

RESEARCH PROJECT

TAKING A CLOSER LOOK AT IRATI

Koen Veelenturf
University of Amsterdam
System and Network Engineering

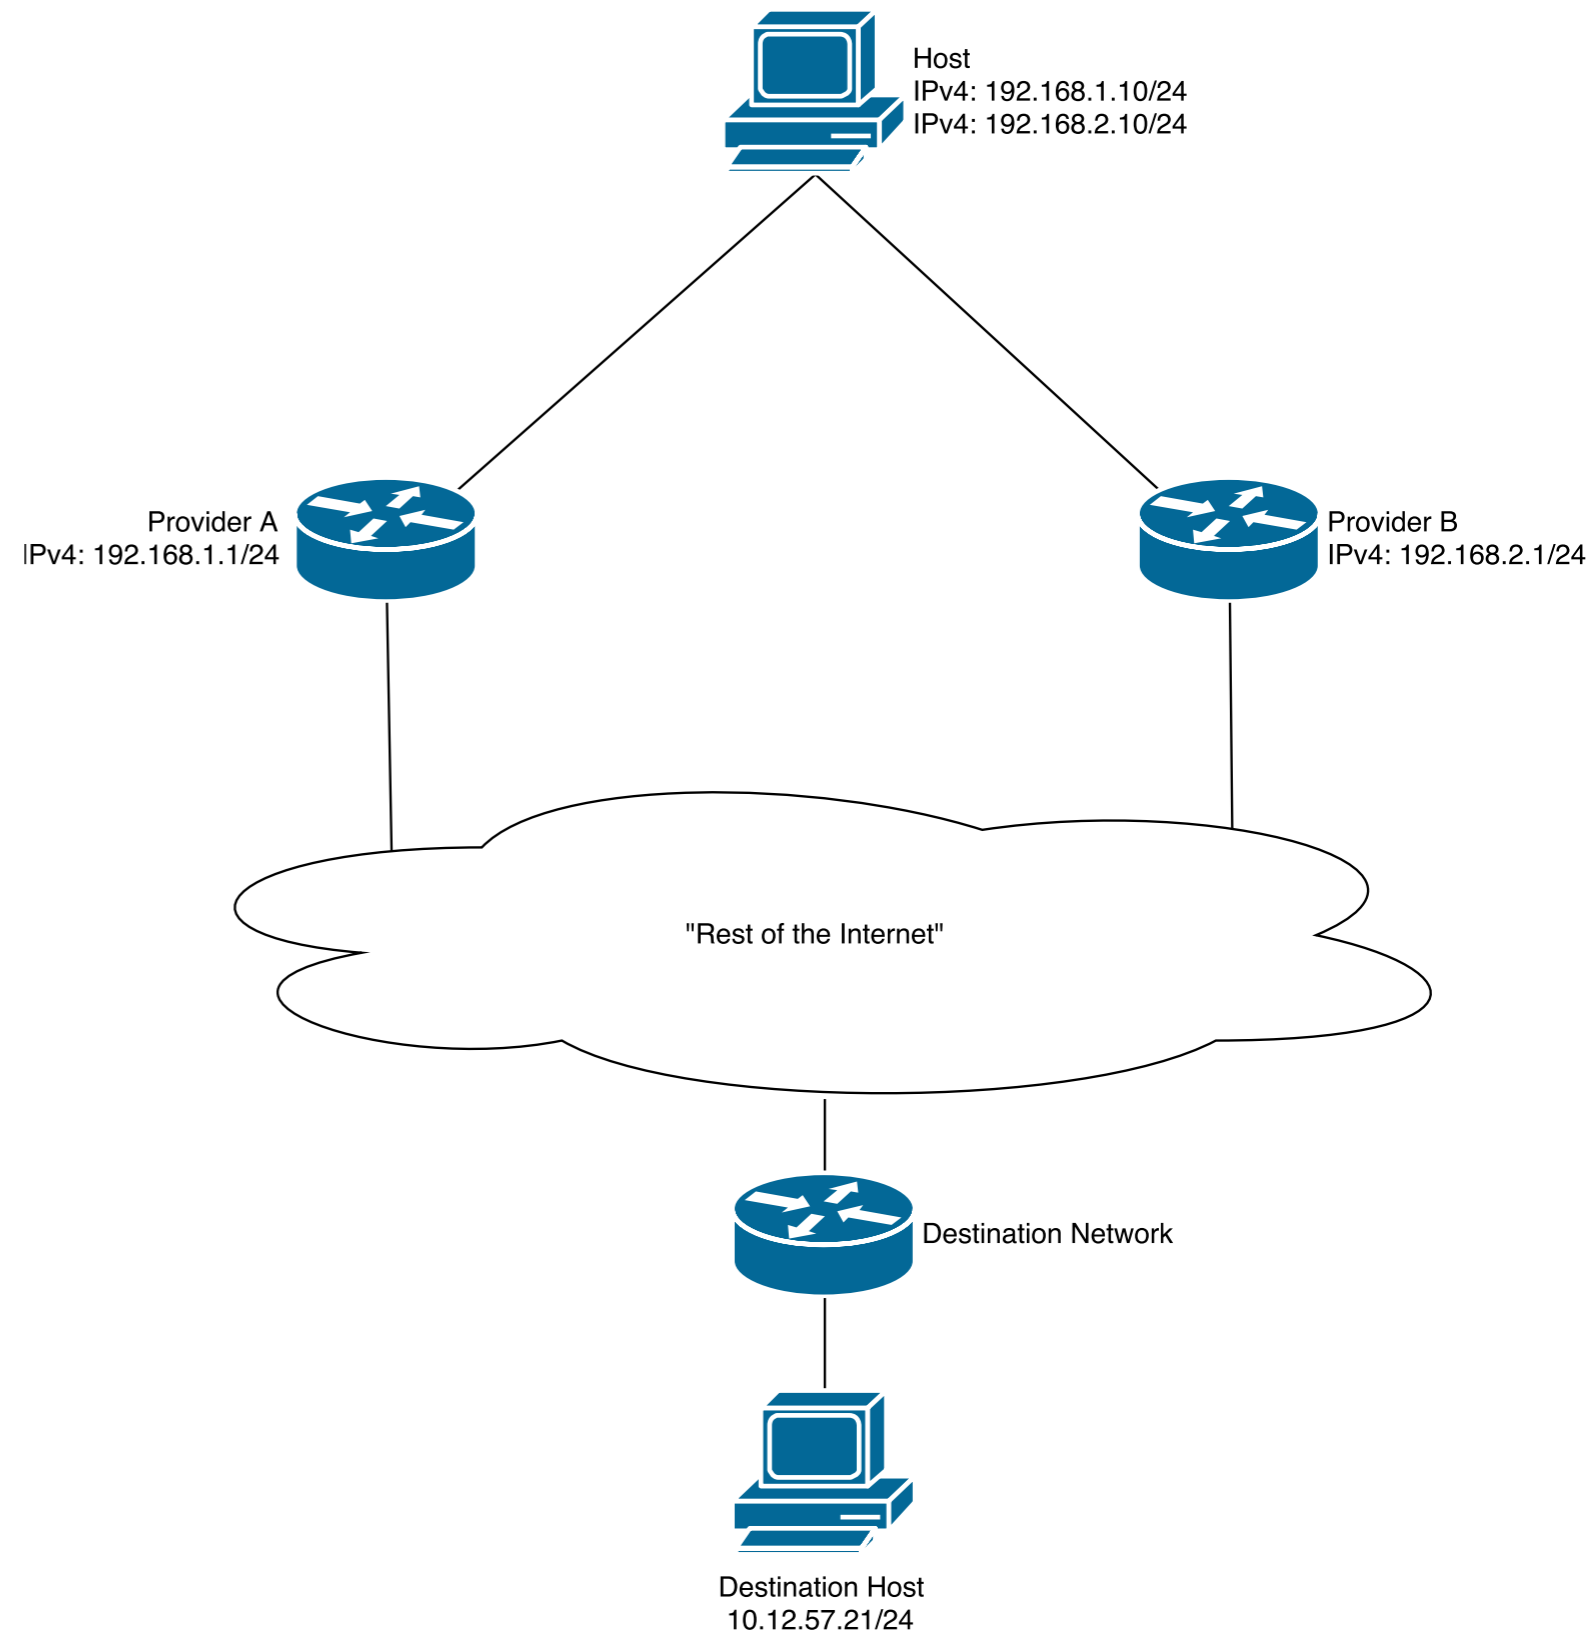
Supervisors:
Marijke Kaat - SURFnet
Ralph Koning - UvA

RESEARCH QUESTIONS

- ▶ How does RINA solve the multihoming problem?
 - ▶ What are the problems with multihoming and mobility in the current Internet?
 - ▶ What kind of solutions are proposed to solve the multihoming/mobility problem in the current Internet?
 - ▶ To what extent is multihoming/mobility implemented in the IRATI implementation?

