this map however, those ops remain classified.

Confirmation for all ops has been preserved by the Internet Archive's WayBack Machine whenever possible.

MOST RECENT UNCLASSIFIED OPS

Tweets by @CyberSquirrel1

CyberSquirrel
@CyberSquirrel1

We can neither confirm nor deny our involvement at this time.

3h

Remarks

- It took 2 years to figure out the cause of the 1996 blackout
- "It is not complex, but complicated"
- "It is not possible to come up with a both interesting and useful result"
- Still engineers seem to care about understanding and preventing blackouts:
  - Tutorial (IEEE PES GM tutorial on how industry deals with planning against blackouts)
  - Very recent survey of IEEE PES WG on cascading failure models
- At least one feature of blackouts is not complicated

Pareto laws in power grids (Hines 09)

$$\mathbb{P}(\boldsymbol{Blackout\ size} \geq \boldsymbol{x}) = \boldsymbol{c}\, \boldsymbol{x}^{-\boldsymbol{\alpha}}$$

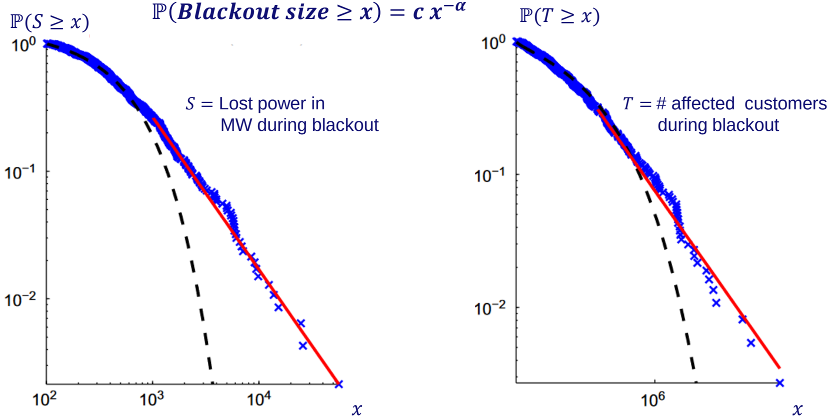

WHY?

Questions and goals

- Two main stylized facts of cascading failures in power grids:
  - propagation is not of nearest neighbor type
  - total size of blackout is heavy-tailed
- Can we explain both using mathematical modeling and probability theory?
- Will number of blackouts increase with percentage of renewable energy?
- We want to connect our model/analysis to structural microscopic models used by engineers
- We will combine a DC power flow model with **rare event analysis**

Rare Events depend on “Tail Behaviors”

Light-Tailed Distributions

- Extreme Values are Very Rare
- Normal, Exponential, etc

Heavy-Tailed Distributions

- Extreme Values are Frequent
- **Pareto Law**, **Weibull**, etc

Rare Events depend on "Tail Behaviors"

Light-Tailed Distributions

- Extreme Values are Very Rare
- Normal, Exponential, etc

Heavy-Tailed Distributions

- Extreme Values are Frequent
- **Pareto Law**, **Weibull**, etc

Heavy tails are not as well understood as light tails.

Rare Events depend on "Tail Behaviors"

Light-Tailed Distributions

- Extreme Values are Very Rare
- Normal, Exponential, etc

Systemwide rare events

arise because

EVERYTHING goes wrong.

(Conspiracy Principle)

Heavy-Tailed Distributions

- Extreme Values are Frequent
- **Pareto Law**, **Weibull**, etc

Heavy tails are not as well understood as light tails.

Rare Events depend on "Tail Behaviors"

Light-Tailed Distributions

- Extreme Values are Very Rare
- Normal, Exponential, etc

Systemwide rare events

arise because

EVERYTHING goes wrong.

(Conspiracy Principle)

Heavy-Tailed Distributions

- Extreme Values are Frequent
- **Pareto Law**, **Weibull**, etc

Systemwide rare events

arise because of

A FEW Catastrophes.

(Catastrophe Principle)

Heavy tails are not as well understood as light tails.

