



UNIVERSITY OF
CAMBRIDGE

Opportunity is the Mother of Invention

How Personal Delay Tolerant Networking led to Data Centric Networking & Understanding Social Networks.

Jon Crowcroft

Jon.crowcroft@cl.cam.ac.uk

Outline Narrative History of Hagggle

1. Hagggle Software Architecture
2. How we got to Declarative Data Driven Nets
3. Why we got diverted into Social Networks

I have 100M bytes of data, who can carry for me?



Don't give to me! I am running out of storage.



Give it to me, I have 1G bytes phone flash.




Thank you but you are in the opposite direction!

I can also carry for you!




Reach an access point.




There is one in my pocket...



Search La Bonheme.mp3 for me




Search La Bonheme.mp3 for me



Search La Bonheme.mp3 for me



Internet



Finally, it arrive...



1. Motivation 2001-2004

- Mobile users currently have a very bad experience with networking
 - Applications do not work without networking infrastructure such as 802.11 access points or cell phone data coverage
 - Local connectivity is plentiful (WiFi, Bluetooth, etc), but very hard for end users to configure and use
- Example: Train/plane on the way to London
 - How to send a colleague sitting opposite some slides to review?
 - How to get information on restaurants in London? (Clue: someone else is bound to have it cached on their device)
- Ad Hoc Networks were a complete washout
 - Failed to account for heavy tailed density distribution
 - Use of 802.11 as radio was at best misguided.

Underlying Problem

- Applications tied to network details and operations via use of IP-based sockets interface
 - What interface to use
 - How to route to destination
 - When to connect
- Apps survive by using directory services
 - Address book maps names to email addresses
 - Google maps search keywords to URLs
 - DNS maps domain names to IP addresses
- Directory services mean infrastructure

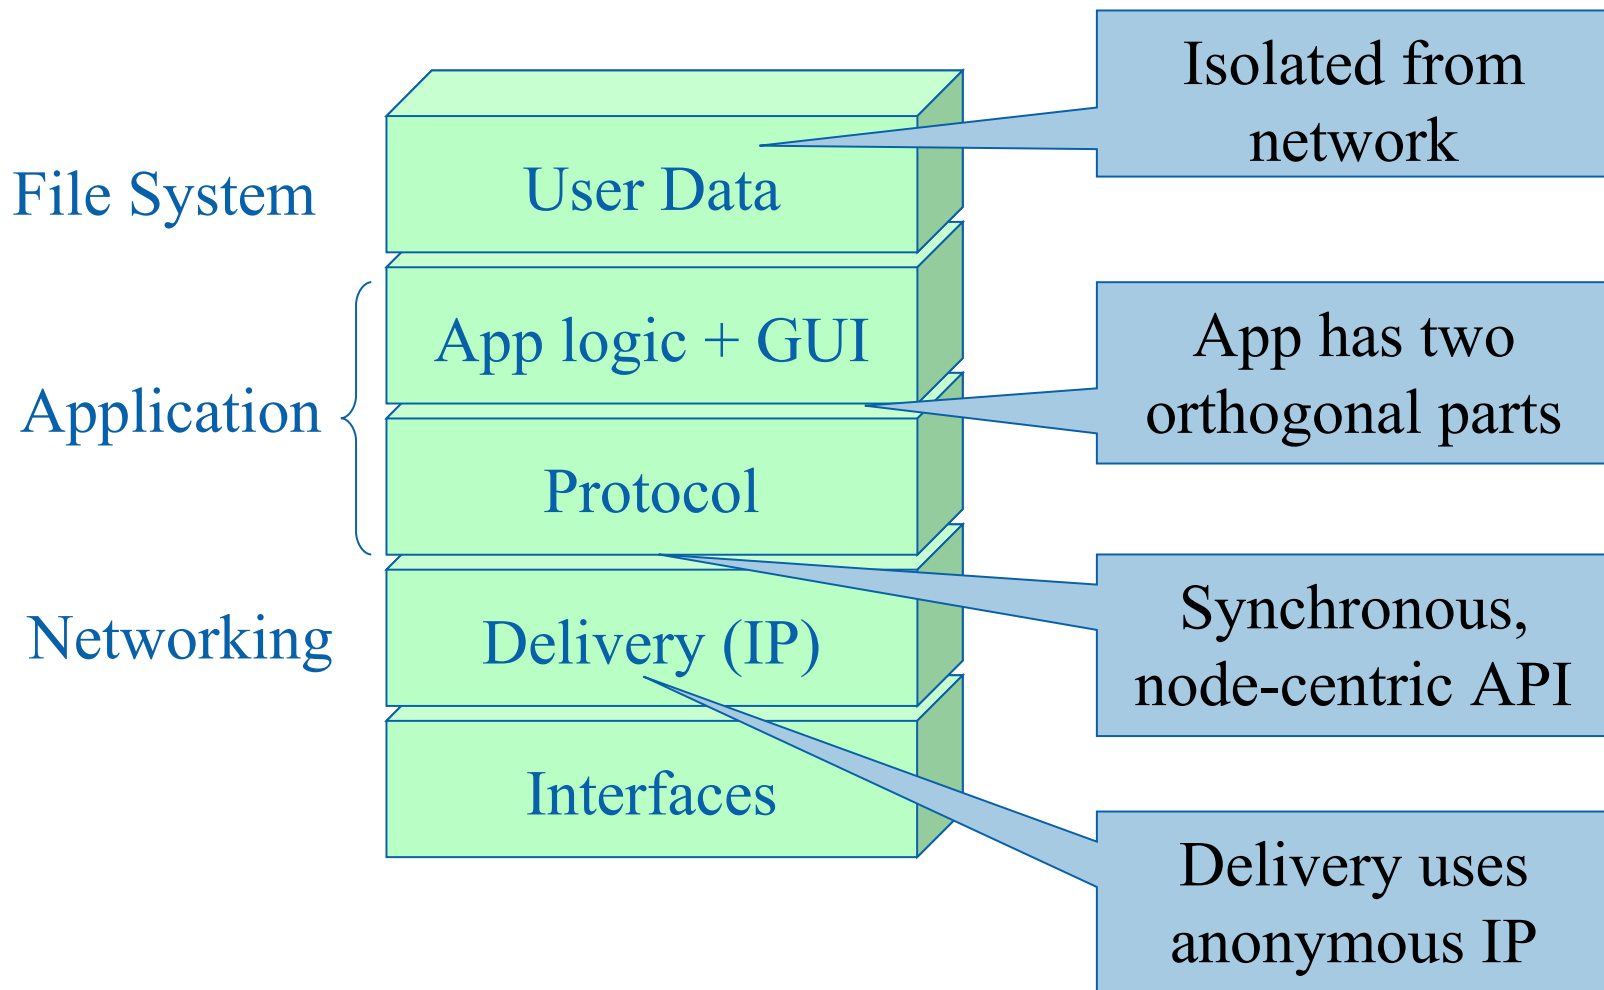
Phase transitions and networks

- Solid networks: wired, or fixed wireless mesh
 - Long lived end-to-end routes
 - Capacity scarce
- Liquid networks: Mobile Ad-Hoc Networking (MANET)
 - Short lived end-to-gateway routes
 - Capacity ok (Tse tricks with power/antennae/coding)
- Gaseous networks: Delay Tolerant Networking (DTN), Pocket Switched Networking (PSN)
 - No routes at all!
 - Opportunistic, store and forward networking
 - One way paths, asymmetry, node mobility carries data
 - Capacity Rich (Grossglauser&Tse) (but latency terrible... ..)
- Hagle targets *all three*, so must work in most general case, i.e. “gaseous”

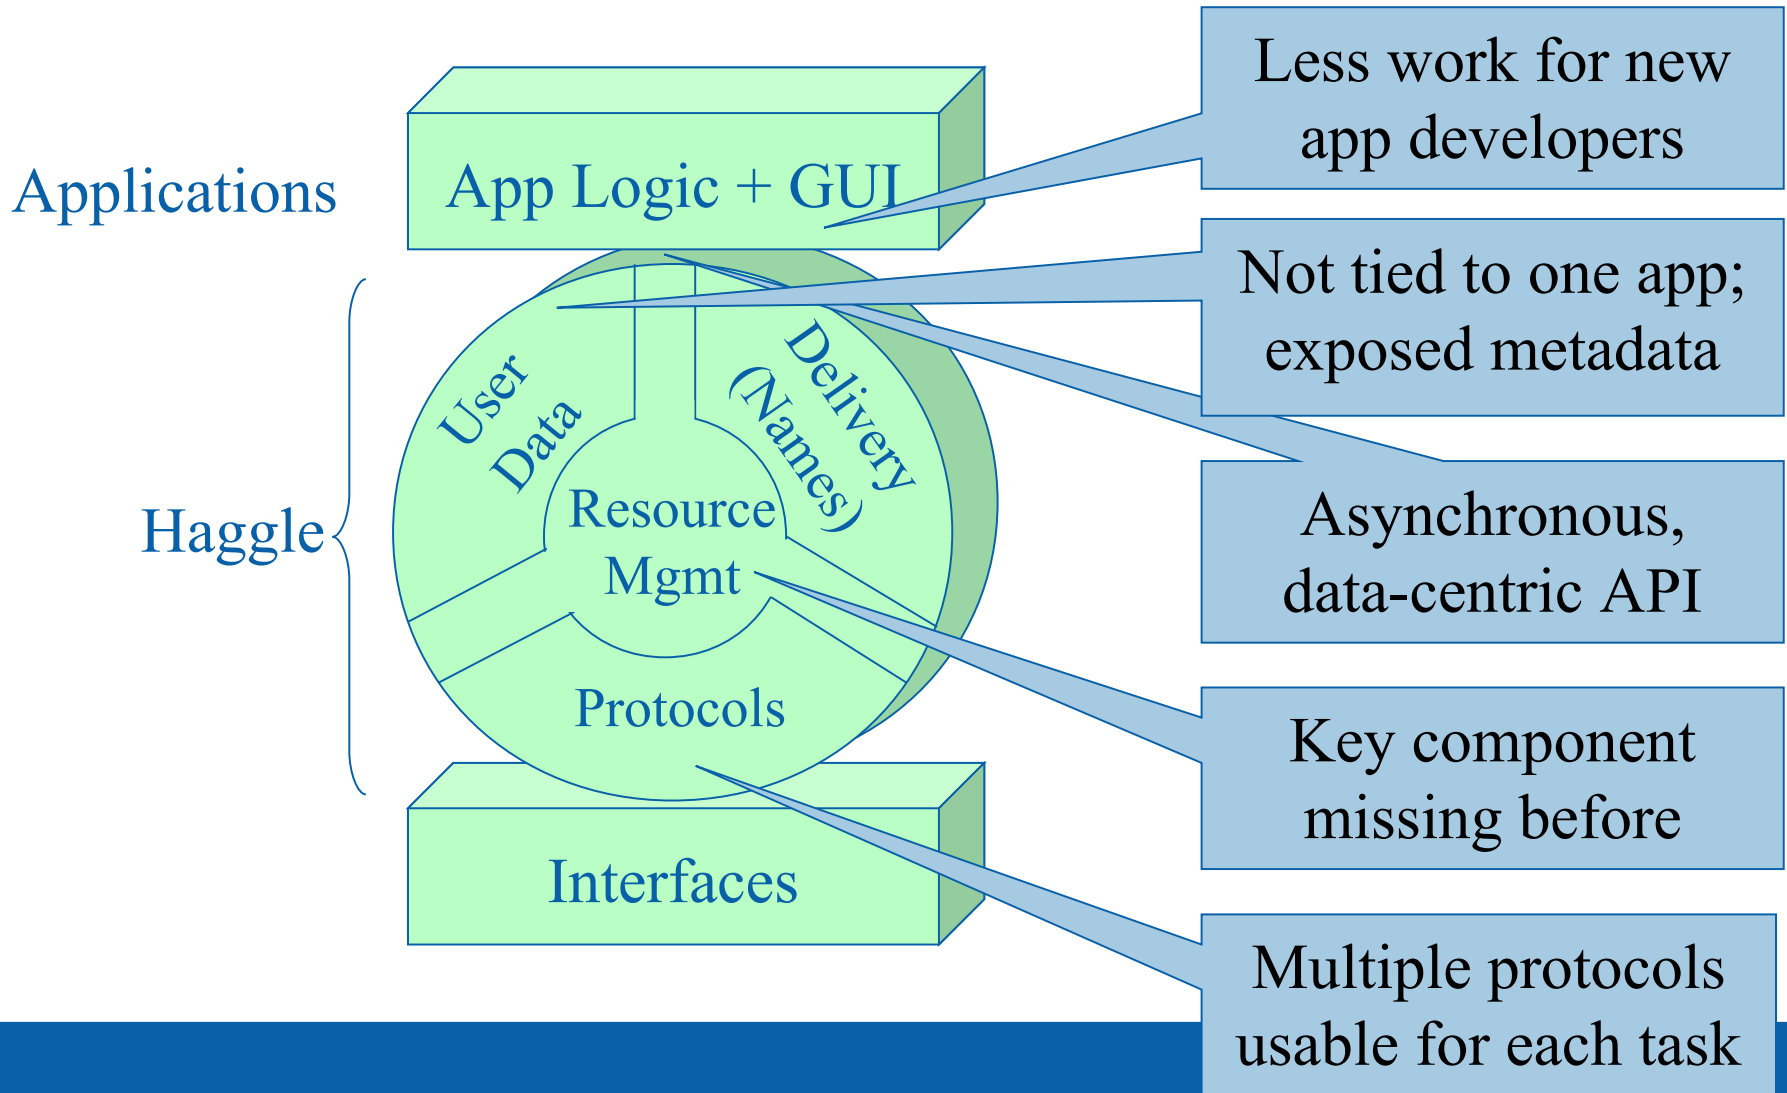
Decentralisation & Disconnectivity

- Absence of infrastructure for
 - Routing, searching, indexing
 - Names, Identity, Currency
- When everything's adhoc, even pagerank has to be
 - Hence "Ad Hoc Google" -> "Haggle" Intel Cam 2004.
 - Bad joke about french pronunciation of "Haddock"
- As early pub/sub systems, interest itself is data
 - So we take event/notify+pub/sub and apply to
 - Discovery of users, nodes, routes, interest
 - everyone soaks it all up and runs ego-centric pagerank

Current device software framework



Haggle framework design



Data Objects (DOs)

- DO = set of attributes = {type, value} pairs
 - Exposing metadata facilitates search
 - Another bad (Diot) joke
- Can link to other DOs
 - To structure data that should be kept together
 - To allow apps to categorise/organise
- Apps/Haggle managers can “claim” DOs to assert ownership

Message

DO-Type	Data
Content-Type	message/rfc822
From	James Scott
To	Richard Gass
Subject	Check this photo out!
Body	[text]

Attachment

DO-Type	Data
Content-Type	image/jpeg
Keywords	Sunset, London
Creation time	05/06/06 2015 GMT
Data	[binary]

DO Filters

- Queries on fields of data objects
- E.g. "content-type" EQUALS "text/html" AND "keywords" INCLUDES "news" AND "timestamp" \geq (now() - 1 hour)
- DO filters are also a special case of DOs
- Huggle itself can match DOFilters to DOs – apps don't have to be involved
- Can be *persistent* or be *sent remotely*...

DO Filter is a powerful mechanism

	<i>One-Off</i>	<i>Persistent</i>
<i>Local</i>	"Desktop" Search (find mp3s with artist "U2")	Listen (wants to receive webpages)
<i>Remote</i>	"Web" Search (find "london restaurants")	Subscribe (send all photos created by user X to X's PC)

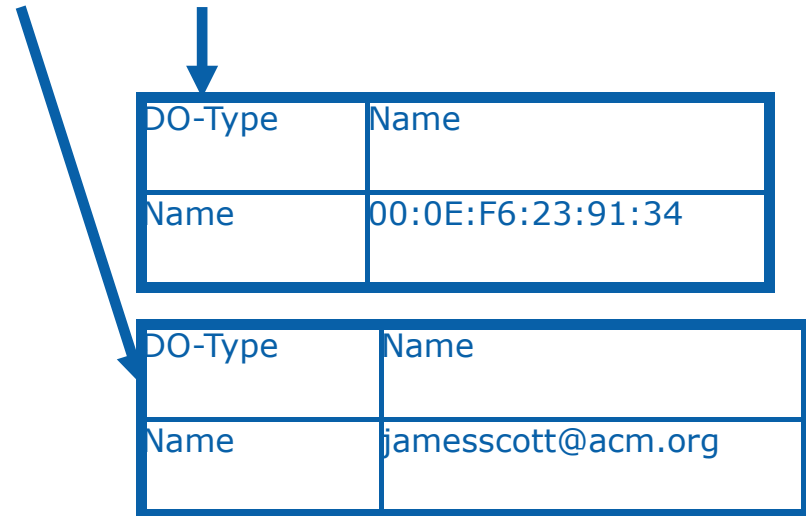
Layerless Naming

- Huggle needs just-in-time binding of user level names to destinations
- Q: when messaging a user, should you send to their email server or look in the neighbourhood for their laptop's MAC address?
 - A: Both, even if you already reached one. E.g. you can send email to a server and later pass them in the corridor, *or* you could see their laptop directly, but they aren't carrying it today so you'd better email it too...
- Current layered model requires ahead-of-time resolution by *the user themselves* in the choice of application (e.g. email vs SMS)

Name Graphs comprised of Name Objects


- Name Graph represents full variety of ways to reach a user-level name
- NO = special class of DO
- Used as destinations for data in transit
- Names and links between names obtained from
 - Applications
 - Network interfaces
 - Neighbours
 - Data passing through
 - Directories

DO-Type	Name
Name	James Scott



Forwarding Objects

- Special class of DO used for storing metadata about forwarding
 - TTL, expiry, etc
- Since full structure of naming and data is sent, “intermediate” nodes are empowered to:
 - Use data as they see fit
 - Use up-to-date state and whole name graph to make best forwarding decision



Connectivities and Protocols

- Connectivities (network interfaces) say which “neighbours” are available (including “Internet”)
- Protocols use this to determine which NOs they can deliver to, on a per-FO basis
 - P2P protocol says it can deliver any FO to neighbour-derived NOs if corresponding neighbour is visible
 - HTTP protocol can deliver FOs which contain a DOFilter asking for a URL, if “Internet” neighbour is present
- Protocols can also perform tasks directly
 - POP protocol creates EmailReceiveTask when Internet neighbour is visible

Forwarding Algorithms

	{Protocol, Name, Neighbour}					
FOs	x		x		x	algorithm 1
				x	x	algorithm 2
		⊗	x			x = scalar "benefit" of forwarding task

- Forwarding algorithms create Forwarding *Tasks* to send data to suitable next-hops
- Can also create Tasks to perform signalling

Aside on security etc

- Security was “left out” for version 1 in this 4-year EU project, but threats were considered
- Data security can reuse existing solutions of authentication/encryption
 - With proviso that it is not possible to rely on a synchronously available trusted third party
- Some new threats to privacy
 - Neighbourhood visibility means trackability
 - Name graphs could include quite private information
- Incentives to cooperate an issue
 - Why should I spend *any* bandwidth/energy on your stuff?
- Did address later (Social Nets 2009-2011)
 - see safebook.us by Eurecom folks...

Haggle S/W Architecture Details #1

- Motivation for doing mobile networking differently
- Problem definition for mobile networking: “Pocket Switched Networking”
- A set of guiding principles behind our PSN solution, named Haggle
- A 39000 feet overview of Haggle’s initial design

Motivation: Mobile users currently have a very bad experience of networking

- Applications do not work without infrastructure
- Local connectivity is plentiful (WiFi, Bluetooth, etc), but very hard to use transparently
- E.g. messaging/file transfer to others in this room?
- E.g. If I had used a modem to get some cached web content (e.g. news) earlier, and you wanted to access it, how can we share it?

The wireless networking research community has failed to support our end users

Pocket Switched Networking (PSN): Scenario for Mobile (DTN) Humans

- Study/define problem before attempting solution
- Pocket Switched Networking: the scenario that the mobile user of consumer IT apps faces every day
 - Humans carry one or more devices with them, each with wireless networking capabilities and storage
 - These devices experience *neighbourhood* (e.g. Bluetooth) and *infrastructure* (e.g. 802.11 AP) connection opportunities (with differing bandwidths, costs, etc)
 - Human *mobility* generates these opportunities as they move around with their normal mobility patterns.

Pocket Switched Networking: Application traffic

- In PSN, we can identify two classes of application traffic:

known-sender where one node needs to transfer data to a user-level destination (not a network-level address), e.g.:

- Another user (who may own many nodes)
- All users in a certain place/with a certain interest
- Users with a certain role (e.g. "police"), etc.

known-recipient in which a device requires content of some sort, but it is irrelevant where the data comes from, e.g.:

- Publicly distributed content such as "current news webpage"
 - Media files, e.g. "songs by the Scissor Sisters"
 - Locally generated information, e.g. traffic news
- For known-sender, there may be many recipients; for known-recipient, there may be many sources.
 - Yes, it's very Pub/Sub (therefore CCN/NDN/PSIRP/LIPSIN etc)

Why status quo (IP) does badly in PSN

- IP doesn't handle many-recipient well, and does not handle many-source at all
- IP's strict layering means infrastructure lookups to find endpoint addresses before data transmission begins
- TCP/IP's stream abstraction means that apps have to implement app-layer protocols; these rely on access to specific infrastructures
 - E.g. email user wants "message to James" but email client implements "message to IP address of MX DNS record of James's email domain"
- IP cannot handle non-contemporaneous connectivity, e.g. use of human mobility as data transfer opportunity
- Packet switching means that app-layer data is lost to the network; further exchanges of the same data means insertion of the data into the network by an application

Design principles for Hagggle

- A. Forward using application layer information
- B. Asynchronous operation
- C. Empower intermediate nodes
- D. Message switching
- E. All user data kept network-visible
- F. Build request-response into the network
- G. Exploit all data transfer methods
- H. Take advantage of brief connection opportunities
- I. Empowered and informed resource management
- J. Use and integrate with existing application infrastructure where possible

A. Forward using application layer information

- Use names meaningful to apps, e.g. human names, keywords for documents, parameters for content wanted (mime-type, etc)
- Delivery of data is accomplished by using the data itself to choose forwarding path, rather than artificial meaningless-to-the-user addresses such as IP.

B. Asynchronous operation

- Apps can indicate network actions when natural to them; actual transfers can happen asynchronously when suitable connectivity occurs
- Late binding of user-level names to network-level addresses means that up-to-date context information can be used, e.g. dynamic IP address
- Support non-contemporaneous, store-and-forward connections

C. Empower intermediate nodes

- Much content is public/sharable – e.g. webpages
- Thus any intermediate node may also be a valid destination, e.g. it's user might also be interested in the webpage later
- Additionally, the intermediate node can be a source for that data too – e.g. it meets another node who is interested in the webpage

D. Message switching

- Message switching is useful for principles A,B,C
- Application layer forwarding information applies to whole messages
- With variability inherent in non-contemporaneous data paths, packet switching would result in lots of useless half-messages arriving, wasting bandwidth
- Intermediate nodes cannot gain benefits unless entire data units are made available to them
- Aside - looks like DTN *Bundles* but it isn't

E. All user data kept network visible

- Data should not be stored privately by applications, but kept in the Huggle framework where it can be shared with other devices.
- Even your most personal data can be shared – with your other devices. Encryption can be used to prevent unauthorised parties snooping.
 - Painful data synchronisation systems (e.g. phone/PC) can be made obsolete!
- Public data is often popular and duplicated locally, so making this visible allows us to find more sources for it
 - {Looks a bit like all browsers are p2p or caches/proxies also}
- Think original IP arch + plutarch

F. Build request-response into the network

- User-level tasks are sometimes inherently two-way, e.g. a request for content and the response including the content
- Supporting this in the network framework itself rather than in apps allows all nodes to be used for the “turnaround point”
- I.e. all nodes can cache application data (even if they do not run any app that can understand it), and respond to a request for that data

G. Exploit all data transfer methods

- Different transfer methods have different properties
 - Synchronous (Bluetooth), asynchronous (SMS)
 - Zero cost (neighbourhood), cost-per-hour (WiFi hotspot), zero cost till monthly limit hit (SMS)
 - Physical-layer bandwidths, latencies, loss rates, etc
- A given transmission may be sent using multiple diverse paths, e.g. by email and later by Bluetooth directly.
- C.f. Plutarch

H. Take advantage of brief connection opportunities

- Connection opportunities can be fleeting
 - E.g. driving past an AP or another car
- Must optimise transmissions to maximum user benefit
- N.B. current protocols such as Web, SMTP, are really bad at this (work to be presented in WMCSA 2006)

I. Empowered and informed resource management

- Hagggle has the potential to use up all your device's resources and really piss you off
- It also has the potential to do resource management correctly, something today's devices don't do
 - Storage: your disk is not full, why not?
 - Networking: your WiFi interface has one queue, why?
 - Battery: why use static levels for "plenty" and "little"?
- Resource management must be put centre stage in Hagggle

J. Use and integrate with existing application infrastructure where possible

- Incremental deployability if a Haggie node can interact with another node's legacy apps
- Can reuse existing (familiar and complex) apps via pretending to be a legacy protocol rather than having to push out new ones
- Leverage vast infrastructures that are already deployed – e.g. email servers, IM servers, the Google index, etc

Delivery using user-level names

- Names come from:
 - Network, e.g. using neighbour discovery (12:AB:23:98:BE:FF)
 - Applications, e.g. Jon Crowcroft » jon.crowcroft@cl.cam.ac.uk
- Some names are also “addresses” i.e. data can be delivered to that name using one of the protocols available
- Delivery engine needs to:
 - Sense “nearby” addresses (e.g. Bluetooth inquiry gives MAC addresses, Internet connectivity means all email addresses are deliverable)
 - *Known-sender*: Map between ADU’s destination name(s) and addresses of suitable next-hop nodes
 - *Known-recipient*: Determine suitable nearby nodes which may be sources or help find sources for requested data
 - Describe these potential transfers and their benefits to the resource manager

Resource management using tasks


- Resource management uses a list of tasks mainly provided by delivery engine
 - E.g. perform discovery on interface I
 - E.g. send ADU X to neighbour Y on interface Z
- Each task has an associated benefit and cost
 - Benefit is specified by task provider. May be time-dependent (i.e. using a pointer-to-function)
 - To get cost, resource use is estimated, and then the "cost" is a function of the resource use * the resource scarcity
 - Resource manager then schedules execution of tasks in order of highest benefit/cost ratio.

D³N*


***2. Programming
Distributed Computation
in Pocket Switched
Networks (CCN/NDN etc)***

PSN: Dynamic Human Networks



- Topology changes every time unit
- Exhibits characteristics of Social Networks



Time unit = t




Time unit = $t+1$





Time Dependent Networks

- Data paths may not exist at any one point in time but do exist *over time*
- Delay Tolerant Communication



Regularity of Network Activity

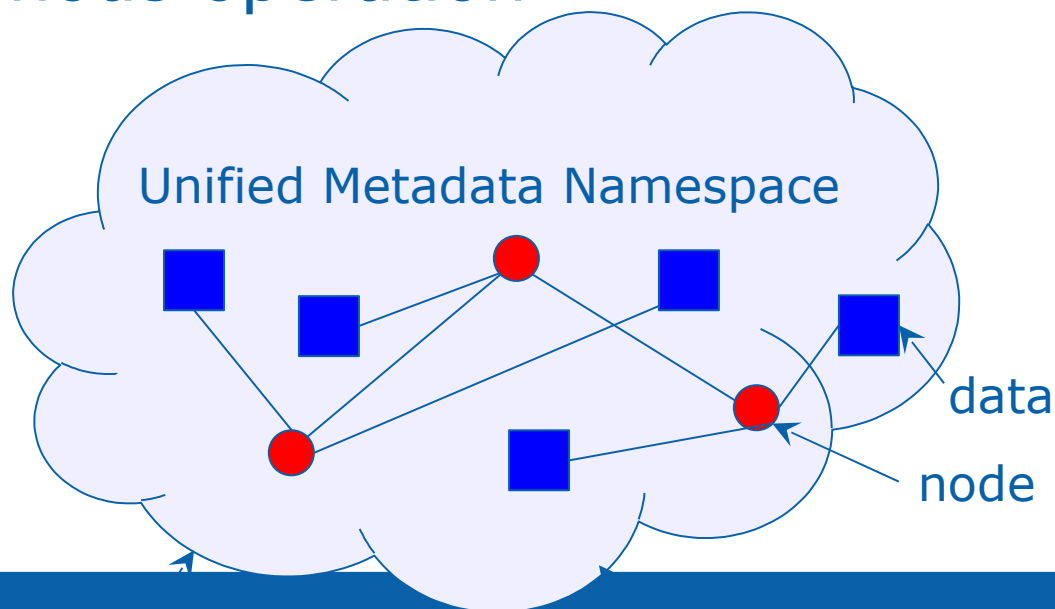
- Size of largest fragment shows network dynamics




```
.register( Event.OnEncounter, fun d:device ->
d.nID = "B" && distance(self,d) < 3 then
ispatch NodeEncountered(d);
```

Haggle Node Architecture

- Each node maintains a data store: its current view of global namespace
 - Persistence of search: delay tolerance and opportunism
- Semantics of publish/subscribe and an event-driven + asynchronous operation
- Multi-platform
(written in C++ and C)
 - Windows mobile
 - Mac OS X, iPhone
 - Linux
 - Android



D³N Data-Driven Declarative Networking

- How to program distributed computation?
 - Use Declarative Networking ?
- The Vodafone Story....
 - Need tested or verified code....so also good...

Declarative Networking

- Declarative is new idea in networking
 - e.g. Search: 'what to look for' rather than 'how to look for'
 - Abstract complexity in networking/data processing
- **P2**: Building overlay using Overlog
 - Network properties specified declaratively
- **LINQ**: extend .NET with language integrated operations for query/store/transform data
- **DryadLINQ**: extends LINQ similar to Google's Map-Reduce
 - Automatic parallelization from sequential declarative code
- **Opis**: Functional-reactive approach in OCaml

D³N Data-Driven Declarative Networking

- How to program distributed computation?
- Use Declarative Networking
 - Use of Functional Programming
 - Simple/clean semantics, expressive, inherent parallelism
 - Queries/Filter etc. can be expressed as higher-order functions that are applied in a distributed setting
- Runtime system provides the necessary native library functions that are specific to each device
 - Prototype: F# + .NET for mobile devices

D³N and Functional Programming I

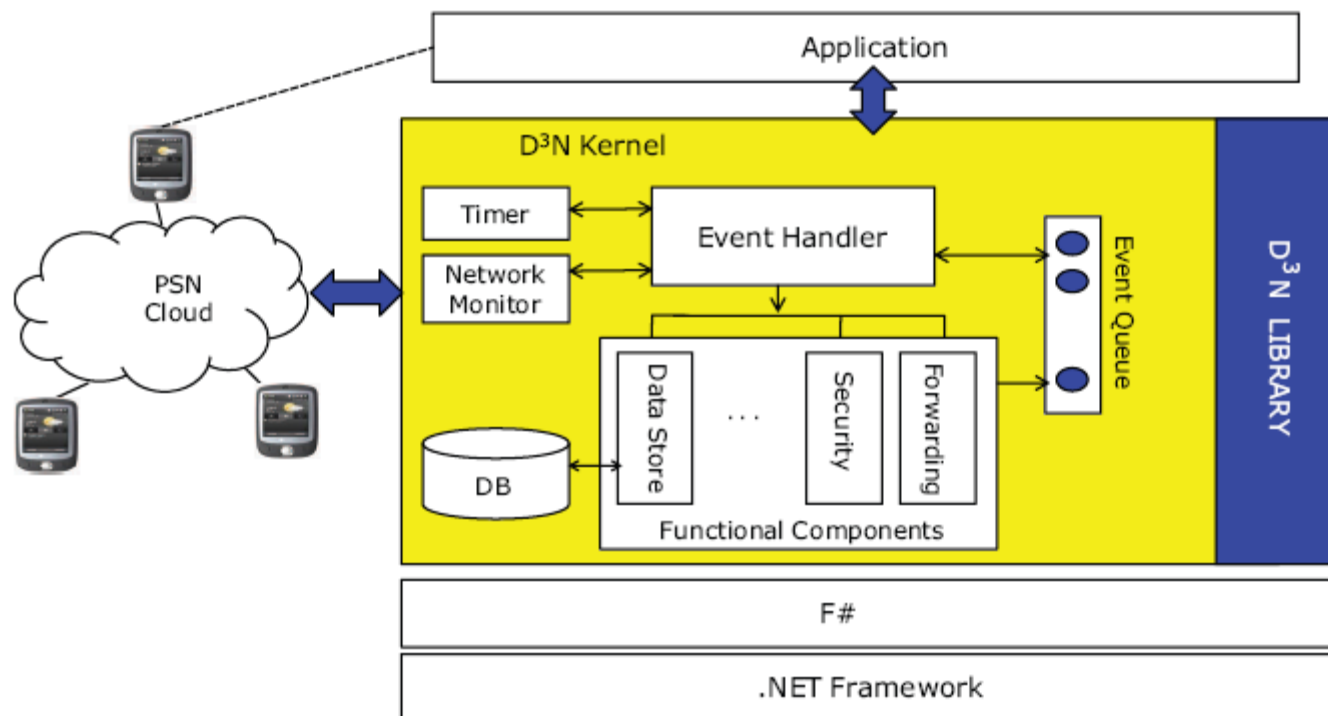
- Functions are first-class values
 - They can be both input and output of other functions
 - They can be shared between different nodes (code mobility)
 - Not only data but also functions flow
- Language syntax does not have state
 - Variables are only ever assigned once; hence reasoning about programs becomes easier
(of course message passing and threads → encode states)
- Strongly typed
 - Static assurance that the program does not 'go wrong' at runtime unlike script languages
- Type inference
 - Types are not declared explicitly, hence programs are less verbose

D³N and Functional Programming II

- Integrated features from query language
 - Assurance as in logical programming
- Appropriate level of abstraction
 - Imperative languages closely specify the implementation details (how); declarative languages abstract too much (what)
 - Imperative – predictable result about performance
 - Declarative language – abstract away

Overview of D³N Architecture

- Each node is responsible for storing, indexing, searching, and delivering data
- Primitive functions associated with core D³N calculus syntax are part of the runtime system
- **Prototype on MS Mobile .NET**



D³N Syntax and Semantics I

- Very few primitives
 - Integer, strings, lists, floating point numbers and other primitives are recovered through constructor application
- Standard FP features
 - Declaring and naming functions through let-bindings
 - Calling primitive and user-defined functions (function application)
 - Pattern matching (similar to switch statement)
 - Standard features as ordinary programming languages (e.g. ML or Haskell)

D³N Syntax and Semantics II

- Advanced features
 - Concurrency (fork)
 - Communication (send/receive primitives)
 - Query expressions (local and distributed select)

```
.register(EventCounter, "node-device")
d.nID = "B" if distance(self, d) < 3 then
  dispatch NodeEncountered(d);
```

BN Language (Core Calculus Syntax)

a, b	::=	<i>Node identifiers</i>
M, N	::=	<i>Value</i>
	x	variable
	$()$	unit
	$\mu f. \lambda \tilde{x}. T$	recursive abstraction
	$c \tilde{M}$	constructor application
S, T	::=	<i>Expression</i>
	M	value
	$M \tilde{N}$	application
	$p \tilde{N}$	primitive application
	let $x = S$ in T	let binding
	match M with	pattern matching
	$c \tilde{x}$ in S else T	
	fork T	fork thread T
	send a M	send value M to a
	receive a	receive a message
	register M N	register event handler
	select M from ...	query
	where N	

Runtime System

- Language relies on a small runtime system
 - Operations implemented in the runtime system written in F#
- Each node is responsible on data:
 - Storing
 - Indexing
 - Searching
 - Delivering
 - Data has Time-To-Live (TTL)
 - Each node propagates data to the other nodes.
 - A search query w/TTL travels within the network until it expires
 - When the node has the matching data, it forwards the data
 - Each node gossips its own metadata when it meets other nodes

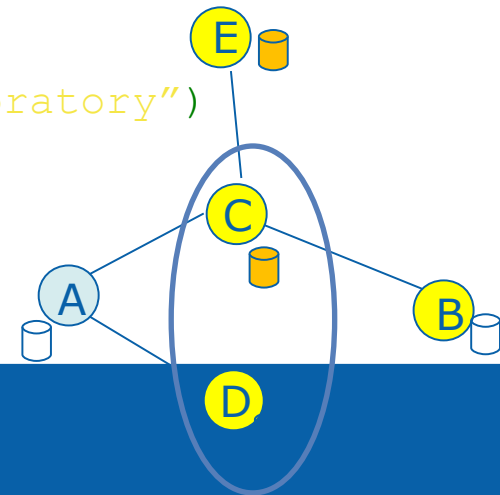
Example: Query to Networks

- Queries are part of source level syntax
 - Distributed execution (single node programmer model)
 - Familiar syntax

D³N: `select name from poll() where institute = "Computer Laboratory"`



F#: `poll()`
`|> filter (fun r -> r.institute = "Computer Laboratory")`
`|> map (fun r -> r.name)`



`(code, nodeid, TTL, data)`

```
.register( Event.OnEncounter, fun d:device ->
d.nID = "B" && distance(self,d) < 3 then
ispatch NodeEncountered(d);
```

Example: Vote among Nodes

- Voting application: implements a distributed voting protocol of choosing location for dinner
- Rules
 - Each node votes once
 - A single node initiates the application
 - Ballots should not be counted twice
 - No infrastructure-based communication is available or it is too expensive
- Top-level expression
 - Node A sends the code to all nodes
 - Nodes map in parallel (pmap) the function `voteOfNode` to their local data, and send back the result to A
 - Node A aggregates (reduce) the results from all nodes and produces a final tally

```
.register( Event.OnEncounter, fun d:device ->
d.nID = "B" && distance(self,d) < 2 then
ispatch NodeEncountered(d):
```

Sequential Map function (smap)


- Inner working
 - It sends the code to execute on the remote node
 - It blocks waiting for a response waiting from the node
 - Continues mapping the function to the rest of the nodes in a sequential fashion
 - An unavailable node blocks the entire computation

```
let rec smap f lst = // Sequential map
match lst with
| [] → []
| n::ns → send f n;receive n :: smap f ns
```

Parallel Map Function (pmap)

- Inner working
 - Similar to the sequential case
 - The send/receive for each node happen in a separate thread
 - An unavailable node does not block the entire computation

```
let rec pmap f lst = // Parallel map
  match lst with
  | [] → []
  | n :: ns →
    fork (fun () →
      send f n; receive n
    ) :: pmap f ns
```



```
.register( Event.OnEncounter, fun d:device ->
d.nID = "B" && distance(self,d) < 3 then
ispatch NodeEncountered(d);
```

Reduce Function

- Inner working

- The reduce function aggregates the results from a map
- The reduce gets executed on the initiator node
- All results must have been received before the reduce can proceed

```
let rec reduce f se lst = // Reduce with starting element
match lst with
| [] → se
| x::xs → f x (reduce f se xs)
```


Voting Application Code

```
type ballot = { locationA : int; locationB : int }  
let emptyBallot = { locationA = 0; locationB = 0 };  
let graph = getSocialGraph();  
let voteForA():ballot = { locationA = 1; locationB = 0 }  
let voteForB():ballot = { locationA = 0; locationB = 1 }
```

```
let rec smap f lst = // Sequential map  
  match lst with  
  | [] → []  
  | n::ns → send f n;receive n :: smap f ns
```

```
let rec pmap f lst = // Parallel map  
  match lst with  
  | [] → []  
  | n :: ns →  
    fork (fun () →  
      send f n;receive n  
    ) :: pmap f ns
```

```
let rec reduce f se lst = // Reduce with starting element  
  match lst with  
  | [] → se  
  | x::xs → f x (reduce f se xs)
```

```
let countVote (b1:ballot) (b2:ballot):ballot =  
  { locationA = b1.locationA + b2.locationA;  
    locationB = b1.locationB + b2.locationB }  
reduce countVote emptyBallot (pmap voteOfNode graph)
```

```
.register( Event.OnEncounter, fun d:device ->  
d.nID = "B" && distance(self,d) < 3 then  
ispatch NodeEncountered(d);
```

Outlook and Future Work

- Current reference implementation:

- F# targeting .NET platform taking advantage of a vast collection of .NET libraries for implementing D³N primitives

- Future work:

- Security issues are currently out of the scope of this paper. Executable code migrating from node to node
- Validate and verify the correctness of the design by implementing a compiler targeting various mobile devices
- Disclose code in public domain

<http://www.cl.cam.ac.uk/~ey204>

Email: eiko.yoneki@cl.cam.ac.uk

3. Connectivity and Routing & How I Got into Social Nets #1

- Motivation and context
- Experiments
- Results
- Analysis of forwarding algorithms
- Consequences on mobile networking

Three independent experiments


- In Cambridge
 - Capture mobile users interaction.
- Traces from Wifi network :
 - Dartmouth and UCSD

User Population	Intel	Cambridge	UCSD	Dartmouth
Device	iMote	iMote	PDA	Laptop/PDA
Network type	Bluetooth	Bluetooth	WiFi	WiFi
Duration (days)	3	5	77	114
Granularity (seconds)	120	120	20	300
Nodes participating	141	238	261	6648
Number of contacts	3,984	8,856	175,105	4,058,284

iMote data sets


- Easy to carry devices
- Scan other devices every 2mns
 - Unsync feature
- log data to flash memory for each contact
 - MAC address, start time, end time
- 2 experiments
 - 20 motes, 3 days, 3,984 contacts, IRC employee
 - 20 motes, 5 days, 8,856 contacts, CAM students

What an iMote looks like



What we measure


- For a given pairs of nodes:
 - contact times and inter-contact times.