MULTIHOMING

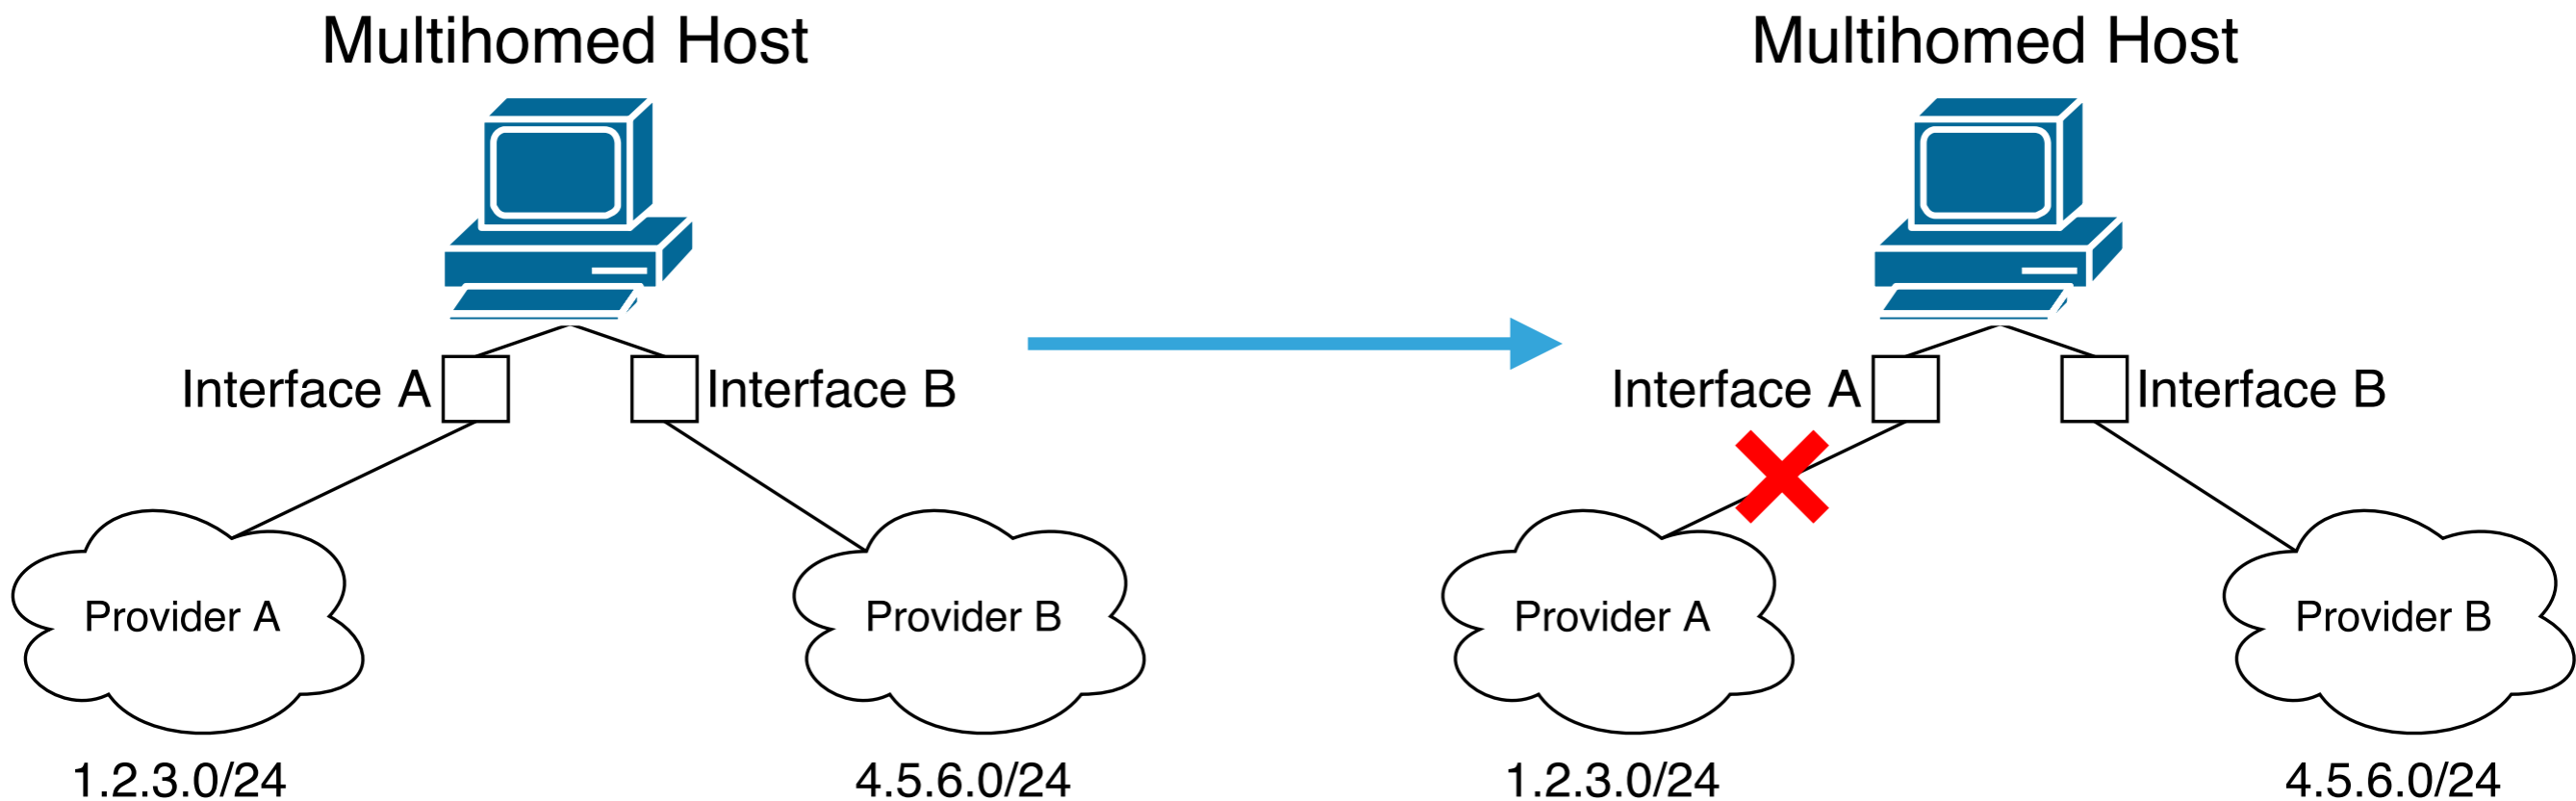
- ▶ *The practise of connecting a host or a computer to more than one network*
 - ▶ *Creating reliability, performance improvements, or to reduce cost*
- ▶ **Mobility**
 - ▶ *“Special kind of multihoming”*



A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM

- ▶ 1972: Tinker Air Force Base joined ARPANET
 - ▶ Connections to two IMPs for redundancy
 - ▶ Problems with routing algorithm
 - ▶ Naming the interface, not the node