Example of heavy tails

As $x \to \infty$:

- Pareto tails (or power tails): $P(X > x) \approx x^{-\alpha} = e^{-\alpha \log x}$
- Lognormal tails: $P(X > x) \approx e^{-\alpha (\log x)^2}$
- Weibull tails: $P(X > x) \approx e^{-\alpha x^{\beta}}$, $\beta \in (0,1)$.

Key properties:

$$E[e^{\varepsilon X}] = \infty, \qquad \varepsilon > 0.$$

$$P(X_1 + \ldots + X_n > x) \sim P(\max_{i=1,\ldots,n} X_i > x).$$

Heavy Tails are Everywhere:

Computer systems

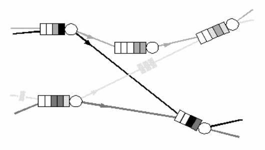

delays, files, …

Finance

losses

Social networks

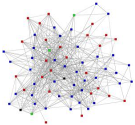

popularity, contagion

Energy Systems

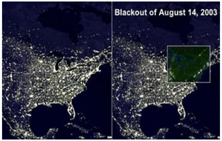

blackouts

How do heavy tails occur?

Heavy tails and queues

How does a large queue length occur?

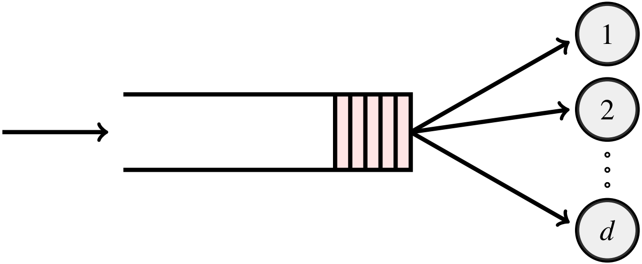

Answer: several large job sizes

Heavy tails and queues

How does a large queue length occur?

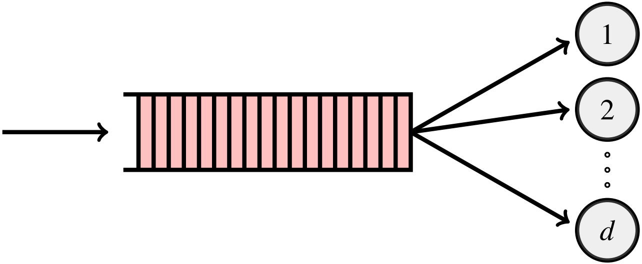

Answer: several large job sizes

Heavy tails and queues

How does a large queue length occur?

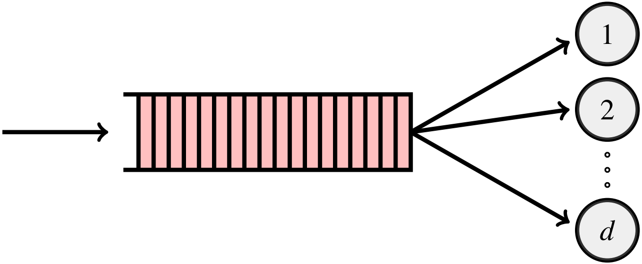

Answer: several large job sizes

Foss & Korshunov (2005, 2012), Bazhba, Blanchet, Rhee, Z (2018)

Heavy tails and queues

How does a large queue length occur?

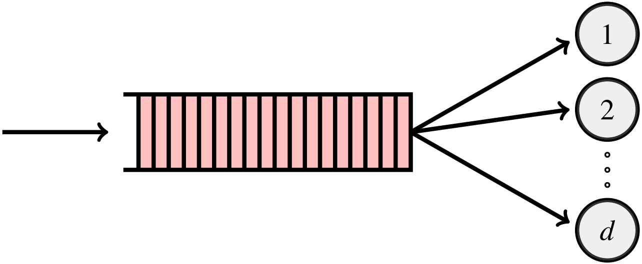

Answer: several large job sizes

Foss & Korshunov (2005, 2012), Bazhba, Blanchet, Rhee, Z (2018)

Take away: heavy tailed output caused by heavy-tailed input

Heavy tails and math finance

Kesten, Harry. Random difference equations and Renewal theory for products of random matrices. Acta Mathematica **131** (1973), 207–248.

