

The logo of the University of Jyväskylä is a vertical blue bar on the left side of the page. It contains a repeating pattern of white stylized symbols: a flame-like shape above a horizontal line, and a vertical line below it.

UNIVERSITY OF JYVÄSKYLÄ

Cheddar

Matthieu Weber (mweber@mit.jyu.fi)

Cheese Factory Project, InBCT 3.2
Agora Center

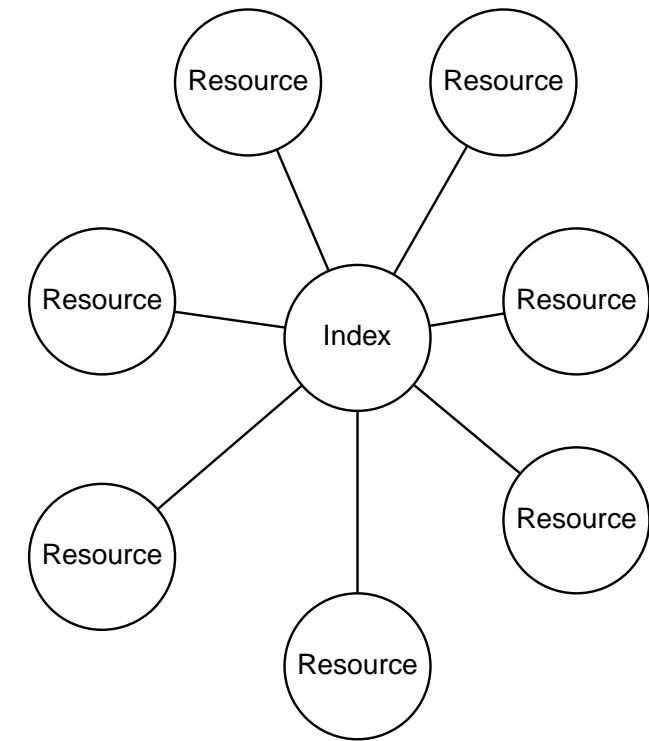


Introduction

Origins

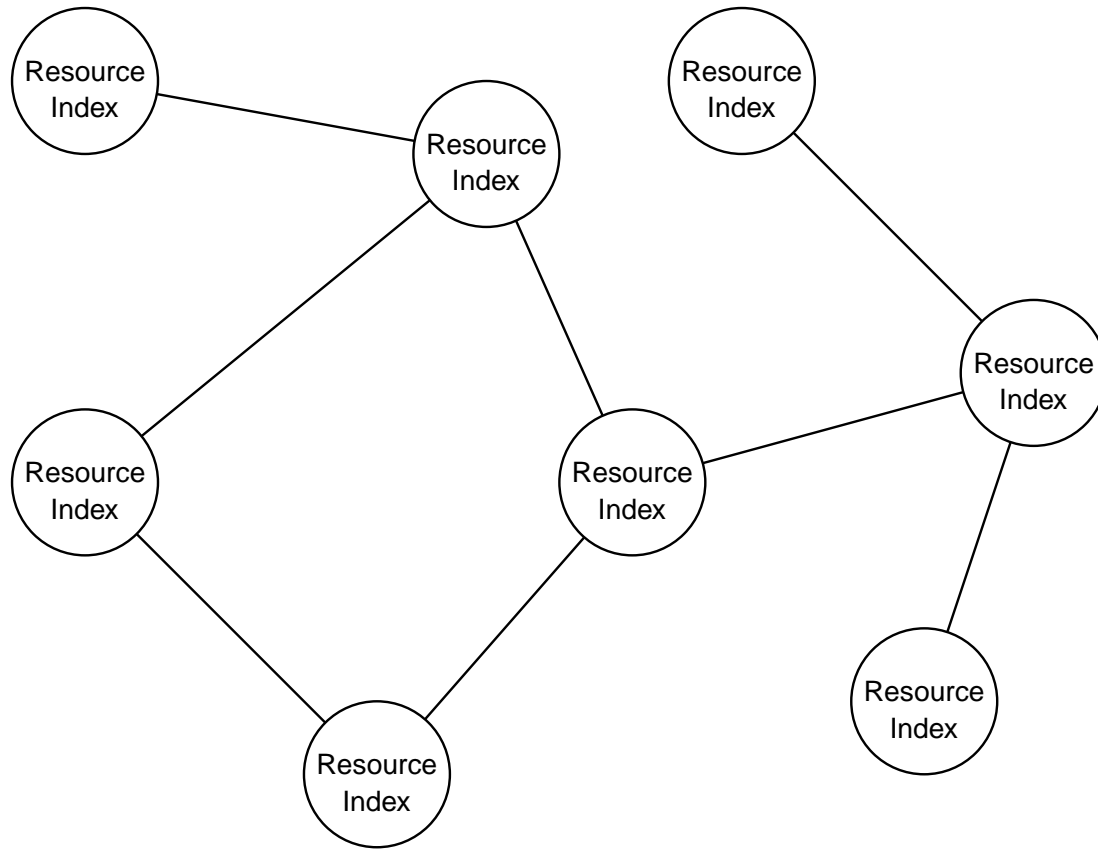
- CheDAR: a CHEap Distributed ARchitecture
 - Completely distributed computing platform
- Globus: grid computing system, with centralized resource discovery
- Chedar refocused to be a fully distributed resource discovery system