A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM



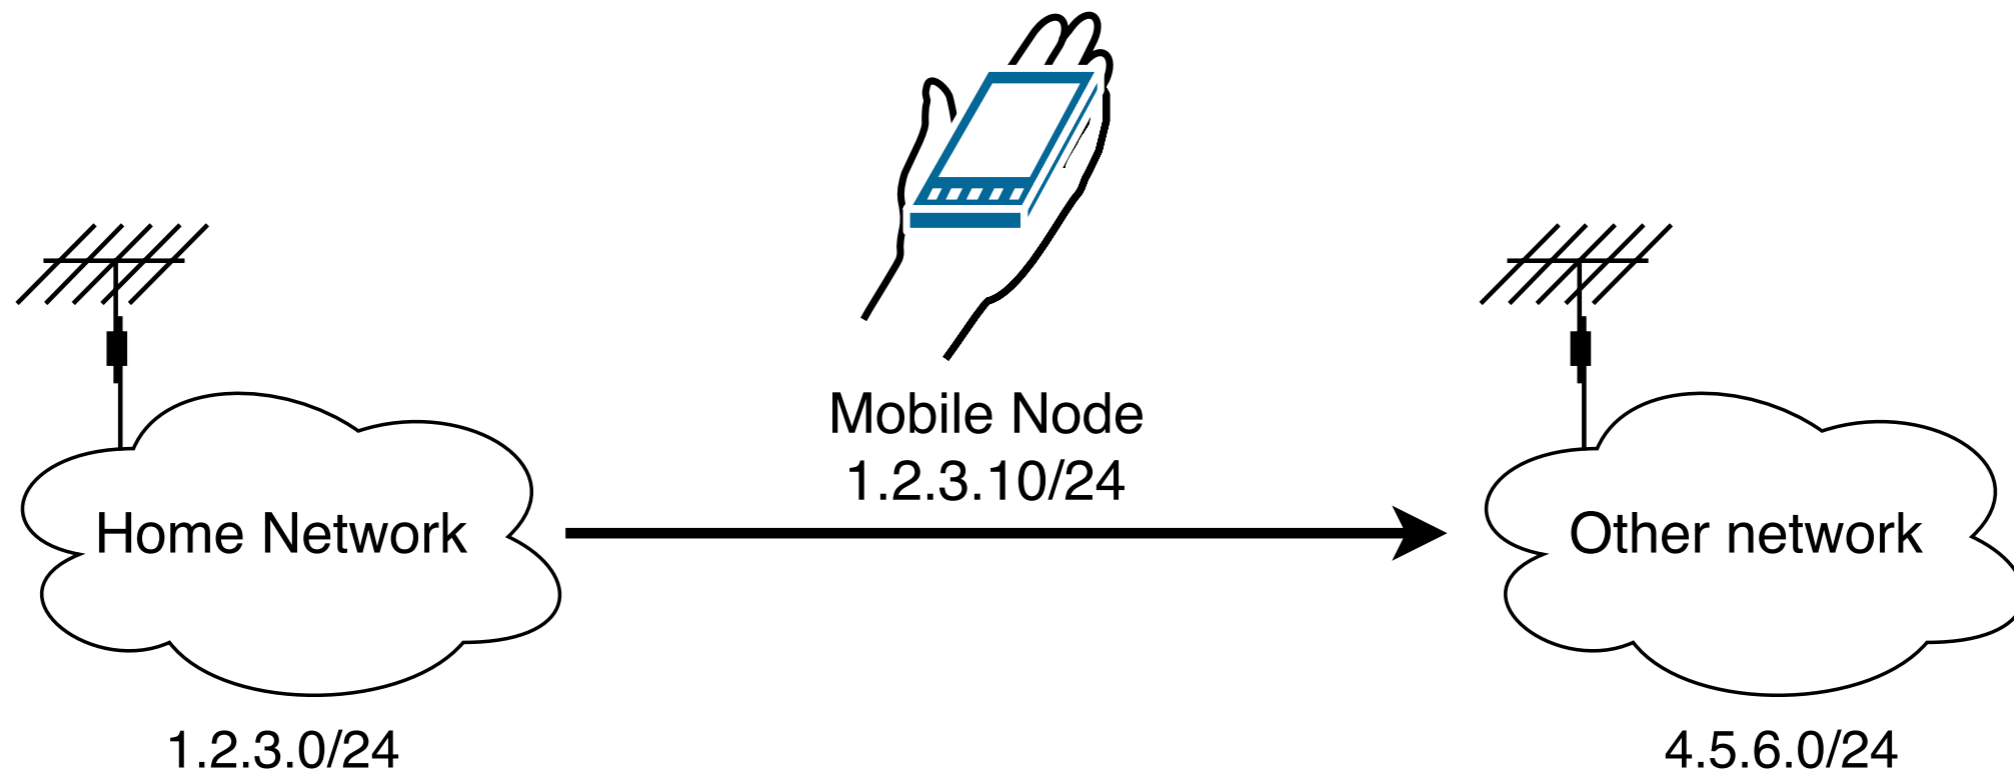
► What happens when one of the interfaces fails?

A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM




- ▶ Mobile Node connected to its own "Home Network"

A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM



- ▶ Mobile Node starts moving around

A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM

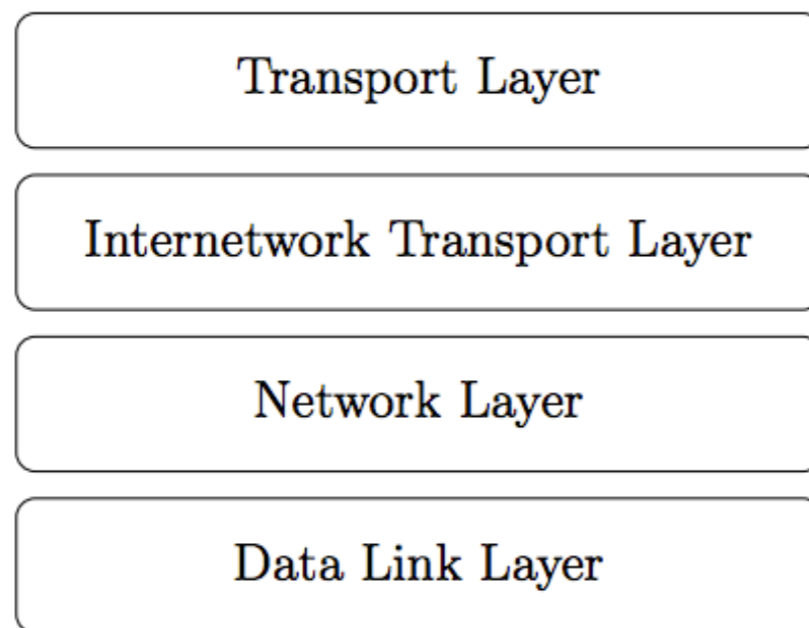


- ▶ Mobile Node connects to “Other Network”. What happens?

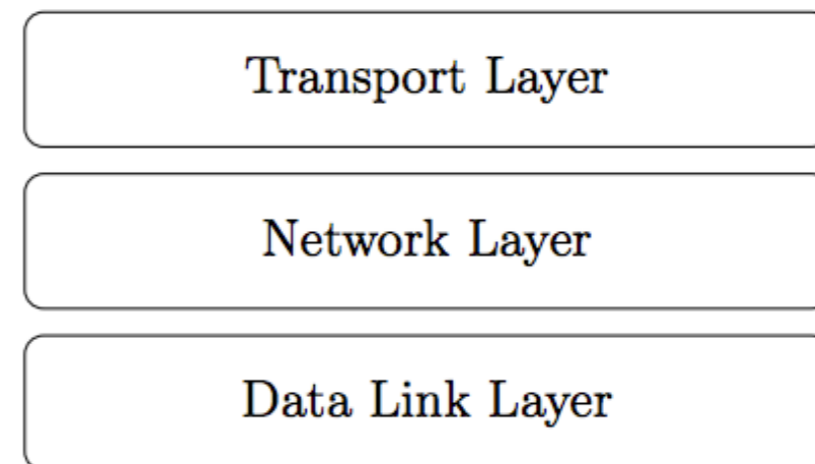
A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM

- ▶ 1978: TCP/IP was split
 - ▶ TCP - Creating segments
 - ▶ IP - Transmitting individual segments
- ▶ Big consequences for the Internet
 - ▶ Layers are dependent on each other

A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM



(a) The Internet layers before getting rid of NCP




(b) The Internet layers after getting rid of NCP

- ▶ Before 1983: TCP/IP running over NCP
- ▶ 1st January 1983 - Flag Day: TCP/IP takes over!

A BRIEF INTRODUCTION OF THE MULTIHOMING PROBLEM

- ▶ 1982: Jerry Saltzer describes naming and binding of network destinations
 - ▶ Left: reality; Right: ideally



<http://rina.tssg.org/docs/FutureNetTutorialPart2-100415.pdf>

CURRENT “SOLUTION” TO THE MULTIHOMING/MOBILITY PROBLEM

- ▶ Multihoming IPv4/IPv6:
 - ▶ Multiple physical connections to two different providers
- ▶ Multihoming IPv6:
 - ▶ SHIM6
- ▶ Other solutions:
 - ▶ MPTCP
 - ▶ SCTP
 - ▶ HIP
- ▶ Mobility:
 - ▶ Mobile IP
 - ▶ Mobile IPv6
 - ▶ LISP

PROBLEMS WITH CURRENT SOLUTIONS

- ▶ More complexity
- ▶ Decrease efficiency
- ▶ Does not actually solve the cause of the problem
- ▶ Will not scale

RINA

Website: PRESENTATION AND MAJOR DESIGN ASPECTS OF THE CYCLADES COMPUTER NETWORK, Louis POUZIN. <http://www.rogermoore.ca/PS/CYCLB.html> en <http://www.cs.utexas.edu/users/chris/nph/>

Paper: On Distributed Communications Networks. Paul Baran. IEEE transactions on communications systems. March 1964.

Paper: SOME PERSPECTIVES ON NETWORKS—PAST, PRESENT AND FUTURE. Paul Baran. 1977 IFIP Congress Proceedings, 1977.

References site: <http://iri.cs.ucla.edu/references.html>:

2008 Constantine Dovrolis, "What would Darwin Think about Clean-Slate Architectures?," CCR, the Computer Communication Review, January 2008.

2007 Ying Zhang, Zheng Zhang, Morley Z. Mao, Charlie Y. Hu, and Bruce M. Maggs, "On the Impact of Route Monitor Selection," IMC, October 2007.

2002 David Clark et al, "Tussle in cyberspace: defining tomorrow's Internet," SIGCOMM 2002, September 2002.

1998 David S. Isenberg, "The Dawn of the Stupid Network," ACM netWorker, February 1998.

1996 S. Deering and R. Hinden, "IPv6 Metro Addressing," INTERNET-DRAFT, March 1996.

Steve Deering, "The Map & Encap Scheme for scalable IPv4 routing with portable site prefixes," Presentation Xerox PARC, March 1996.

1977 Paul Baran, "Some Perspectives on Networks - Past Present and Future," IFIP, January 1977.

1964 Paul Baran, "On Distributed Communications Networks," IEEE Transactions of the Professional Technical Group on Communications Systems, January 1964.