In one dimension: if

$$X \stackrel{d}{=} AX + B,$$

then (under some conditions)

$$P(X > x) \sim Cx^{-\kappa},$$

with $\kappa > 0$ solving the equation $E[A^{\kappa}] = 1$.

Take away: heavy tails occur by multiplication (nonlinearity)

Heavy tails and networks

A.-L. Barabási, R. Albert. Emergence of Scaling in Random Networks. Science 15 Oct 1999: Vol. 286, Issue 5439, pp. 509-512.

6.3: AN EXAMPLE: PREFERENTIAL ATTACHMENT AND YULE PROCESSES 119

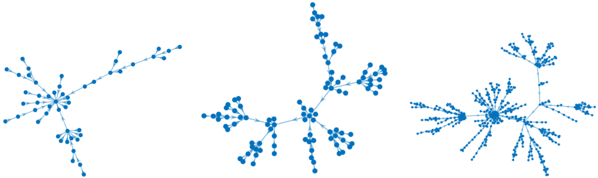

(a) 50 nodes (b) 100 nodes (c) 500 nodes

Figure 6.4: *Illustrations of graphs generated by the preferential attachment model. For clarity, self-edges are not shown. Note that there are a few 'core' high-degree nodes, surrounded by a periphery of many low-degree nodes.*

Take away: power law degrees created by preferential attachment.

Known since Yule (1924)

Heavy tails and critical phenomena

Consider a Branching Process with $Z_0 = 1$ and

$$Z_{n+1} = \sum_{i=1}^{Z_n} C_{ni}.$$

If $\mathbf{E}[C_{ni}] = 1$ the branching process is said to be critical. The total size and depth of the tree are heavy-tailed.

Self-organized criticality: many natural and man-made systems appear to behave like critical systems

Take-away: heavy tails occur if system operates in critical regime

Summary: heavy tails can occur in many ways

- Exogenous factors (e.g. job sizes in queueing)
- Multiplication (e.g. gains or losses in finance)
- Preferential attachment (social, and other networks)
- (Self-organized) criticality

Existing work on blackouts, based on model simulation output data, show long-range correlations in outages, and attributes this to criticality. Earlier work suggests the usage of critical branching processes.

A different explanation

Let C be the size of a city, in terms of number of people, and let T be the size of a blackout, in terms of number of customers affected
Both have statistically significant, almost identical power law for US:

$$P(C > x) \approx x^{-1.37} \qquad\qquad P(T > x) \approx x^{-1.31}.$$

German city sizes: power law with index 1.28

2015 rank	City	State	2015 Estimate	2011 Census	Change	2015 land area	2015 po
1	***Berlin***	Berlin	3,520,031	3,292,365	+6.91%	891.68 km^2 344.28 sq mi	
2	***Hamburg***	Hamburg	1,787,408	1,706,696	+4.73%	755.3 km^2 291.6 sq mi	
3	***Munich*** (München)	Bavaria	1,450,381	1,348,335	+7.57%	310.7 km^2 120.0 sq mi	
4	*Cologne* (Köln)	North Rhine-Westphalia	1,060,582	1,005,775	+5.45%	405.02 km^2 156.38 sq mi	
5	*Frankfurt am Main*	Hesse	732,688	667,925	+9.70%	248.31 km^2 95.87 sq mi	

log-log plots and Hill plots

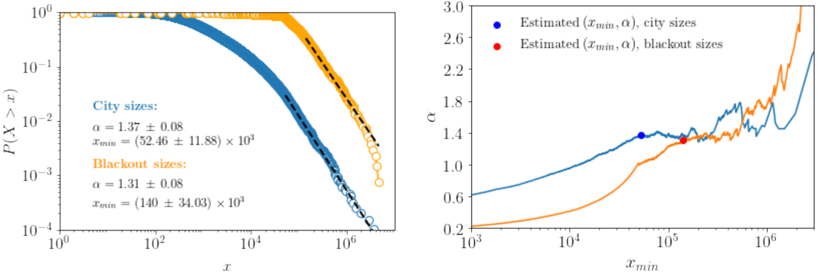

US city size data (2000 census) and US outage data (NERC, 2002-2018). Estimate is done according to the PLFIT method of Clauset et. al (2009). Standard deviation is 0.08 for both estimates.