What we measure (cont'd)

- Distribution per event.
 \neq seen at a random instant in time.
- Plot log-log distributions.
- We aggregate the data of different pairs.
 (see the following slides).


Example: a typical pair




Examples : Other pairs



Aggregation (1): for one fixed



Aggregation (2) : among iMotes



Summary of observations

- Inter-contact time follows an approximate power-law shape in all experiments.
- $\alpha < 1$ most of the time (very heavily tailed).
- Variation of parameter with the time of day, or among pairs.

Problem

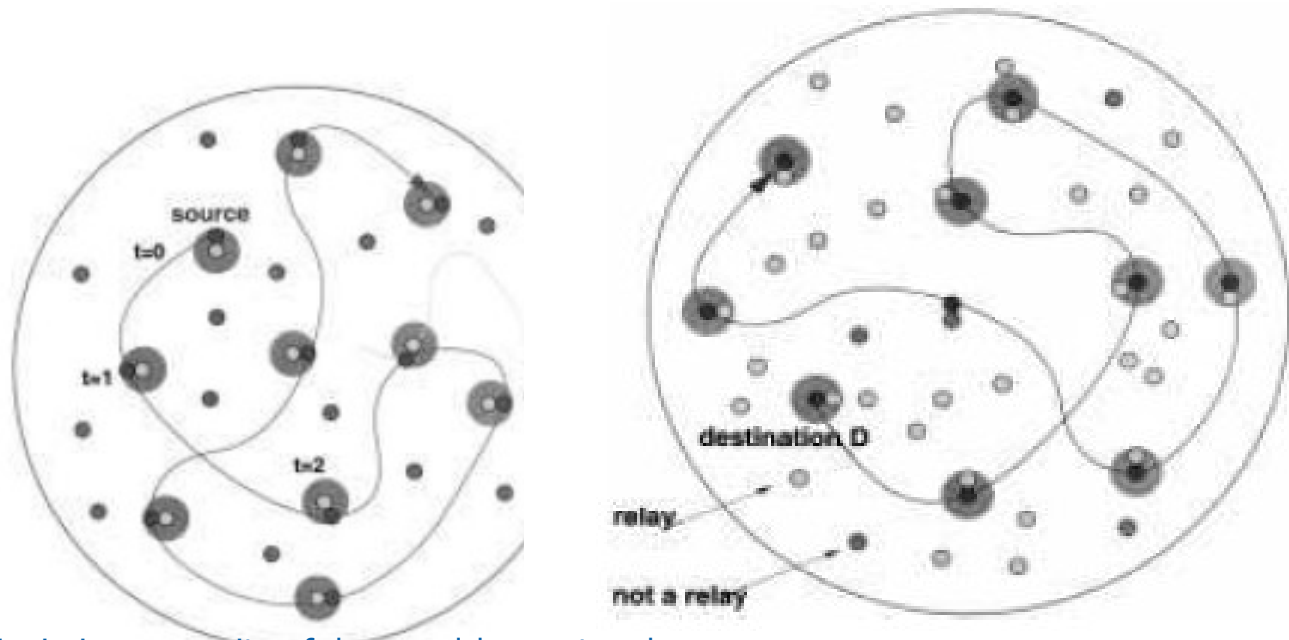
- Given that all data set exhibit approximate power law shape of the inter-contact time distribution:
 - Would a purely opportunistic point-to-point forwarding algorithm converge (i.e. guarantee bounded transmission delays) ?
 - Under what conditions ?

Forwarding algorithms

- Based on opportunities, and “Stateless” :
 - Decision does not depend on the nodes you meet.
- Between two extreme relaying strategies :
 - Wait-and-forward.
 - Flooding.
- Upper and Lower bounds on bandwidth:
 - Short contact time.
 - Full contact time (best case, treated here).

Two-hop relaying strategy

- Grossglauer & Tse (2001) :



- Maximizes capacity of dense ad-hoc networks.
- Authors assume nodes location i.i.d. uniform.

Our assumptions on Mobility

- Homogeneity

- Inter-contact for every pairs follows power law.


- No cut-off bound. $P[X \geq t] = t^{-\alpha}$

- Independence

- In “time”: contacts are renewal instants.
- In “space”: pairs are independent.

Two-hop: stability/instability

- $\alpha > 2$
The two hop relaying algorithm converges, and it achieves a finite expected delay.
- $\alpha < 2$
The expected delay grow to infinity with time.



Two-hop: extensions

- Power laws with cut-off:
 - Large expected delay.
- Short contact case:
 - By comparison, all the negative results hold.
 - Convergence for $\alpha > 3$ by Kingman's bound.
 - We believe the same result holds for $\alpha > 2$.

The Impact of redundancy

- The Two-hop strategy is very conservative.
 - What about duplicate packet ? Or epidemics forwarding ?
- This comes to the question:

D_1, D_2 independent, with same law ,
if $\mathbb{E}[D_1] = \mathbb{E}[D_2] = \infty$ what is $\mathbb{E}[\min(D_1, D_2)]$?

Forwarding with redundancy:

- For $\alpha > 2$
Any stateless algorithm achieves a finite expected delay.
- For $\alpha \geq \frac{m+1}{m}$ and $n \geq 2m$:
There exists a stateless algorithm with n nodes and m copies and a finite expected delay.
- For $\alpha < 1$
No stateless algorithm (even flooding) achieve a bounded delay (Orey's theorem).

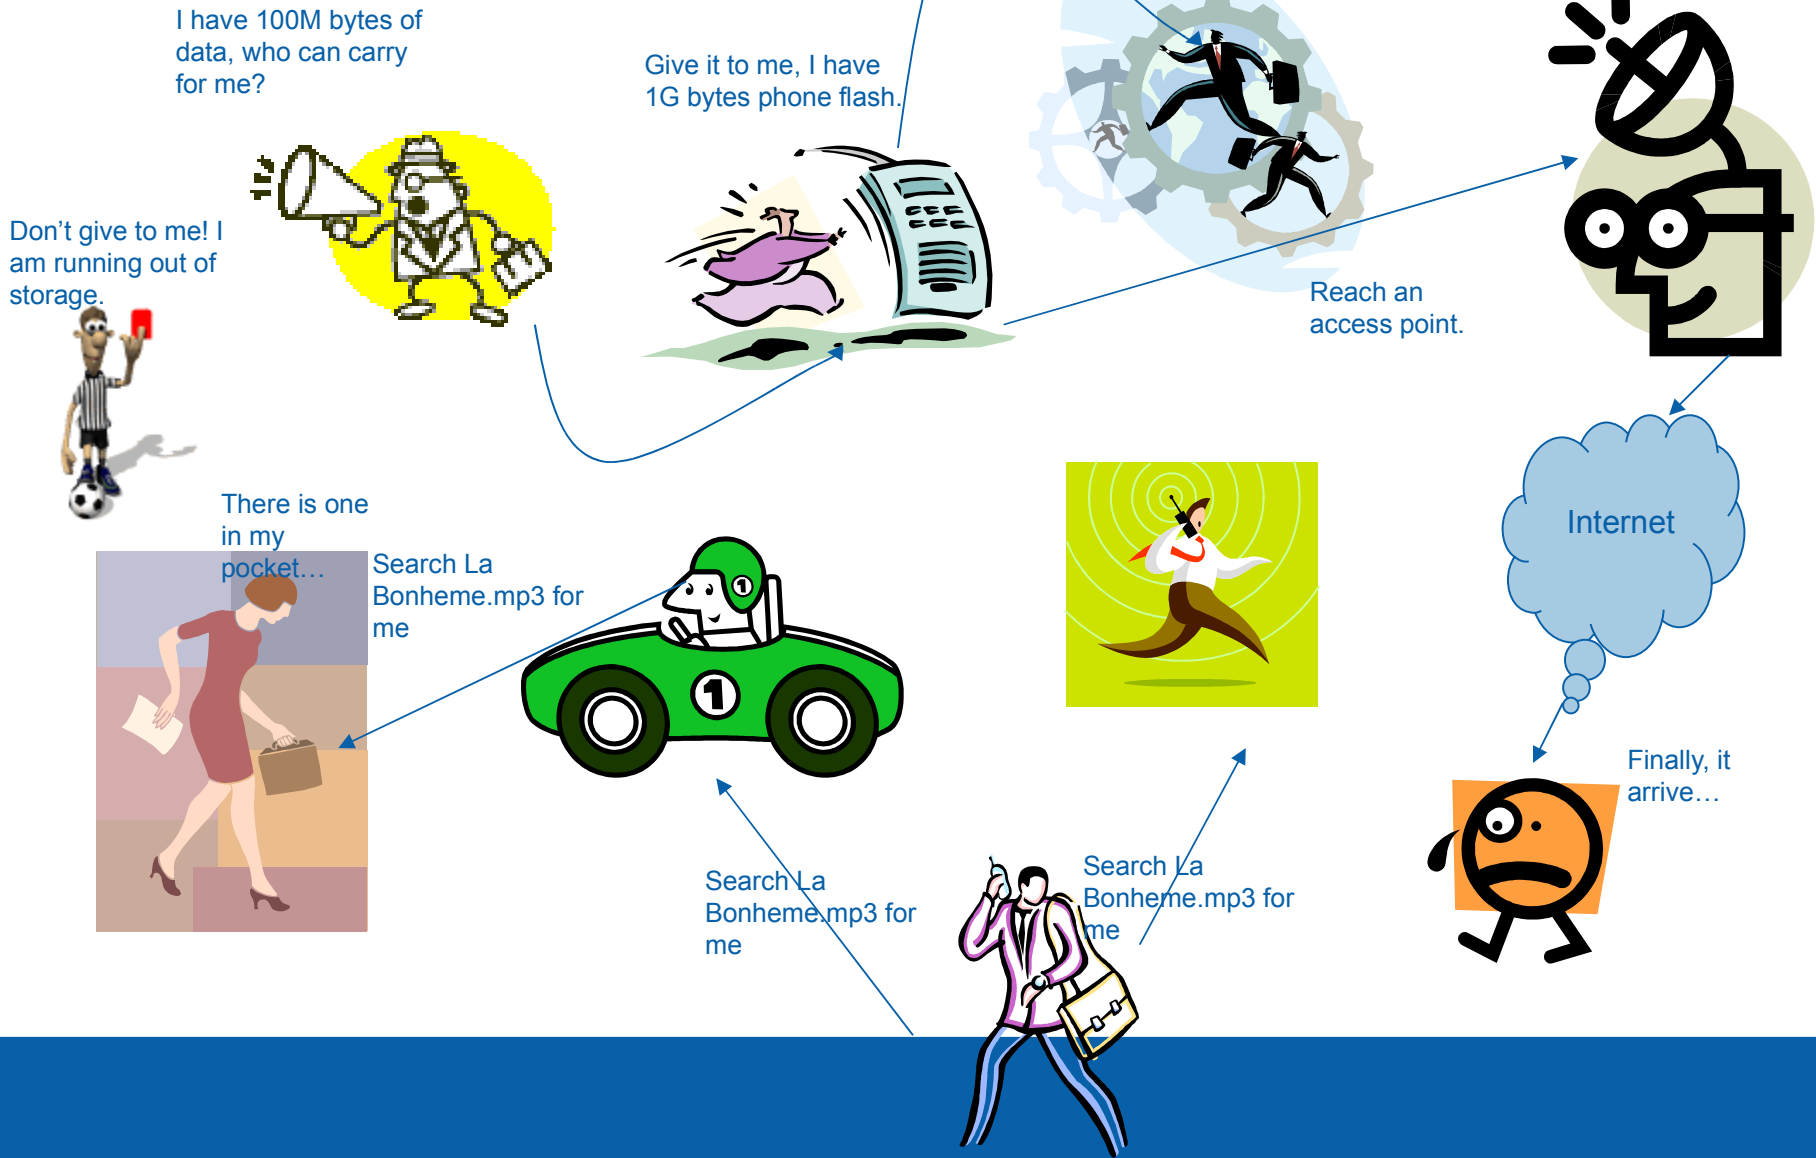
Forwarding w. redundancy (cont'd)

- Further extensions:
 - The short contact case is open for $1 < \alpha < 2$.
 - Can we weaken the assumption of independence between pairs ?


Consequences on mobile networking

- Mobility models needs to be redesigned
 - *Exponential decay of inter contact is wrong.*
 - *Mechanisms tested with that model need to be analyzed with new mobility assumptions.*
- Stateless forwarding does not work
 - *Can we benefit from heterogeneity to forward by communities ?*
 - *Scheme for peer-to-peer information sharing.*


3b Connectivity & Routing Ever More Social




K-clique Communities in Cambridge Dataset





k=3




k=4



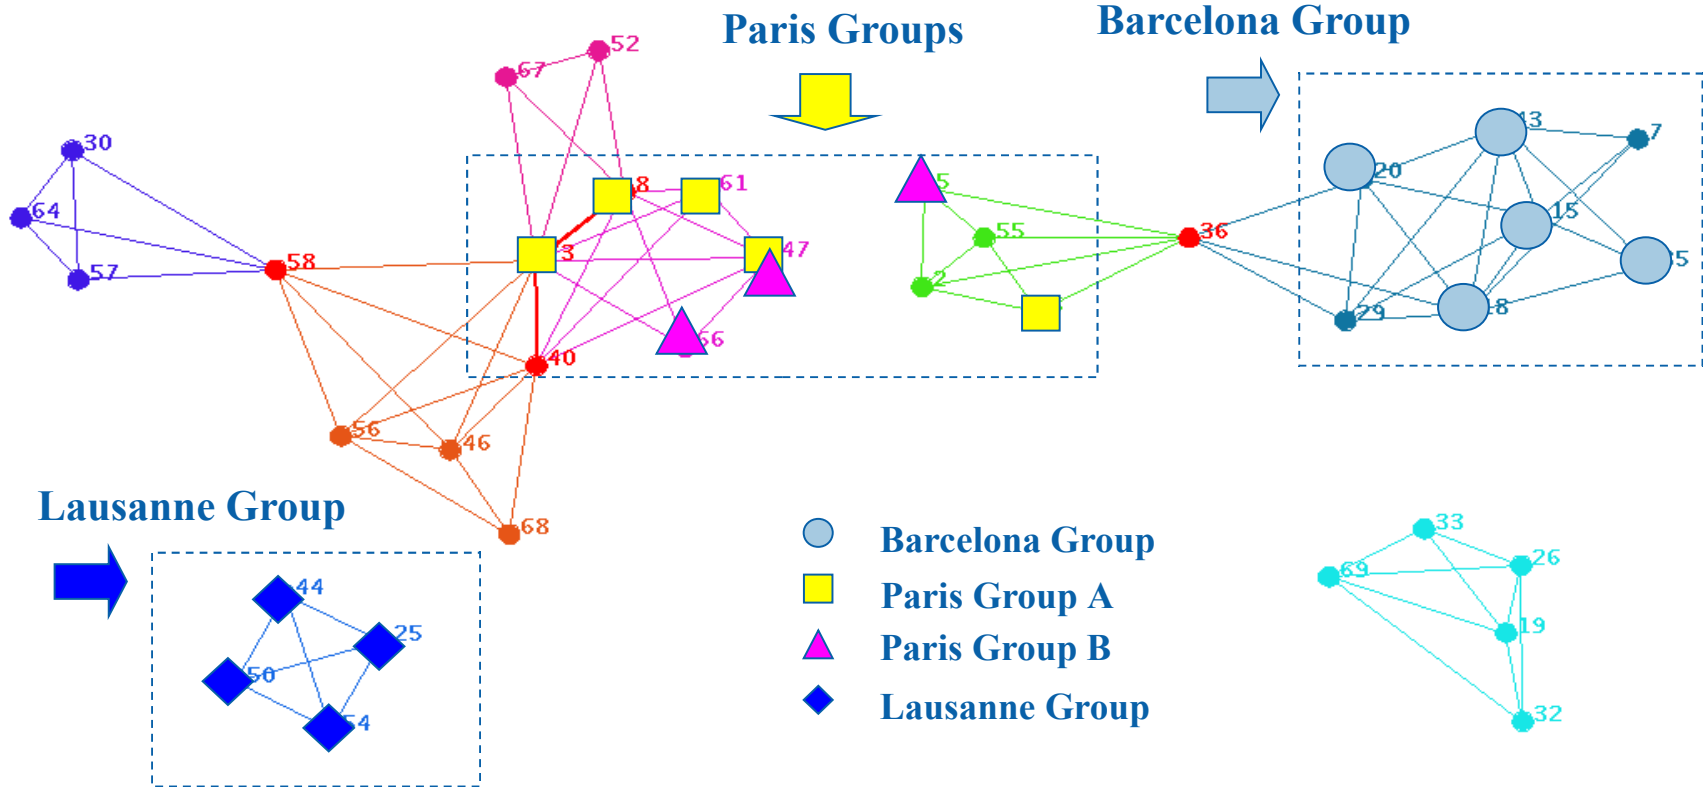
k=5



k=10



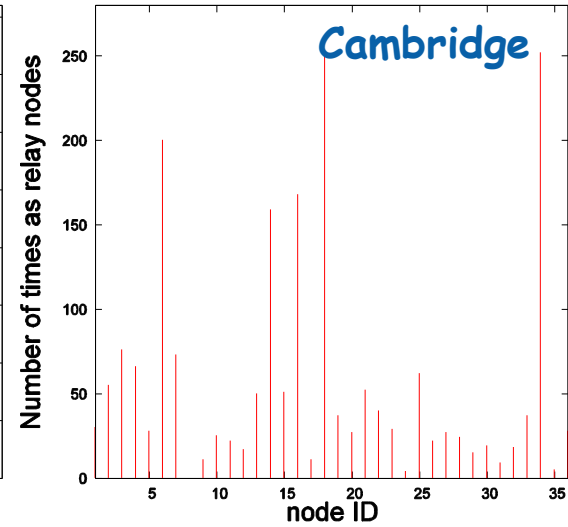



K-clique Communities in Infocom06 Dataset

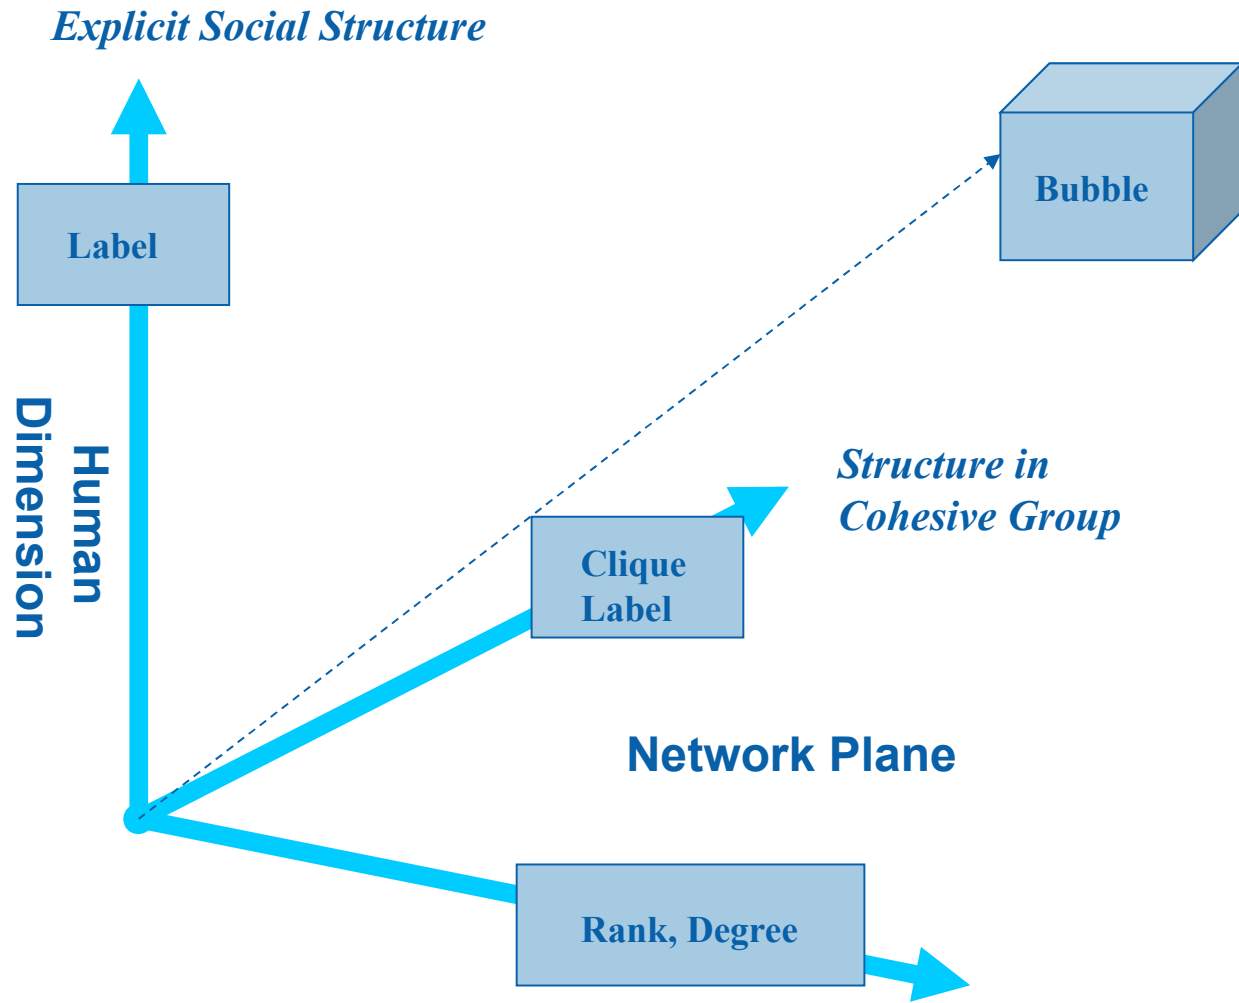


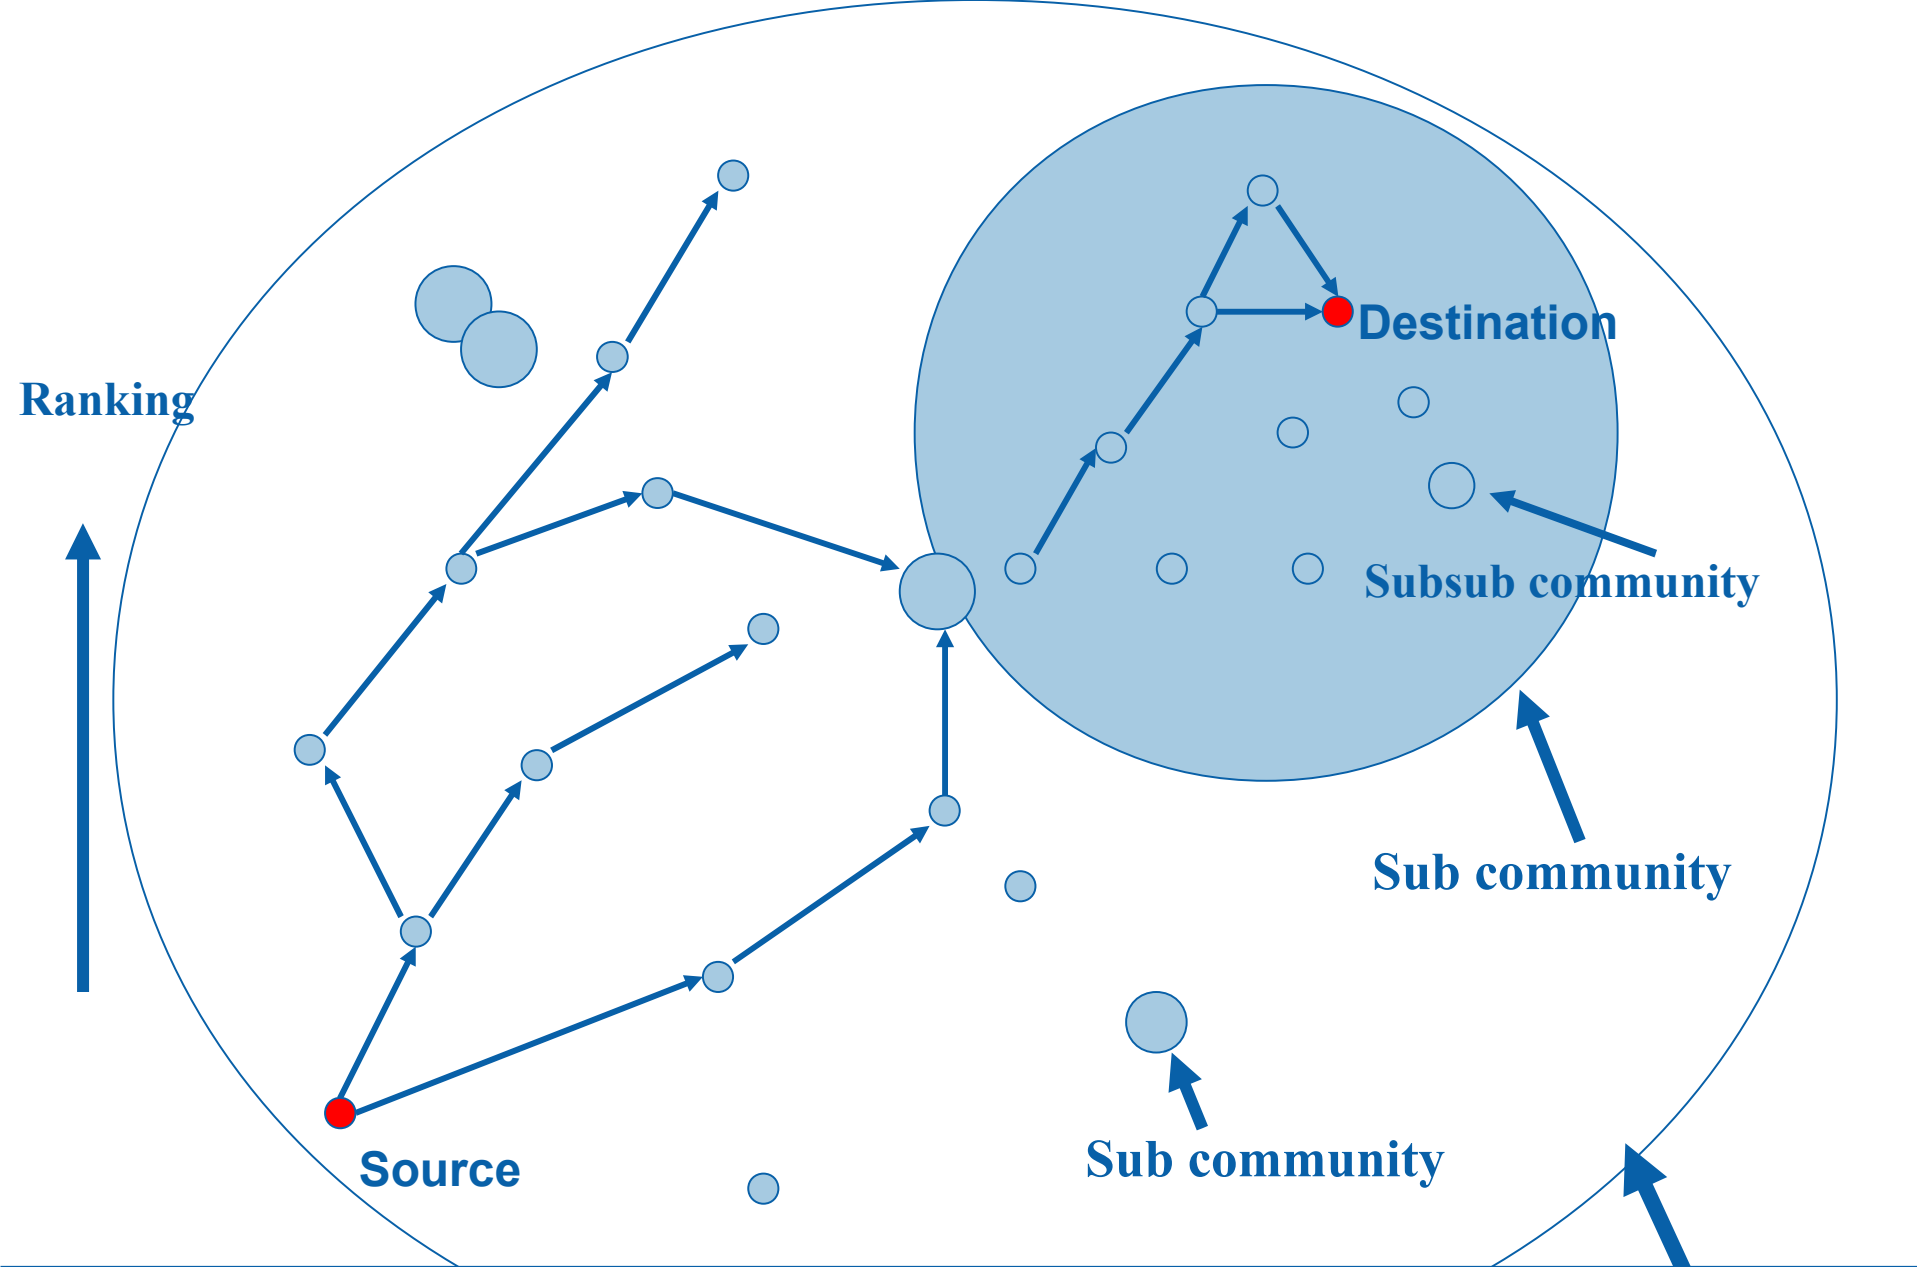
K=4

Human Hubs: Popularity




Forwarding Scheme Design Space





Use affiliation+hubs to fwd inter+intra cliques



3c Connectivity & Routing 3 - Community Detection

I have 100M bytes of data, who can carry for me?



Don't give to me! I am running out of storage.



Give it to me, I have 1G bytes phone flash.




Thank you, but you are in a bad position.

I can also carry for you!




Reach an access point.




There is one in my pocket...



Search La Bonheme.mp3 for me




Search La Bonheme.mp3 for me



Search La Bonheme.mp3 for me



Internet




Finally, it arrive...




Community improves forwarding

- Identifying communities (e.g. affiliations) improves forwarding efficiency. [label]
- Evaluate on Infocom06 data.