1928 J. B. S. Haldane, "On Being the Right Size," Grabbed from the reading list of a course Prof. Kurose, January 1928.

Paper: THE END OF THE END-TO-END ARGUMENT, David P. Reed April, 2000.

Paper: A critical review of "End-to-end arguments in system design", Tim Moors, Polytechnic University, 2000.

Paper: A Protocol for Packet Network Intercommunication. VINTON G. CERF AND ROBERT E. KAHN. IEEE Trans on Comms. 1974.

[IEN-1] Issues in the interconnection of datagram networks. University College London. <<http://www.rfc-editor.org/ien/ien1.pdf>>. July 1977.

[IEN19] Shoch, J., "Inter-Network Naming, Addressing, and Routing" <<http://www.rfc-editor.org/ien/ien19.txt>>. January 1978.

[IEN23] Cohen, D., "On Names, Addresses, and Routings", <<http://www.rfc-editor.org/ien/ien23.pdf>>. January 1978.

[IEN31] Cohen, D., "On Names, Addresses, and Routings (II)". <<http://www.rfc-editor.org/ien/ien31.pdf>>. April 1978.

Paper: The New Routing Algorithm for the ARPANET. IEEE TRANSACTIONS ON COMMUNICATIONS, VOL. COM-28,NO.5, MAY 1980

Paper: 48-bit Absolute Internet and Ethernet host numbers. Dalal en Printis. Xerox company. July 1981

[RFC-1498]: J. Saltzer: "On the Naming and Binding of network destinations", August 1993.

[RFC-814]: "NAME, ADDRESSES, PORTS, AND ROUTES." David D. Clark. July 1982.

Webpage: Will The Real "End-End Principle" Please Stand Up? Noel Chiappa. http://mercury.lcs.mit.edu/~jnc/tech/end_end.html

IETF draft: Endpoints and Endpoint Names: A Proposed Enhancement to the Internet Architecture. Noel Chiappa, 1999.

Paper: Developing a Next-Generation Internet Architecture. Robert Braden, David Clark, Scott Shenker, and John Wroclawski July 15, 2000.

Paper: Rethinking the design of the Internet: The end to end arguments vs. the brave new world. Marjory S. Blumenthal, David D. Clark. ACM Transactions on Internet Technology, August 2001.

Paper: New Arch: Future Generation Internet Architecture. David Clark, Mark Handley, Noel Chiappa, etc. December 2003.

Paper: Towards A New Internet Routing Architecture: Arguments for Separating Edges from Transit Core. Lixia Zhang Oct 2008 [Seventh ACM Workshop on Hot Topics in Networks]

Paper: Networking Named Content. Van Jacobson. CoNEXT'09, December, Rome, Italy.2009.

Paper: A Dynamic Recursive Unified Internet Design (DRUID). Joe Touch etc. Elsevier Comput. Netw. 2011.

Paper: Robustness and the Internet: Design and evolution. Walter Willinger and John Doyle, March 2002. (*)

Report: Remaking the Internet: Taking Network Architecture to the Next Level, Richard Bennett. Summer 2011. www.twcresearchprogram.com

Report: Designed for Change: End-to-End Arguments, Internet Innovation, and the Net Neutrality Debate. Richard Bennett Information Technology and Innovation Foundation. September 2009.

Paper: Why Loc/Id Split Isn't the Answer. John Day 2008.

IETF Draft: Architectural Implications of Locator/ID Separation, D. Meyer, D. Lewis. draft-meyer-loc-id-implications-00.txt, Jan 2009.

Paper: Declarative Transport No more transport protocols to design, only policies to specify. G.Gursun, Ibrahim Matta, K. Mattar, John Day Technical Report BUCS-TR-2008-014. August 26, 2008

Paper: On the Performance and Robustness of Managing Reliable Transport Connections. G.Gursun, Ibrahim Matta, K. Mattar. Boston University. 2009

Presentation: Handle with Care: Naming, Layering and Caching "Broken" on the Internet. Saleem Bhatti, University of St Andrews. October 2009.

Paper: Revisiting A Soft-State Approach to Managing Reliable Transport Connections. G. GURSUN IBRAHIM MATTIA, K. MATTAR. 2010.

Paper: OOI Recursive Inter-Network Architecture (RINA). John Day etc. december 2009.

Presentation: Networking is IPC and only IPC or How to Clean a Slate. John Day 2010

Presentation: Things They Never Taught You About Naming and Addressing. John Day. Feb 2010.

Paper: On Supporting Mobility and Multihoming in Recursive Internet Architectures. Ibrahim Matta. Technical Report BUCS-TR-2010-035. 2010.

Presentation: An Introduction to RINA Or How I Learned to Stop Worrying and Love the Internet FutureNet Tutorial Part I. John Day. May 2010

Paper: IS THE INTERNET AN UNFINISHED DEMO? MEET RINA! Eduard Grasa, John Day, Ibrahim Matta, etc. TERENA. october 2010.

Paper: Design principles of the Recursive InterNetwork Architecture (RINA). Eduard Grasa, John Day, Miguel Ponce de Leon, John Day, Ibrahim Matta, etc. 2010?

Paper: How in the Heck Do You Lose a Layer!? John Day. 2011 International Conference on the Network of the Future, 2011.

Paper: Network Forensics in a Clean-Slate Internet Architecture, Strauss, Olivier, Univ of Pretoria.

Paper: Demonstrating RINA using the GENI Testbed. Ibrahim Matta. Boston University.

Paper: Experimenting with Routing Policies using ProtoRINA over GENI. Ibrahim Matta, Boston University

Cook Report: Recursive Internet Architecture. An introduction to a Possible New Foundation for Global Telecommunications. Gordon Cook, oktober 2013.

Presentation: Recursive InterNetwork Architecture - A policy-based recursive approach to building a better Internet. Abraham Matta, Boston Univ. March 2014

"Networking is IPC": A Guiding Principle to a Better Internet Internet 1.0 is broken - Internet 2.0 is a repeat with more b/w How about Internet 3.0? Ibrahim Matta, Boston University

Paper: IRINA - Milestone M5: Comparative analysis of clean slate architectures and final use case for the NREN environment. Geant Open Call, may 2014.

Paper: Why the Internet only just works. M Handley, BT Technology Journal • Vol 24 No 3 • July 2006

Presentation: Re-thinking the Control Architecture of the Internet. Mark Handley, University College London, Dec 2009. [<http://www0.cs.ucl.ac.uk/staff/m.handley/slides/>]

Paper: Moving beyond TCP/IP. Fred Goldstein and John Day for the Pouzin Society April 2010.

Paper: Should Specific Values Be Embedded In The Internet Architecture? Ian Brown, David Clark, Drik Trossen, Nov 2010. [Conference: Re-Architecting the Internet workshop (CoNEXT).]

Presentation: The Delta-t Transport Protocol: Features and Experiences. Pathapati. February 2011.

Paper: Fundamental Limitations of current Internet and the path to Future Internet. EC FIArch Group. March 2011.

Website: <http://pnanetworks.com/> PSOC - Pouzin Society

Website: <http://irati.eu/introduction-to-rina/>

Website: <http://www.martingeddes.com/think-tank/network-architecture-research-tcp-ip-vs-rina/>

Website: <http://spectrum.ieee.org/computing/networks/osi-the-internet-that-wasnt>

Website: NDN - Named Data Networking - <<http://named-data.net/>>

Website: NSF FUTURE INTERNET ARCHITECTURE PROJECT <<http://www.nets-fia.net/>>

Website: NewArch Project: Future-Generation Internet Architecture (2003). <http://www.isi.edu/newarch/>

