

COMO PROJECT

**AN OPEN INFRASTRUCTURE
FOR NETWORK MONITORING**

GIANLUCA IANNACCONE

JOINT WORK WITH
CHRISTOPHE DIOT

**RIPE 49
THURSDAY**

**MANCHESTER, UK
SEPTEMBER 23 2004**

What is CoMo?

- High speed network monitoring system
- Large community working on monitoring
 - Many projects: IPMON, Gigascope, Hyperion, NProbe, Smart (UPC, Barcelona), NLANR, Abilene observatory, CAIDA
 - uncoordinated efforts, different datasets, difficult to validate other people's research, and data is never "enough"...

What is CoMo? (cont'd)

- CoMo as a building block of an open “trusted” infrastructure
 - Allows to share data and to run analysis methods on many diverse datasets
- Two major requirements:
 - Open, to provide access to data and to compute new metrics
 - Robust, to be always-on, handle anomalies in traffic and user behaviors

Challenge #1:

Ease of deployment

- It must be easy for users to deploy part of the infrastructure and access the rest
 - implement new method to analyse data traffic
 - availability of diverse datasets
- Many current design choices are driven by the need to keep things simple
- These aspects will drive the success of CoMo

Challenge #2: Resource Management

- Conflicting needs:
 - Open to users to run traffic analysis
 - Always available and give minimum performance guarantees
- How to define priority among user requests?
- Strict division of labor among system components
 - separate functions that need to run in real-time
 - push functions down to hardware (e.g., IXPs)

Challenge #3: Security, Privacy and Confidentiality

- Two levels of security
- Data security
 - How to preserve privacy of network users?
 - How to preserve confidentiality in network data?
 - Difficult to identify confidential information
 - Diverse sets of users require rich policy language
- Infrastructure security
 - Attacks on the resource management (denial-of-service)
 - Attacks on the access policies

Challenge #4:

