

## MongoDB

- a popular document database
  - stores collections of JSON-like documents with the identifier `_id`.
- For instance, documents about prominent computer scientists in a `bios` collection:

```
{
  "_id": 4,
  "awards": [ {"award": "Rosing Prize", "year": 1999, "by": "Norwegian Data Association"},
              {"award": "Turing Award", "year": 2001, "by": "ACM"},
              {"award": "IEEE John von Neumann Medal", "year": 2001, "by": "IEEE"} ],
  "birth": "1926-08-27",
  "contribs": ["OOP", "Simula"],
  "death": "2002-08-10",
  "name": {"first": "Kristen", "last": "Nygaard"}
}
```

We formalize JSON-documents as trees and define a relational view over them.

events collection:

```
{
  "_id": 1,
  "year": 1997,
  "event": "Deep Blue defeats Garry Kasparov"
},
{
  "_id": 2,
  "year": 1999,
  "event": "Melissa virus outbreak"
},
{
  "_id": 3,
  "year": 1999,
  "event": "Jeff Bezos is person of the year"
}
```

## Aggregation Framework: MQuery



MongoDB provides a powerful querying mechanism in the form of the **aggregation framework**, where a query is a multi-stage pipeline evaluated over one collection. We formalized this query language as MQuery, or  $\mathcal{M}^{\text{mupgl}}$ . It includes five stages:

<p><b>Match</b> <math>\mu_{\text{criterion}}</math></p>	<pre>db.bios.aggregate([   { \$match: { \$and: [     {"awards.year": { \$eq: 1999 }},     {"name.first": { \$eq: "Kristen" }}   ] } },</pre>	<table border="1"> <thead> <tr> <th>_id</th> <th>award</th> <th>year</th> <th>by</th> <th>birth</th> <th>contribs lit</th> <th>death</th> <th>name.first</th> <th>name.last</th> </tr> </thead> <tbody> <tr> <td>4</td> <td>Rosing Prize</td> <td>1999</td> <td>Norwegian Data Association</td> <td>1926-08-27</td> <td>OOP Simula</td> <td>2002-08-10</td> <td>Kristen</td> <td>Nygaard</td> </tr> <tr> <td></td> <td>Turing Award</td> <td>2001</td> <td>ACM</td> <td></td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td></td> <td>IEEE John von Neumann Medal</td> <td>2001</td> <td>IEEE</td> <td></td> <td></td> <td></td> <td></td> <td></td> </tr> </tbody> </table> <p>Selection <math>\sigma</math></p>	_id	award	year	by	birth	contribs lit	death	name.first	name.last	4	Rosing Prize	1999	Norwegian Data Association	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard		Turing Award	2001	ACM							IEEE John von Neumann Medal	2001	IEEE					
_id	award	year	by	birth	contribs lit	death	name.first	name.last																														
4	Rosing Prize	1999	Norwegian Data Association	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard																														
	Turing Award	2001	ACM																																			
	IEEE John von Neumann Medal	2001	IEEE																																			
<p><b>Unwind</b> <math>\omega_{\text{path}}</math></p>	<pre>{ \$unwind: "\$awards" },</pre>	<table border="1"> <thead> <tr> <th>_id</th> <th>awards.award</th> <th>awards.year</th> <th>awards.by</th> <th>birth</th> <th>contribs lit</th> <th>death</th> <th>name.first</th> <th>name.last</th> </tr> </thead> <tbody> <tr> <td>4</td> <td>Rosing Prize</td> <td>1999</td> <td>Norwegian Data Association</td> <td>1926-08-27</td> <td>OOP Simula</td> <td>2002-08-10</td> <td>Kristen</td> <td>Nygaard</td> </tr> <tr> <td>4</td> <td>Turing Award</td> <td>2001</td> <td>ACM</td> <td>1926-08-27</td> <td>OOP Simula</td> <td>2002-08-10</td> <td>Kristen</td> <td>Nygaard</td> </tr> <tr> <td>4</td> <td>IEEE John von Neumann Medal</td> <td>2001</td> <td>IEEE</td> <td>1926-08-27</td> <td>OOP Simula</td> <td>2002-08-10</td> <td>Kristen</td> <td>Nygaard</td> </tr> </tbody> </table> <p>Unnest <math>\chi</math></p>	_id	awards.award	awards.year	awards.by	birth	contribs lit	death	name.first	name.last	4	Rosing Prize	1999	Norwegian Data Association	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard	4	Turing Award	2001	ACM	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard	4	IEEE John von Neumann Medal	2001	IEEE	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard
_id	awards.award	awards.year	awards.by	birth	contribs lit	death	name.first	name.last																														
4	Rosing Prize	1999	Norwegian Data Association	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard																														
4	Turing Award	2001	ACM	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard																														
4	IEEE John von Neumann Medal	2001	IEEE	1926-08-27	OOP Simula	2002-08-10	Kristen	Nygaard																														
<p><b>Project</b> <math>\rho_{\text{path/definition}}</math></p>	<pre>{ \$project: {   "awards": true,   "firstName": "\$name.first",   "calledJohn": {     \$eq: ["\$name.first", "John"]   },   "invisible": "\$abc" } },</pre>	<table border="1"> <thead> <tr> <th>_id</th> <th>awards.award</th> <th>awards.year</th> <th>awards.by</th> <th>firstName</th> <th>calledJohn</th> </tr> </thead> <tbody> <tr> <td>4</td> <td>Rosing Prize</td> <td>1999</td> <td>Norwegian Data Association</td> <td>Kristen</td> <td>false</td> </tr> <tr> <td>4</td> <td>Turing Award</td> <td>2001</td> <td>ACM</td> <td>Kristen</td> <td>false</td> </tr> <tr> <td>4</td> <td>IEEE John von Neumann Medal</td> <td>2001</td> <td>IEEE</td> <td>Kristen</td> <td>false</td> </tr> </tbody> </table> <p>Extended projection <math>\pi</math></p>	_id	awards.award	awards.year	awards.by	firstName	calledJohn	4	Rosing Prize	1999	Norwegian Data Association	Kristen	false	4	Turing Award	2001	ACM	Kristen	false	4	IEEE John von Neumann Medal	2001	IEEE	Kristen	false												
_id	awards.award	awards.year	awards.by	firstName	calledJohn																																	
4	Rosing Prize	1999	Norwegian Data Association	Kristen	false																																	
4	Turing Award	2001	ACM	Kristen	false																																	
4	IEEE John von Neumann Medal	2001	IEEE	Kristen	false																																	
<p><b>Group</b> <math>\gamma_{\text{grouping:aggregation}}</math></p>	<pre>{ \$group: {   "_id": {"year": "\$awards.year"},   "awardNames": {     \$addToSet: "\$awards.award"   } } },</pre>	<table border="1"> <thead> <tr> <th>_id.year</th> <th>awardNames lit</th> </tr> </thead> <tbody> <tr> <td>1999</td> <td>Rosing Prize</td> </tr> <tr> <td>2001</td> <td>Turing Award IEEE John von Neumann Medal</td> </tr> </tbody> </table> <p>Nest <math>\nu</math></p>	_id.year	awardNames lit	1999	Rosing Prize	2001	Turing Award IEEE John von Neumann Medal																														
_id.year	awardNames lit																																					
1999	Rosing Prize																																					
2001	Turing Award IEEE John von Neumann Medal																																					
<p><b>Lookup</b> <math>\lambda_{\text{local=coll.foreign result}}</math></p>	<pre>{ \$lookup: {   from: "events",   localField: "_id.year",   foreignField: "year",   as: "joinedDocs" } } ])</pre>	<table border="1"> <thead> <tr> <th rowspan="2">_id.year</th> <th rowspan="2">awardNames lit</th> <th colspan="3">joinedDocs</th> </tr> <tr> <th>_id</th> <th>year</th> <th>event</th> </tr> </thead> <tbody> <tr> <td rowspan="2">1999</td> <td rowspan="2">Rosing Prize</td> <td>2</td> <td>1999</td> <td>Melissa virus outbreak</td> </tr> <tr> <td>3</td> <td>1999</td> <td>Jeff Bezos is person of the year</td> </tr> <tr> <td rowspan="2">2001</td> <td rowspan="2">Turing Award IEEE John von Neumann Medal</td> <td></td> <td></td> <td></td> </tr> </tbody> </table> <p>Left outer equijoin <math>\bowtie</math></p>	_id.year	awardNames lit	joinedDocs			_id	year	event	1999	Rosing Prize	2	1999	Melissa virus outbreak	3	1999	Jeff Bezos is person of the year	2001	Turing Award IEEE John von Neumann Medal																		
_id.year	awardNames lit	joinedDocs																																				
		_id	year	event																																		
1999	Rosing Prize	2	1999	Melissa virus outbreak																																		
		3	1999	Jeff Bezos is person of the year																																		
2001	Turing Award IEEE John von Neumann Medal																																					

## Expressivity of MQuery

We have shown that **well-typed** MQuery is equivalent to nested relational algebra (NRA):

- well-typed  $\mathcal{M}^{\text{mupgl}}$   $\equiv$  NRA over a single collection, hence it is possible to express **joins** without lookup
- well-typed  $\mathcal{M}^{\text{mupgl}}$   $\equiv$  NRA

Well-typedness is required as arbitrary MQueries may produce forests for which a relational view cannot be defined.

## Complexity of MQuery (and NRA)

Fragment	Query complexity	Combined complexity
$\mathcal{M}^m$	LogSpace-complete	
$\mathcal{M}^{mp}, \mathcal{M}^{mupgl}$	PTime-complete	
$\mathcal{M}^{mu}$	LogSpace-complete	NP-complete
$\mathcal{M}^{mup}, \mathcal{M}^{mul}, \mathcal{M}^{mupl}$		NP-complete
$\mathcal{M}^{mug}$		PSpace-hard
$\mathcal{M}^{mupg}, \mathcal{M}^{mupgl}$	$TA[2^{n^{O(1)}}, n^{O(1)}]$ -complete*	
NRA	$TA[2^{n^{O(1)}}, n^{O(1)}]$ -complete	

\* The class of problems solvable by an alternating Turing machine running in exponential time with polynomially many alternations.