(a) Delivery



(b) Cost

Centralized Community Detection

K-clique Detection[Palla04]

Weighted Network Analysis[Newman05]

Betweenness [Newman04]

Modularity [Newman06]

Information theory[Rosvall06]

Statistical mechanics[Reichardt]

Survey Papers[Danon05][Newman04]

K-clique Detection

Union of k-cliques reachable through a series of adjacent k-cliques

Adjacent k-cliques share k-1 nodes

Members in a community reachable through well-connected well subsets

Examples

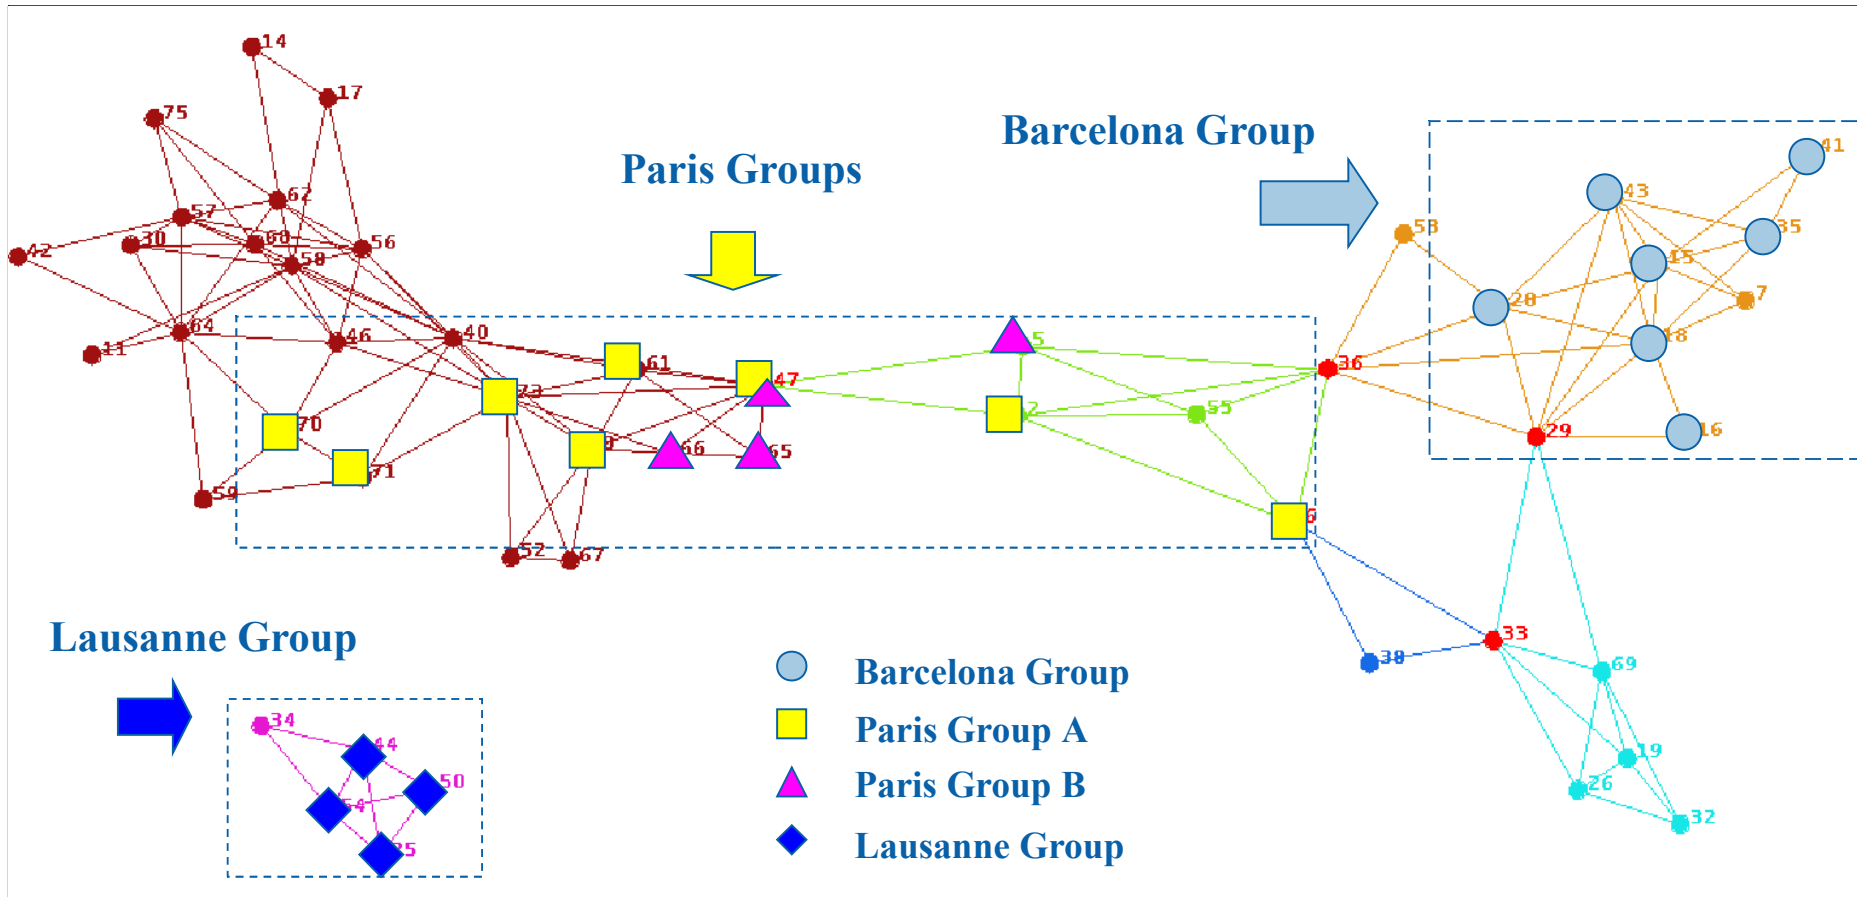
- 2-clique (connected components)
- 3-clique (overlapping triangles)

Overlapping feature

Percolation threshold

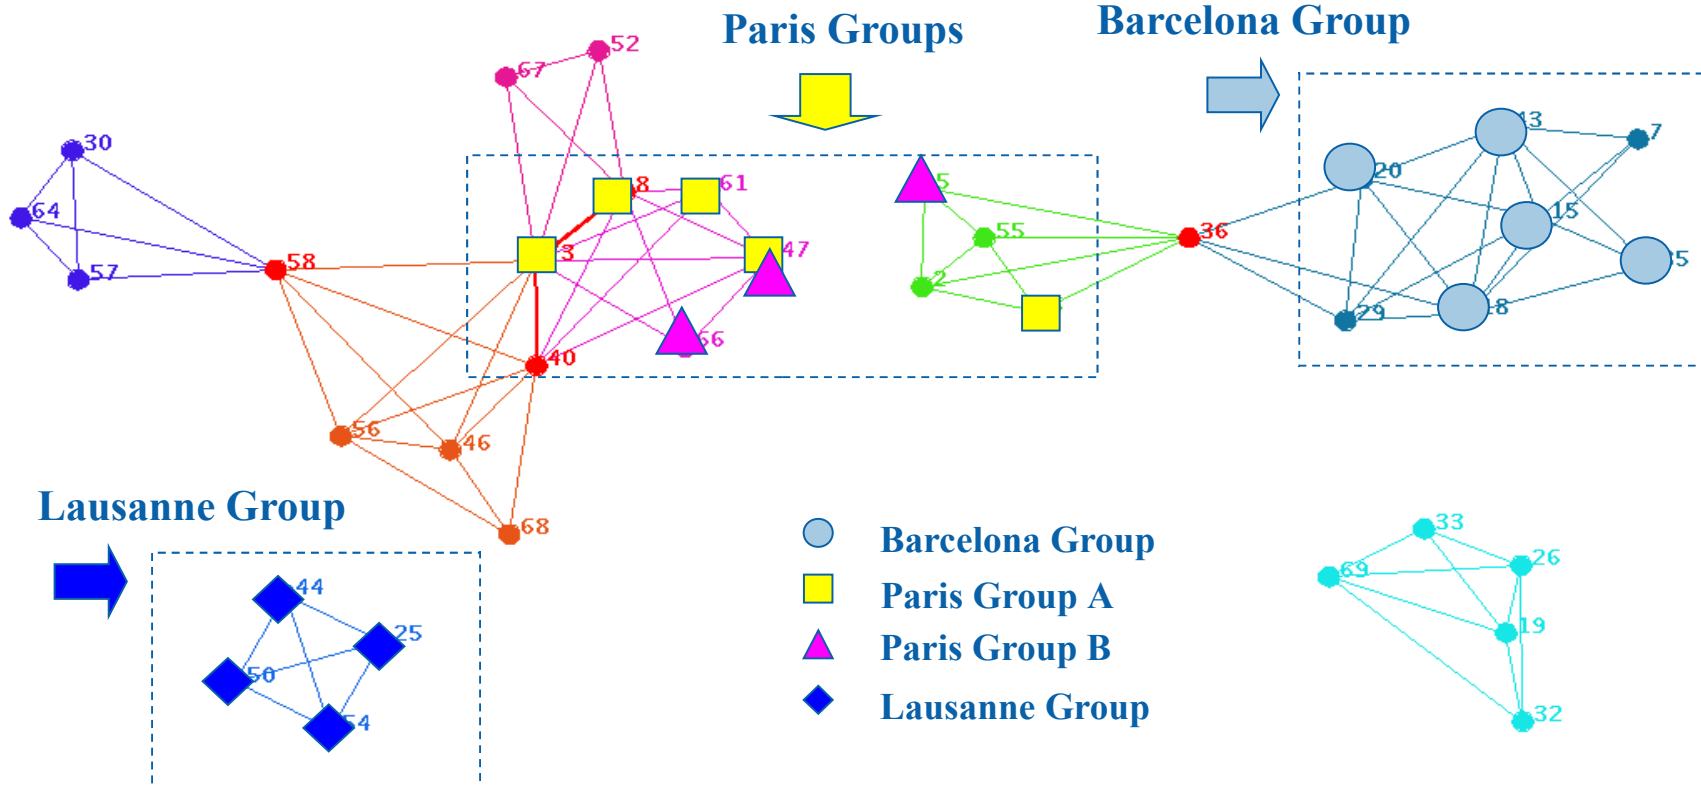
$$p_c(k) = 1/[(k-1)N]^{1/(k-1)}$$

K-clique Communities in Infocom06 Dataset



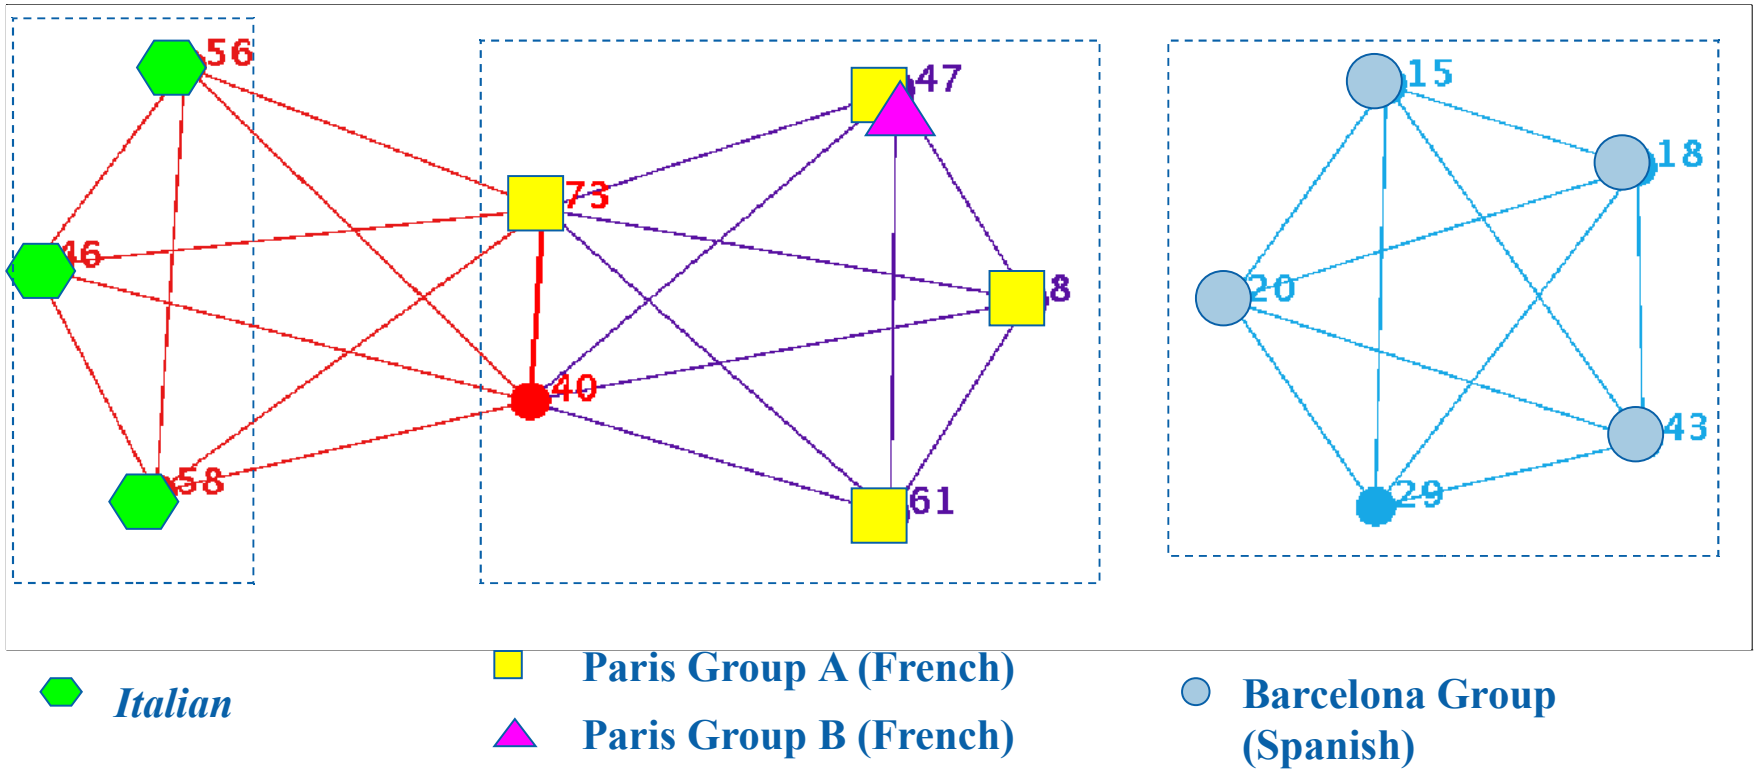
K=3

K-clique Communities in Infocom06 Dataset



K=4

K-clique Communities in Infocom06 Dataset



K=5

Weighted network analysis (WNA)

Calculate the unweighted edge betweenness.

Divide each calculated betweenness value by its weight.

Remove the edge with the highest edge betweenness. and repeat from 1 until there are no more edges in the network.

Recalculate the modularity value of the network with the current community partitioning. Select those splitting with local maxima of modularity.

Community Detection using WNA

Experimental dataset	Infocom06	Cambridge	Reality	Infocom05
Q_{max}	0.2280	0.4227	0.5682	0.3039
Max. Community Size	13	18	23	13
No. Communities	4	2	8	4
Avg. Community Size	8.000	16.500	9.875	6.5
No. Community Nodes	32	33	73	26
Total No. of Nodes	78	36	97	37

Distributed Community Detection

SIMPLE, K-CLIQUE, MODULARITY


Terminology : Familiar Set (F), Local Community (C)

Update and exchange local information during encounter


Build up Familiar Set and Local Community

- *CommunityAccept(), MergeCommunities()*


SIMPLE




We want to see whether v_i should be added to C_0



So we first count the number of vertices in v_i 's familiar set = $(\text{C} + 0)$



Then we count the number of vertices in both v_i 's familiar set and also in the local community of $v_0 = \text{O}$



We only consider merging the two communities C_0 & C_i if the fraction of them in common $\text{shaded} > \gamma$


And we admit v_0 to C_0 iff $\text{O} > (\text{C} + 0) \times \lambda$

MergeCommunities (C_0 , C_i)

K-CLIQUE

CommunityAccept (v_i) : $|F_i \cap C_0| \geq k - 1$


MergeCommunities(C_0, C_j): $|\tilde{F}_j \cap C_0| \geq k - 1$



MODULARITY

Boundary Set

$$B_0 = \{v_i \mid (v_i \in C_0) \text{ and } ((F_i \setminus C_0) \neq \emptyset)\}$$



Local Modularity


$$R_0 = \frac{I}{|T|}$$

Measure of the sharpness of local community

MODULARITY

CommunityAccept (v_i) :

$(F_i \subseteq C_0 \text{ and } B_0 \neq \emptyset)$ or $\Delta R_0^i > 0$
 MergeCommunities(C_0, C_i): for each $v_k \in C_0 \cap C_i$,
 $\tilde{F}_k \subseteq C_0$ or $\Delta R_0^k > 0$



Results and Evaluations



Data Set	SIMPLE	K-CLIQUE	MODULARITY
Reality	0.79/0.81	0.87	0.89
UCSD	0.47/0.56	0.55	0.65
Cambridge	0.85/0.85	0.85	0.87
<i>Complexity</i>	$O(n)$	$O(n^2)$	$O(n^4)/O(n^2k^2)$

Newman weighted analysis


Palla et al. k-Clique

$$\sigma_{Jaccard} = \frac{|\Gamma_i \cap \Gamma_j|}{|\Gamma_i \cup \Gamma_j|}$$

Results and Evaluations



Distributions
of
Local Community Views



Evolution of communities

More general Familiar Set threshold (e.g. hours per day)

Detection of different categories of relationship by specifying contact duration and number of contacts

Dynamic selection of Familiar Set threshold (e.g. fuzzy logic)

Aging effect

Temporal communities

Evaluation on more data sets (e.g. Dartmouth WiFi, iMote experiments)

The End

- With much thanks&acknowledgements to
- James Scott, Ebon Upton, Menghow Lim, Pan Hui
- Eiko Yoneki, Ioannis Baltopoulos, Shu-yan Chan
- Jing Su, Ashvin Goyal, Eyal de Lara
- Christophe Diot, Augustin Chaintreau, Richard Gass



UNIVERSITY OF
CAMBRIDGE

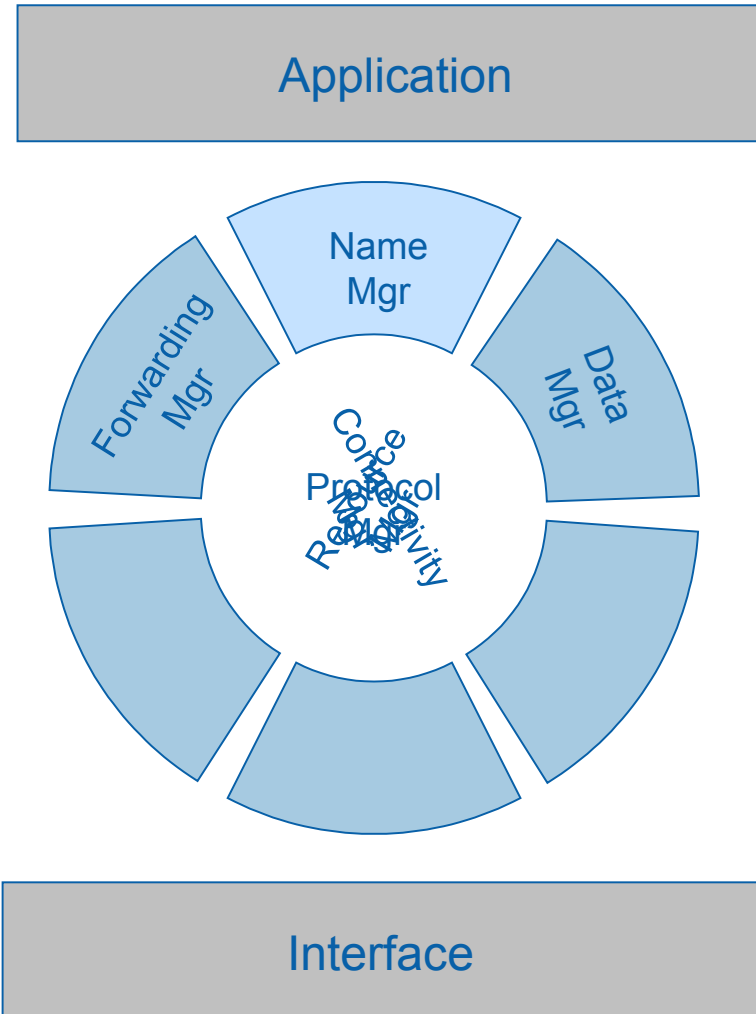
Backup Slides

S/W Arch + Contact Graphs

Haggle S/W Architecture Details #2

- Proposed Haggle Framework (changed)
- Data Object (DO) app interface
- Names and their use in delivery
- Neighbour discovery
- Send DO / DO reception
- Solicit DO

Proposed Hagggle Framework

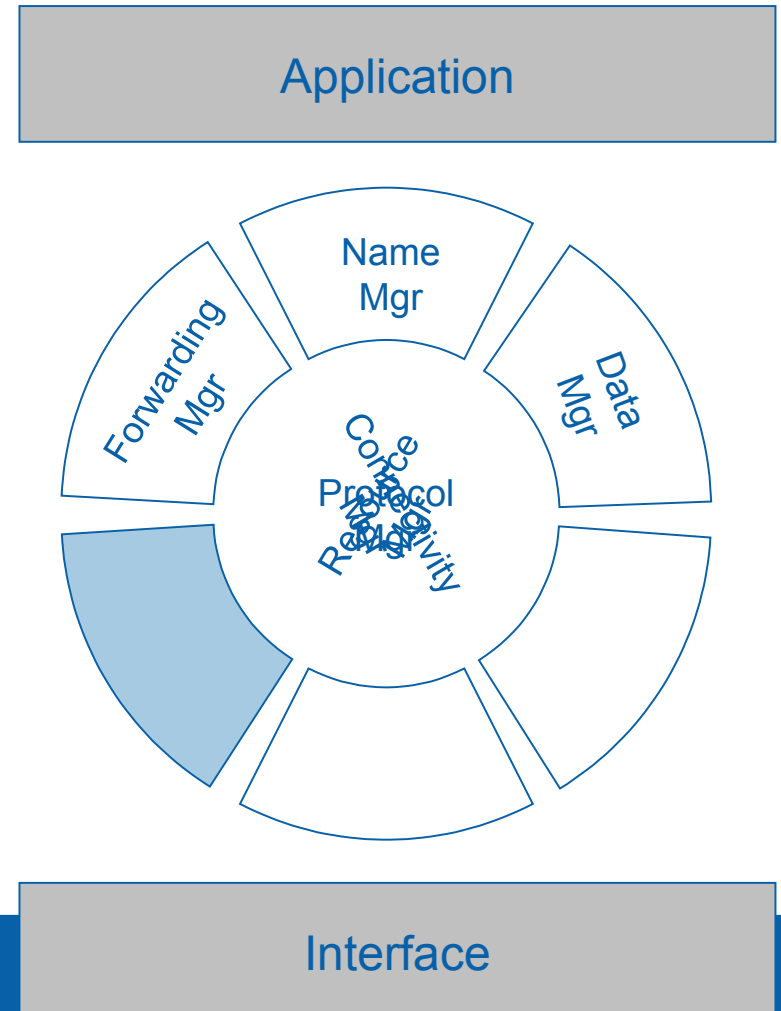


Connectivity Manager

Discover neighbours and instruct interfaces to provide needed connectivity

Interfaces with:

- Name Mgr to insert Names for discovered neighbours
- Protocol Mgr to indicate that neighbours are nearby (thus turning some names into deliverable addresses)
- Resource Mgr to add tasks asking for neighbour discovery to occur

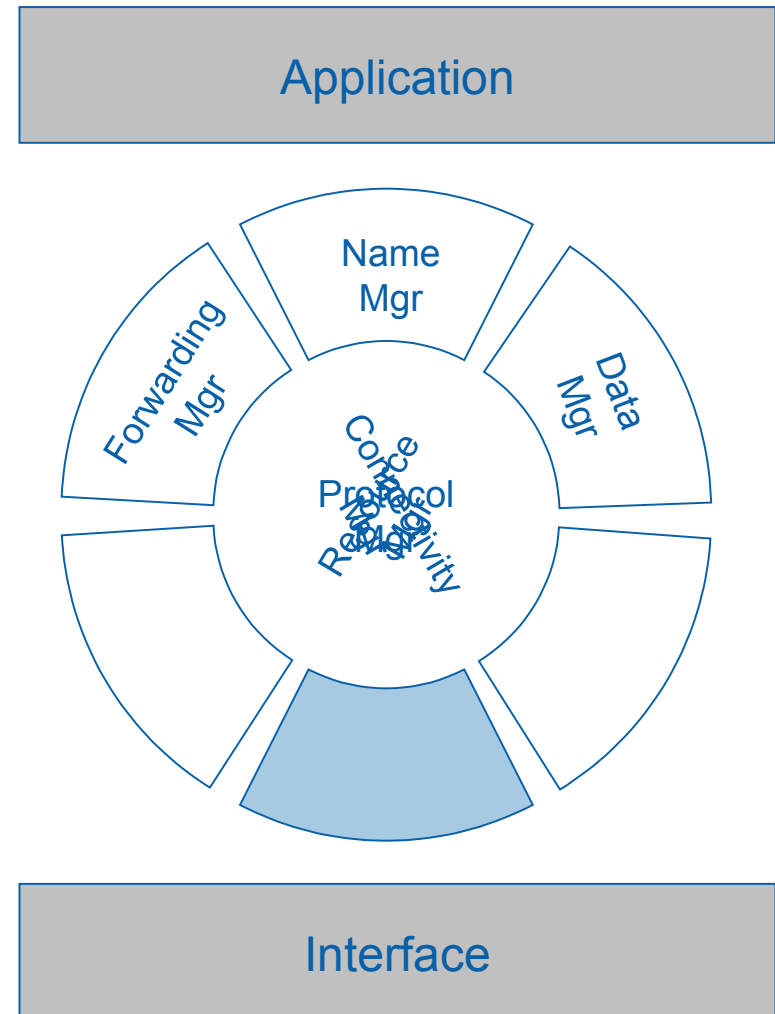


Protocol Manager

Encapsulates all protocols available for forwarding data to next hops.

Interface with:

- Conn Mgr to detect neighbours and determine which names can be delivered to
- Forward Mgr to inform it about deliverable addresses and accept requests for forwarding data to next hops
- Name Mgr if it discovers any new names

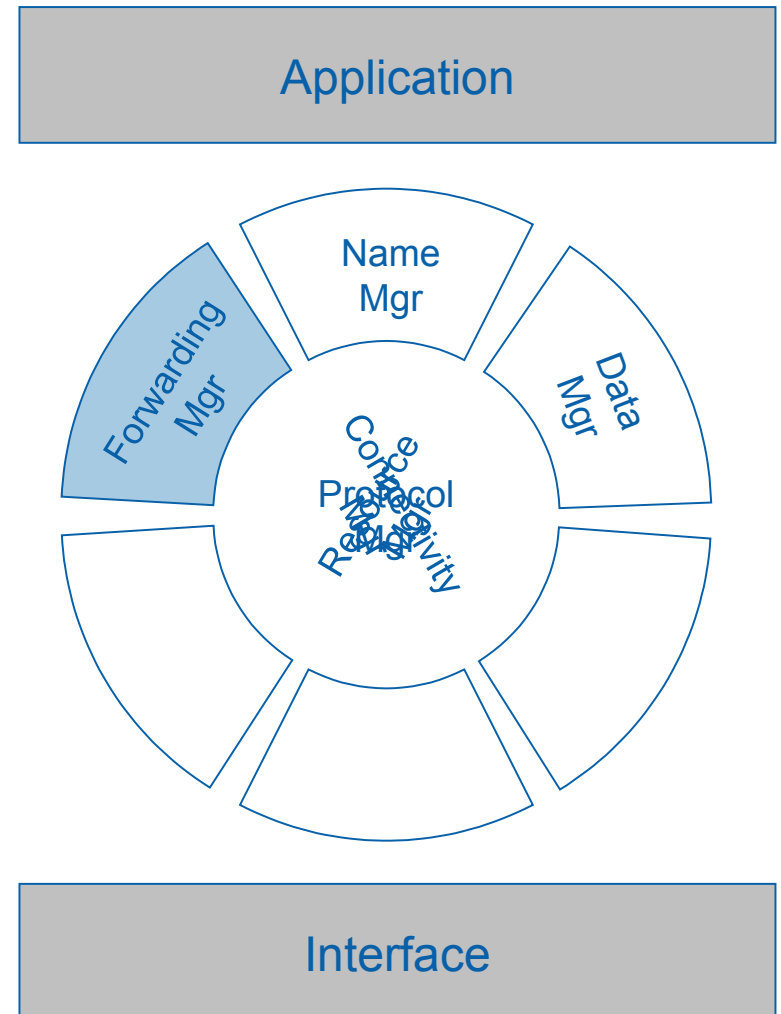


Forwarding Manager

The Forwarding Mgr is responsible for deciding suitable next hops for data in transit, and interfacing with the app to allow it to ask for data to be sent or solicited

Interfaces with:

- Name mgr to determine destinations for data and possible paths by which it can get there
- Protocol mgr to find out which next-hops are nearby and thereby what forwarding can be done, as well as to send out data to a next-hop
- Resource Mgr to propose forwarding tasks and get approval
- Data mgr to store persistent forwarding state

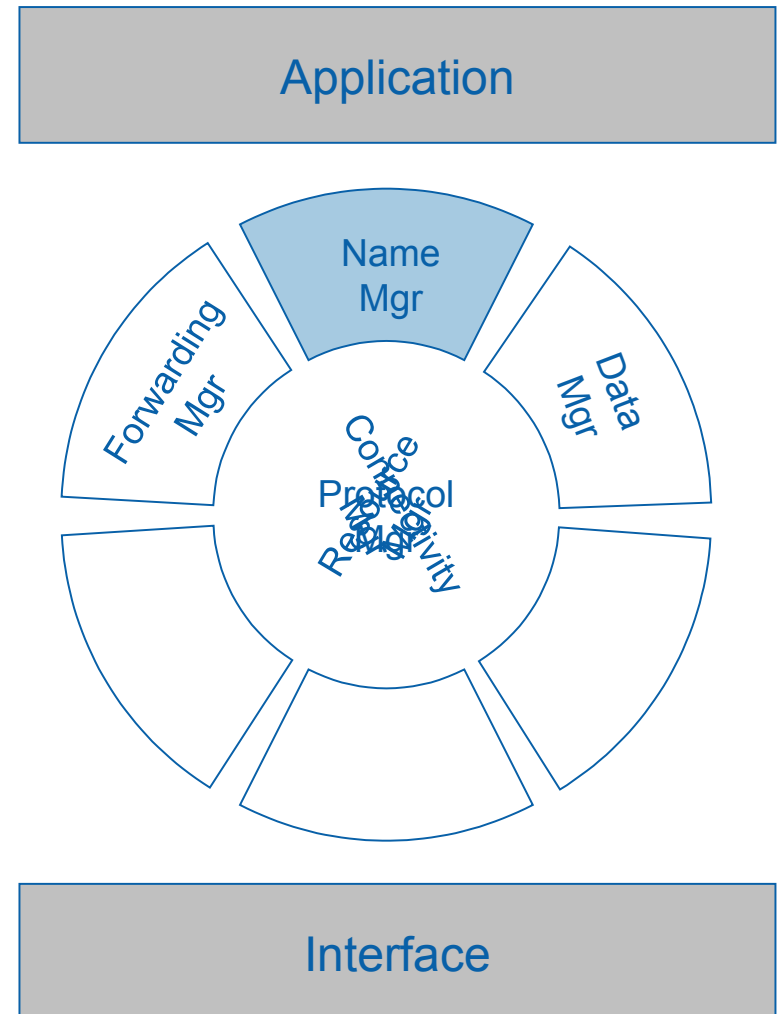


Name Manager

Allows Names to be created, linked to each other, and queried

Interfaces with:

- Application or any other Mgr which needs to use Name information
- Data Mgr to store Names persistently

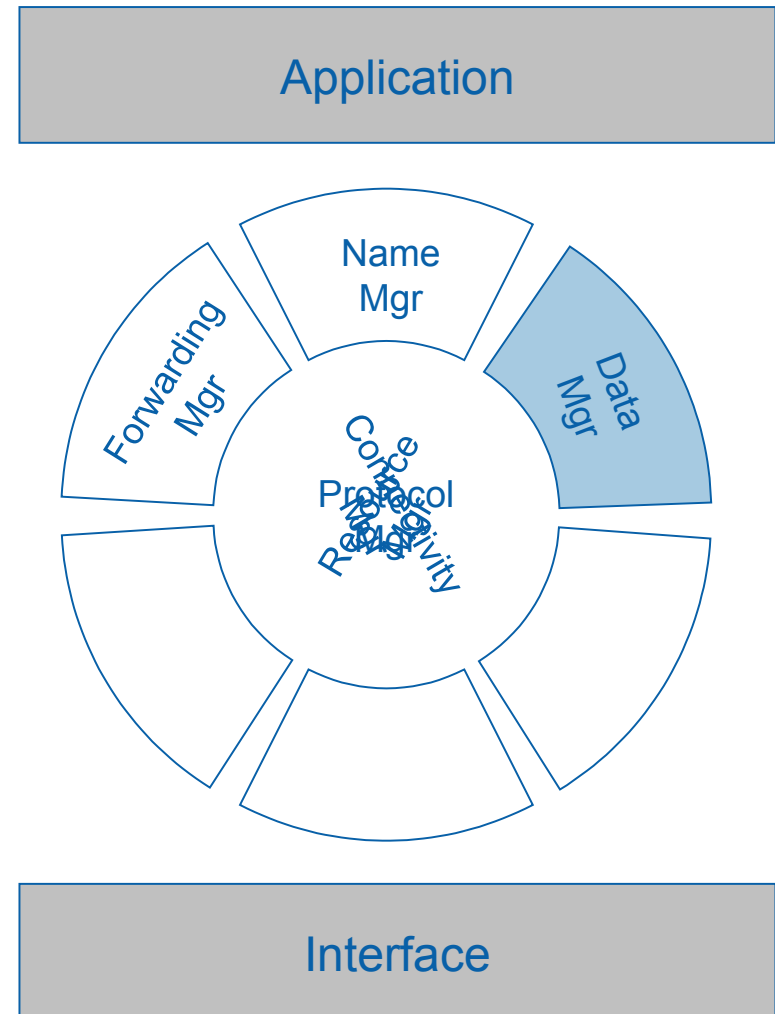


Data Manager

Data Mgr provides persistent storage of Data Objects (DOs), and allows the database to be queried. Also provides a callback interface so that arriving DOs matching a DOFilter can be proactively found.

Interfaces with:

- Application or any other manager needing to use persistent storage or search/register interest in DOs.

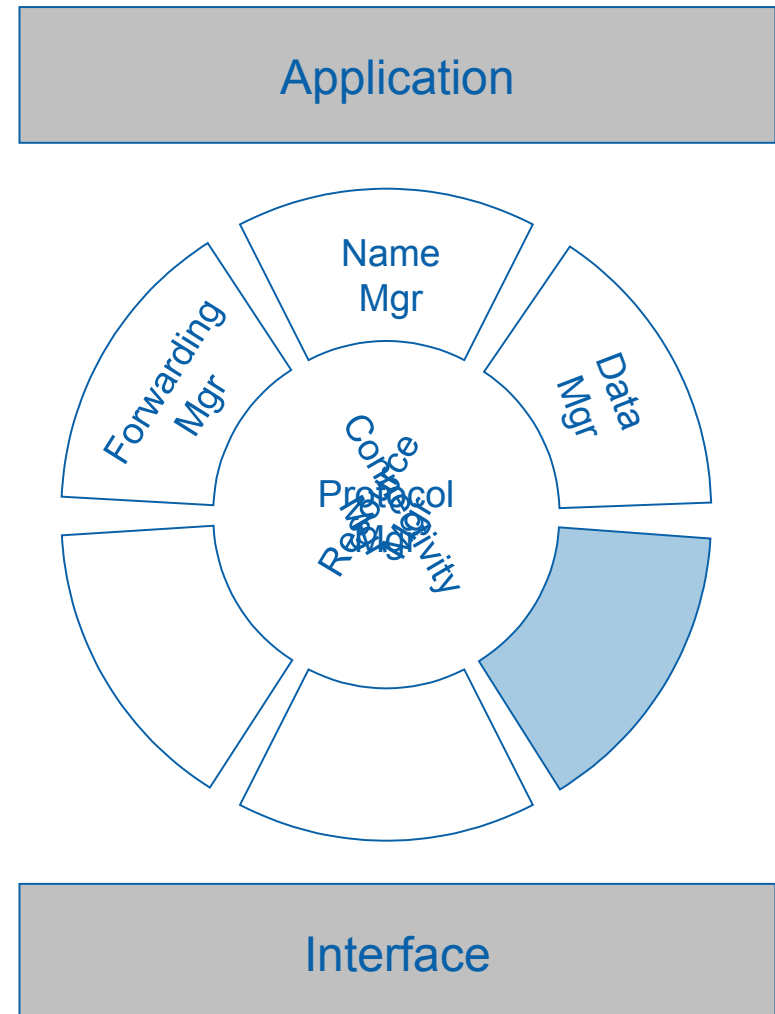


Resource Manager

All network usage is approved by the resource manager. Provides a “task” abstraction and compares costs of tasks with benefit in order to decide next task.

