

Fine-grained and efficient lineage querying of collection-based workflow provenance

Paolo Missier, Norman Paton, Khalid Belhajjame

Information Management Group

School of Computer Science, University of Manchester, UK

EDBT Conference

Lausanne, Switzerland, March 2010

- **Setting:**
 - *Black box* provenance of **workflow** data products
- **Fine-grained provenance:**
 - tracking provenance through **collections**: motivation
 - **functional model** of collection-oriented workflow processing
- **Efficient query processing:**
 - leveraging the functional model to achieve efficient processing for a simple query model
- **Experimental evaluation**

- Provenance graph is an unfolding of the workflow graph structure
 - large: grows with size of input
 - lineage queries involve graph traversal

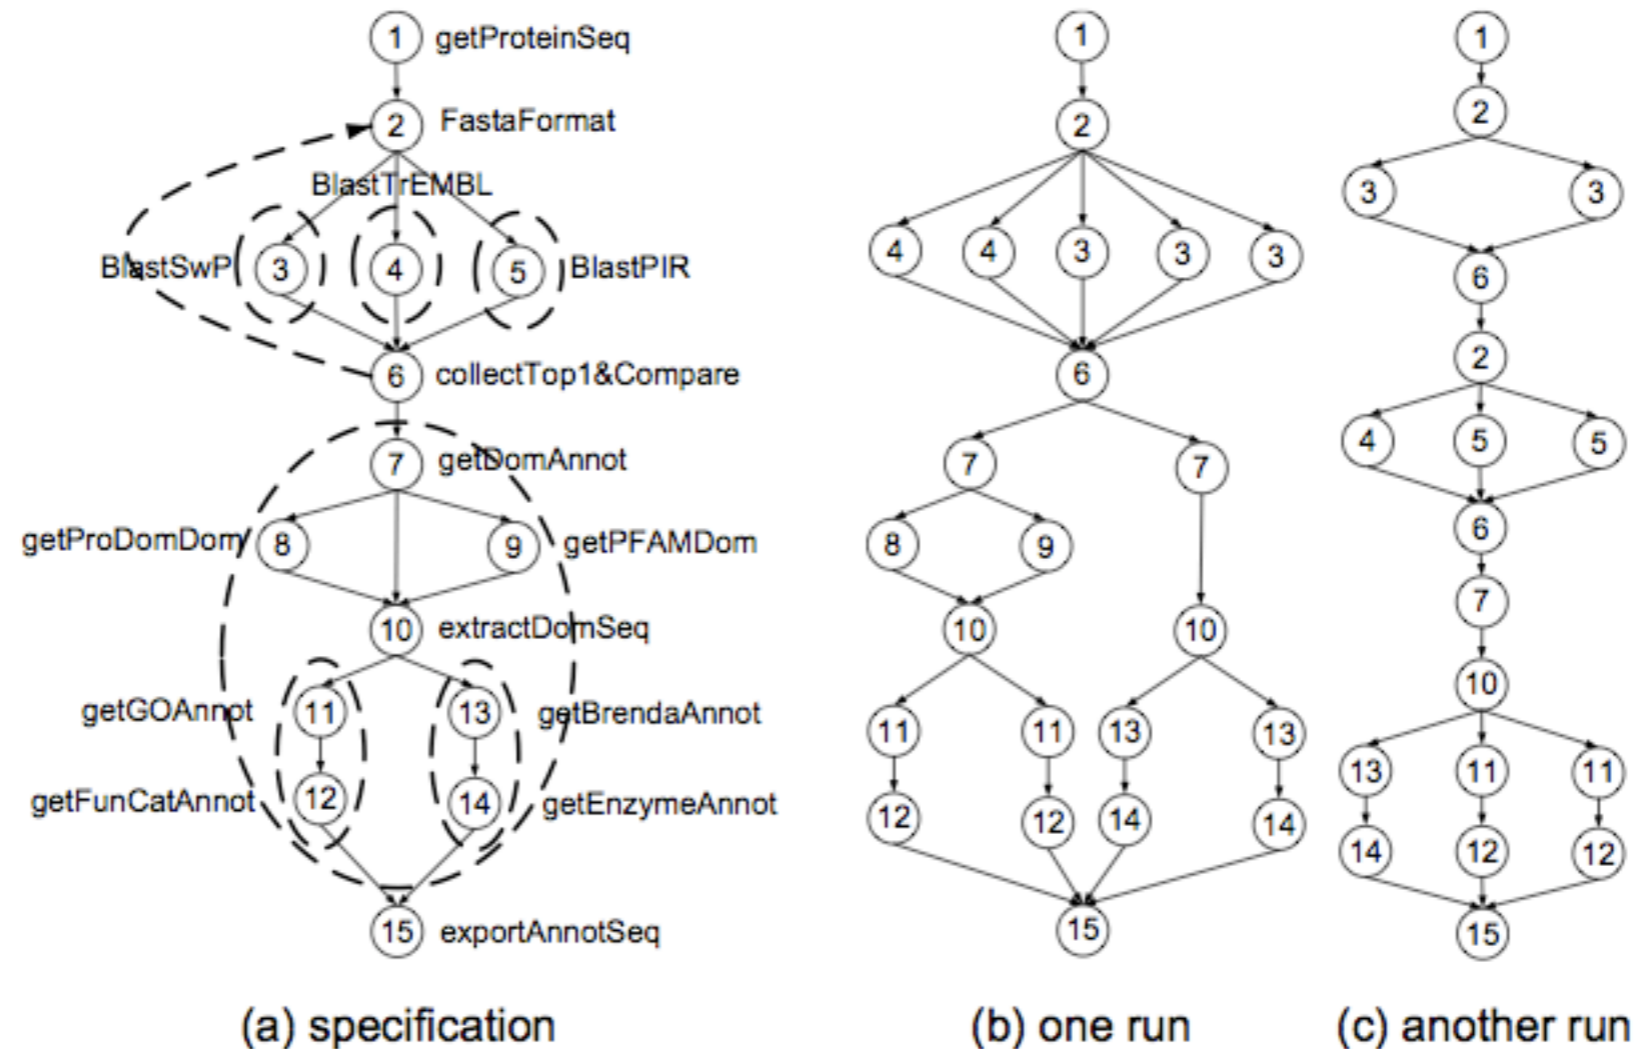



Fig. 1. Protein annotation workflow specification and runs

From:


Z. Bao, S. Cohen-Boulakia, S. Davidson, A. Eyal, and S. Khanna, "Differencing Provenance in Scientific Workflows," Procs. ICDE, 2009.

Workflow graph




$$y = [[y_{11} \dots y_{1n}], \dots [y_{m1} \dots y_{mn}]]$$

Provenance graph




- Query the provenance of individual collections elements
- But, avoid computing transitive closures on the provenance graph
 - potentially very large
- Traverse the workflow graph instead -- much smaller
- This results in substantial performance improvement for typical queries

