

# COLLABORATION IN LINUX KERNEL MAILING LISTS

**Dawn M. Foster**

@geekygirldawn

dawn@dawnfoster.com

[fastwonderblog.com](http://fastwonderblog.com)

PhD Student  
University of Greenwich  
London, UK

Community and Open  
Source Consultant  
The Scale Factory





# WHOAMI

**Geek, traveler, reader**

**Community and open source consultant  
at The Scale Factory**

**20+ year tech career focused on  
community & open source  
(Intel, Jive, Puppet Labs, etc.)**

**PhD at the University of Greenwich  
researching Linux kernel (September)**




Photos by [Josh Bancroft](#), [Don Park](#)



# ABOUT MY RESEARCH

**Interviews with 16 Linux kernel developers**

**Quantitative analysis focused on  
kernel mailing list collaboration**




# 16 KERNEL INTERVIEWS

<b>Gender</b>	Men	13	Women	3						
<b>Maintainer</b>	Yes	13	No	3						
<b>Company HQ Region</b>	N. America	9	S. America	0	Europe	5	Asia	2		
<b>Participant Region</b>	N. America	9	S. America	1	Europe	5	Asia	1		
<b>Company Size (employees)</b>	< 100	2	100 - 1,000	2	1,000 - 10,000	6	10,000 - 100,000	4	100,000+	2



# LOCATION DOESN'T MATTER

"Similar time zones can be helpful because I can get a reply immediately. But it is not super important."




"The Linux community doesn't care where you're located, ever. You can be on the moon as long as you have a good internet connection."



# KERNEL FIRST, EMPLOYER SECOND

"In 5 years, I'm not sure if I'll be in [Company] or not, but I'm pretty sure I'll be working on the kernel."



"At the core I'm a Linux kernel guy. ... At some point, I'm probably going to have the inkling to try something else, and then ... I'll be a Linux kernel guy at the next place."



# TRUST AND FREEDOM

"I have considerable freedom to decide what to work on and which direction to take."



"There's no real direction, no one really telling me do this or do that, I just keep stuff running."

"They're fine with whatever as long as they have feature support."



# RELATIONSHIPS ARE IMPORTANT

"There are many community developers who I feel very comfortable with at a social level, ... who I look forward to spending time with at conferences. In some cases, I'd say they're genuine friendships."



"Having a functional social relationship makes it much easier to feel that asking them for a favor is justifiable and with a strong expectation that you'll be able to return that favor at some point in the future."



# WORK WITH COMPETITORS



"When we're dealing with other Linux kernel developers, we are dealing with them as kernel developers, rather than as competitors, per se."

"I think there's an effective social contract that you are willing to help people who work for competitors on the assumption that you'll get about the same amount of benefit from them."



# COMPANY AFFILIATION

"Whether they are fresh out of school in Hungary or whether they've been working for Google for 20 years or somewhere else. Personally, I don't really care."

"If I don't know them personally or if they don't use their work email, I don't necessarily know."



"I would be a bit more forgiving and give them maybe a little bit more kid glove treatment if I knew they were not being paid to do it."