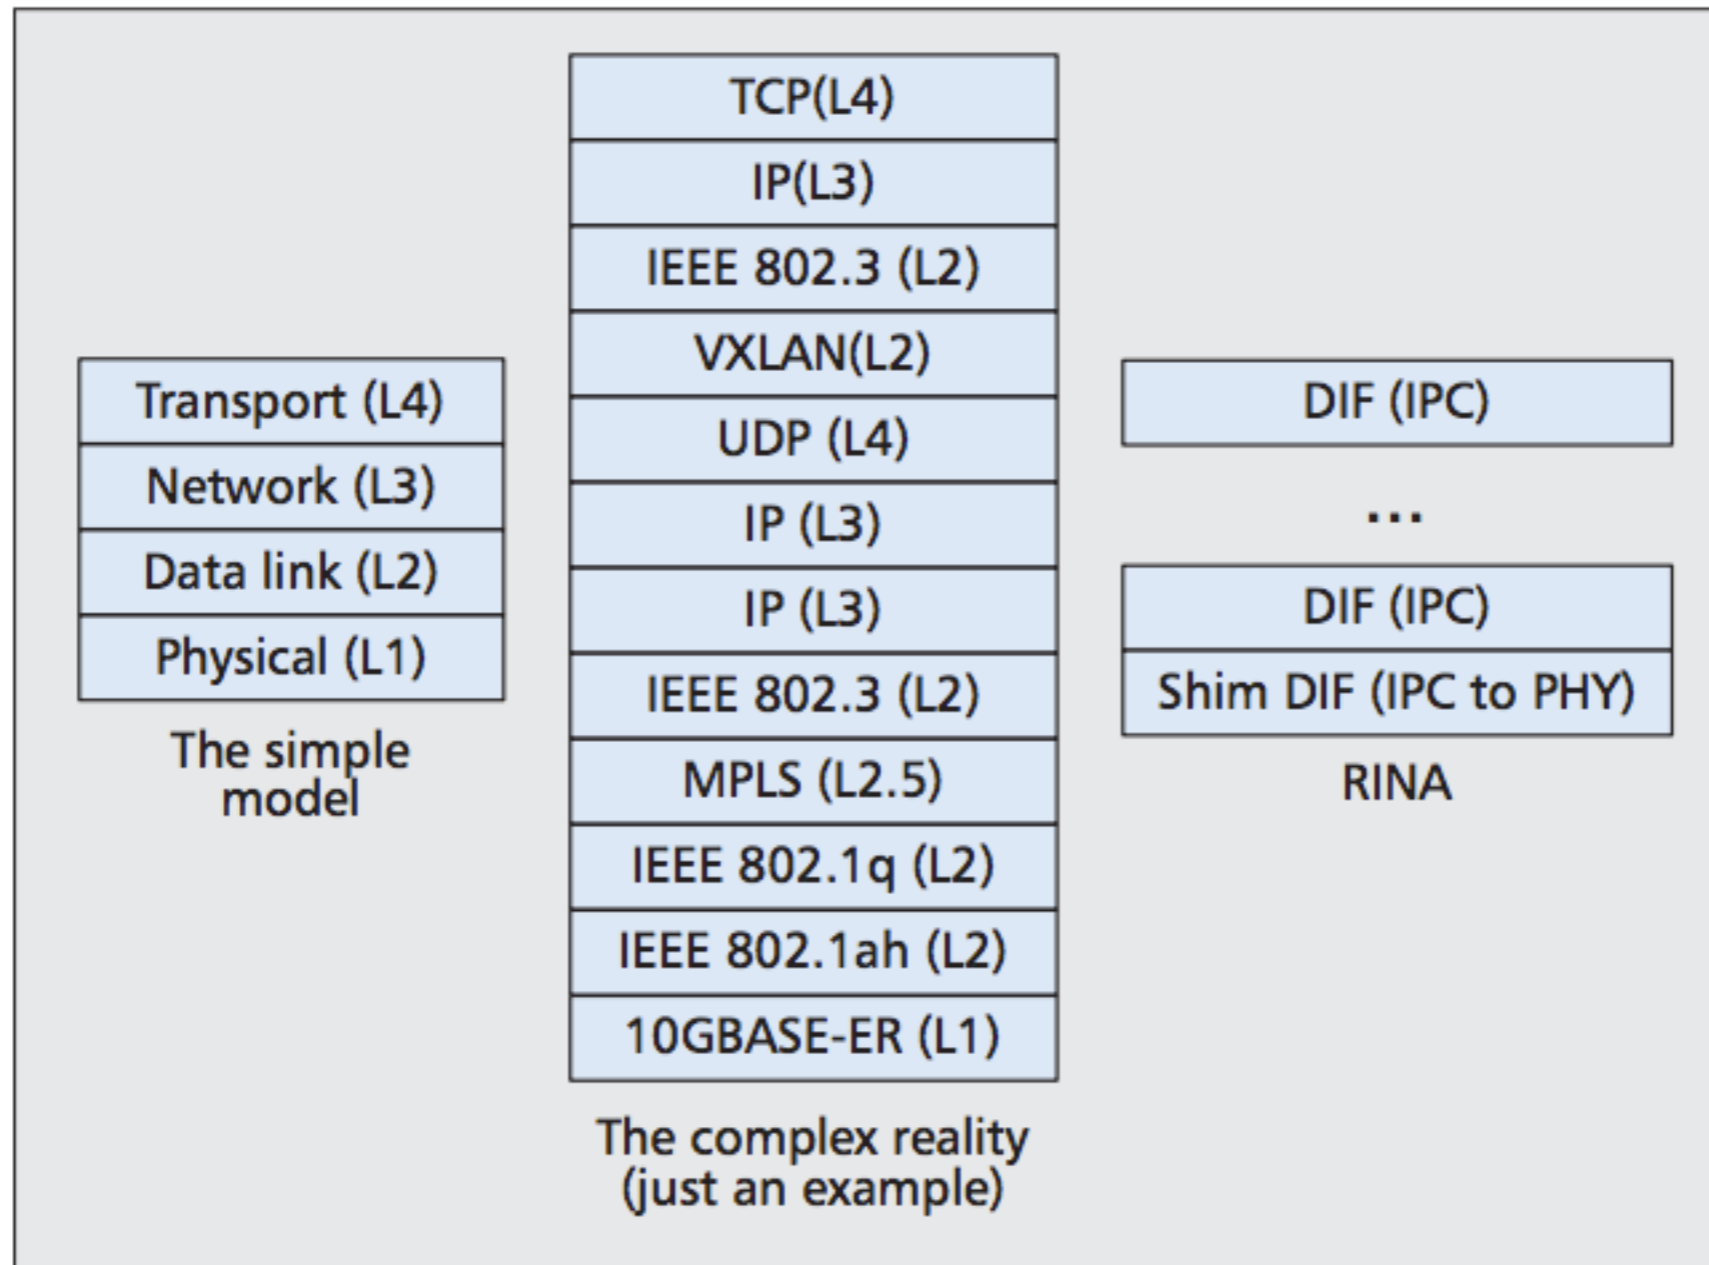
RECURSIVE INTERNETWORK ARCHITECTURE

- ▶ Programmable networking approach based on IPC
 - ▶ High scalability
 - ▶ **Multihoming**
 - ▶ Built-in Security
 - ▶ Seamless access to real-time information
 - ▶ And more...

RINA: IPC MODEL

- ▶ Inter Process Communication (IPC)
 - ▶ “Networking provides the means by which processes on separate computer systems communicate, generalising the model of local inter-process communications”

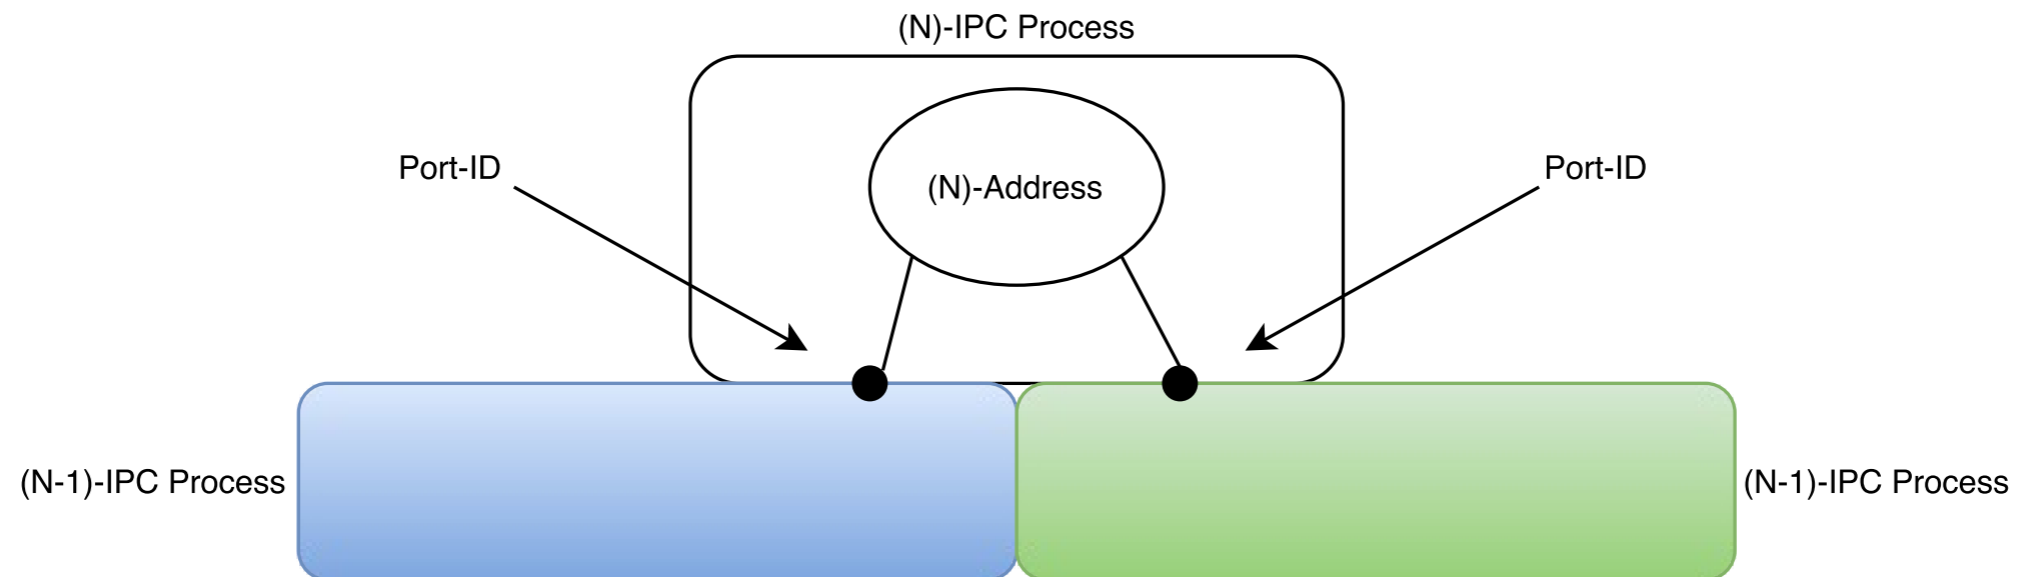
RINA



<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=6786609>

RINA: MULTIHOMING

- ▶ Each node has its own address
- ▶ A node can have multiple connections, each having a separate address
- ▶ A route to a destination node address is a sequence of intermediate node addresses



