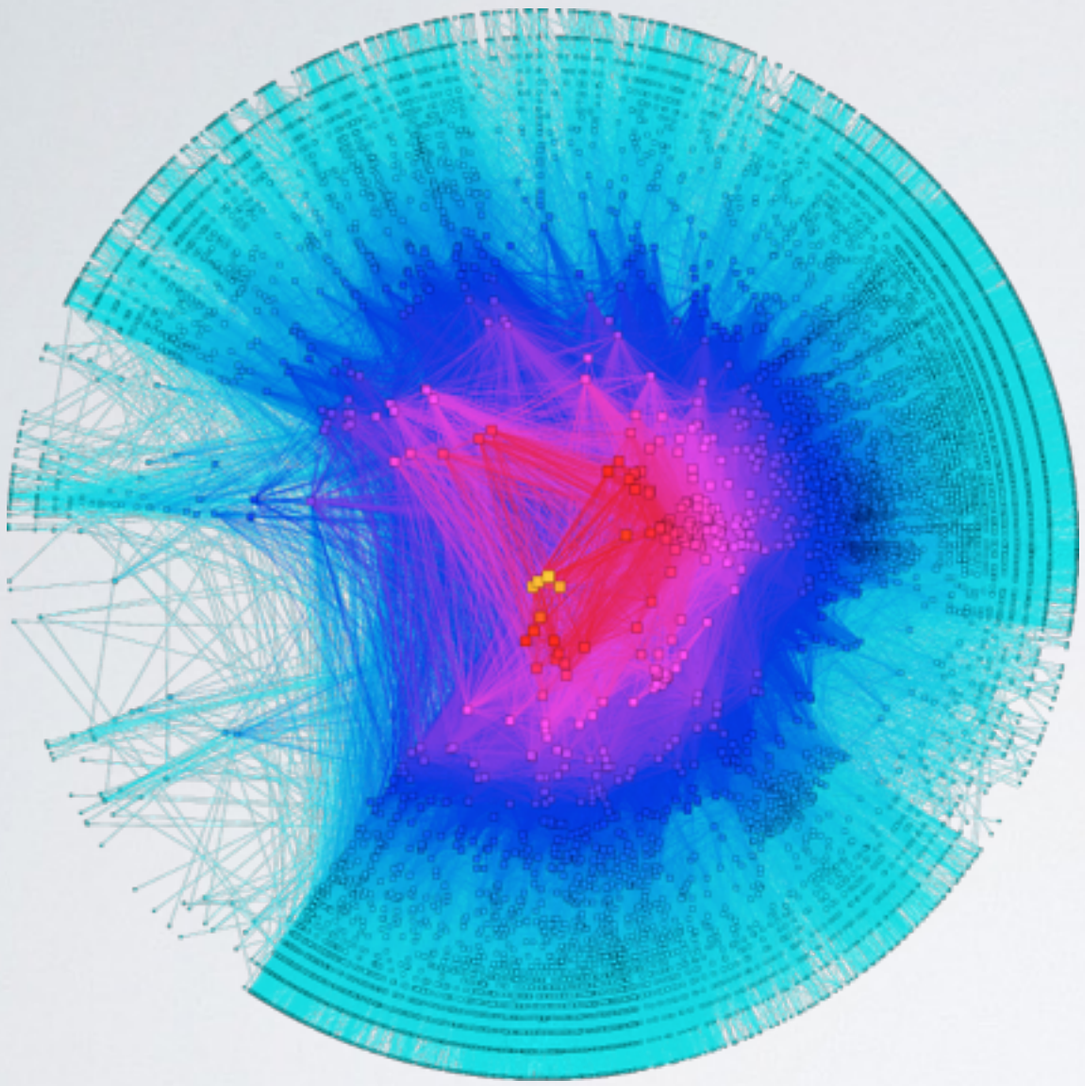

IMPROVING CAIDA'S AS LEVEL TOPOLOGY

Bradley Huffaker
CAIDA, UC San Diego

AIMS 2015, March 31st 2015

GOAL

Evaluate how the **AS** level **topology** is affected by **adding peering links** observed via Ark (**traceroutes**) and Internet eXchanges **(route servers)**.

DATA

- **traceroute**
  - inferred peering links observed one hop from an ark monitor
  - collected at 106 ark monitors
- **Internet eXchange (IXs)**
  - inferred from IXs' route server data and BGP communities
  - collected at 13 large European IXs

COMBINED GRAPH

graphs	links			nodes
	c/p	peer	total	total
BGP-derived	93,539	83,852	177,391	46,177
traceroute-derived	0	7,166	7,166	2,432
IX-derived	0	264,803	264,803	1,555
combined	93,539	325,312	418,851	46,320

- some overlap between traceroute and IX links
  - number of **ASes** increased **0.3%** (over BGP-derived)
  - number of **AS links** increased **136%**

COMPARISON METRICS

- **distance:** shortest (policy-free) distance between AS pair
  - **eccentricity:** longest shortest path
- **betweenness:** number of shortest path crossing an AS
- **degree:** number of ASes directly connected to an AS
- **coreness:** minimum degree value needed to recursively strip the AS from the graph
- **clustering:** fraction of possible triangles formed between an AS and its neighbors
- **customer/peer cone:** number of ASes recursively reachable from an AS along customer/peer links

SHORTEST DISTANCE

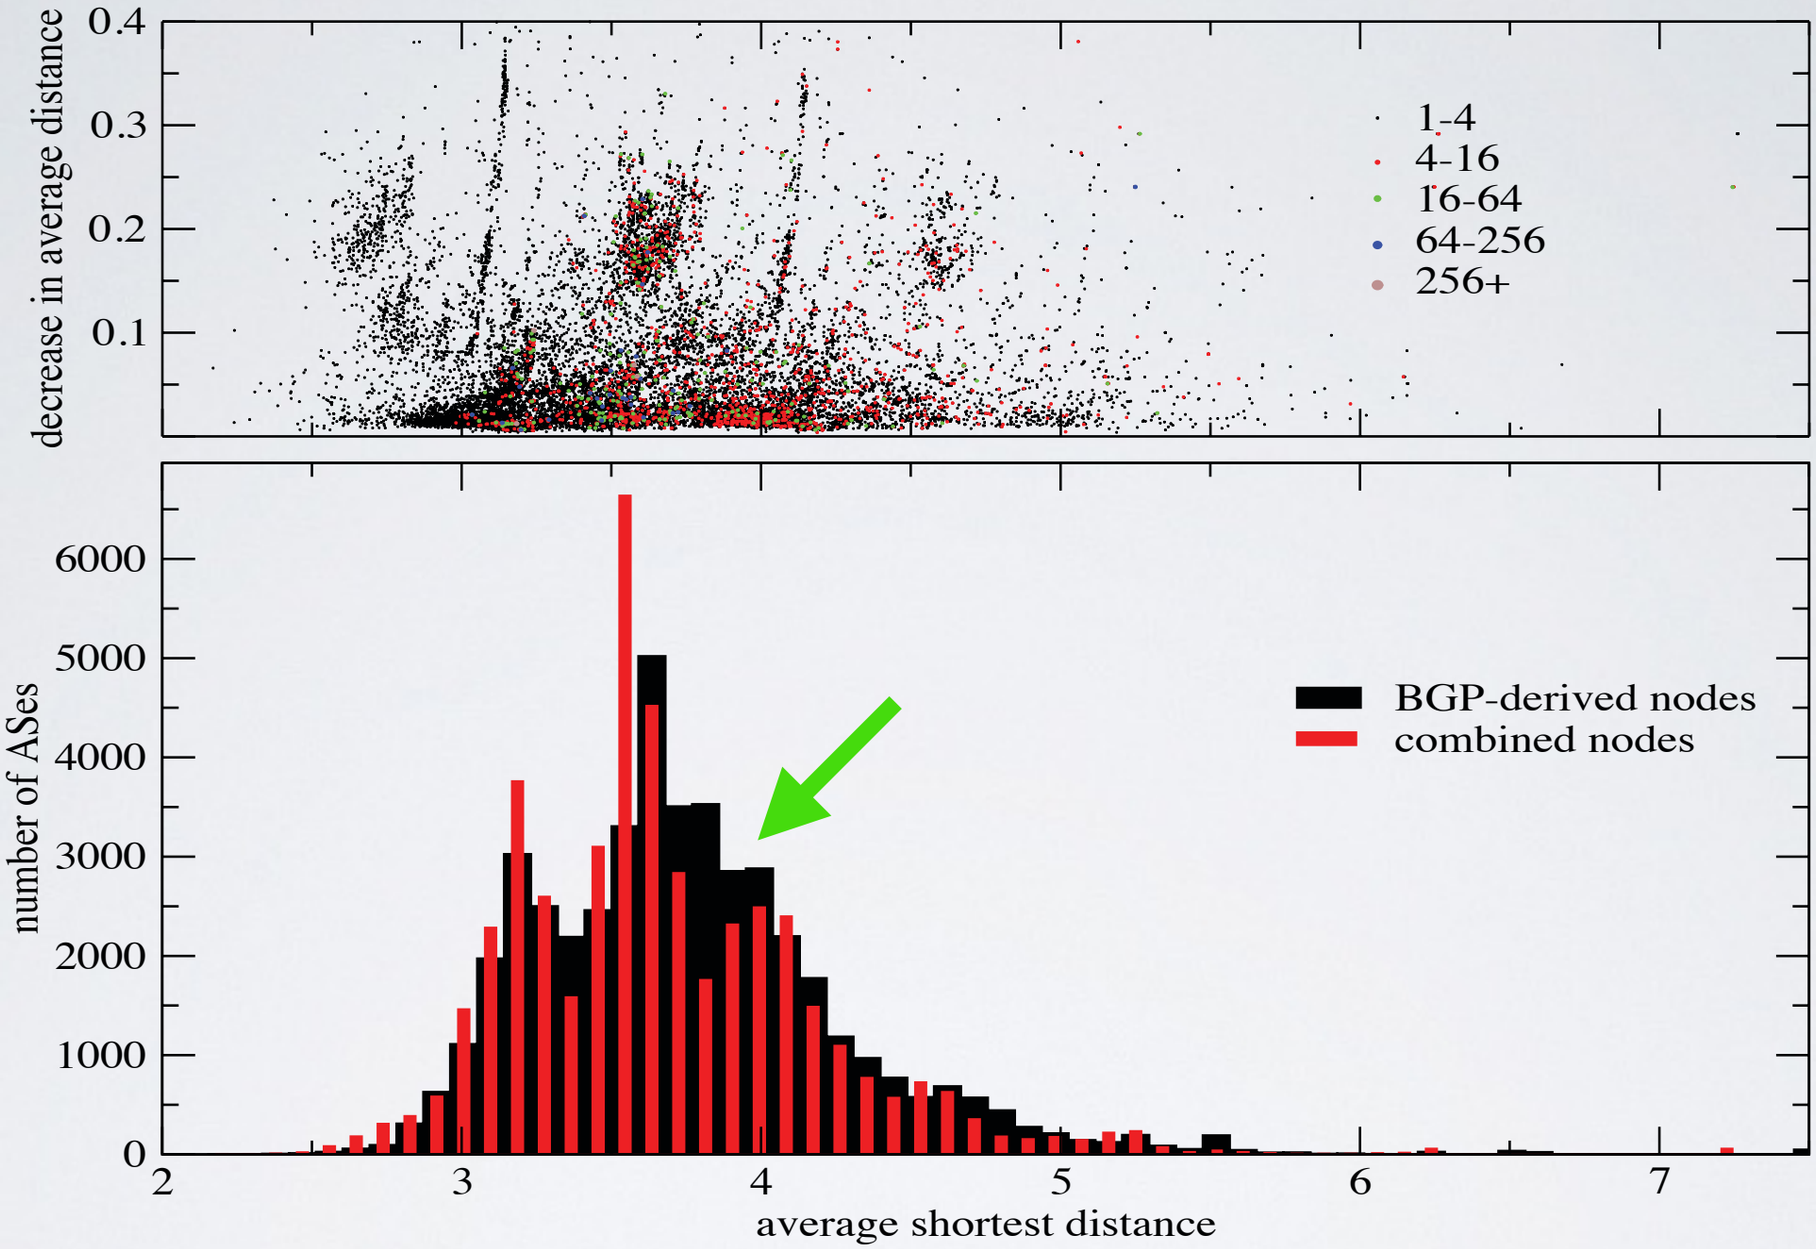

- **average** distance dropped from **3.75 to 3.68**
- switched from **bipolar to tripolar** distribution