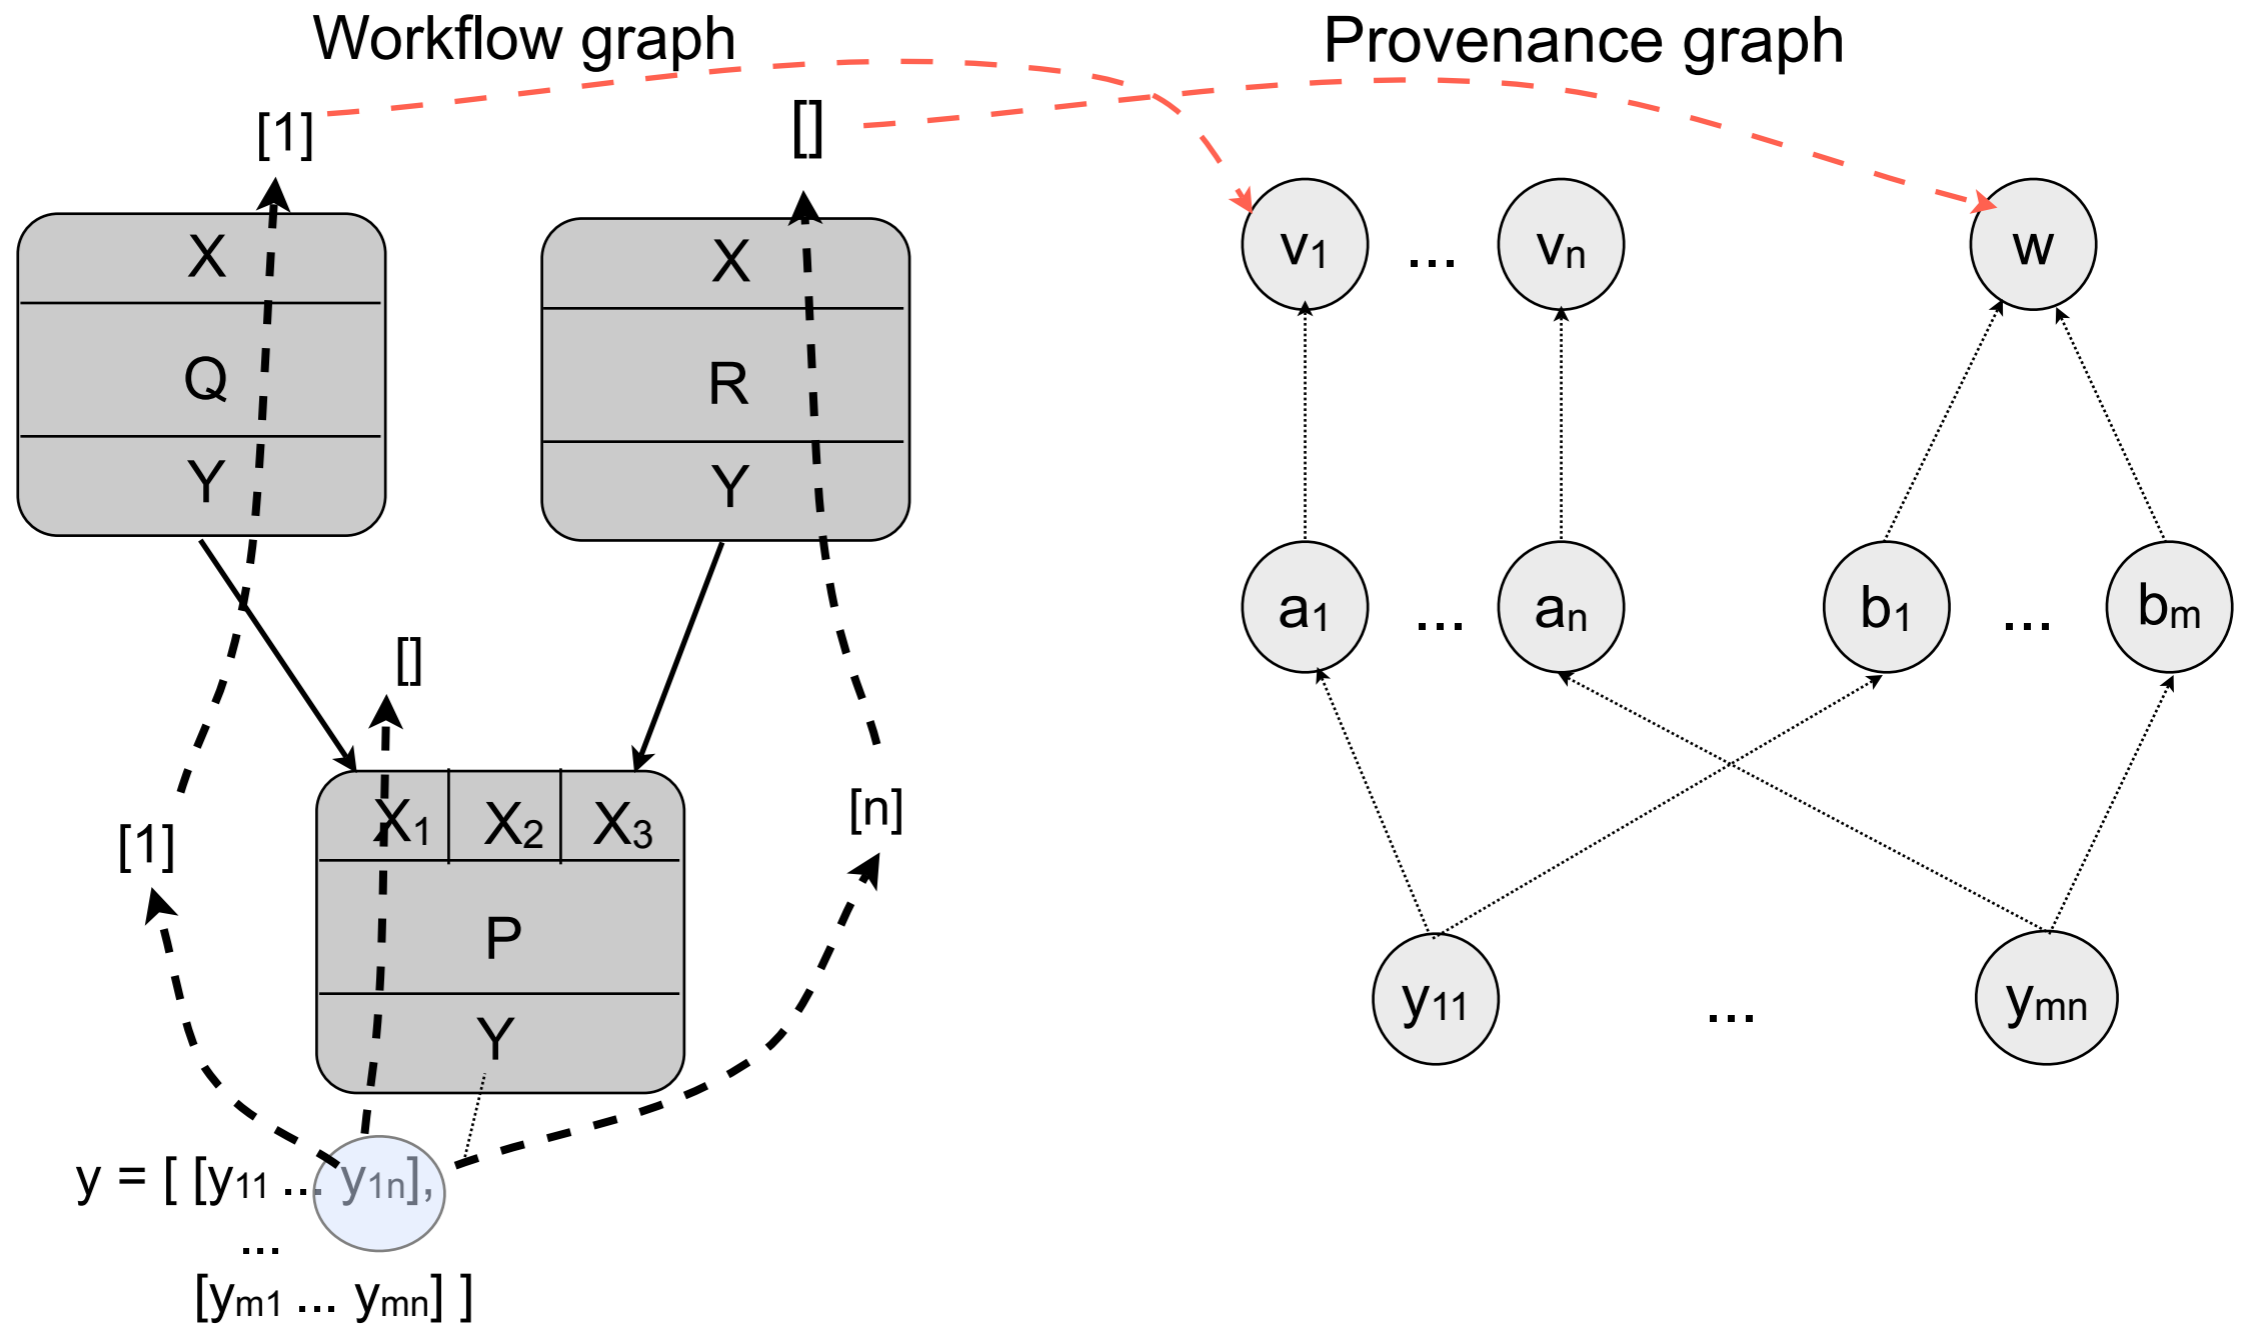
Workflow graph



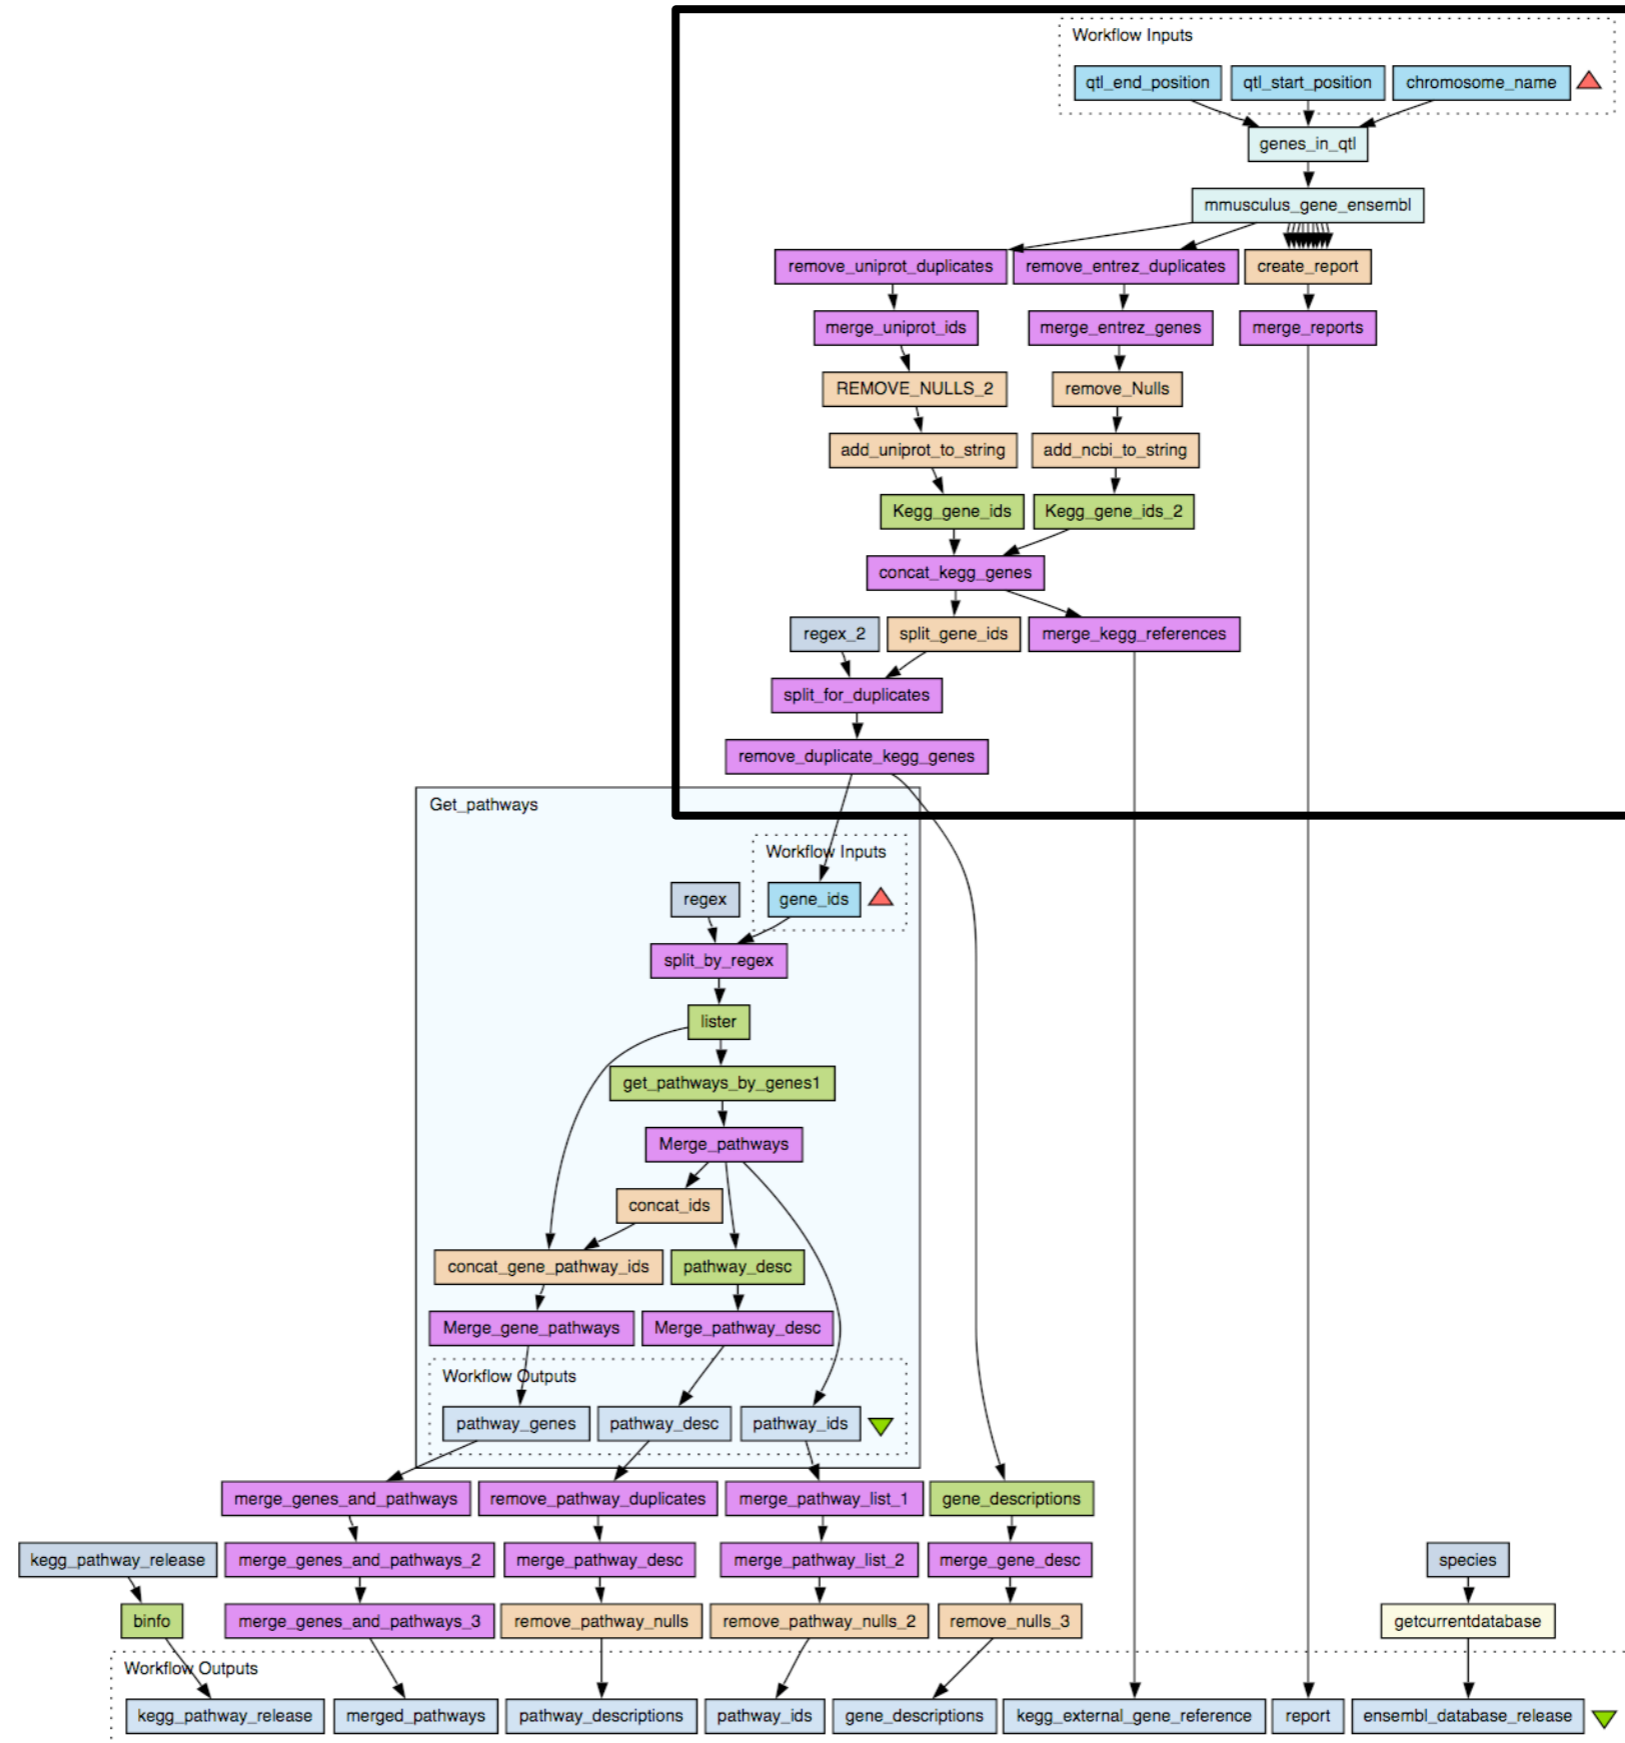
Provenance graph

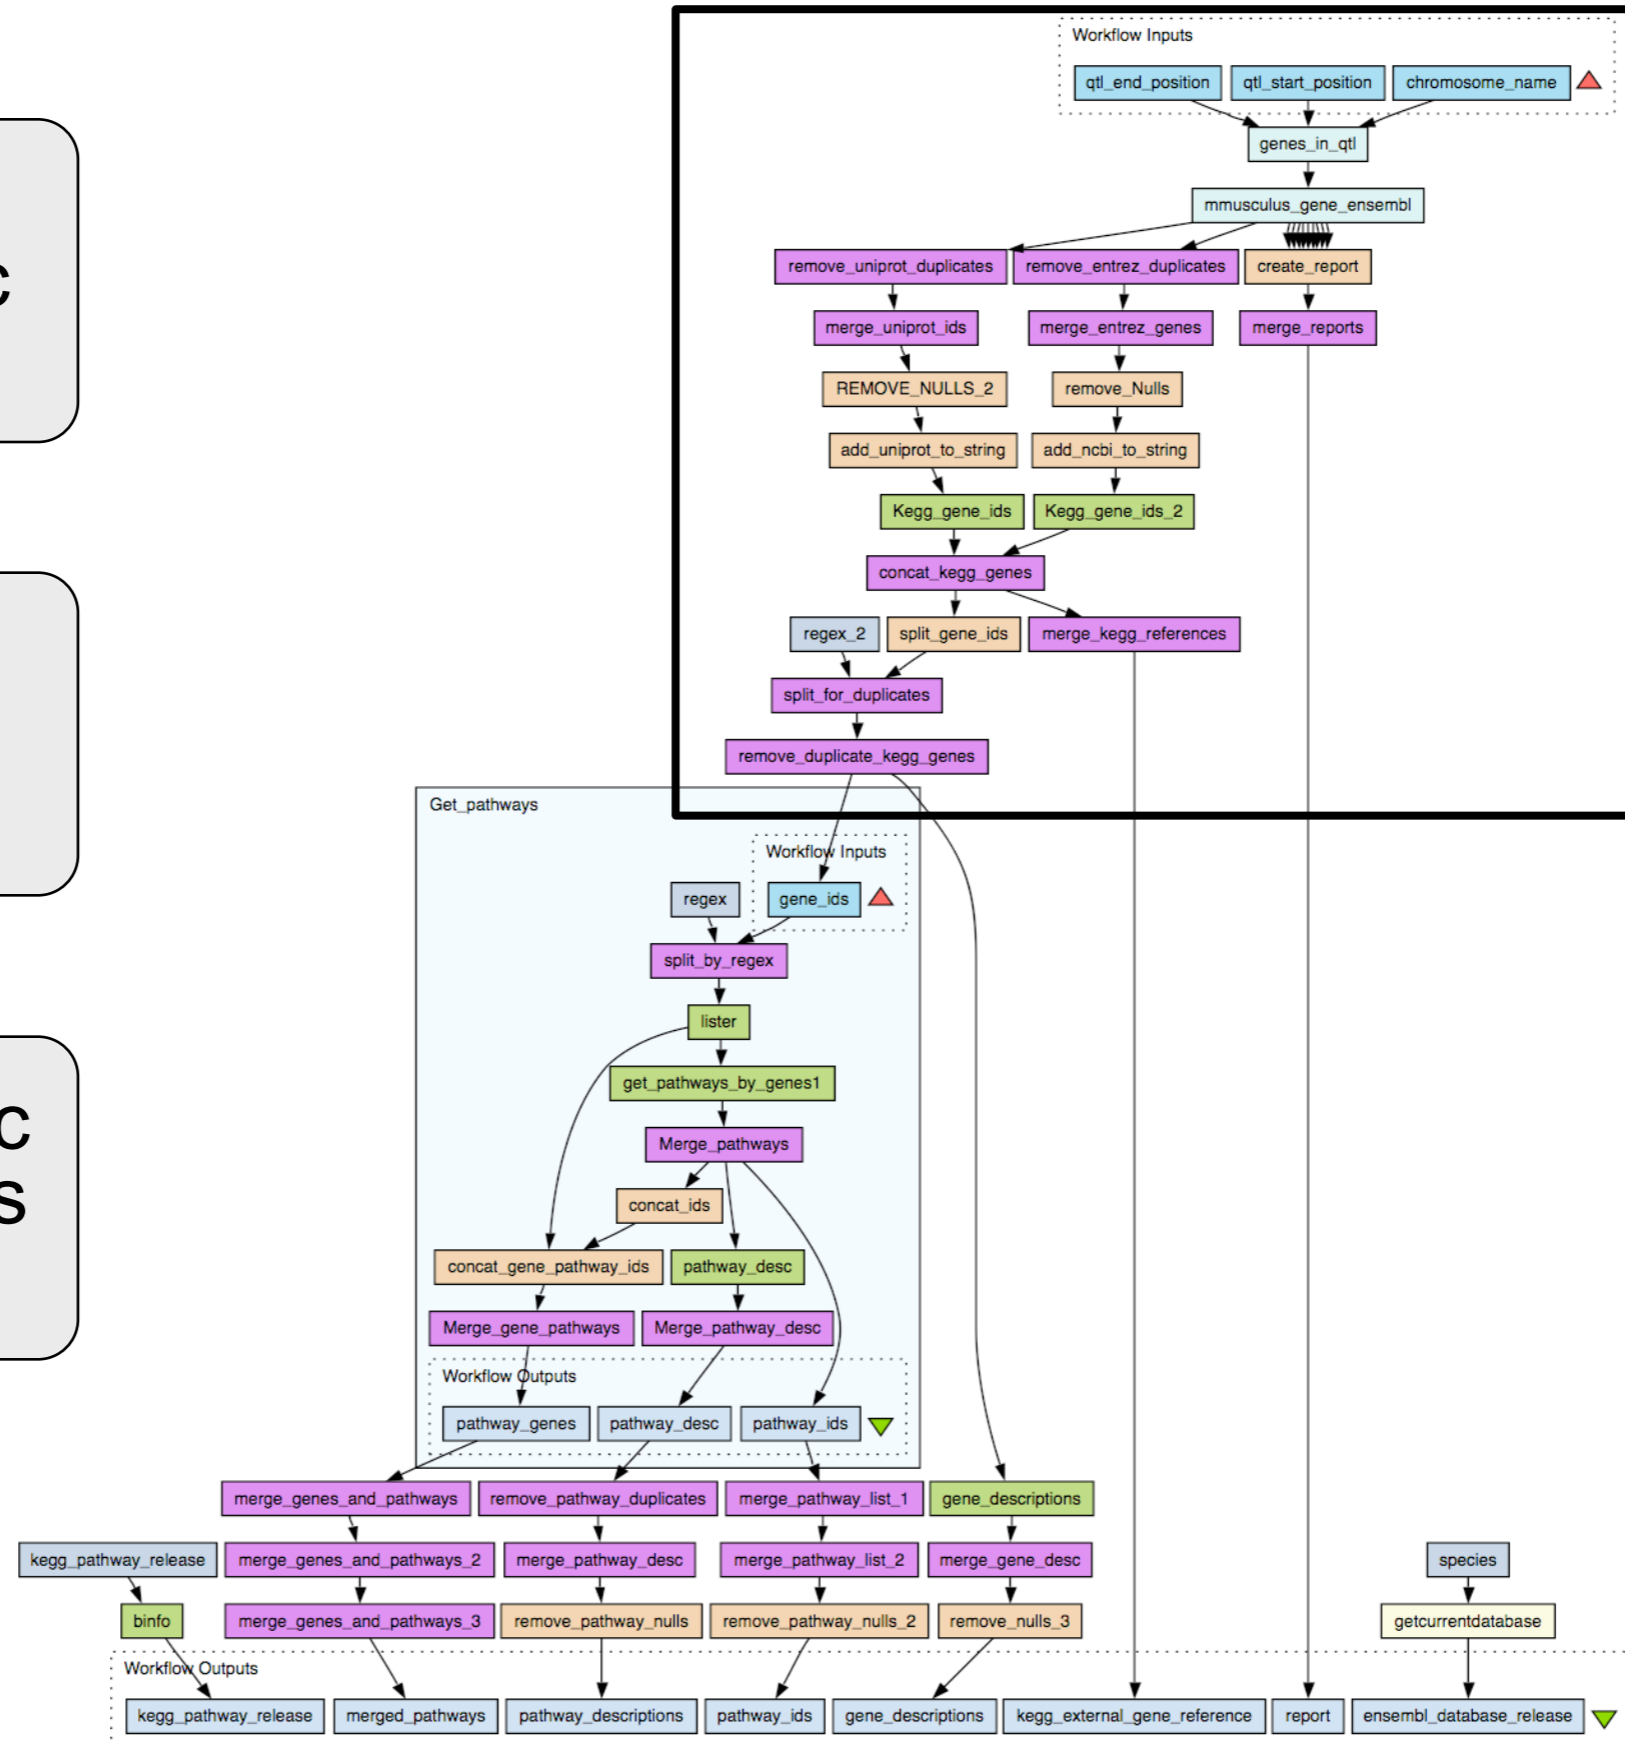




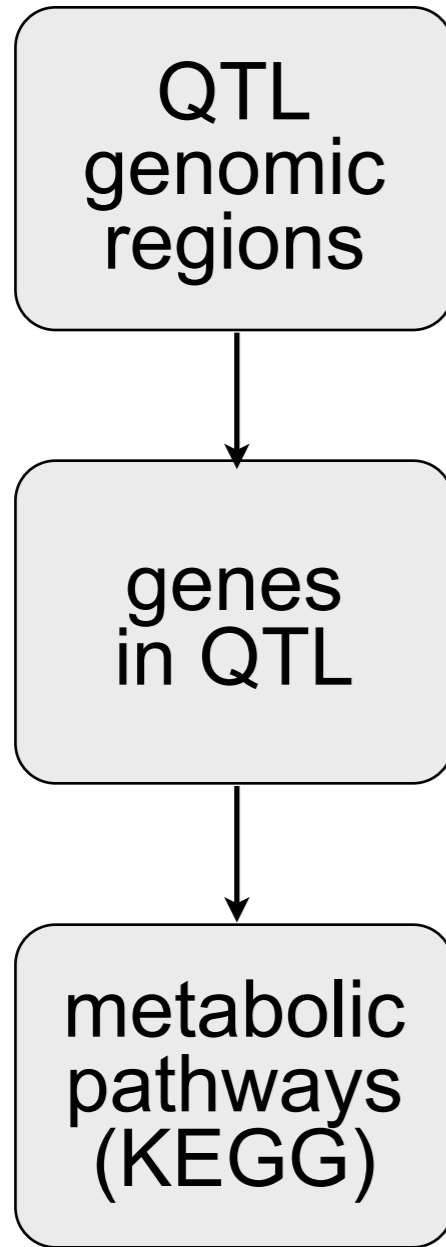
- Query the provenance of individual collections elements
- But, avoid computing transitive closures on the provenance graph
 - potentially very large
- Traverse the workflow graph instead -- much smaller
- This results in substantial performance improvement for typical queries