IRATI

- ▶ Investigating RINA as an Alternative to TCP/IP
- ▶ European Project (Framework Programme 7)
- ▶ Make enhancements of RINA architecture reference model and specification
- ▶ Focussing on DIFs over Ethernet

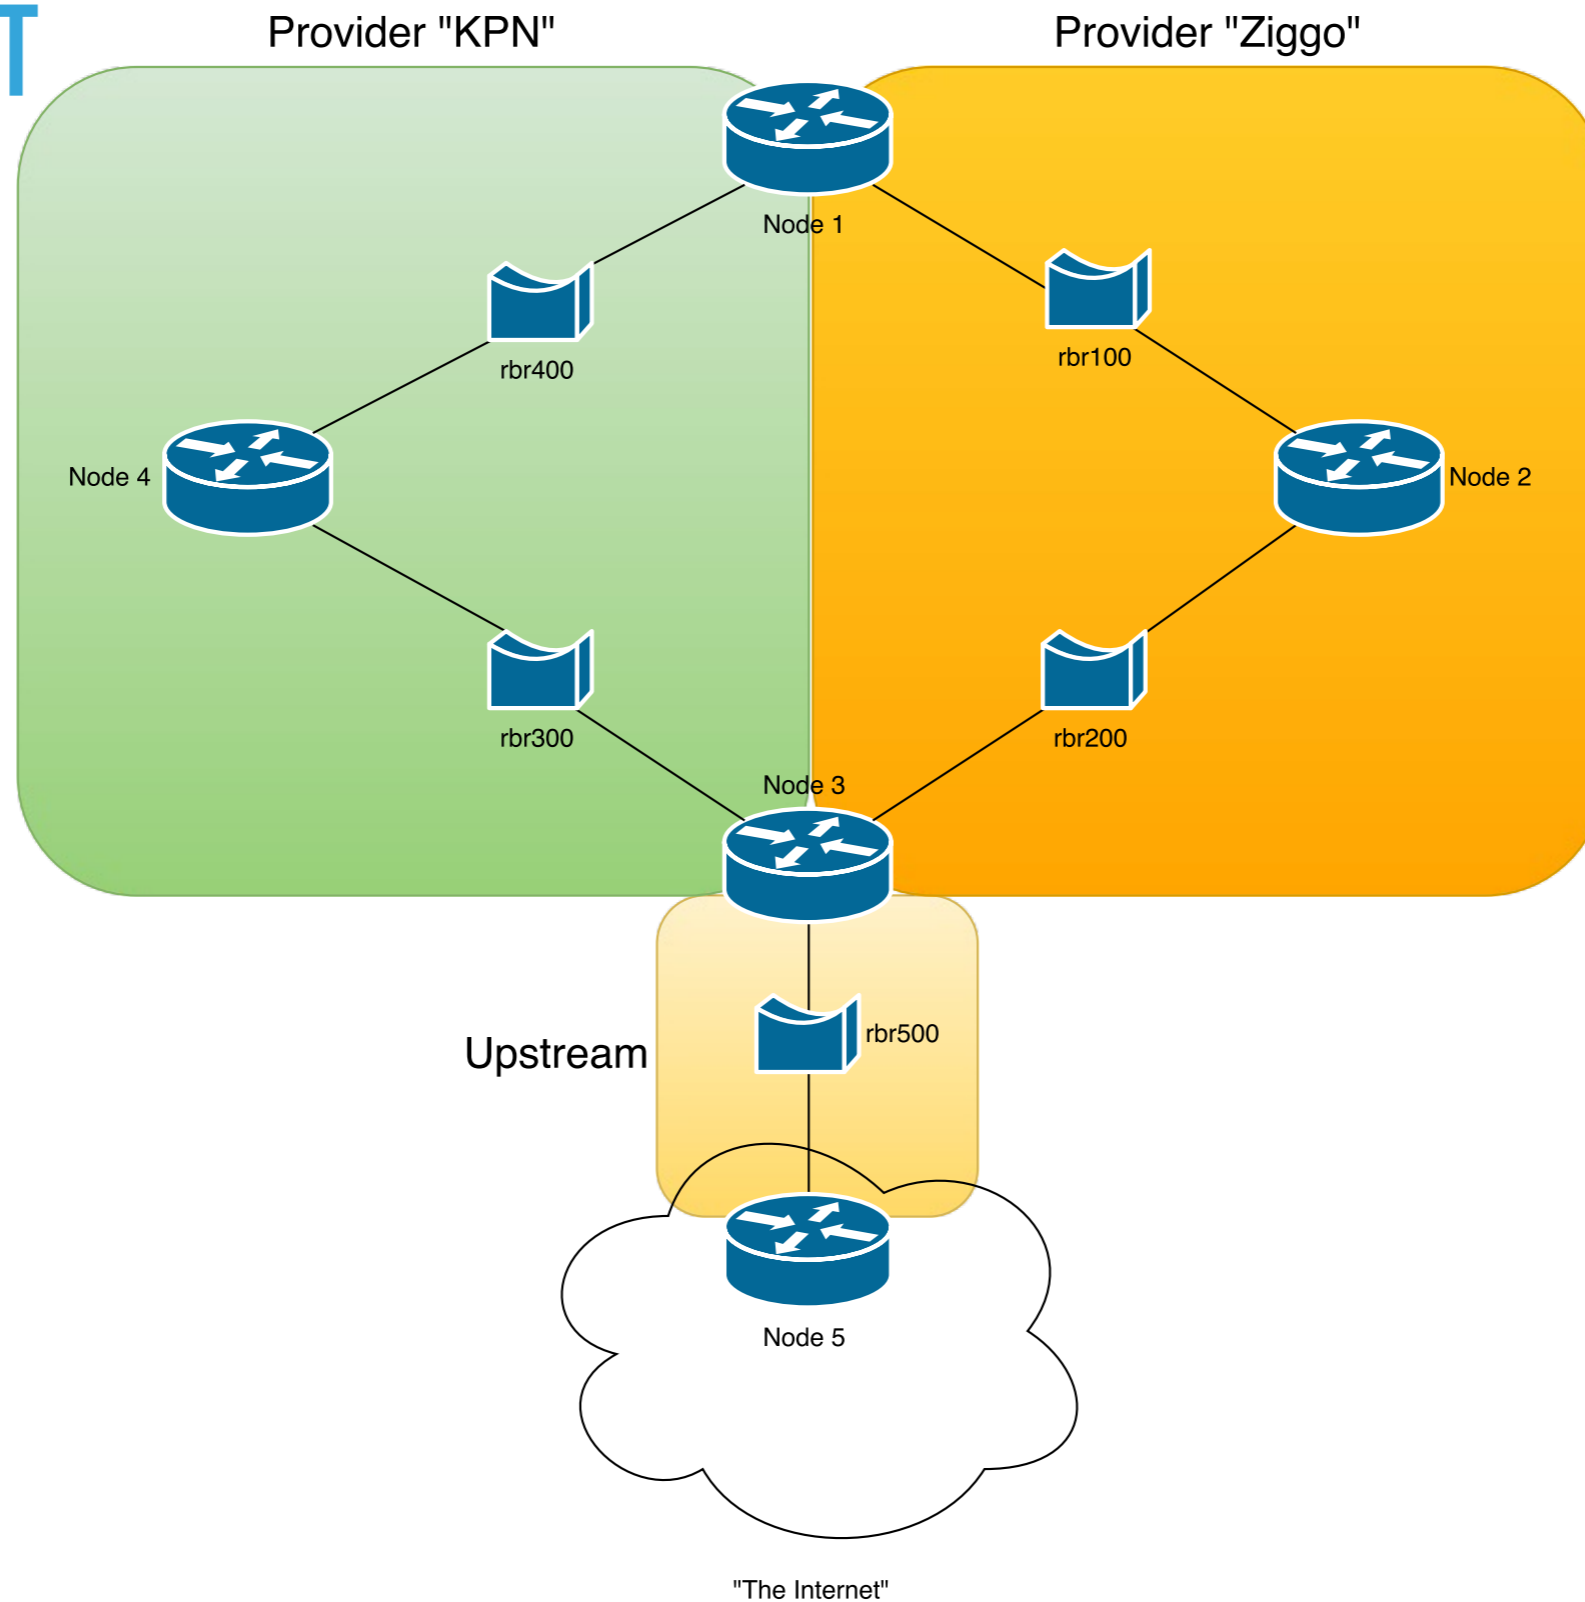
PRISTINE

- ▶ Closely related to IRATI
- ▶ European Project
- ▶ Creates programmable functions for congestion control
- ▶ Facilitating more efficient topological routing and multi-layer management
- ▶ Continuation of work IRATI