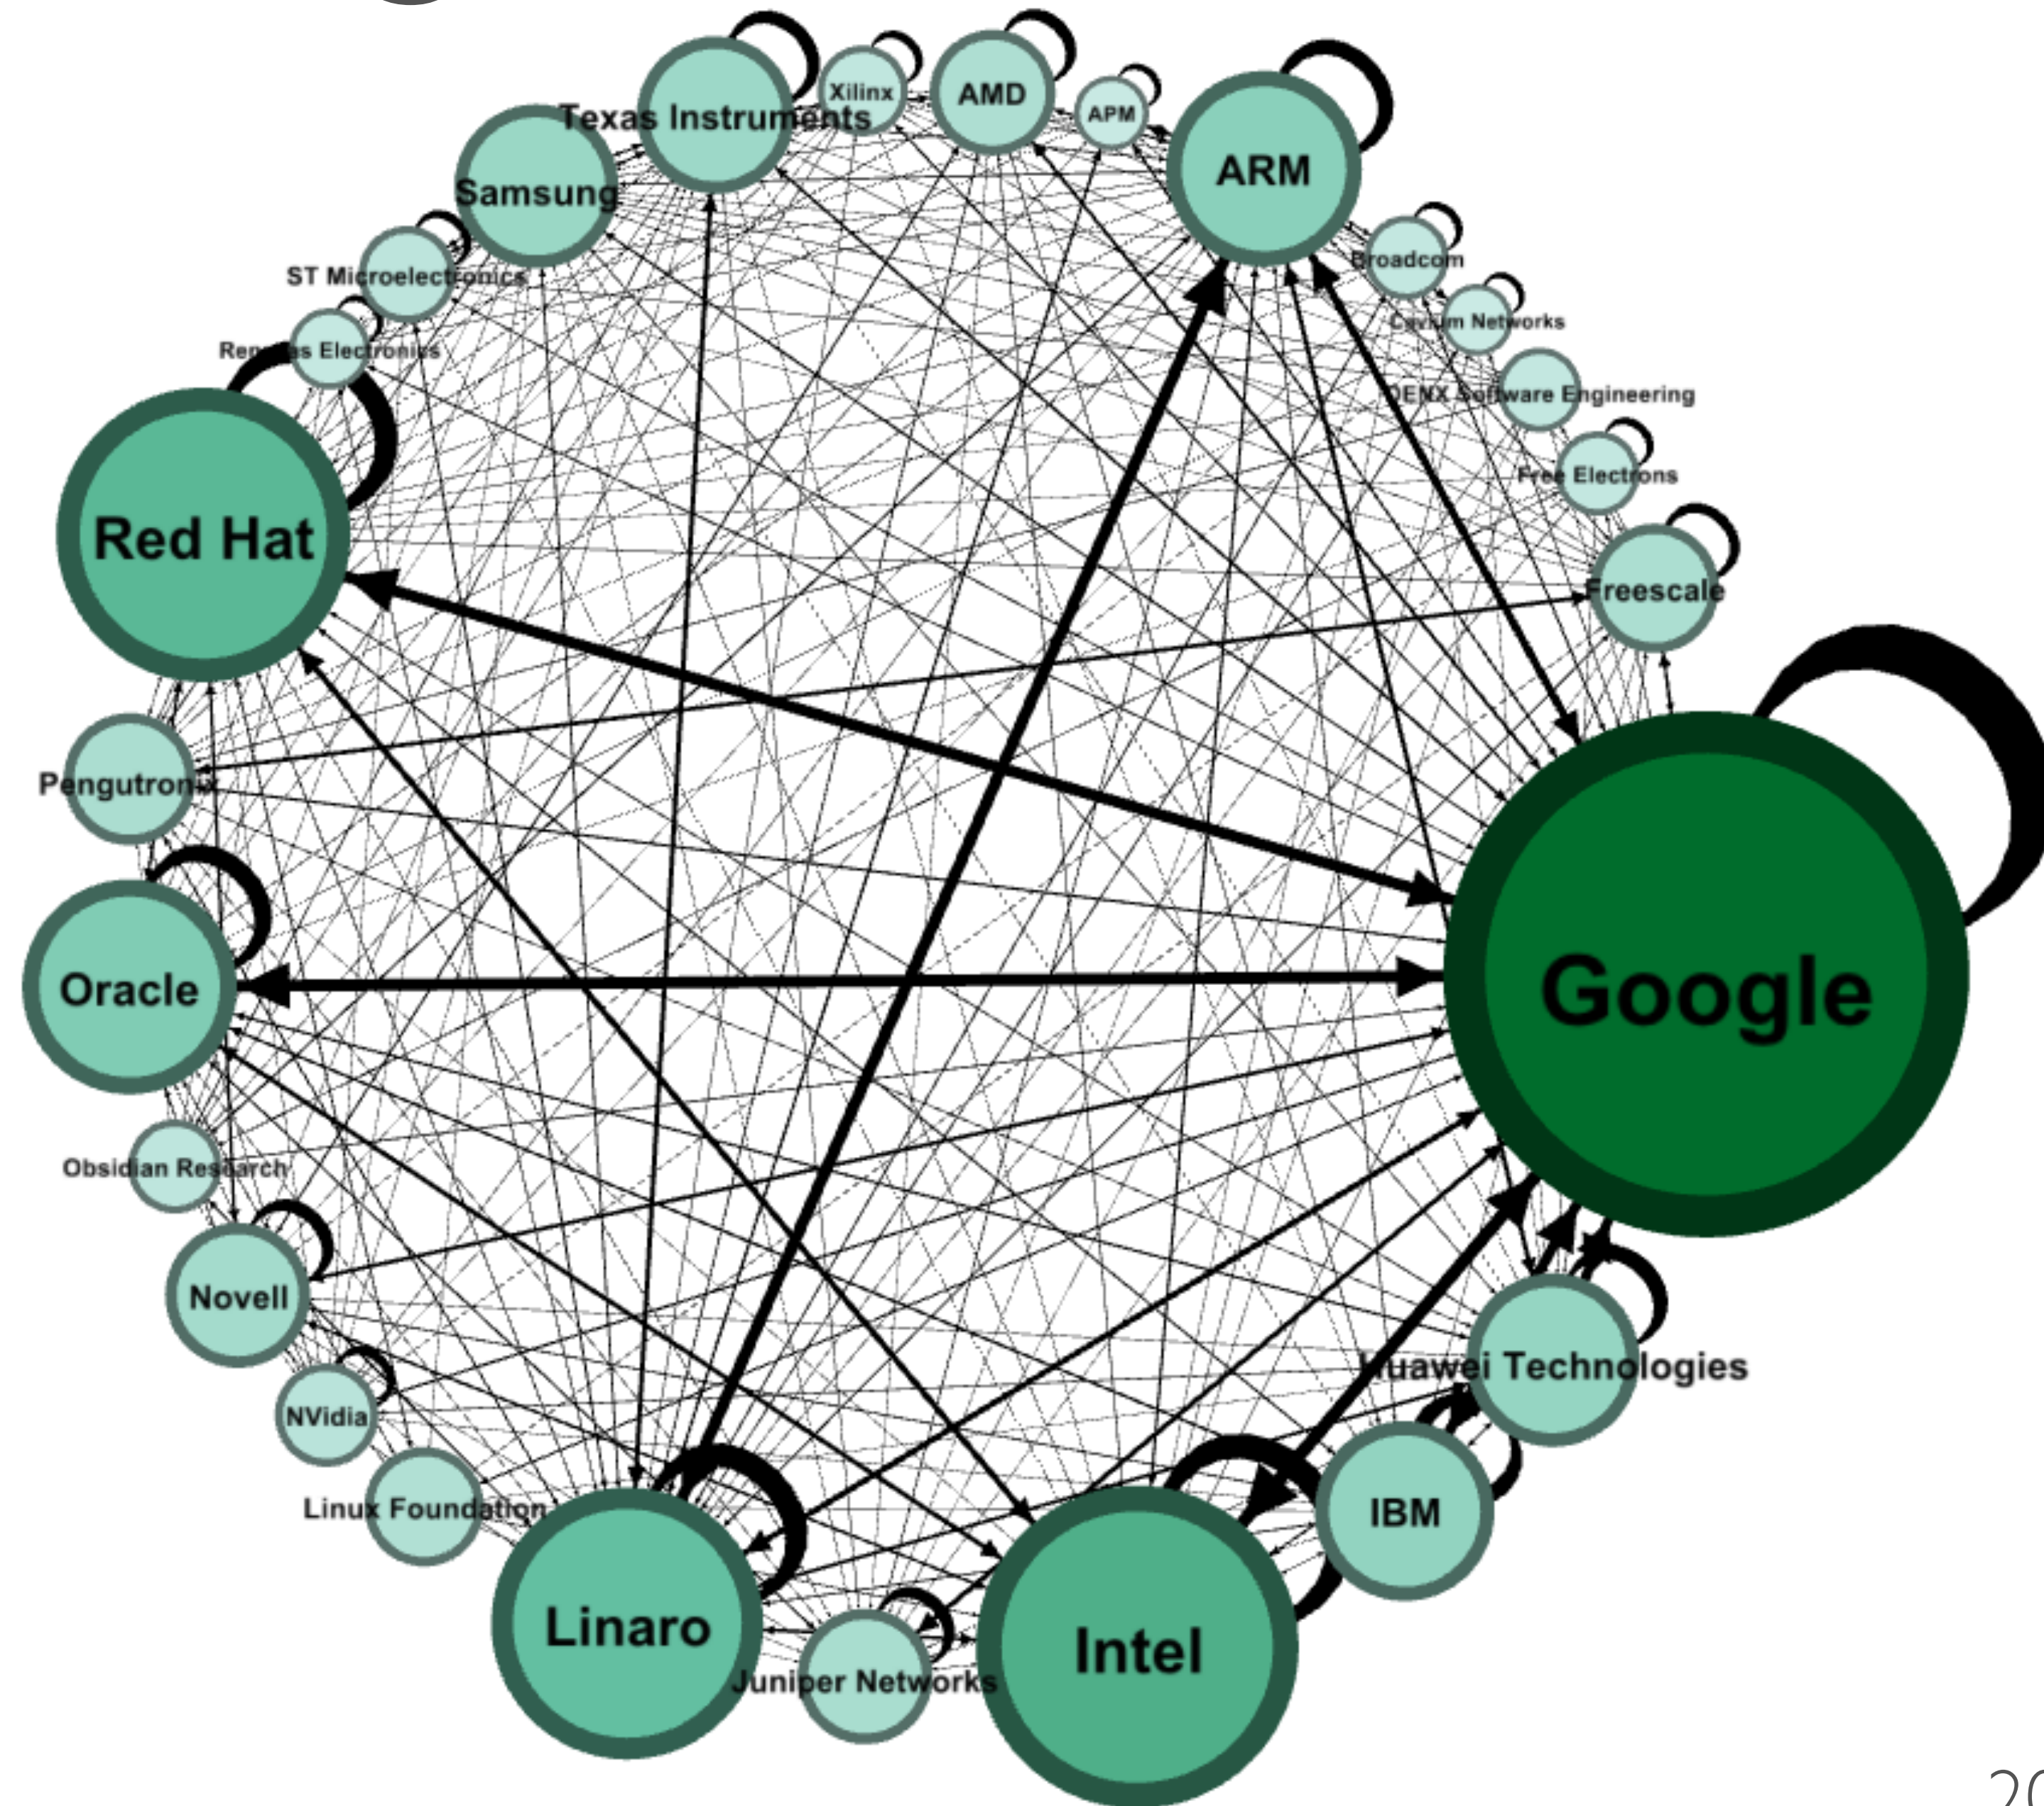
# MAILING LISTS

**Subsystem lists vs. LKML**  
**VGER vs. MAINTAINERS**





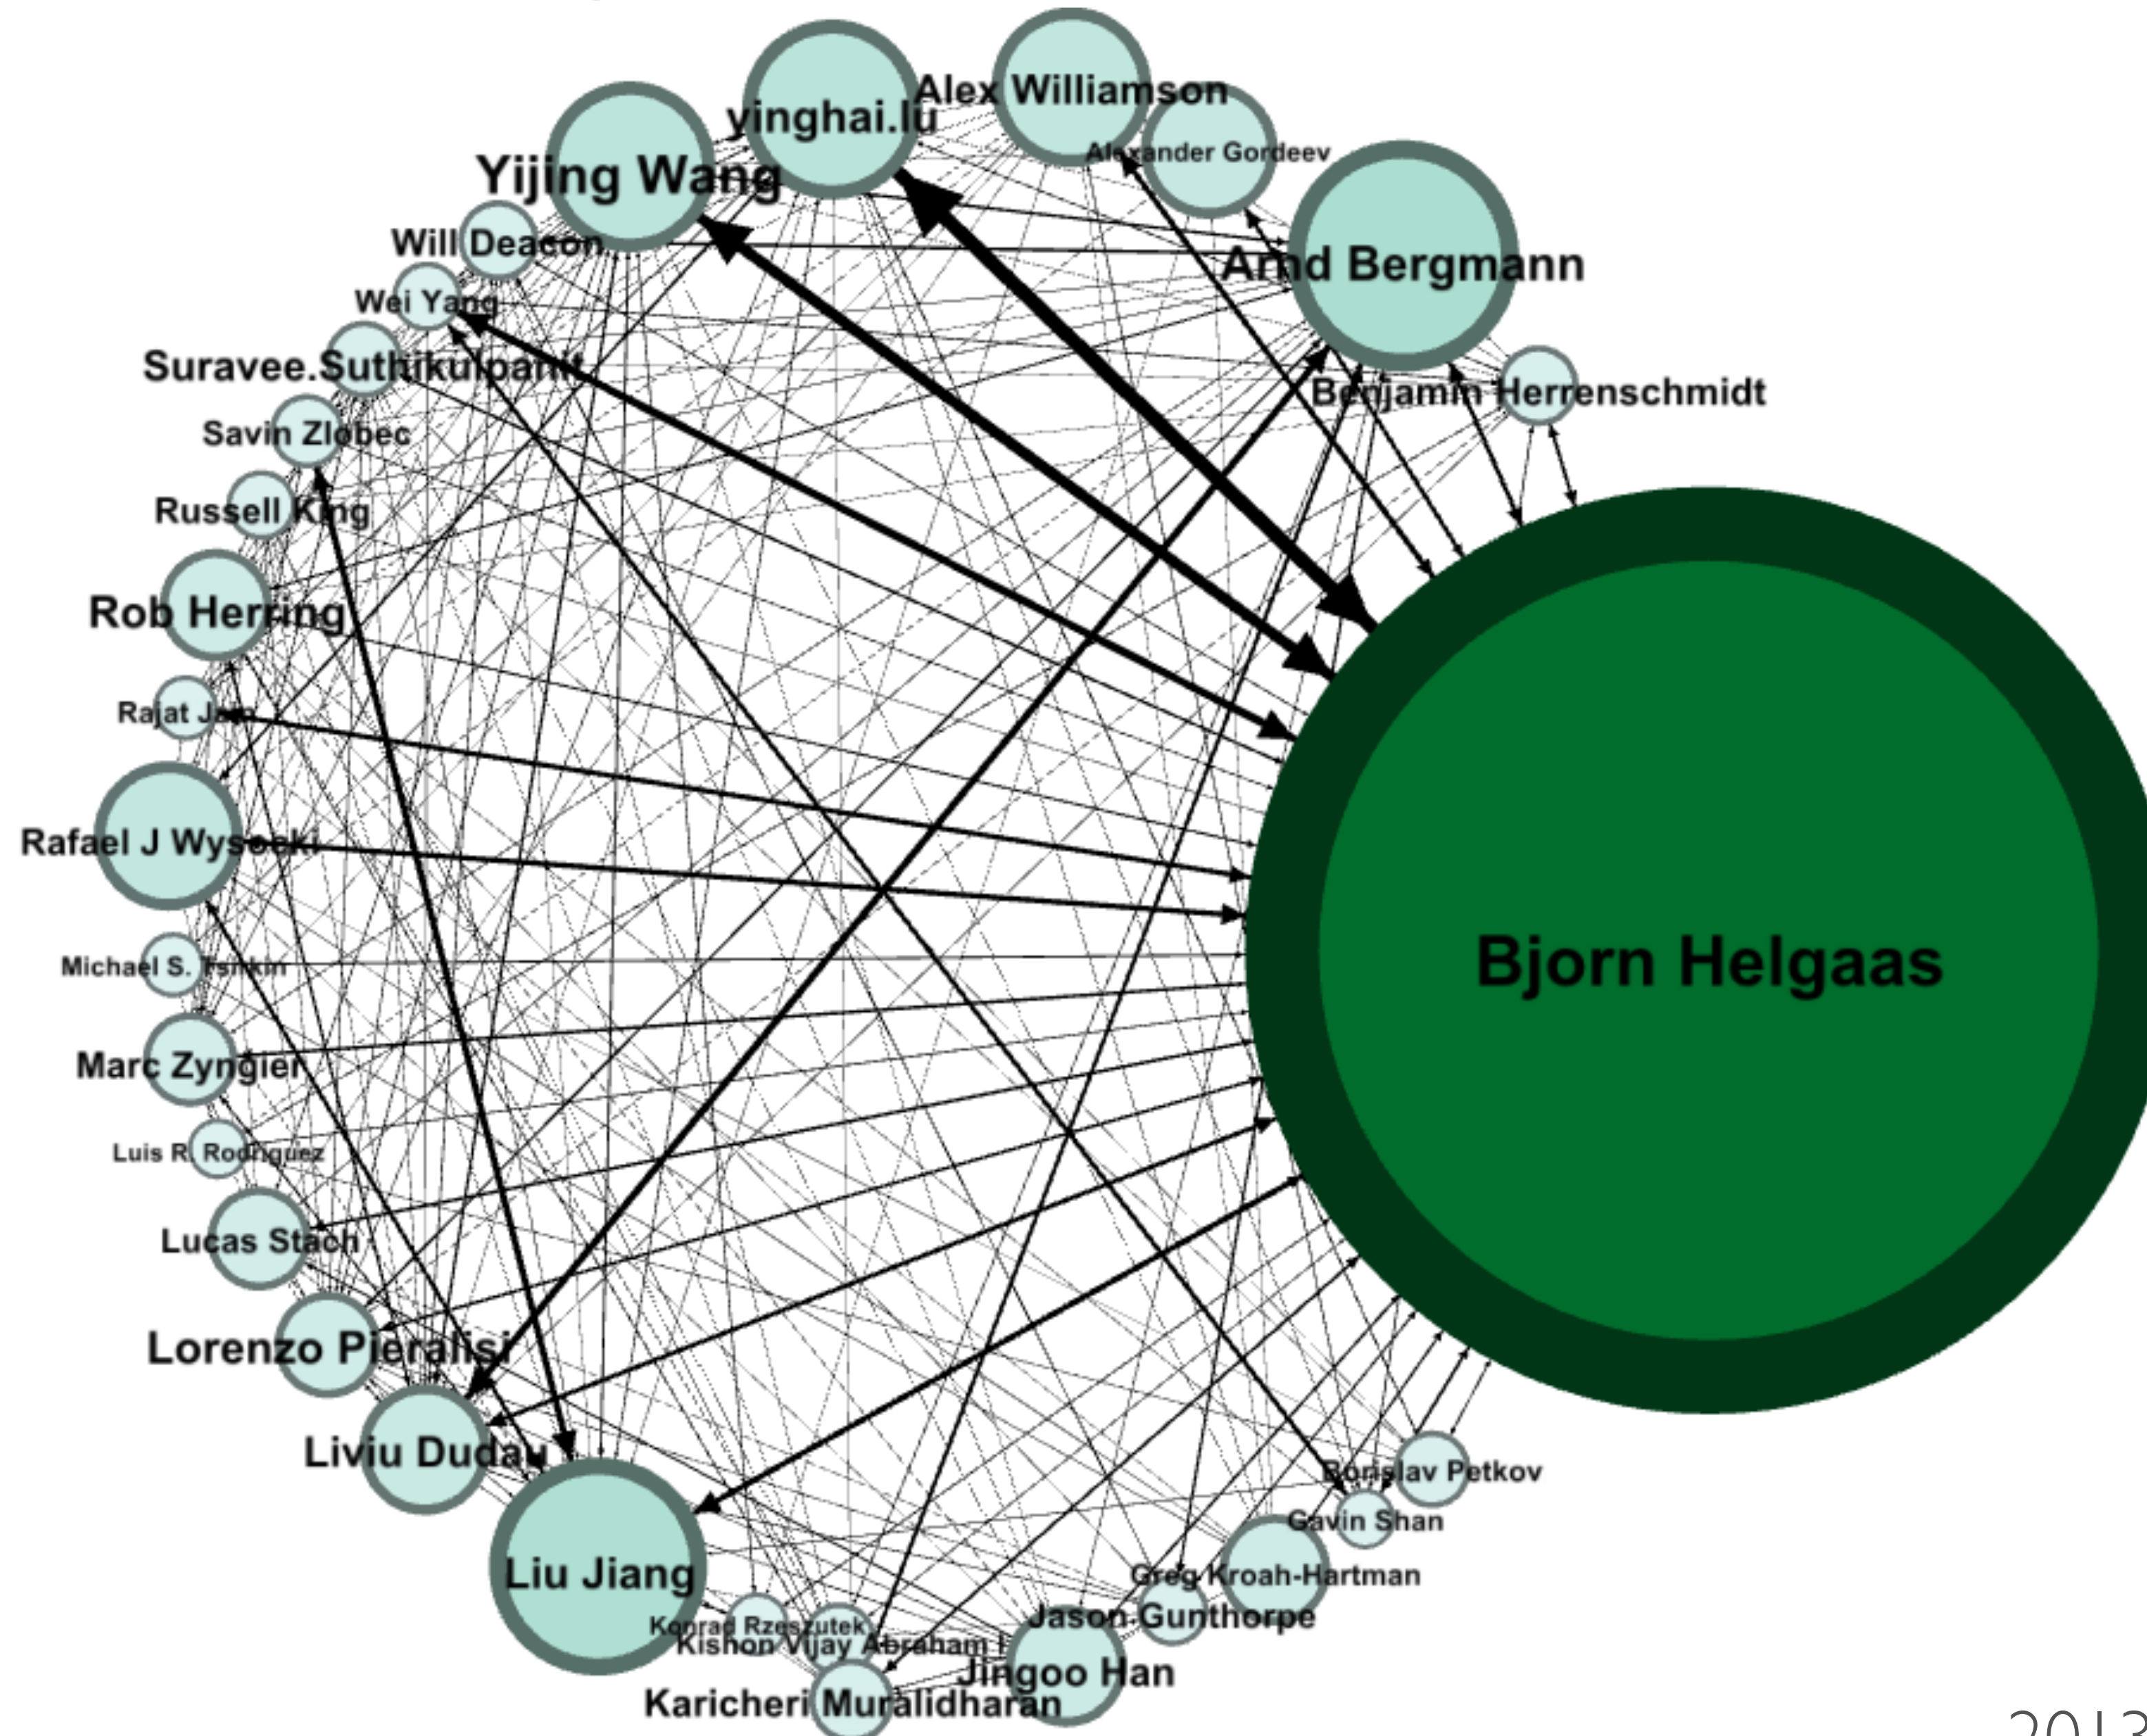
# LINUX-PCI@VGER.KERNEL.ORG COMPANIES



2013-10-31 to 2015-10-31

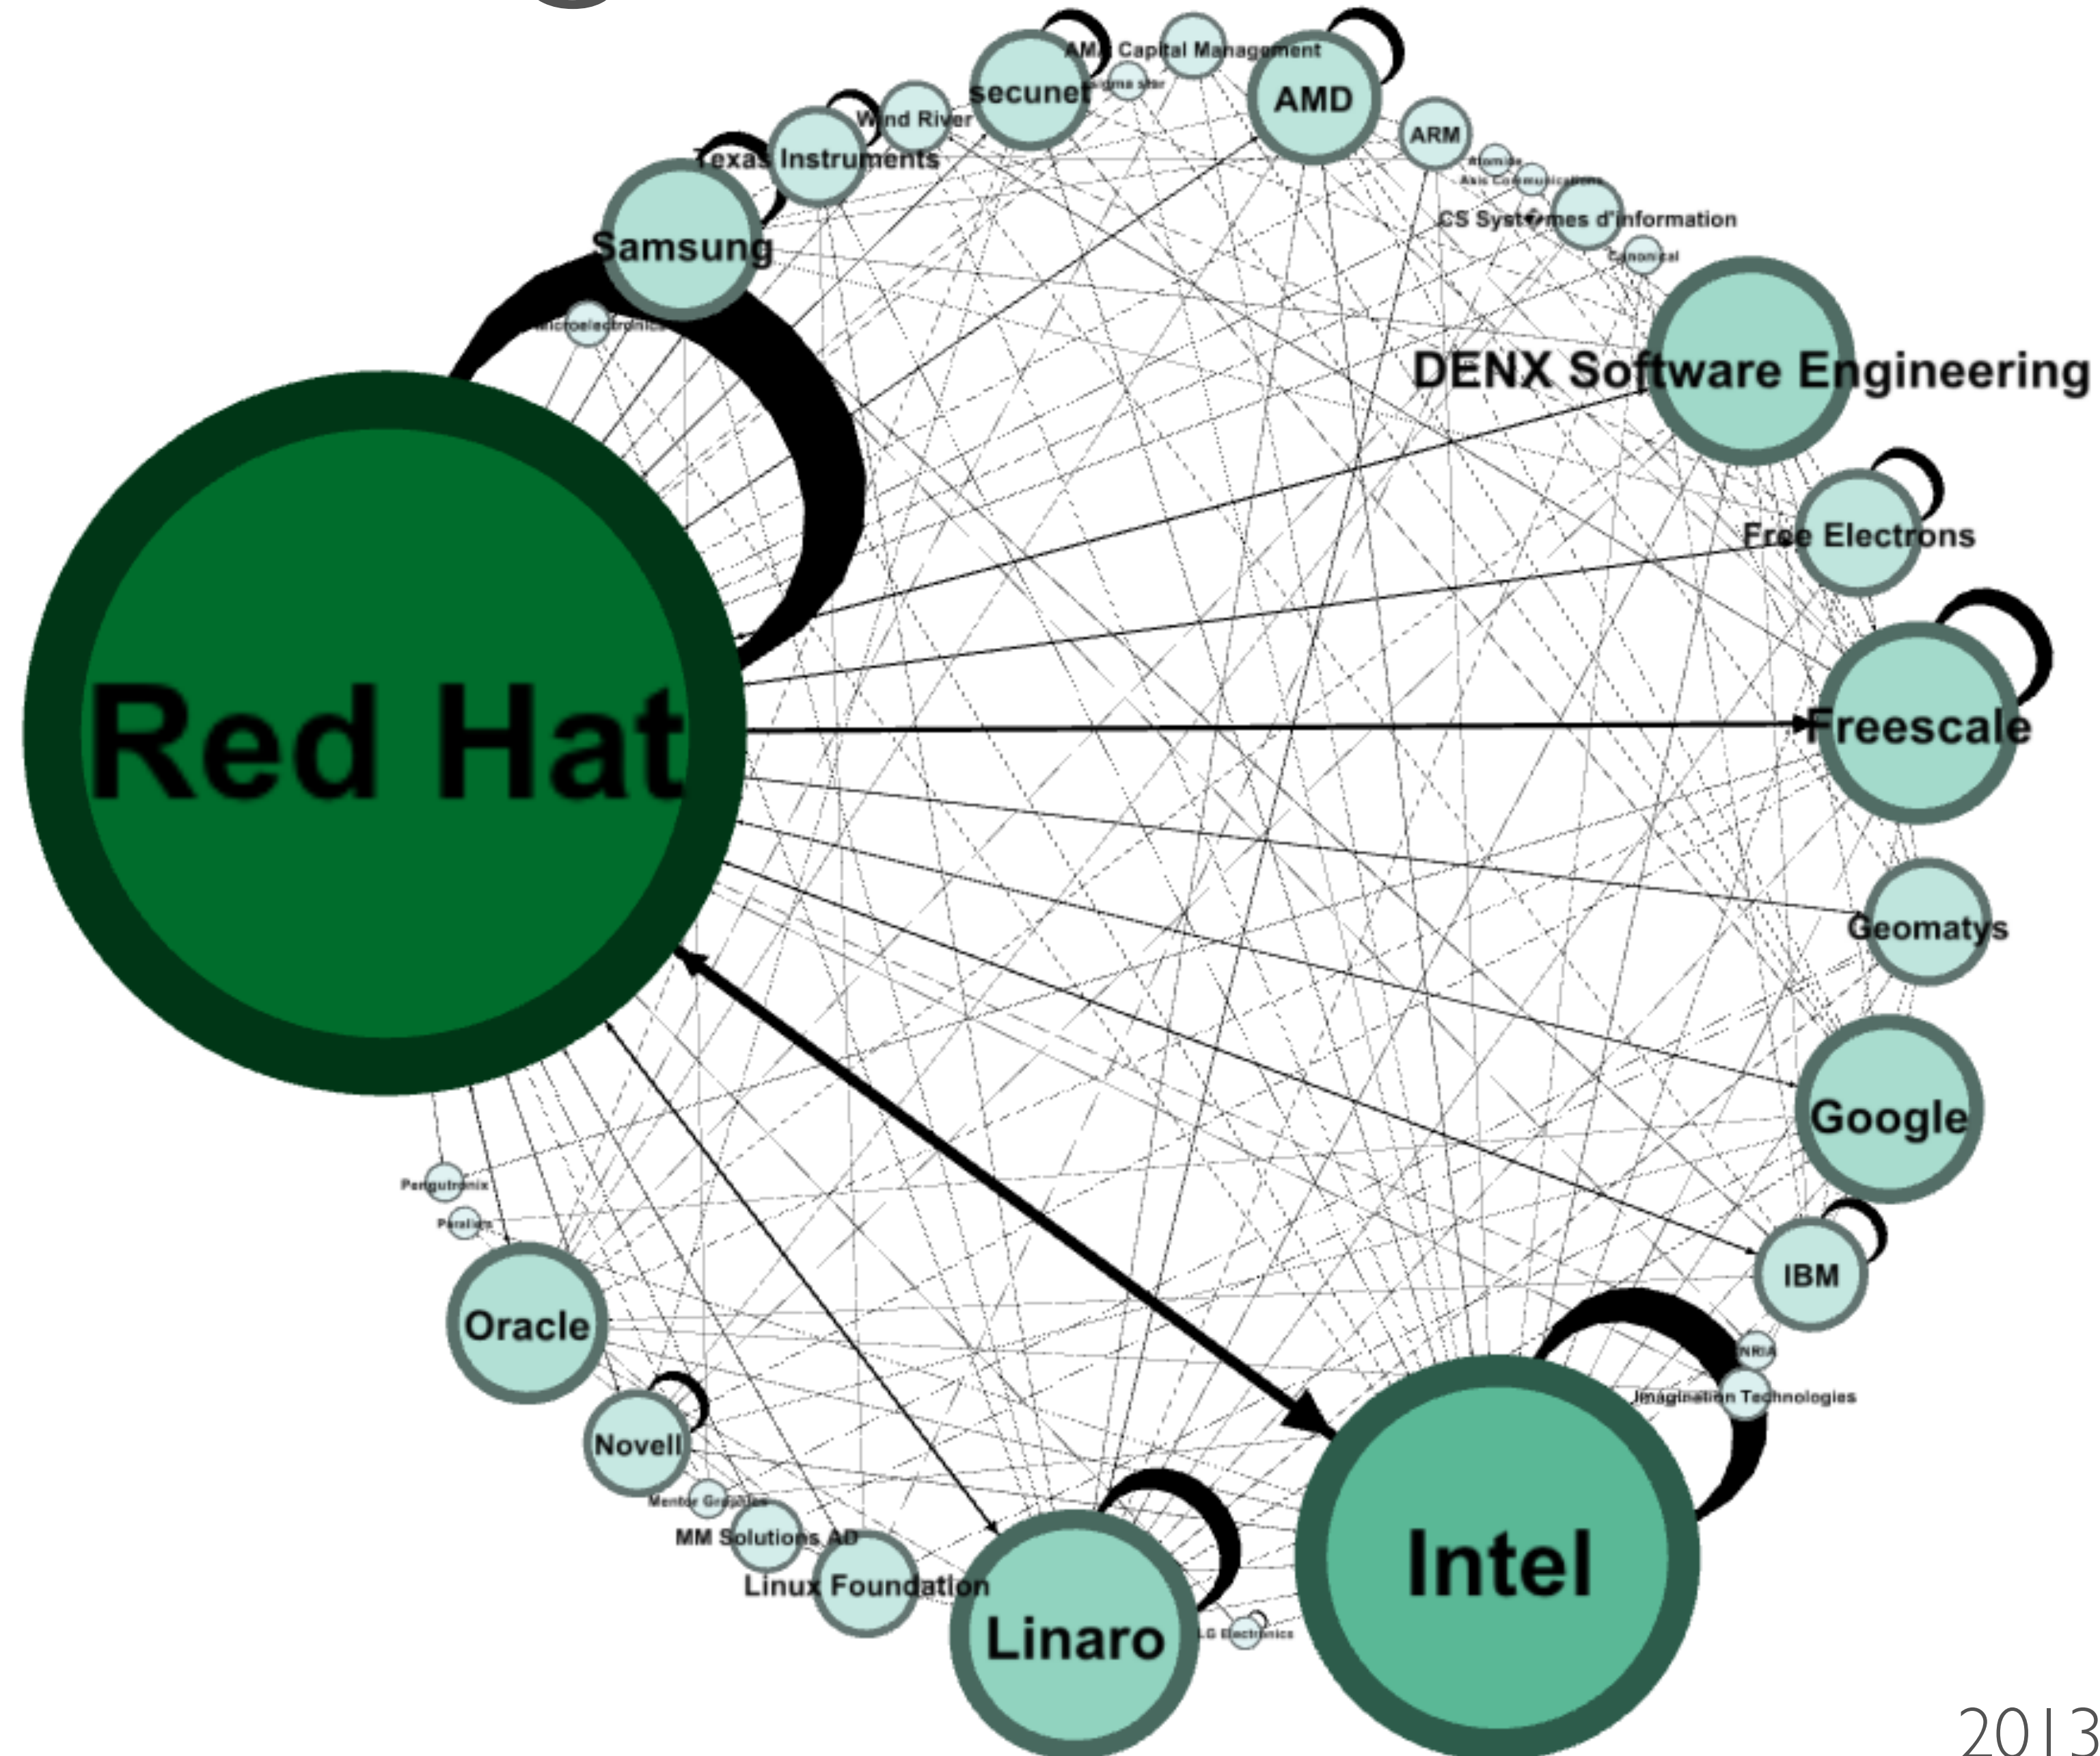


# LINUX-PCI@VGER.KERNEL.ORG PEOPLE



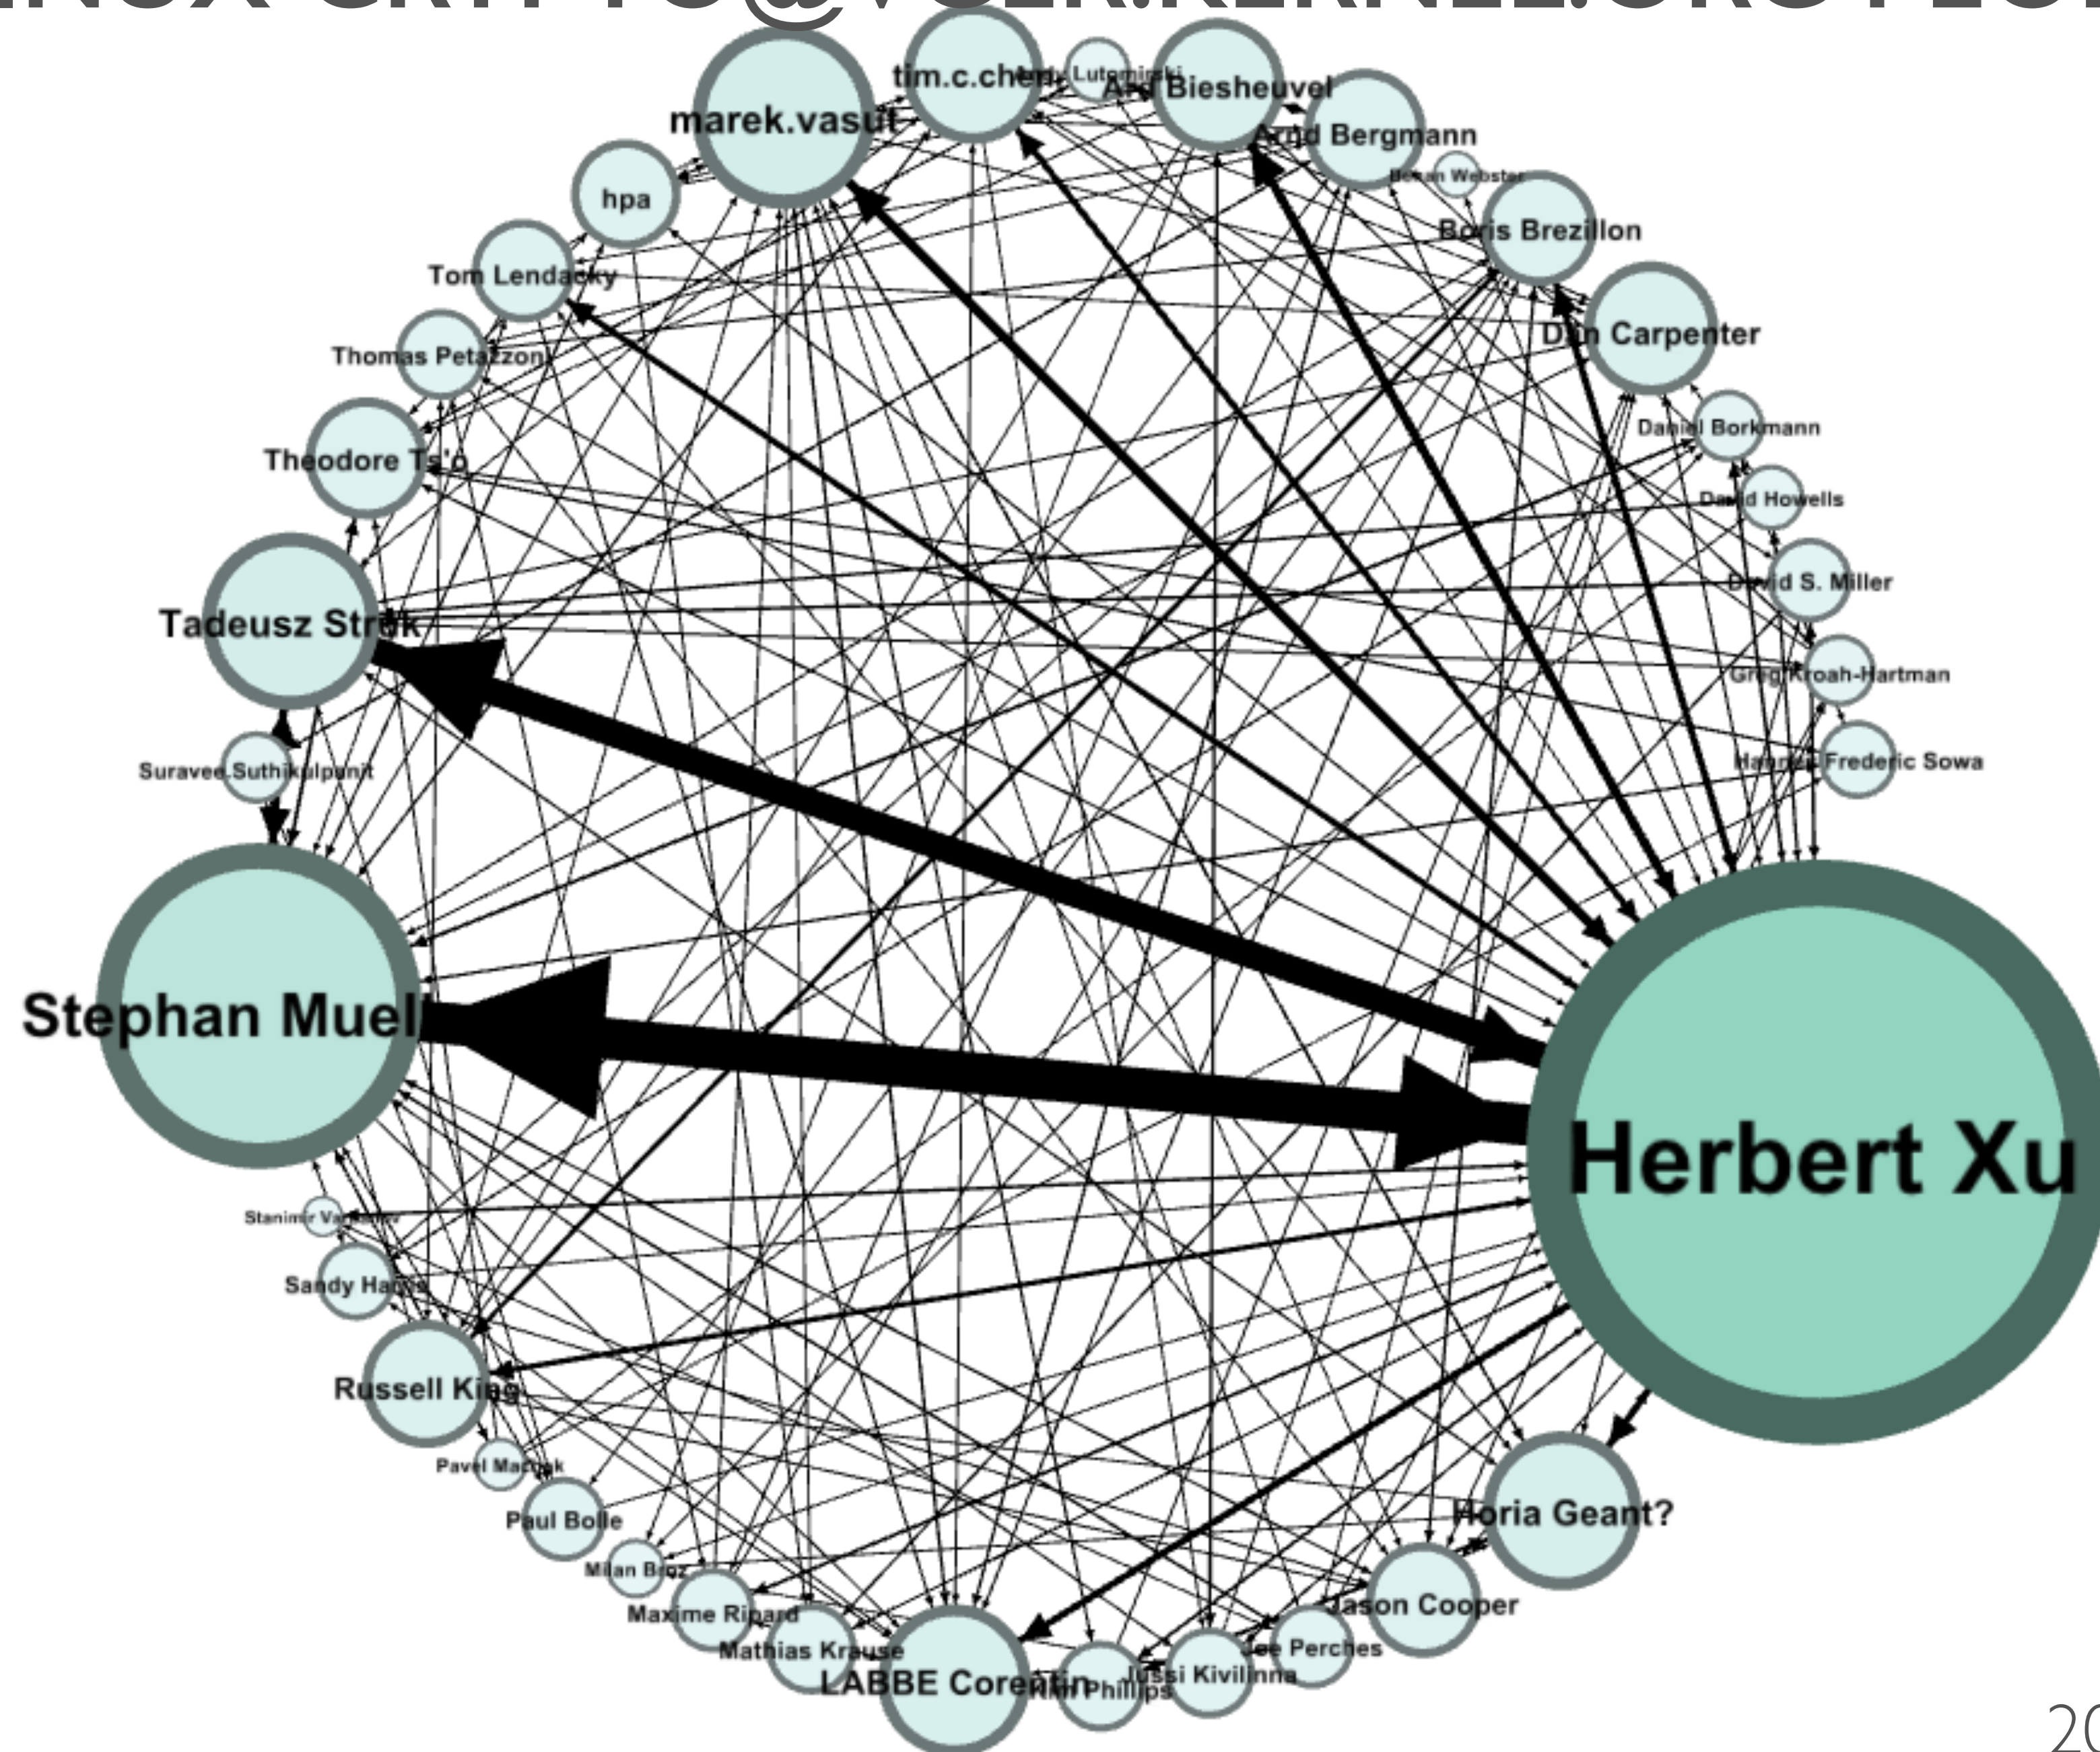


# LINUX-CRYPTO@VGER.KERNEL.ORG COMPANIES





# LINUX-CRYPTO@VGER.KERNEL.ORG PEOPLE





# PCI ML STATISTICS (AKA MAGIC)

Variables	Coef	Std Err	Sig