Interfaces with:

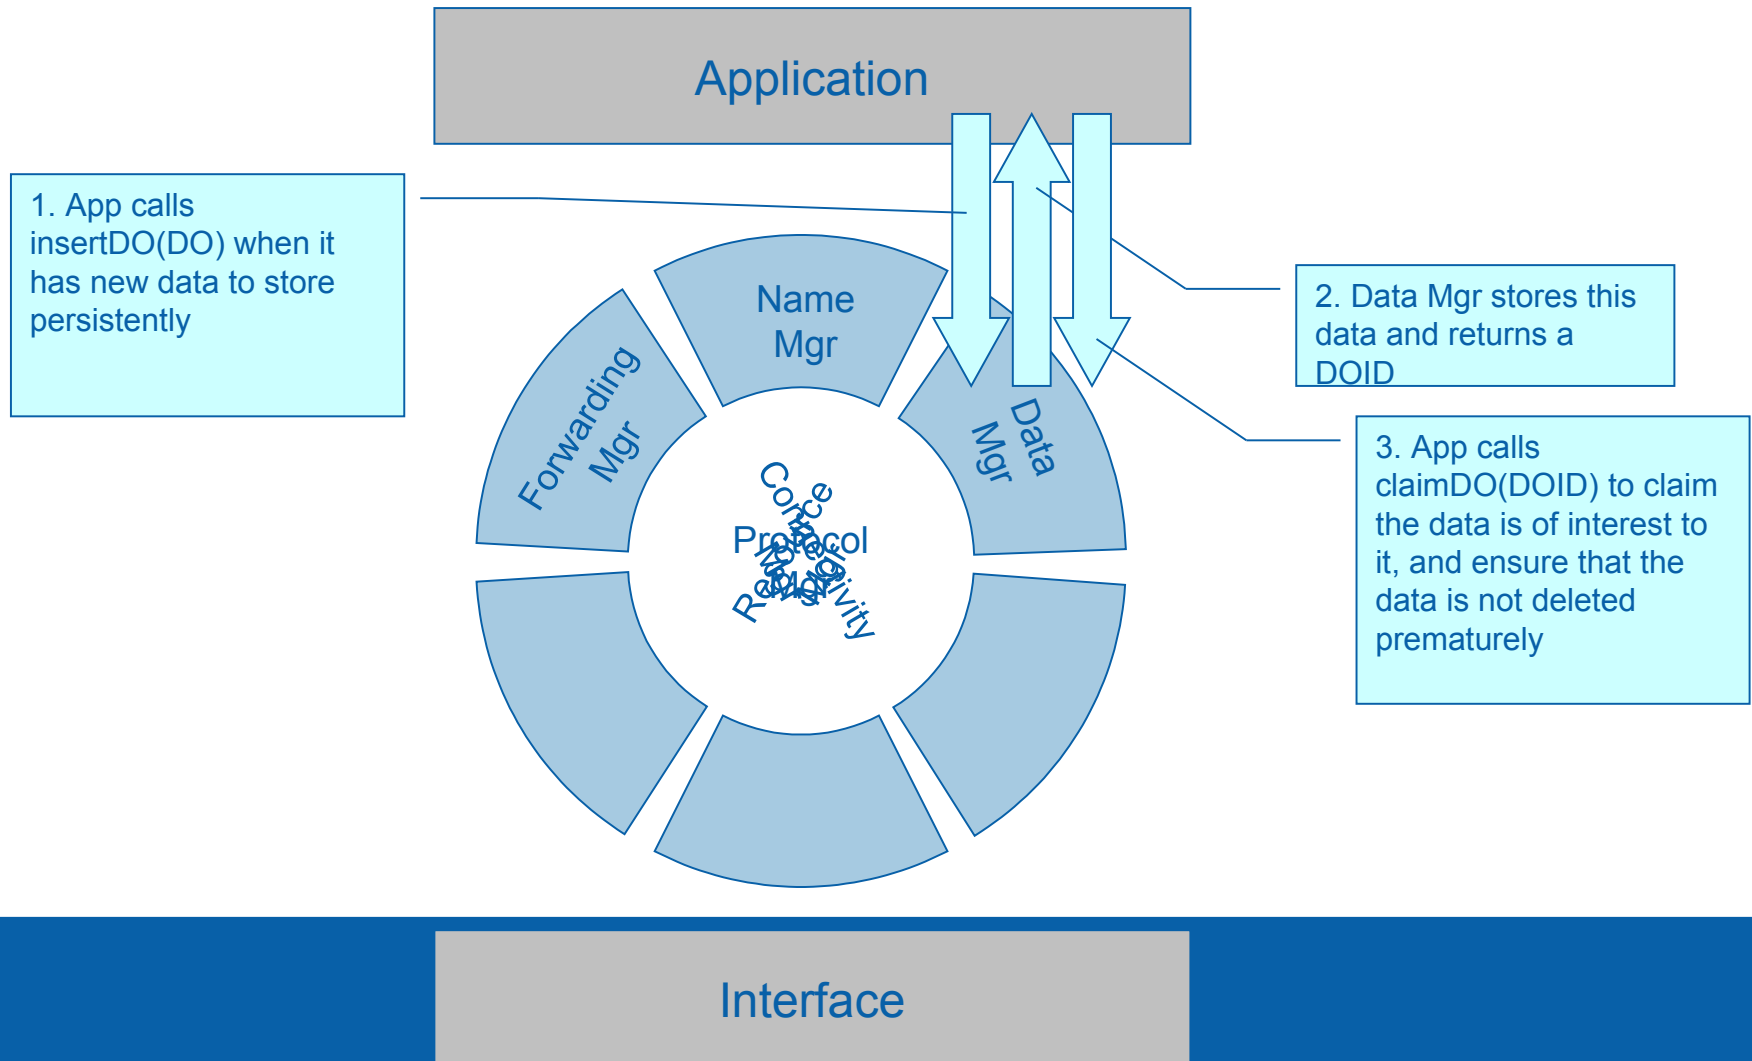
- Connectivity Mgr (for neighbour discovery tasks)
- Forwarding Mgr (for forwarding tasks)



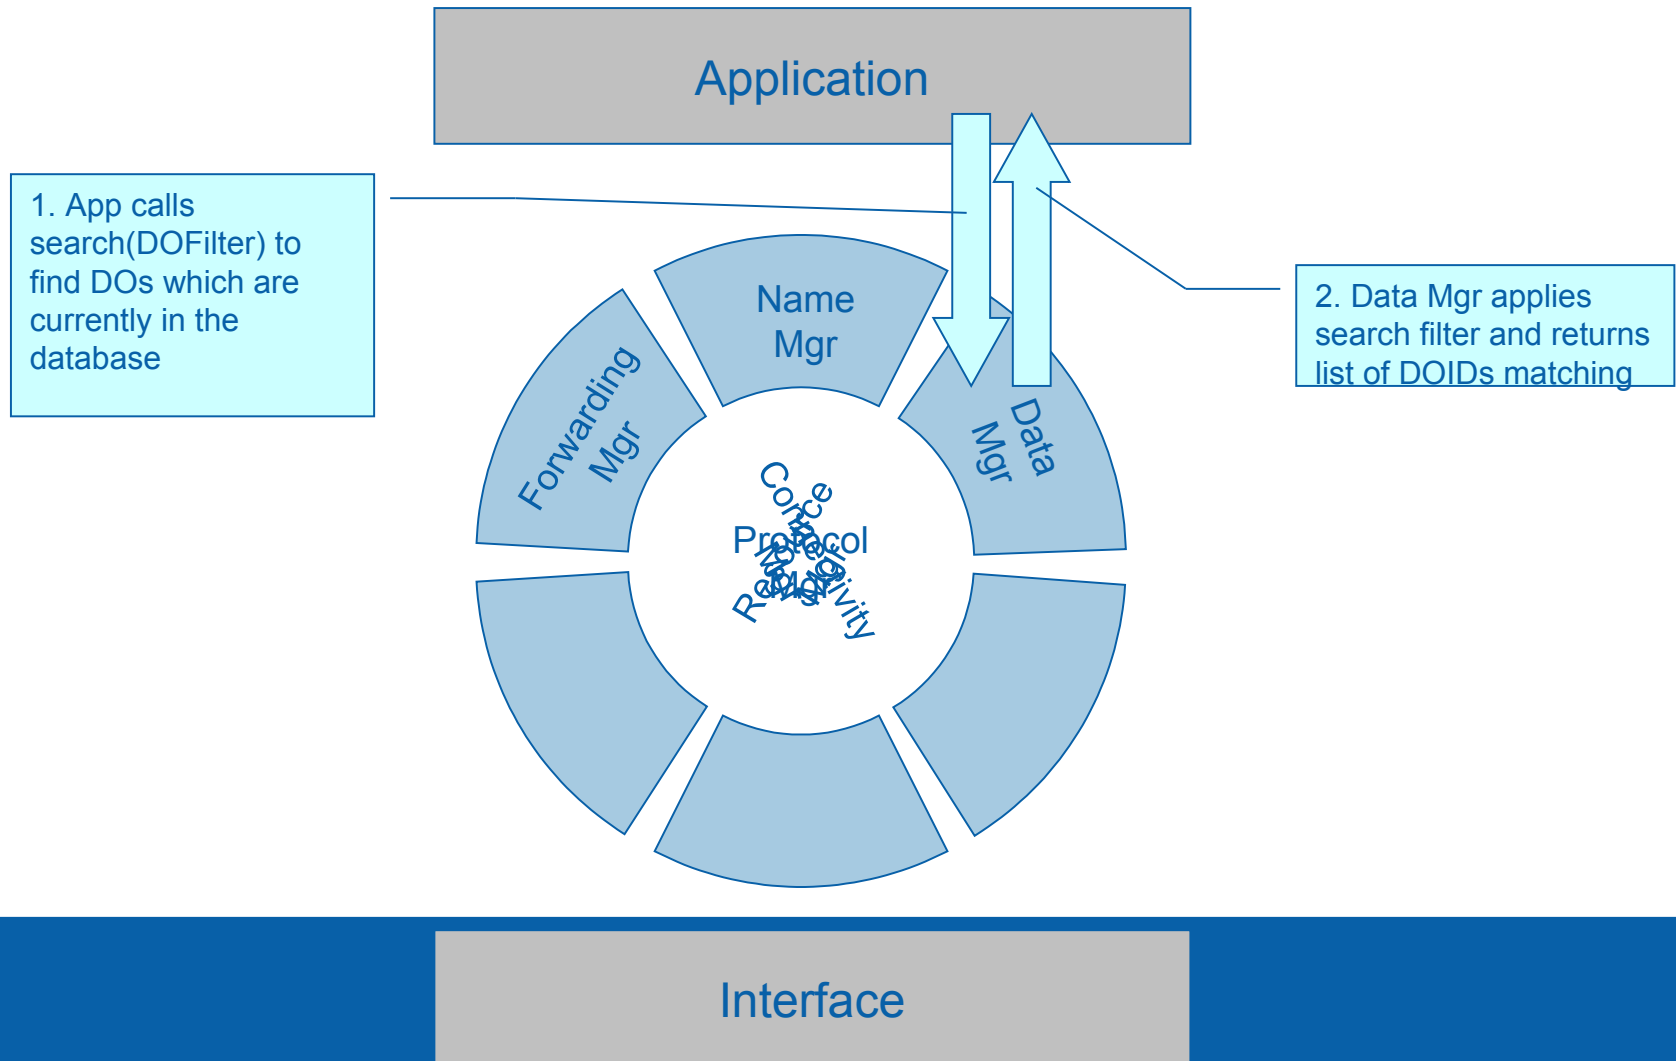
Outline

- Data Object (DO) app interface
- Names and their use in delivery
- Neighbour discovery
- Send DO / DO reception
- Solicit DO

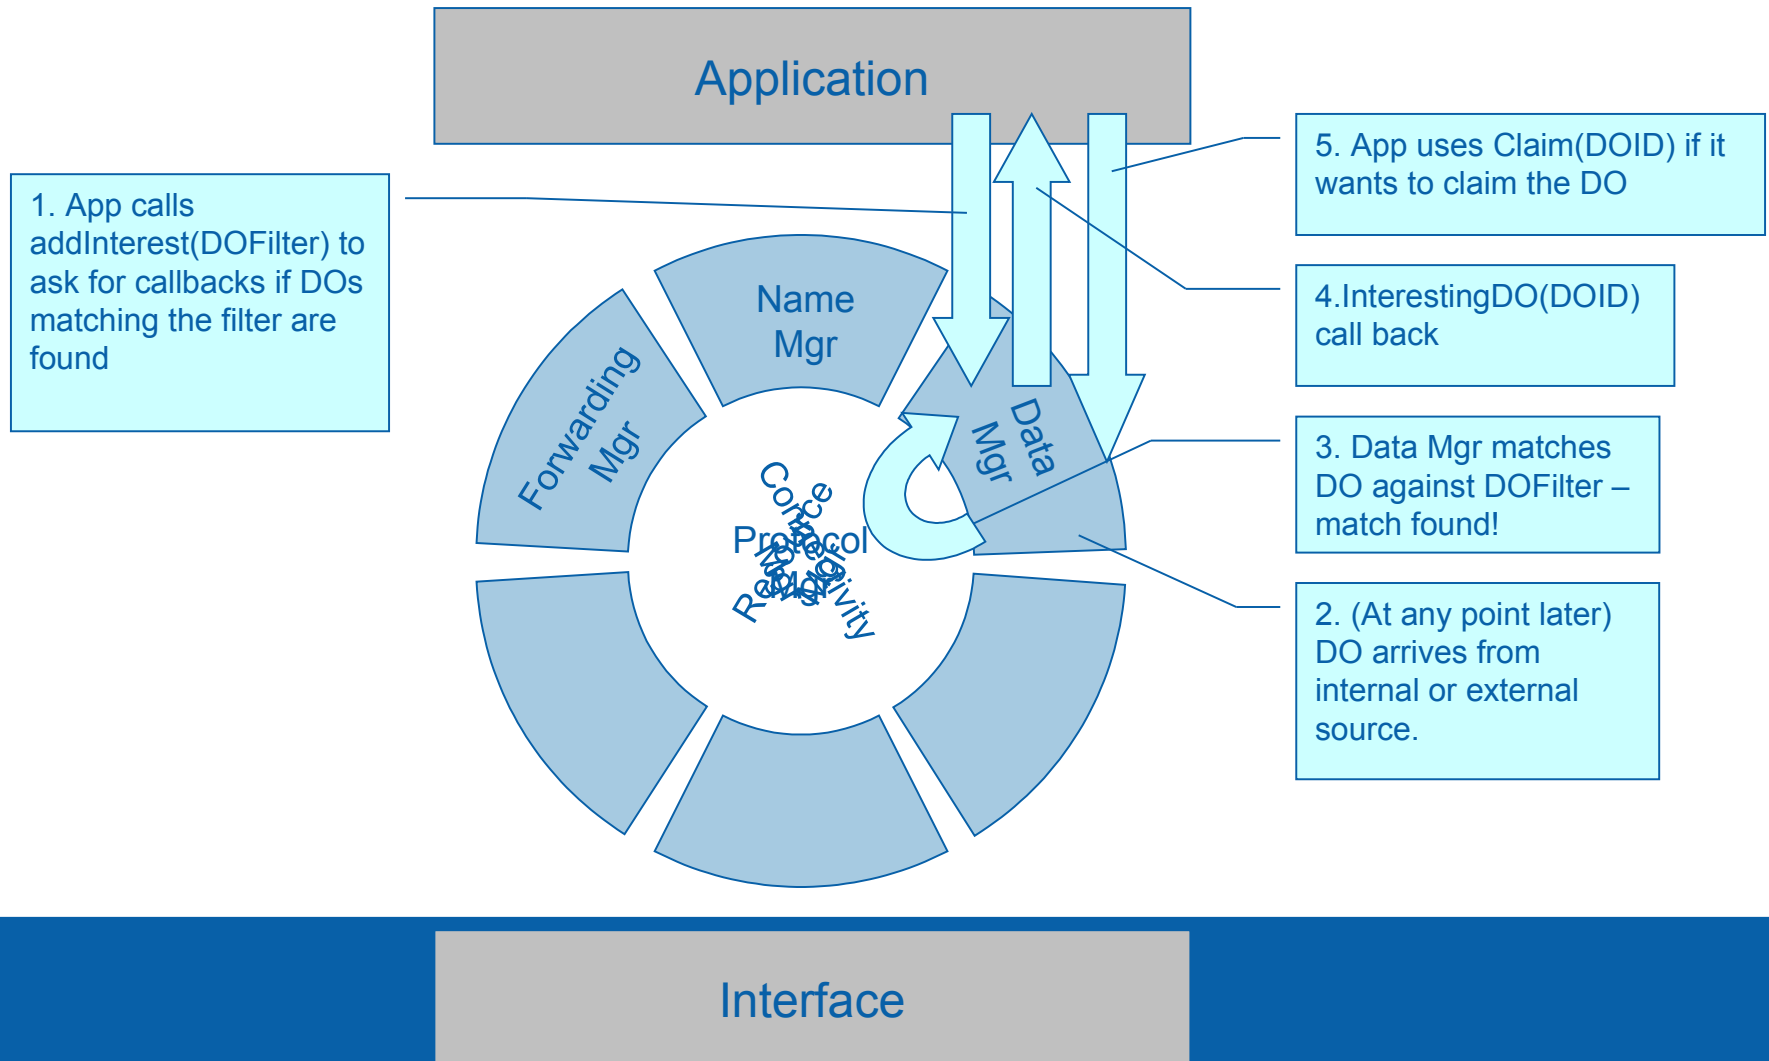
Data Mgr interface (insert/claim)



Data Mgr interface (search)



Data Mgr interface (callback)



Outline

- Data Object (DO) app interface
- Names and their use in delivery
- Neighbour discovery
- Send DO / DO reception
- Solicit DO

Name Objects

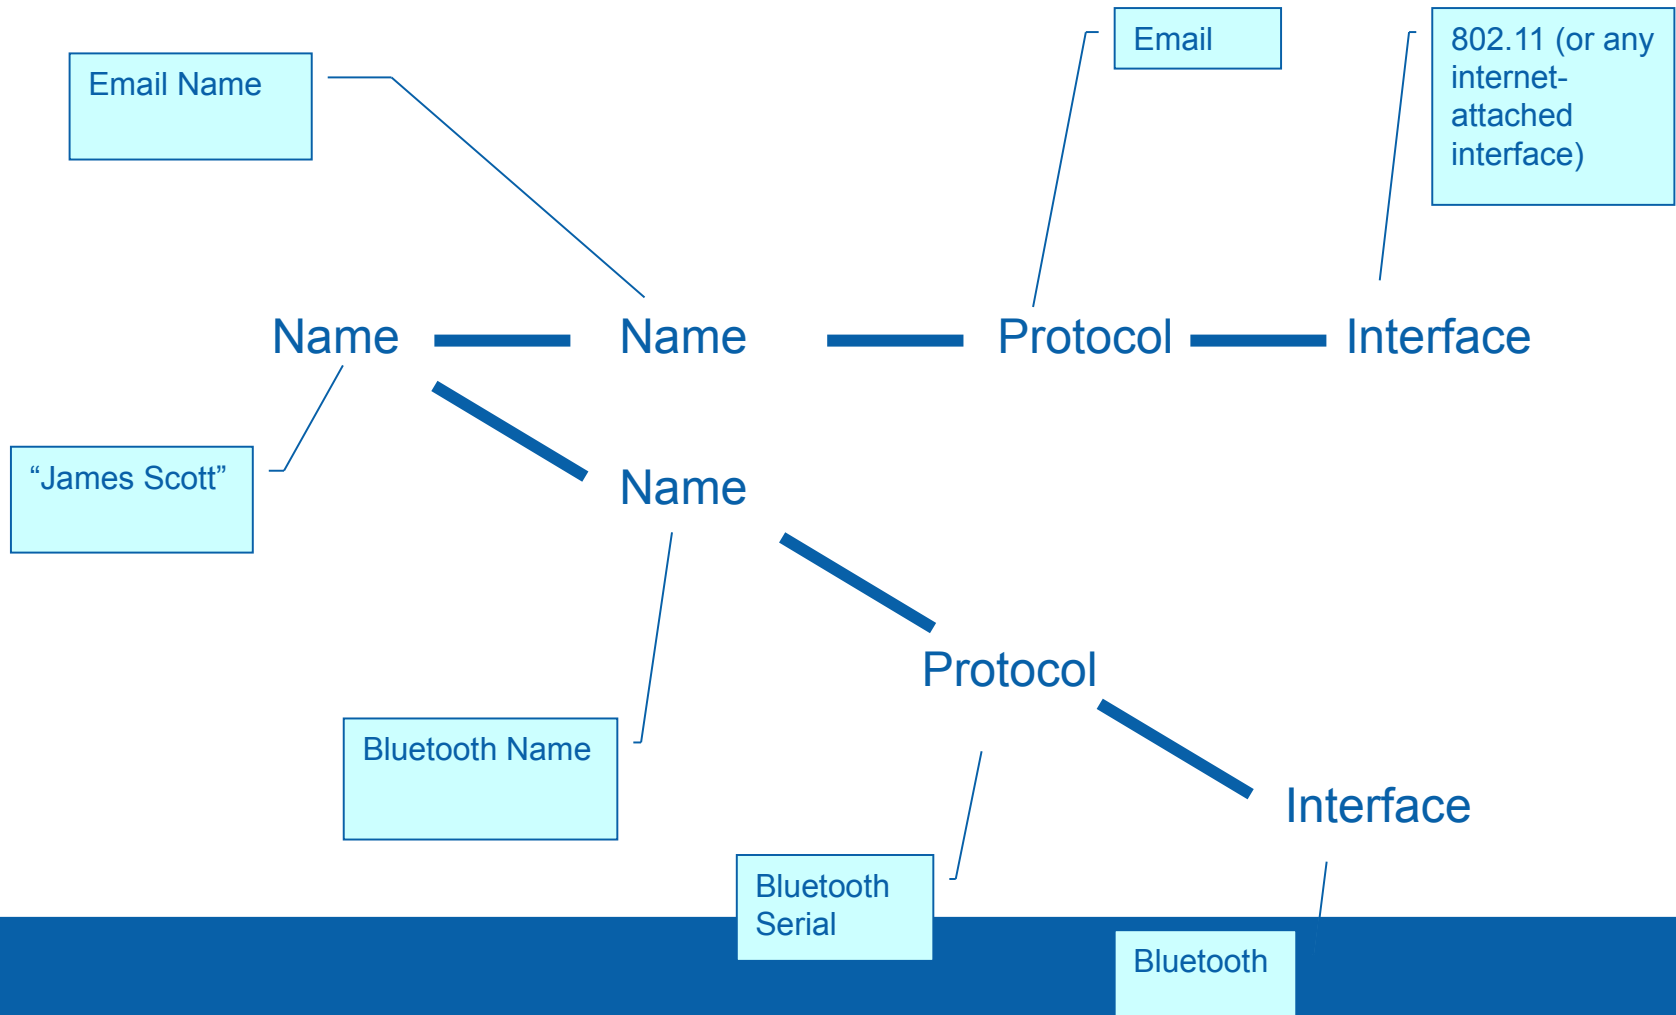
DO-ID	001
DO-Type	Name
Schema	Email
Identifier	James.w.scott@intel.com
Link	null

DO-ID	003
DO-TYPE	Name
Schema	String
Identifier	James Scott
Link	001,002

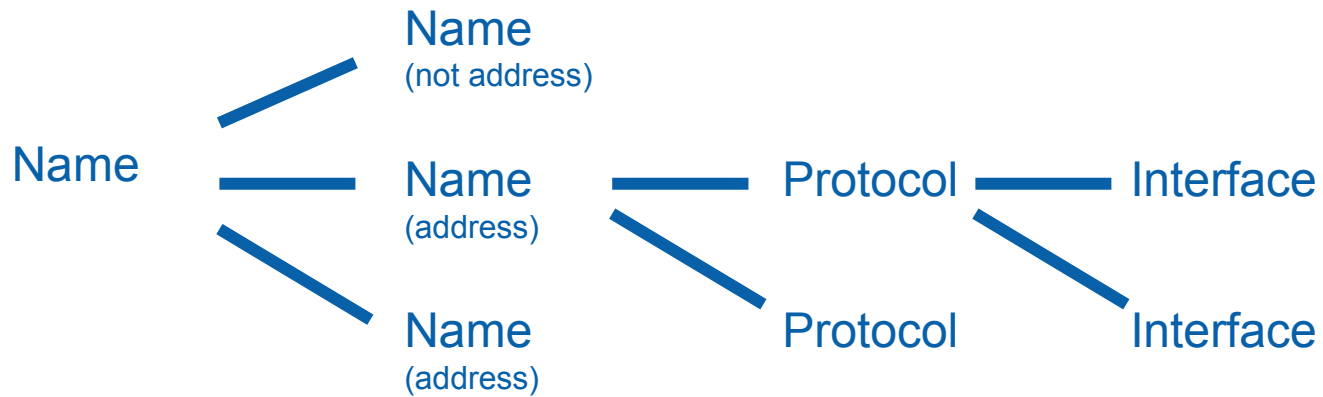
DO-ID	002
DO-Type	Name
Schema	Bluetooth
Identifier	ab:cd:ef:01:23:45
Link	null

- Name objects comprise at least a Schema, Identifier and Link type entries. All values can be null.
- Name records can be created by the Application, the Connectivity Mgr, the Forwarding Mgr, the Protocol Mgr, etc via an interface provided by the Name Mgr
- Name objects can be annotated

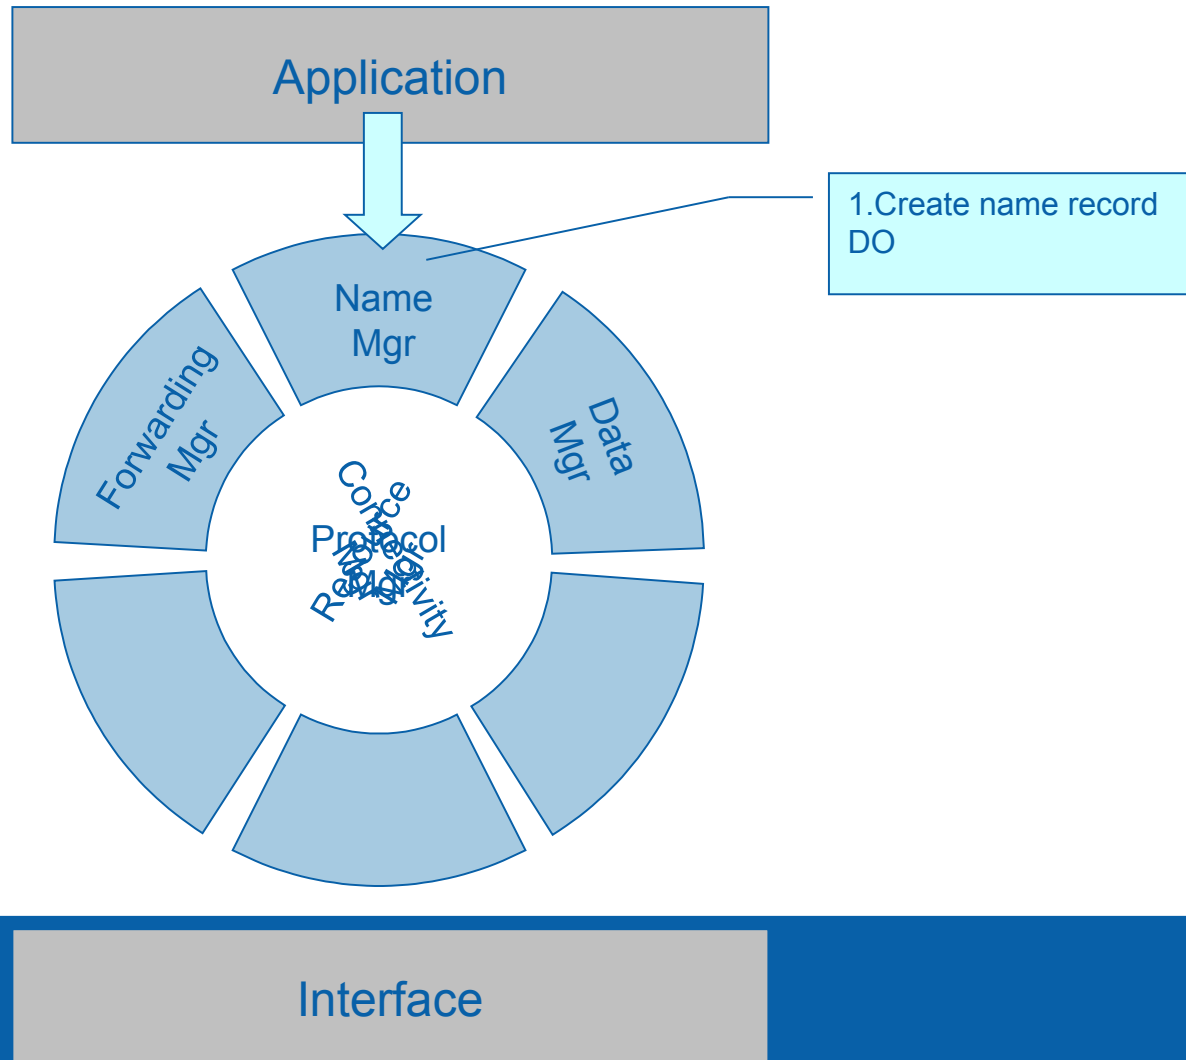
Example use of naming in delivery



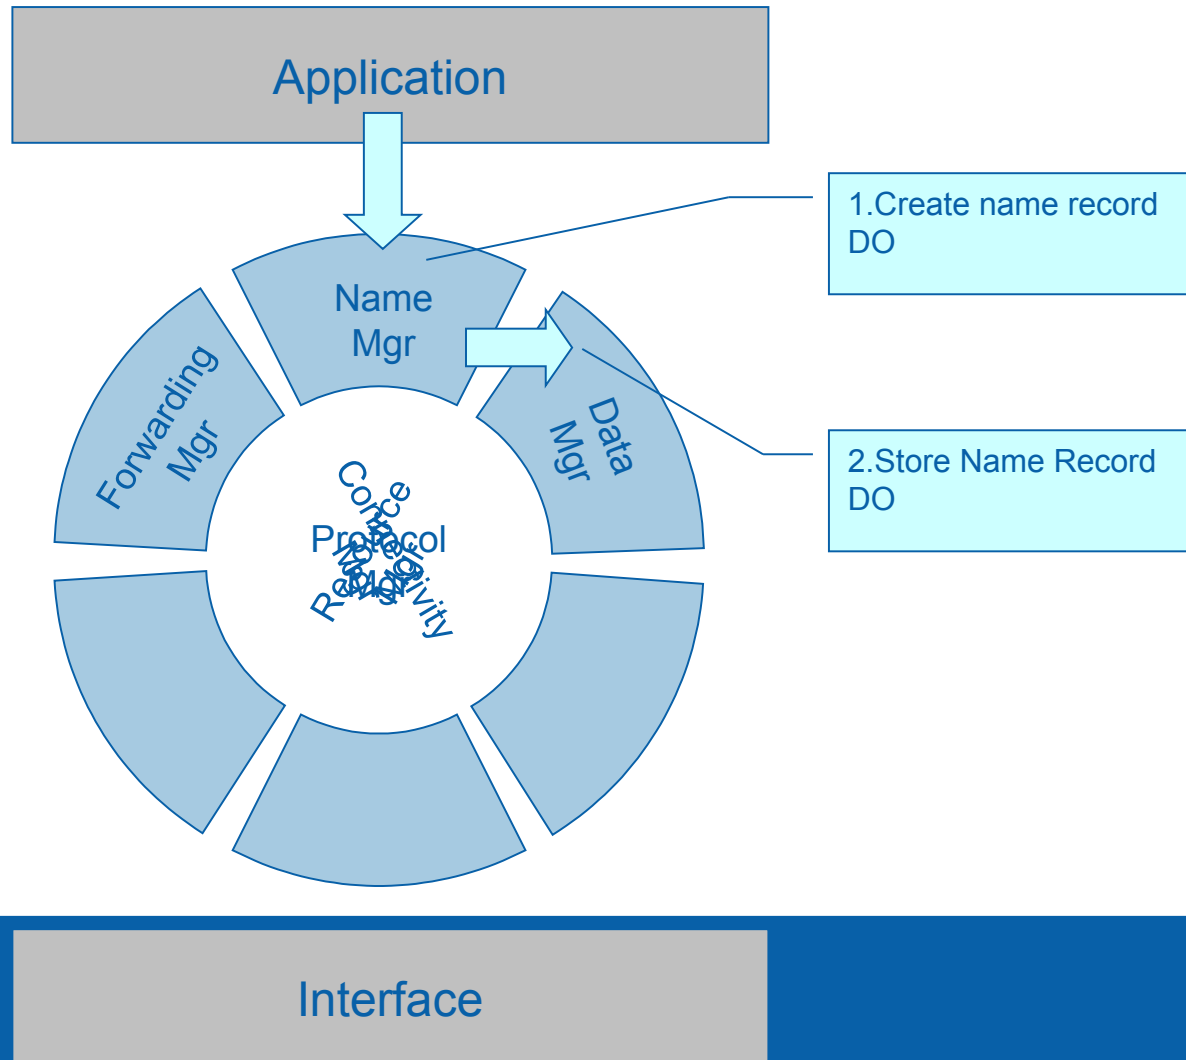
General use of naming in delivery



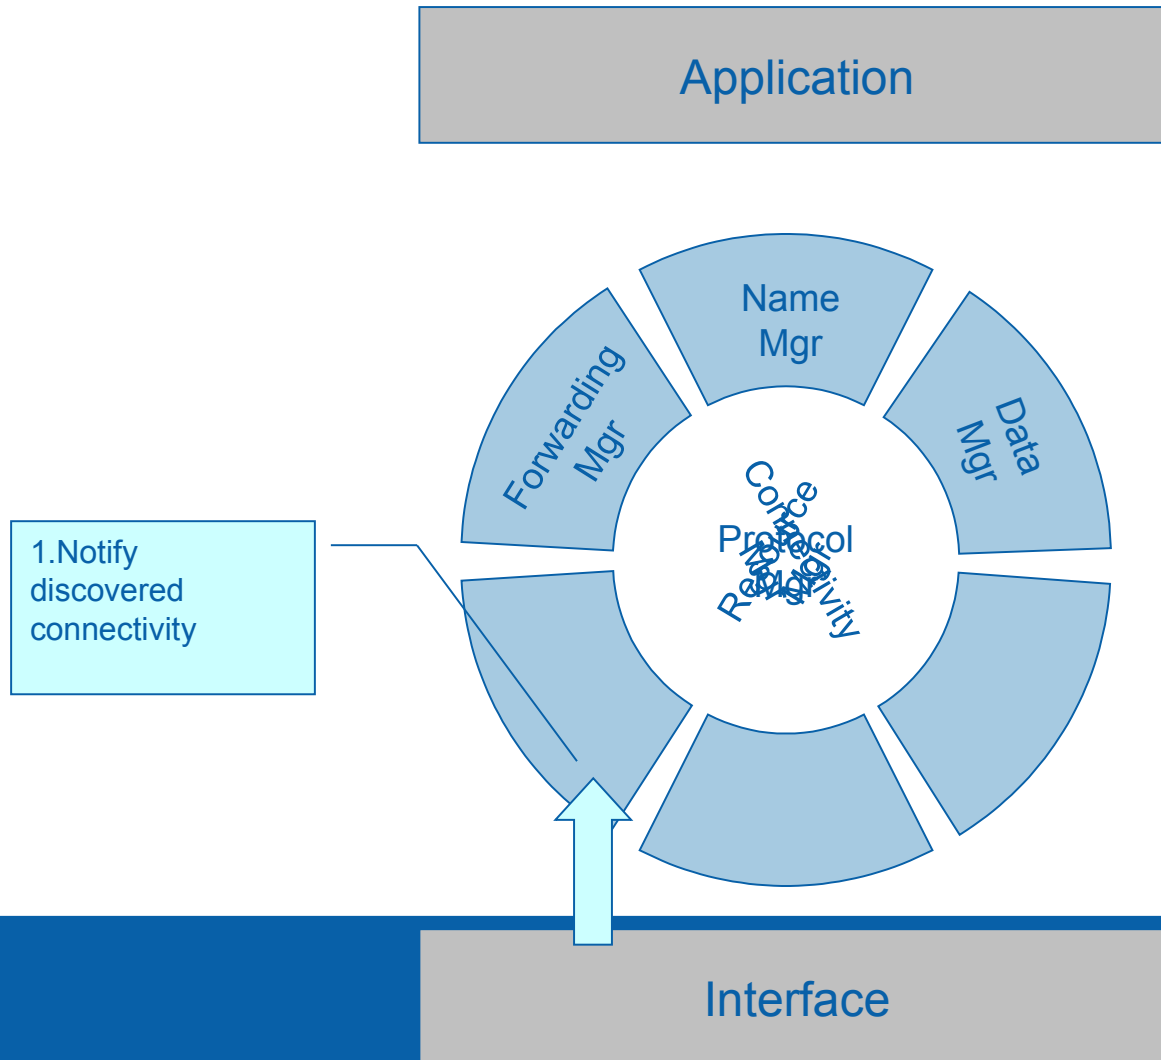
Create Name Data (App)



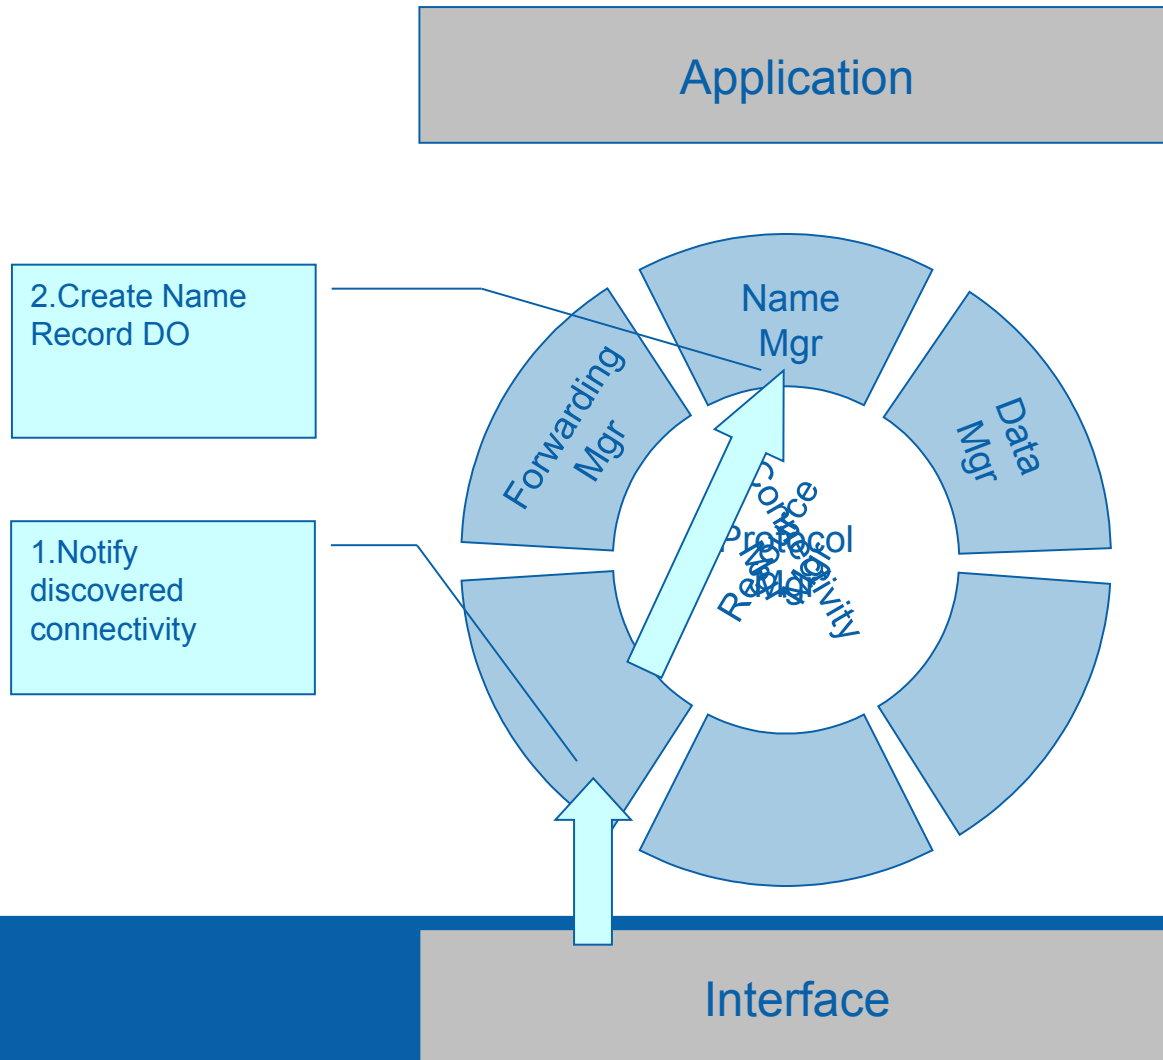
Create Name Object (App)



