Interdisciplinary Experiences with GridStat: Pub-Sub Status Dissemination for the Power Grid

Prof. Dave Bakken

School of Electrical Engineering and Computer Science
Washington State University
Pullman, Washington USA

51st IFIP 10.4 WG Meeting, Guadeloupe
January 12, 2007

http://gridstat.net

GridStat Team

- Faculty: Dave Bakken, Carl Hauser, Anjan Bose
- Current Students:
  - Graduate: Stian Abselsen, Erlend Viddal, Jim Kusznir, Sunil Muthuswamy, Erik Solum, Wendy Maiden (PNNL), Joel Hekley, Kim Swenson
  - Undergraduate: Eric Paige, Loren Thompson, Nathan Schubkegel
- Alumni:
  - Graduate students: Ioanna Dionysiou (PhD 2006), Kjell "Harald" Gjermundrød (PhD 2006), Venkata Irava (PhD 2006), Ryan Johnston (MS 2005), Ping Jiang (MS 2004), Suprith Sheshadri (MS 2005)
  - Undergraduates: about a dozen on senior projects (Avista Utilities)
- Note: all above students are computer science
  - Also working with Sudipto Bhowmik (PhD EE soon, almost MS CS)

Outline of Talk

- **The Problem**
- Interdisciplinary Observations
- GridStat Rationale & Overview
- GridStat Framework

Rationale for Better Communications

- US Electric Power Communications System is aging
  - SCADA & ICMP are 1960s technology
  - Not updated meaningfully (no industry investment)
  - Much star-connected, inflexible, slow, crude SCADA "polling"
  - **Very little between electric utilities**
- Data collection has increased many fold at substations
  - Faster measurement rates, often time synchronized
  - Communications not there to move this data where needed

Rationale for Better Communications (cont.)

- Clark Gellings, EPRI* (**emphasis** mine)
  - *"The **ultimate challenge** in creating the power delivery system of the 21st century is in the development of **a communications infrastructure that allows for universal connectivity**."*
  - *"In order to create this new power delivery system, what is needed is a **national electricity-communications superhighway** that links generation, transmission, substations, consumers, and distribution and delivery controllers."*

*EPRI≡Electric Power Research Institute, www.epri.com, an industry-funded US R&D org.

Rationale for Better Communications (cont.)

- Mechanisms for protection and control are >99% local
  - Poor communications infrastructure does not allow otherwise!
  - But dynamic phenomena are grid-wide
  - Special communication links needed for SPS/RAS
  - Special links and data concentrators for PMUs
- Power grid landscape is changing!
  - More “miles x megawatts”: little new transmission lines
  - More participants that can affect grid stability
  - More heterogeneity of devices
  - Heightened security concerns
- **Resulting situation awareness is bleak (“flying blind”)**
  - Strong contributing factor in all recent blackouts (US, Italy, …)
  - Greatly limits better control and protection schemes

Outline of Talk

- The Problem
- **Interdisciplinary Observations**
- GridStat Rationale & Overview
- GridStat Framework

Interdisciplinary Observations on Power R&D

- Electric power in USA spends less on R&D than pet food industry (IEEE, EPRI)
  - Similar problems & culture in Europe, too (sabbatical & US-EU)
  - Starting to change: now bimodal/bipolar
- Different terms: security, N-1, status
- Power (and other) engineers
  - Tend to "lock on" to a particular technology …
  - Tend to be unaware of state in the art and practice of applied distributed systems
    - *x*BB example
    - Analogous to "security is just encryption"
  - Tend to hard code things at many levels
  - Ignorant/underestimate cyber security vulnerabilities (Idaho Krings & Oman)

Other Misc. Remarks

- Power industry has a tendency to latch onto a given technology (bridged ethernet, IPv6, ….)
  - Then stuck with it for decades
  - Much better to focus on what (non-functional/QoS) **requirements** you have, then have a middleware layer above the technology
  - This is EXACTLY why many industries (aerospace, trains, etc.) have been using middleware heavily the last decade or more
  - Good programs in DARPA & EC in last 10 years on this (QoS-managed middleware)
- First Energy like problems can be detected with derived values & triggers
  - Subscribe to trigger on a minimum value of a derivative: among a set of variables, something should be changing over time…
  - More inter-utility data can be shared if auto-enabled only when nearing a crisis

Opinion: Joint IT-Power Research Needed!

- Premises
  1. Continued piecemeal expansion of the grid's communication capability (RAS/SPS) is unnecessarily expensive and does not meet even today's requirements
  2. Modernizing the grid must include communications
  3. Modernizing grid communications involves focused IT research
  4. This IT research should be done jointly with power researchers
- Without #4, we keep doing the same old things …
  - CS researchers publish, claiming to solve part of problem
    - Never integrated into any complete end-to-end IT solution & fully evaluated in real environment
  - Power researchers publish with control and protection strategies assuming today's inflexible communications
- Prediction: #3,#4 will never happen without DoE/DHS or EC leading

Outline of Talk

- The Problem
- Interdisciplinary Observations
- **GridStat Rationale & Overview**
- GridStat Framework

Reality Check & Focused Opportunity

- Unsolved problem: providing
  - Multi-dimensional QoS guarantees (latency, jitter, bandwidth, …)

  for a
  - Mixture of sophisticated and arbitrary application programs

  running on
  - A dynamic network with arbitrary topology and subscriptions

  Likely to be unsolved 20-30 years from now (general case)!
- More solvable problem: providing
  - Multidimensional QoS (softer) guarantees augmented by redundant paths and specialized routers
  - Delivering status updates and alerts for simple and predictable power grid applications
  - Static (almost) and predictable network topologies & subscriptions
- GridStat is working on this more solvable problem

GridStat Approach

- Build pragmatic, comprehensive end-to-end framework
  - Extensibility & customizability are key (lots of hooks…)
  - Intended to extend to capabilities & scope of large power grid
- "Outside-In" not "Inside-Out"
  - lay down all the end-to-end plumbing, *a la* QuO
- Start with simple QoS & sub-optimal mechanisms
  - Hard QoS guarantees only if we control all access points
  - Provide QoS APIs & hooks to capture requirements to enable many more optimizations and more extensive management
- Extend over time for more coverage of
  - QoS guarantees
  - Adaptability
  - Security

  With more QoS mechanisms, policy languages, validation, ….

GridStat is Publish-Subscribe Middleware

- Delivers status value updates (sensors, control outputs…)
- Simple, CORBA APIs for both publishers and subscribers, management/control infrastructure, etc. (.NET pubs/subs)
- Network of internal status routers (SRs) managed for QoS – timeliness, redundancy and security
  - Middleware-level store-and-forward with rate filtering & multicast
  - Data plane kept separate from management plane
  - Forwarding latency ~0.5ms (Java) and 50K/sec on 3-year-old HW
- Optimized for semantics of status items
  - Not just arbitrary event delivery like generic publish-subscribe
  - Different subscribers (subtrees) can get different rates, latencies, #paths
  - Designed to allow many adaptations assuming semantics of status updates
- Goal: provide data availability via managed QoS & data load shedding