SHORTEST DISTANCE

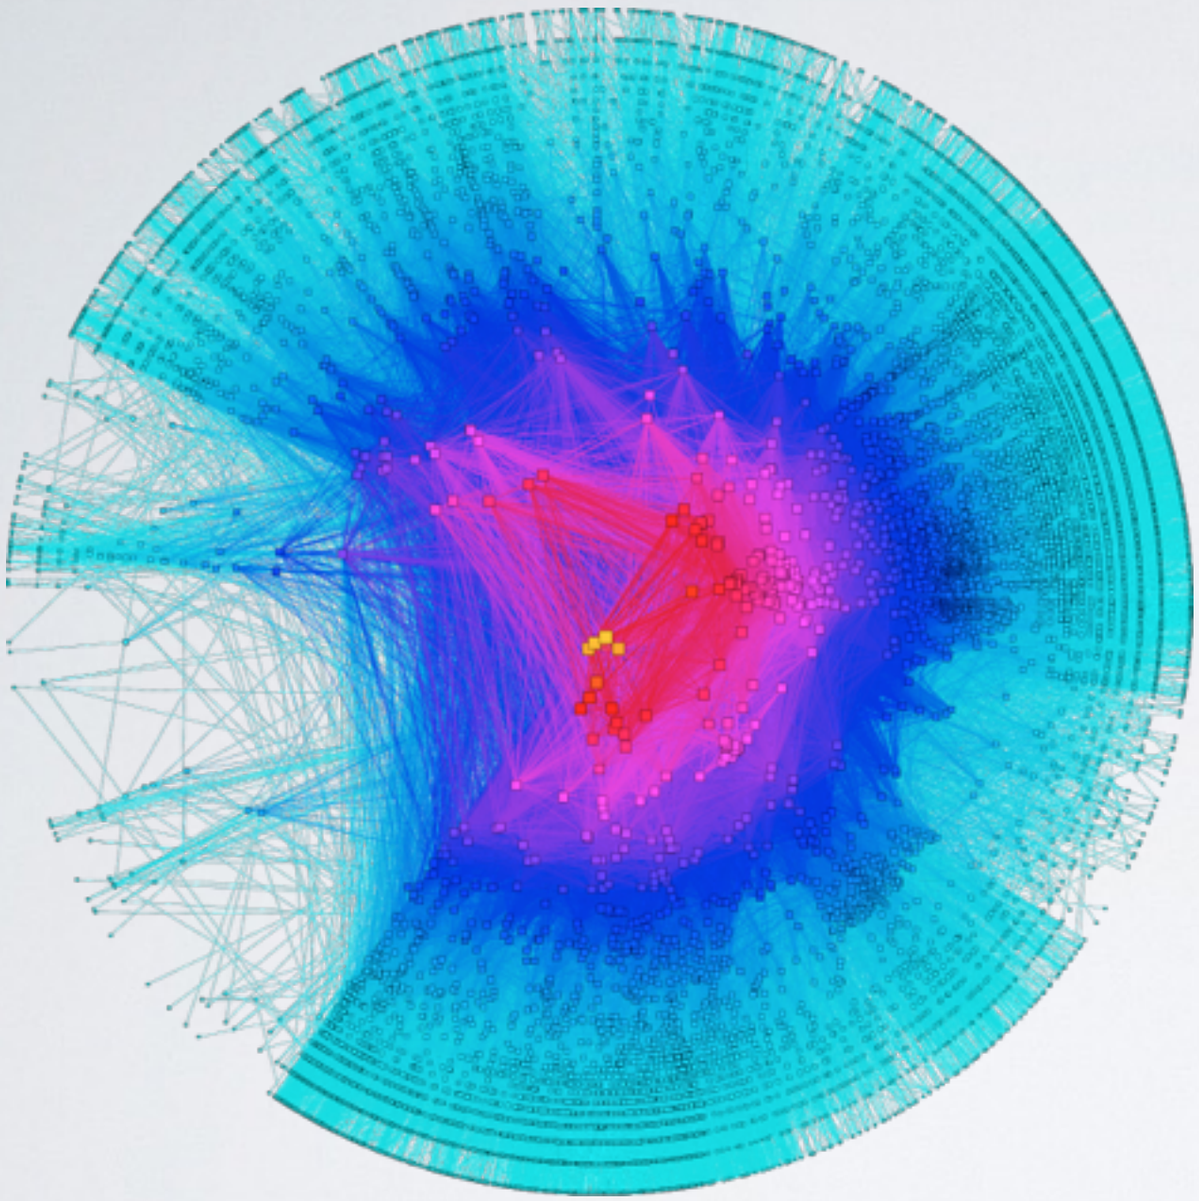

BGP-derived

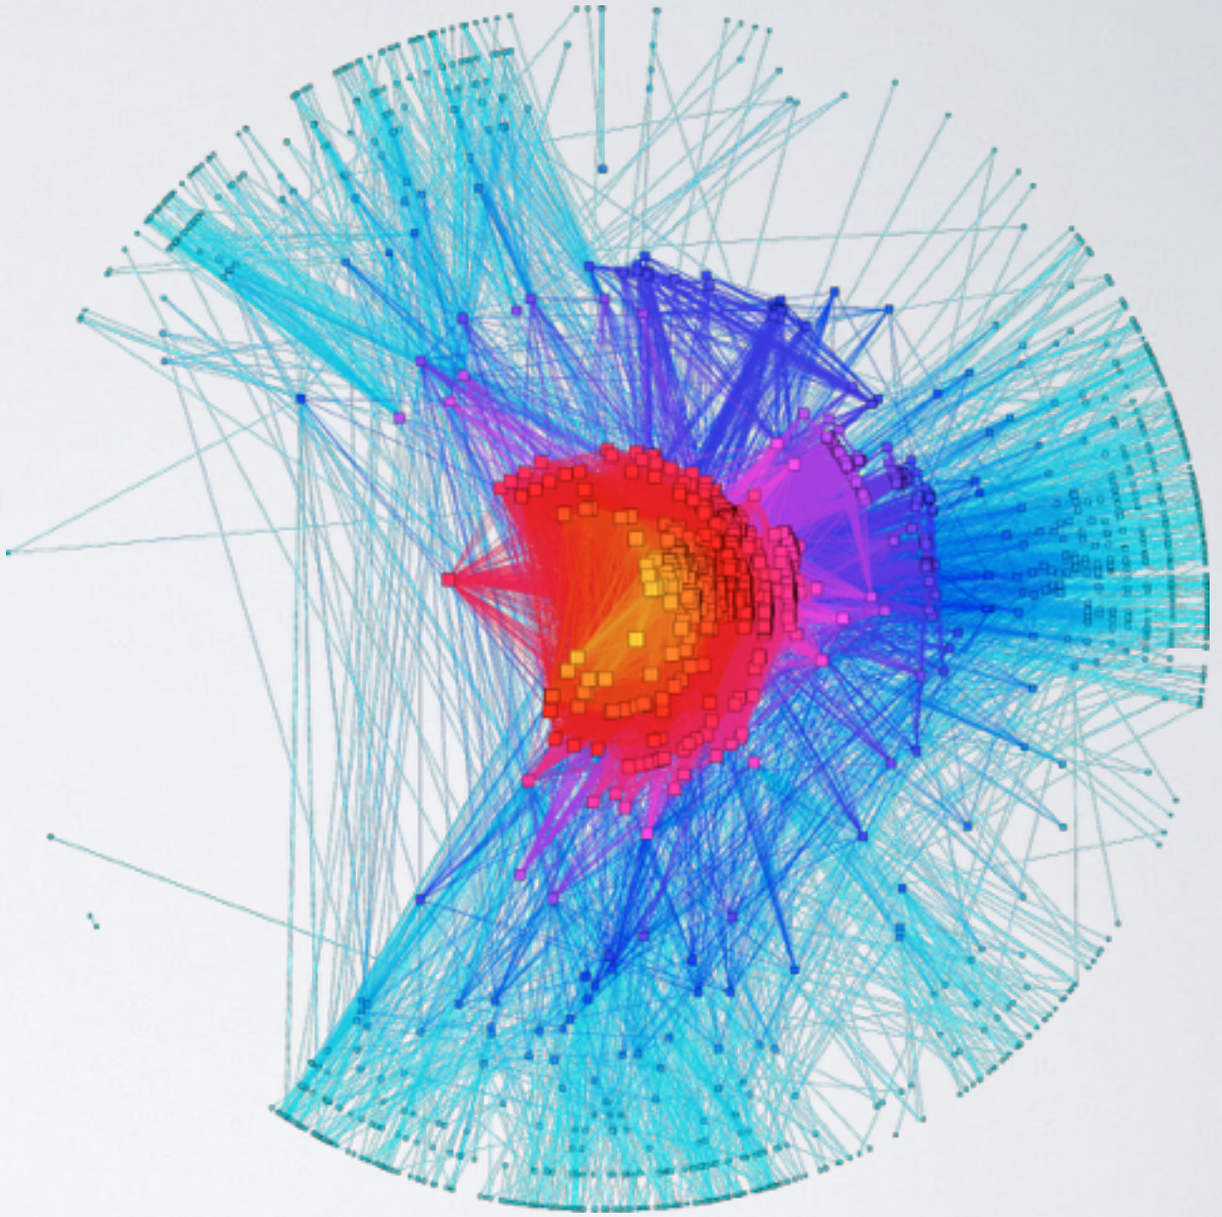

new links and ASes with at least one new link

$$\textbf{radius} = 1 - \log\left(\frac{\text{degree(AS)}+1}{\text{maxium.degree}+1}\right)$$

$$\textbf{angle} = (\text{ longitude of the AS's BGP prefixes })$$

caida

ECCENTRICITY

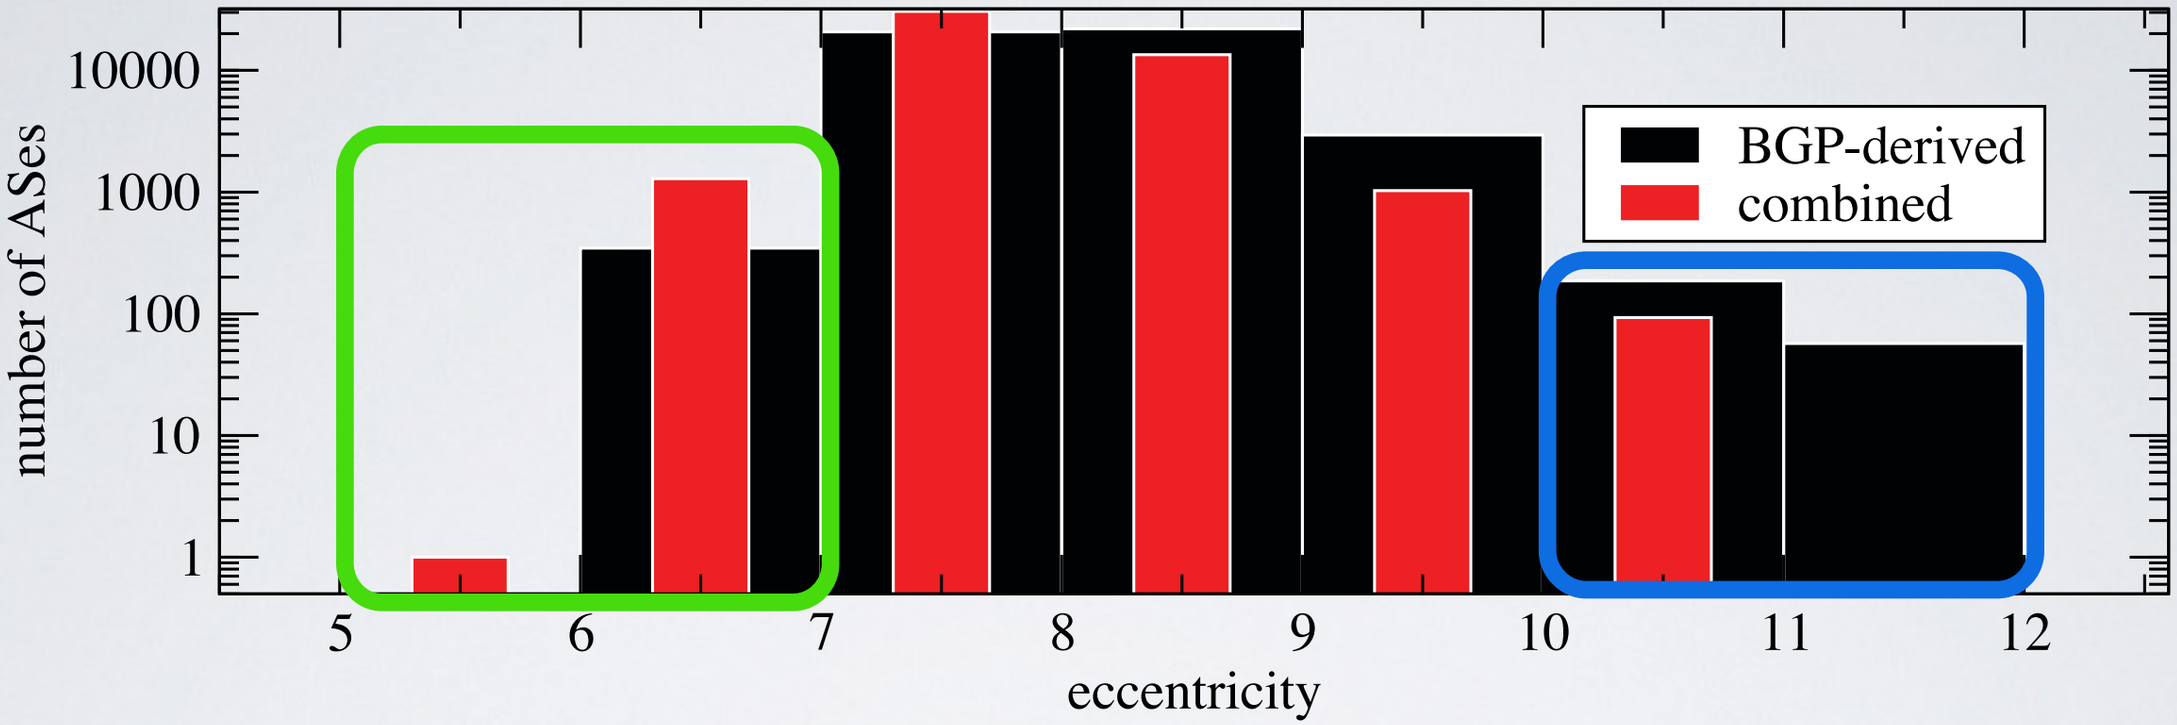

- **average** eccentricity dropped from **7.61 to 7.31**
- graph **diameter**, longest, dropped from **11 to 10**
- graph **radius**, shortest, dropped from **6 to 5**
  - but only for a single AS