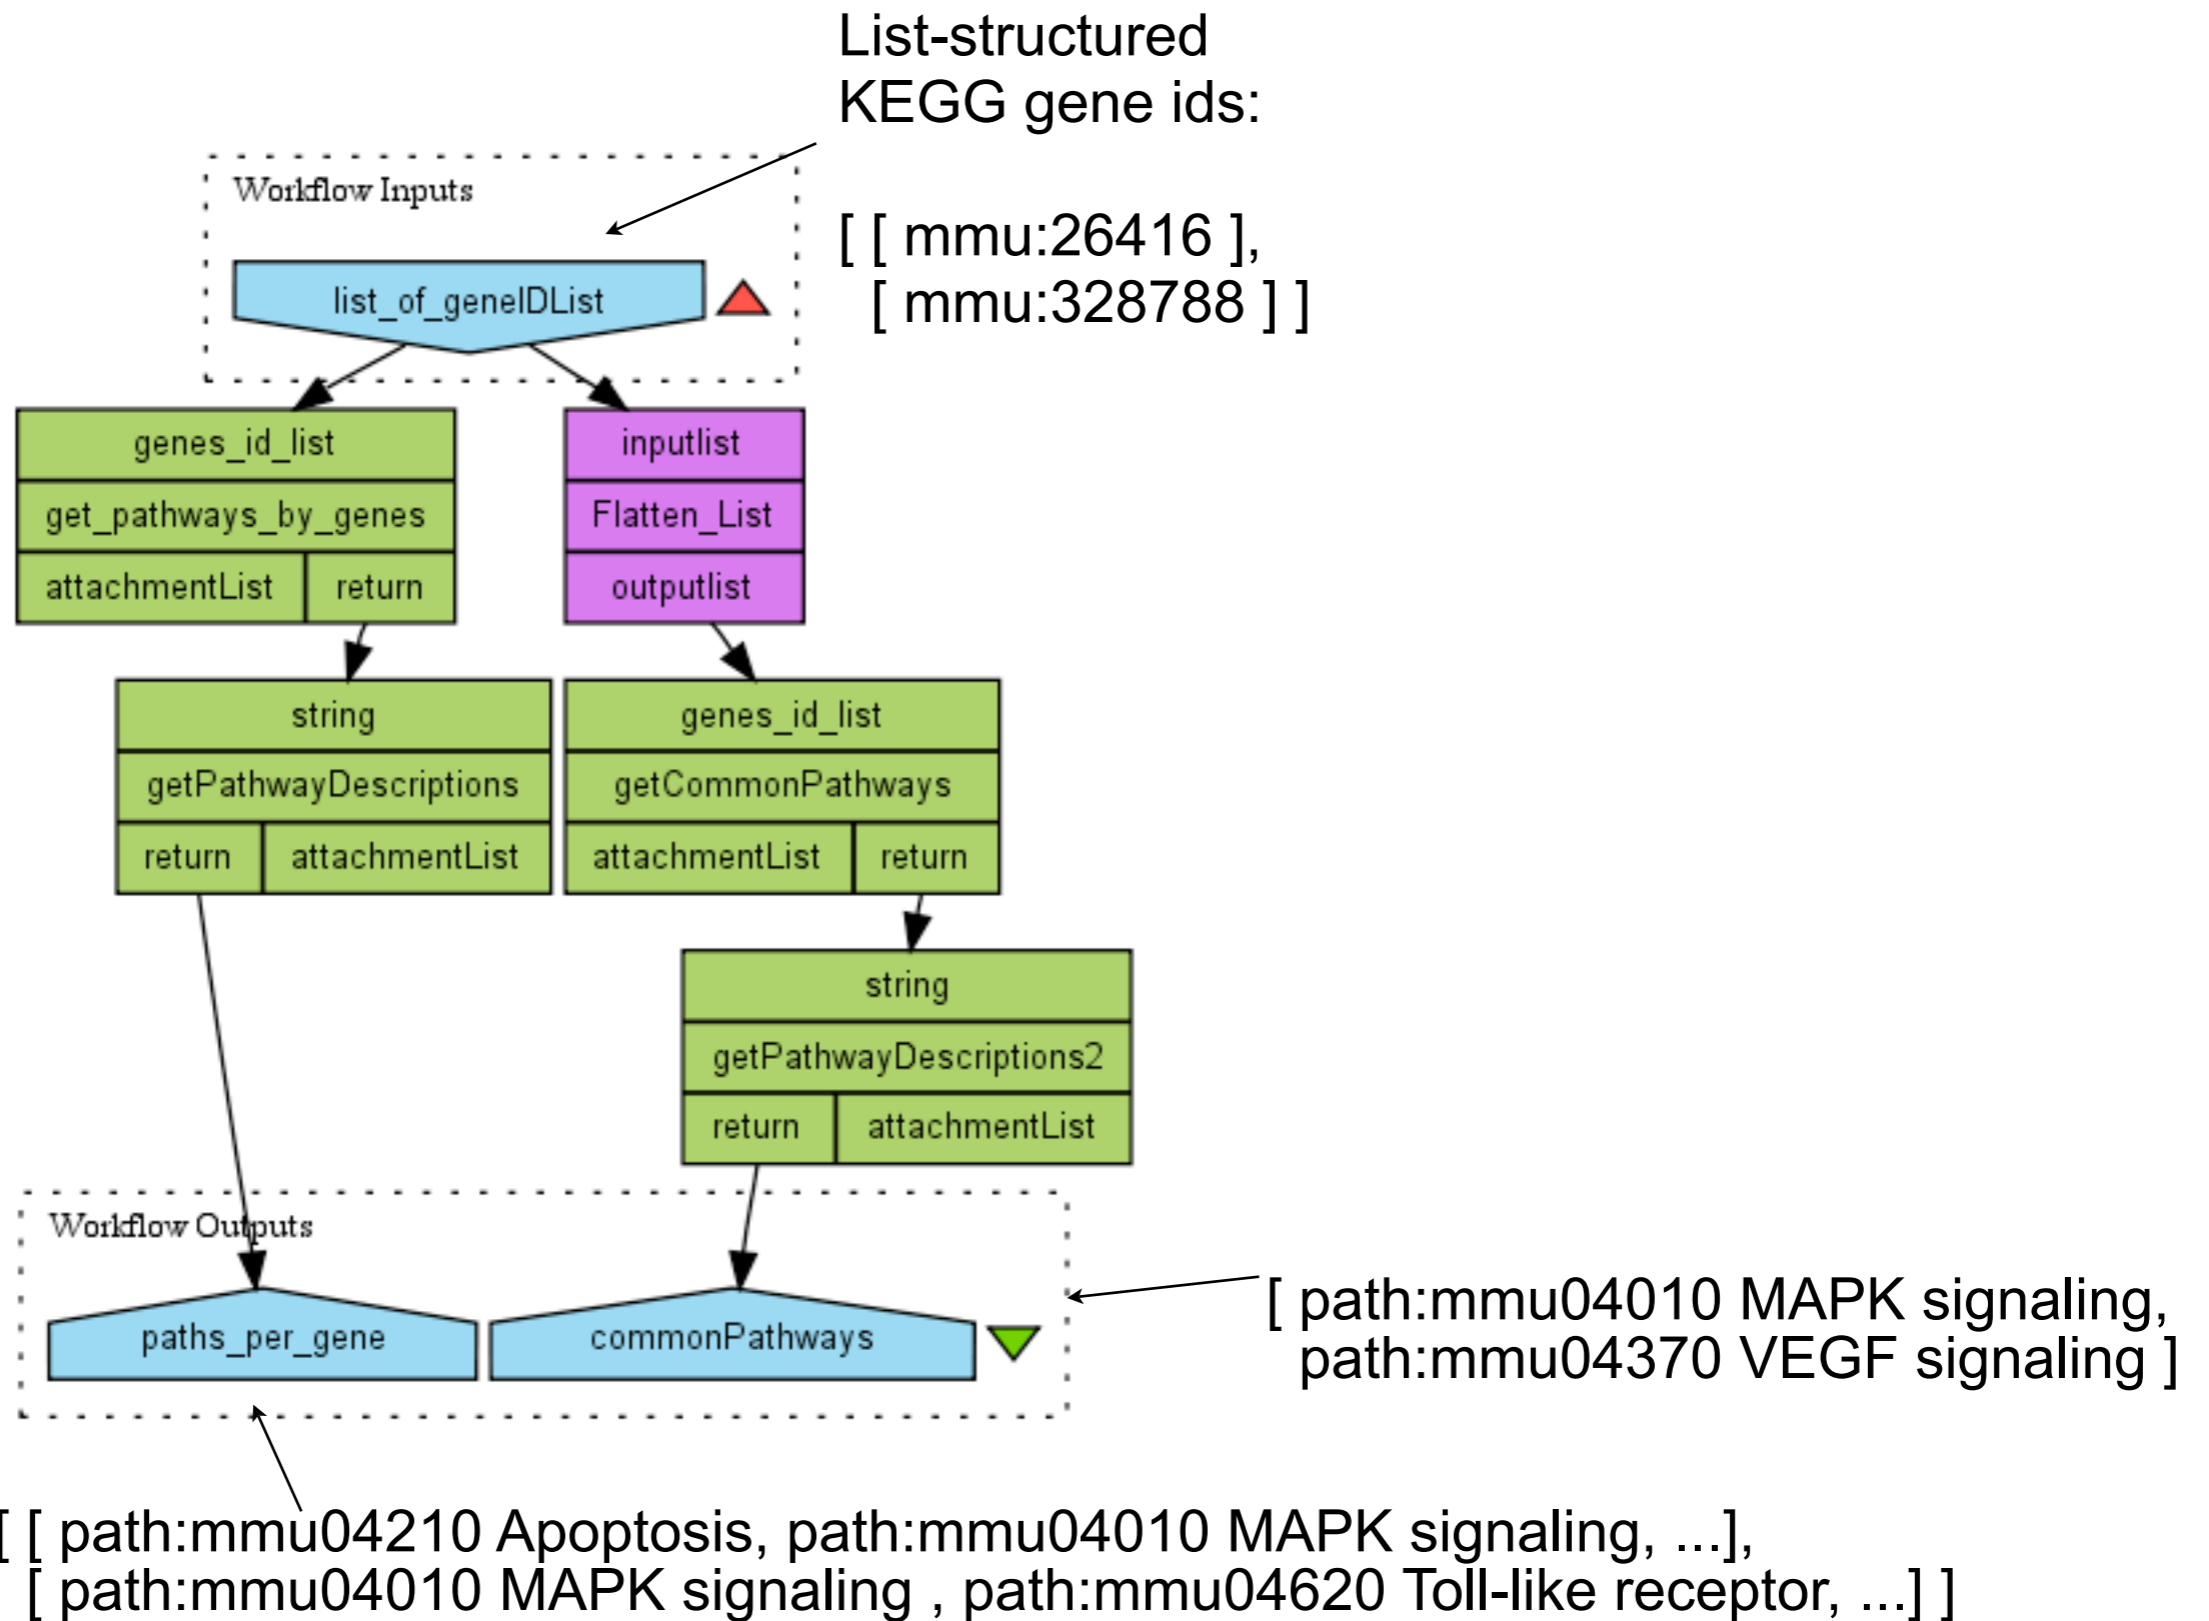


- Query the provenance of individual collections elements
- But, avoid computing transitive closures on the provenance graph
 - potentially very large
- Traverse the workflow graph instead -- much smaller
- This results in substantial performance improvement for typical queries



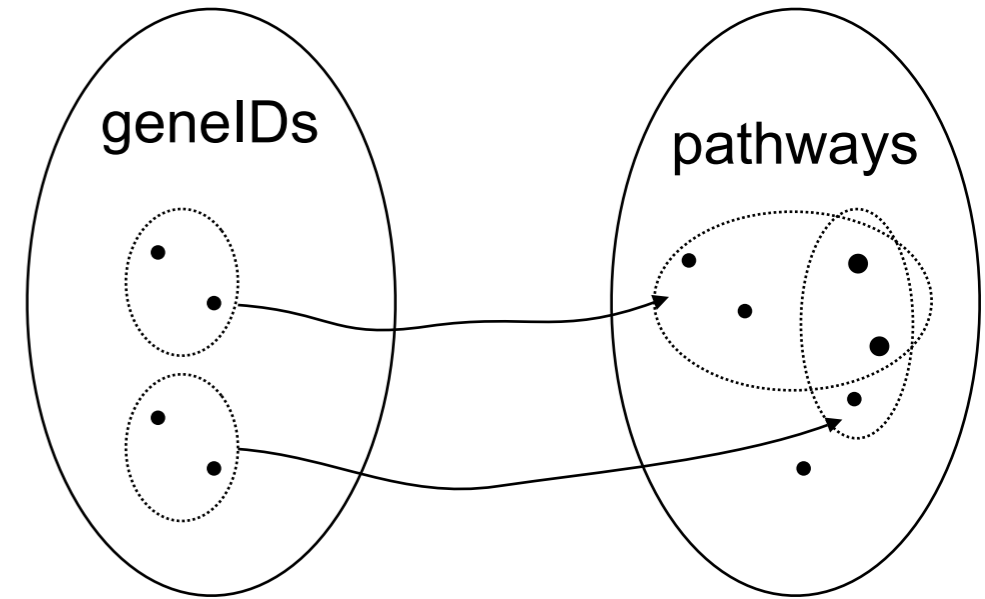
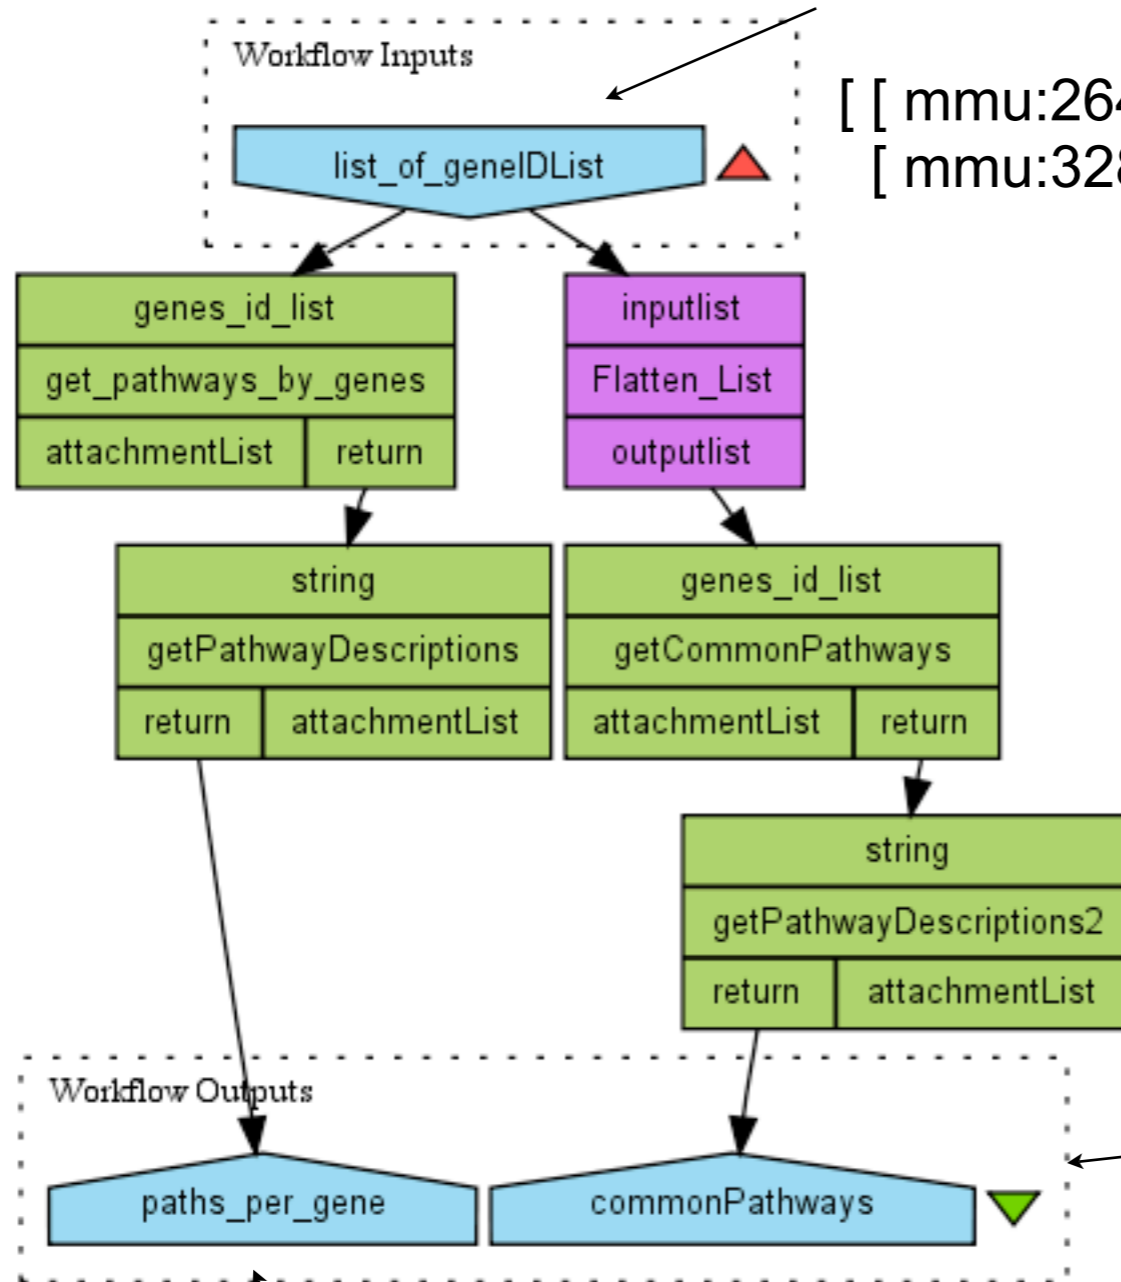