GridStat Middleware (MW) in Context

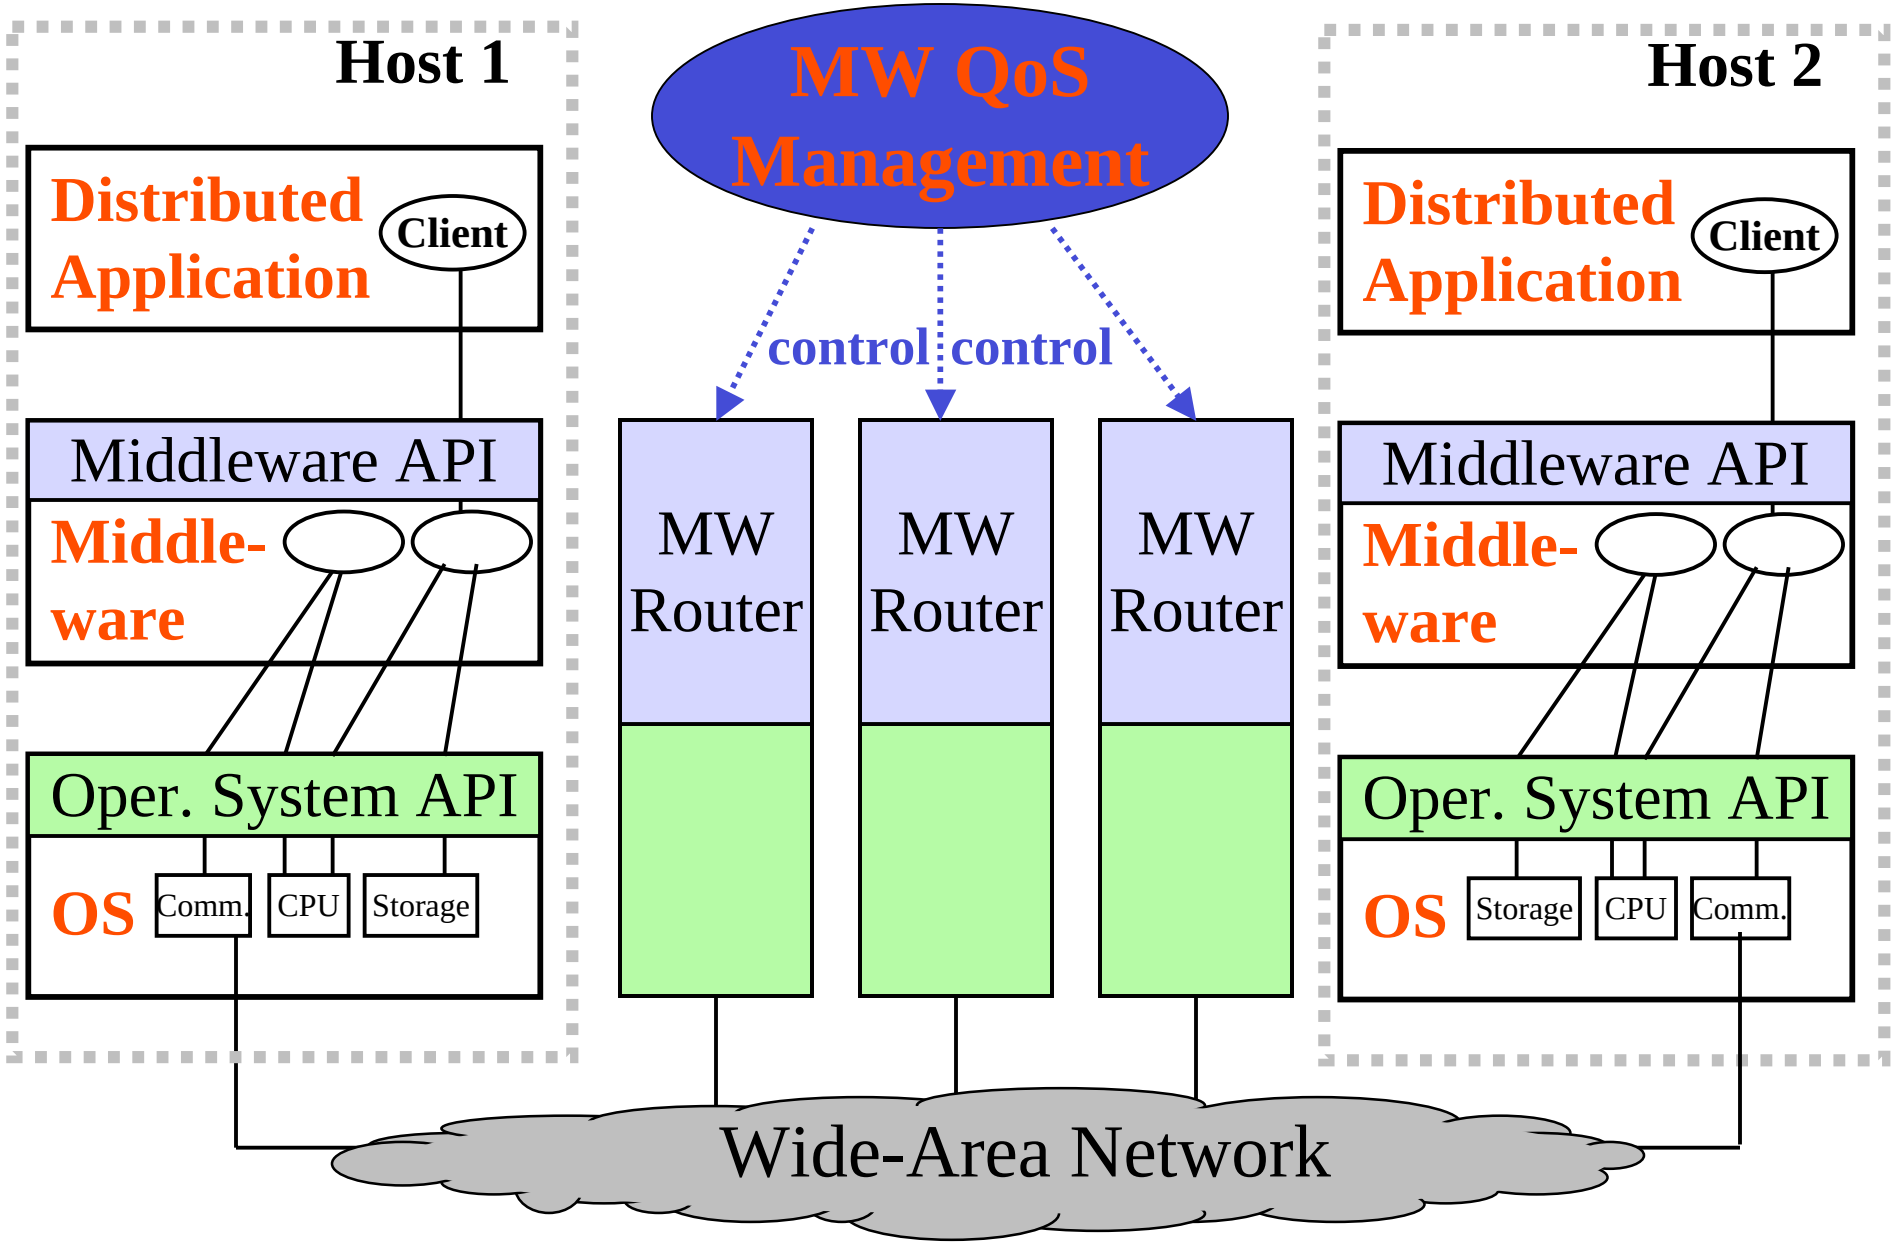

GridStat Capabilities: Subscriber-Side

- Subscribe to status variable updates or alerts
- Subscribe to baseline status value or derived values
- APIs for status variable subscription:
  - **Pull-from-Cache**: (use directly in computations)
  - **Direct Push**: update via callback object
  - **QoS Push** [optional]: callback if specified QoS violated
- QoS specified: desired & worst-case latency, rate, redundant paths
- Extrapolation functions (preconfigured or customizable) compensate for omission failures of update delivery

GridStat Capabilities: Status Routers

- Rate filtering mechanisms at SRs and subscriber proxies
- Multicast with link reuse for efficiency
- Temporally synchronized rate filtering across different status update flows …
- Condensation functions: user-extensible aggregation logic
- Preconfigured modes & mode transitions supporting "subscription bundles"
- Network transparent: run over multiple COTS networking technologies
  - IP, ATM (or lower fiber), network processors, …
  - Run over dedicated lines, shared Internet, …
    - Some baseline has to be dedicated (!!)