BETWEENNESS

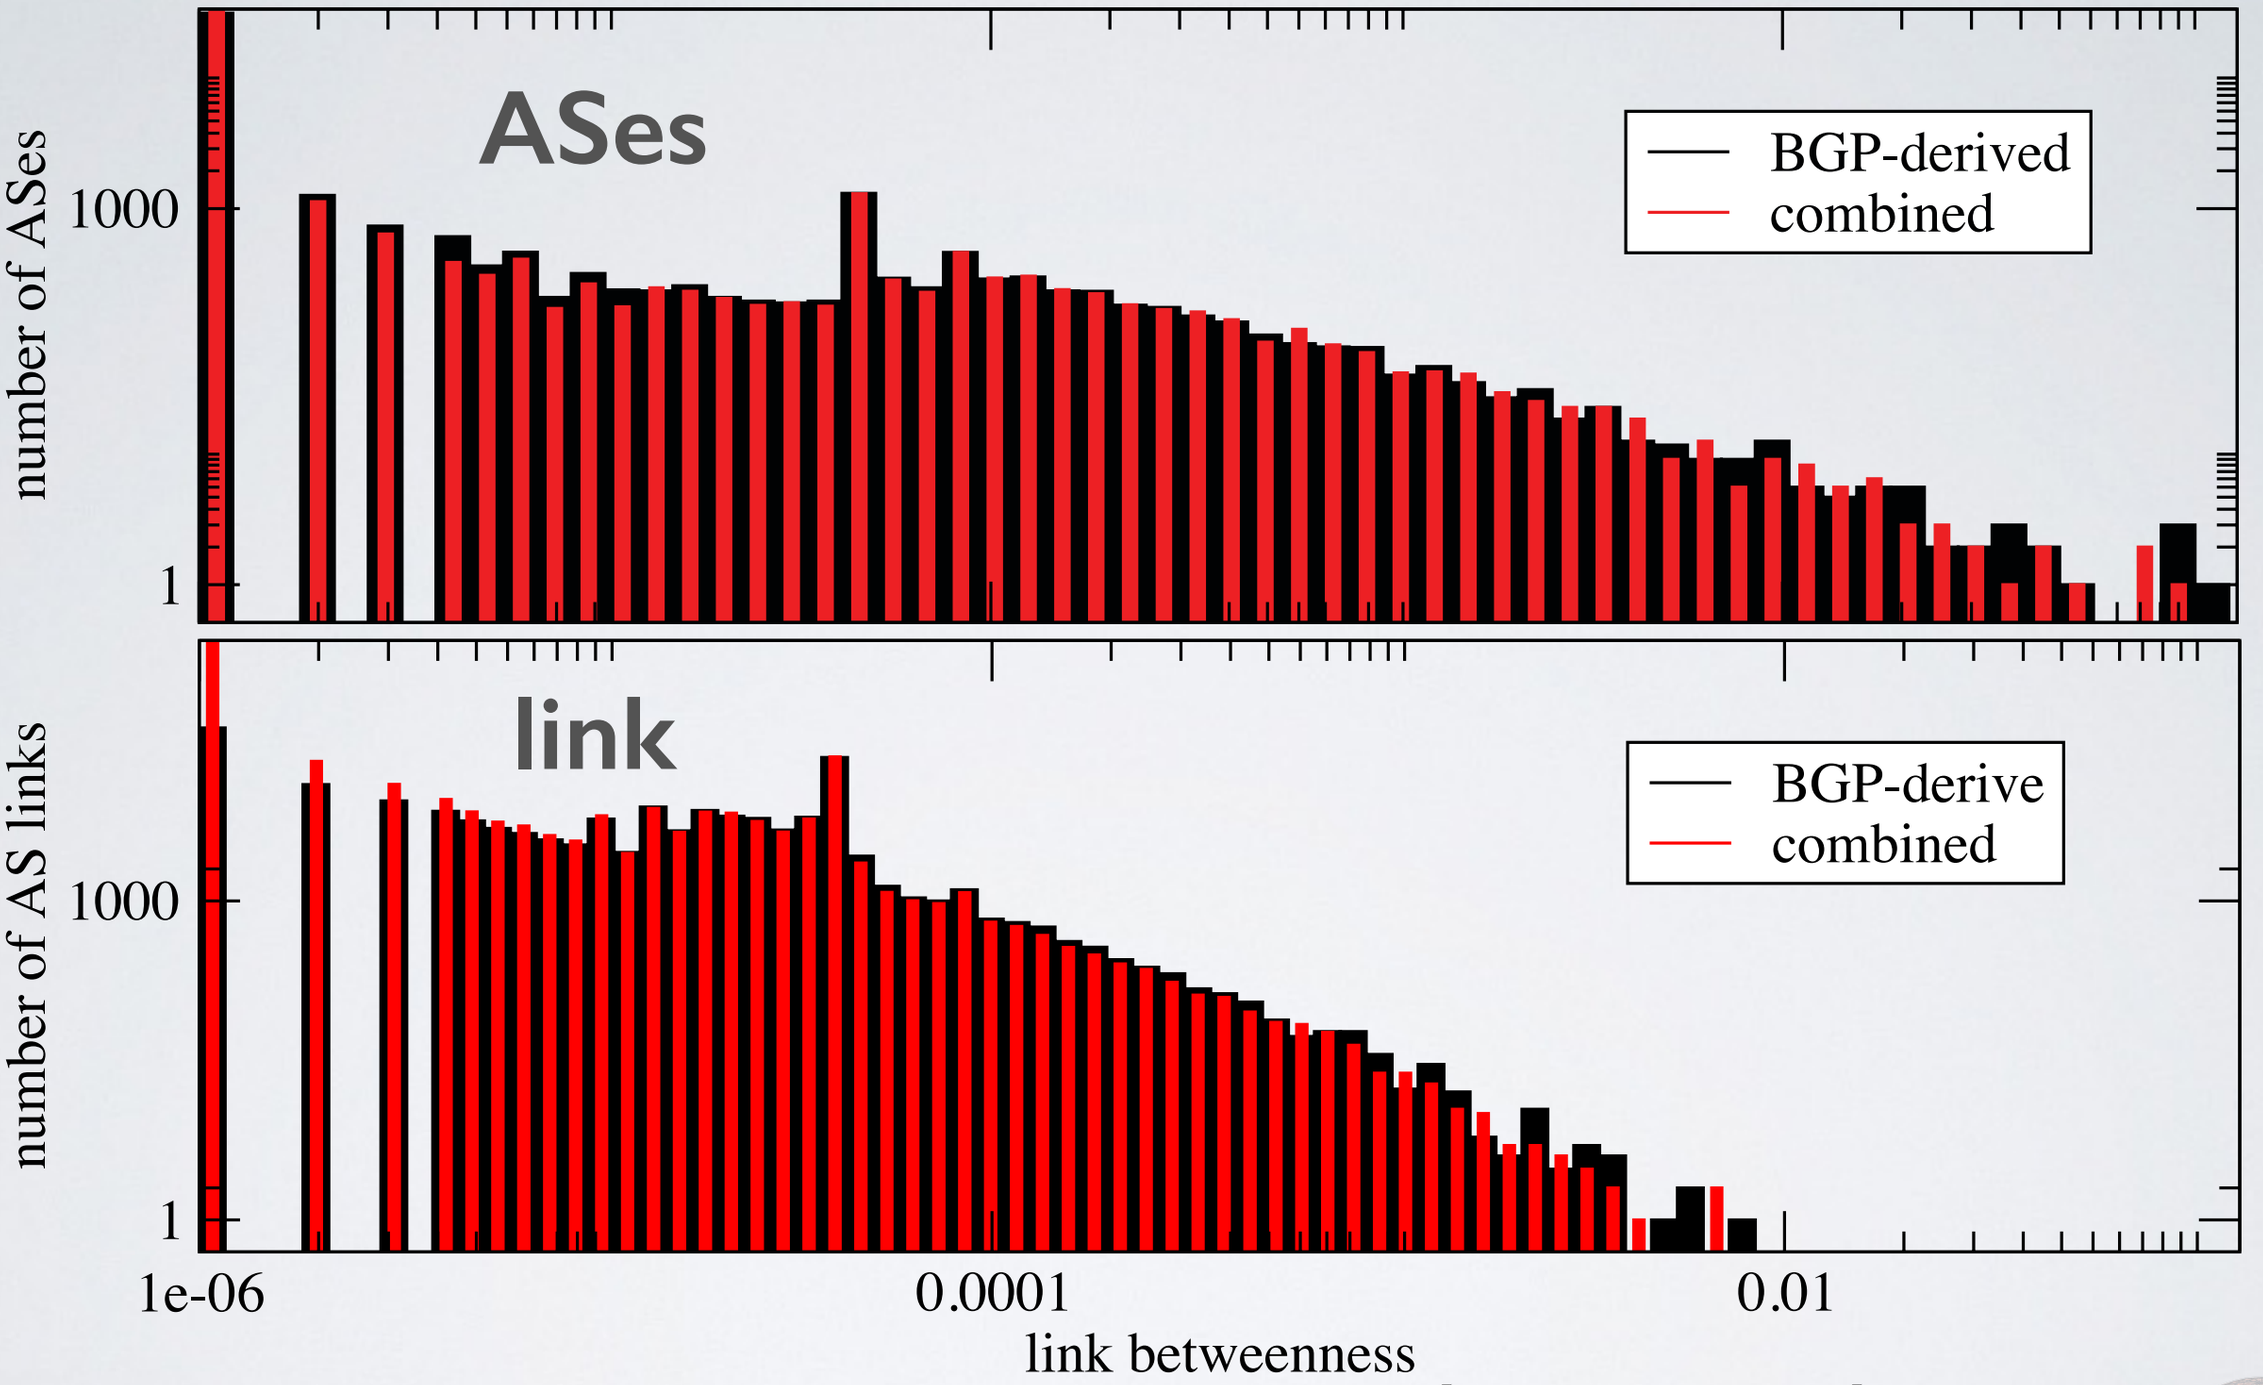

- **link:** average dropped from **21x10^{-6} to 9x10^{-6}**
- **AS:** average dropped from **6.0x10^{-5} to 5.8x10^{-5}**