Alter Maintainer	-0.044	(0.014)	**
Either Maintainer	0.161	(0.075)	*
Alter Committer	-0.253	(0.061)	***
Either Committer	0.638	(0.148)	***
Ego To CC	2.597	(0.685)	***
Geographic Proximity	0.152	(0.088)	.
Organizational Proximity	0.944	(0.176)	***
Institutional Proximity	-0.078	(0.048)	
Social Proximity	1.072	(0.242)	***
Social Proximity Squared	-0.050	(0.012)	***
Cognitive Proximity	4.528	(0.920)	***
Cognitive Proximity Squared	-1.991	(0.513)	***
Repeated Effect	0.006	(0.002)	**
Participation Shift	-0.088	(0.156)	
Recency Effect	0.760	(0.217)	***
Transitive Closure	0.025	(0.008)	**
Cyclic Closure	0.078	(0.018)	***
Shared Partnership Inbound	-0.066	(0.016)	***
Shared Partnership Outbound	-0.126	(0.031)	***
Social x Cognitive Prox	-0.613	(0.134)	***
Organizational x Social Prox	-0.269	(0.056)	***
Organizational x Cog Prox	0.582	(0.286)	*
Geographic x Social Prox	0.153	(0.048)	**
Geographic x Cog Prox	-2.725	(0.738)	***
Sig p < 0.001 '***' p < 0.01 '**' p < 0.05 '*' p < 0.1 '.'			