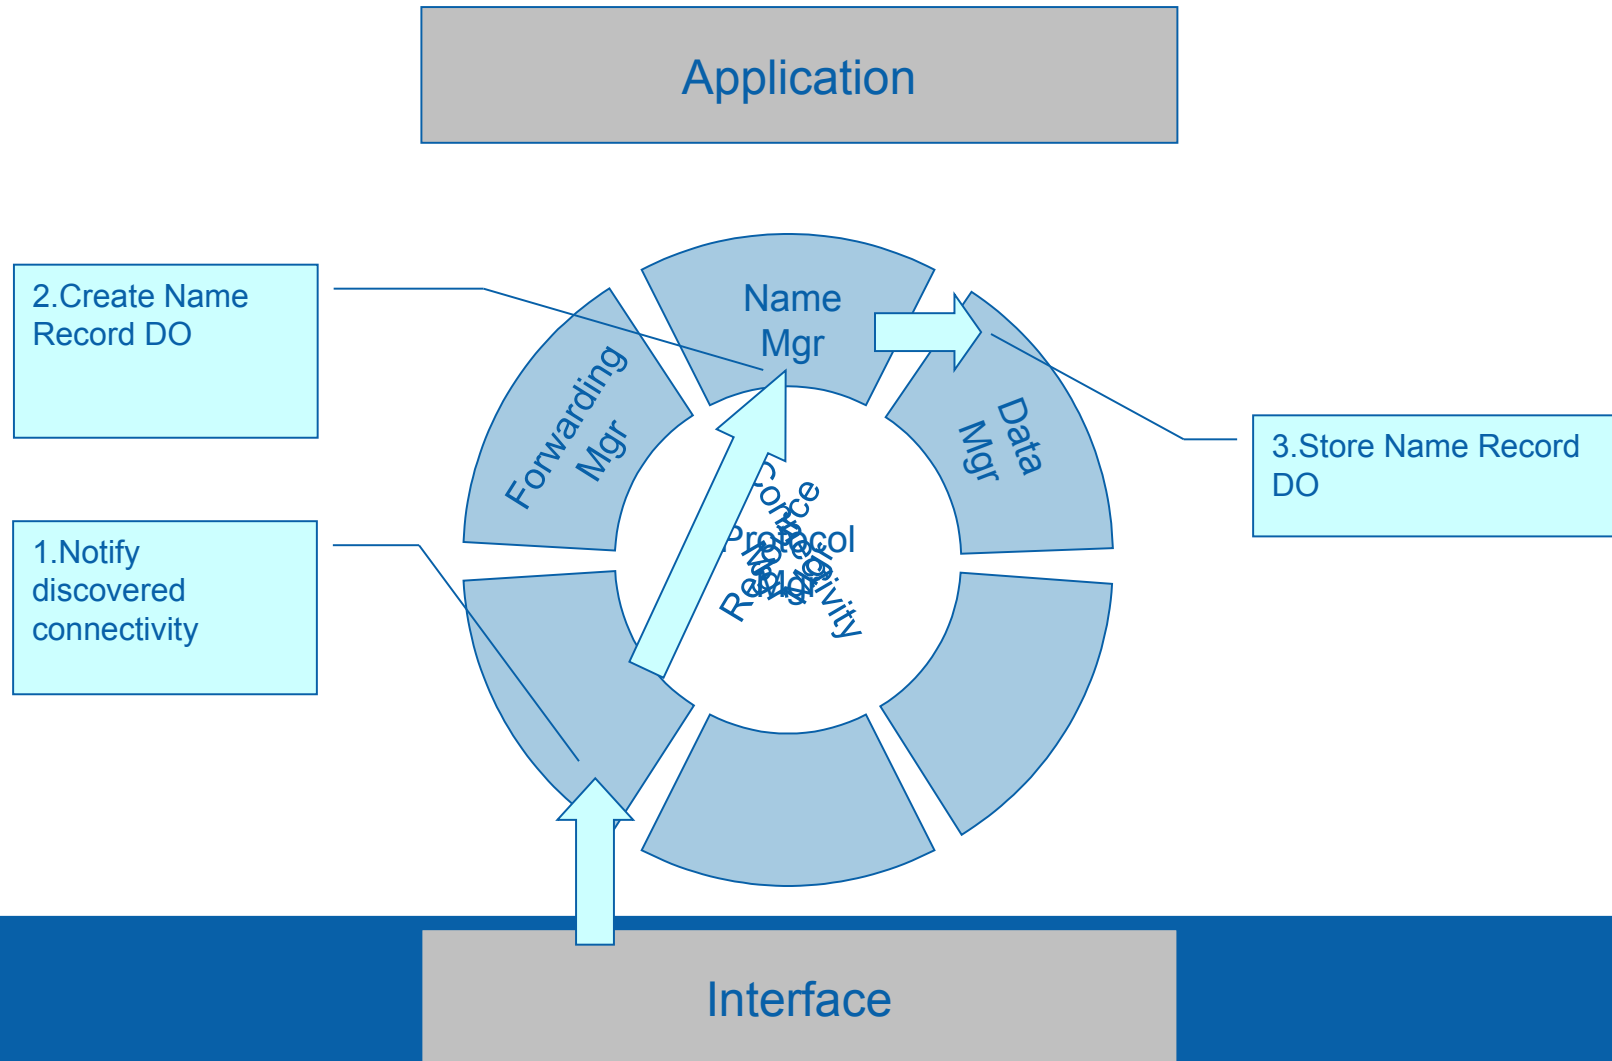
Create Name Object (Network)



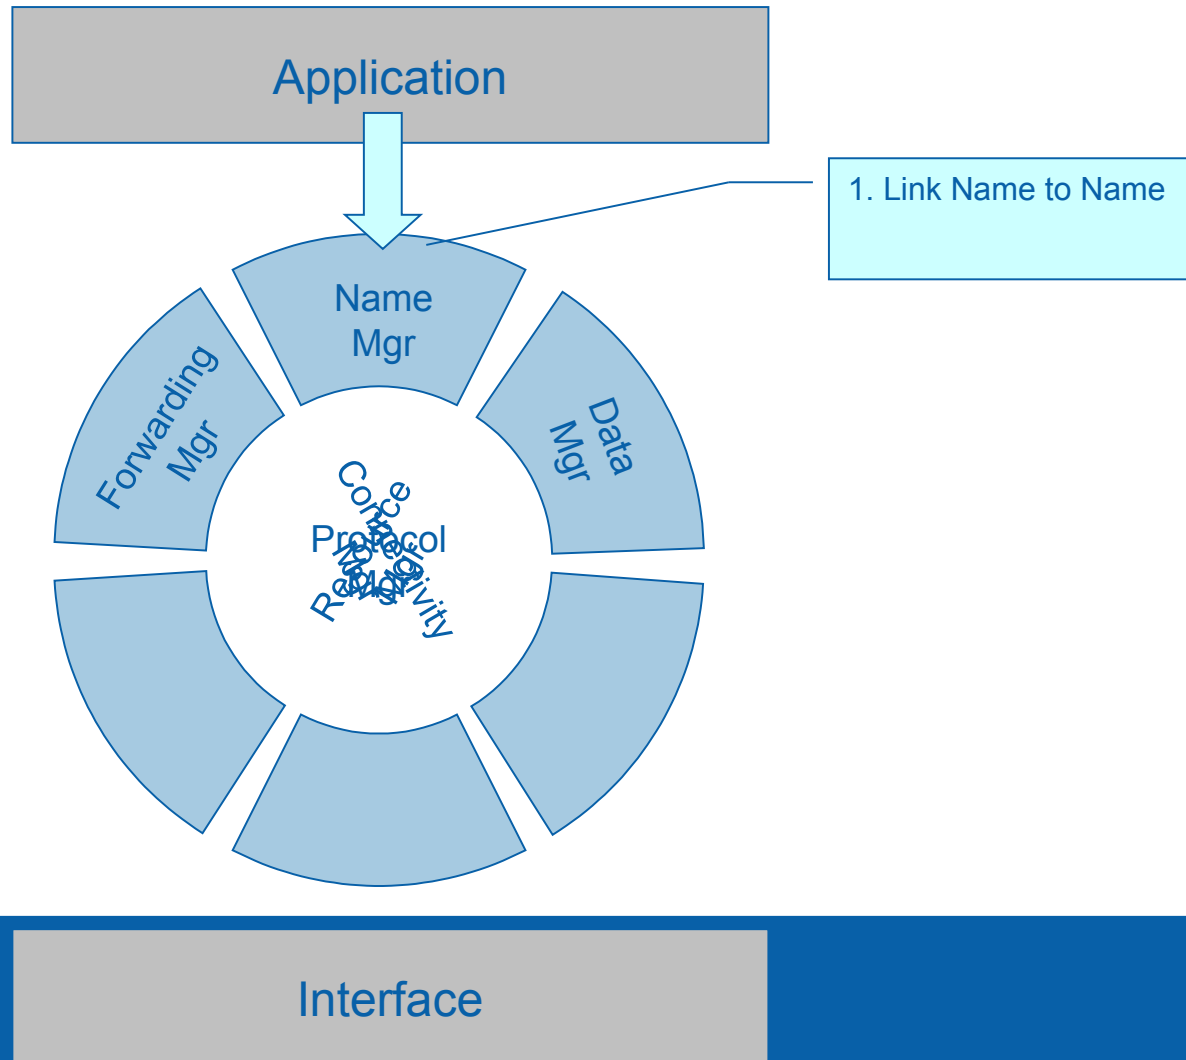
Create Name Object (Network)



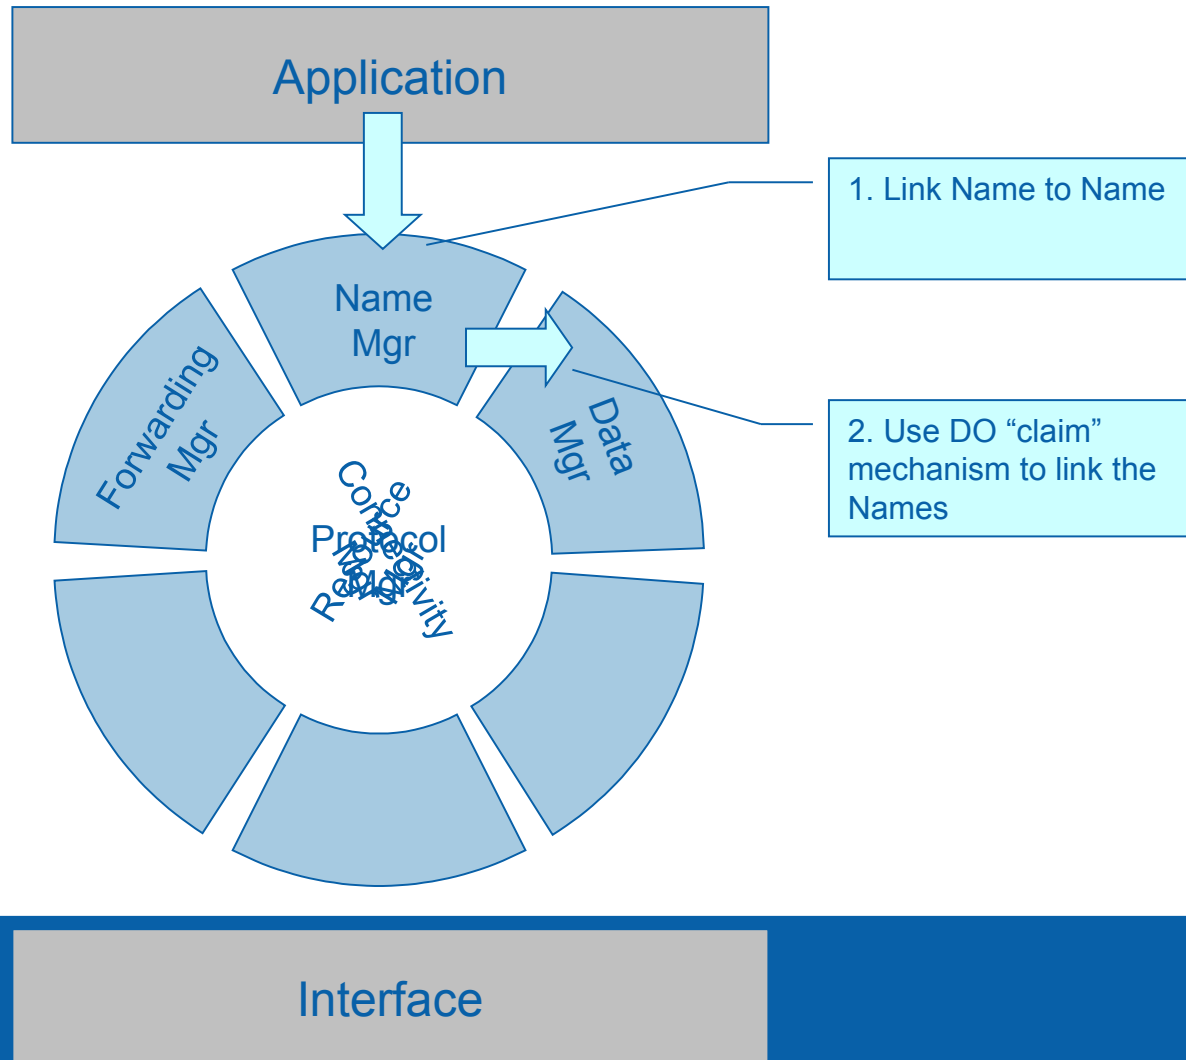
Create Name Object (Network)



Name Linking (App)




Name Linking (App)



Name Linking (Forward Mgr)

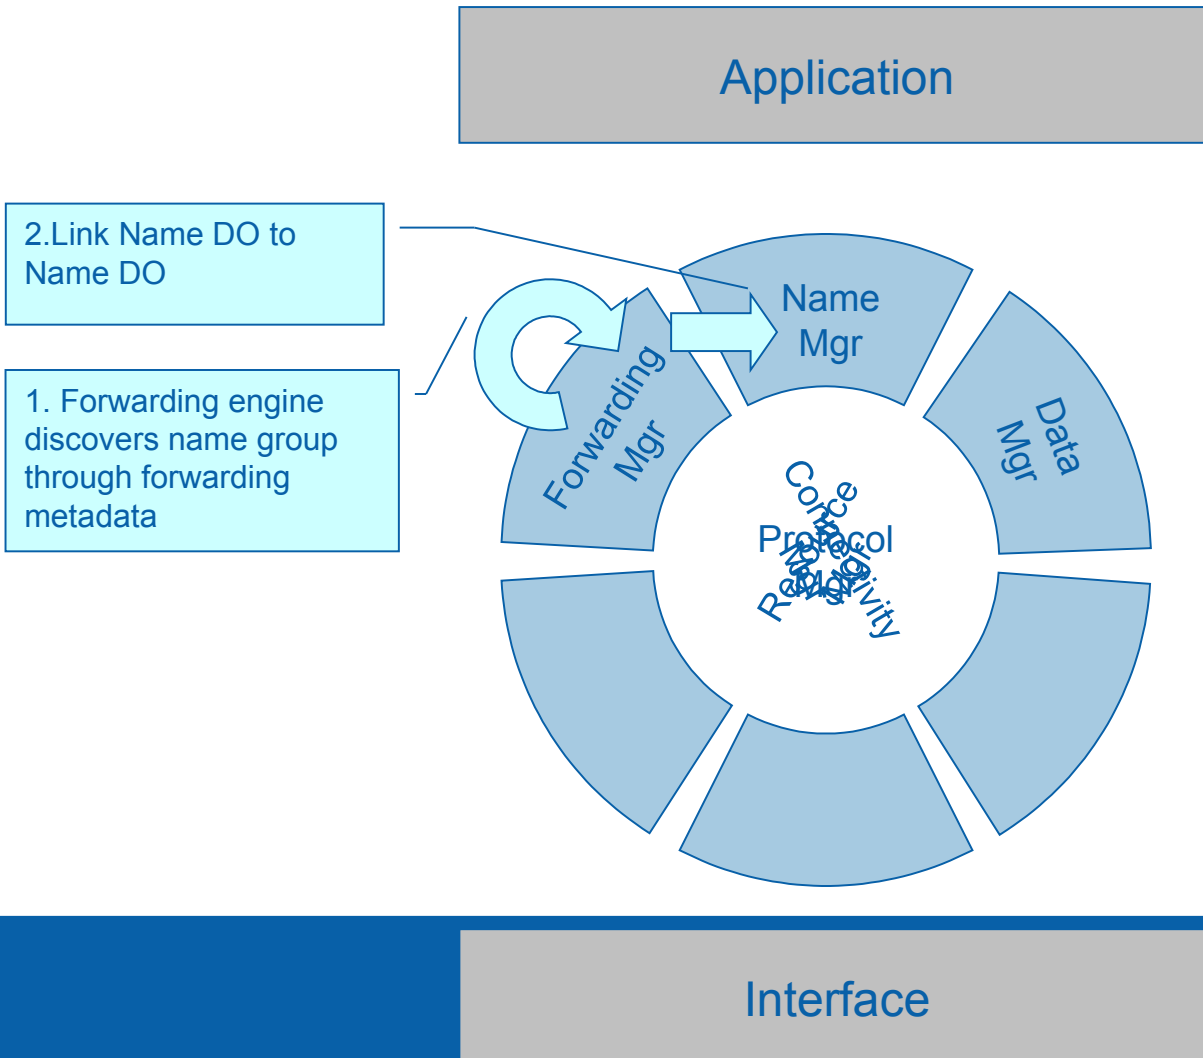
Application

1. Forwarding engine discovers name mapping through forwarding metadata

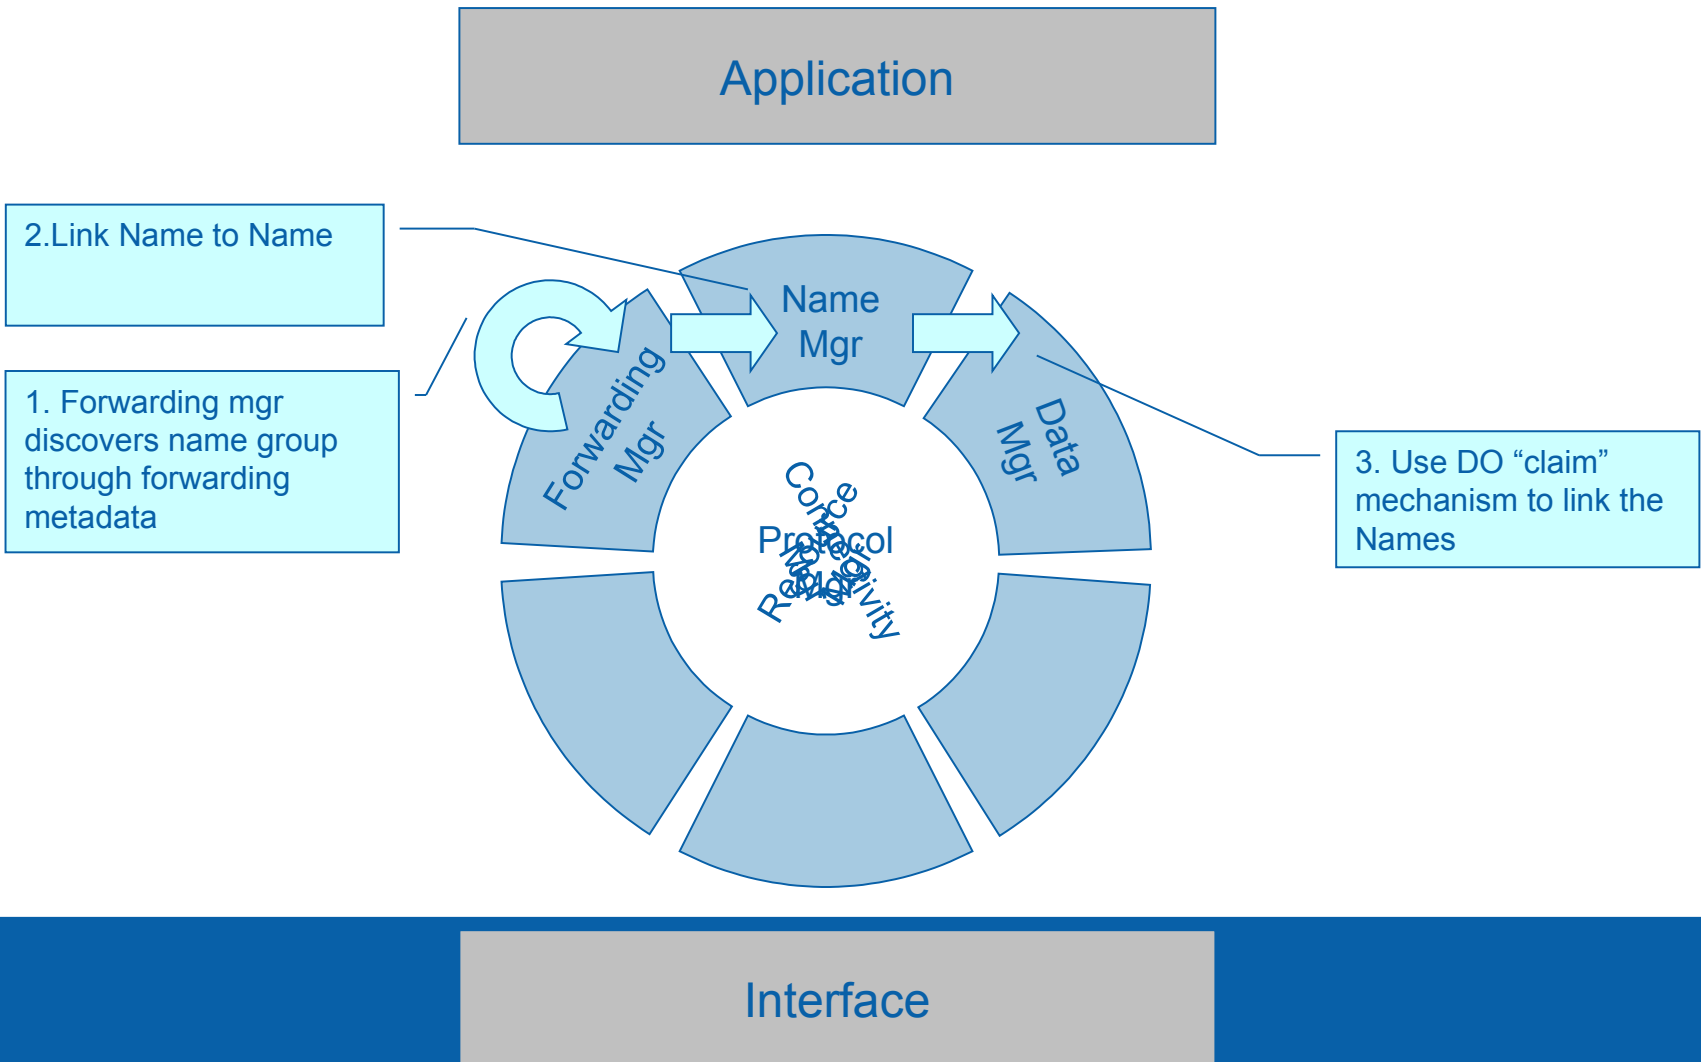


Interface

Name Linking (Forward Mgr)



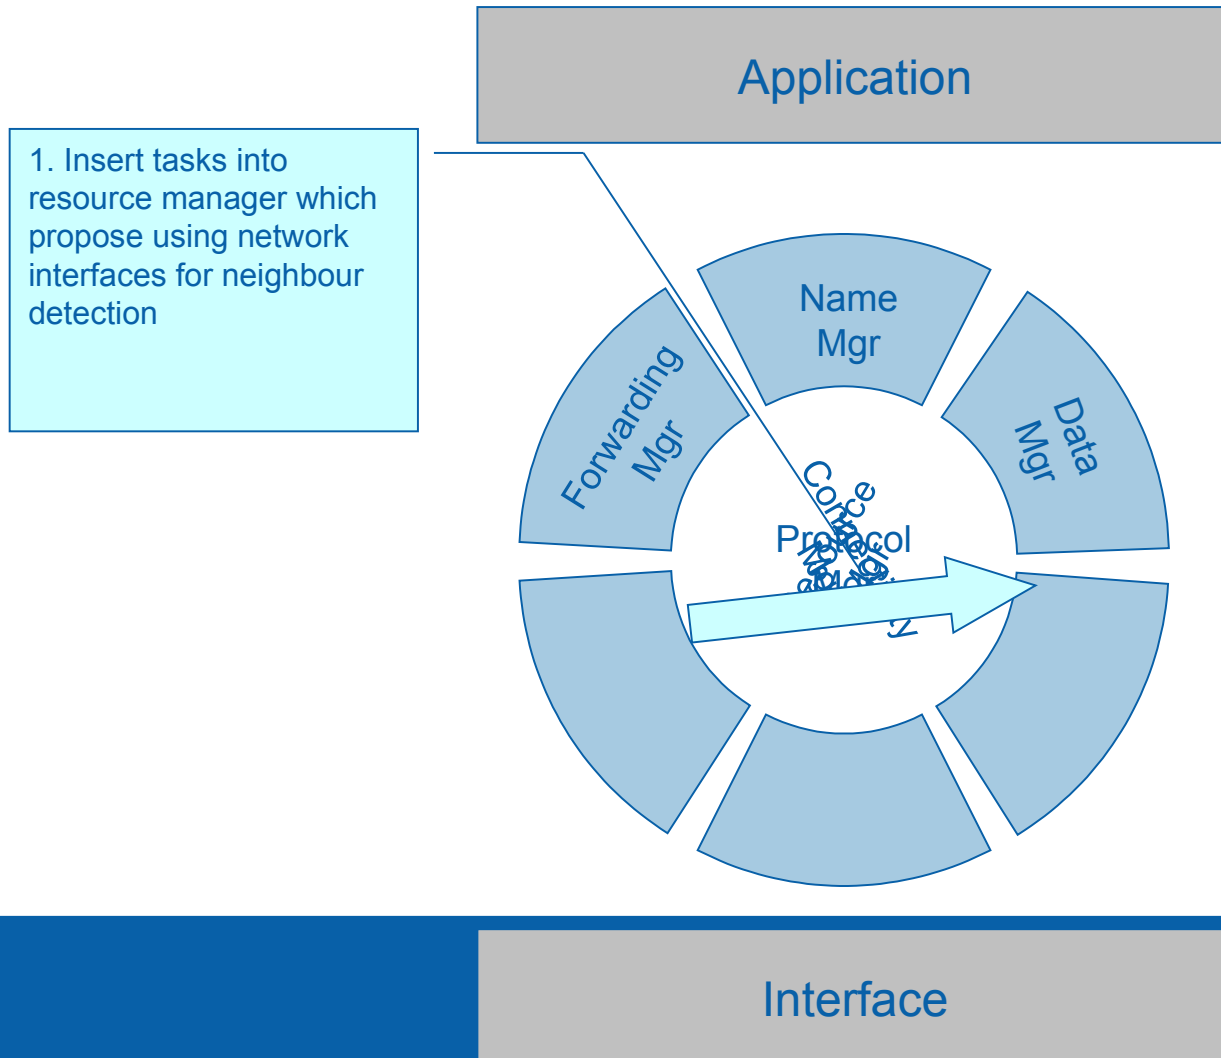
Name Linking (Forward Mgr)



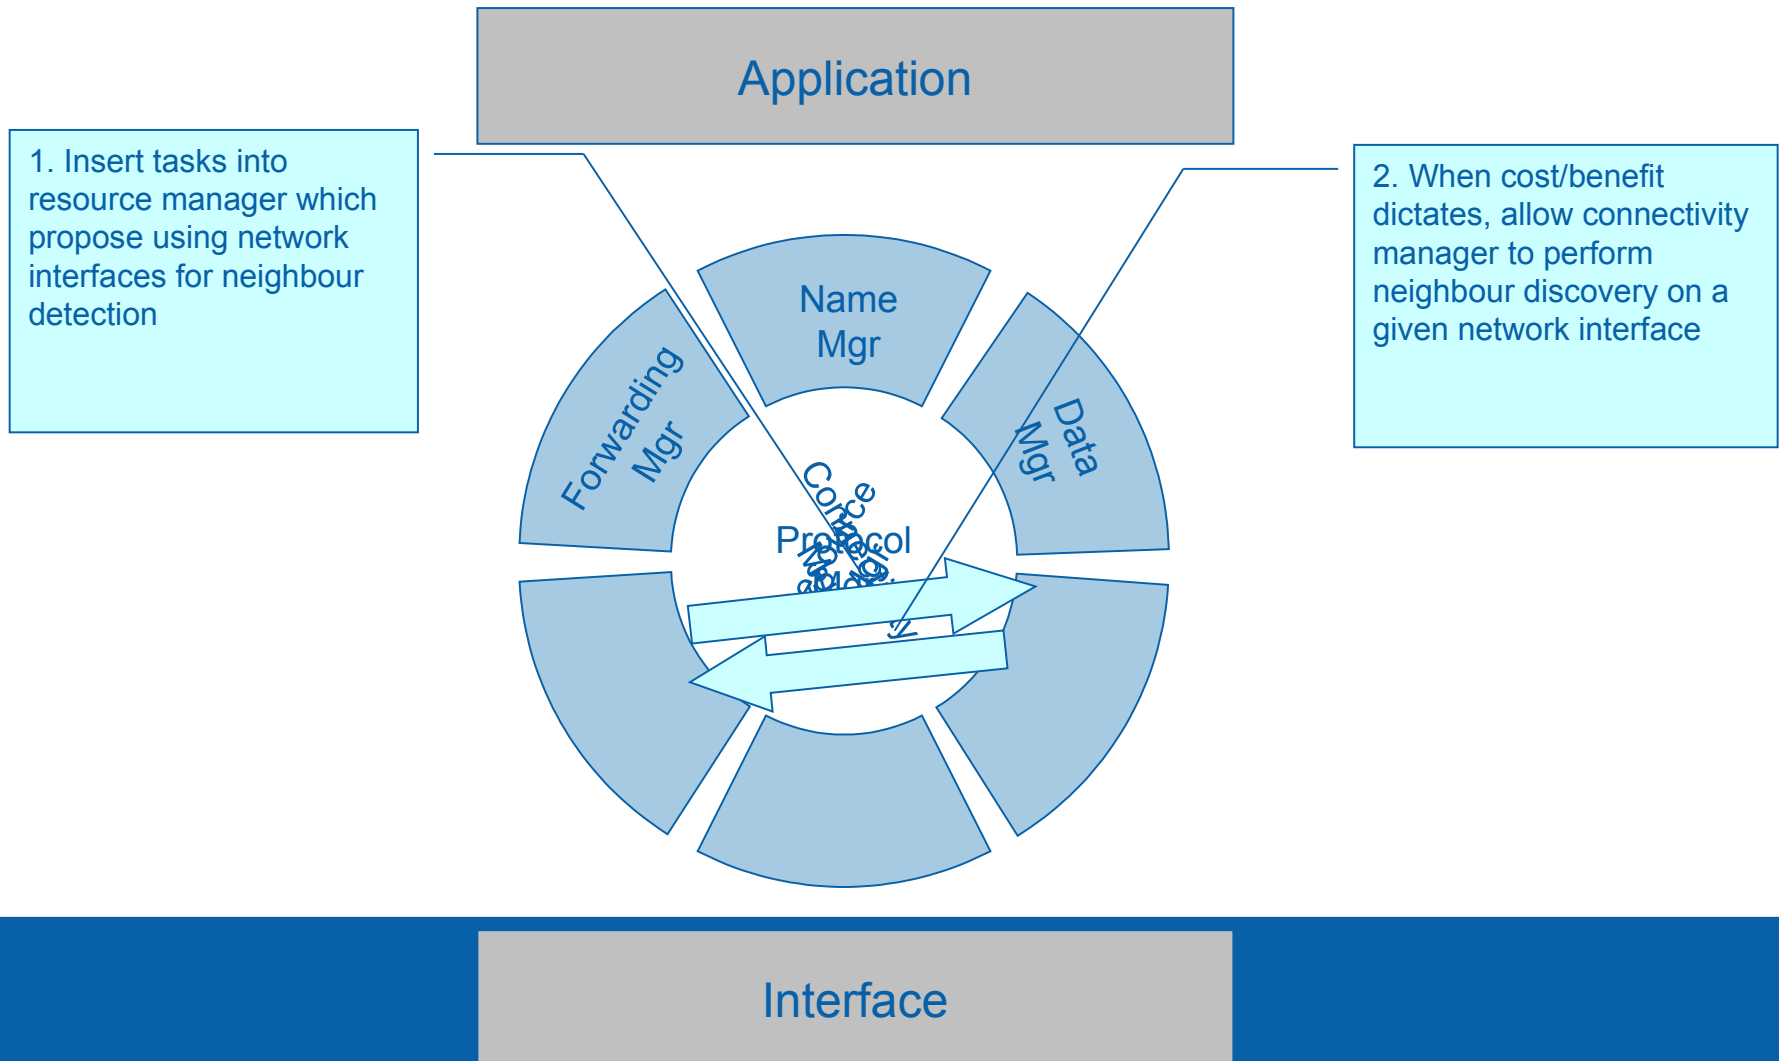
Outline

- Proposed Hagggle Framework (changed)
- Data Object (DO) app interface
- Names and their use in delivery
- Neighbour discovery
- Send DO / DO reception
- Solicit DO