Centralized Vs. Distributed Systems




Centralized Resource Discovery

Centralized Vs. Distributed Systems



Distributed Resource Discovery



Centralized Resource Discovery

Resources

- Resource discovery, based on specific attributes
- Various types of resources:
 - Computing resource (available computing time)
 - File resource (a given file, free storage space)
 - Display resource (a screen)
 - Printing resource (a printer)
- Query/reply messages



Architecture

Chedar's Components

- ChedarClient
- Connection
- Connection Manager
- Topology Manager
- Propagation Engine

ChedarClient

- Control interface of a Chedar node
- Applications and user interfaces communicate with Chedar through the ChedarClient component
- Used by:
 - P2PStudio
 - Data fusion application

Connection

- Represents a connection with another node
- Based on JXTA's pipes
- Monitors the “health” of the connection

Connection Manager

- Keeps a track of former neighbors (connections)
- Establishes/drops connections based on:
 - the connection's health
 - the topology manager's decisions

Topology Manager

- Optimizes the network's topology locally
- Decides where to establish new connections (overtaking)
- Based on the quality of the relations with the neighbors (node rating)

Propagation Engine


- Handles the resource queries and replies
- Answers queries when they match the local resources
- Forwards queries to other nodes when it cannot answer it



Algorithms


Basic Search

- Generates a lot of traffic




Basic Search

- Generates a lot of traffic




Basic Search

- Generates a lot of traffic



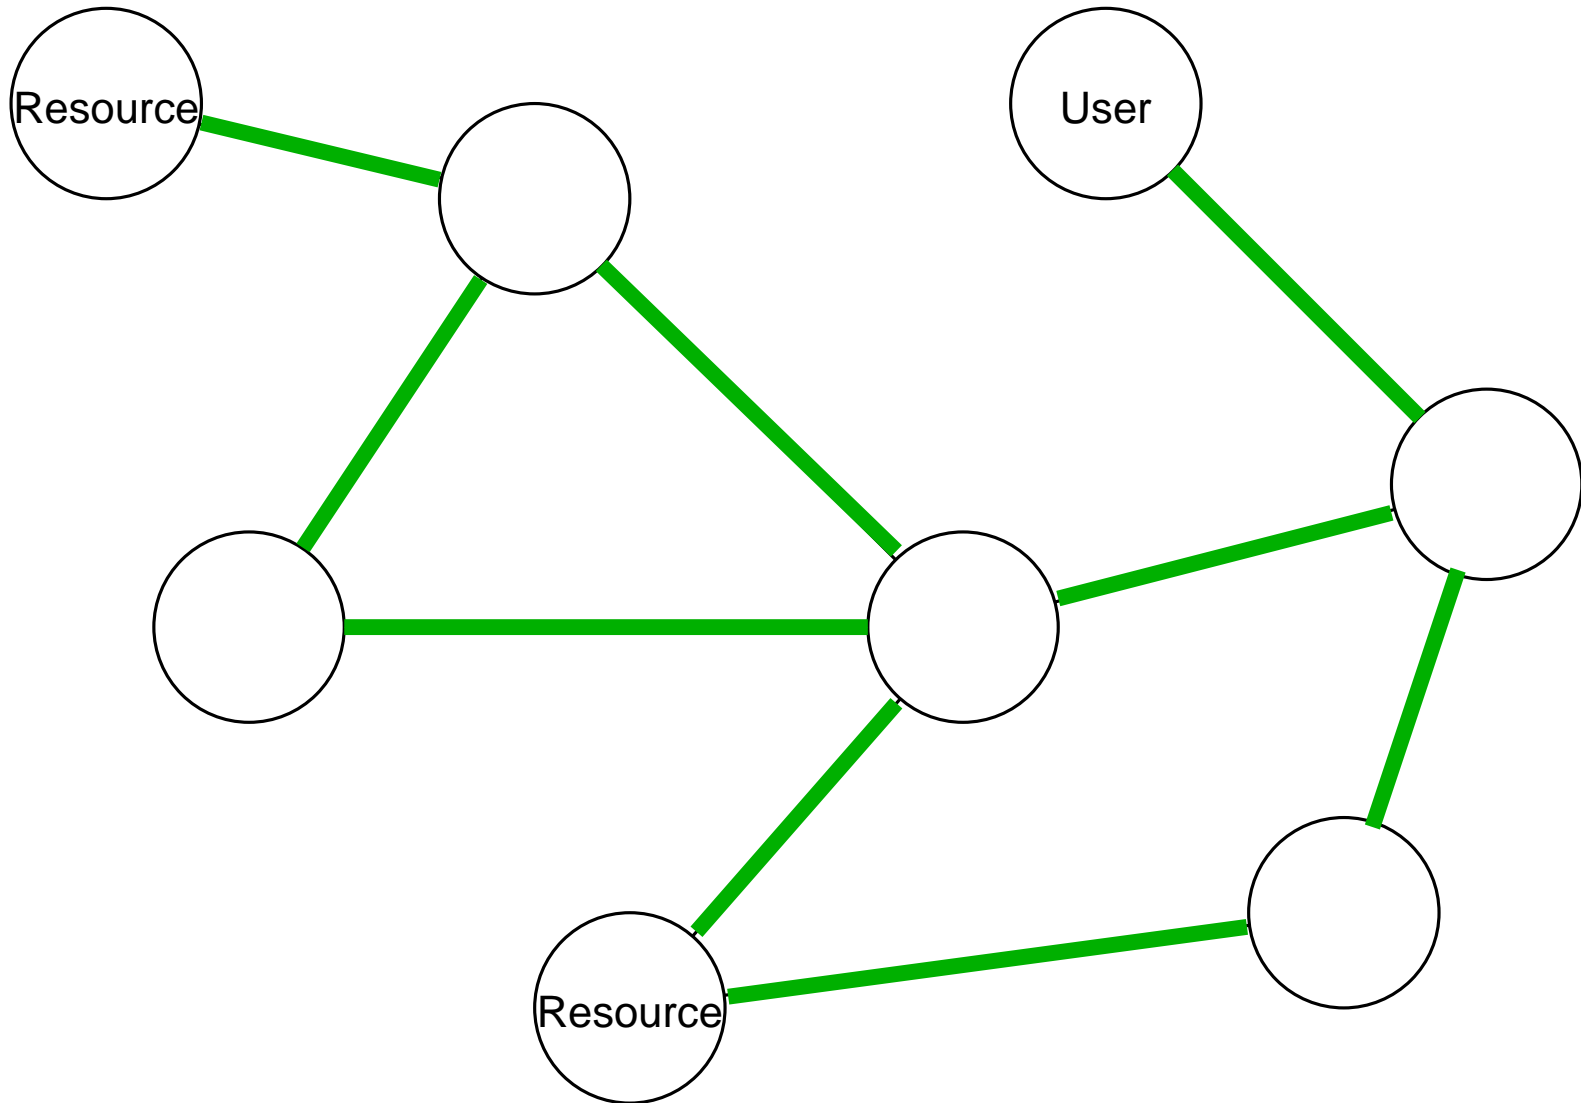
Basic Search

- Generates a lot of traffic




Basic Search

- Generates a lot of traffic



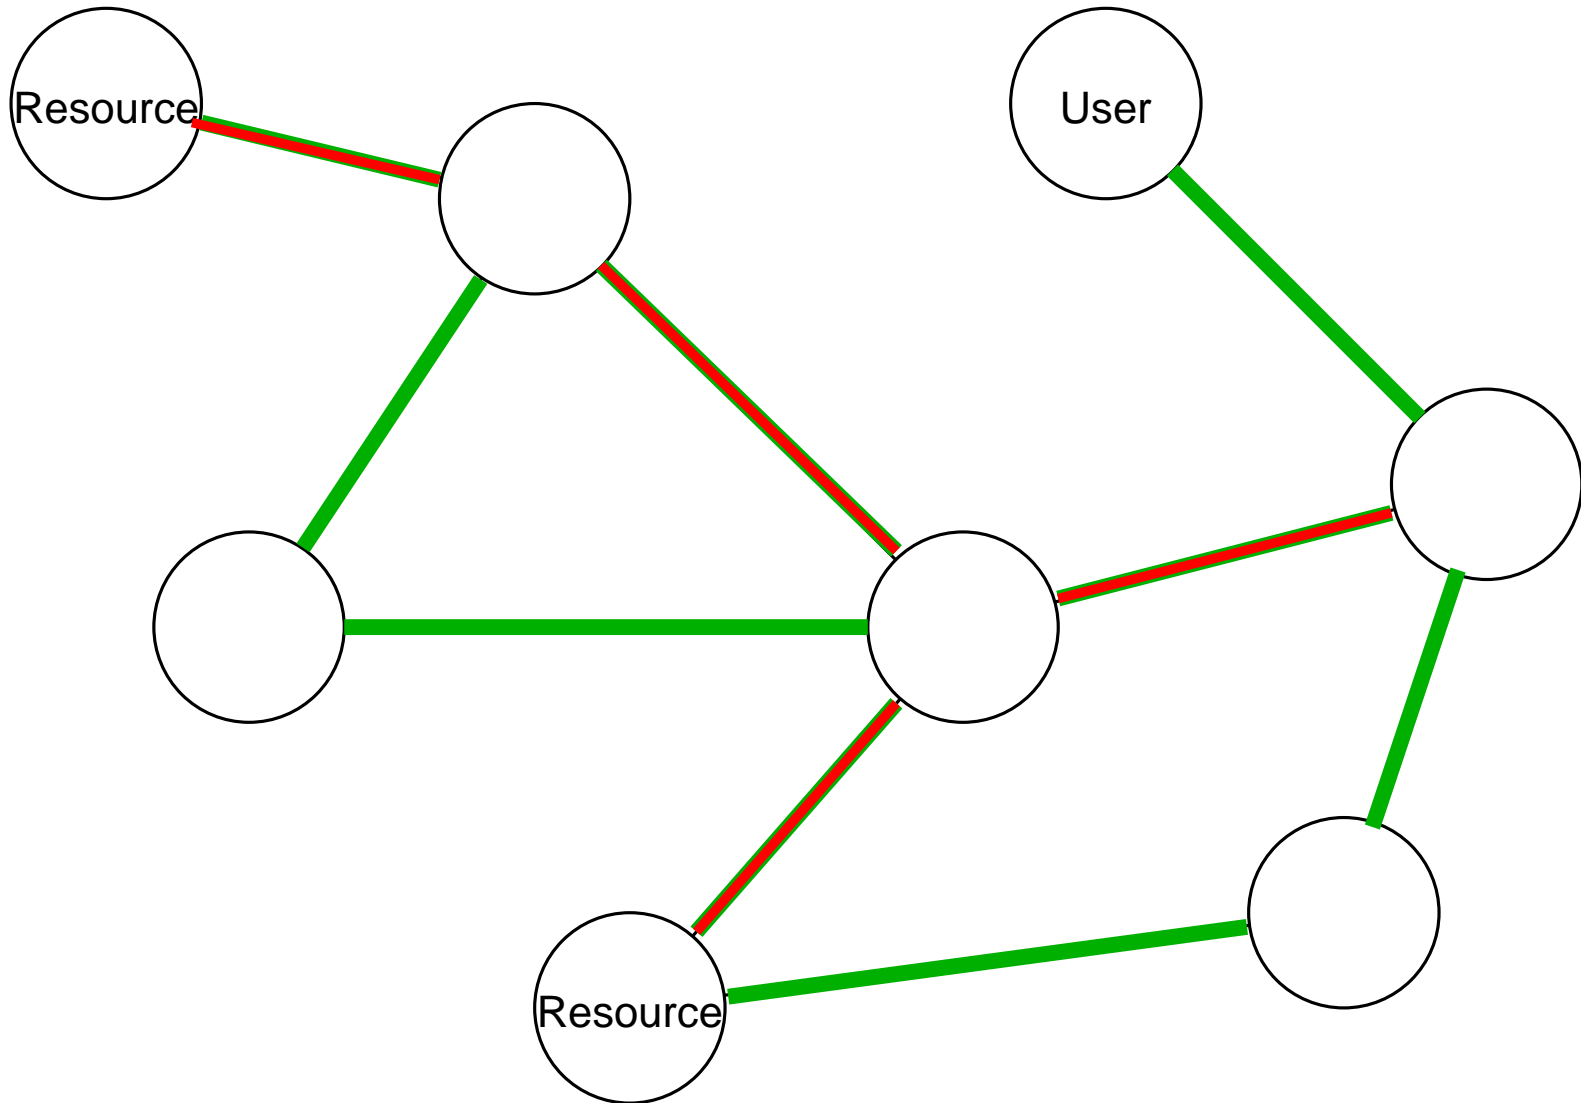
Basic Search

- Generates a lot of traffic




Basic Search

- Generates a lot of traffic



Basic Search

- Generates a lot of traffic



Node Rating

- Ratings are given to the neighbors:
 - amount of replies
 - number of neighbors
 - reply speed
- Node selection based on those ratings
- Improves the search's efficiency

Overload Limitation


- Limits the propagation of messages
- More adaptive than time-to-live
 - Each message contains a hop count
 - Each node accepts messages whose hop count is below a given limit
 - That limit is broadcast to the neighbors, which must respect it

Return Path

- 3 ways to reply to a query:
 - reverse the query's path
 - reverse path with “jumping” over unavailable nodes
 - direct connection to the requester
- Increases the fault-tolerance

Overtaking Algorithm

- The nodes that are not valuable are overtaken



Overtaking Algorithm

- The nodes that are not valuable are overtaken