# PCI MAILING LIST RESULTS

- People in TO / CC *much* more likely to reply
- People who commit to same areas of source code more likely to reply
- Maintainers / committers more likely to reply, less likely to be replied to

```
PCI SUBSYSTEM
M: Bjorn Helgaas <bhelgaas@google.com>
L: linux-pci@vger.kernel.org
Q: http://patchwork.ozlabs.org/project/linux-pci/list/
T: git git://git.kernel.org/pub/scm/linux/kernel/git/helgaas/pci.git
S: Supported
F: Documentation/devicetree/bindings/pci/
F: Documentation/PCI/
F: drivers/pci/
F: include/linux/pci*
F: arch/x86/pci/
F: arch/x86/kernel/quirks.c

PCI ENDPOINT SUBSYSTEM
M: Kishon Vijay Abraham I <kishon@ti.com>
L: linux-pci@vger.kernel.org
T: git git://git.kernel.org/pub/scm/linux/kernel/git/kishon/pci-endpoint.git
S: Supported
F: drivers/pci/endpoint/
F: drivers/misc/pci_endpoint_test.c
F: tools/pci/

PCI DRIVER FOR ALTERA PCIE IP
M: Ley Foon Tan <lftan@altera.com>
L: rfi@lists.rocketboards.org (moderated for non-subscribers)
L: linux-pci@vger.kernel.org
S: Supported
F: Documentation/devicetree/bindings/pci/altera-pcie.txt
F: drivers/pci/host/pcie-altera.c


PCI DRIVER FOR ARM VERSATILE PLATFORM
M: Rob Herring <robh@kernel.org>
L: linux-pci@vger.kernel.org
L: linux-arm-kernel@lists.infradead.org
S: Maintained
F: Documentation/devicetree/bindings/pci/versatile.txt
F: drivers/pci/host/pci-versatile.c
```



# PCI MAILING LIST RESULTS

**More likely to reply:**

- **Work for same organization**
- **Previous interactions in same mailing list threads**





# PCI MAILING LIST RESULTS

**Not significant:**

- **Physical location (time zone)**
- **Type of organization (company, non-profit, hobby, academic)**





# NEXT STEPS

**Submit dissertation corrections**

**Get real job again**





# THANK YOU



**Dawn Foster**

**@geekygirl dawn**

**dawn@fastwonder.com**

**fastwonderblog.com**

**PhD Student**

**Centre for Business Network Analysis  
University of Greenwich**

**Consultant**

**Open Source / Community Building  
The Scale Factory**