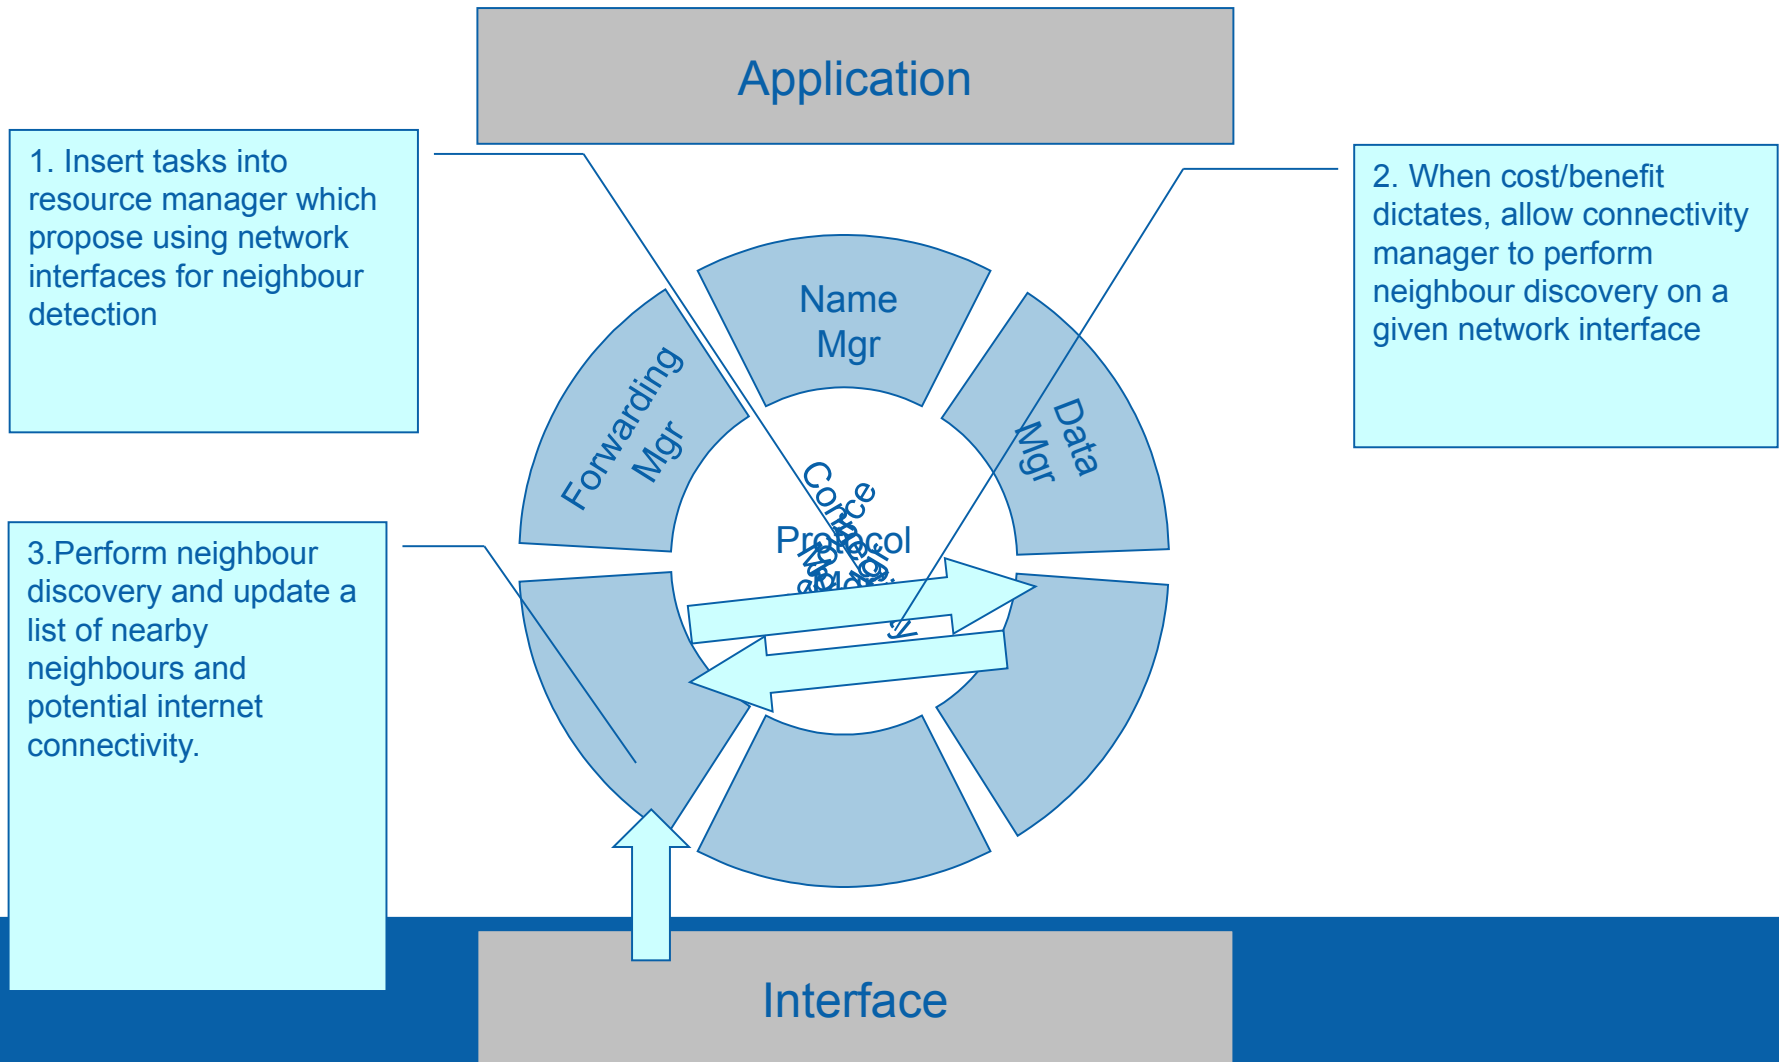
Neighbour Discovery



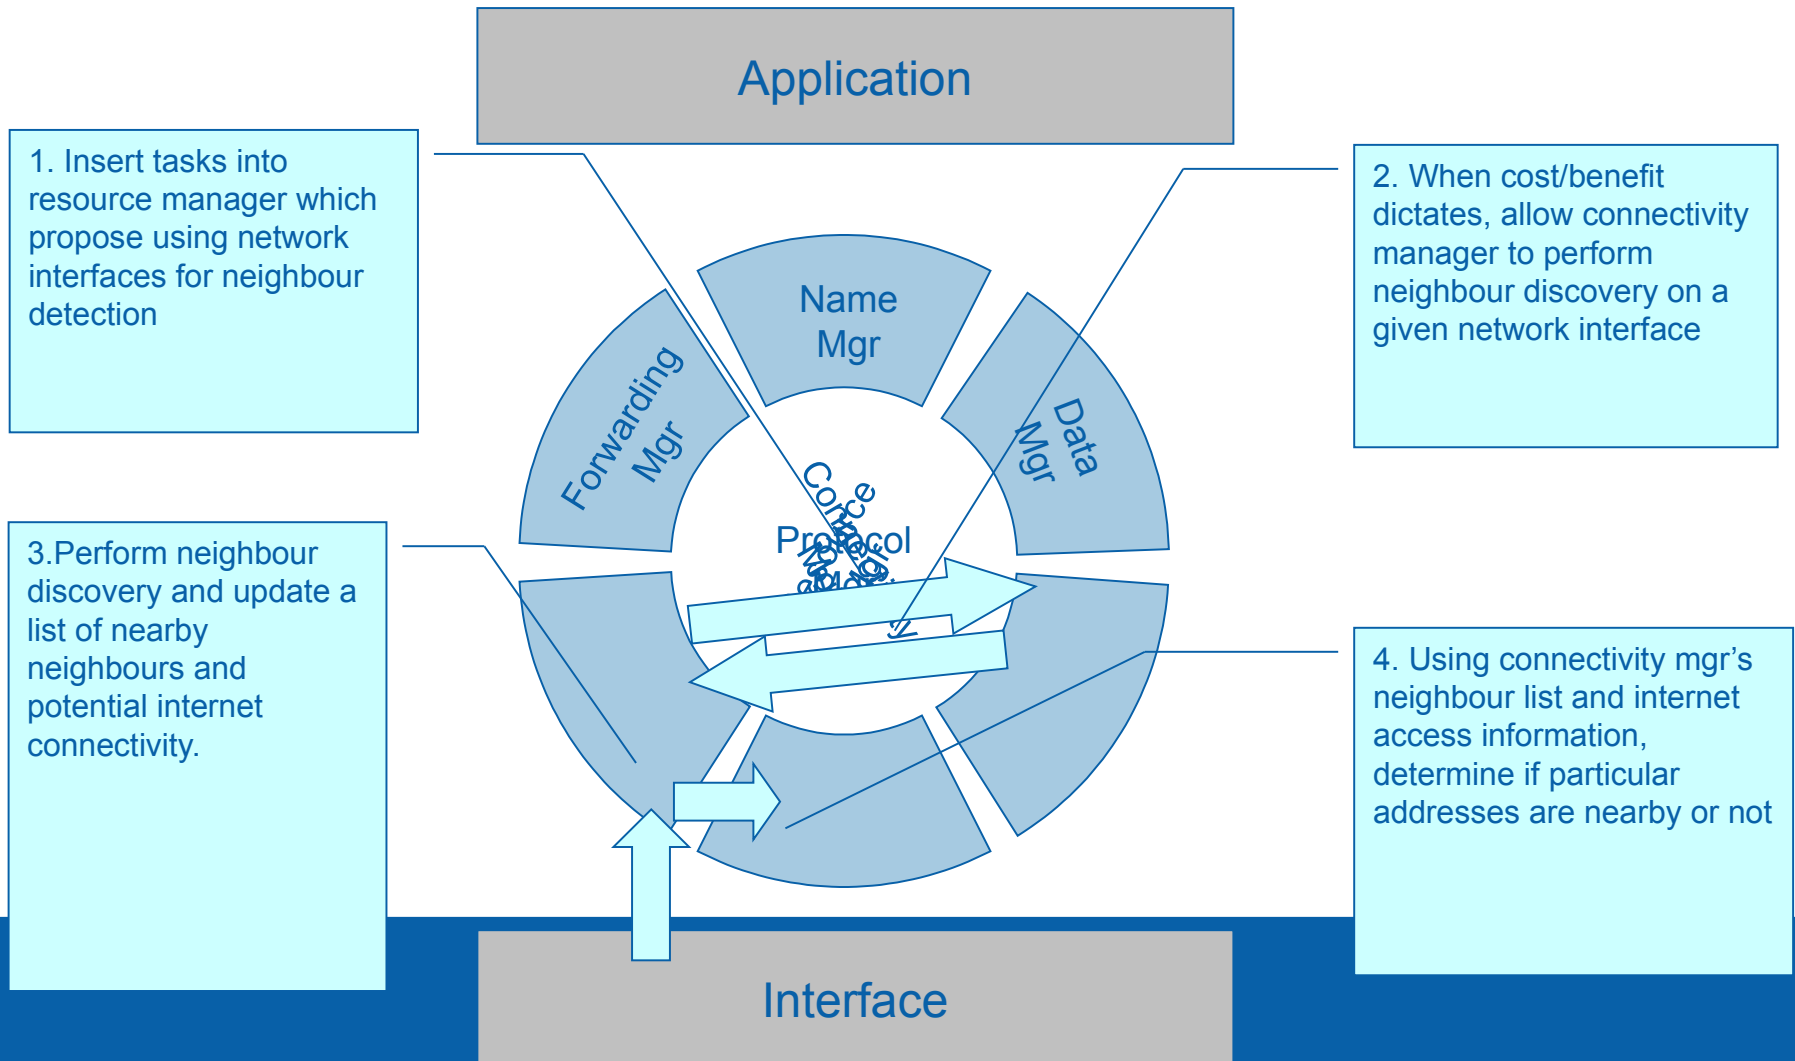
Neighbour Discovery



Neighbour Discovery



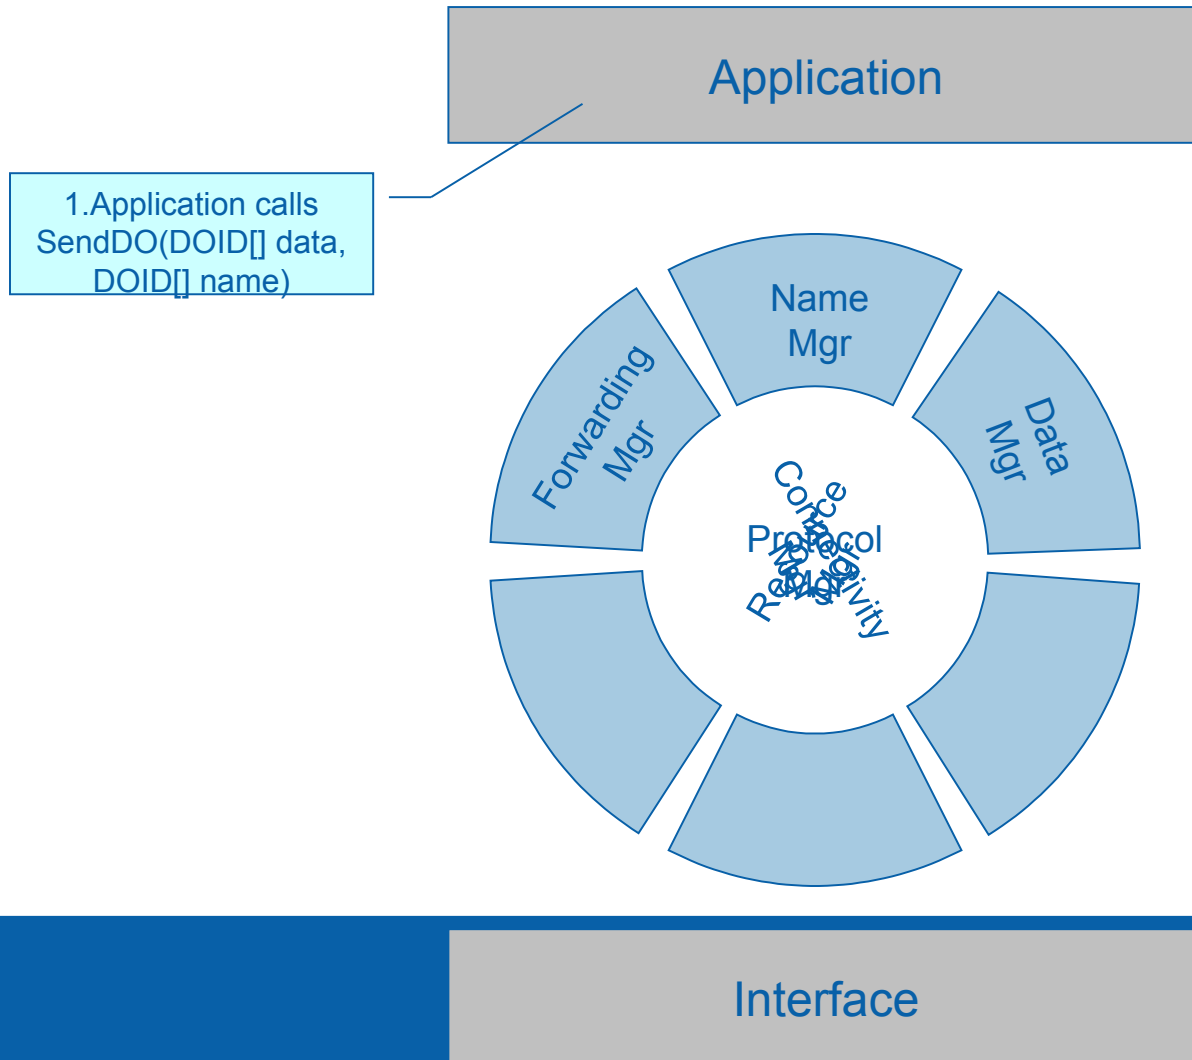
Neighbour Discovery



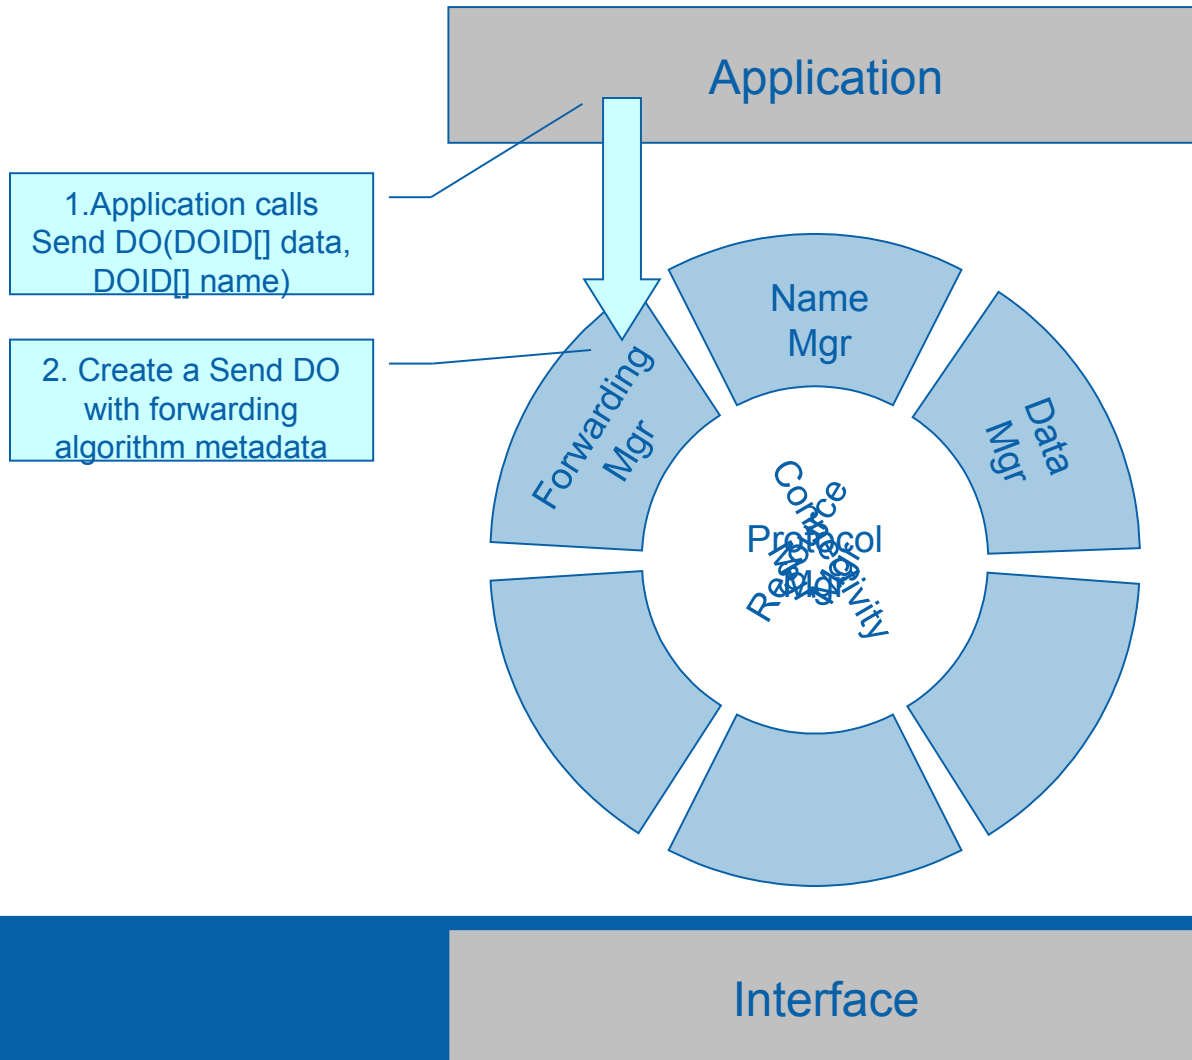
Outline

- Data Object (DO) app interface
- Names and their use in delivery
- Neighbour discovery
- Send DO / DO reception
- Solicit DO

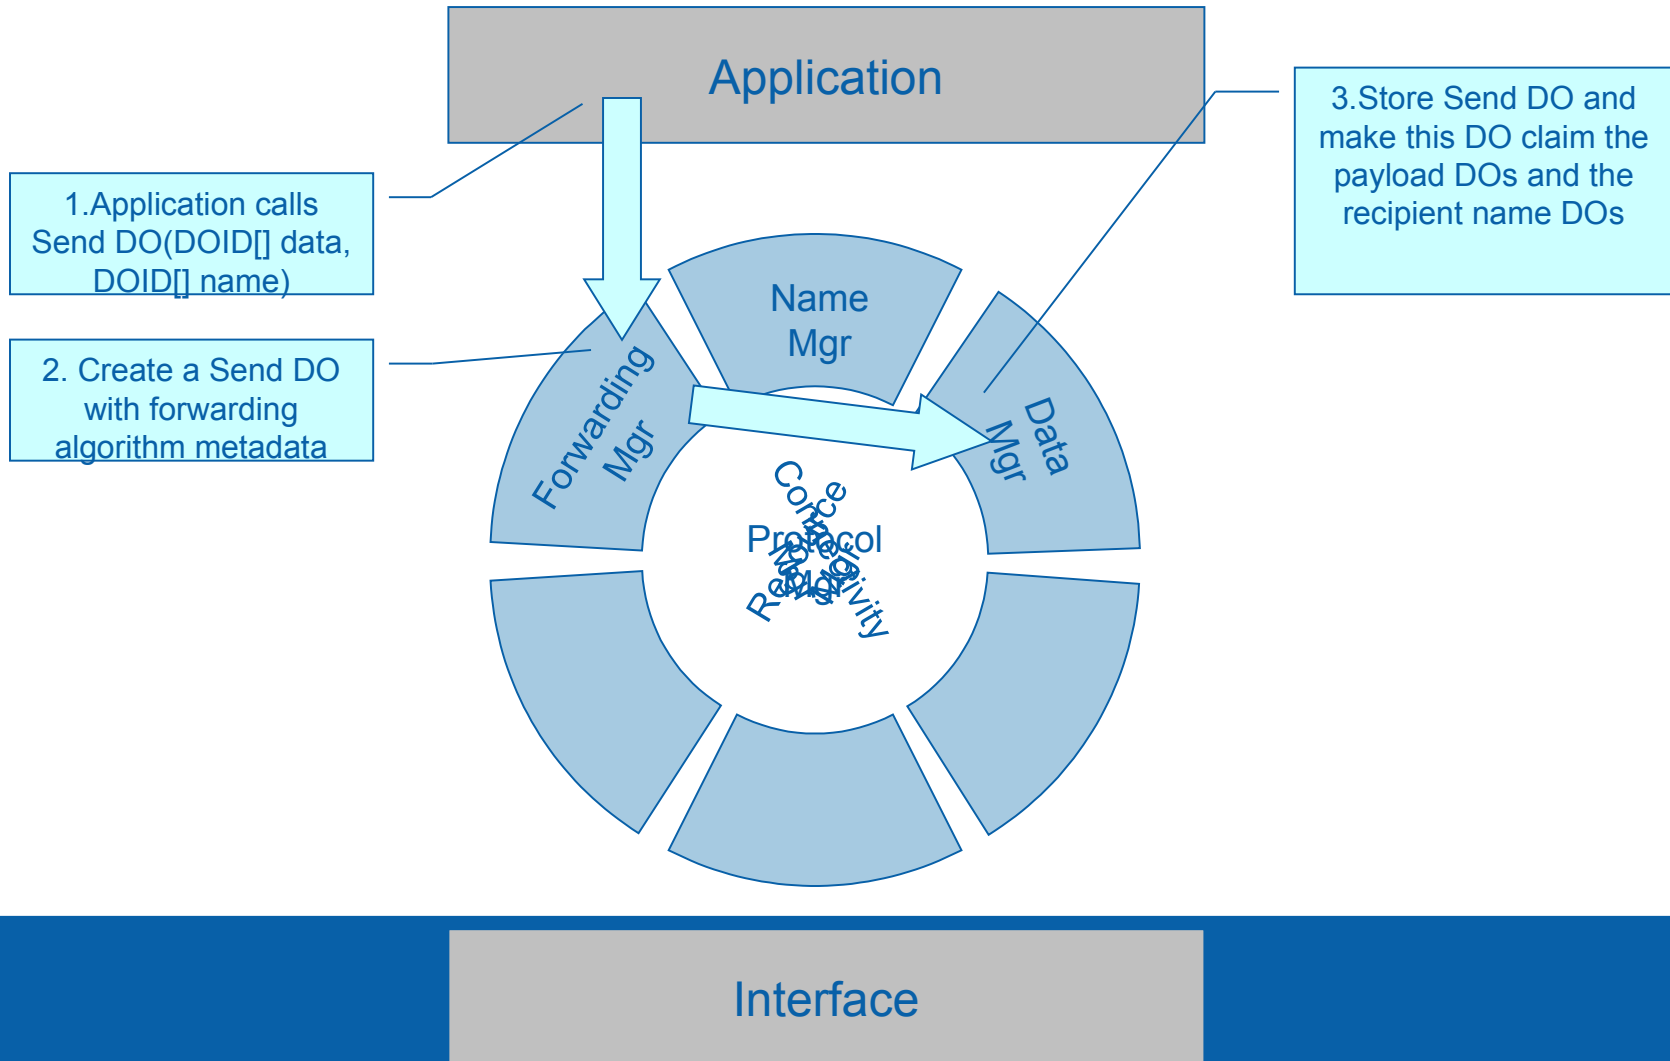
Send DO (App insertion)



Send DO (App insertion)




Send DO (App insertion)



Next-Hop decisions

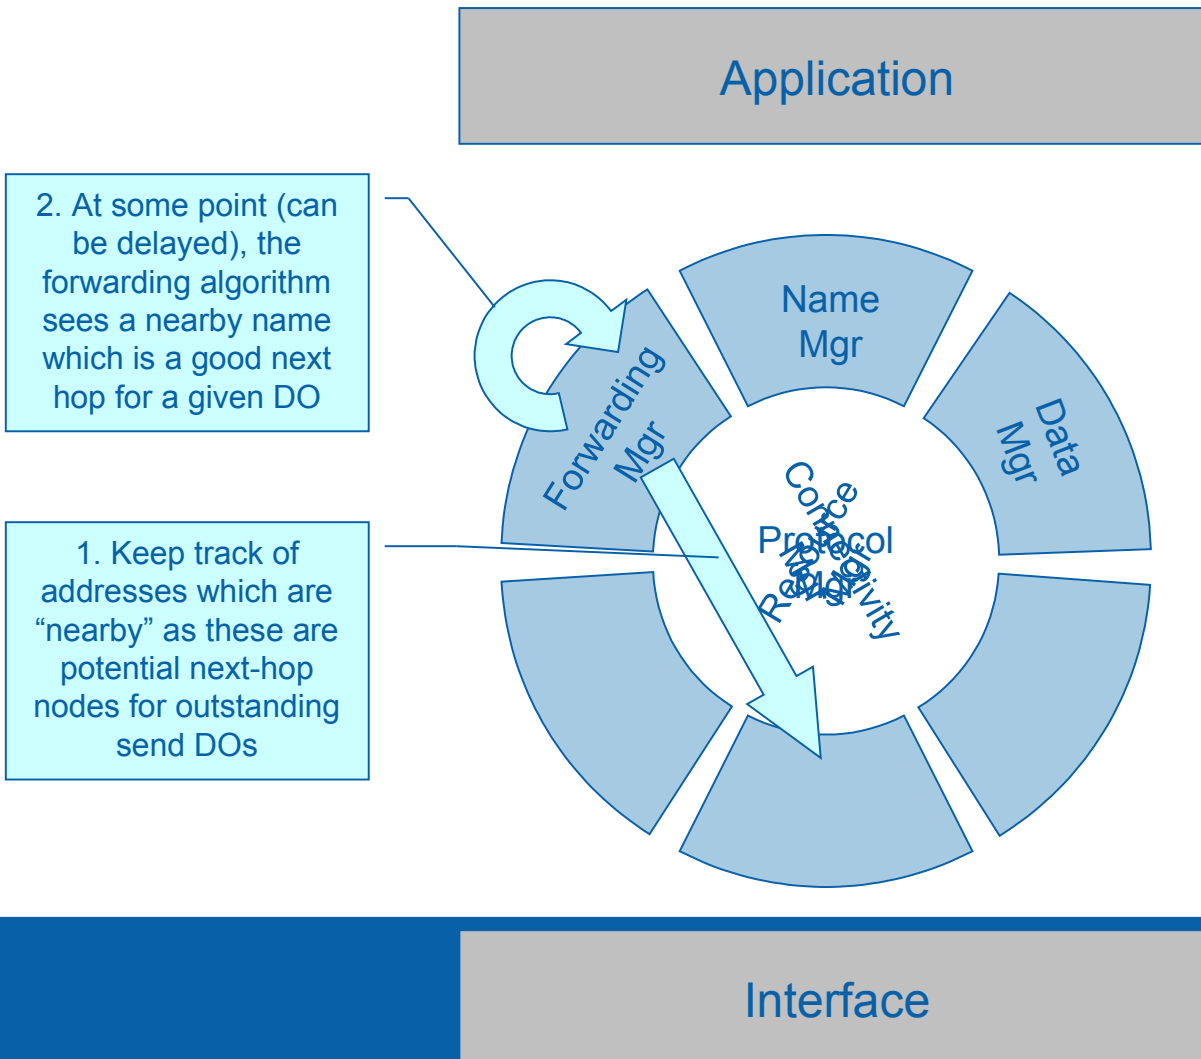
Application



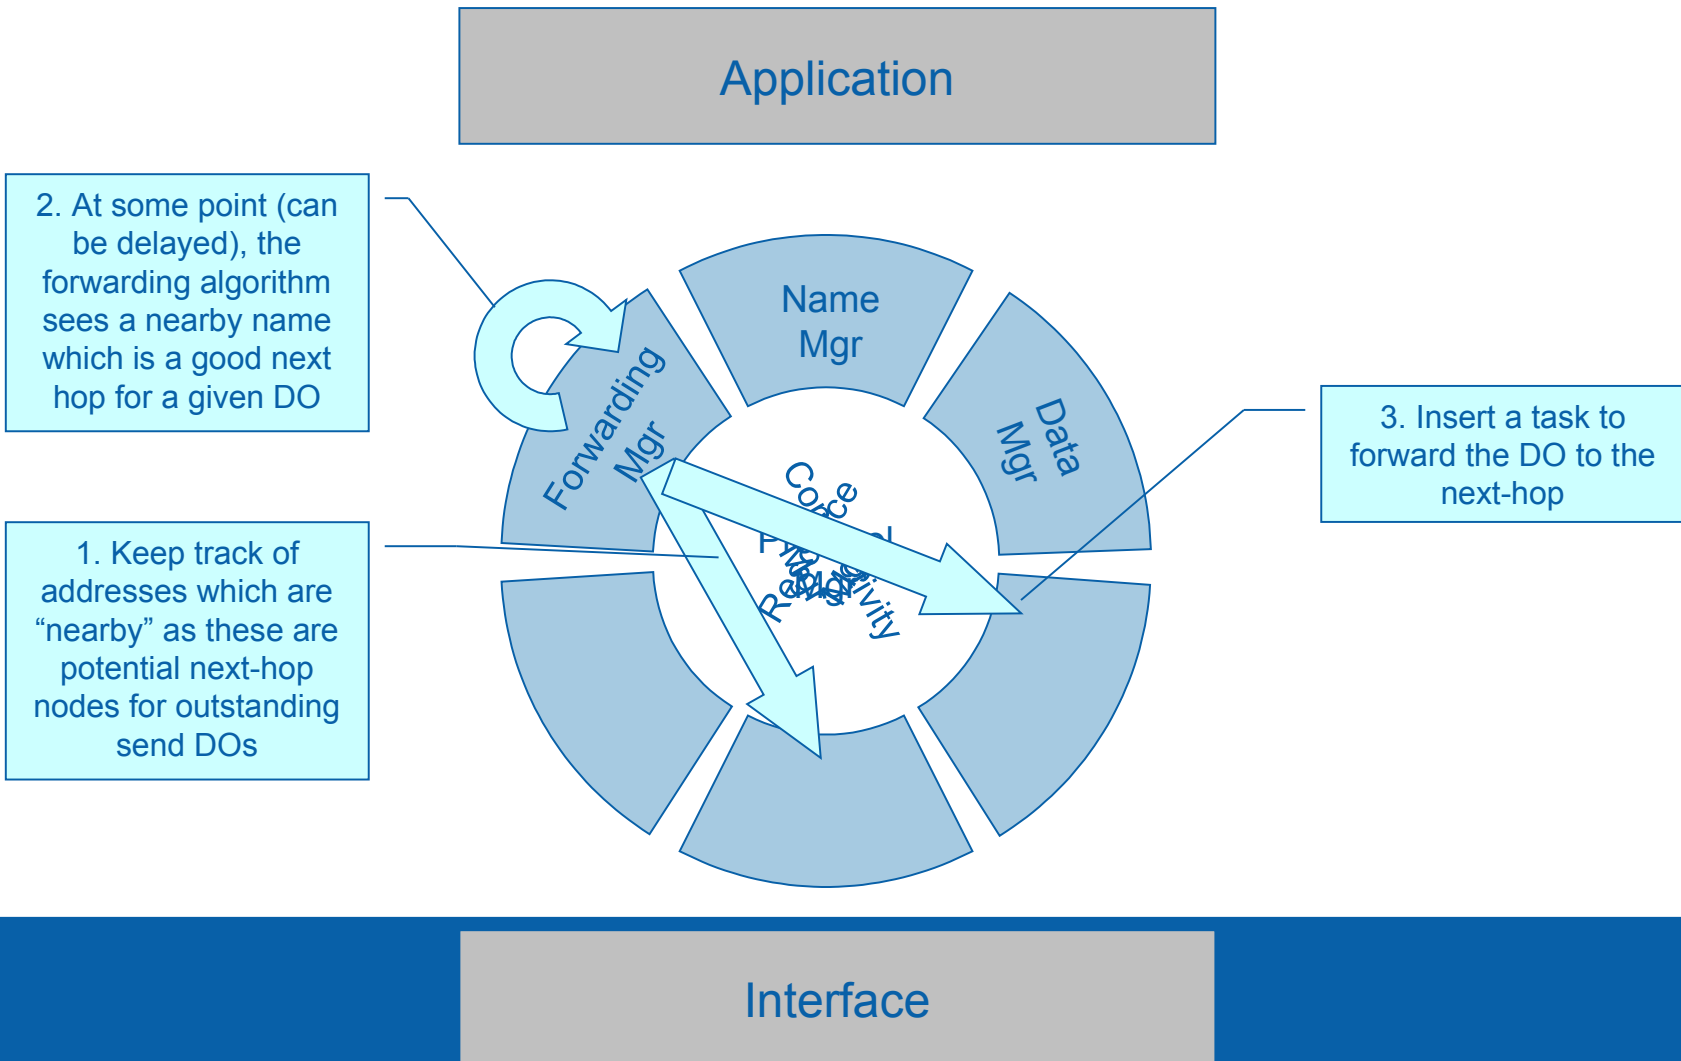
1. Keep track of addresses which are "nearby" as these are potential next-hop nodes for outstanding send DOs

Interface

Next-Hop decisions




Next-Hop decisions



Send DO (Network tx)

Application




1. If task's costs are acceptable, and cost/benefit shows this is the most worthwhile task, respond with a "go"

Interface

Send DO (Network tx)

Application

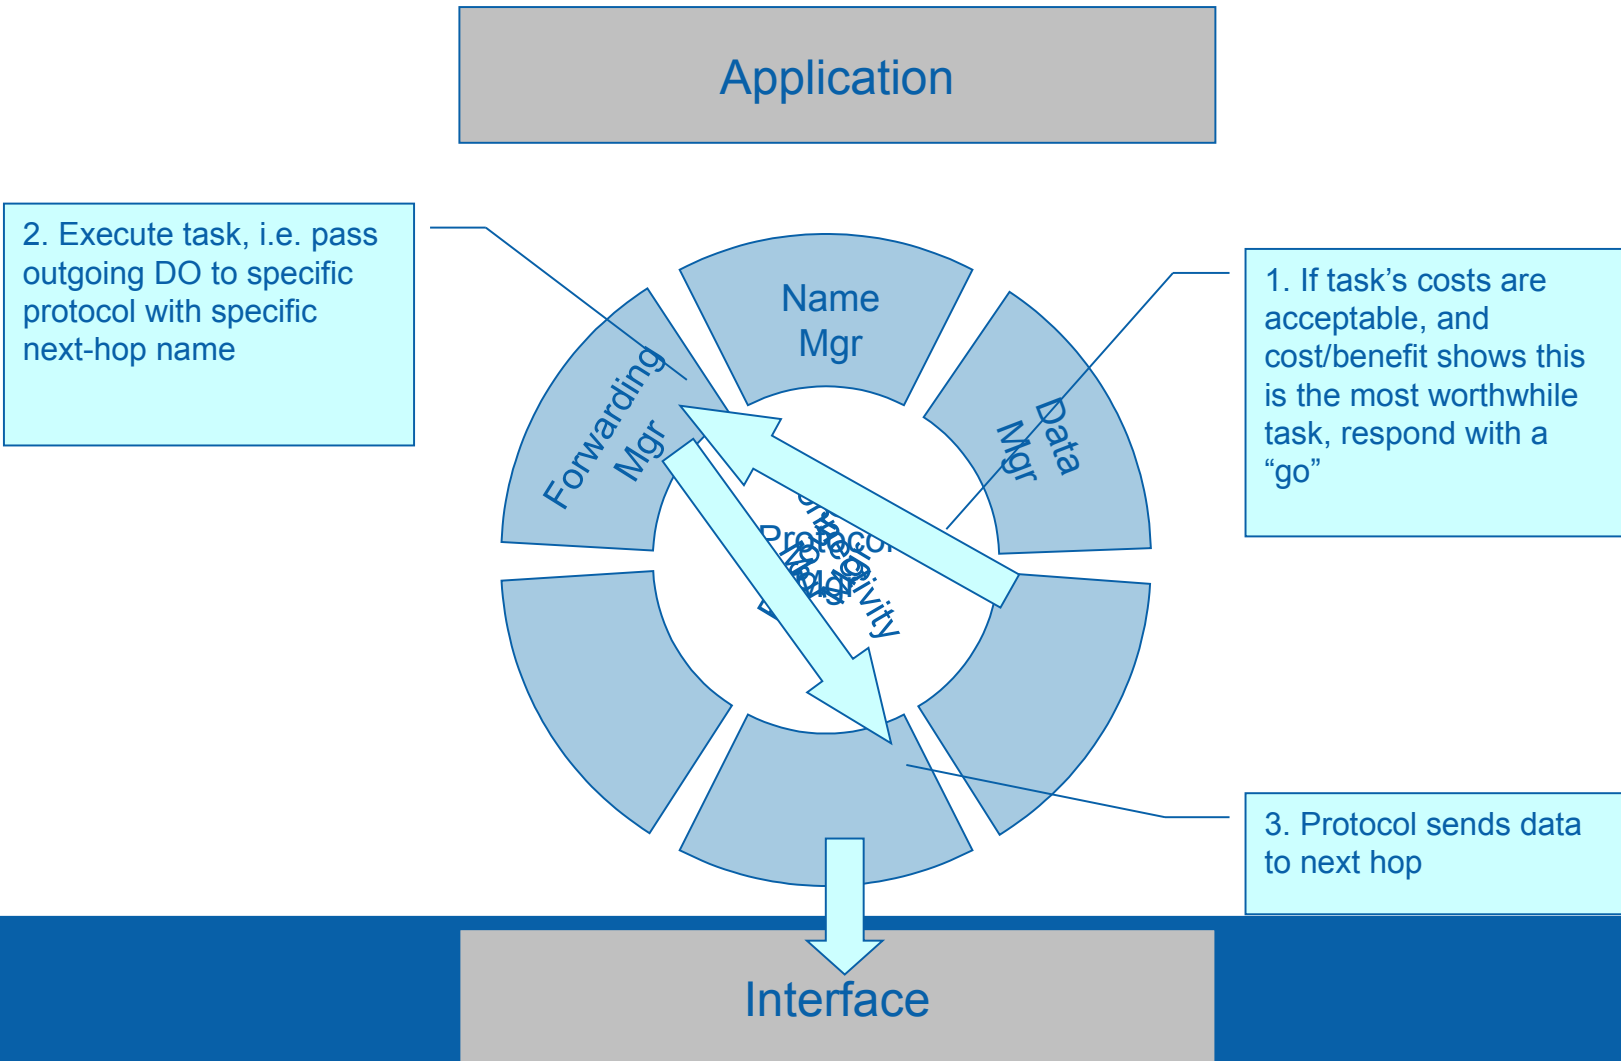
2. Execute task, i.e. pass outgoing DO to specific protocol with specific next-hop name



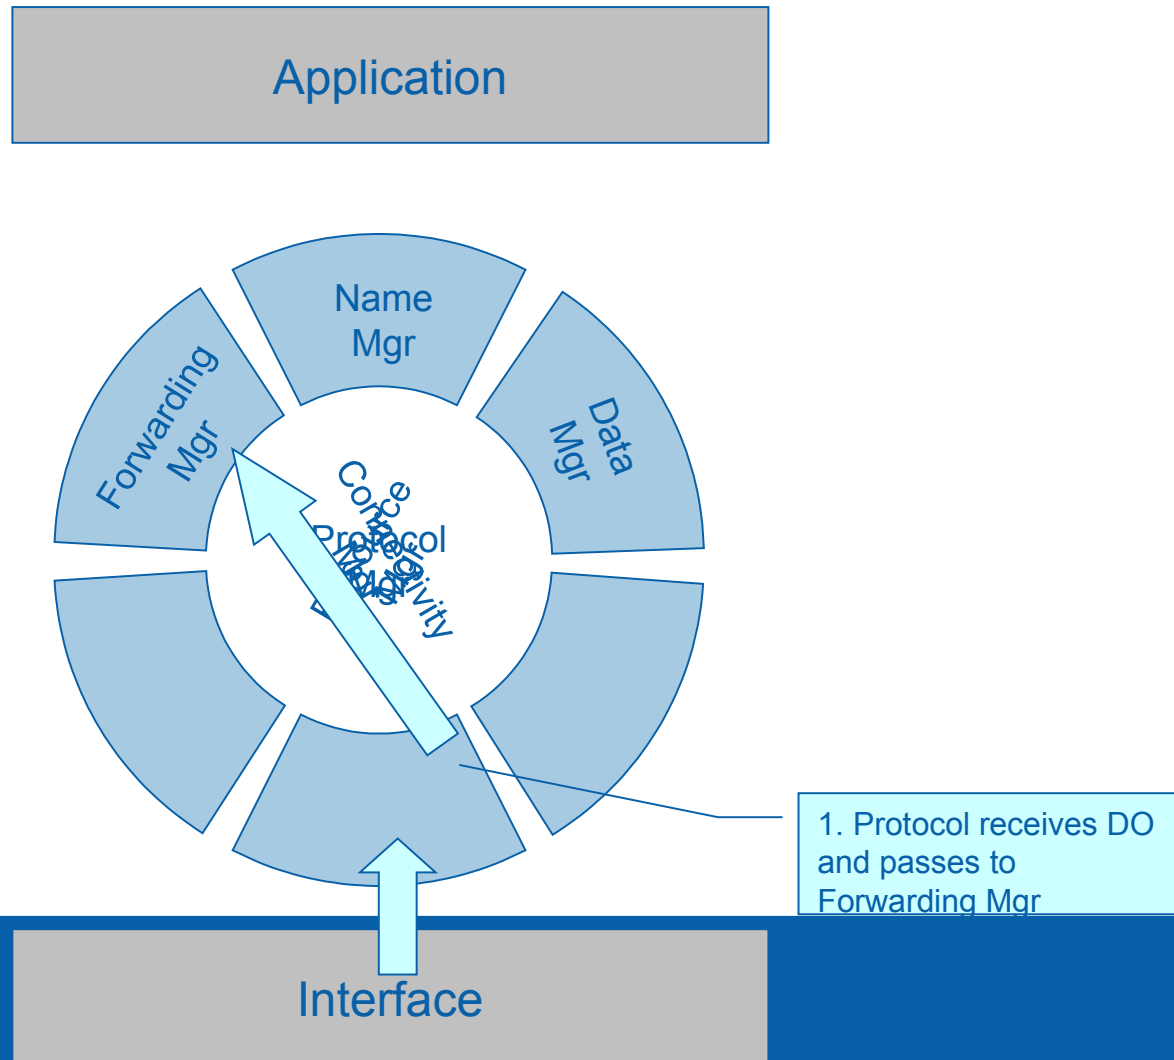
1. If task's costs are acceptable, and cost/benefit shows this is the most worthwhile task, respond with a "go"

Interface

Send DO (Network tx)