caida

BETWEENNESS

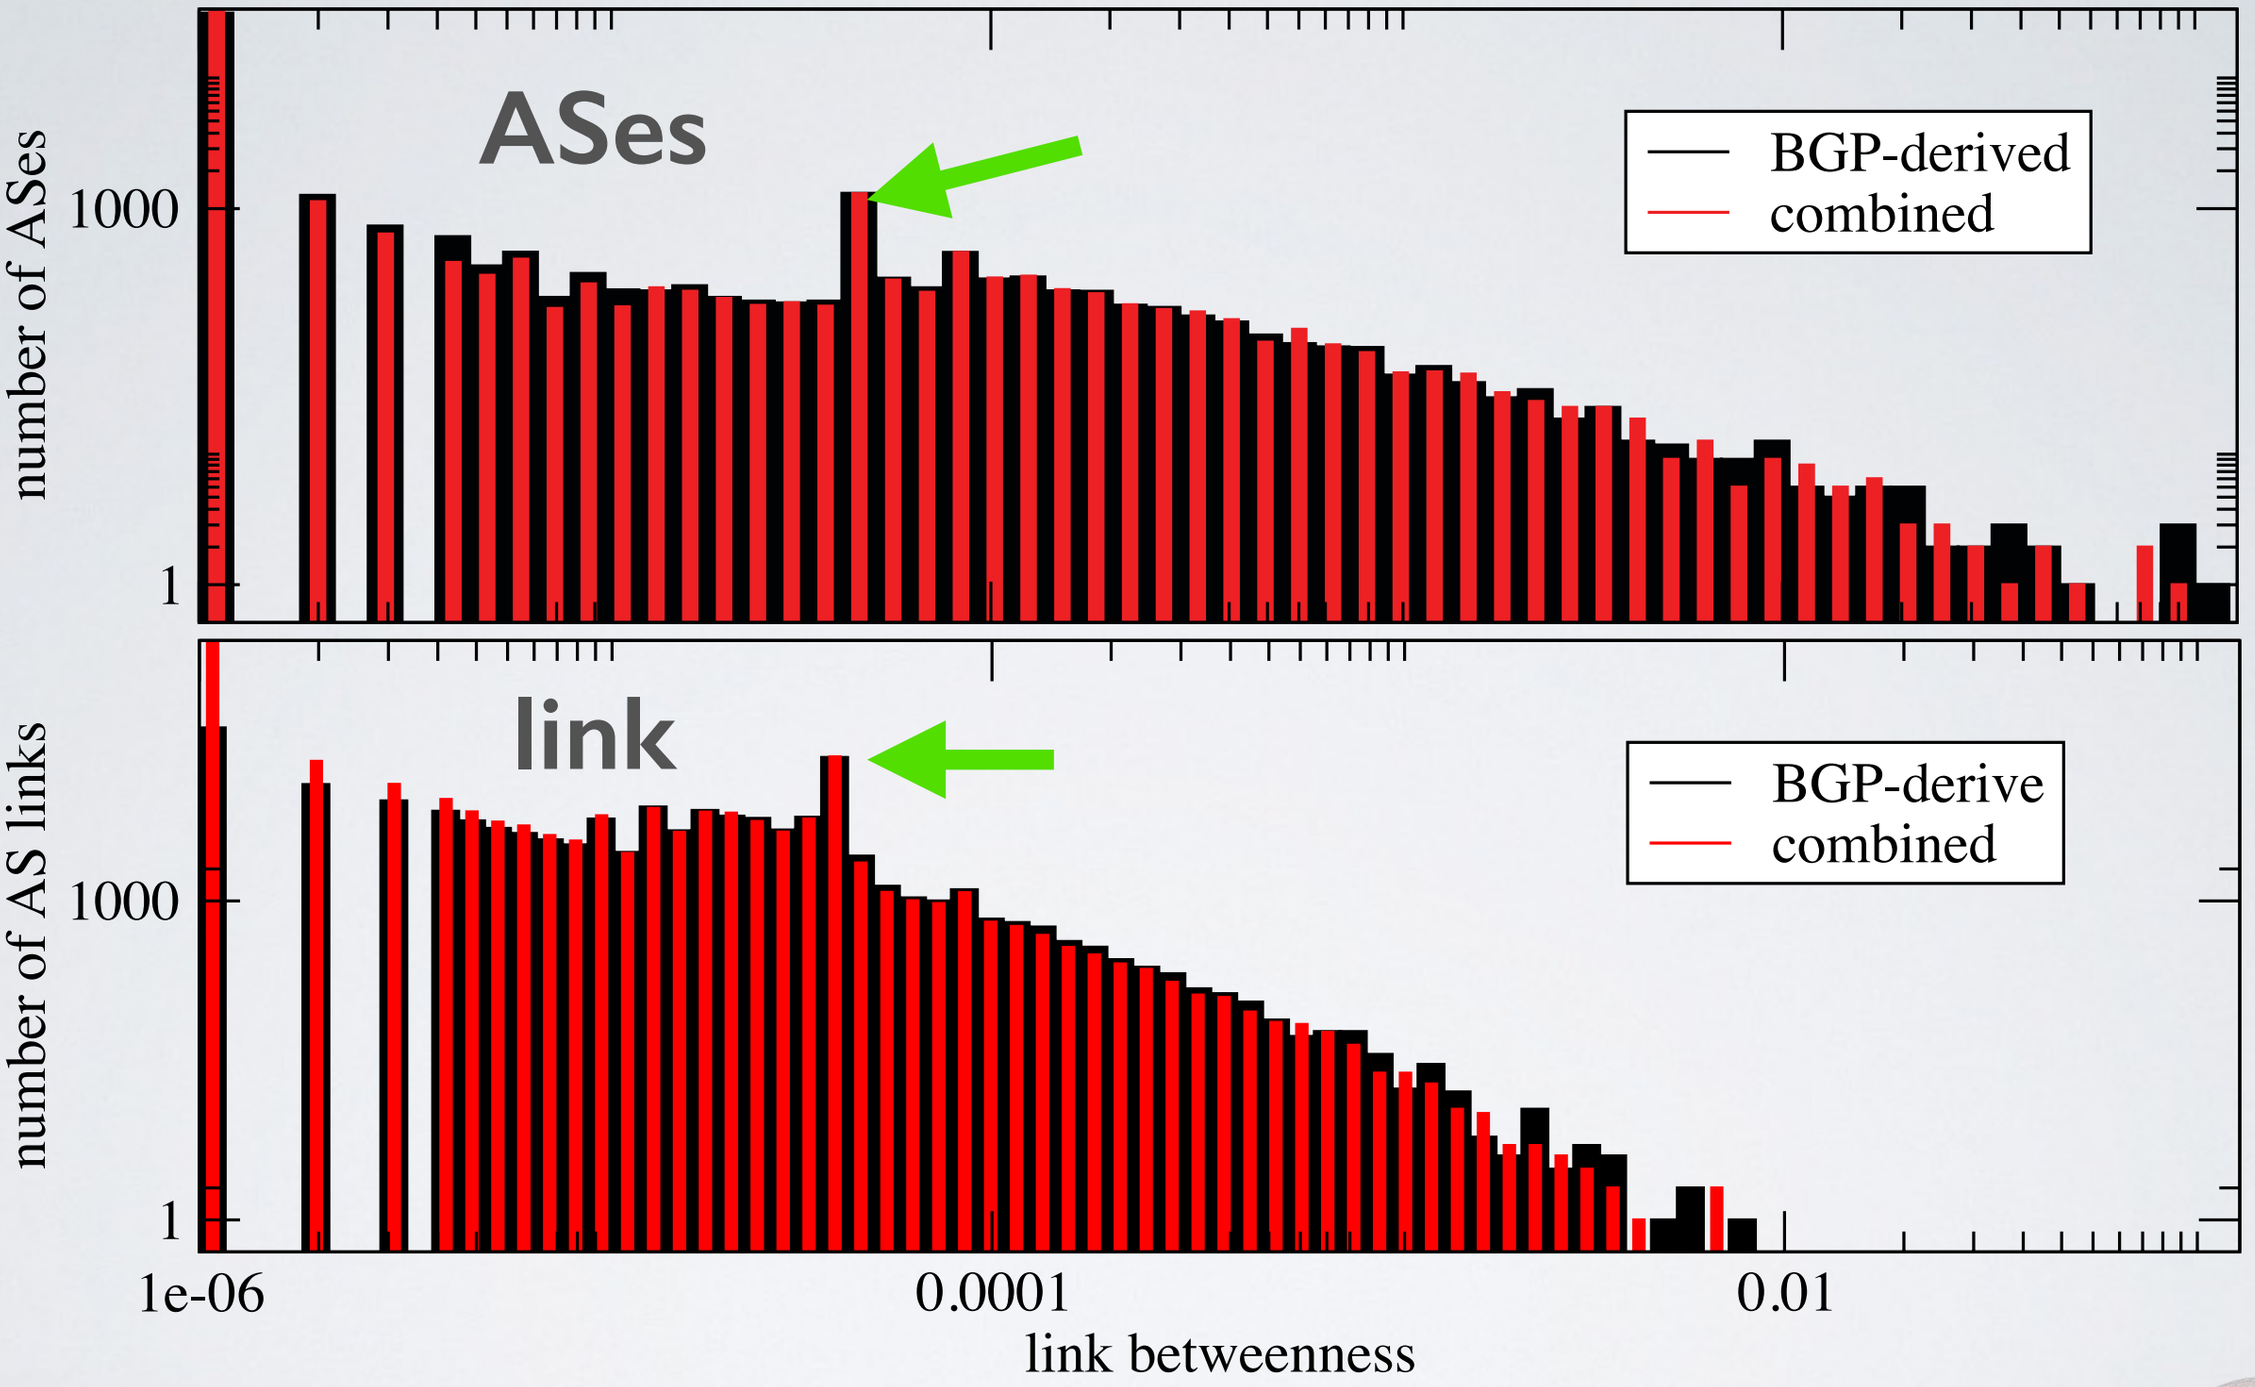

- **spike for (some) ASes** that are **neighbors of the Tier-1**

BETWEENNESS

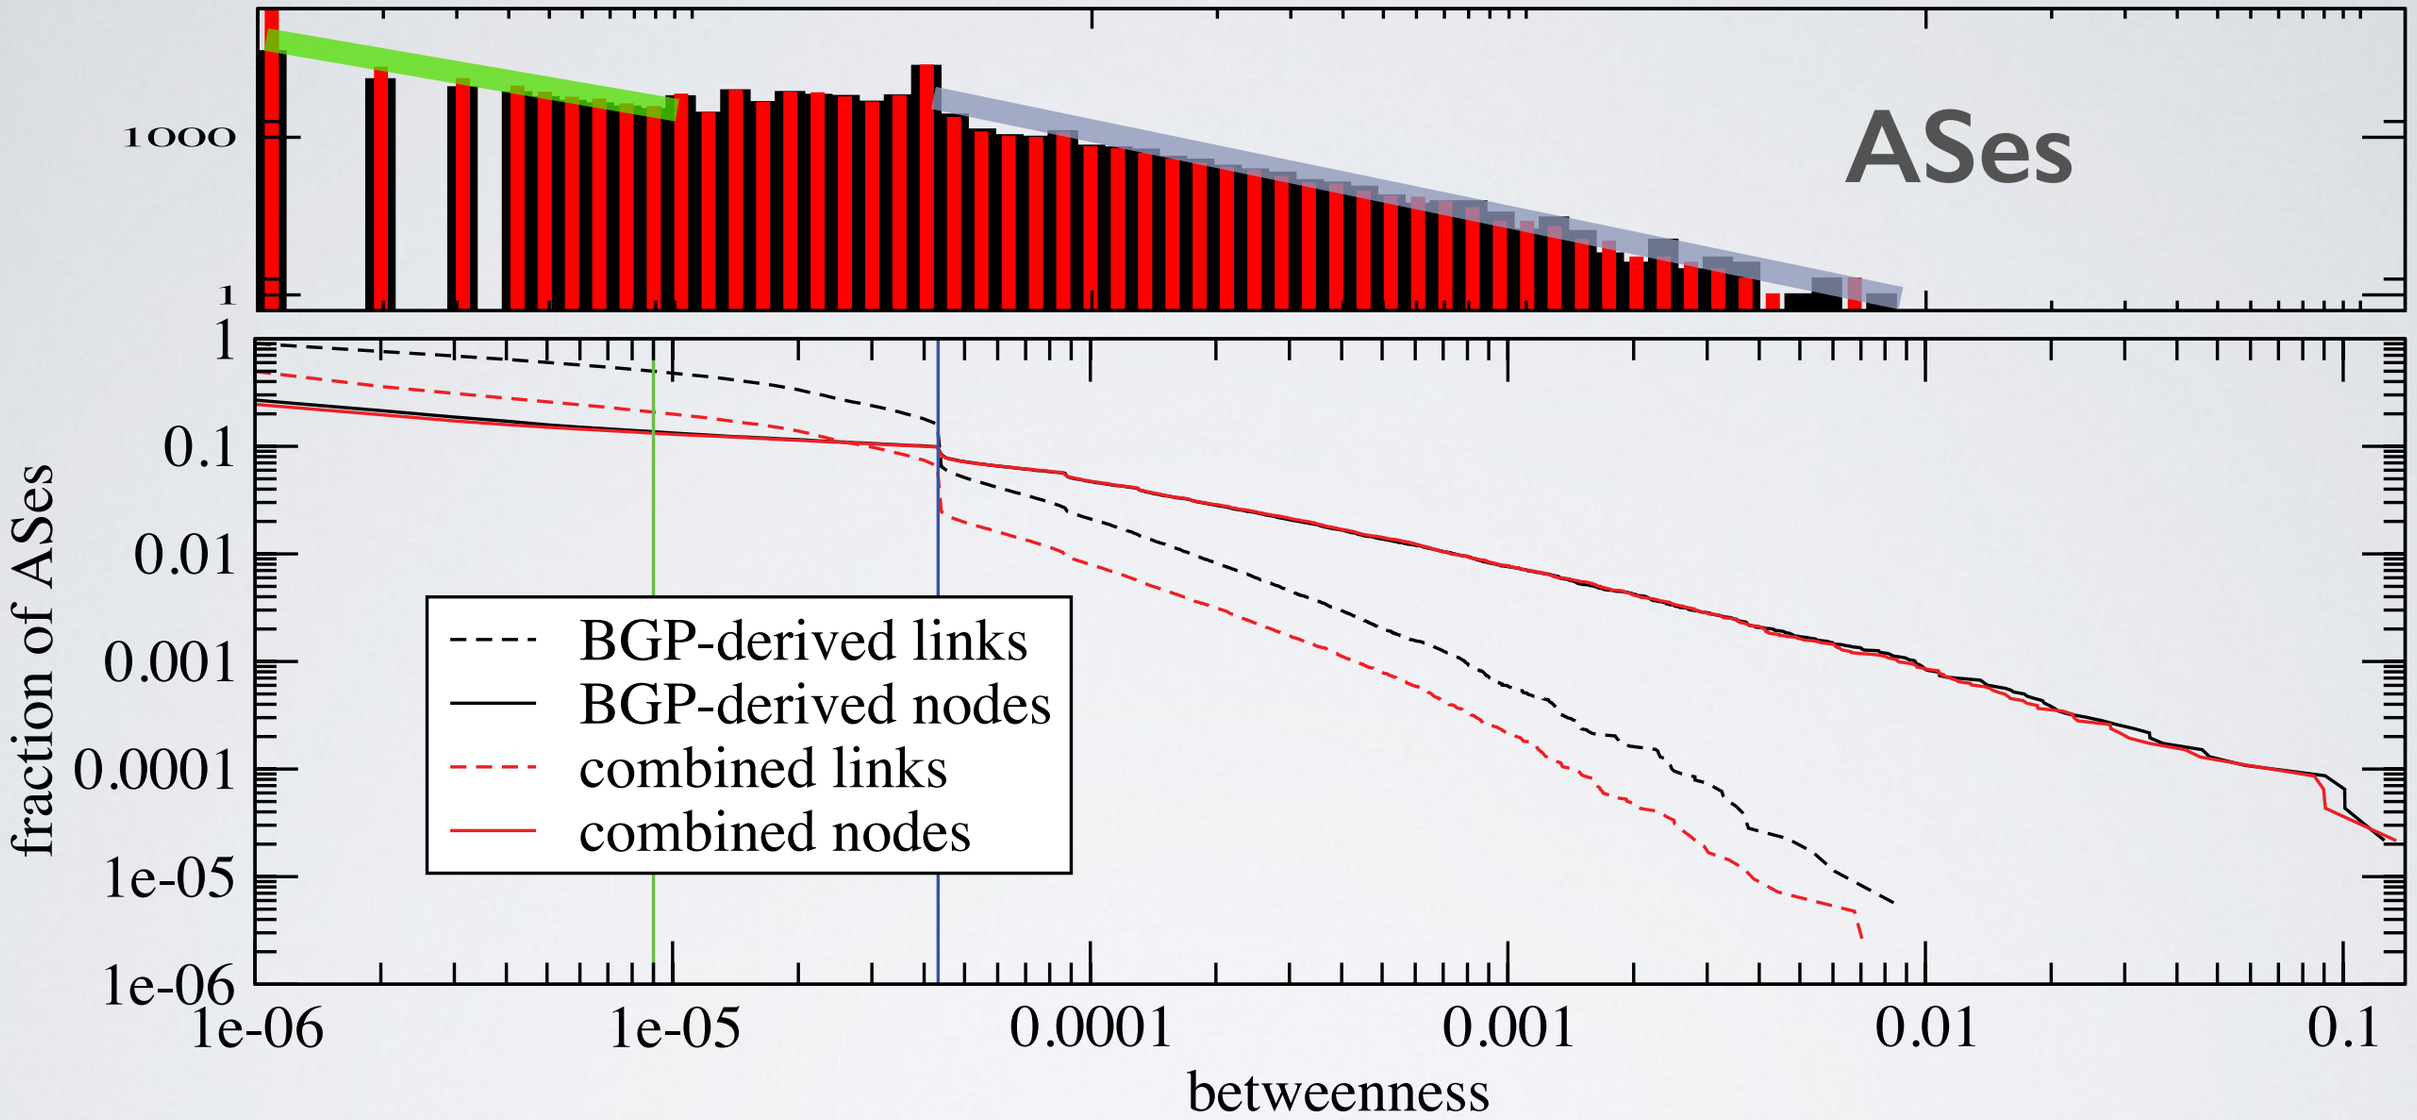

- **50%** of ASes have **betweenness <= 9x10^{-6}** (green line)
- **7%** of ASes have **betweenness >= 4x10^{-5} (blue line)**

DEGREE

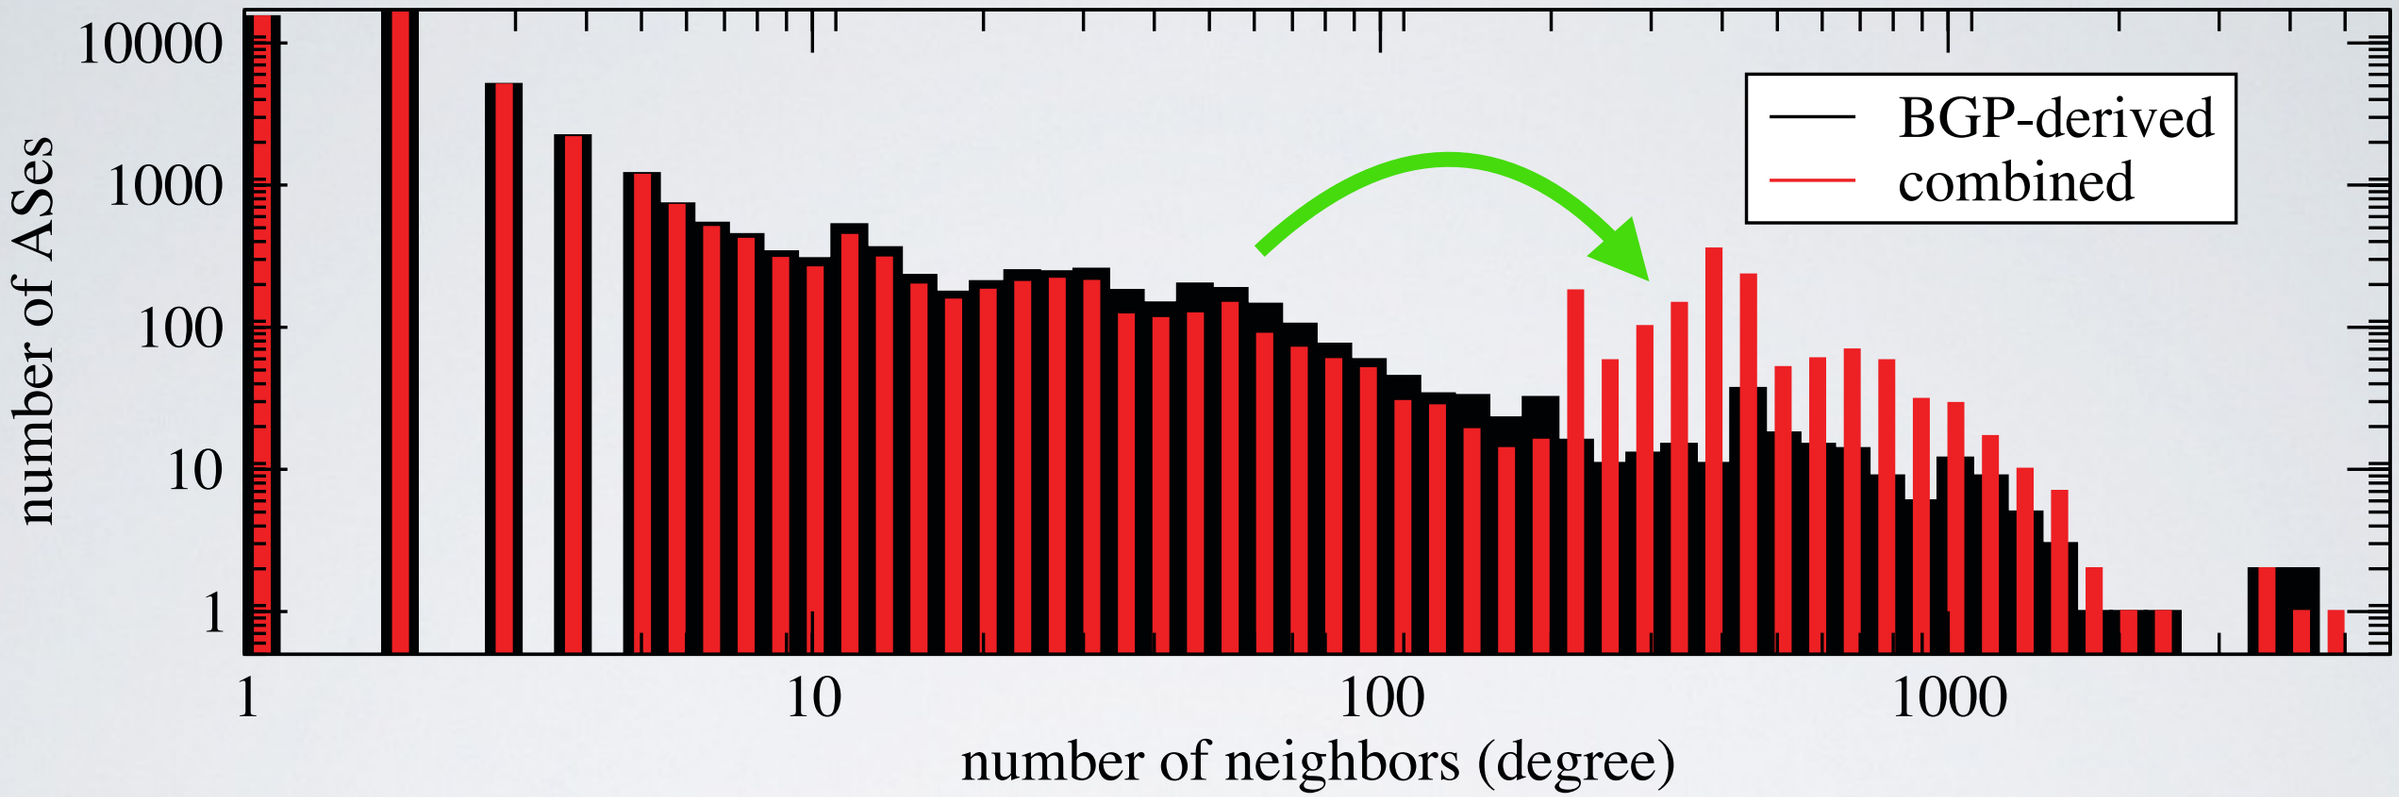

- ASes of **degree >= 230** increased from **2.9% to 4.0%**
- **average** degree rose from **7.7 to 18.0**
- **majority of links** were added to **"middle" tier**