Mathematical model

To explain the reason why the tail behavior of city sizes and blackout sizes is similar, we develop a mathematical model.

- We will combine power systems models with extreme value theory
- Our model is a graph with multiple heavy-tailed sinks. We use the DC load flow model. Network topology and reactances are all encoded in the load-flow matrix $\mathbf{V}$
- We consider three stages in our model:
  - Planning
  - Operation
  - Emergency

Model: operational stage (DC-OPF)

Production in each node is computed by solving

$$\min \frac{1}{2}\sum_{i=1}^{n} g_i^2$$

$$\sum_i g_i = \sum_i X_i$$

$$-\bar{\mathbf{f}} \leq \mathbf{V}(\mathbf{g}-\mathbf{X}) \leq \bar{\mathbf{f}}.$$

This determines the network flows $\mathbf{F} = \mathbf{V}(\mathbf{g}-\mathbf{X})$ which play a role in the cascade.

We determine the line limits $\bar{\mathbf{f}}$ in a planning problem.

Model: planning stage

Given n cities with random sizes $X_1, \ldots, X_n$ and given a network topology, we determine line limits $\bar{\mathbf{f}}$ by solving an unconstrained OPF problem:

$$\min \frac{1}{2} \sum_{i=1}^{n} g_i^2$$

subject to the balance constraint

$$\sum_i g_i = \sum_i X_i.$$

The planning problem has solution $g_i = \bar{X}_n$ for $i = 1, \ldots, n$, with $\bar{X}_n = (1/n) \sum_{i=1}^{n} X_i$ the average city size. We now let $\lambda \in (0,1)$ and set

$$\bar{\mathbf{f}} = \lambda \mathbf{V}(\bar{X}_n \mathbf{e} - \mathbf{X}) = -\lambda \mathbf{V}\mathbf{X},$$

This vector will be used in the operational stage

Model: operational stage (DC-OPF)

$$\min \frac{1}{2}\sum_{i=1}^{n} g_i^2$$

$$\sum_i g_i = \sum_i X_i$$

$$-\bar{\mathbf{f}} \leq \mathbf{V}(\mathbf{g}-\mathbf{X}) \leq \bar{\mathbf{f}}$$

with

$$\bar{\mathbf{f}} = \lambda \mathbf{V}(\bar{X}_n \mathbf{e} - \mathbf{X}) = -\lambda \mathbf{V}\mathbf{X}.$$

This leads to actual line flows $\mathbf{F} = \mathbf{V}(\mathbf{g}-\mathbf{X})$.

Model: emergency stage

Given: line flows $\mathbf{F} = \mathbf{V}(\mathbf{g} - \mathbf{X})$

- Start with one random line outage.
- Recompute power flows.
- Additional lines fail if new line flow exceeds $\lambda^* \bar{f}_i$ for some $\lambda^* > 1$ (e.g. $\lambda_* = 1/\lambda$).
- If islands occur, load or generation is shed proportionally.
- T: size of total load shed once cascade is over.

Main result

Let X be a generic city size, with $P(X > x) \sim C_X x^{-\alpha}$.
Note that T is the blackout size [in terms of number of customers affected]

$$P(T > x) \sim C_T x^{-\alpha}, \qquad x \to \infty, \tag{1}$$

$$C_T = C_X n \sum_{j=1}^{n} P(|A_1| = j)(1 - j\lambda/n)^{\alpha}. \tag{2}$$

A_1 denotes the (random) set of nodes making up the island with the largest city in the network.

Proof idea: heavy-tailed large deviations theory allows us to consider the case of a single big city, and many small cities, reducing the analysis of the cascade to a single-sink network.