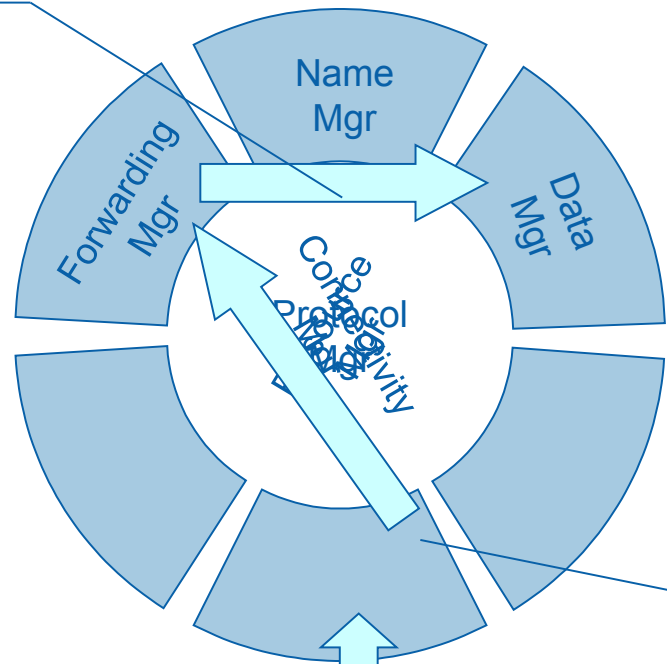
Receiving DO



Receiving DO

Application

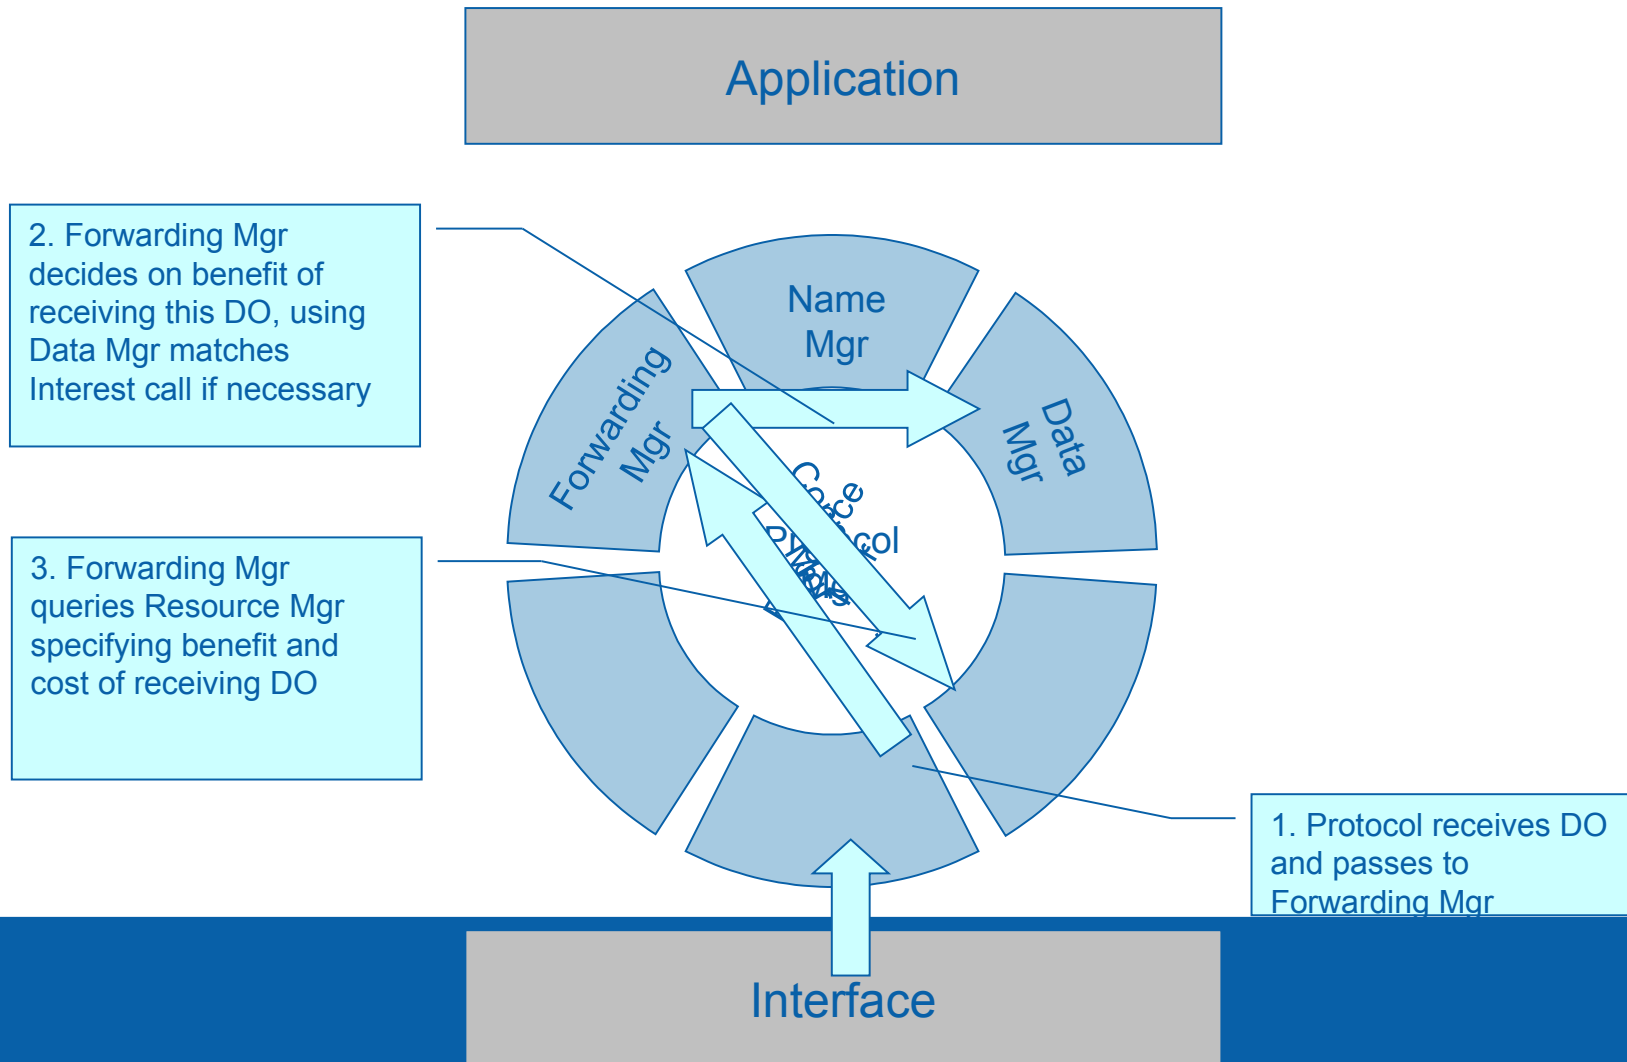
2. Forwarding Mgr decides on benefit of receiving this DO, using Data Mgr matches Interest call if necessary



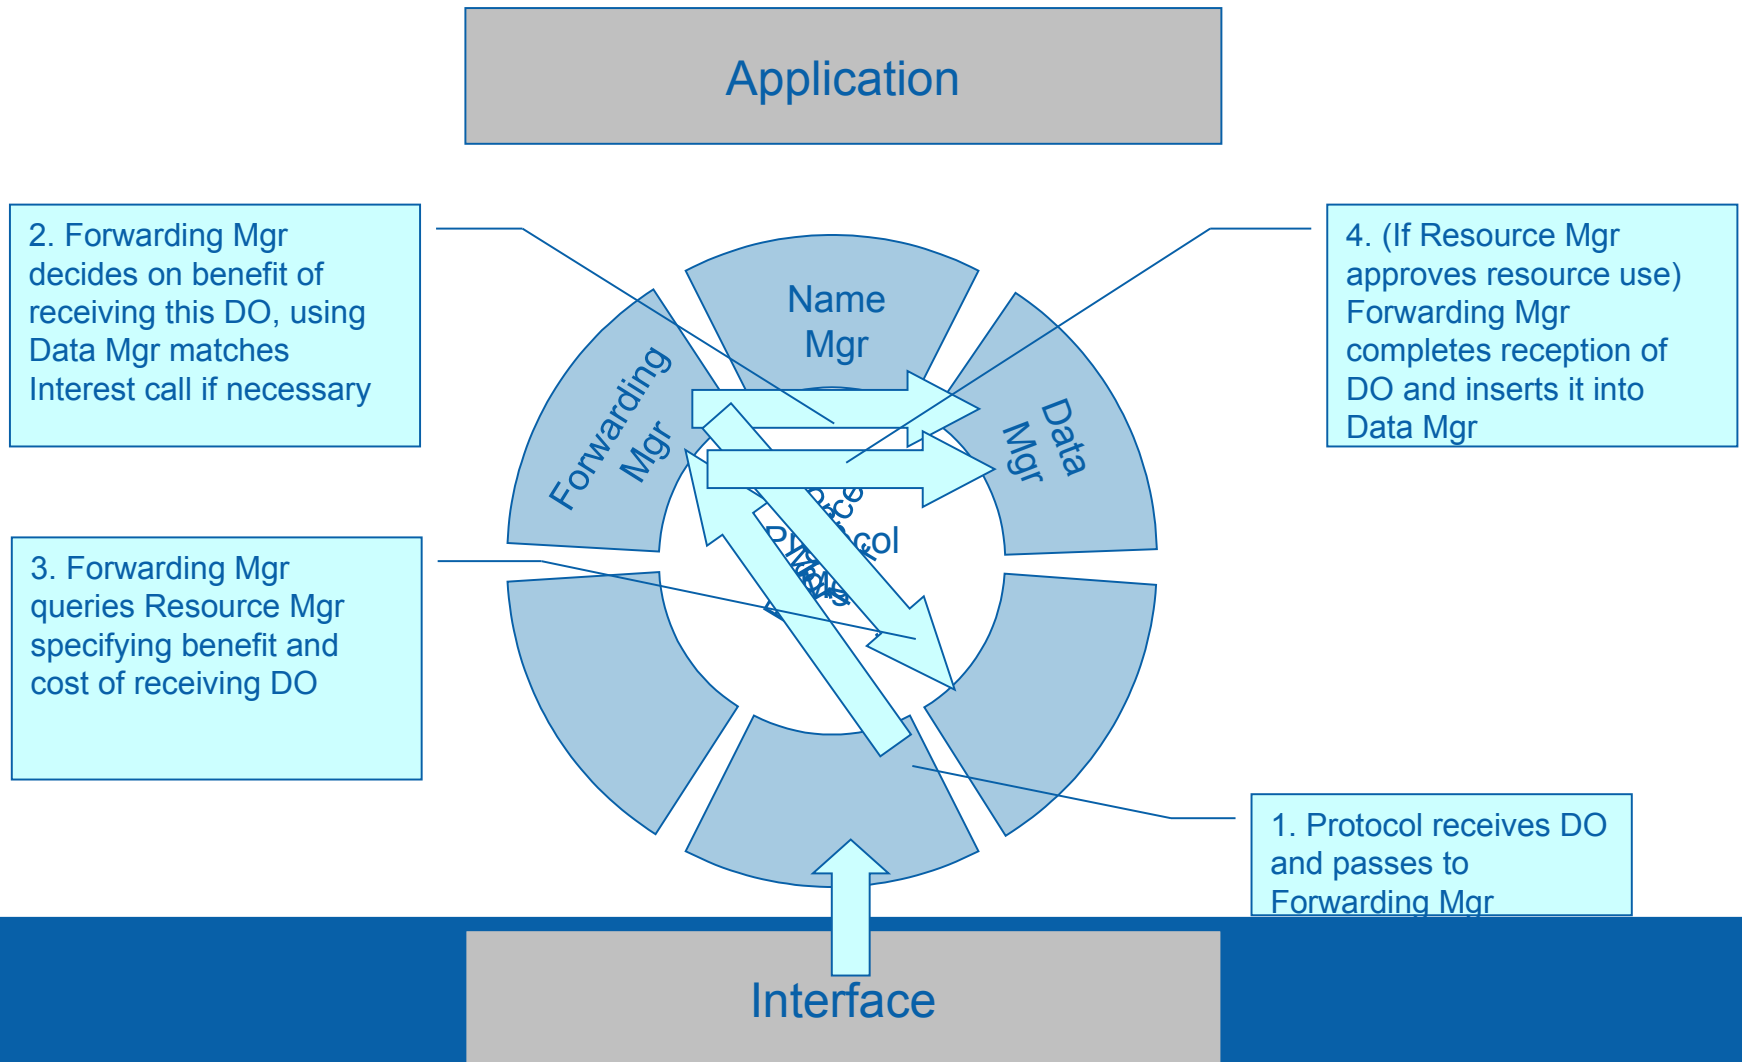
1. Protocol receives DO and passes to Forwarding Mgr

Interface

Receiving DO



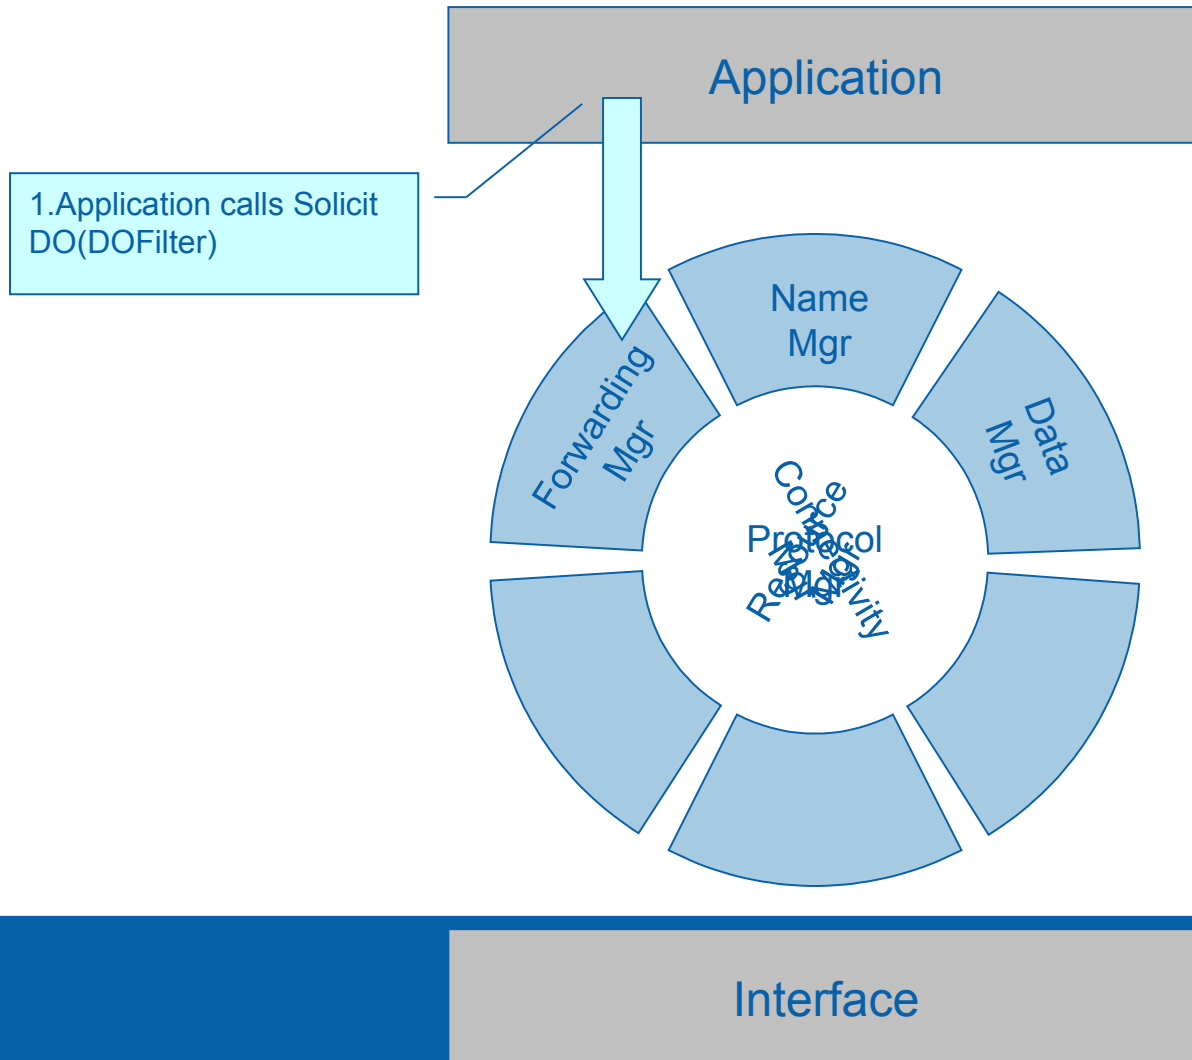
Receiving DO



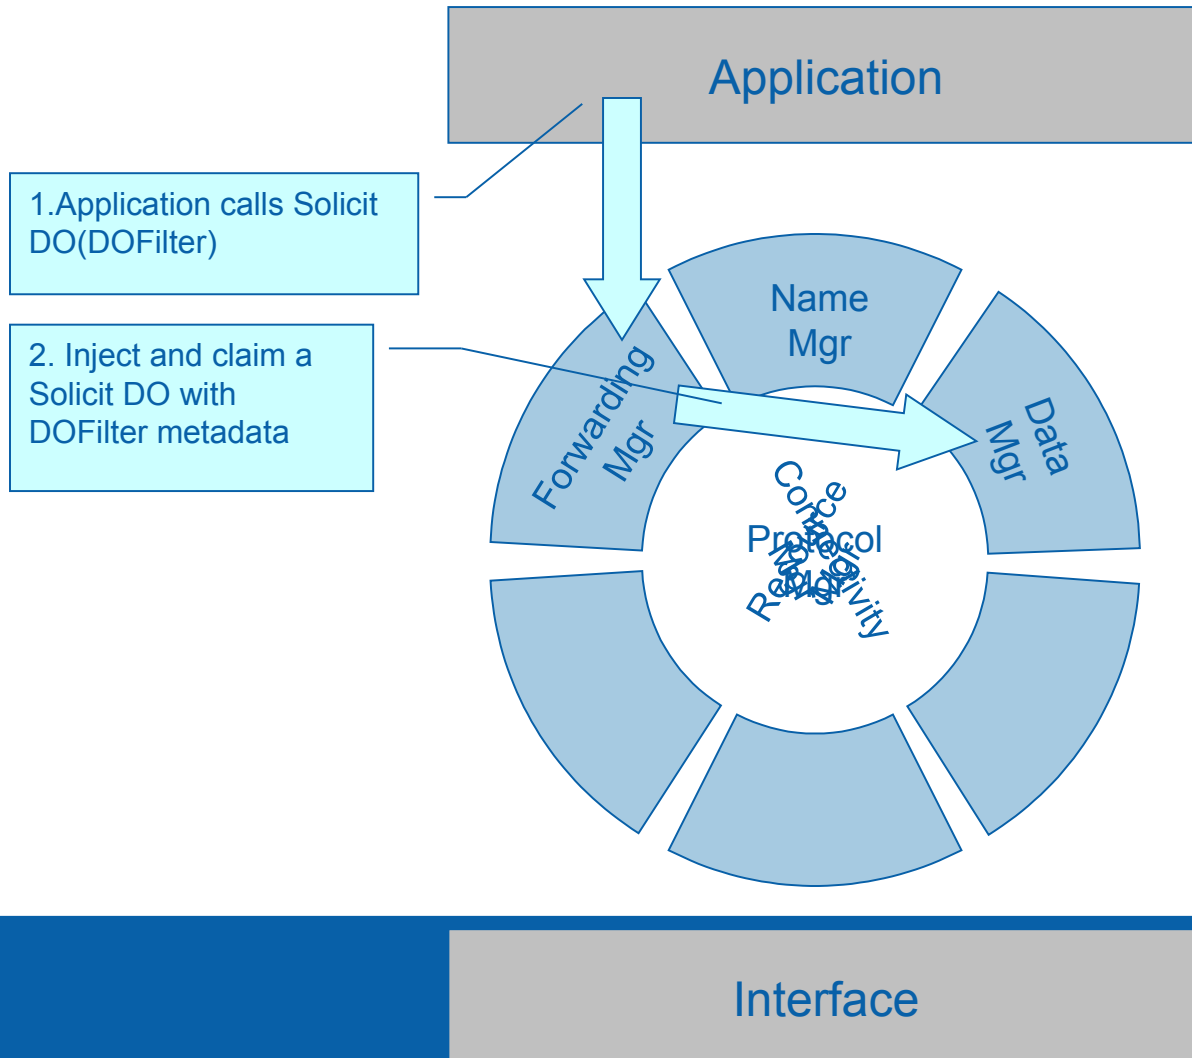
Outline

- Data Object (DO) app interface
- Names and their use in delivery
- Neighbour discovery
- Send DO / DO reception
- Solicit DO

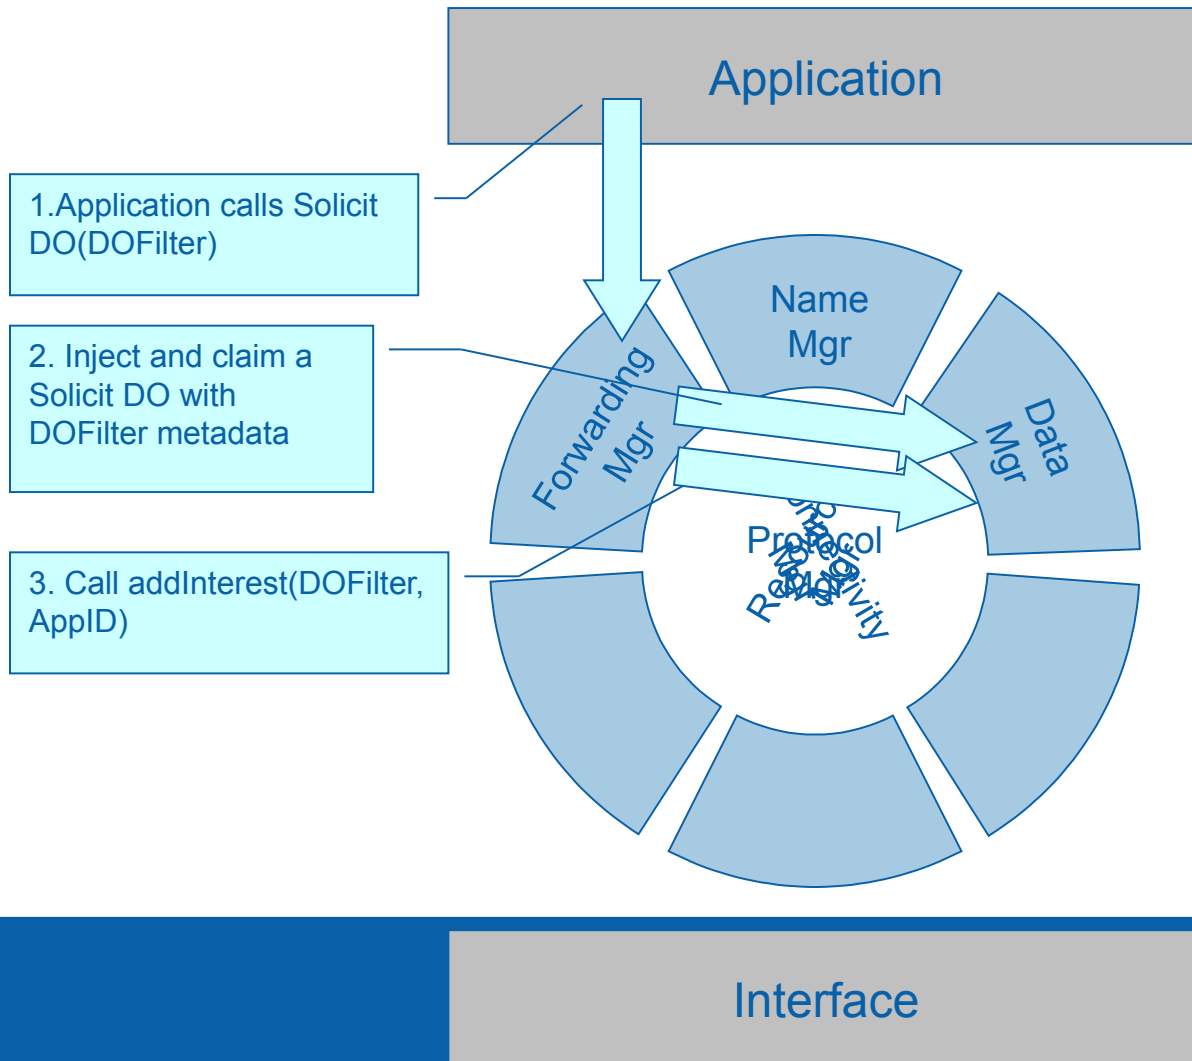
Solicit DO (App insertion)



Solicit DO (App insertion)

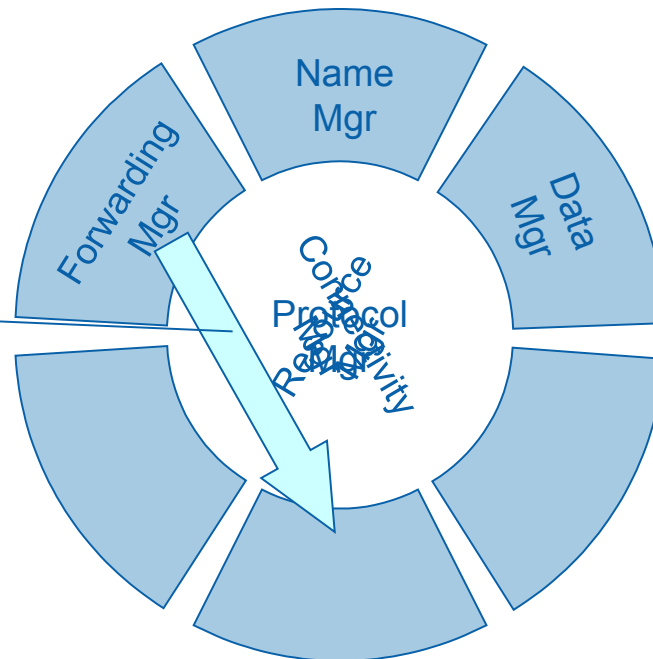


Solicit DO (App insertion)



Solicit DO (Propagate)

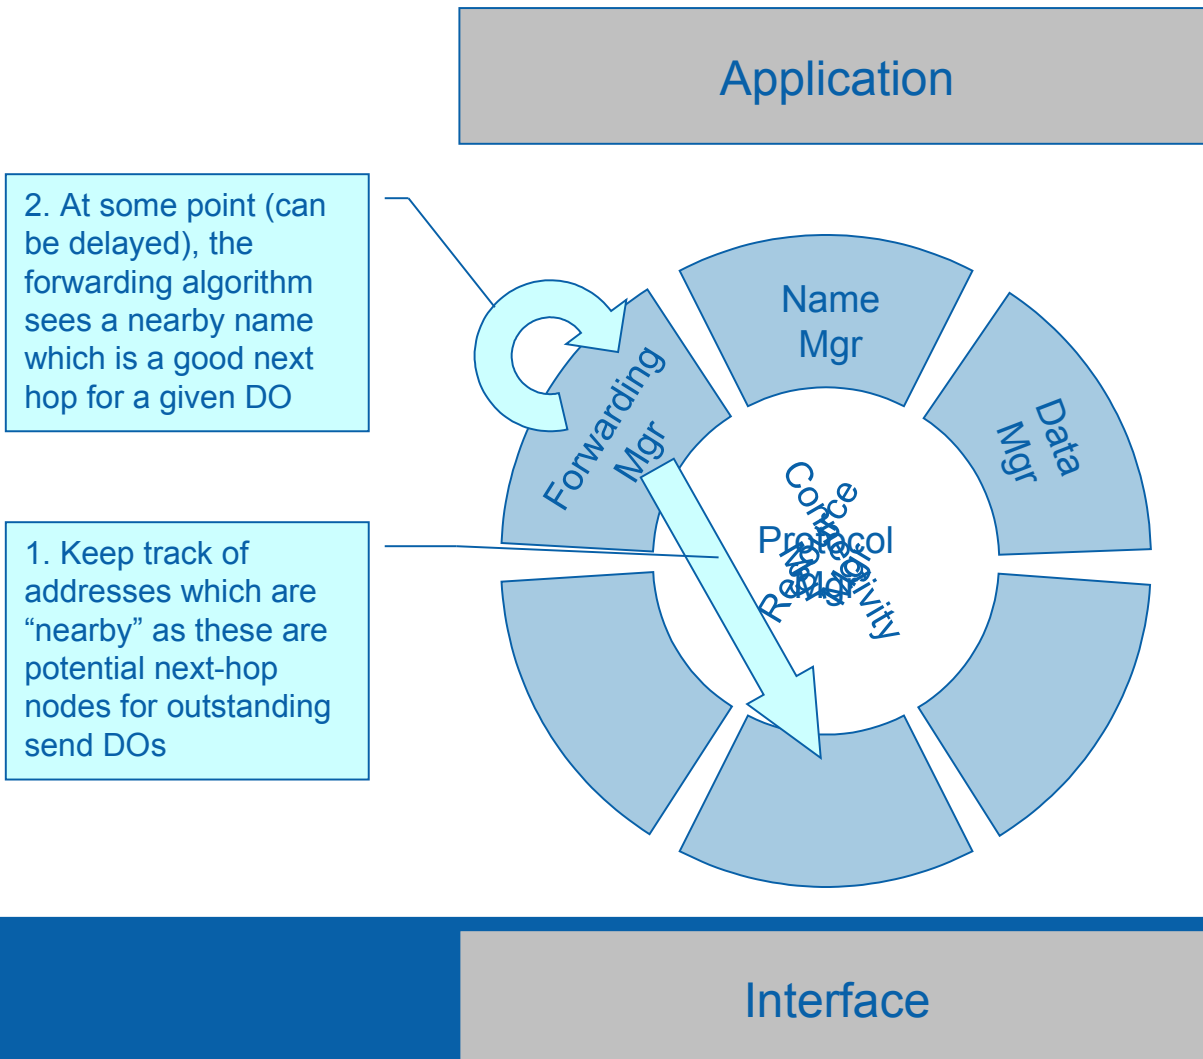
Application



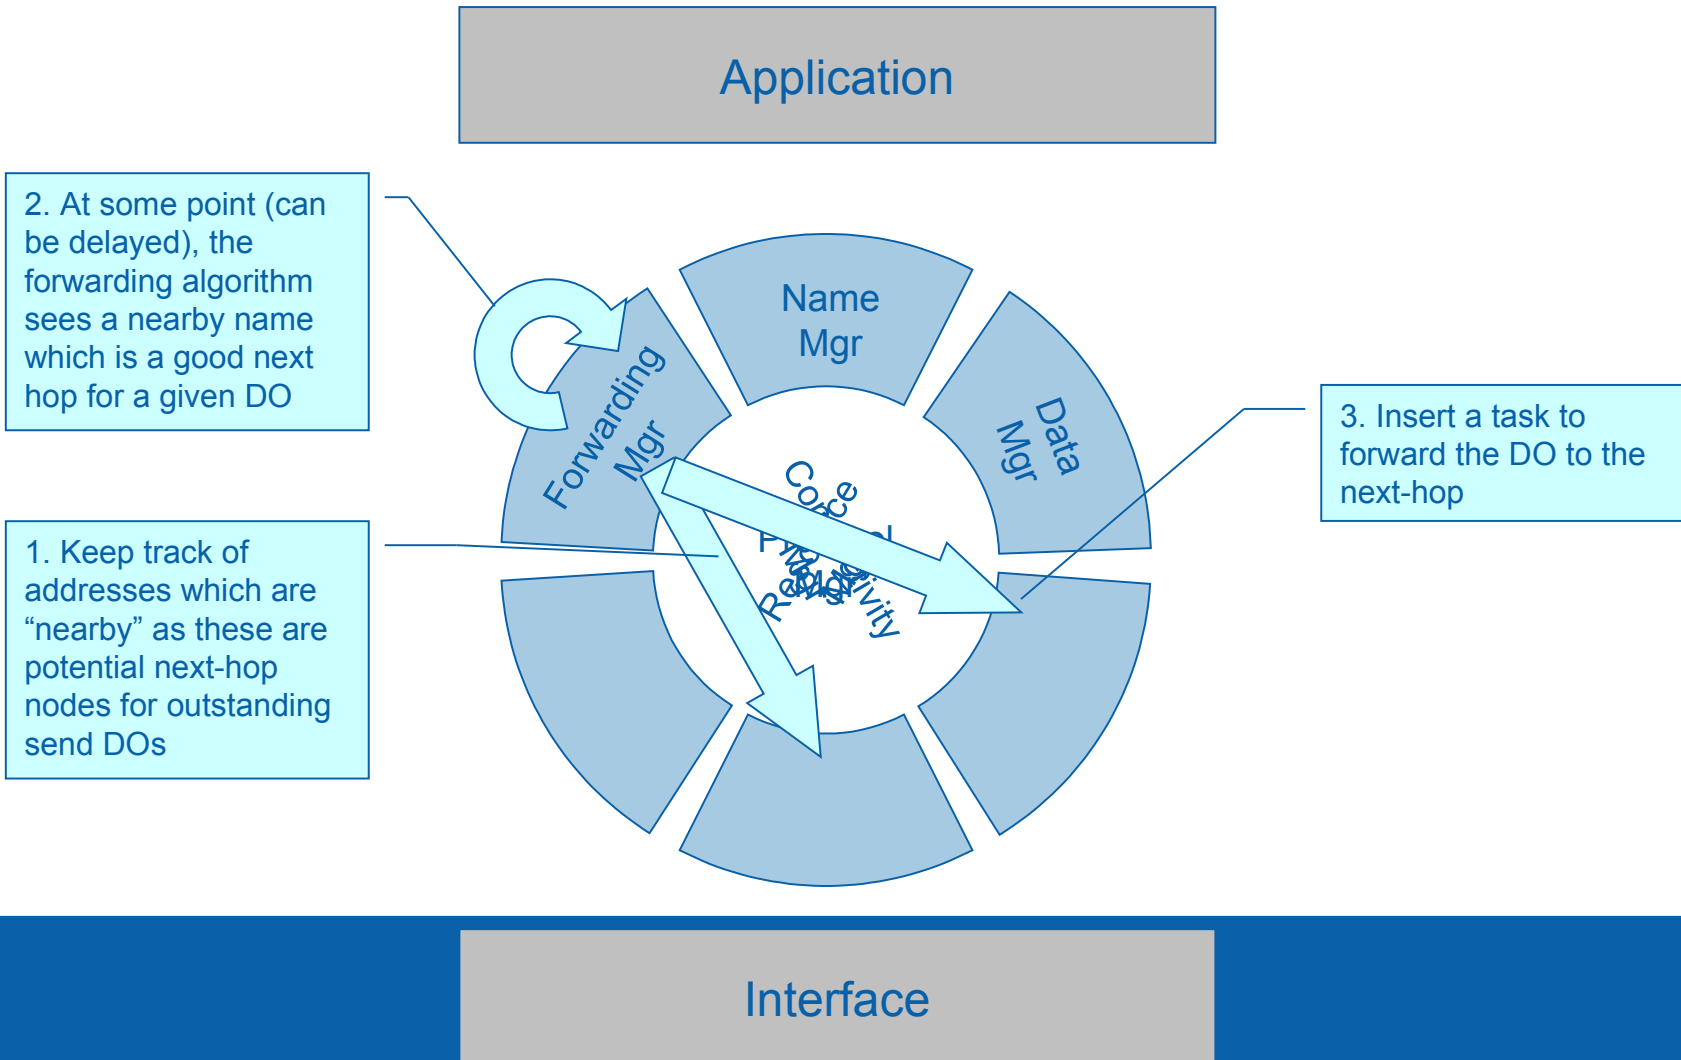
1. Keep track of addresses which are "nearby" as these are potential next-hop nodes for outstanding send DOs