CORENESS

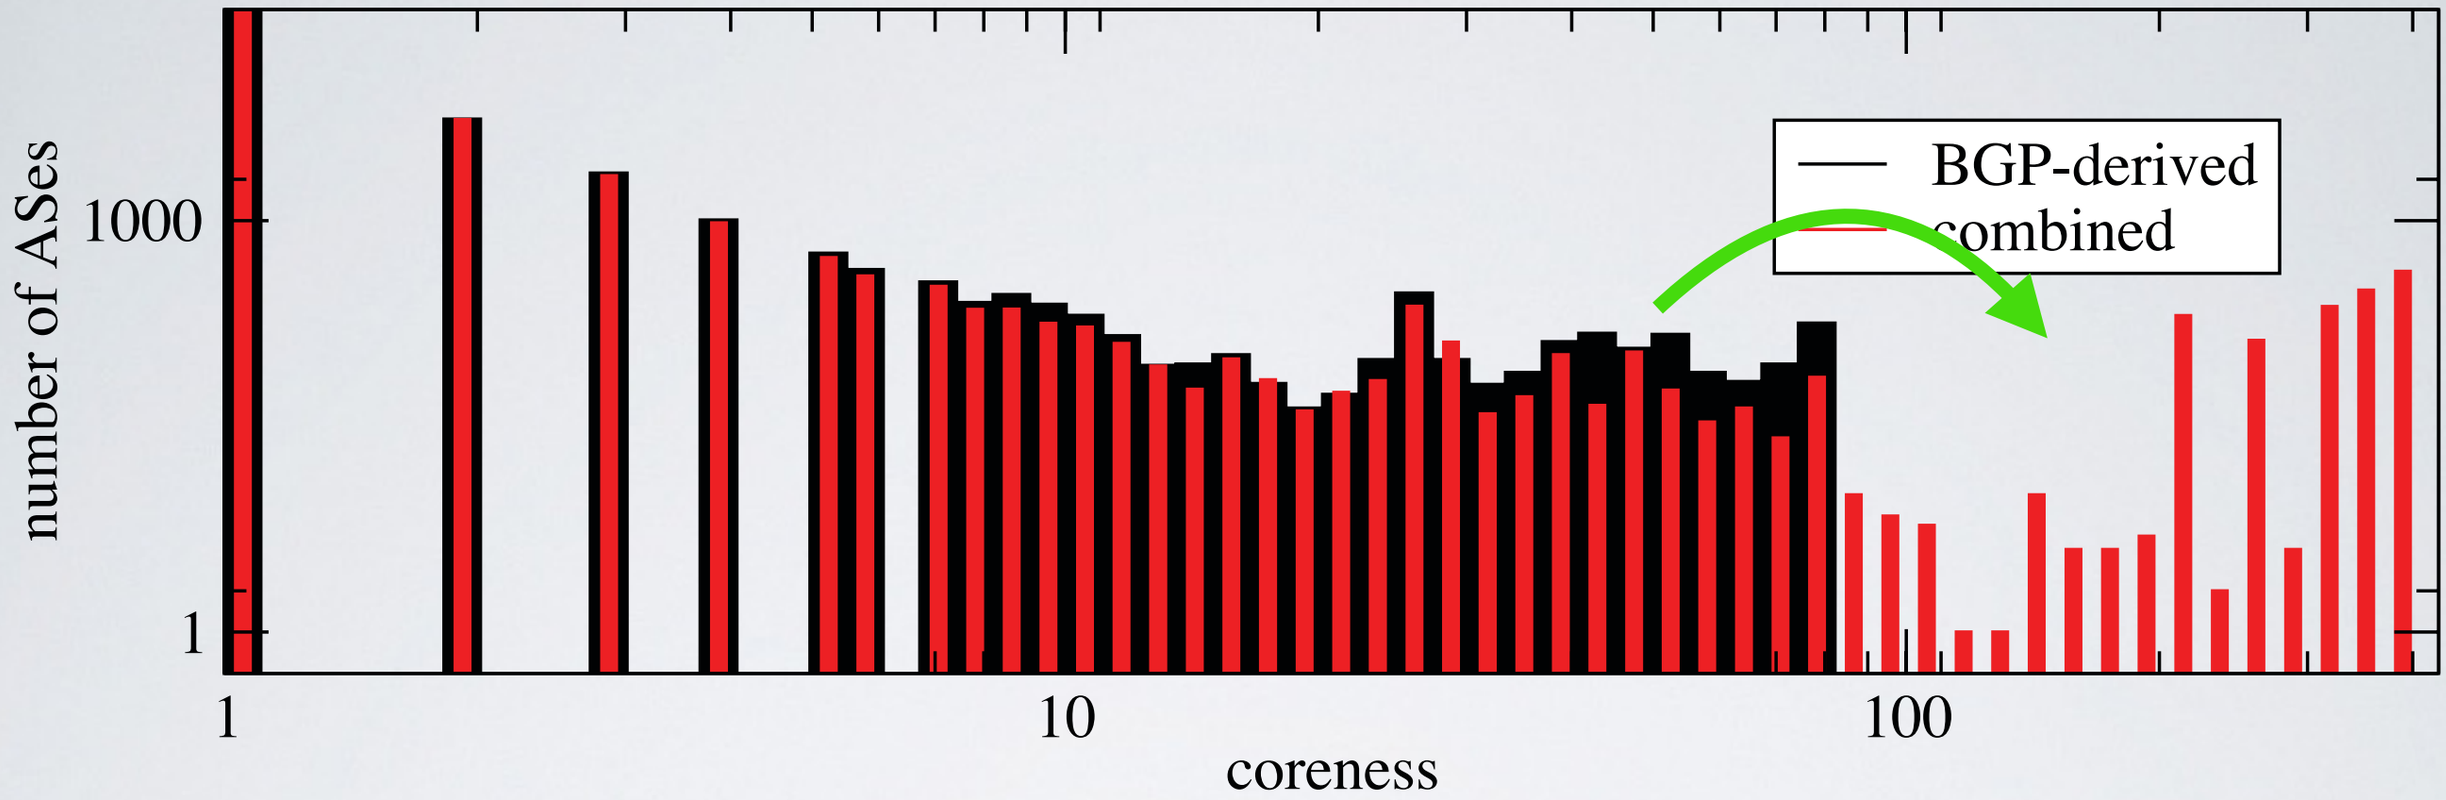

- **maximum** grew from **76 to 409**
- **3.2%** of ASes had **coreness > 76** in combined graph
- **average** grew from **2.96 to 11.93**
- **fringe left mostly unchanged, increase in core**

CLUSTERING

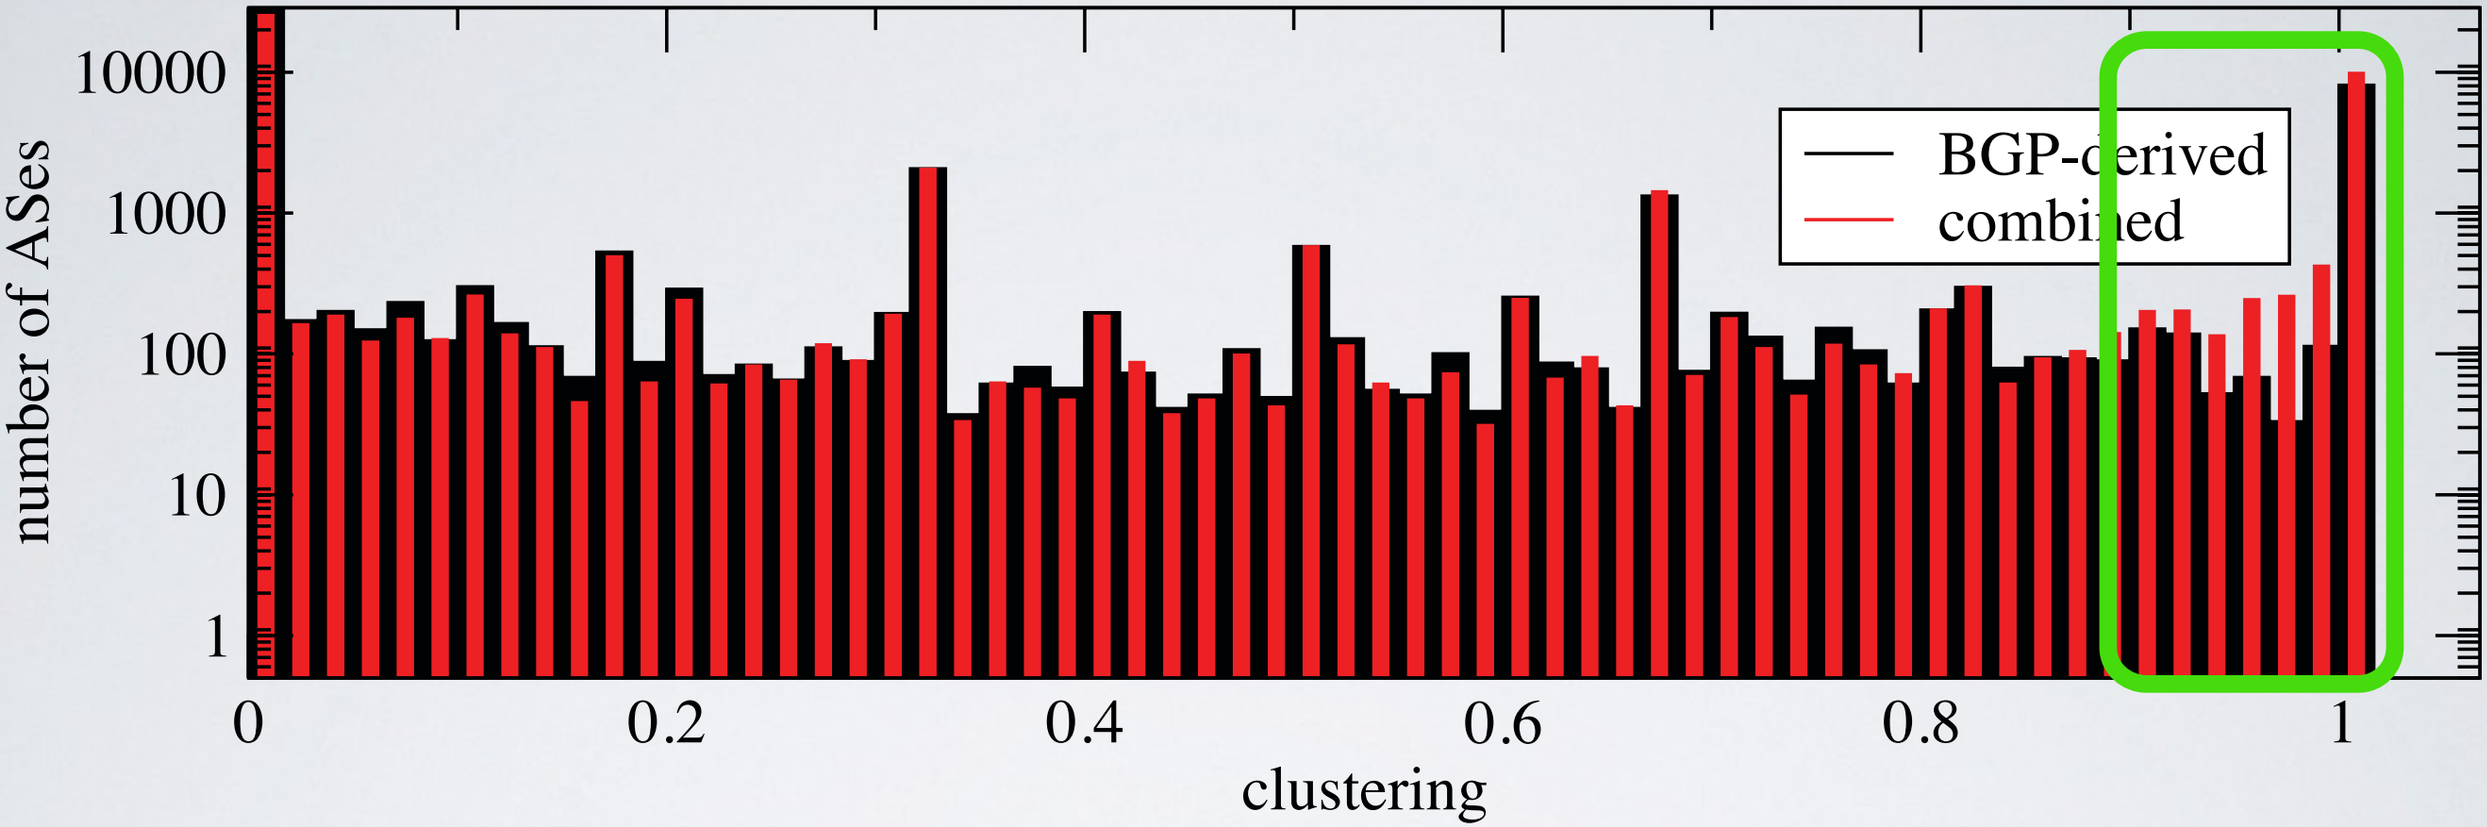

- unlike other metrics, **more uniform distribution**
- **increase** in number of ASes with **higher values**
- **average** increase from **0.28 to .33**
- coefficient**, global,** increase from **0.05 to .47**

caida

PEERING & CUSTOMER CONE

	customer	peer
A	A,C,B,D,E,F	A,B,C,D,E,F,G,H,I
B	B,E,F	B,C,D,E,F,H,I
C	C,D,E	C,B,D,E,F
D	D	D
E	E	E
F	F	F
G	G,H,I	A,B,C,D,E,F,G,H,I
H	H,I	B,E,F,H,I