Basic GridStat Functionality

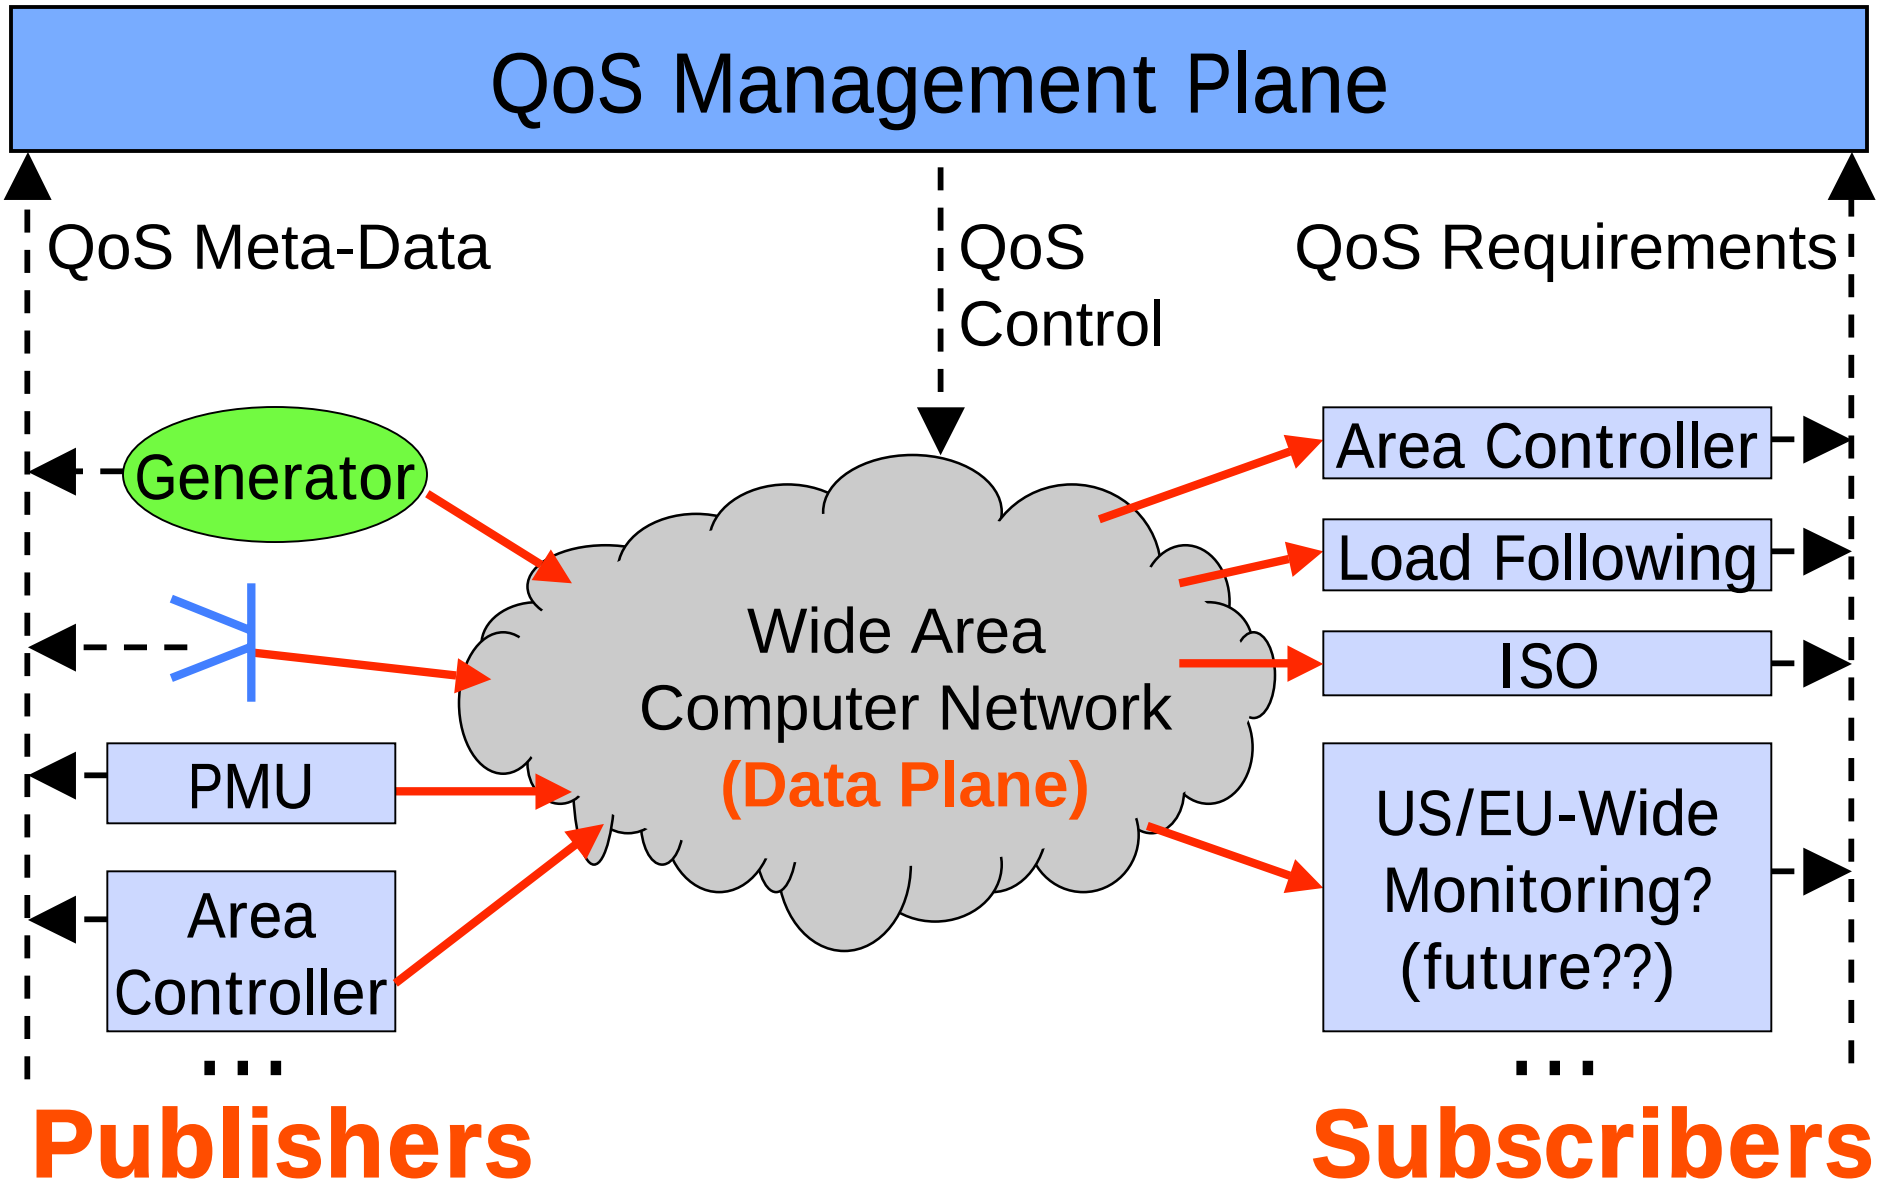

Outline of Talk

- The Problem
- Interdisciplinary Observations
- GridStat Rationale & Overview
- **GridStat Framework**

