Generic Router Assist

Presenter: Tony Speakman
Co-authors:
Brad Cain
Ken Calvert
Christos Papadopoulos
Don Towsley
Swapna Yelamanchi

Outline

- Model of Operation
- Filter Definitions
- GRA Headers
- Principles of Operation
- Drafts

GRA - Model of Operation

- GRA-capable routers in some fraction of the routers in a source-specific multicast distribution tree
- Each of those routers discovers a unique upstream GRA neighbour
- Pre-defined filter definitions reside in those routers
- filter definitions define (amongst other things) matching GRA headers
- The source and receivers in the transport session direct packets with GRA headers into the distribution tree
- Those packets are matched by GRA routers against the filter definitions and processed accordingly

GRA - Filter Definitions

Each filter definition specifies:

- Filter Identifier (FID)
- Housekeeping functions (such as a life timer for the filter definition itself)
- Action Specifications (sub-filter types):
  - Action Identifier (SFTYPE)
  - Housekeeping functions (such as a life timer for the action specification itself)
  - A GRA header format
  - Steps (sub-actions):
    - Predicates
    - Functions
    - State
    - Housekeeping functions

GRA - Schematic Filter Definition

The handling of parity NAKs and parity retransmissions in PGM can be described as a predicate eliminating and subcasting filter augmented by a packet operand, the number of parity packets requested.

- Filter ID
  - ELIM_SCAST
- Housekeeping Functions
  - FILTER_LT - life timer for the filter spec
- Action Specifications

GRA - Schematic Action Specification

- Action Identifier
  - RCVR_UPDATE (i.e., a NAK)
- Housekeeping Functions
  - RCVR_SVC_LT - life timer for the action
- GRA Header Operands
  - SQN (KEY in this case)
  - RQST_COUNT
- Key-specific state - KEY_STATE
  - KEY_ET - elimination timer
  - KEY_LT - life timer for the key-specific state; discard KEY_STATE upon expiry
  - HIGH_COUNT - maximum RQST_COUNT seen on any IIF
  - OIF_LIST - list of IIFs seen and their OIF_COUNTs (acts as an OIF list selector)

GRA - Schematic Action Specification

- Unconditional step
  - multicast packet on incoming interface (suppression)
- Predicate on a KEY miss
  - NOOP - action is unconditional
- Steps
  - HIGH_COUNT = MAX(RQST_COUNT, HIGH_COUNT)
  - OIF_LIST = IIF
  - OIF_COUNT for IIF = RQST_COUNT
  - start KEY_ET, start KEY_LT,
  - reverse forward packet to upstream neighbour
- Predicate on a KEY match
  - KEY_ET is running or RQST_COUNT LEQ HIGH_COUNT?
- Steps on TRUE
  - OIF_COUNT for IIF = MAX(RQST_COUNT, OIF_COUNT for IIF)
  - restart KEY_LT
  - discard packet
- Steps on FALSE
  - OIF_COUNT for IIF = MAX(RQST_COUNT, OIF_COUNT for IIF)
  - restart KEY_ET, HIGH_COUNT = RQST_COUNT
  - reverse forward to upstream neighbour

GRA - Schematic Action Specification

- Action Identifier
  - FORWARD (i.e., subcast)
- Housekeeping Functions
  - FWD_SVC_LT - life timer for the action
- GRA Header Operands
  - SQN (KEY in this case)
- Predicate on a KEY miss
  - NOOP - action is unconditional
- Steps
  - discard packet
- Predicate on a KEY match
  - (for all OIF_COUNTs), OIF_COUNT NE 0
- Steps on TRUE
  - decrement OIF_COUNT
  - forward packet on OIF
  - OIF_COUNT == 0?
    - delete OIF_STATE from OIF_LIST
  - OIF_LIST == NULL?
    - discard KEY_STATE
- Steps on FALSE
  - discard packet