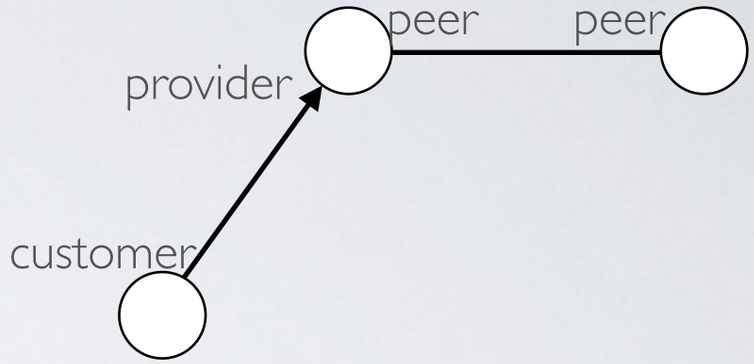

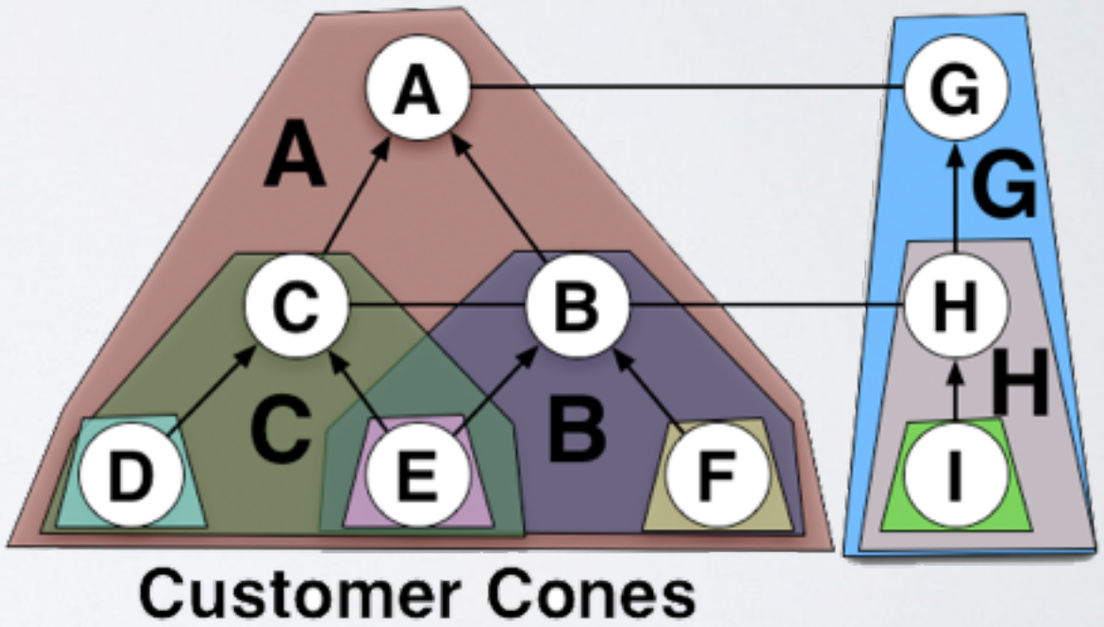

Customer Cones

CUSTOMER/PEER CONE

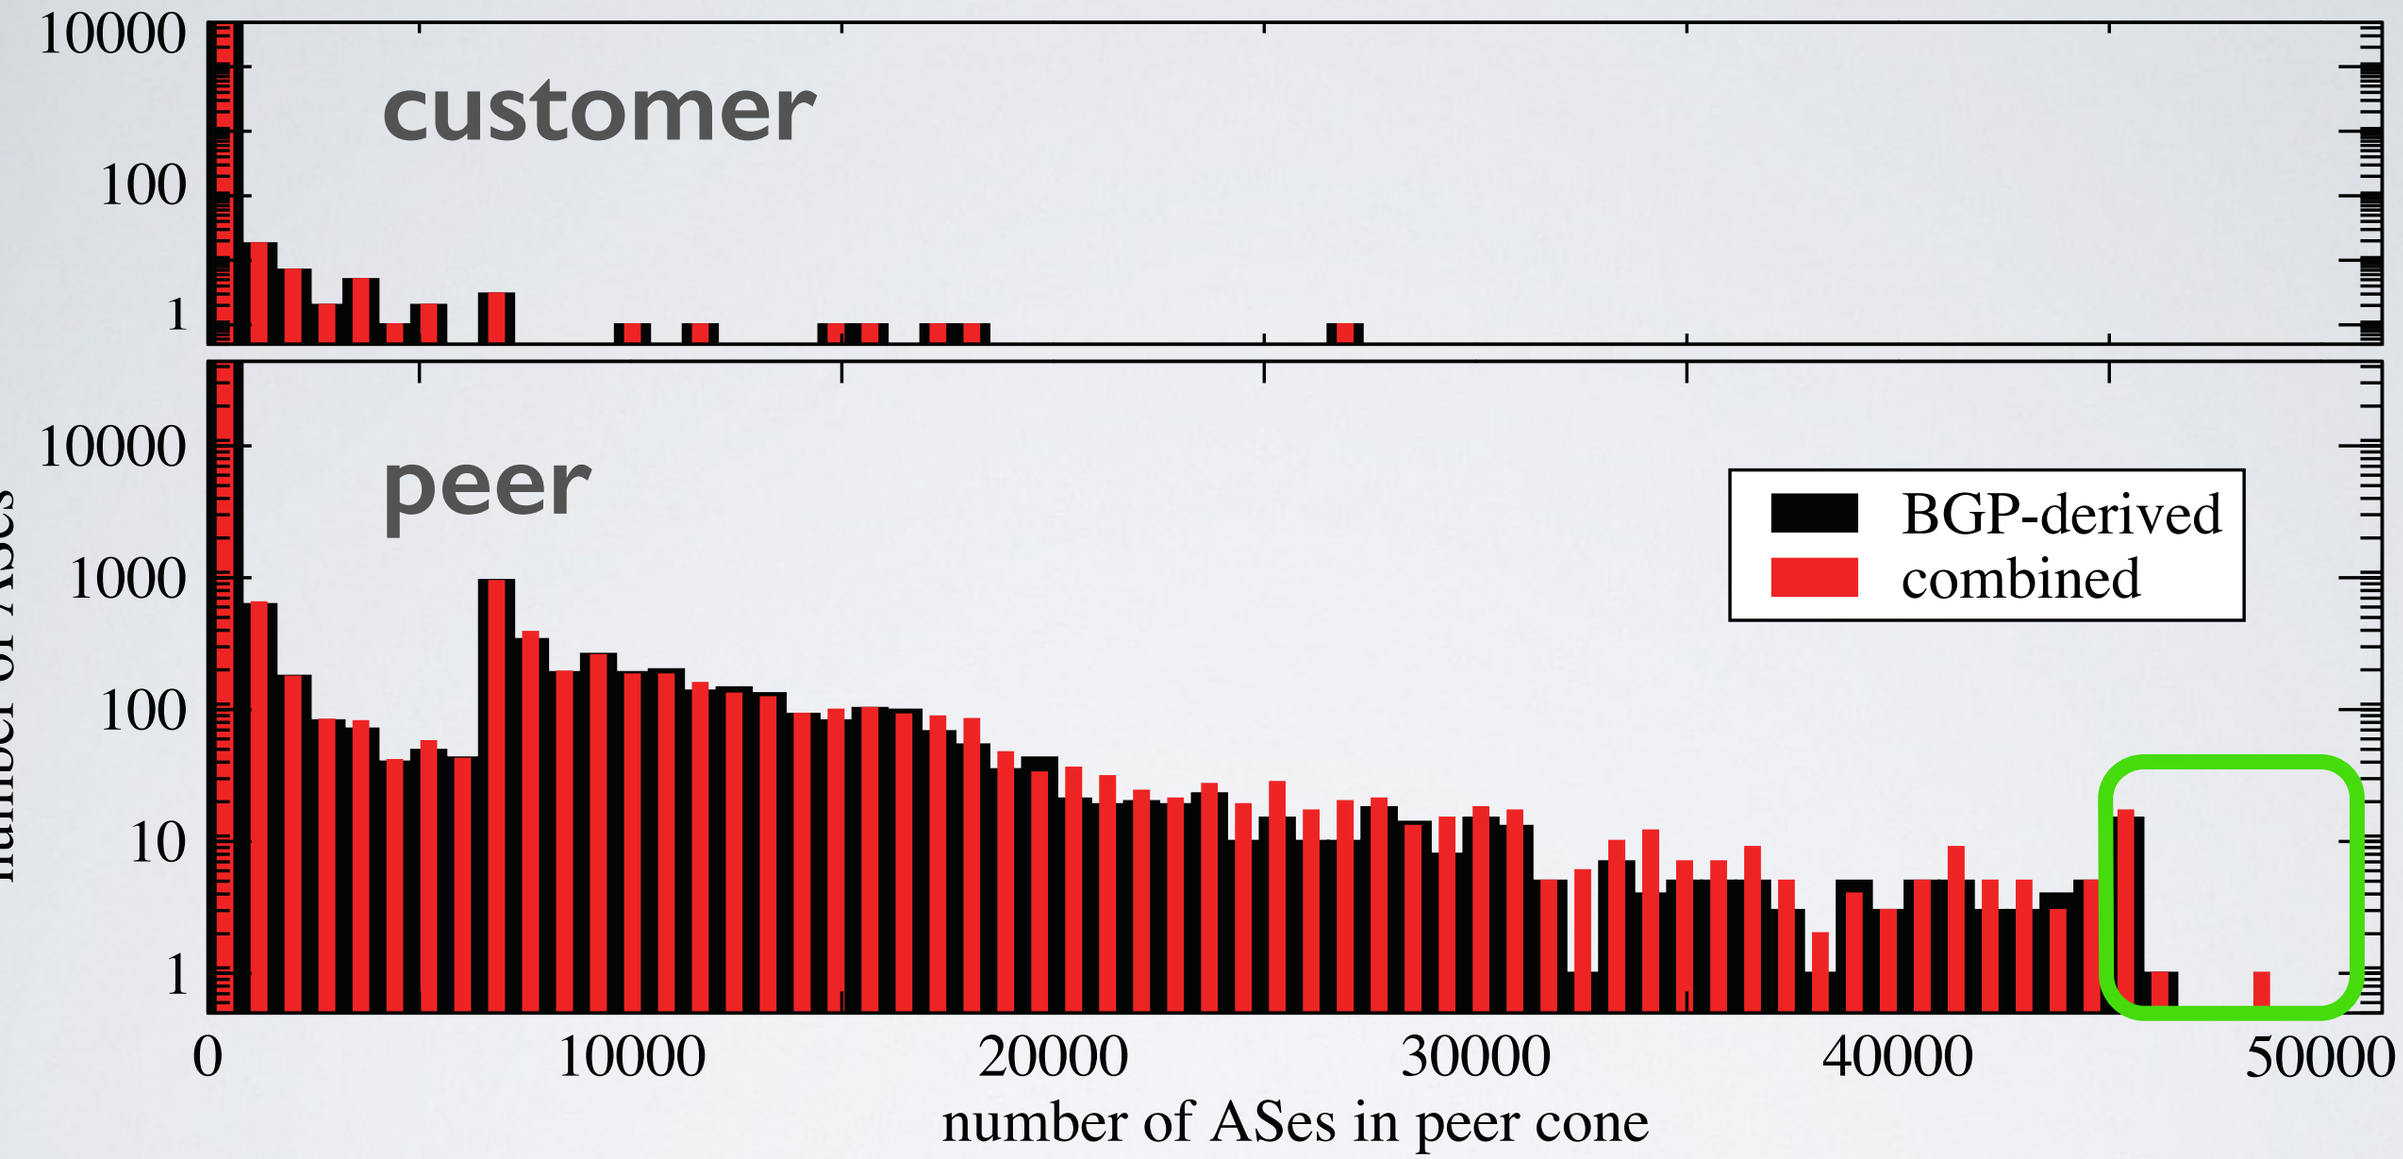

- **peer cone is not complete** (missing peering links)