GridStat Architecture

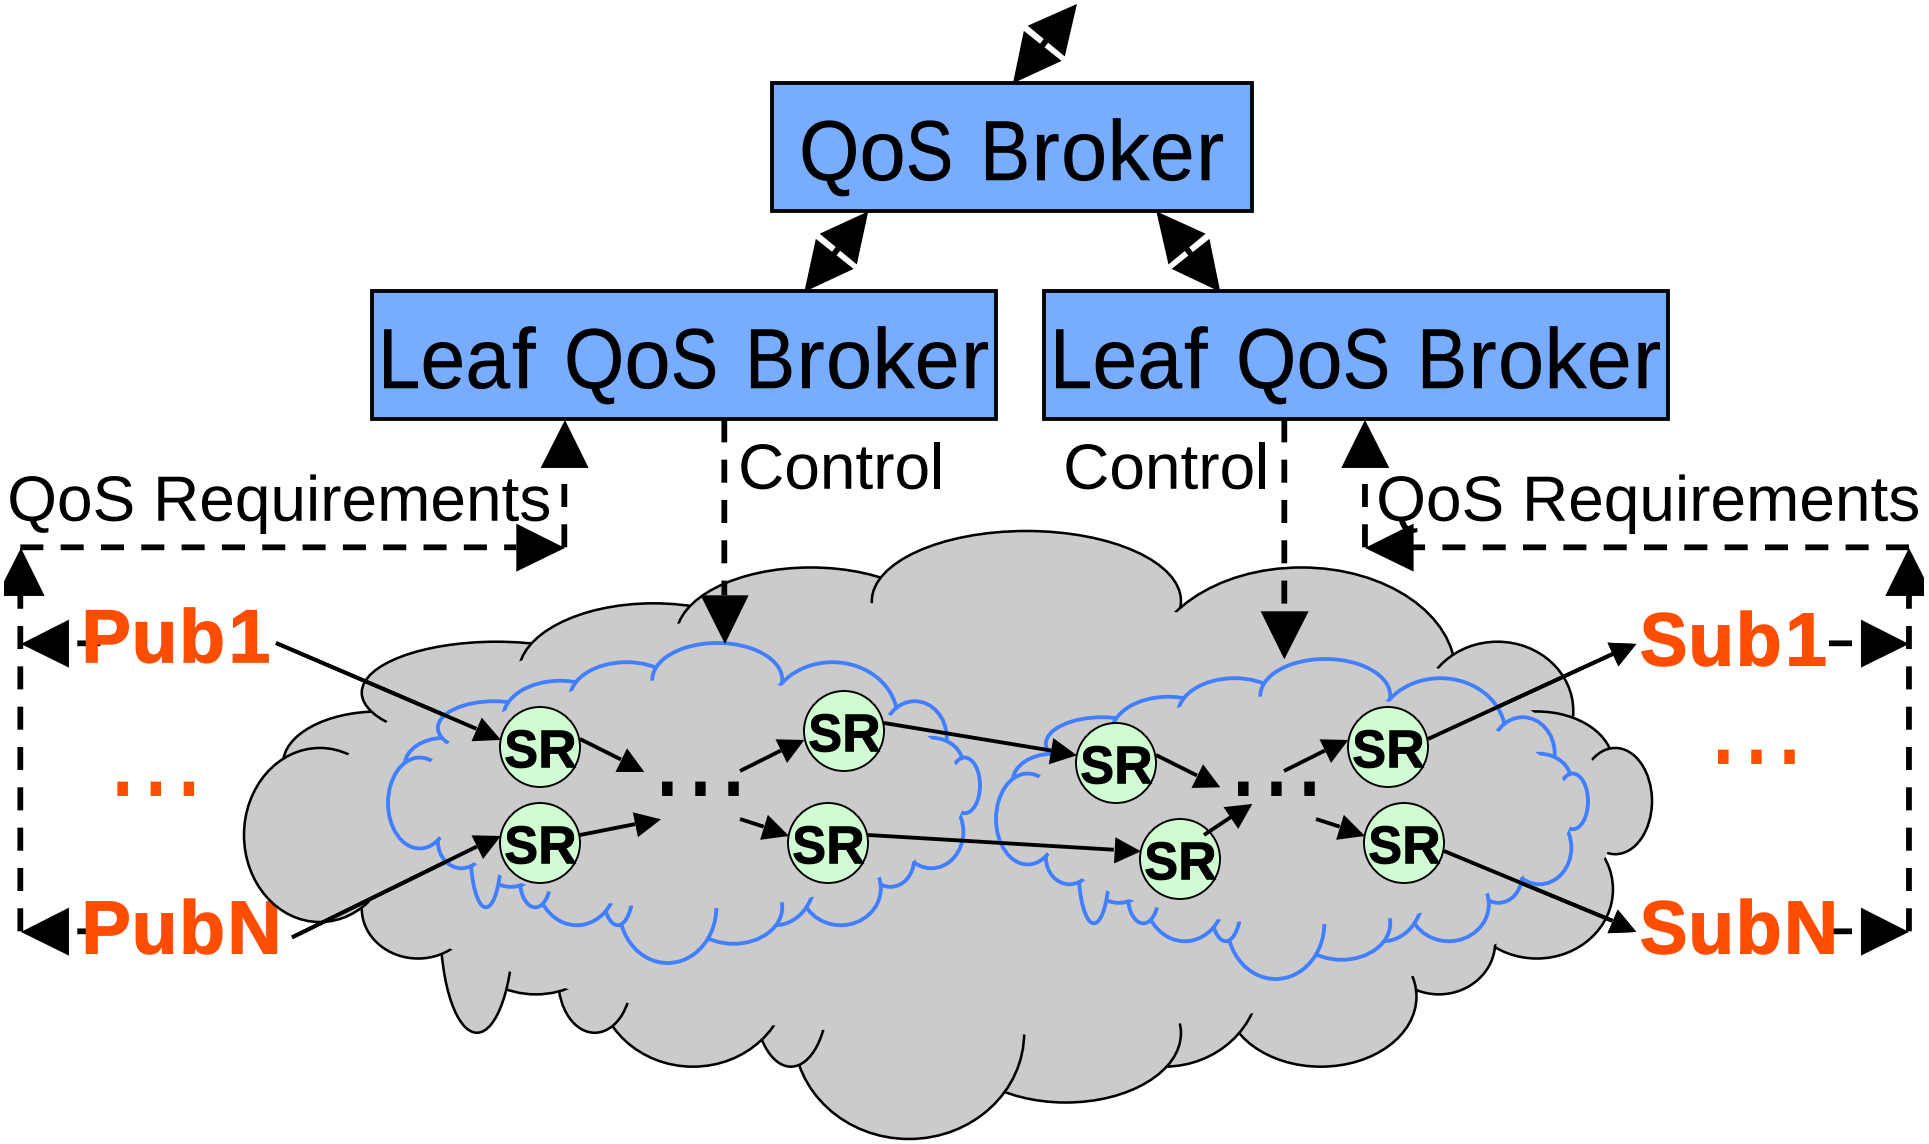

Note: GridStat handles routing decisions

Route Allocation to Subscriber 1

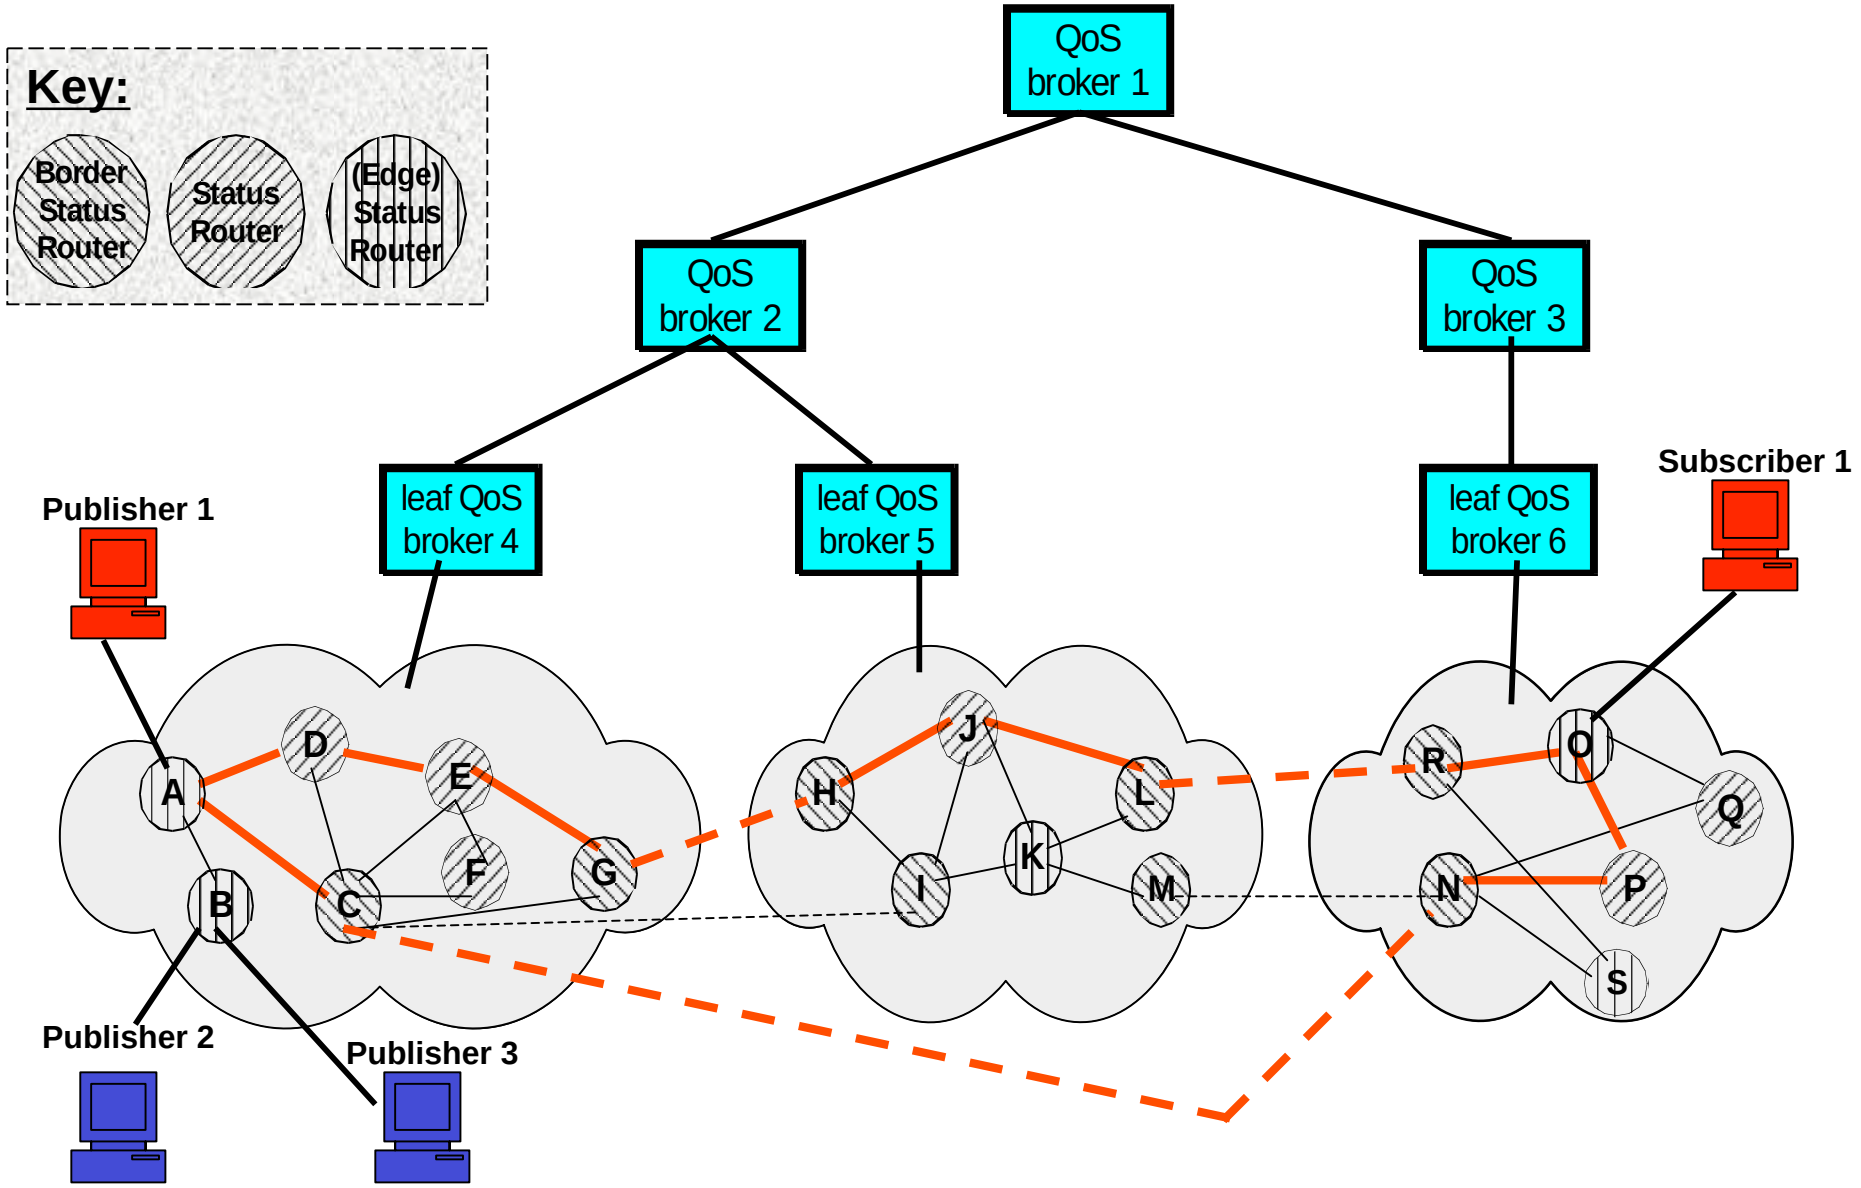

Route Allocation to Subscriber 2

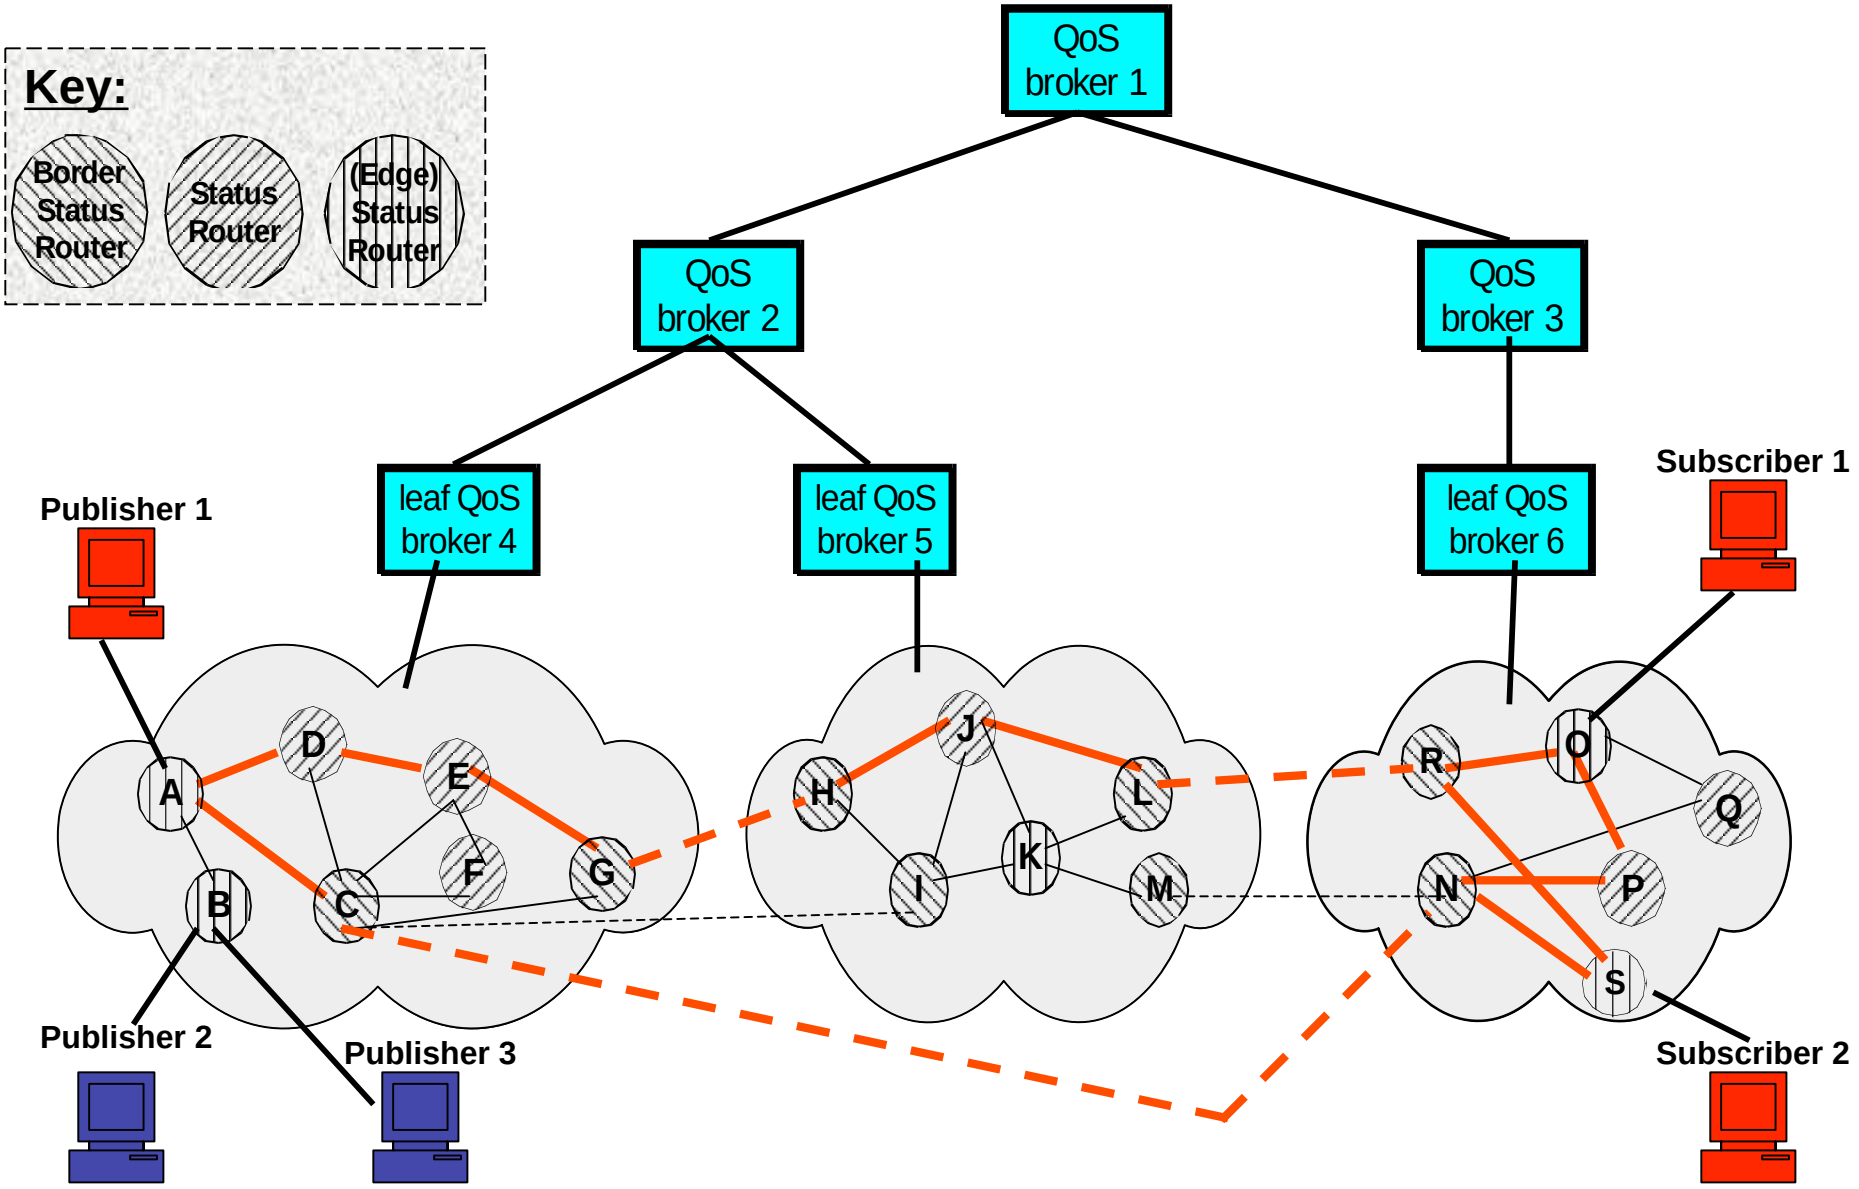

Note: Sub2 may have a different rate, latency, or redundancy than Sub1

Filtering and Multicast

- Proxy in publisher filters stream of update events down to highest subscribed rate
- Status routers provide both rate filtering and multicast path sharing
  - Filtering drops status update events while preserving downstream subscribers' delay and rate requirements
  - Multicast ensures that a given status update even only sent out once for all downstream subscribers
- Rate filtering and multicast together both preserve *temporal synchronism* across multiple status update flows
  - Required by phasor measurement units (PMUs) for power grid state estimation (really measurement!)
  - E.g., pass thru update #1, #11, #21, … for required flows
    - Must have same publisher rate (or multiple) & schedule (GPS)
- Future: filtering on change (% or Δ)

Status Semantics & Data Load Shedding

- Electric Utilities can do **load shedding** (I call **power load shedding**) in a crisis (but can really hurt/annoy customers)
- GridStat enables **Data Load Shedding**
  - Subscriber's desired & worst-acceptable QoS (rate, latency, redundancy) are already captured; can easily extend to add priorities
  - In a crisis, could shed data load: move most subscribers from their desired QoS to worst case they can tolerate (based on priority, and eventually maybe also the kind of disturbance)
  - Works very well using GridStat's operational modes
  - Note: this can prevent **data blackouts**, and also does not irritate subscribers
- Example research needed: systematic study of *data load shedding* possibilities in order to prevent *data blackouts* in contingencies and disturbances, including what priorities different power apps can/should have…

