Multicast Support

Nimrod approaches multicast with the same ideas used elsewhere: try and break the problem up into pieces, and put as much of the functionality as possible *outside* the architecture, to allow flexibility in algorithms, etc.

This last is especially important for multicast, where group sizes can vary by many orders of magnitude.

Nimrod separates several distinct phases of creating a multicast group:

- Determining membership
- Deciding what kind of data distribution you want (per source, or whatever)
- Calculating one or more spanning trees which connect the members
- Installing the state about those spanning trees in the routers
- Forwarding user data

Multicast Support, Continued

Nimrod provides some very fundamental multicast building block(s), such as multicast flow setup. All the rest (like the mechanisms to maintain the state about group membership, calculate the spanning tree, etc) are outside the core architecture. Users can select locally whichever algorithm makes sense in their particular application.

This has all the advantages of the Nimrod approach to unicast:

- Makes it a simpler design
- Makes it less likely there will be something wrong
- Allow experiments with new algorithms
- Allows incremental deployment of new algorithms

Nimrod will also distinguish between a *multicast group* (i.e. a name for the set of members) and a particular *multicast flow* (i.e. a particular data distribution channel to that group). There can be multiple flows which go to the same group, but controlled independently.

Flow Aggregation

Flow aggregation:

- Provides no extra capability to the users
- Adds complexity

So, why add it?

- Because it is needed to allow a positive economy of scale in high-speed switches, since the unit (i.e. per-flow) cost of per-flow memory is higher there.
- Flow aggregation also has the nice property that it works well with virtual links.

Note that you don't have to have a *lot* of aggregation, just enough to keep the unit cost of per-flow state memory declining at an appropriate rate.

Other Topics

There are many other topics which are not covered here. They include:

- Virtual links and routers, and recursive instantiation
- Renumbering
- Abstraction and advertisement
- Interaction with Resource Allocation
- Advanced topics in flow aggregation, such as asymmetric aggregation
- Mobility
- Topology database exchange
- Nimrod entities
- Flow setup
- Protocols
- Deployment plans