GRA - Schematic Header

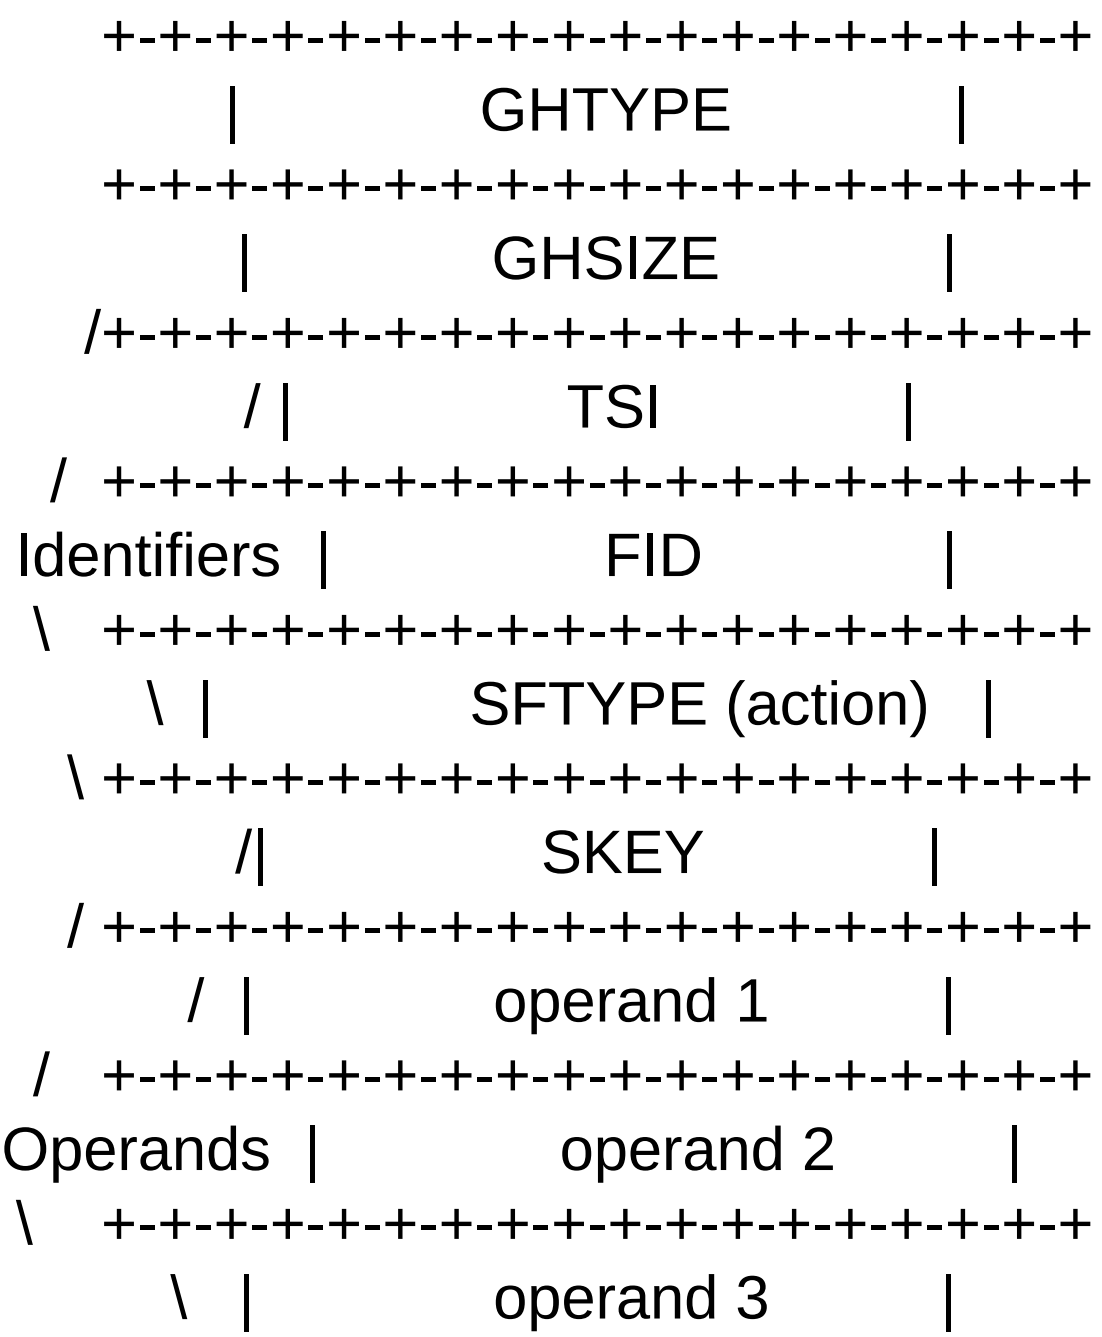

GRA - Header Field Descriptions

- GHTYPE
  - GRA Header Type
- GHSIZE
  - GRA Header Size
- TSI
  - Transport Session Identifier
- FID
  - Filter Identifier
- SFTYPE
  - Action Identifier - sub-filter type
- SKEY
  - Sub-session-specific label corresponding to key-specific state
- operands
  - other stuff ...

GRA - Principles of Operation

- Filter Definitions
  - Half the FSID space to be static and standard
  - Half the FSID space reserved so as not to preclude dynamic/custom filter definitions
    - which explains in part why FIDs are scoped by TSIs
- Header Specifications
  - Fixed identifier part
  - Fixed operands
  - Variable operands - TLV
- Storage
  - (in addition to state specified in the filter definition)
  - not to exceed one complete copy of the GRA header operands per session/sub-session
  - to accommodate time-triggered forwarding
  - precludes accumulation
- Packet Modifications
  - restricted to (over)writing GRA header operands as specified
- Packet Formatting
  - none, specifically no encapsulation/decapsulation, no accumulation

GRA - Principles of Operation

- Forwarding Functions
  - Multicast
    - NLA: S and G are the source and multicast destination pair associated with the TSI
    - on a known route possibly with an interface selector to select some subset of the OIF list
  - Unicast
    - NLA: S is the GRA element, D is variable
    - routed w.r.t. conventional unicast routing information
- Packet Generation
  - preclude other than forwarding a GRA packet in hand or a stored GRA header
- Control Protocol
  - in-band session information (GRA neighbour info)
  - (out-of-band?) administration
- Stringing it all together
  - Network-layer header contains a transport protocol number (serves, amongst other things, to scope TSIs), and a GRA-present indicator (Router Alert)
  - Implication of GRA-present indicator is that immediately following the network-layer header, there's a GRA header
  - Transport header immediately follows GRA header
  - Note that nothing precludes referencing the GRA header (the TSI and the operands specifically) in the implementation of the transport protocol

GRA - Drafts

- Architecture spec (Informational)
- Functional spec will morph into:
  - Filter Definition Language spec
  - Filter Definition specs
  - GRA Protocol spec (including in-band session control)
- GRA Control Protocol spec