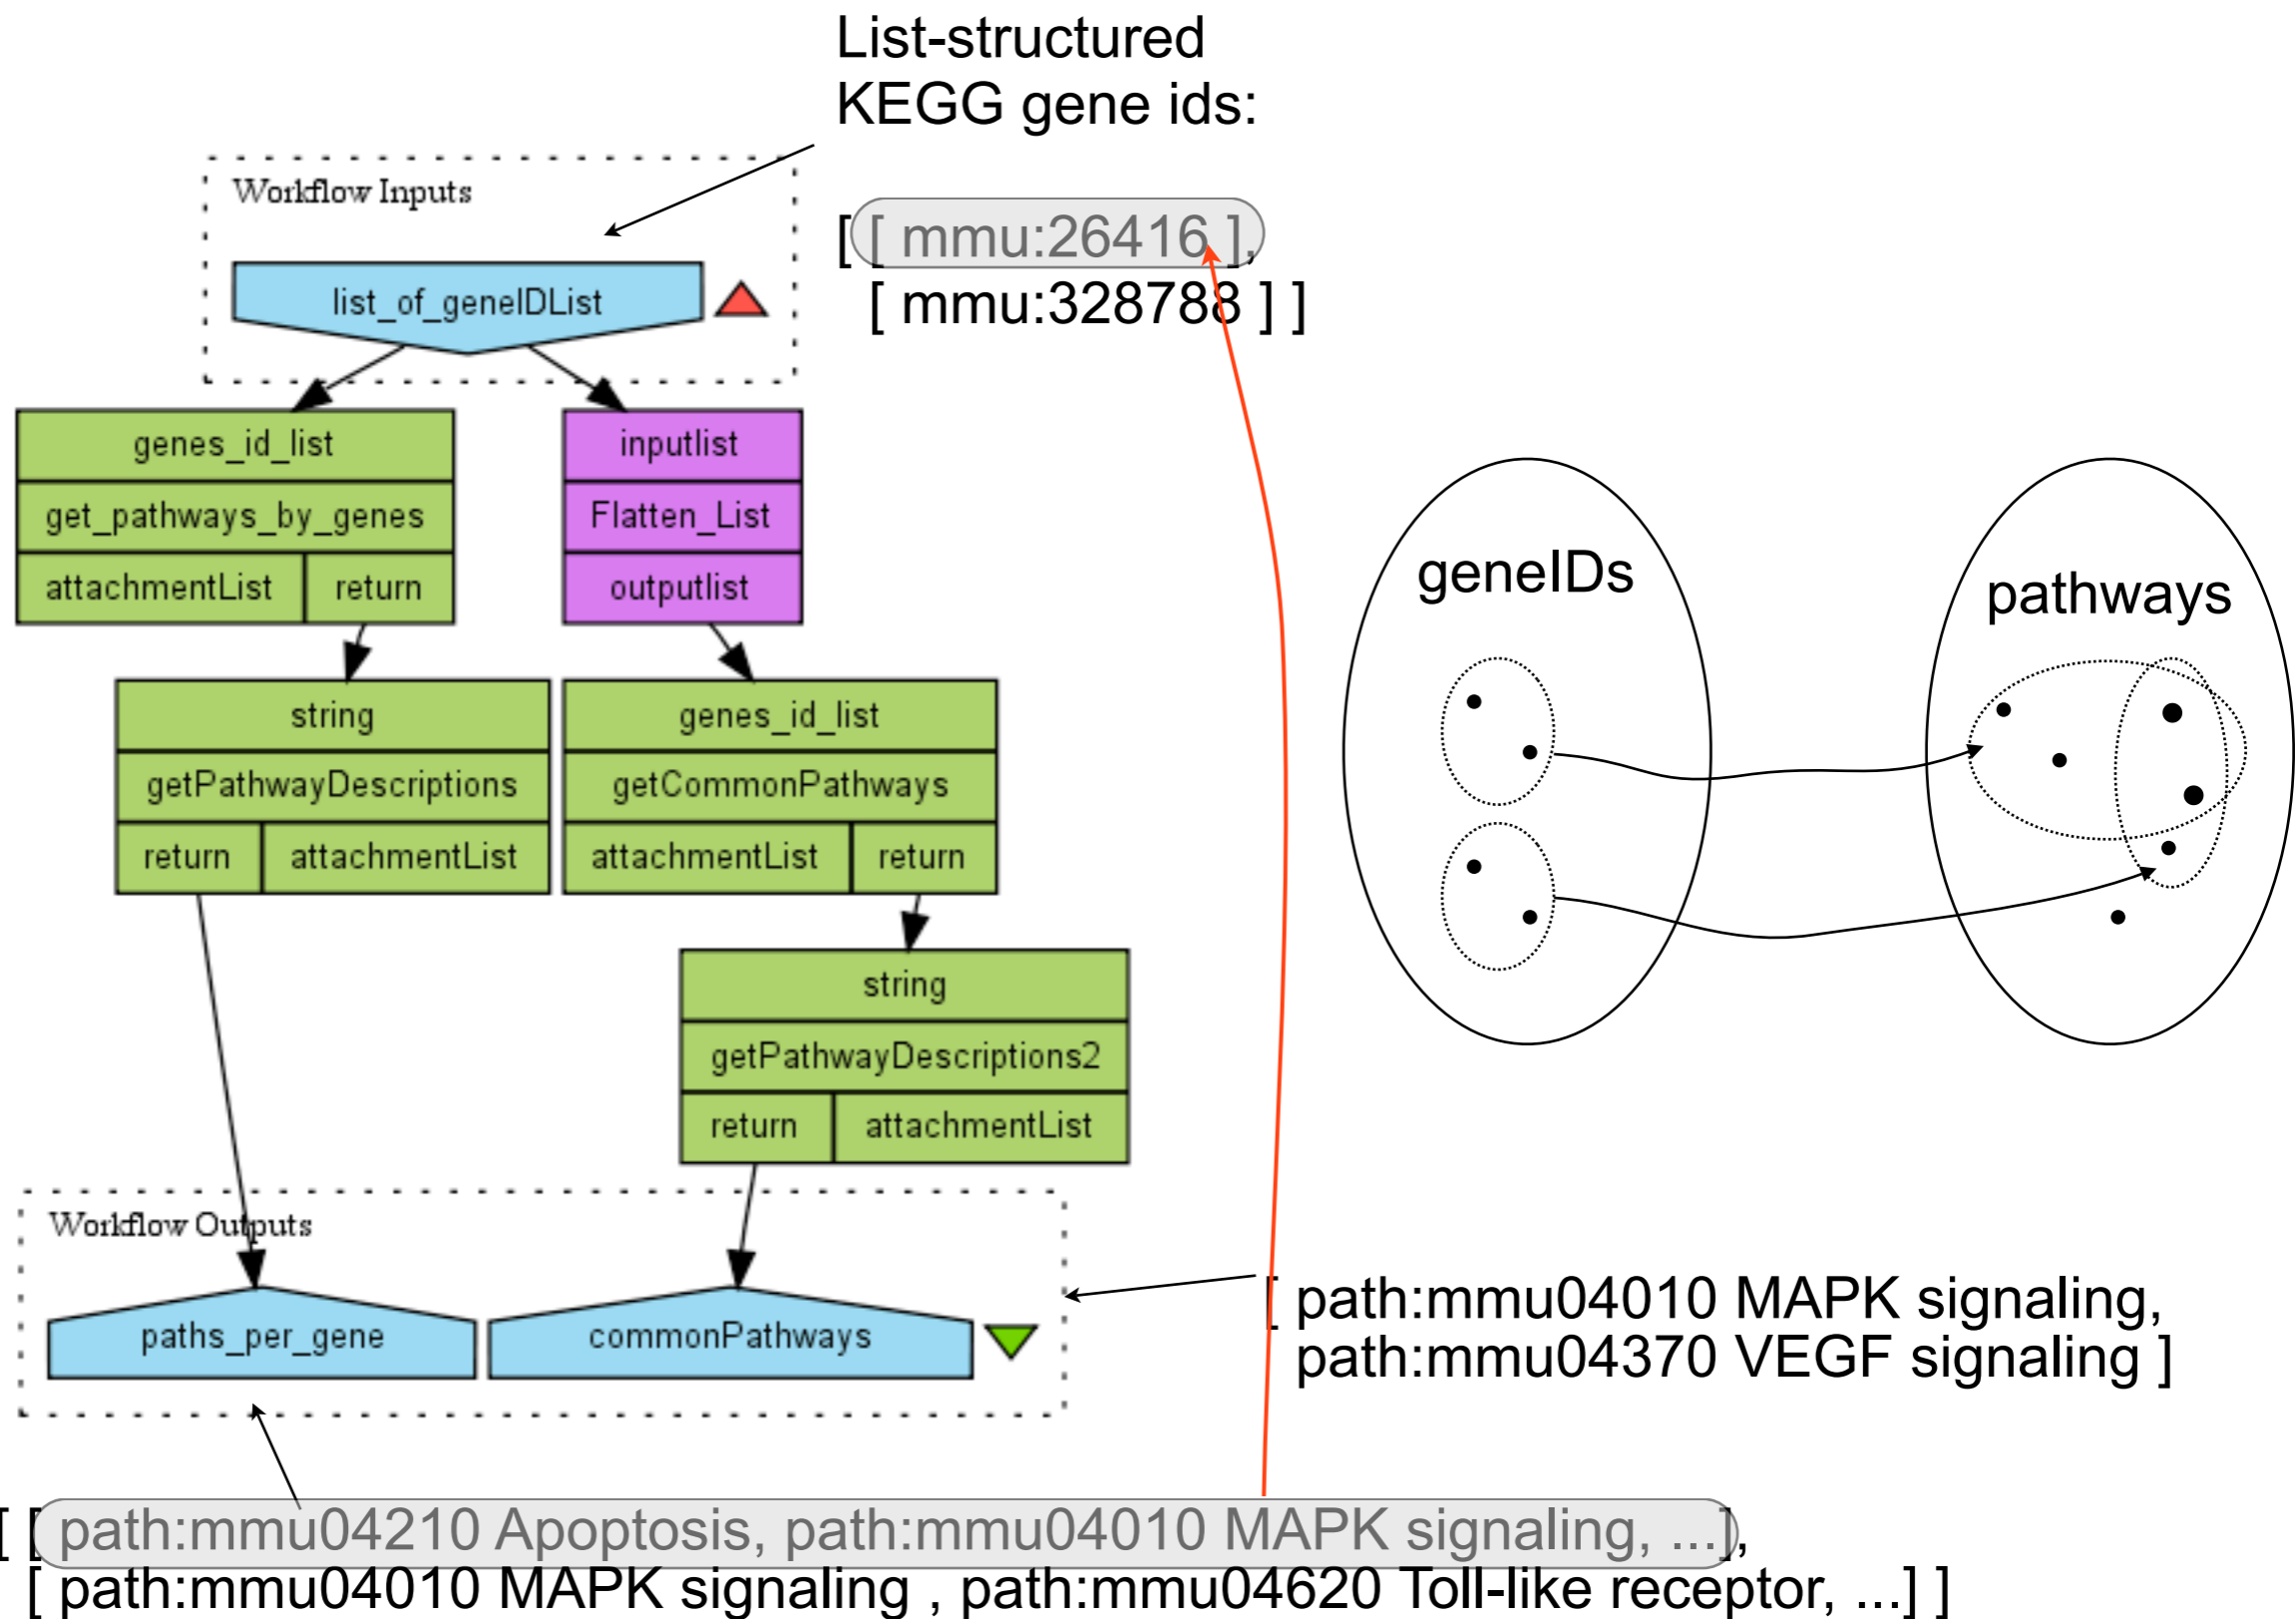
List-structured KEGG gene ids:

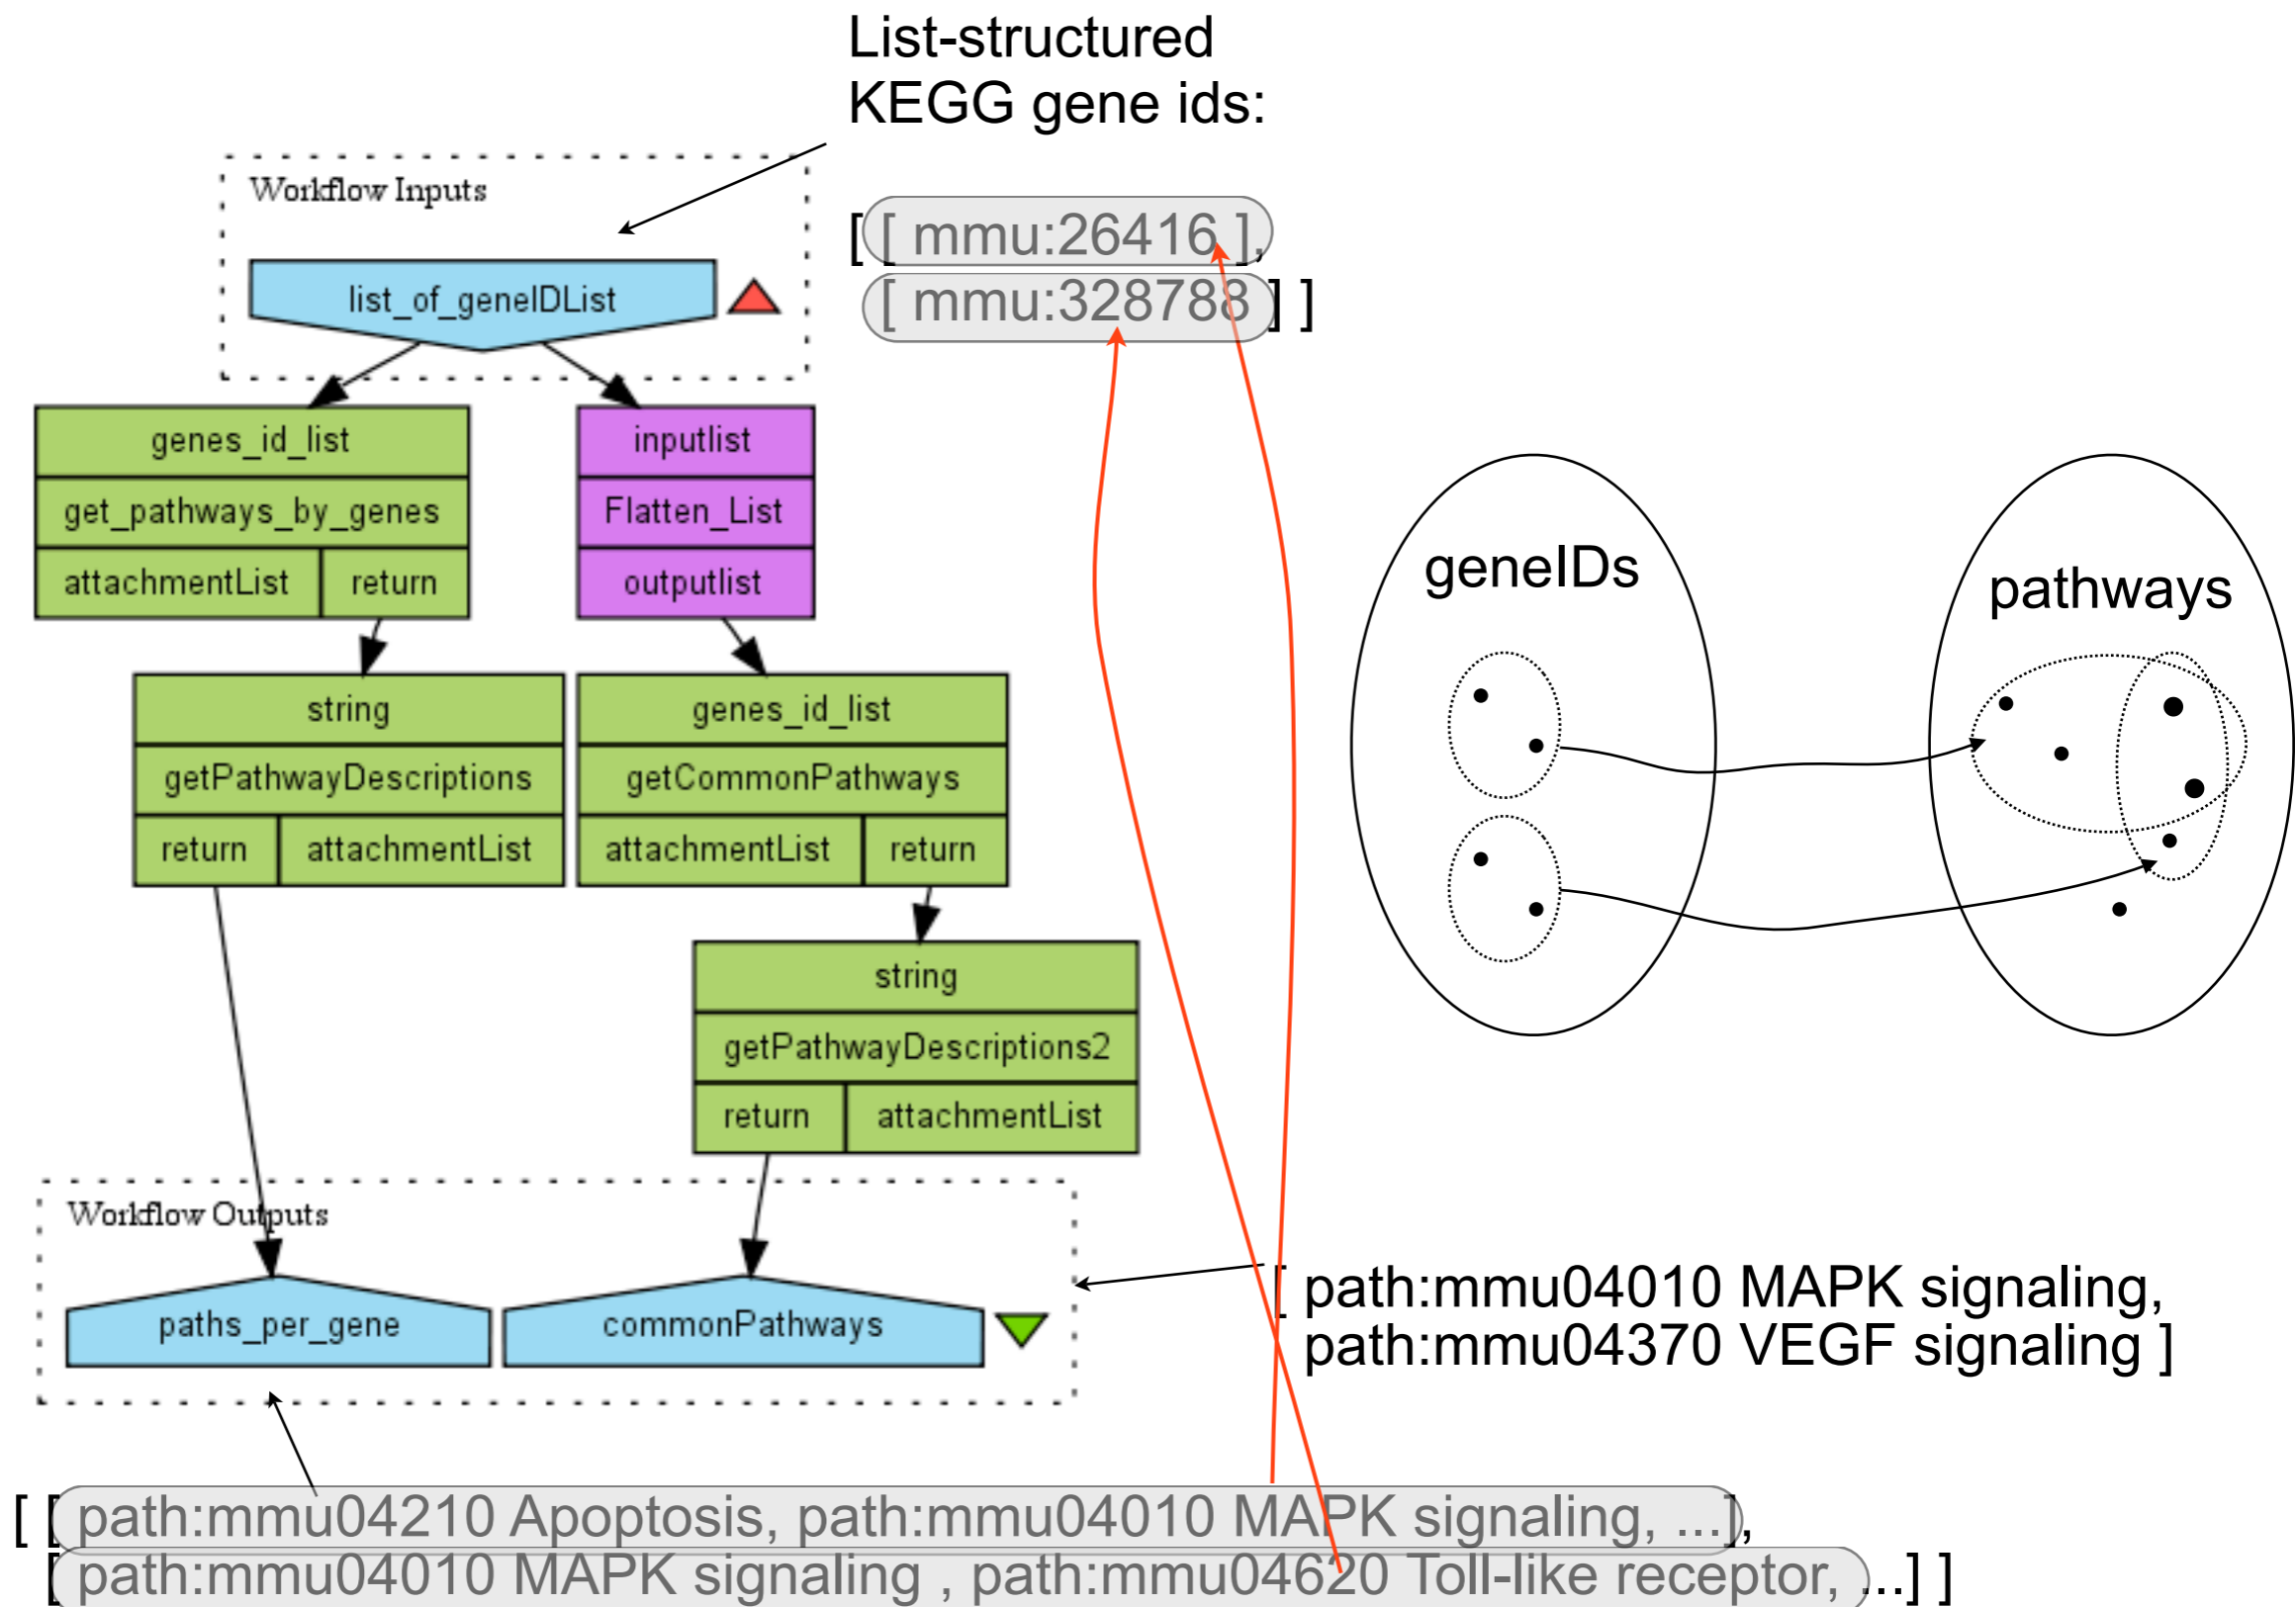
```
[ [ mmu:26416 ],
  [ mmu:328788 ] ]
```