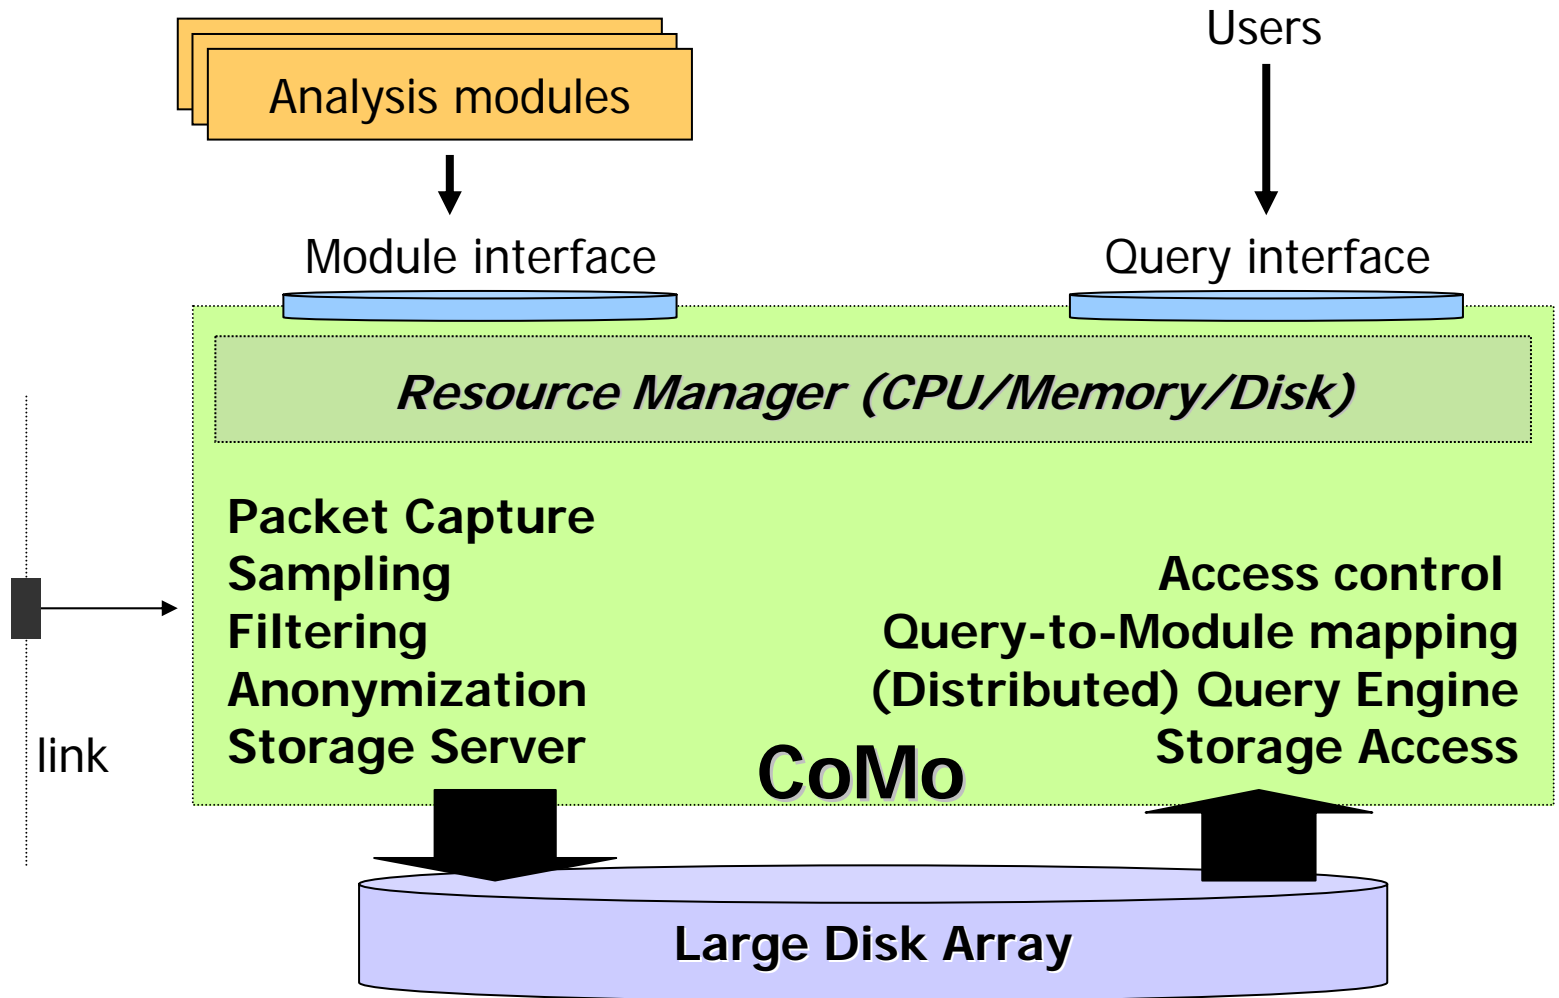
Querying the network data

- Express the query and run it without explicit built-in support
 - Difficult to predict the type of queries in advance
 - Need to allow custom-built queries that are not easy to express in a typical query language
- Managing the queries
 - Query will interest the entire CoMo network
 - CoMo system maintains large amount of data
 - Need methods to find the data (raw or pre-processed)