Condensation Functions

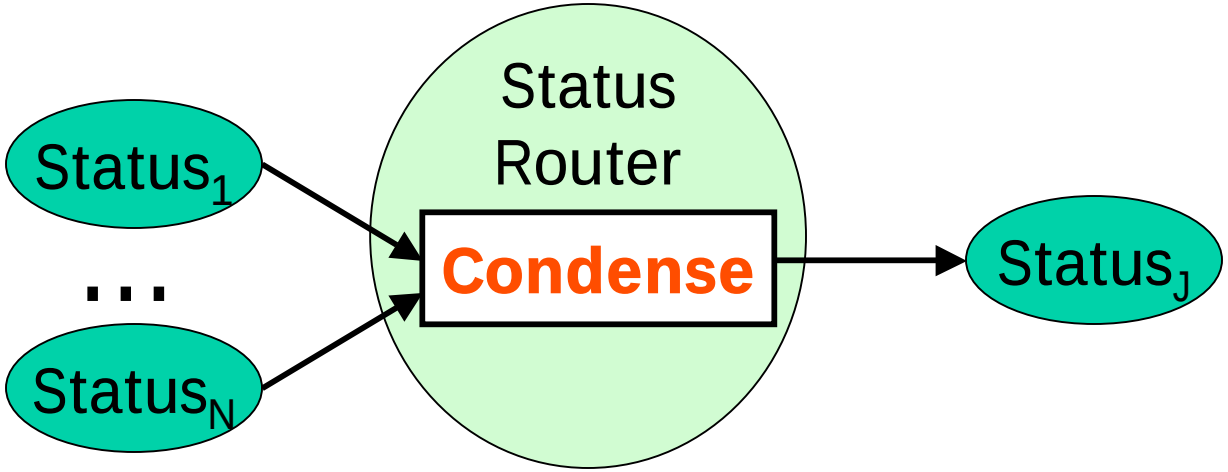

- *Condensation functions* allow applications to define new derived status variables
  - Sometimes subscribers just read a large set of status items once to calculate a derived variable
  - Supported by allowing user-defined condensation functions to be loaded in status routers
  - Building block for other mechanisms/capabilities
- Can be dynamically loaded into SRs

Condensation Functions (cont.)

- Creation with GUI-based tool
  - Specify input variables & four modules
- Modules
  - Input filter [optional]: filter status update events by value range
  - Trigger: initiates calculation; builtins:
    - Time triggered
    - Event triggered: received update events from x input variables
    - Alert triggered: received alerts from x of the subscribed input alert variables
  - Calculator
    - Init method
    - Calculation method
  - Output filter [optional]: like input filter
- Placed in cloud with input variables (present limitation)
- May evolve to status expressions (w/typing & QoS & inferrences)

Ongoing GridStat Research

- Ongoing GridStat Research
  - RPC over pub-sub with QoS & safety pre+post-conditions
  - Making modes global and hierarchical
  - Securing the multicast data plane
  - Securing the management plane
- Likely near-term work
  - Lots of likely short-term collaborations with other TCIP colleagues
  - Value error detection across multiple update paths
  - EC Framework Programme 7 collaboration

Related Work

- Key GridStat differentiators
  - Semantics of status updates
  - QoS management for rate, delay, redundancy
  - Rate filtering with multicast preserving temporal synchronism
  - Extensibility with application logic
- Pub sub frameworks (lots)
  - Real-time event channels
  - Content-based
- Power industry: IntelliGrid, UCA/IEC 61850,
- Probabilistic multicast (esp. gravitational gossip)
- CRUTIAL

Ongoing and Emerging Parnerships/Interest

- SEL
- Avista Utilities
- DoE EIPP (Eastern Interconnect Phasor Project)
- PNNL Electricity Infrastructure Operations Center (EIOC)
- INL SCADA Testbed
- TCIP Center (NSF CyberTrust, DoE, DHS August 2005)
  - Computer science award, working with power researchers
  - U. Illinois (headquarters)
  - Washington State University
  - Dartmouth College
  - Cornell University

Conclusions

- Interdisciplinary CIP research
  - Takes time and patience
  - Takes evangalization/outreach (and obvious learning)
  - Can be both frustrating and rewarding (usually lots of both!)
- GridStat is a flexible pub-sub middleware framework
  - Architected to be very flexible
  - Semantics of Status Dissemination
  - Managed for QoS
  - Demo in 2002, trial utility deployment since 2003
- Backup Slides:
  - Flexibility Needed for Grid Communications
  - A Few Examples of What GridStat Enables (above net. level)
  - More GridStat Details

For More Info

- Carl Hauser, David Bakken, and Anjan Bose. "A Failure to Communicate: Next-Generation Communication Requirements, Technologies, and Architecture for the Electric Power Grid", *IEEE Power and Energy*, 3(2), March/April, 2005, 47–55. Available via www.gridstat.net/intro.pdf
- David E. Bakken, Anjan Bose, Carl H. Hauser. *EC Efforts in SCADA-Related Research: Selected Projects.* Technical Report EECS-GS-008, Washington State University, 20 October, 2006. Available via http://www.gridstat.net/EC/EC-SCADA-CIP-Report.pdf
- IEEE Standard 1646, "IEEE Standard Communication Delivery Performance Requirements for Electric Power Substation Automation", 2004.

Backup Slides

- **Flexibility Needed for Grid Communications**
- A Few Examples of What GridStat Enables
- More GridStat Details

Next-Generation Grid Comms. Requirements

- **In summary: Flexibility and QoS!!!!**
- Status information can easily be made available to any legitimate participant at any location
- Status information is **predictably timely and reliable**
- Status information is protected against illegitimate use
  - Subscriber getting unauthorized status item
  - Subscriber "leaking" status info to others
  - Publisher sending bad status data (accidentally or otherwise)
- Crucial point: **you can't just "plug in a network"**
  - When you need it most it will be least available
  - Higher-level software needed for quality of service (QoS) management, IT failure recovery, adapting to cyberattacks, …
    - Dedicated fiber alone is insufficient … not an "end-to-end" solution

Flow of Operational Status Data [A. Bose]

- Much status data on the power grid exchanged for operation and control
  - Breaker status
  - Voltages (and some angles) at all buses
  - MW, MVAr at generators, load feeders and transmission lines
- Increasingly needs to go to multiple entities
  - Control center (plus backup control center)
  - Regional security coordinator (ISO/RTO)
  - Control centers of neighbors
  - Partners in ancillary services
  - Special controls or monitoring (SPS, WAMS, etc.)
- Data availability and usage depends on
  - Data measurement frequencies
  - Data transmission rates

Problems with Recent Trends [A. Bose]

- SPS/RAS is too expensive for widespread use
  - Hardwired communication is inflexible, changes require new installation
  - Even the settings require continual updating, which requires expensive off-line studies
  - Coordination of such piecemeal SPS installation is complex and error-prone
- WAMS design cannot be sustained for dramatic increase of PMU installation
  - PMUs are getting cheaper and will become part of local protection systems (e.g. SEL421)
  - Much higher bandwidth needed to move all that data
  - Data has to be moved to where the control is determined (rather than to some central controller)

Why Gridstat Flexibility Needed [A. Bose]

- All data collected at high frequencies cannot be brought into the central EMS/SCADA of one Control Area (let alone for the whole interconnection)
- The right data needs to go to the right computer at the right frequency depending on the function
- The functions and the data needs change over time and this arrangement for moving data must be very flexible
- The monitoring, operation, control and protection of the power grid should be changeable by software alone

Monitoring and Control w/Gridstat [A. Bose]

- Consider SPS/RAS
  - An existing SPS can be updated or a new SPS installed **solely by software**
    - Change input data
    - Change logic
    - Change output (control) signals
  - Instead of using off-line studies to set the controls every few months, use on-line computation to adapt the controls continually
    - Such on-line computation can be done using real-time data
    - Will need dedicated computer to do so
- Consider PMU and WAMS
  - PMU data could be handled just like any other data (the distinction is already blurring)
  - Monitoring of today can be extended to control tomorrow

The Crux of the Matter

- **Continued piecemeal expansion of the grid's communication capability is unnecessarily expensive and does not meet even today's requirements**
  - Lack of situation awareness major contributor to slow blackout response (US, Italy)
  - SPS/RAS deployment is very expensive
  - New control and protection schemes infeasible without better communications
- Desire: A more flexible alternative that can meet evolving communication needs of the grid
  - Without cheaper and more flexible communications, power researchers are unlikely to experiment with new communications topologies and control/protection schemes utilizing them
  - Without better control & protection schemes the full benefit of improving the grid's communications cannot be realized