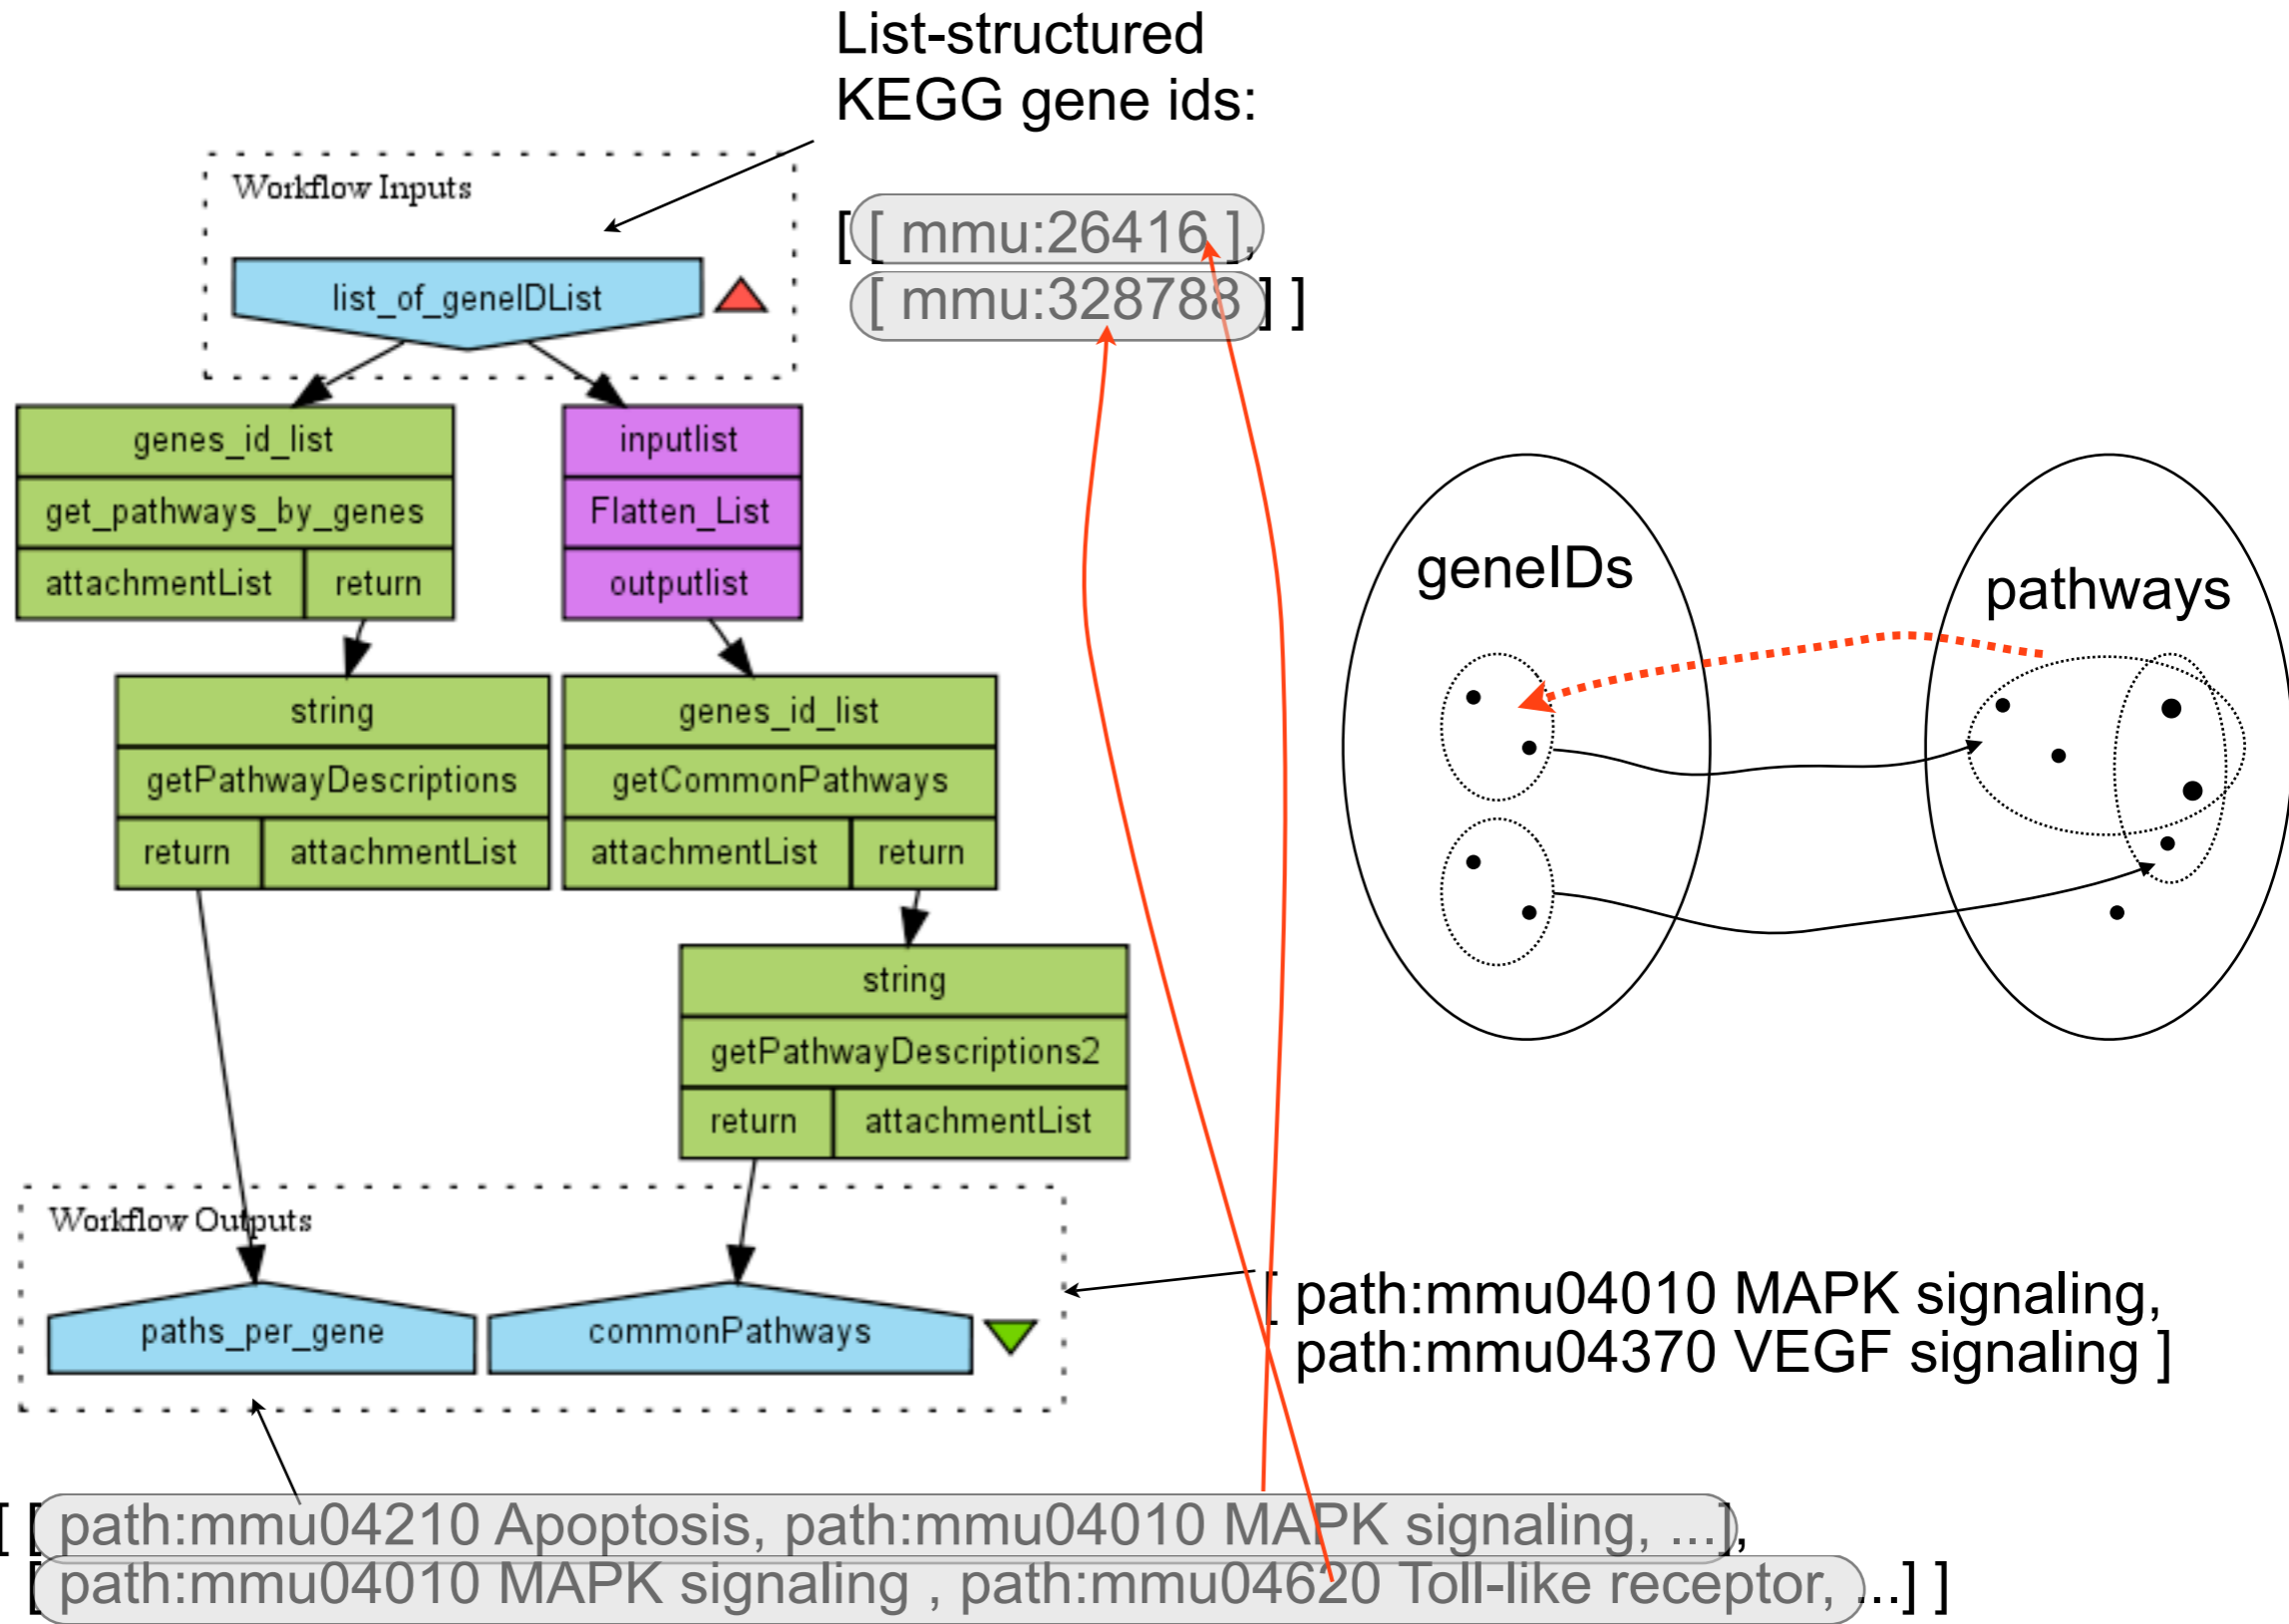


[path:mmu04010 MAPK signaling,
path:mmu04370 VEGF signaling]

```
[ [ path:mmu04210 Apoptosis, path:mmu04010 MAPK signaling, ...],
  [ path:mmu04010 MAPK signaling , path:mmu04620 Toll-like receptor, ...] ]
```