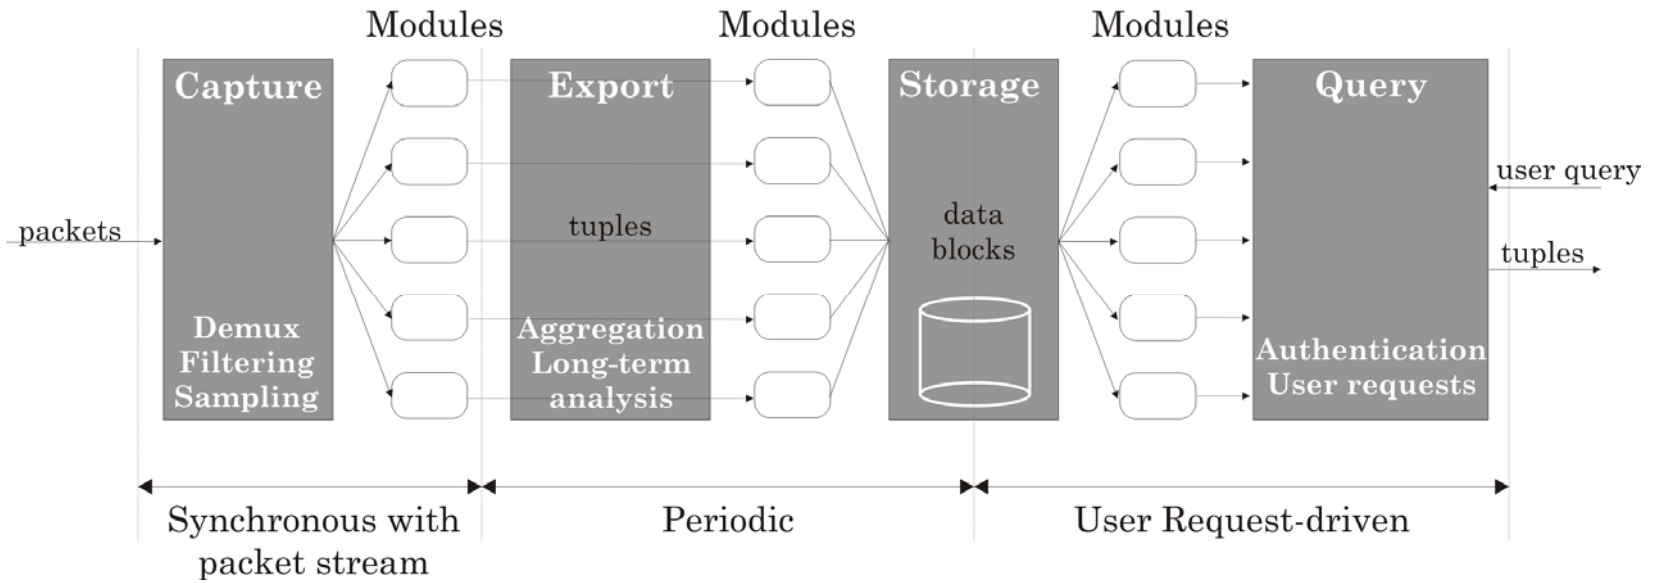
CoMo Architecture

- Core engine provides basic primitives
 - an interface to plug-in the module
 - an interface to access the data computed by the modules
- Traffic metrics reside in plug-in modules
 - modules compute one specific metric on the traffic data
 - modules activated on demand
 - modules pre-compute queries
 - users can “push” module to infrastructure

Functional view



Data flow view



- Strict decoupling between processes
 - Capture & Storage meant to isolate processes
 - Modules organized as set of callbacks
 - Each process could run on different hardware (e.g., Capture on IXP platform)

Module design and requirements

- Modules are independent of each other
 - No state information passed between modules
 - Implemented independently by users
- Modules are very restricted on what they can do
 - limited access to system calls, memory, etc.
- Modules can be shutdown at any time.
 - Under the assumption that packet trace is always collected
 - Module can run at later time on the store trace

Query support

- CoMo maintains a large amount of data
 - Modules pre-compute as much as possible
 - Random query needs to find the module that is pre-computing the data
- Static query: 'sendto <ip>:<port>'
 - it knows the module by configuration
- On-demand query: carries the module
 - query goes with source code
 - if no module matches, run the code
- Ad-hoc query: module compiled on the fly
 - query written in specific language (SQL-like)

Partner Projects

- NProbe (Cambridge)
- SMARTxAC (UPC, Barcelona)
- Hyperion (UMass)
- Diamond (Intel Research Pittsburgh)
- iDIM (USC/Intel)
- GEANT Monitoring (DANTE/Intel)

Related Projects

- DANTE Monitoring Infrastructure
- CAIDA's Internet Measurement Data Catalog
- Query support
 - IRISNET (Intel Research Pittsburgh)
 - TelegraphCQ (Berkeley)
 - Aurora (MIT)

CoMo Donation Program

- CoMo systems to universities and other private organizations
 - Intel-based platform + Endace DAG card
 - Current budget for 10+ systems
- Constitute core of CoMo infrastructure
 - Access to the system will be shared with other CoMo participants
 - Systems will also be partially open to public (details defined on case-by-case basis)

Status & Deployment

- Software released under BSD license
- IRC is online
- Intel IT pilot project
 - EBC datacenter (FM7) in Folsom, CA
 - WAN link on Intel Corporate Network
- Core set of CoMo nodes
 - Intel Research Lablets
 - UMass, Univ. Pisa, Cambridge
 - working on anonymization

More info @
www.cambridge.intel-research.net/como

CoMo PROJECT
@IntelResearch
Cambridge

*All truths are easy to understand
once they are discovered;
the point is to discover them!*
-- Galileo

Home
People
Publications
Talks
Software
Related Projects

Overview

CoMo (Continuous Monitoring) is a passive monitoring system. CoMo has been designed to be the fundamental building block for a network monitoring infrastructure that will allow researchers and network operators to easily process and share network traffic statistics over multiple sites. The architecture of CoMo is designed to compute and report various performance metrics while sustaining high speed traffic collection. There is a strict decoupling between the CoMo core system and the measurement modules. The core system moves the data from the network to memory and storage and manages the system resources. The measurement modules simply pick packets of interest and compute a specific traffic metric. CoMo is open in the sense that modules can be implemented independently and from different developers and then plugged into the systems dynamically. CoMo also provides a query interface to allow users to elicit the system to export the results of the measurement performed.

Announcements

September 10, 2004
First software release
The software comes with a BSD type license. It supports standard NIC

CoMolive!

Friday, September 10th, 2004
Systems: 1
Storage space: 350GB
Packets observed: 23,276
Bytes observed: 31,435,400

[go to CoMolive!](#)

News

August 24th, 2004
CoMo system installed on Intel Research Cambridge access link.

Upcoming Events

RIPE Meeting
Manchester, UK
September 20th-24th, 2004
<http://www.ripe.net>

Contact Info

christophe.diot@intel.com
gianluca.iannaccone@intel.com