EXPERIMENT

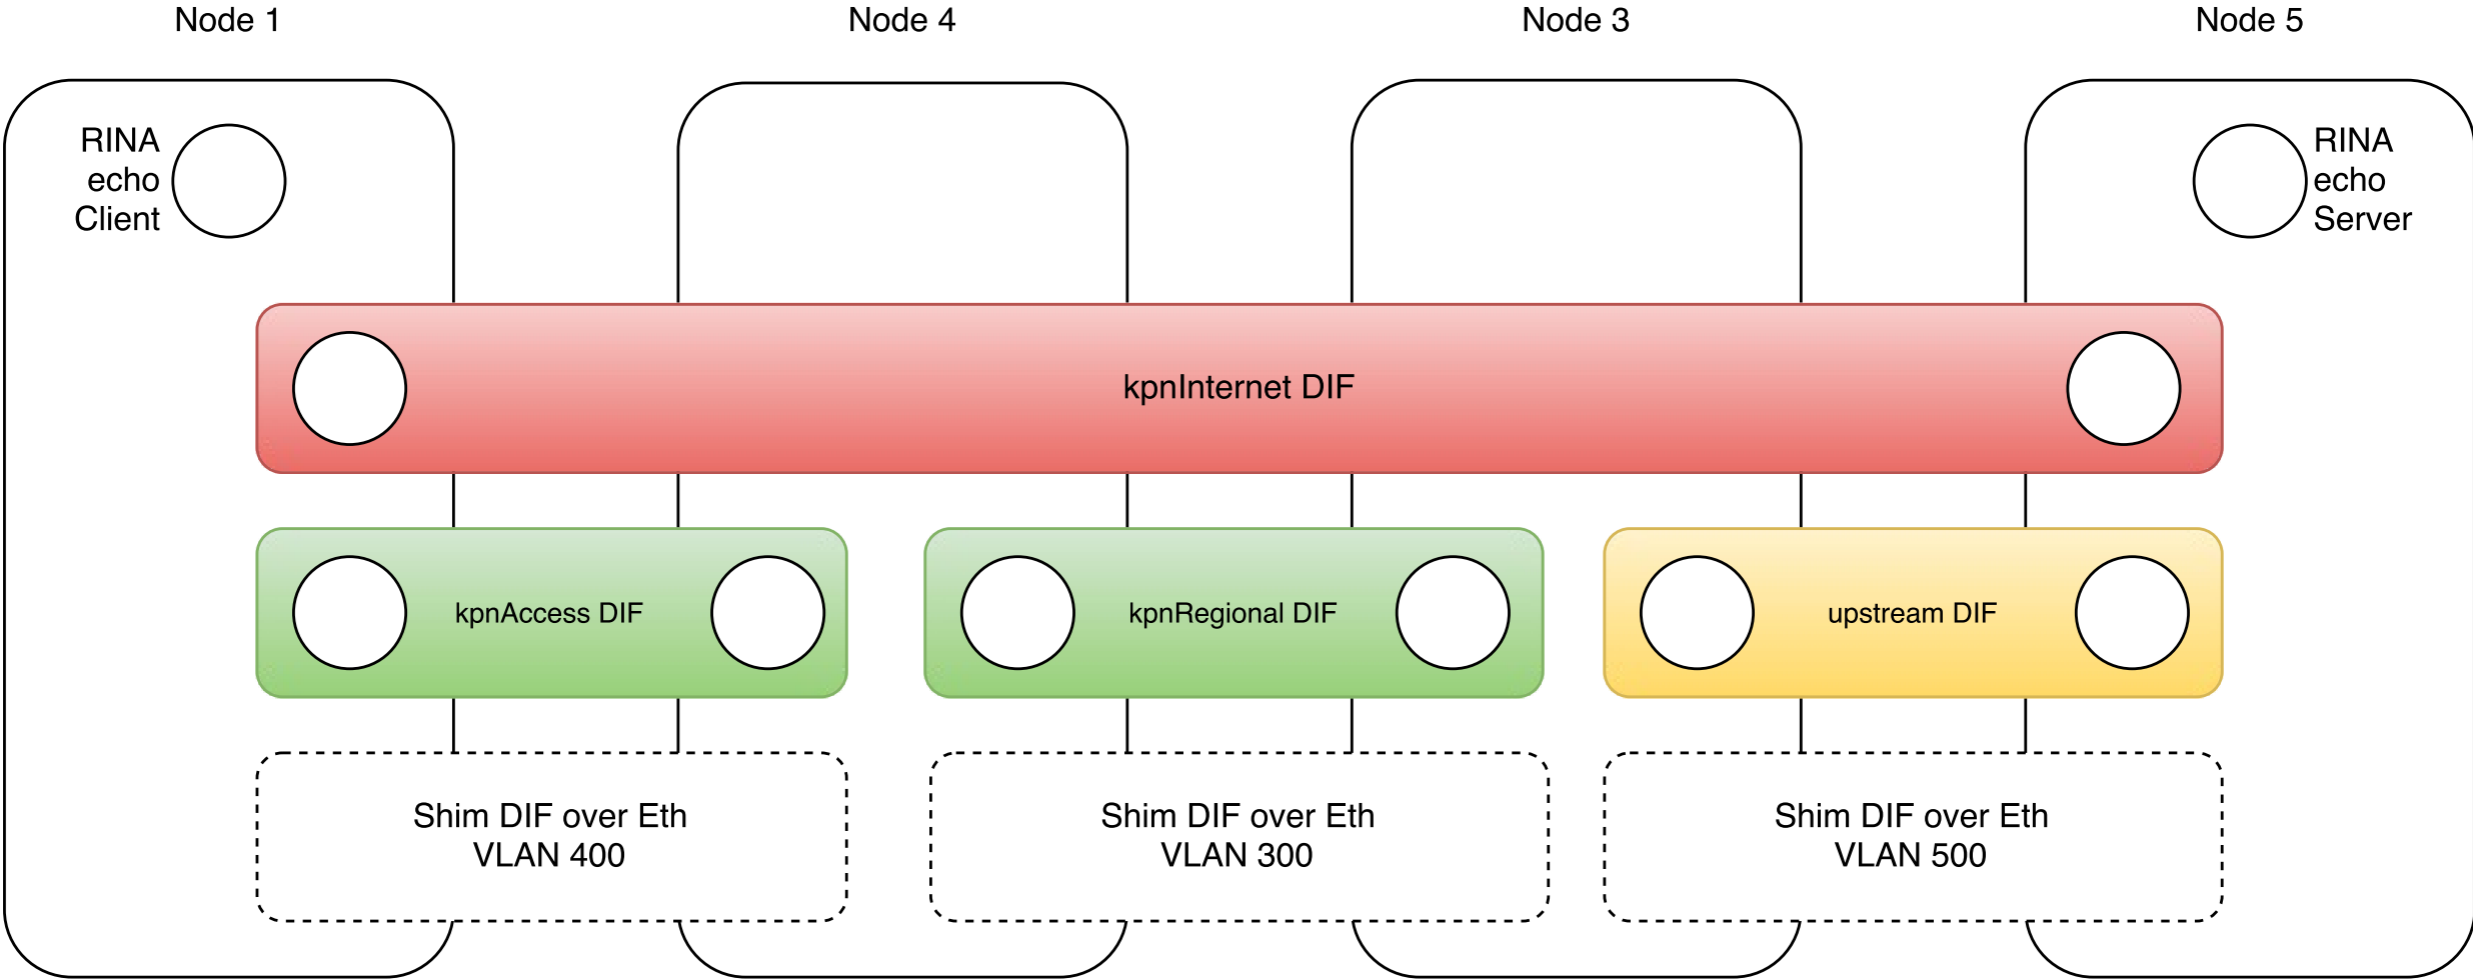

- ▶ Testing multi-DIF environment
- ▶ Testing with IRATI Demonstrator and VMs with compiled stack (pristine-1.5 branch)
- ▶ Testing if IRATI supports multihoming