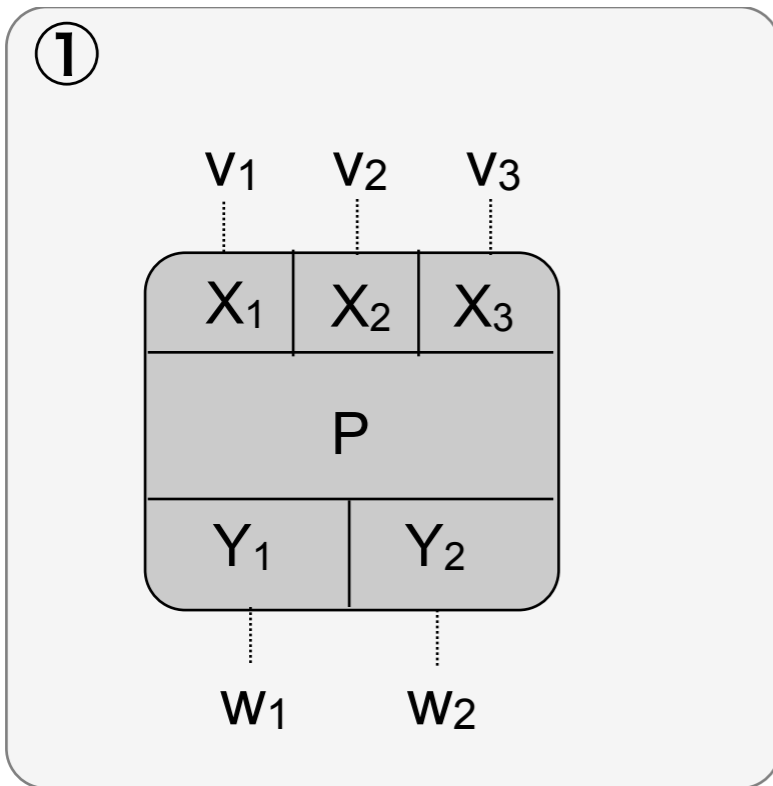


- **Setting:**
 - *Black box* provenance of workflow data products
- **Fine-grained provenance:**
 - tracking provenance through collection elements
 - motivation
 - ➔ **functional model** of collection-oriented workflow processing

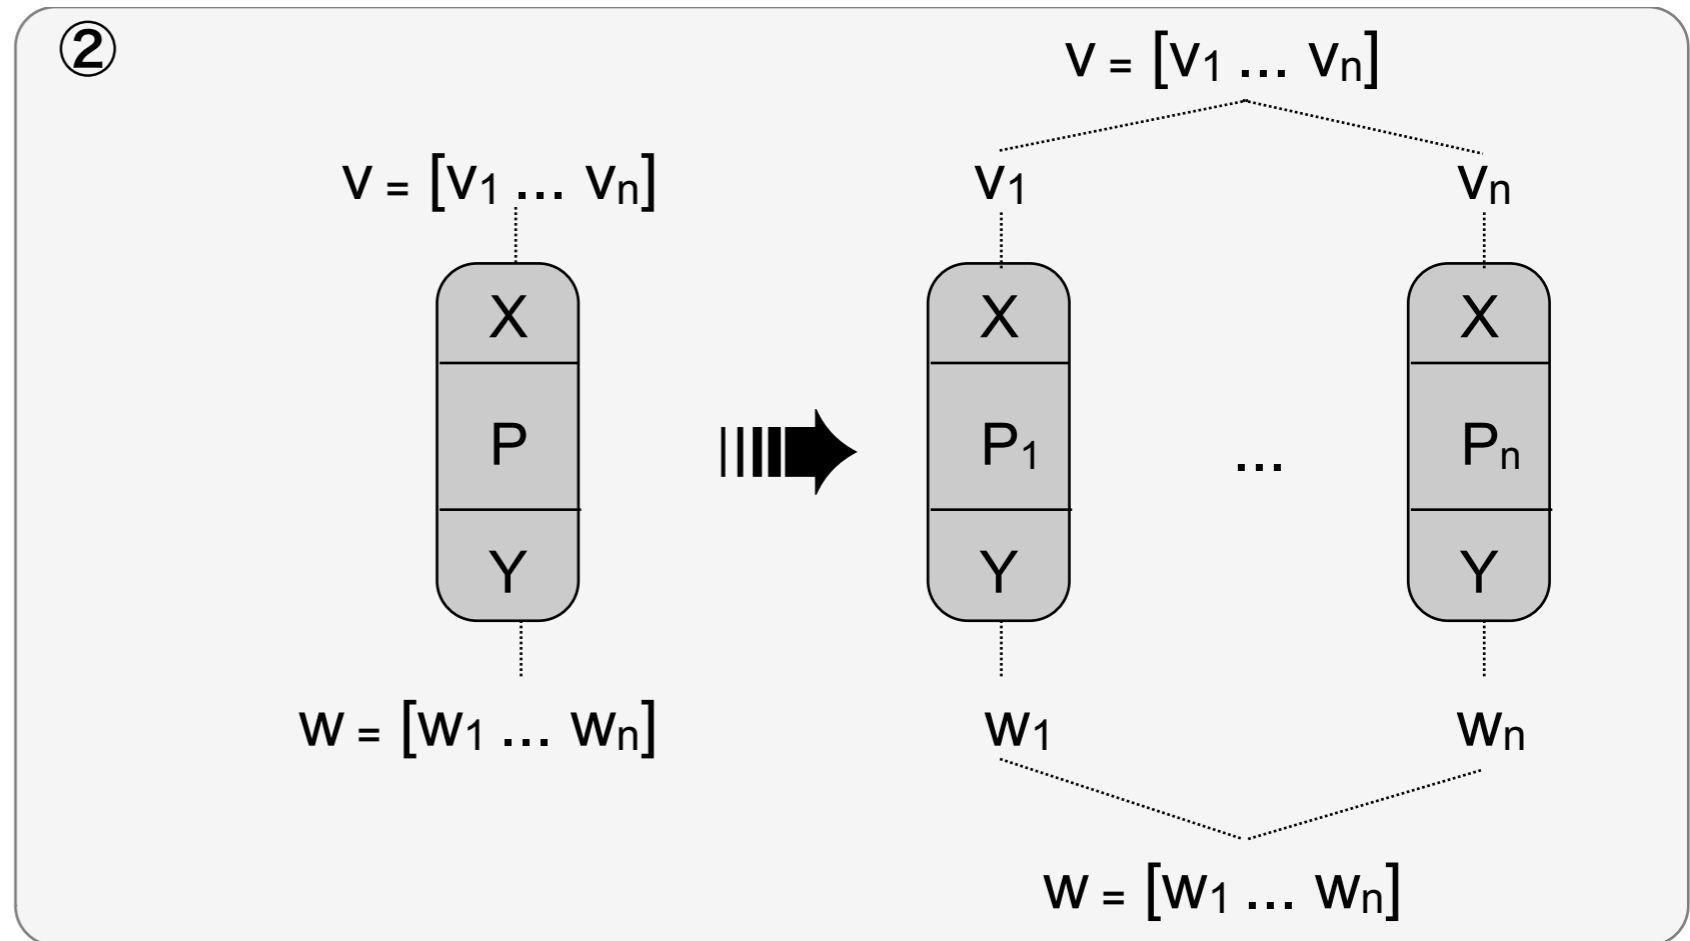
Simple processing:
service expects atomic values,
receives atomic values