CUSTOMER/PEER CONE

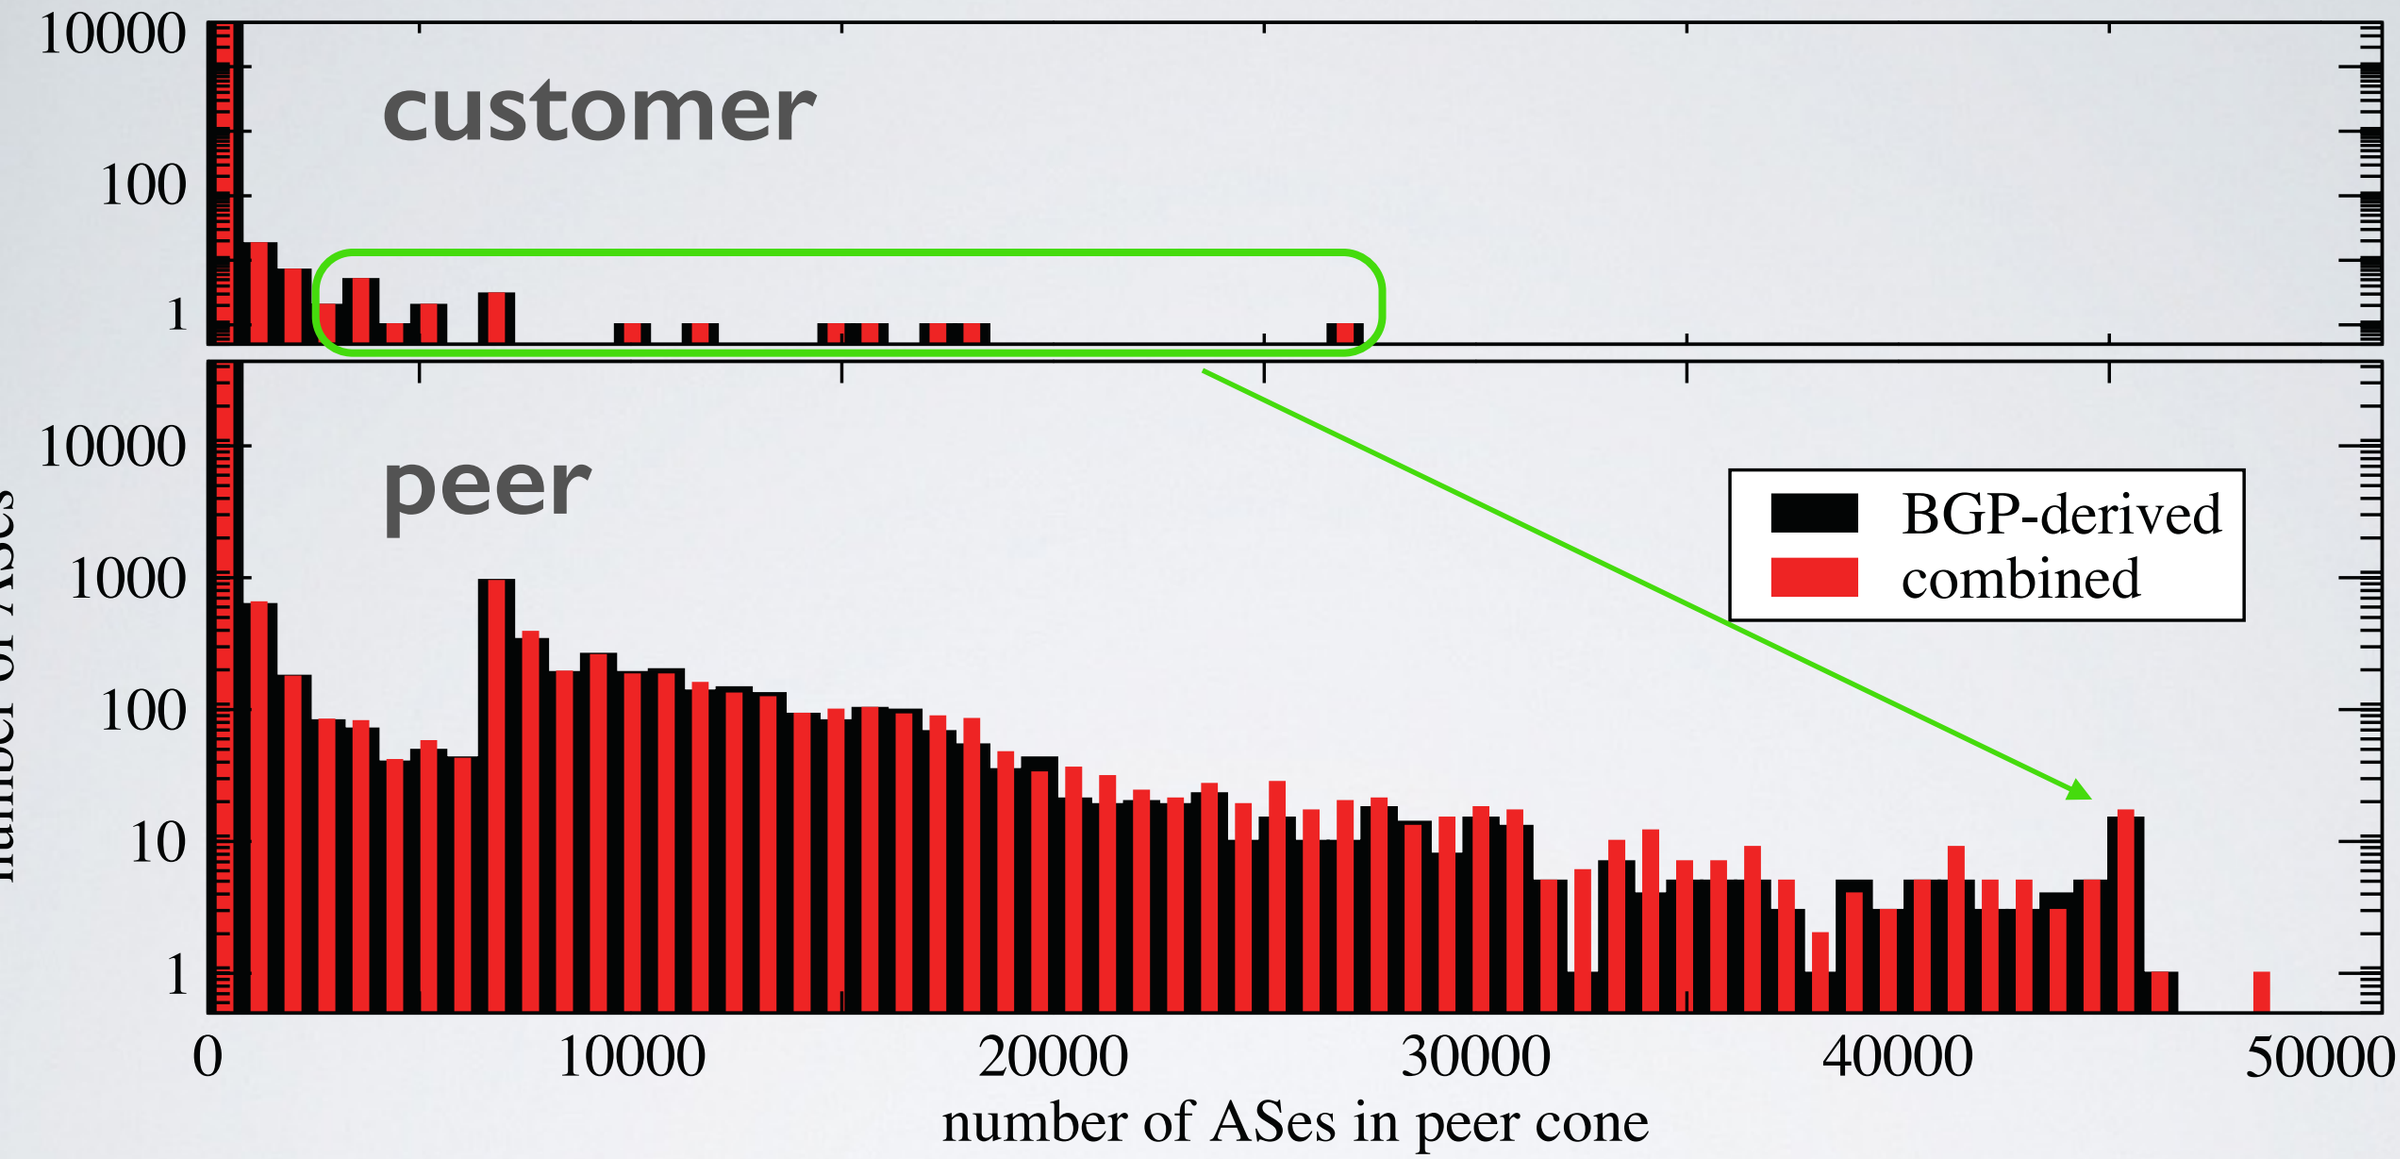

- **peer cone is not complete** (missing peering links)
- **customer Tier1s merged into single peer bin**

CUSTOMER/PEER CONE

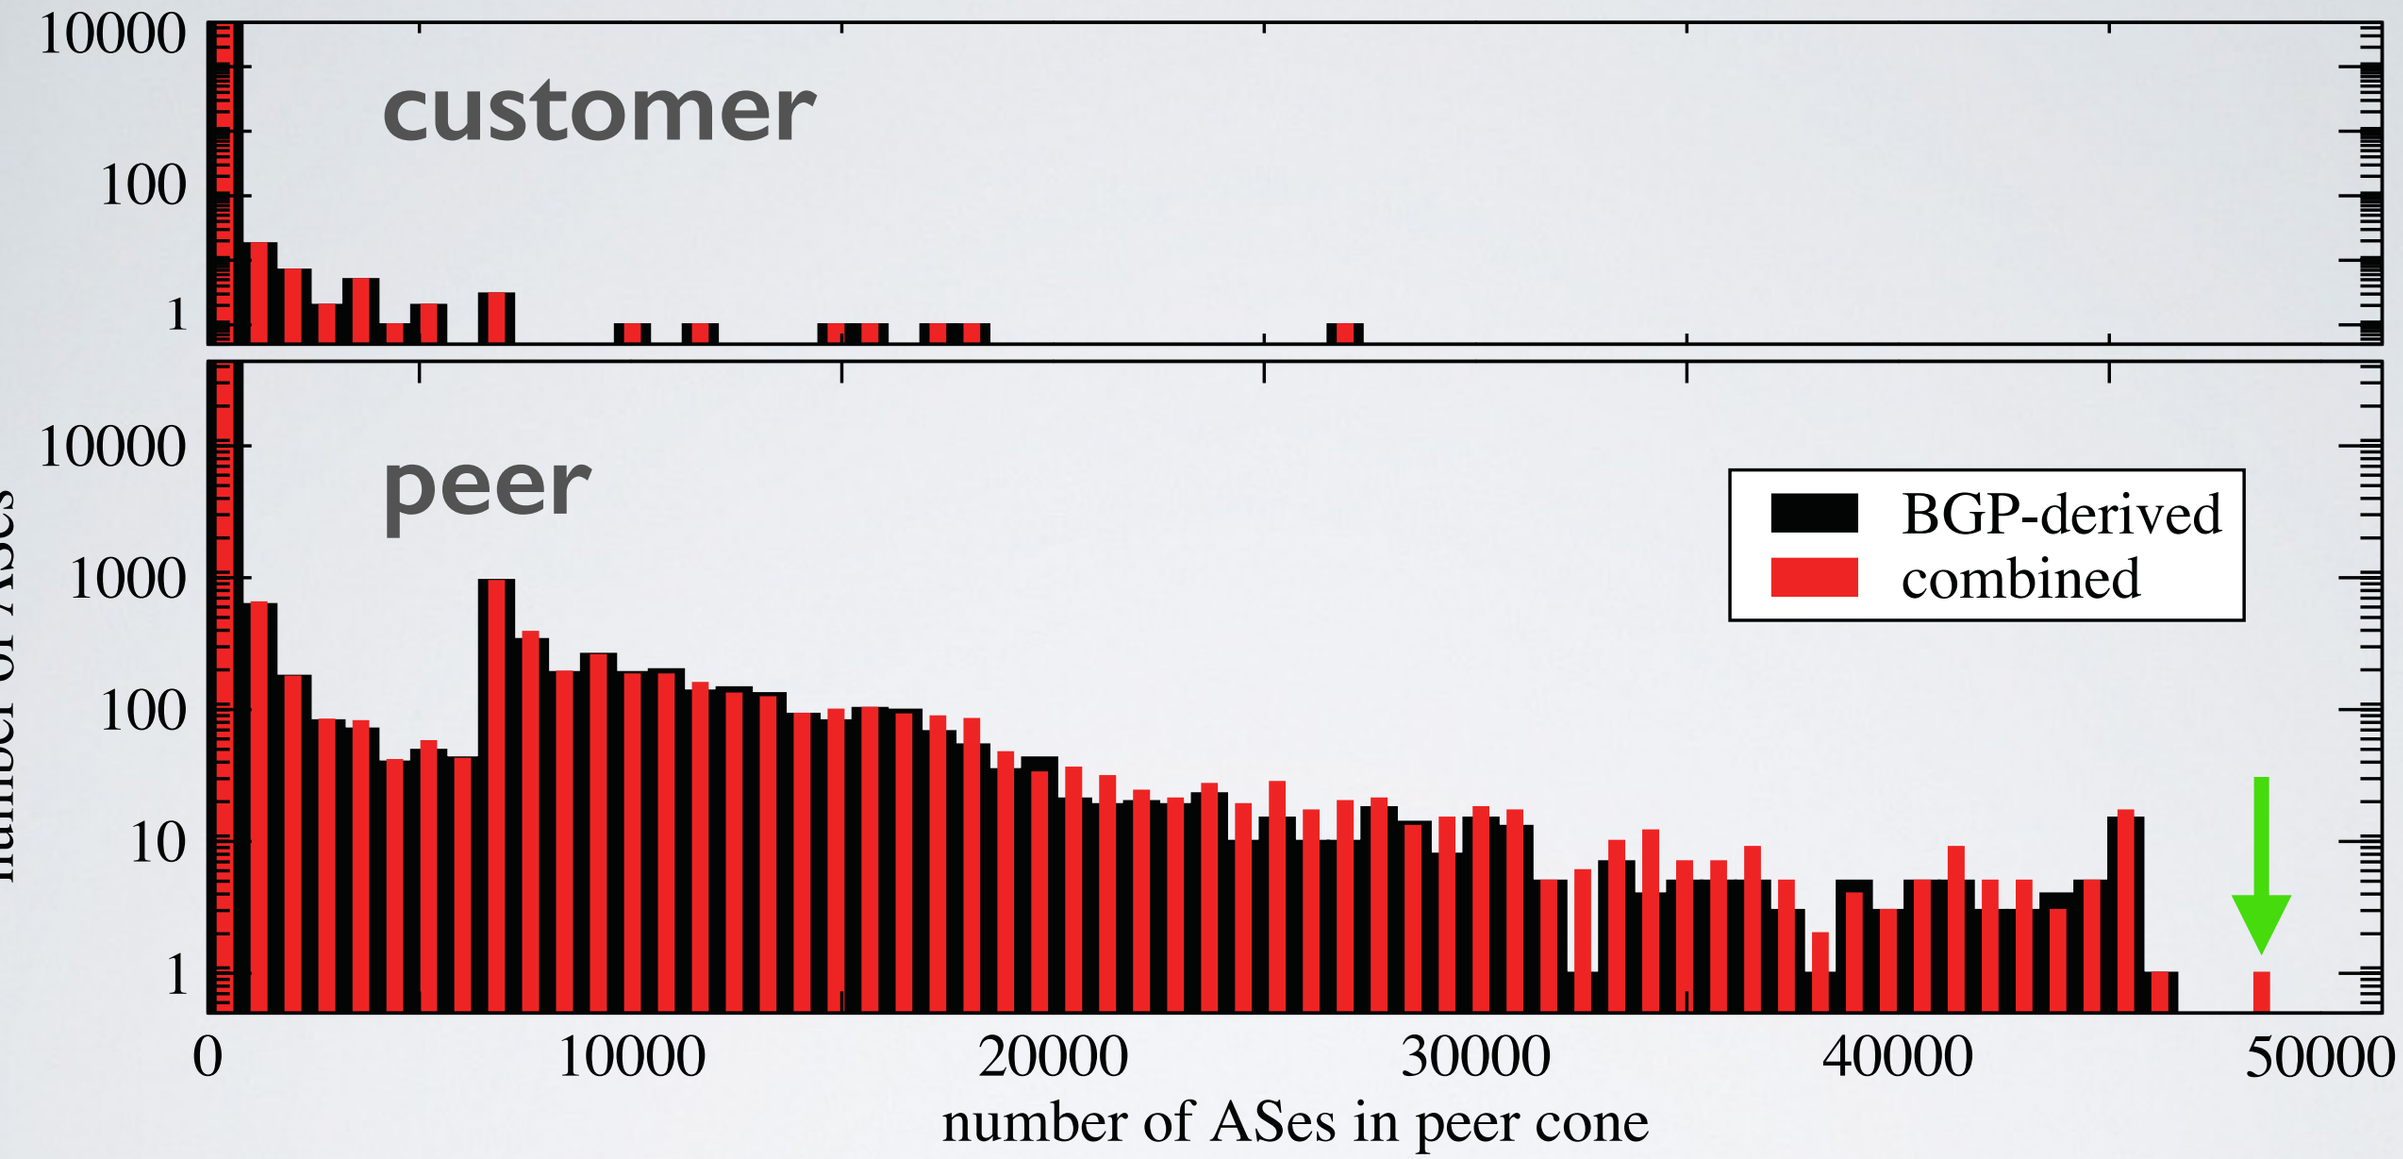

- **peer cone is not complete** (missing peering links)
- **customer Tier 1s merged into single peer bin**
- **Hurricane Electric** stands **alone** in **peer cone**