EXPERIMENT



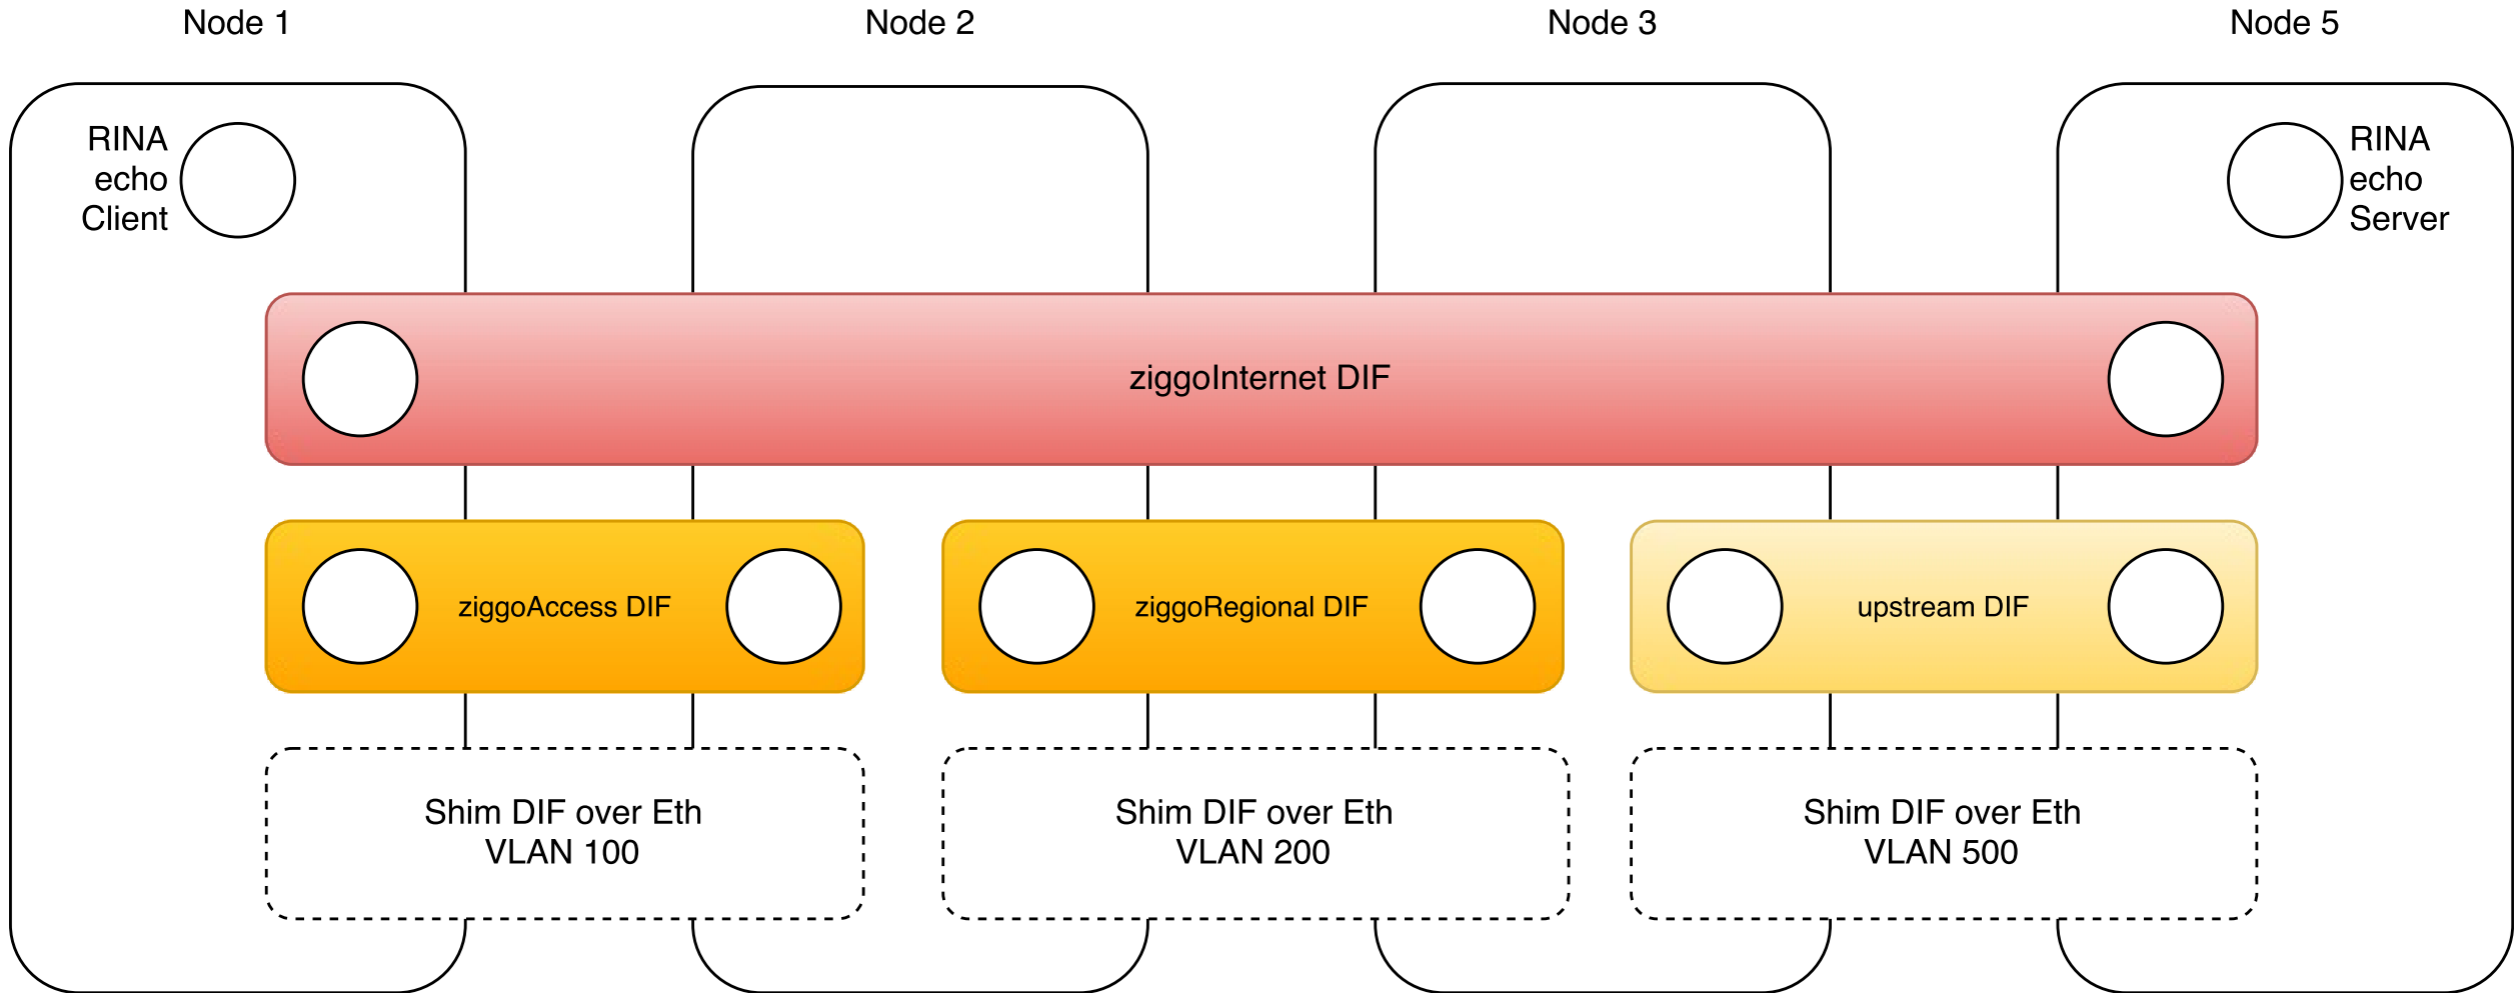

TAKING A CLOSER LOOK AT IRATI

EXPERIMENT



TAKING A CLOSER LOOK AT IRATI


EXPERIMENT



RESULTS

- ▶ IRATI Demonstrator keeps consuming memory until crash
- ▶ VMs with compiled stack did not show this behaviour
- ▶ Mapping of applications is static
- ▶ RINA tools do not support registering at two DIFs at the same time (yet?)

IRATI Demonstrator: Free Memory



RESULTS

- ▶ DIF allocator will be properly implemented by H2020 ARCFIRE Project
- ▶ PRISTINE project further investigates and develops routing algorithms and routing information dissemination strategies that **optimally** exploits RINA's support of multihoming for load balancing and rapid recovery

CONCLUSION

- ▶ Current Internet has some serious problems, including the multihoming problem
- ▶ Current solutions for multihoming/mobility create a more complex Internet and do not solve the actual problem(s)
- ▶ Since Klomp/van Leur's research in January a lot of improvements have been made in IRATI stack, yet still very experimental
- ▶ At this moment IRATI is not able to do multihoming yet
 - ▶ Mapping of applications still static

FUTURE WORK

- ▶ Documentation is still quite poor
- ▶ DIF allocation is still static...
- ▶ RINA tools are not able to register to multiple DIFs at the same time (yet)
- ▶ PRISTINE and H2020 ARCFIRE are actively improving IRATI

QUESTIONS?