Example - Single sink with scaled demand 1

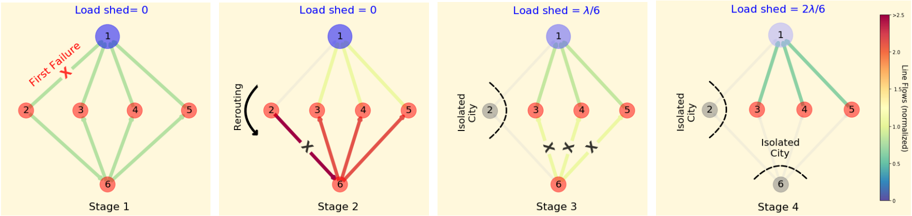

The nominal flows are $\lambda/24$ for each the four lower lines and $5\lambda/24$ for each of the four upper lines, with corresponding line limits of $1/24$ and $5/24$, respectively.

Phase transion at $\lambda = 3/4$.

SciGRID case study - Impact of λ

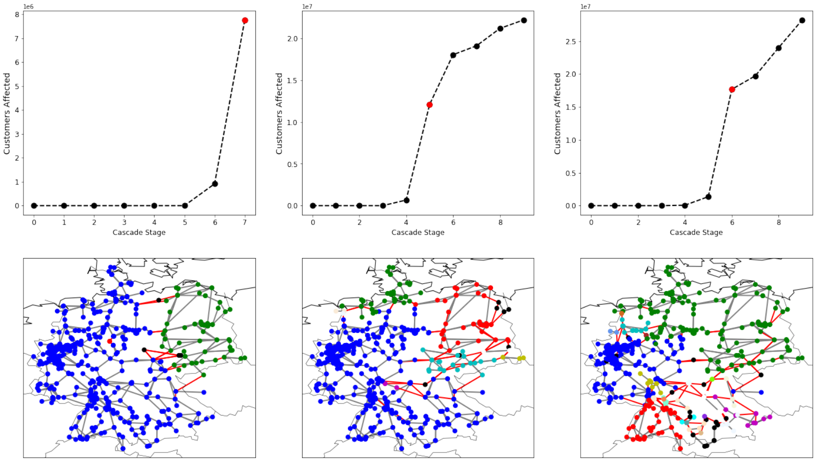

Figure: Dissection of biggest blackout for loading factors $\lambda = 0.7$ (left panels), $\lambda = 0.8$ (middle) and $\lambda = 0.9$ (right) in terms of the cumulative number of affected customers at each consecutive stage as displayed in the top charts with the biggest jump colored red.

Number of load shedding events during cascade

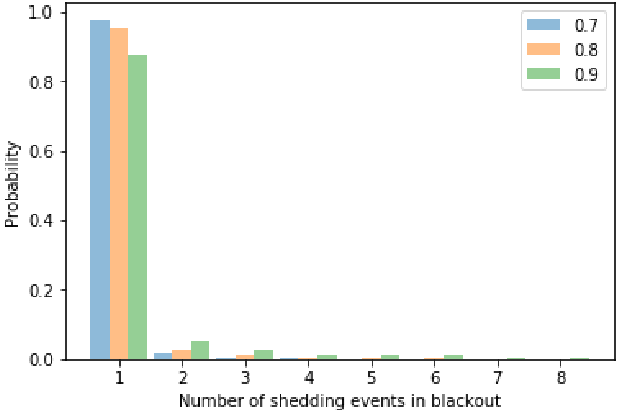

Figure: Histogram of the total number of load shedding events in the SciGRID network. For a moderate loading factor $\lambda = 0.7$, nearly 99% of the blackouts only involved a single jump. Even for a high loading factor $\lambda = 0.9$, 87% of the blackouts involve just a single jump. The fraction of blackouts with four or more jumps remains below 5%

Concluding remarks

- Rare event analysis is an area where proving a theorem could be easier than running an insightful simulation.
- Main benefit of rare event analysis is to determine the typical way a rare event happens, given that it occurs
- Main goal of our work: description why power laws occur in blackouts using rare event analysis.
- Insight: main source of variability are city sizes (nodal demands).
- Will not change with more renewables, though it could change with massive offshore windparks.

Book in preparation

The fundamentals of heavy tails

properties - emergence - identification

J. Nair, A. Wierman, B. Zwart

To receive sample chapters register at

http://users.cms.caltech.edu/~adamw/heavytails.html