EDGE/CORE

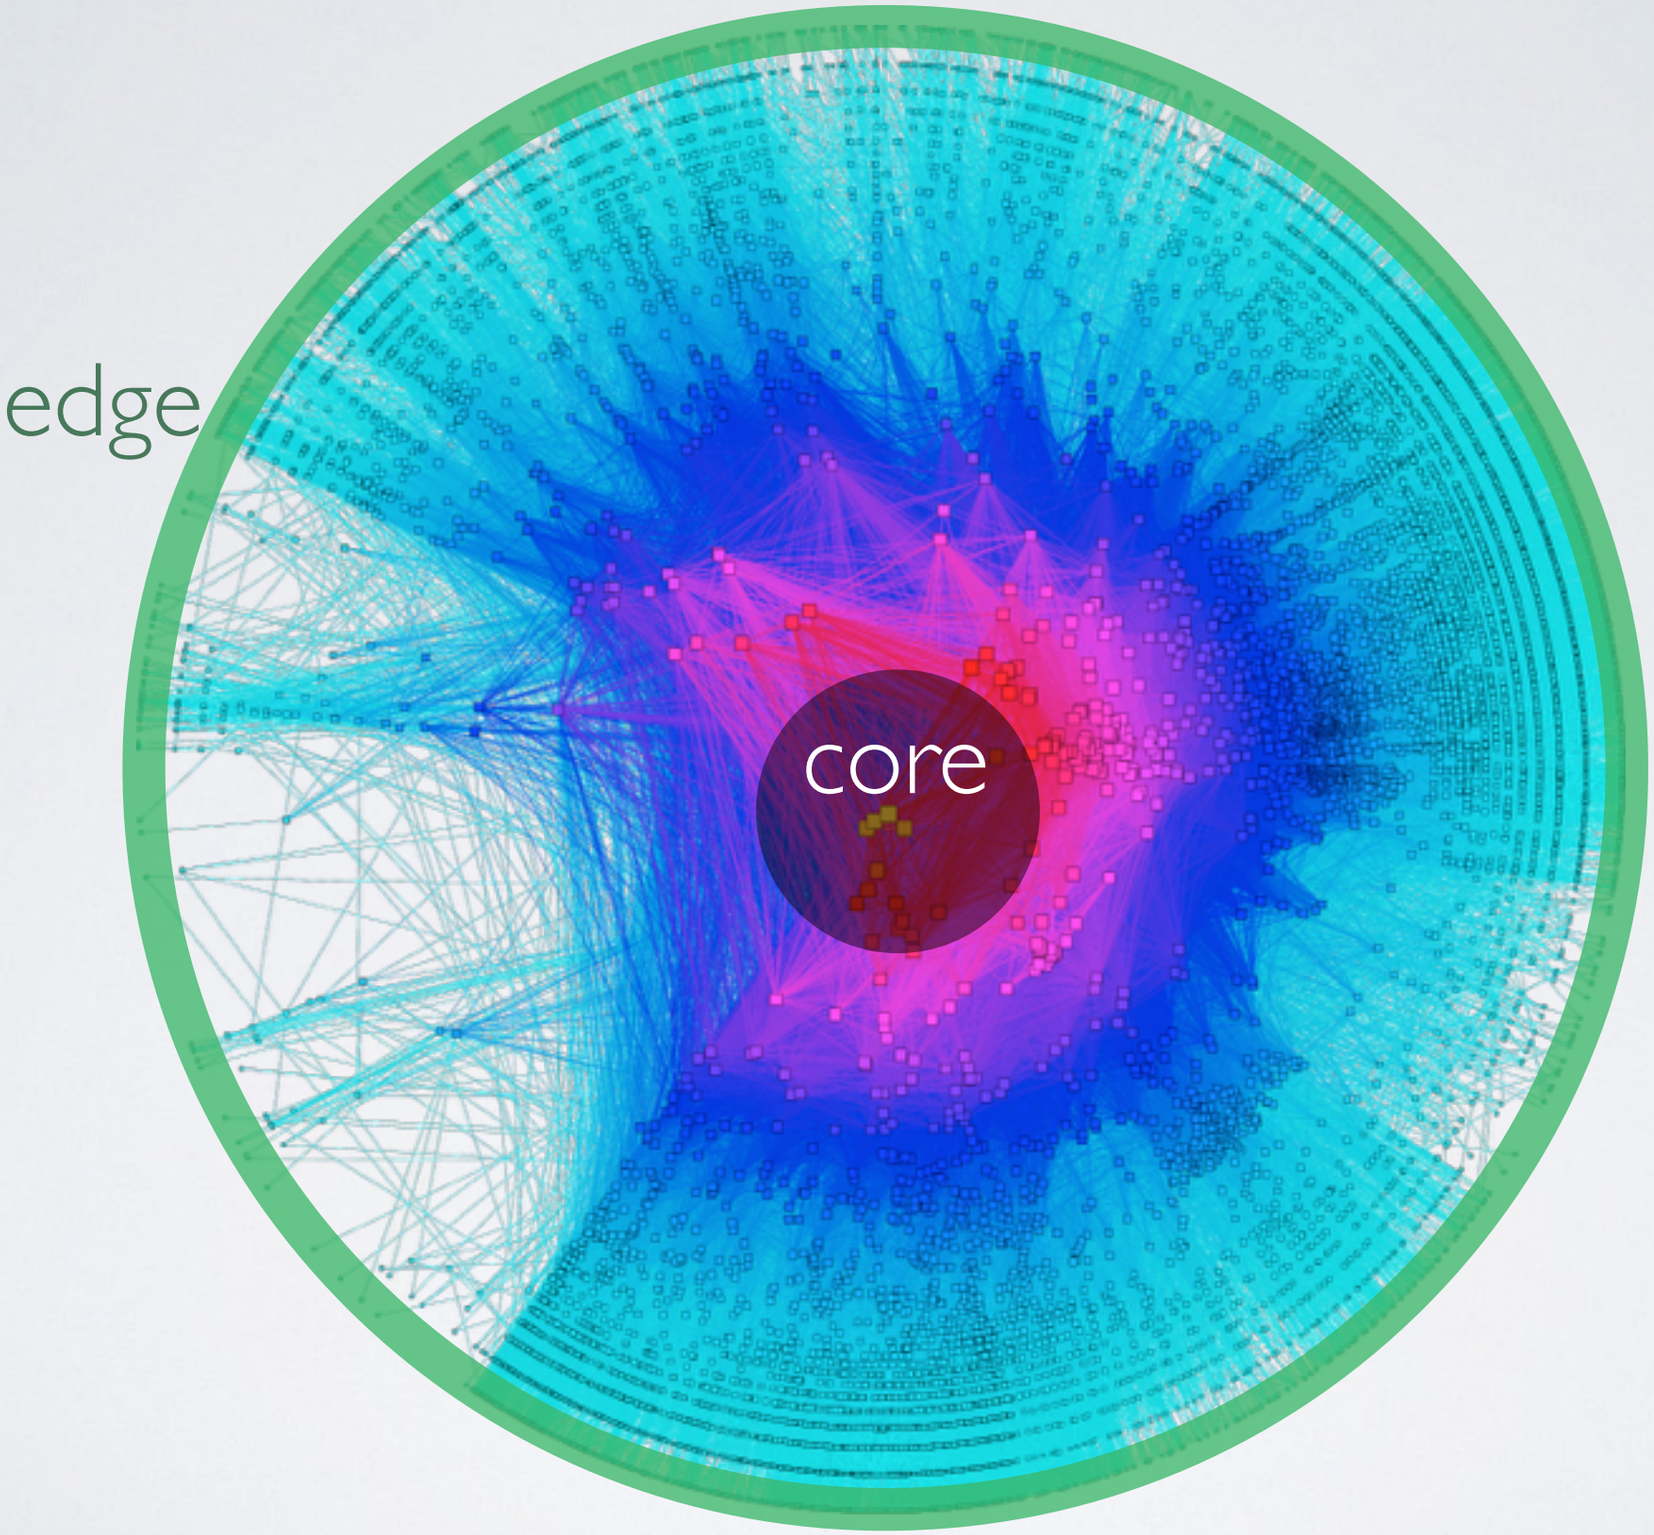

EDGE/CORE

	edge	**core**
eccentricity	ASes with the **largest** value	ASes with **smallest** value
betweenness (fuzzy)	ASes with **smaller** values	ASes with **larger** values
coreness	ASes with **smallest** value	ASes with **largest** value
customer cone	ASes with **smallest** value	**largest clique** in ASes with largest **transit degree***

*defined in AS Relationship inference algorithm

CORE/EDGE

	number of ASes			
eccentricity	periphery		center	
BGP-derived	346	0.073%	57	0.0012%
combined	93	0.020%	1	0.0002%
betweenness	edge		central	
BGP-derived	39,867	86%	4,651	10%
combined	40,178	87%	4639	10%
coreness	fringe		core	
BGP-derived	15,533	34%	151	0.32%
combined	15,514	33%	422	0.91%
customer cone	edge		clique	
BGP & combined	39,142	85%	16	0.035%

- edge as defined by (**betweenness, customer cone)** are most similar in size (and membership)

- unlike other metrics, **eccentricity** has **most** ASes in **middle**

AS RANK

An AS's **rank** is equal to the **number of ASes** that have a **more "central" value** than the AS.

CLIQUE AS RANK

	ASes in transit degree clique															
	3356	**174**	**1299**	**2914**	**3257**	**6453**	**6762**	**2828**	**3549**	**701**	**7018**	**1239**	**3320**	**209**	**12956**	**5511**
peer cone size																
BGP-derived	11	10	14	1	5	9	4	6	0	15	12	8	2	3	6	13
combined	14	2	7	3	5	4	11	10	1	16	12	13	9	6	8	15
customer cone size																
BGP-derived	0	1	2	3	4	5	6	8	9	10	11	13	15	17	29	32
combined	0	1	2	3	4	5	6	8	9	10	11	13	15	17	29	32
eccentricity																
BGP-derived	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
combined	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1
betweenness																
BGP-derived	1	0	11	10	16	22	27	14	3	9	4	25	46	5	71	154
combined	1	0	8	9	17	21	31	15	3	7	4	24	46	6	78	183
degree																
BGP-derived	1	0	36	17	32	64	158	28	3	7	4	56	85	8	236	261
combined	1	0	61	26	47	195	1059	44	3	10	4	165	300	11	1432	1435
coreness																
BGP-derived	0	151	0	0	151	197	233	442	0	562	760	562	223	442	487	562
combined	970	973	978	1316	1115	1334	1348	1506	978	1563	1563	1540	1335	1510	1510	1523
clustering																
BGP-derived	18986	19003	18764	18837	18849	18723	18377	18963	18989	19000	19013	18935	18698	18994	17977	18126
combined	21106	21127	20614	20896	20867	20751	20343	21120	21125	21150	21159	21091	20568	21146	20001	20111

- members of the transit degree clique
- (which strongly correlate with Tier 1s)

caida

CLIQUE AS RANK

	ASes in transit degree clique															
	3356	**174**	**1299**	**2914**	**3257**	**6453**	**6762**	**2828**	**3549**	**701**	**7018**	**1239**	**3320**	**209**	**12956**	**5511**
peer cone size																
BGP-derived	11	10	14	1	5	9	4	6	0	15	12	8	2	3	6	13
combined	14	2	7	3	5	4	11	10	1	16	12	13	9	6	8	15
customer cone size																
BGP-derived	0	1	2	3	4	5	6	8	9	10	11	13	15	17	29	32
combined	0	1	2	3	4	5	6	8	9	10	11	13	15	17	29	32
eccentricity																
BGP-derived	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
combined	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1
betweenness																
BGP-derived	1	0	11	10	16	22	27	14	3	9	4	25	46	5	71	154
combined	1	0	8	9	17	21	31	15	3	7	4	24	46	6	78	183
degree																
BGP-derived	1	0	36	17	32	64	158	28	3	7	4	56	85	8	236	261
combined	1	0	61	26	47	195	1059	44	3	10	4	165	300	11	1432	1435
coreness																
BGP-derived	0	151	0	0	151	197	233	442	0	562	760	562	223	442	487	562
combined	970	973	978	1316	1115	1334	1348	1506	978	1563	1563	1540	1335	1510	1510	1523
clustering																
BGP-derived	18986	19003	18764	18837	18849	18723	18377	18963	18989	19000	19013	18935	18698	18994	17977	18126
combined	21106	21127	20614	20896	20867	20751	20343	21120	21125	21150	21159	21091	20568	21146	20001	20111

54 ASes same value

- **eccentricity** and **betweenness contains** clique **in highest ranked ASes**

caida

CLIQUE AS RANK

	ASes in transit degree clique															
	3356	**174**	**1299**	**2914**	**3257**	**6453**	**6762**	**2828**	**3549**	**701**	**7018**	**1239**	**3320**	**209**	**12956**	**5511**
peer cone size																
BGP-derived	11	10	14	1	5	9	4	6	0	15	12	8	2	3	6	13
combined	14	2	7	3	5	4	11	10	1	16	12	13	9	6	8	15
customer cone size																
BGP-derived	0	1	2	3	4	5	6	8	9	10	11	13	15	17	29	32
combined	0	1	2	3	4	5	6	8	9	10	11	13	15	17	29	32
eccentricity																
BGP-derived	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
combined	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1
betweenness																
BGP-derived	1	0	11	10	16	22	27	14	3	9	4	25	46	5	71	154
combined	1	0	8	9	17	21	31	15	3	7	4	24	46	6	78	183
degree																
BGP-derived	1	0	36	17	32	64	158	28	3	7	4	56	85	8	236	261
combined	1	0	61	26	47	195	1059	44	3	10	4	165	300	11	1432	1435
coreness																
BGP-derived	0	151	0	0	151	197	233	442	0	562	760	562	223	442	487	562
combined	970	973	978	1316	1115	1334	1348	1506	978	1563	1563	1540	1335	1510	1510	1523
clustering																
BGP-derived	18986	19003	18764	18837	18849	18723	18377	18963	18989	19000	19013	18935	18698	18994	17977	18126
combined	21106	21127	20614	20896	20867	20751	20343	21120	21125	21150	21159	21091	20568	21146	20001	20111

- **soreness** and **clustering failed** to contain clique

CONCLUSION

- **peering links mostly added to "middle" ASes**
- **shortest "policy free" paths mostly unchanged**
- **eccentricity and betweenness closest to cone ranking**