Backup Slides

- Flexibility Needed for Grid Communications
- **A Few Examples of What GridStat Enables**
- More GridStat Details

Multi-Level Contingency Plannning & Adapting

- GridStat supports **operational modes**
  - Can switch routing tables very fast
  - Avoids overloading subscription service in a crisis
- Example: Applied R&D on coordinated
  1. Power dynamics contingency planning
  2. Switching modes to get new data for contingency
  3. New PowerWorld visualization specific for the contingency

  involving contingencies with
  A. Power anomalies
  B. IT failures
  C. Cyber-attacks
- Note: state of art and practice today: 1 & A only, offline

Example: Early-Warning System w/Triggers

- Simple benefit of GridStat: allow selective sharing of some key status variables, decided dynamically and
- Example: simple early-warning system
  - Cooperating power companies publish key leading indicators of problems, for cross-checking
  - Ideally: choose good indicators but not market sensitive
  - Virtually everything could be market sensitive, caution not sharing
- Solutions
  - #1: publish derived values (rate of change, …) not direct values
  - #2: alert-triggered temporary subscriptions for contingencies
  - #3: Add aggregation in QoS broker and policies to allow simple specification of thresholds of #alerts, etc.

  Note: #2 and #3 are not yet implemented, but quite doable in a year with 2 Computer Science MS projects & 1 companion EE power MS project.

Backup Slides

- Flexibility Needed for Grid Communications
- A Few Examples of What GridStat Enables
- **More GridStat Details**

GridStat Entities

- **Publisher**:
  - Publishes status events, the value of the status variable at that time.
  - Most of the publications are periodic, published at a given rate.
  - Some of the publications are alerts, which are only published when something unusual happens.
- **Subscriber**:
  - Subscribes to status variables by giving the name of the status variable along with its QoS requirements.
  - Will only receive updates for variables that it has subscribed to.
  - Multiple subscribers may subscribe to the same status variable.
  - Can request multiple redundant (disjoint) paths from publisher

GridStat entities (cont.)

- **Status Router**:
  - It forwards status events according to its routing table. The routing table is populated by the leaf QoS broker.
  - Like an IP router in that it forwards messages, but with additional ability for:
    - Optimized multicast
    - Operational modes
    - Filtering
    - Adaptive message packing
    - …
  - The set of status routers can be viewed as a message bus for status events

GridStat entities (cont.)

- **Leaf QoS broker**:
  - Controls one administration domain (its resources), called a cloud.
  - Allocates paths from the publishers to the subscribers that will satisfy the specified QoS requirements.
    - Does this by issuing commands to the SR in its domain to add/remove routing entries.
  - Communicates with the rest of the management through a connection to its parent QoS broker.

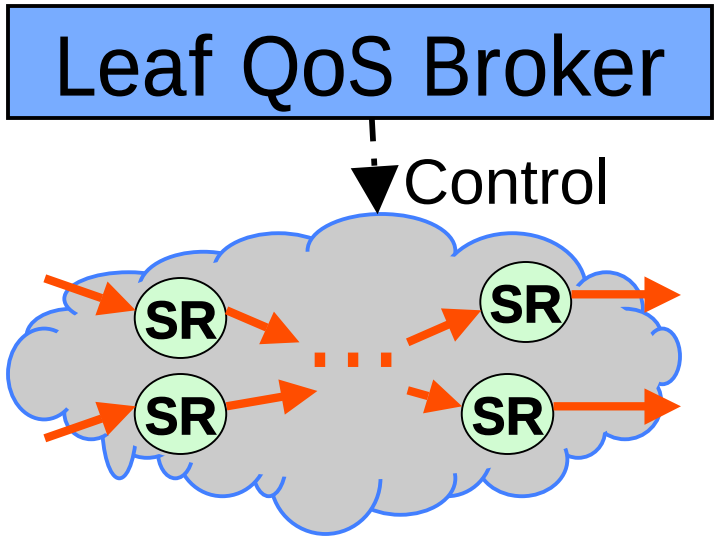

GridStat entities (cont.)

- **QoS broker**:
  - Hierarchical supervision of the leaf QoS brokers.
  - Controls the global resources provided by the different domains, through policies.
  - Allocates paths from the publishers to the subscribers (that are in different administration domains) that will satisfy the specified QoS requirements.
- Future: natural point for
  - Aggregation: more global view
  - More global decisions/adaptations

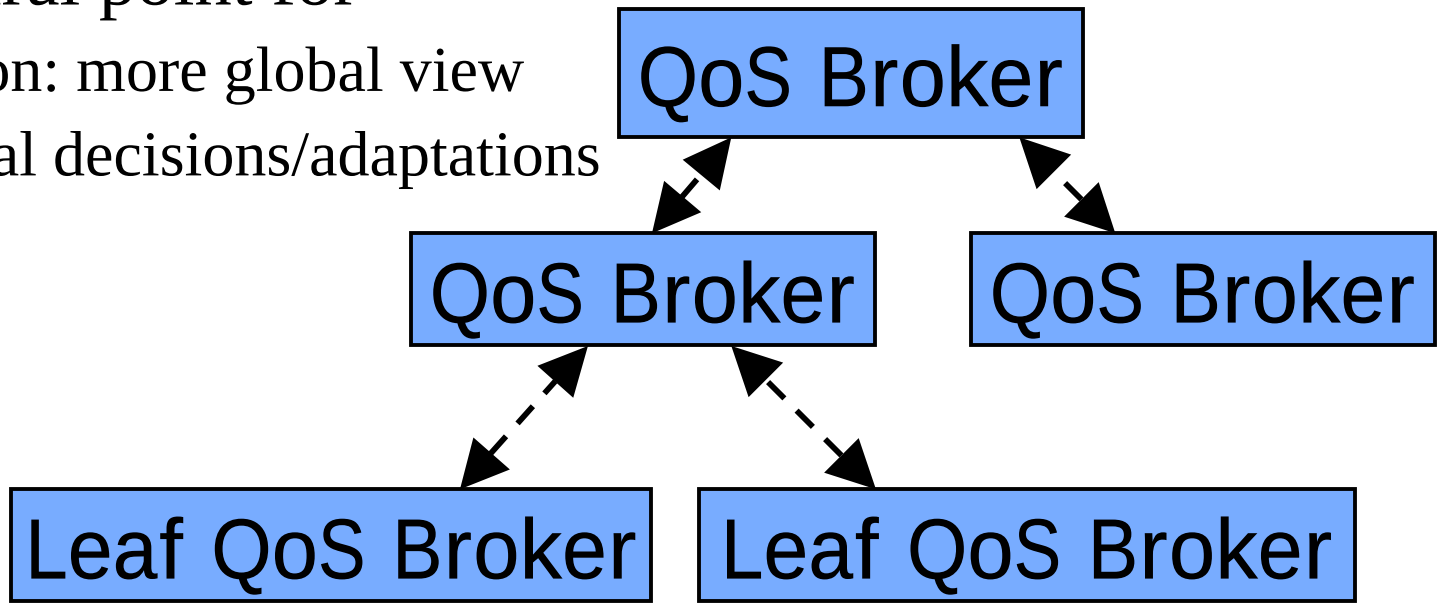

Alerts

- Report abnormal conditions requiring attention (or tracking)
- Bypass status variable queues at SRs with high priority
- Two types
  - Subscribed alert
  - Flooded alert
- Subscribed alert: similar to boolean status variable
- Flooded alert
  - Not subscribed to
  - Flooded USENET-style to a given level in QoS Broker hierarchy
- Trigger mechanisms presently implemented with condensation functions (later direct impl.)