Interface

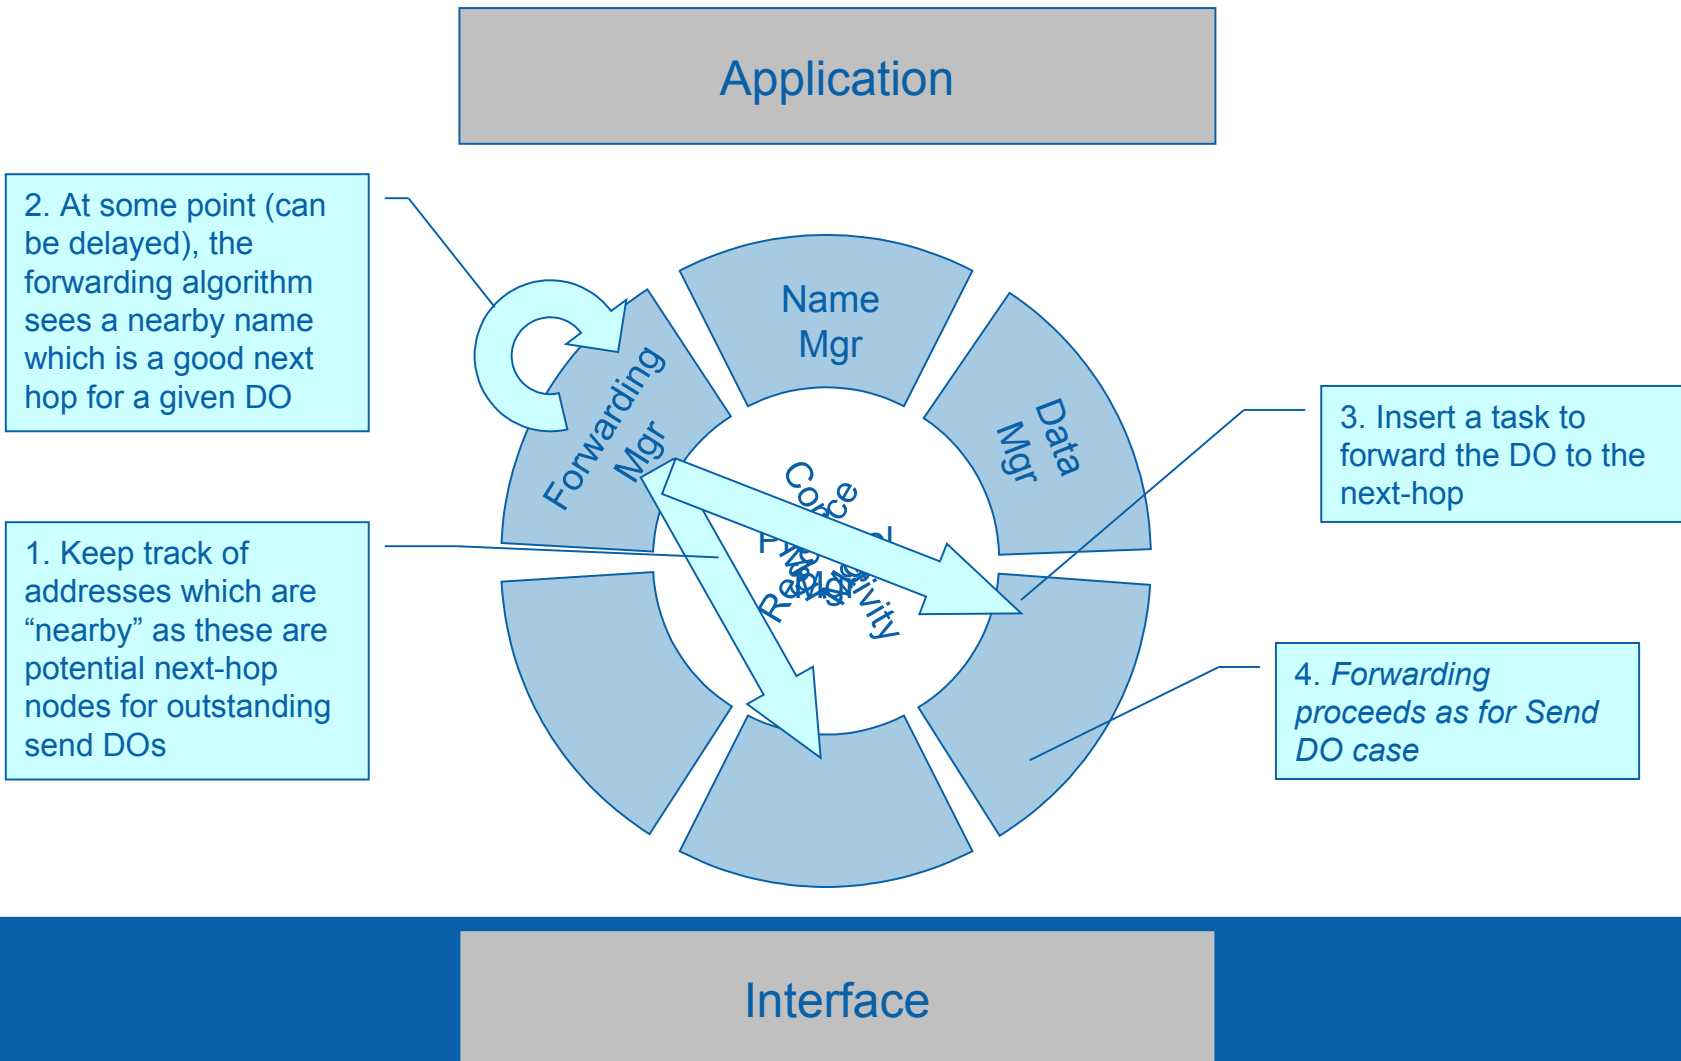
Solicit DO (Propagate)





Solicit DO (Propagate)



Solicit DO (Propagate)



Solicit DO (Network arrival)

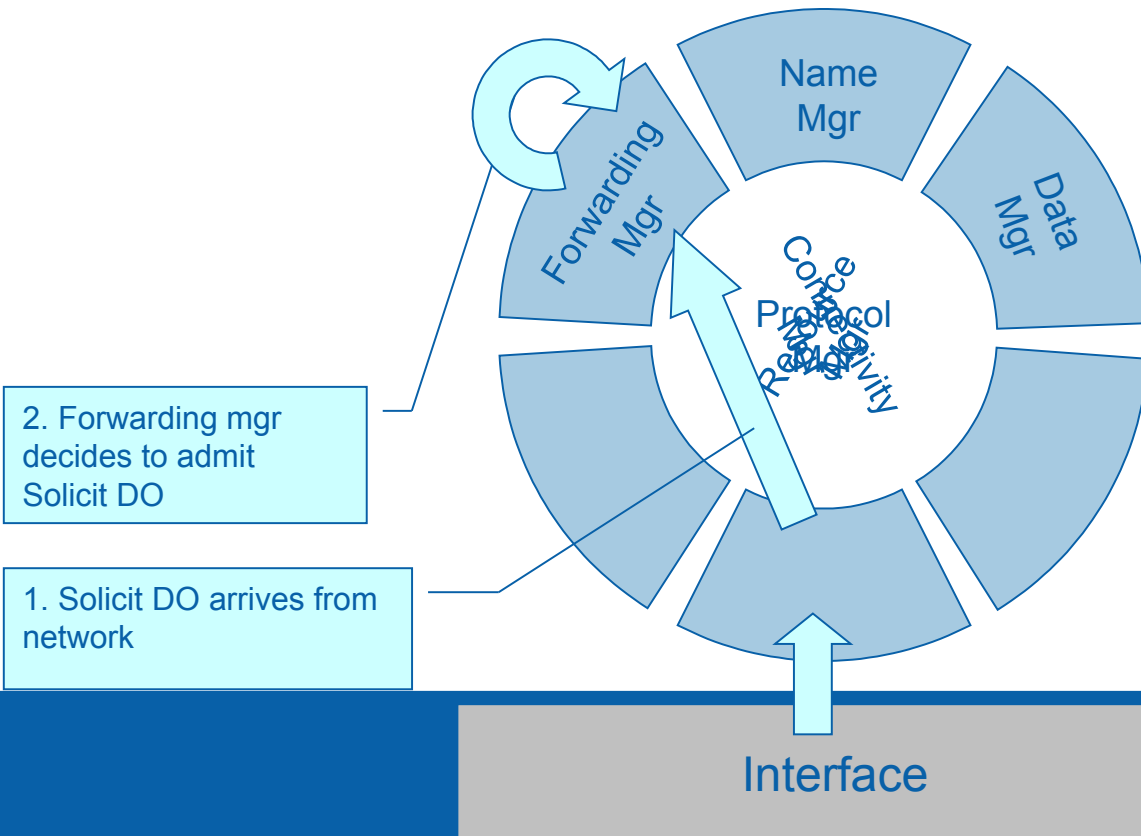


1. Solicit DO arrives from network



Solicit DO (Network arrival)

Application

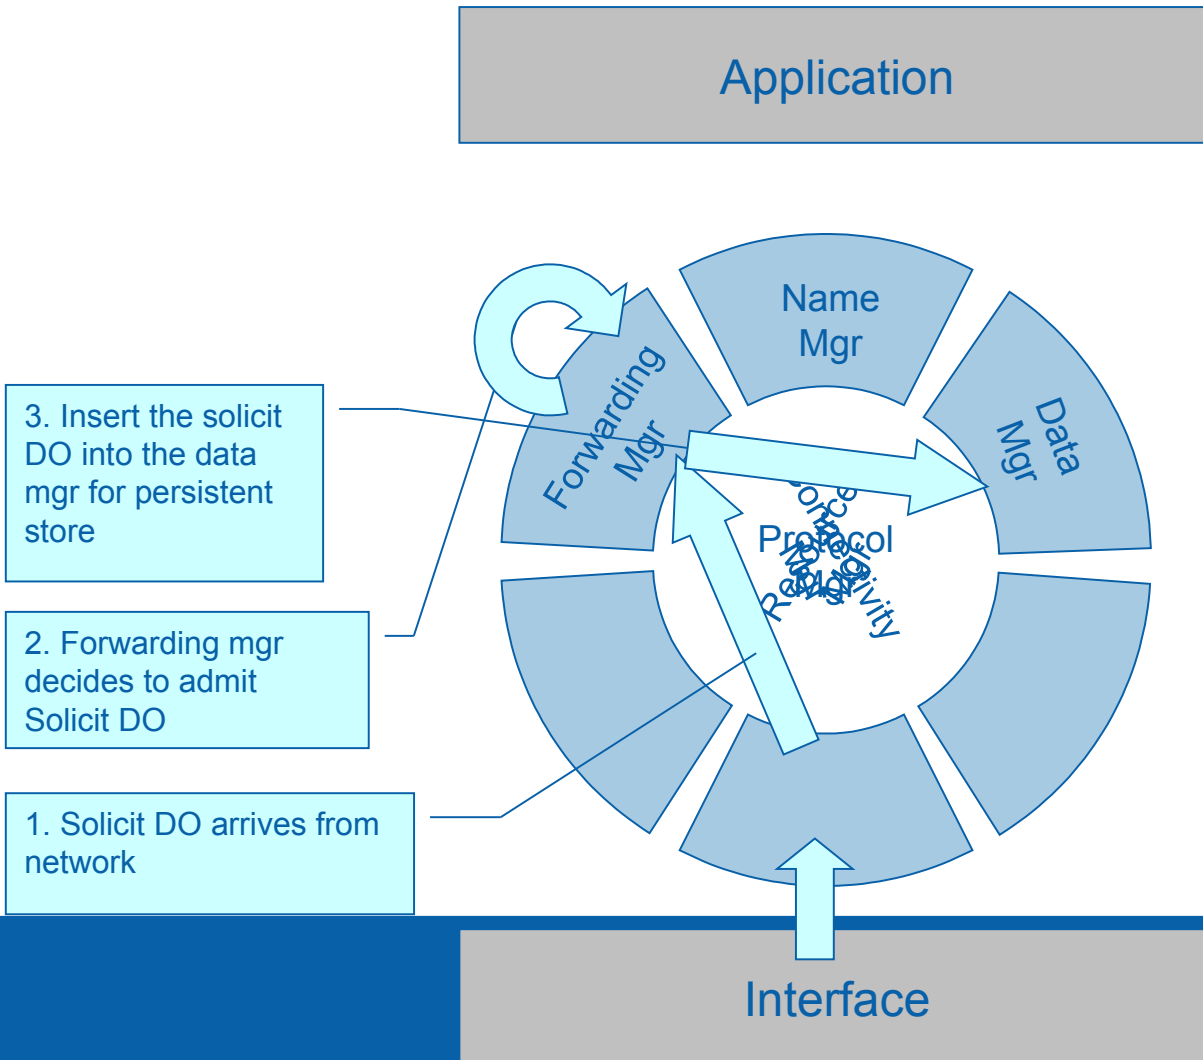


2. Forwarding mgr decides to admit Solicit DO

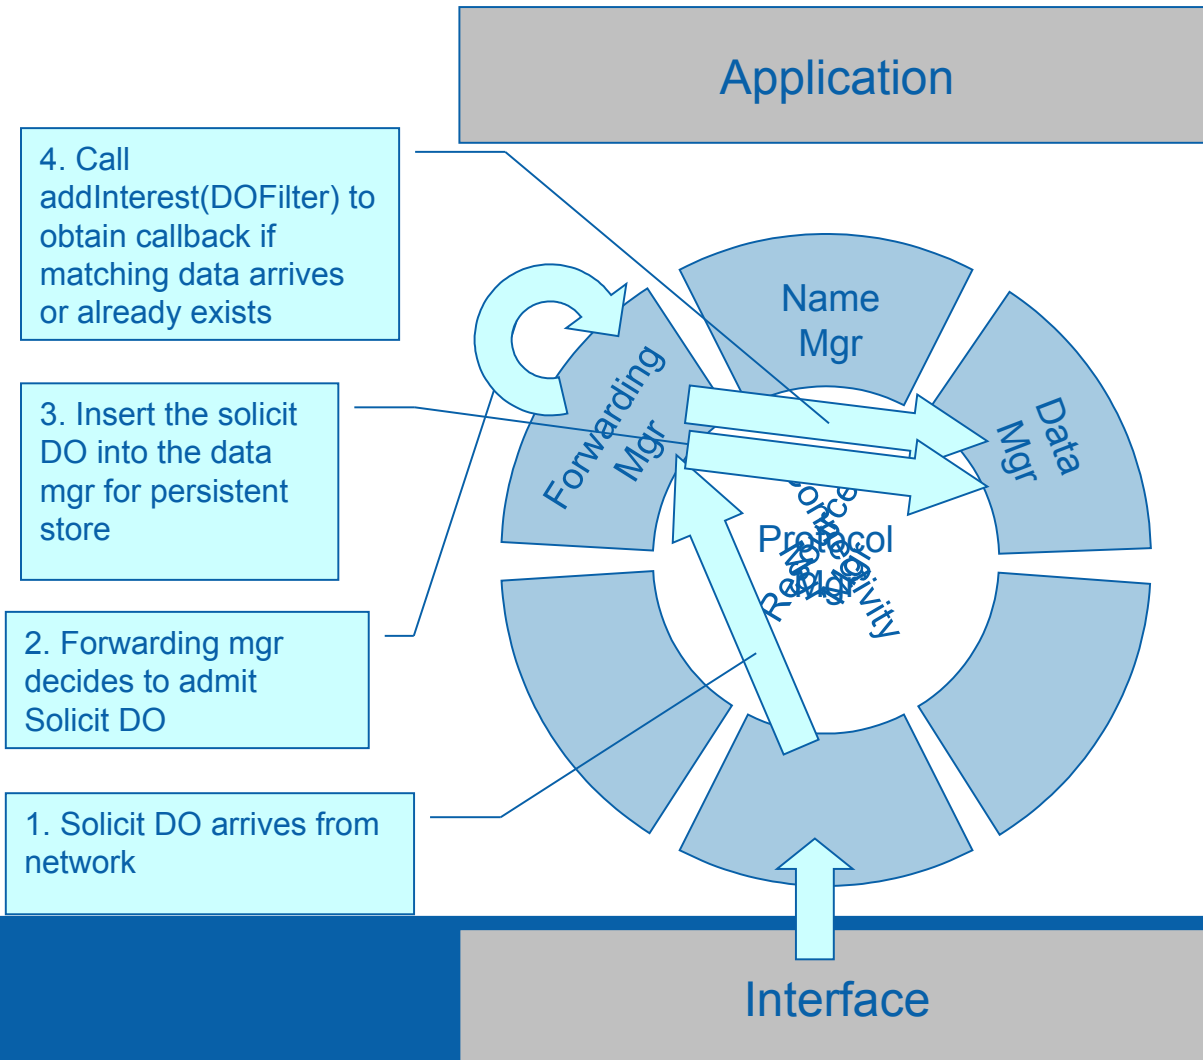
1. Solicit DO arrives from network

Interface

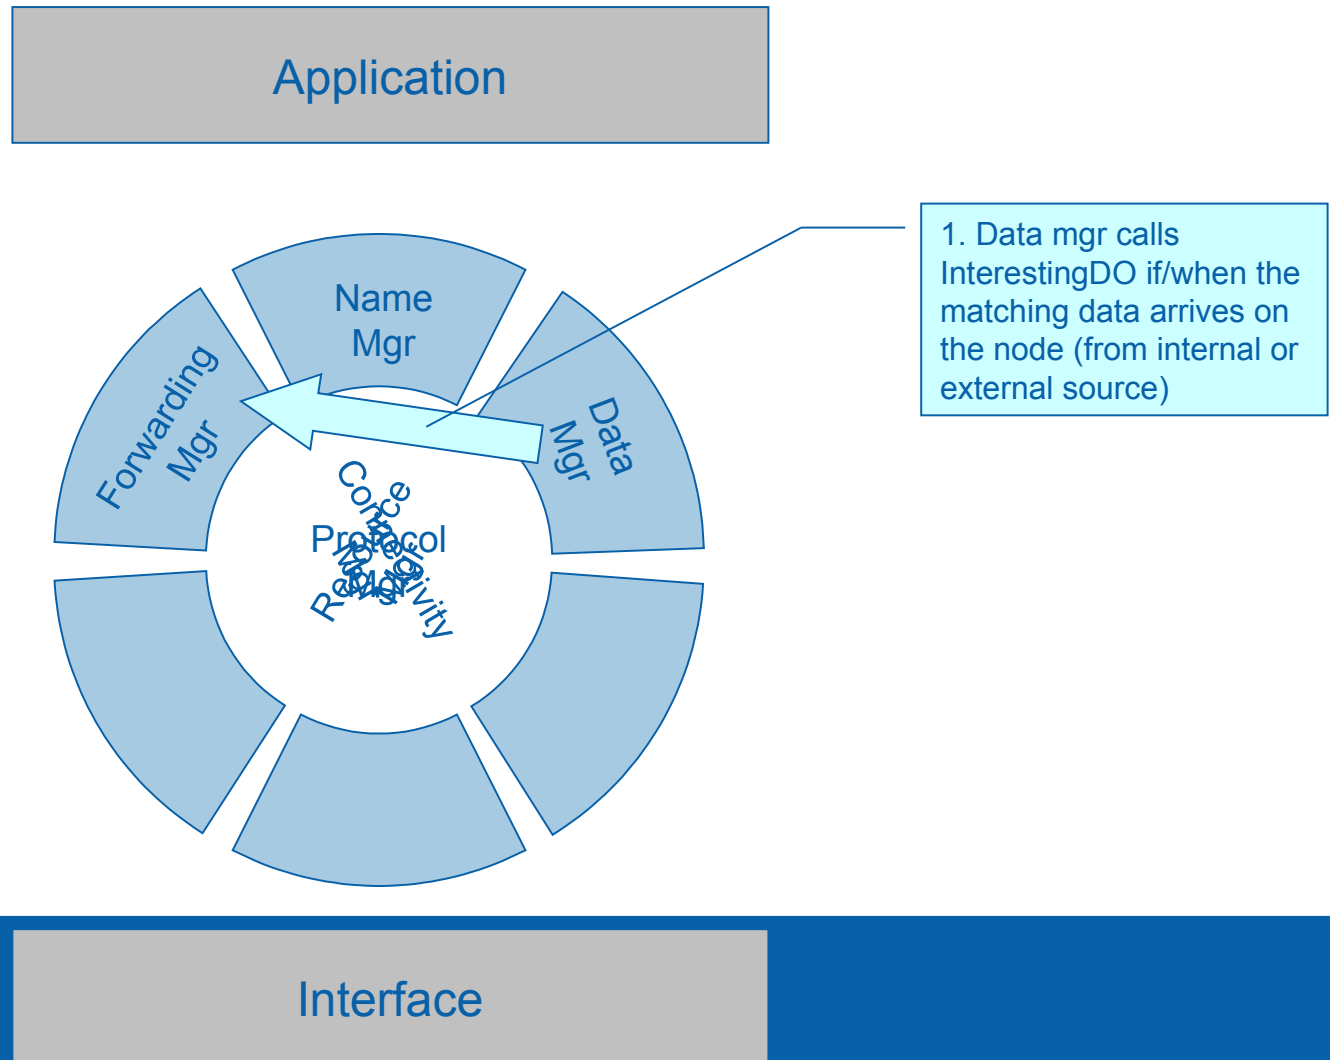
Solicit DO (Network arrival)



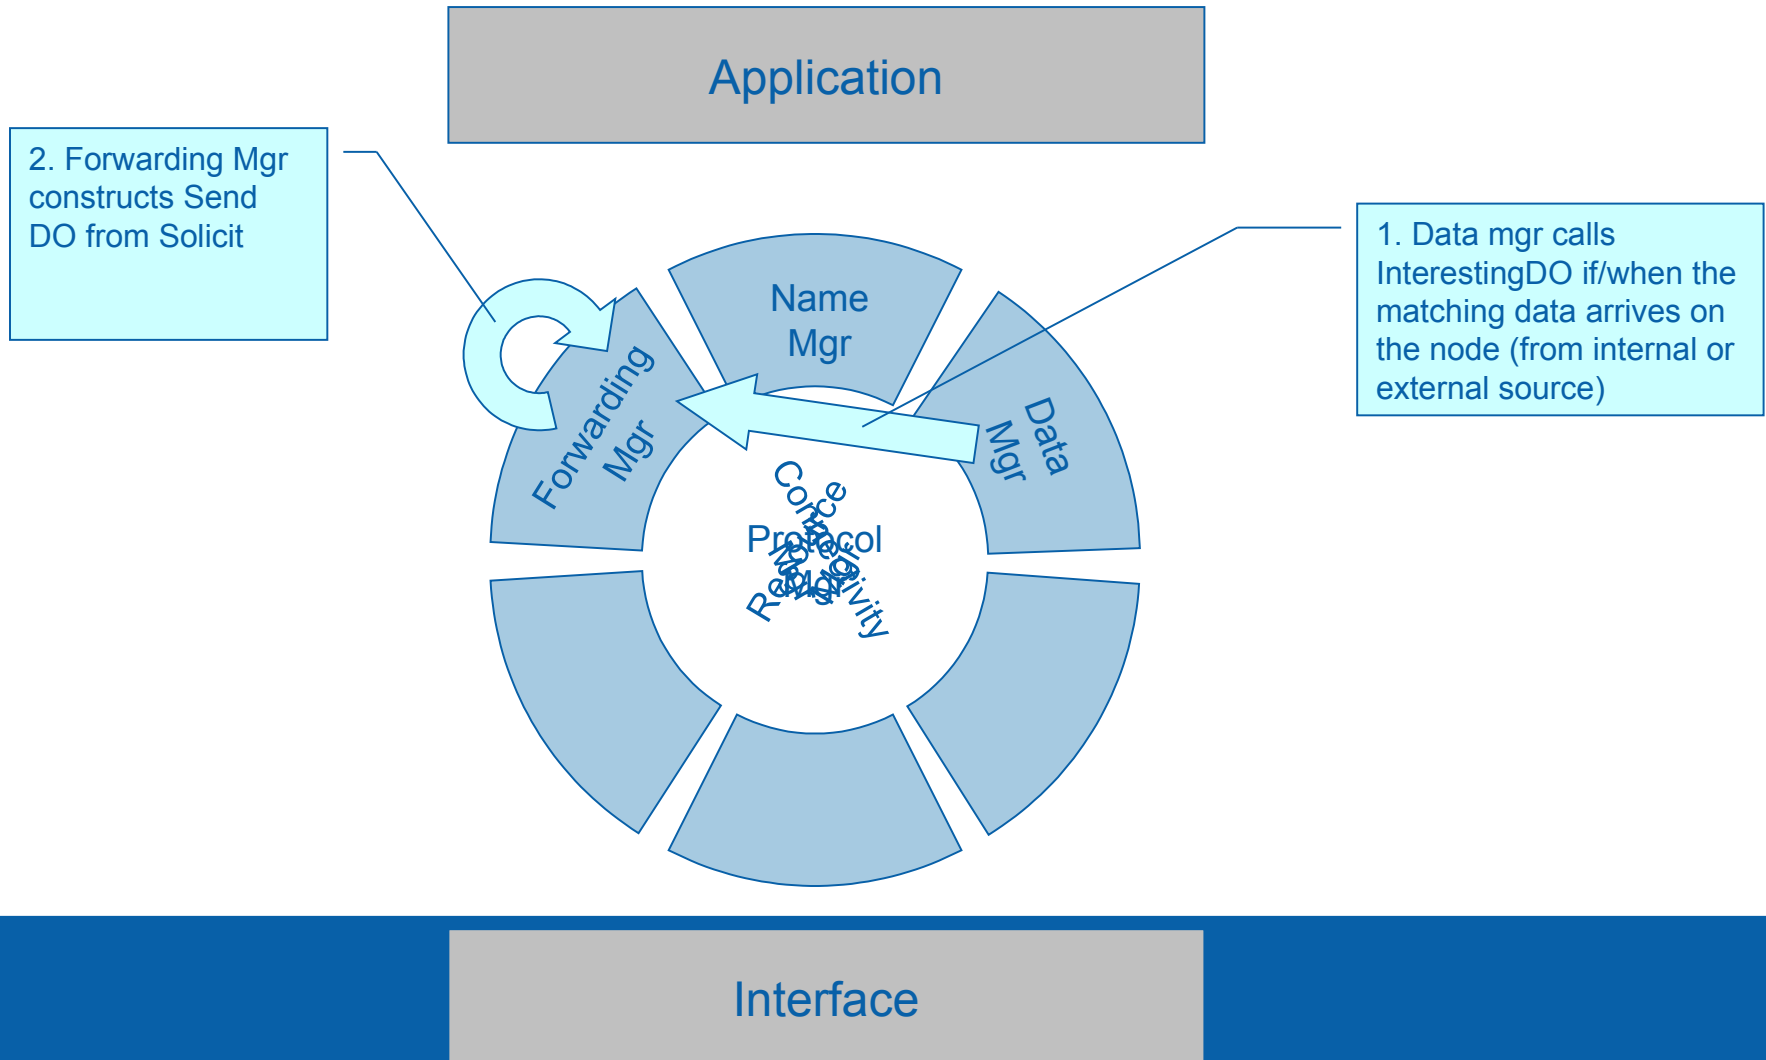
Solicit DO (Network arrival)



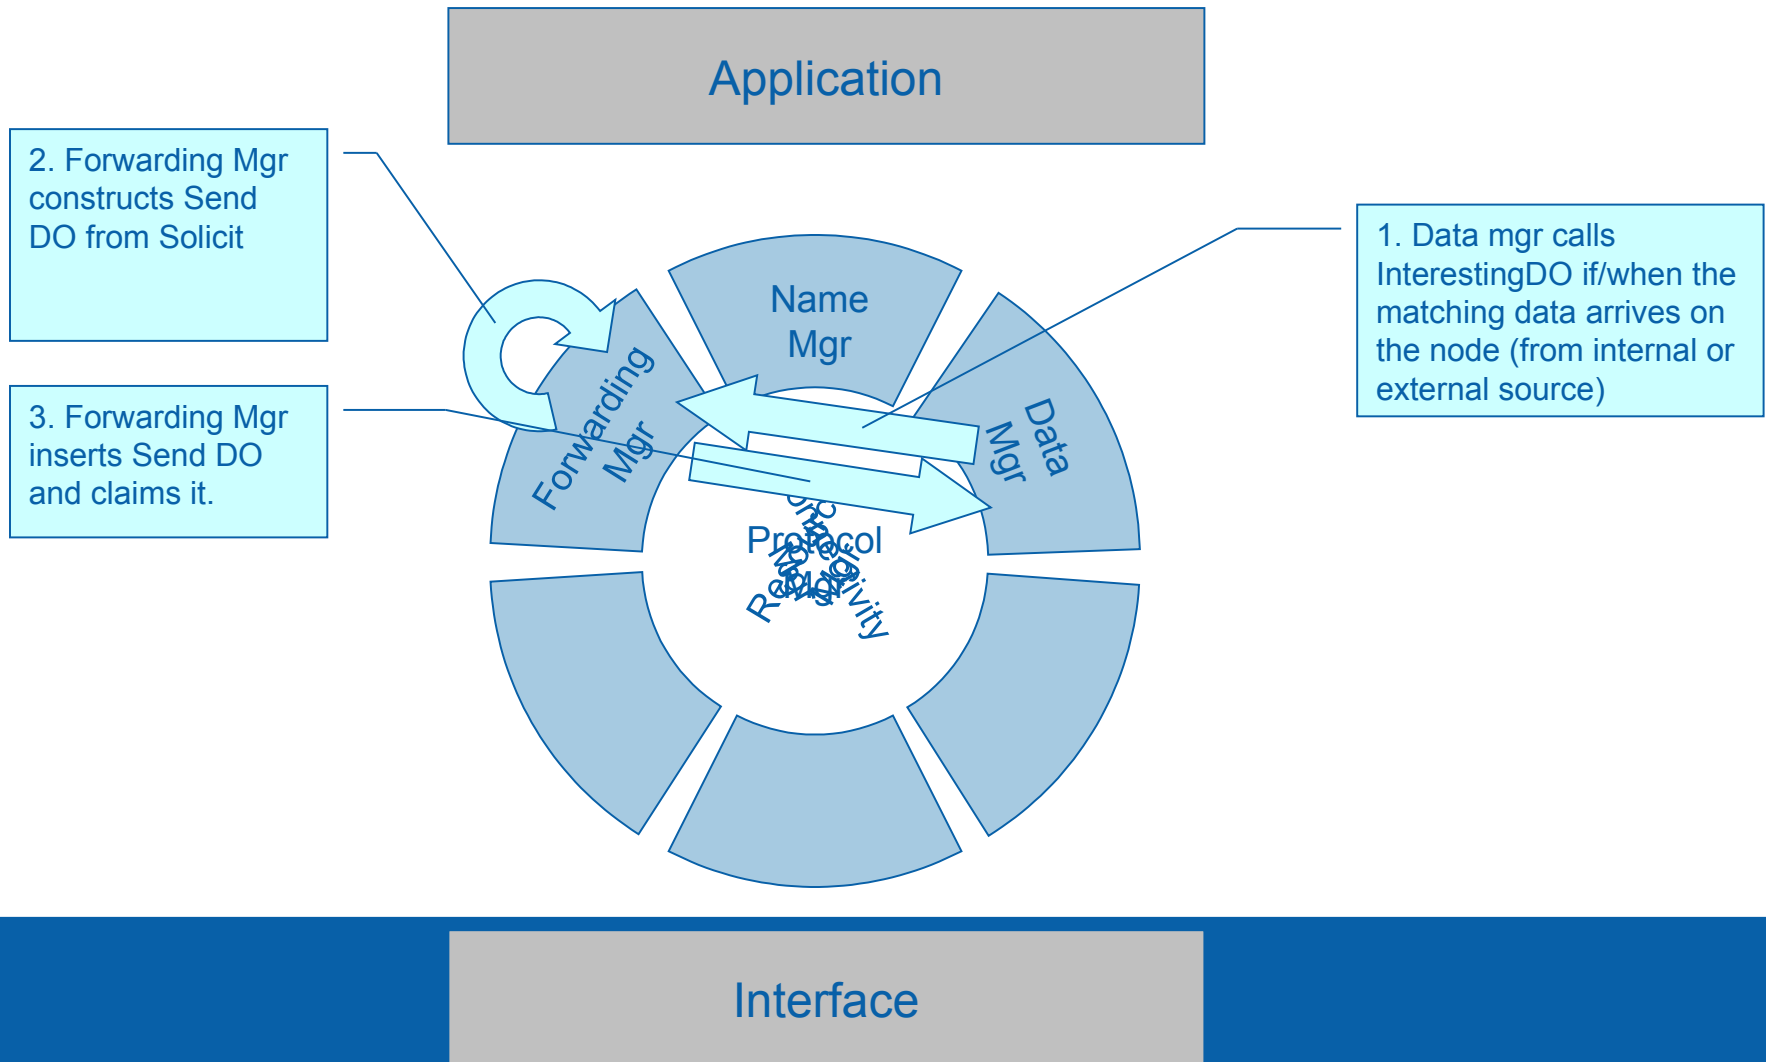
Solicit DO (Data found)



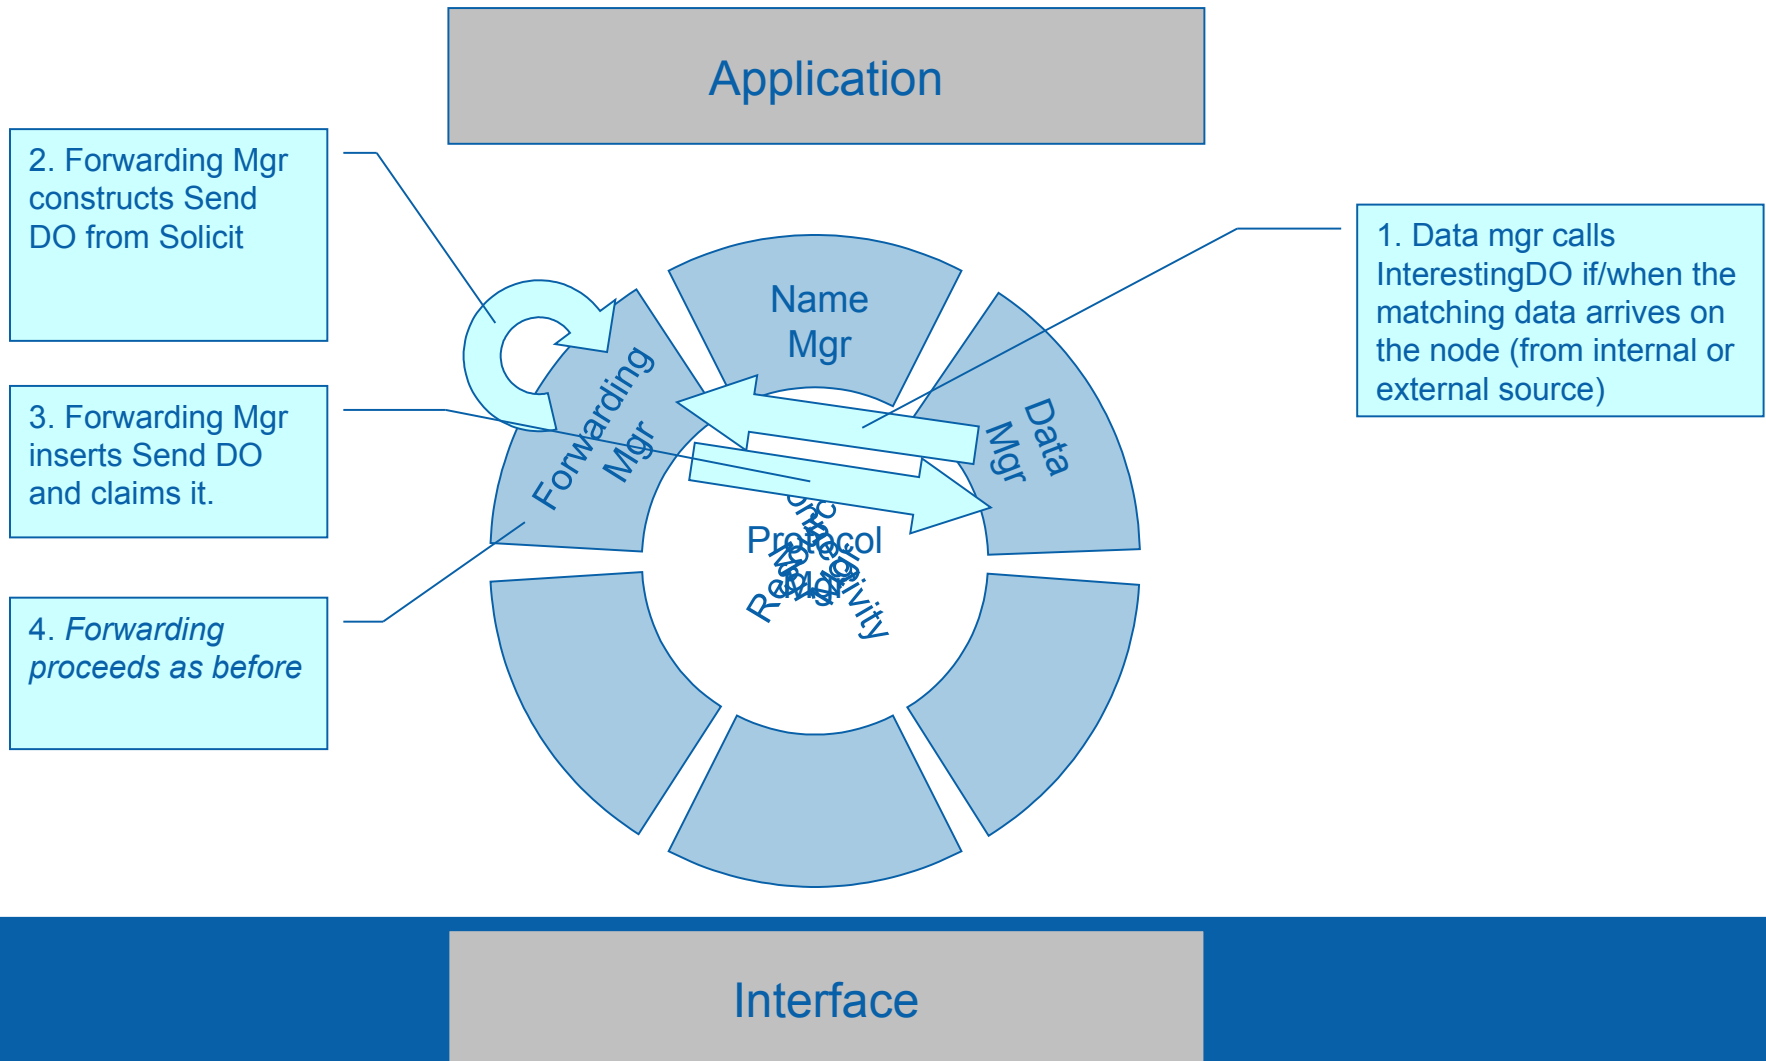
Solicit DO (Data found)



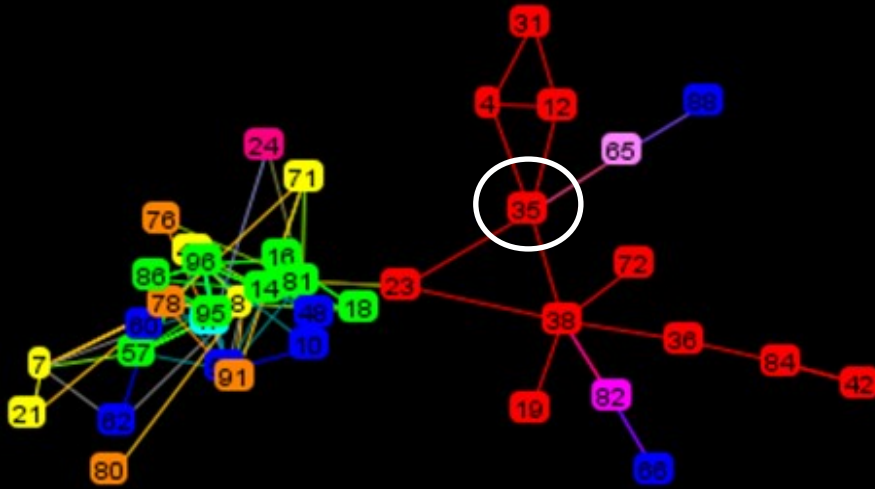
Solicit DO (Data found)



Solicit DO (Data found)




Pocket Switched Networks




- Topology changes every time unit
- Node 35 is a hub

Human-to-Human


Use of dynamic human




Inter-contact with External nodes



Inter-contact time for WiFi traces



Inter-contact time during the day



Inter-contact time during the day

