Core Networking Panel

*

Sally Floyd

January 9, 2003

NSF ANIR Principle Investigator Meeting

With input from members of the End-to-end Research Group, and with material ”borrowed” from Craig Partridge, Dave Rossetti, and others.

#1: Progress in Core Networking Research Areas?

*

- There is a great need for further research on core networking.
- We can't rely on industry for all of the research in core networking.
- Research is needed on the evolution of the Internet infrastructure as well as on longer-time-scale grand challenges.
- A key consideration for NSF:

– Whether the research is science (e.g., contributing to our body of knowledge of networking) as well as engineering (the practical application of that knowledge).

– This includes the body of knowledge about the human-created complex system of the global Internet.

#1: The Need for Further Research in Routing:

*

Smaller issues:

- Convergence.
- Multimetric routing (e.g., with both hop-count and delay as metrics).
- How do users choose the ISPs that their packets will traverse?

Larger and longer-time-scale issues:

- What replaces BGP?
- How do routing algorithms deal with intermittent connectivity, scheduled connectivity, and unidirectional links.
- ...

#1: The Need for Further Research in Transport:

*

Smaller issues:

- Transport-level congestion control mechanisms:

– Multicast as well as unicast.
– Congestion control for streaming media or for other applications.
– Congestion control for very-high-speed or high delay environments.

Larger and longer-time-scale issues:

- Congestion control mechanisms in routers:
  – Responding to aggregates such as DDoS attacks.
  – More fine-grained feedback from routers?

- New forms of communication between layers?

- Understanding global dynamics.

- ...

#2: The Evolution of the Internet?

*

- We need a more scientific understanding of the evolutionary potentials and difficulties of the Internet infrastructure:

  – of the inherent forces and dynamics of evolution;
  – of the past difficulties and successes;
  – of the future range of possibilities.

- We don't want to give up on the evolutionary development of the Internet infrastructure just because it is harder than it used to be ... where that is the appropriate path.

#2: Overlay Networks:

*

- Overlay networks are a great transition technology.
- Capabilities from overlay networks should move into the core if:
  – they have an inherent need to be everywhere;
  – there are scaling issues;
  – they are general purpose and widely useful.

#3: The challenges of scalability, manageability, robustness, and more:

*

- The Internet architecture gives a reasonably sound foundation (though not perfect):

  – We don't want to start again from scratch, unnecessarily.

  – The best track records so far have come from combining architecture and analysis with a deeply empirical approach.

- At the same time, the scalability, manageability, and robustness of the Internet are at risk.

#3: Revisiting the core problems in this context:

*

- We could start by understanding better:
  – the current scaling limits;
  – the new stresses from the environment;
  – the tradeoffs between complexity and robustness;
  – the challenges of manageability;
  – how to add robustness in terms of robust performance in known extreme environments, without sacrificing robustness in unforseen circumstances.
  – ...

- We probably also have to explore a range of new directions.