Simple processing:
service expects atomic values,
receives atomic values

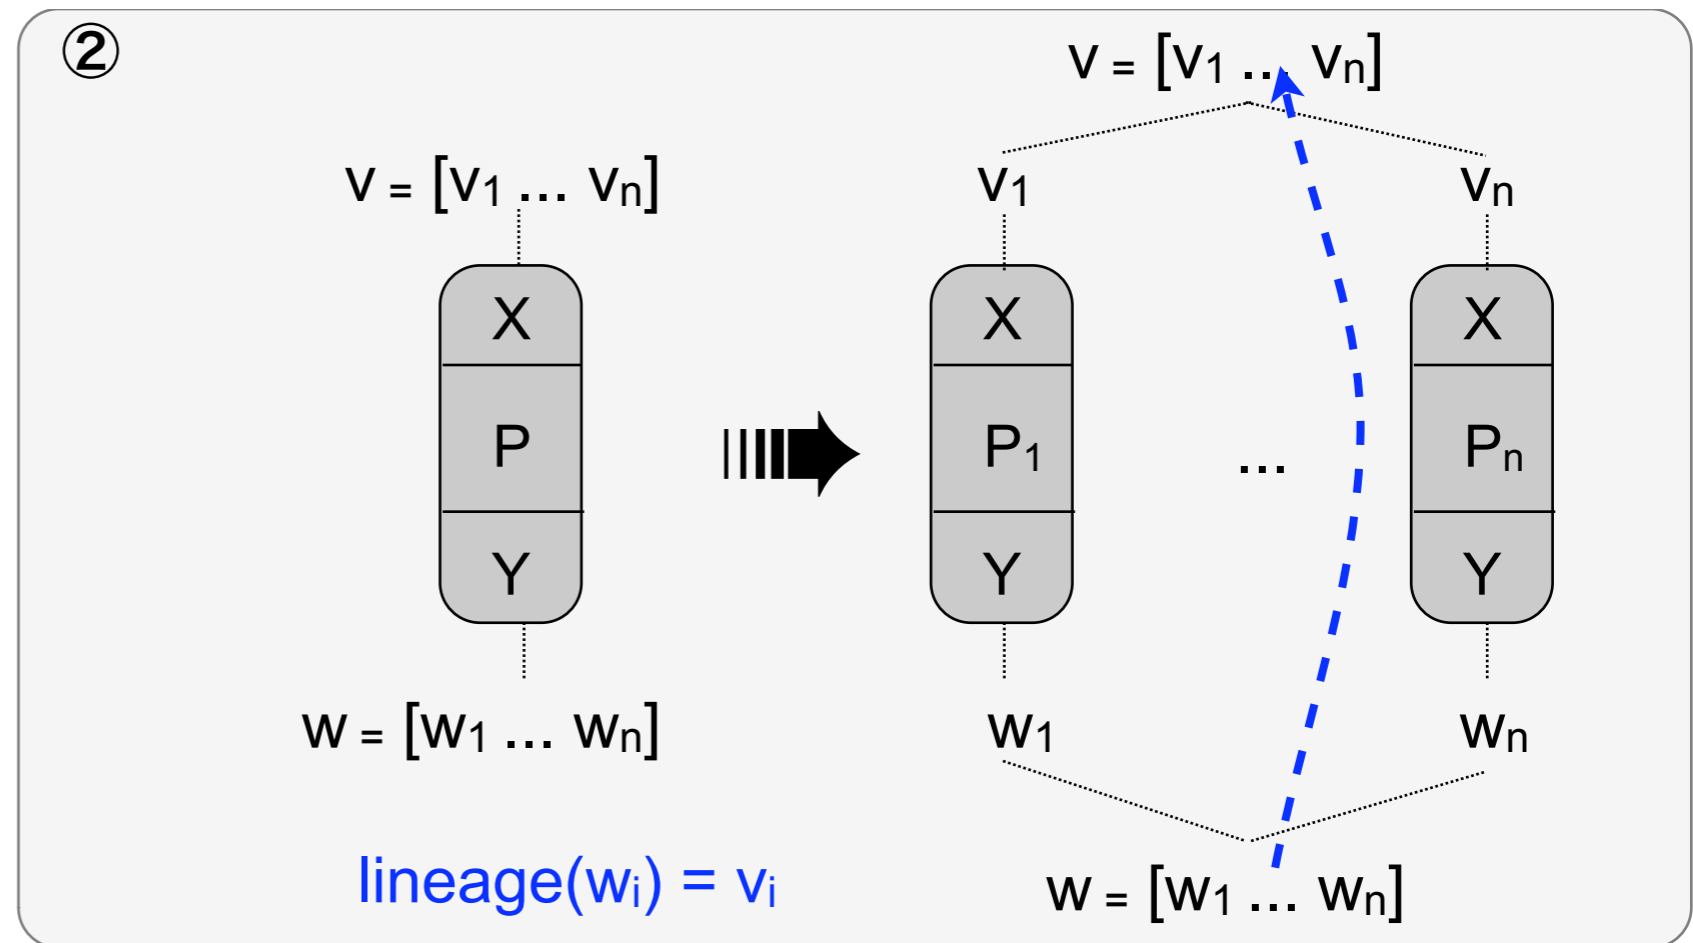



Simple iteration:
service expects atomic values,
receives input **list**



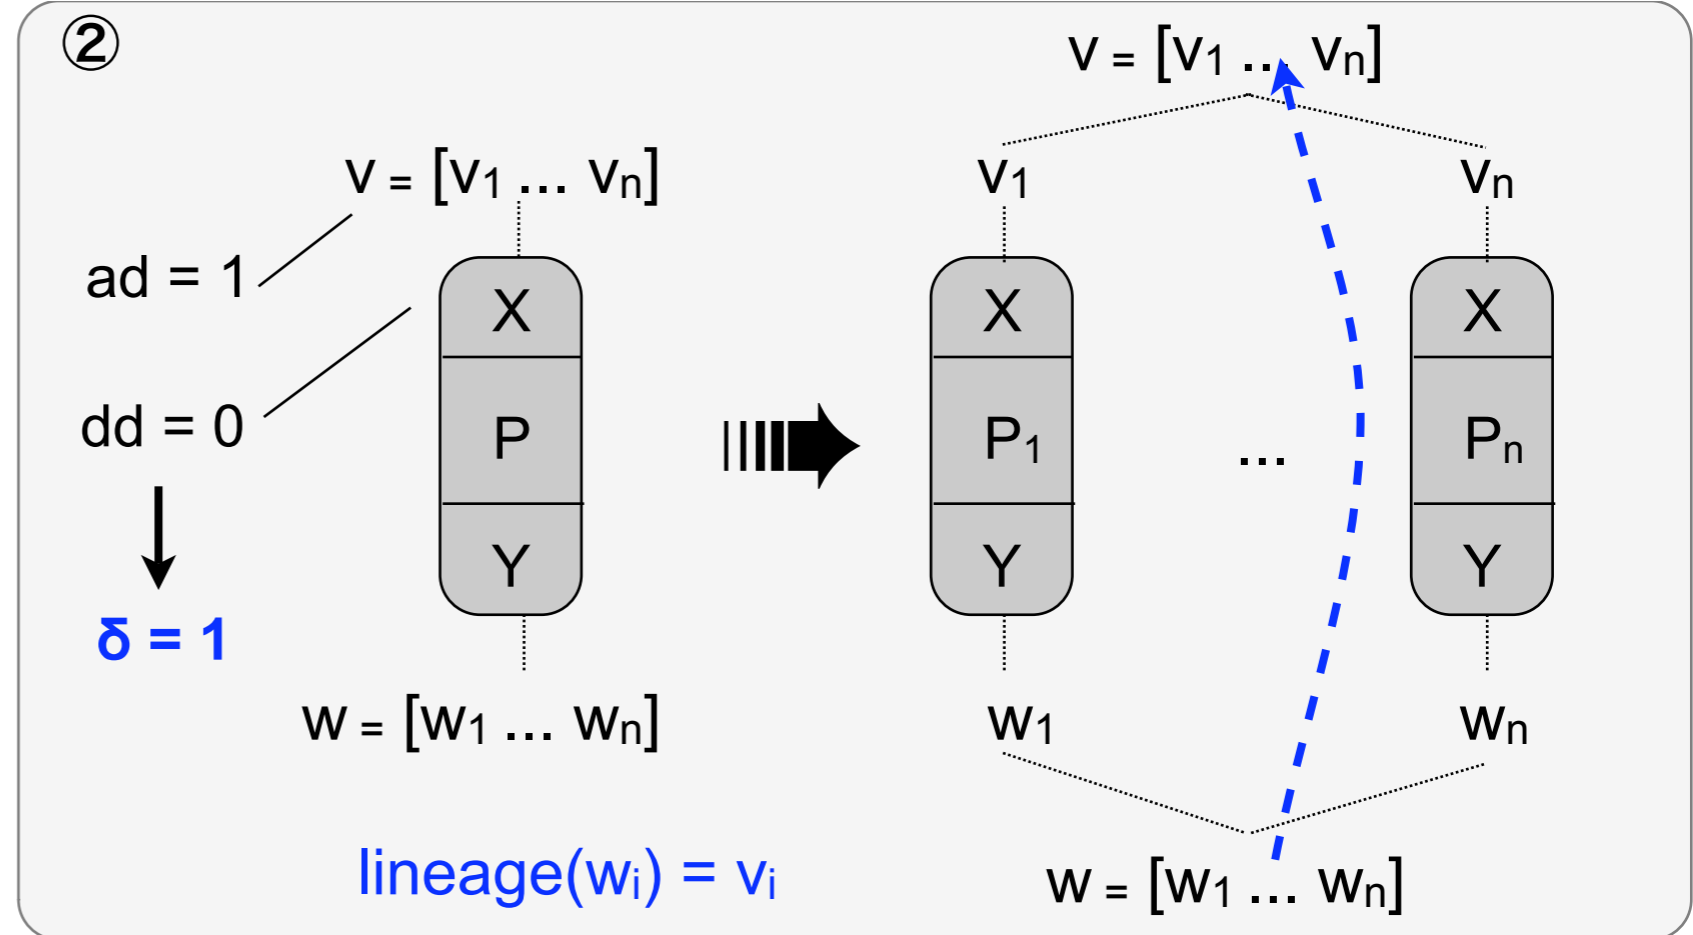

Simple processing:
service expects atomic values,
receives atomic values

Simple iteration:
service expects atomic values,
receives input **list**



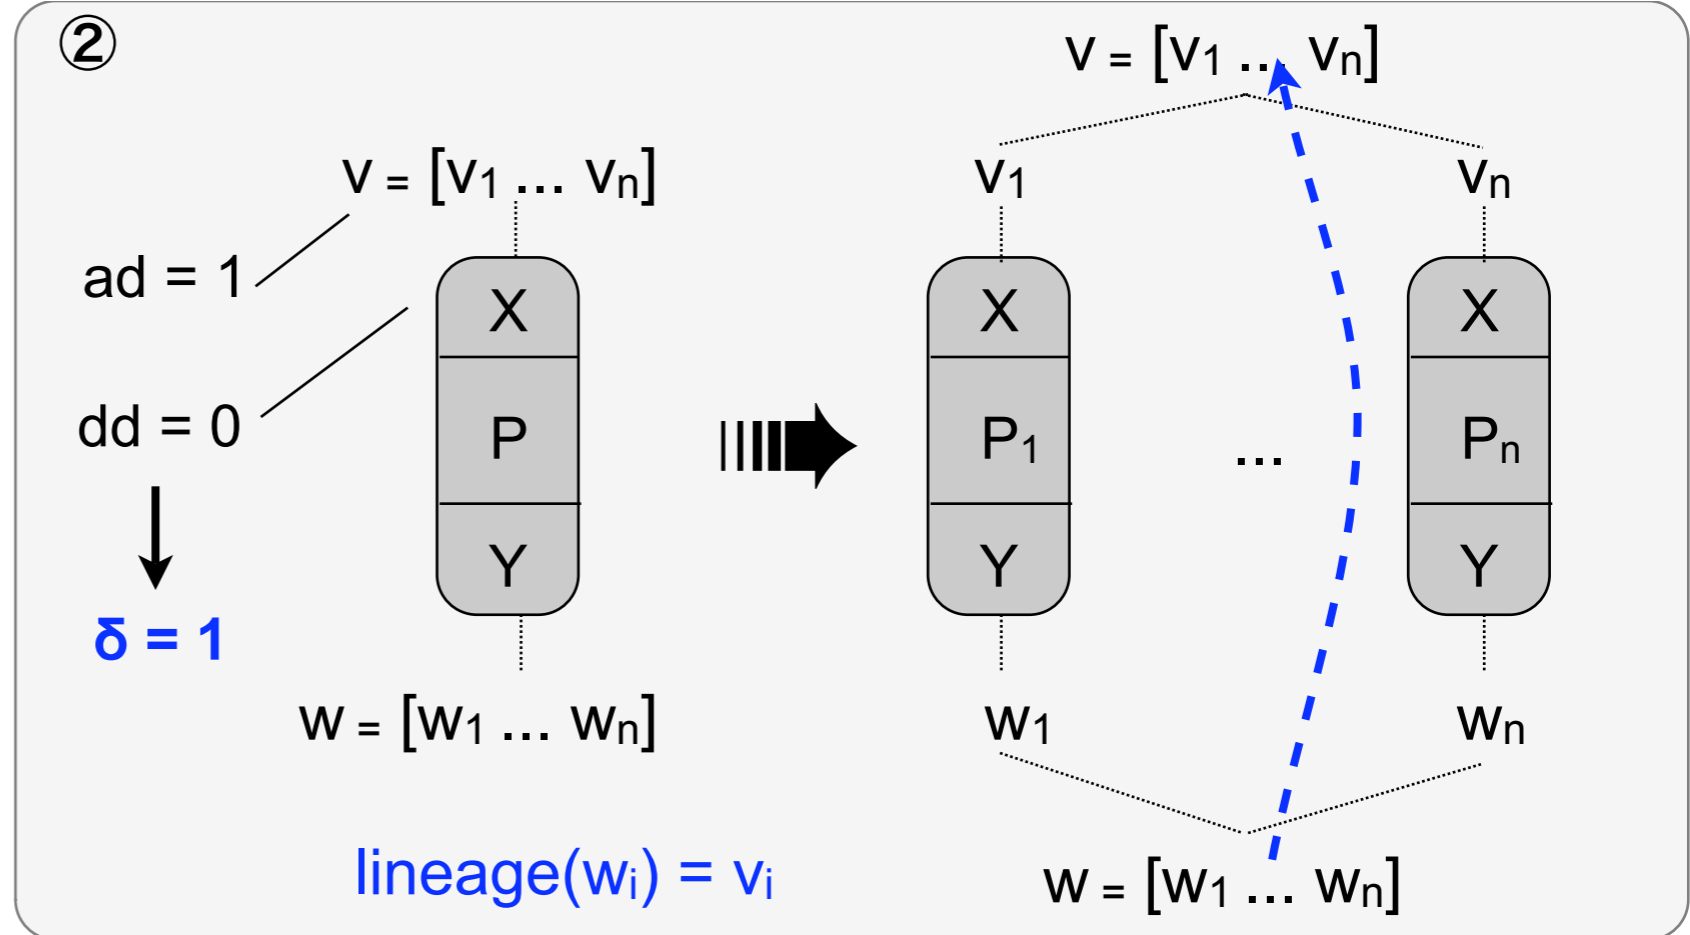

Simple processing:
service expects atomic values,
receives atomic values

Simple iteration:
service expects atomic values,
receives input **list**

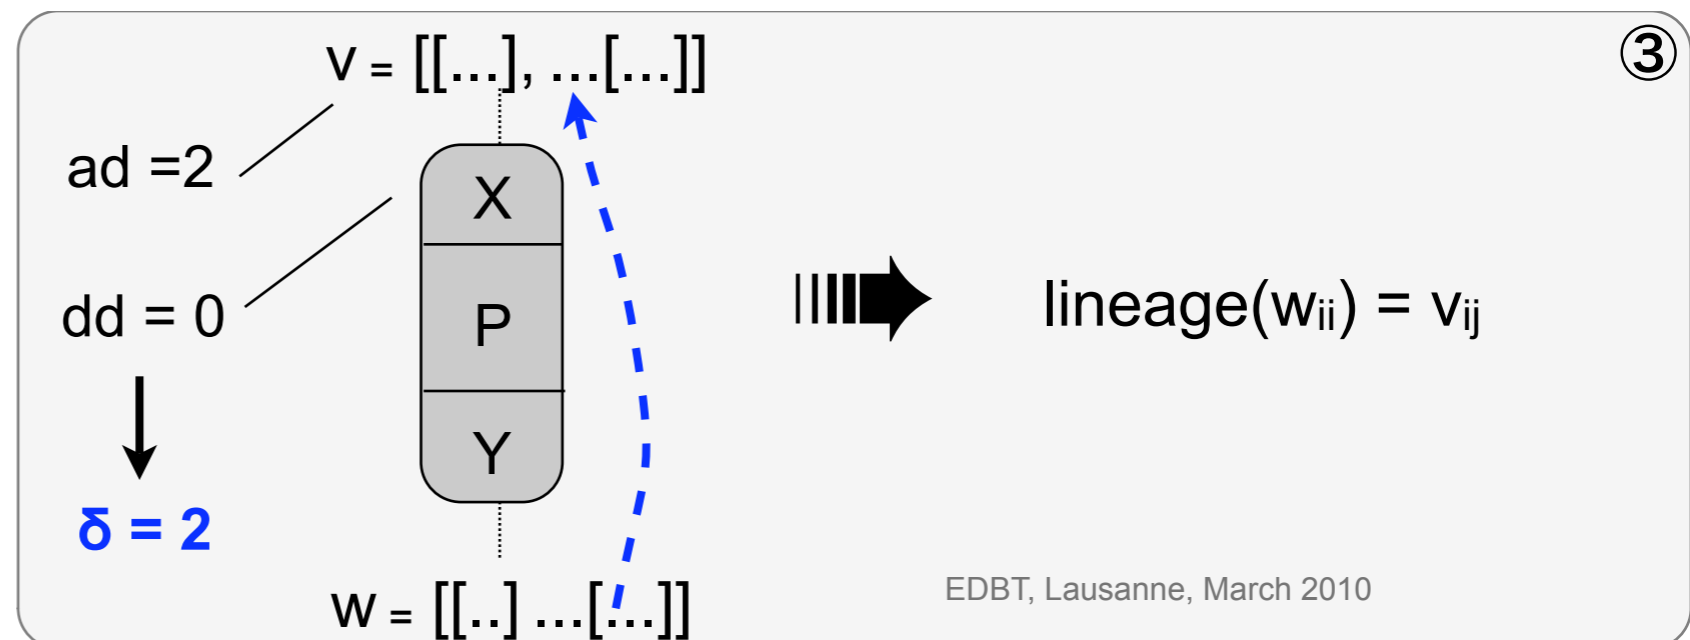


Simple processing:
service expects atomic values,
receives atomic values


Simple iteration:
service expects atomic values,
receives input **list**




Extension:
service expects atomic values,
receives input **nested list**



The simple iteration model generalises by induction to a generic $\delta = n - m$




The simple iteration model generalises by induction to a generic $\delta = n - m$




This leads to a recursive functional formulation for simple collection processing:

$$v = \langle a_1 \dots a_n \rangle$$

$$(\text{eval}_l P v) = \begin{cases} (P v) & \text{if } l = 0 \\ (\text{map } (\text{eval}_{l-1} P) v) & \text{if } l > 0 \end{cases}$$



$dd(X_1) = 0, ad(v_1) = 1 \implies \delta_1 = 1$
$dd(X_2) = 1, ad(v_2) = 1 \implies \delta_2 = 0$
$dd(X_3) = 0, ad(v_3) = 1 \implies \delta_3 = 1$




$dd(X_1) = 0, ad(v_1) = 1 \implies \delta_1 = 1$
$dd(X_2) = 1, ad(v_2) = 1 \implies \delta_2 = 0$
$dd(X_3) = 0, ad(v_3) = 1 \implies \delta_3 = 1$

Cross-product involving v_1 and v_2 (but not v_3):

$$v_1 \otimes v_3 = [[\langle v_{1i}, v_{3j} \rangle \mid j:1..m] \mid i:1..n] \quad // \text{ cross product}$$

$$\text{and including } v_2: [[\langle v_{1i}, v_2, v_{3j} \rangle \mid j:1..m] \mid i:1..n]$$



$dd(X_1) = 0, ad(v_1) = 1$	\implies	$\delta_1 = 1$
$dd(X_2) = 1, ad(v_2) = 1$	\implies	$\delta_2 = 0$
$dd(X_3) = 0, ad(v_3) = 1$	\implies	$\delta_3 = 1$

$$\text{lineage}(w_{ii}) = \langle v_{1i}, v_2, v_{3j} \rangle$$

Cross-product involving v_1 and v_2 (but not v_3):

$$v_1 \otimes v_3 = [[\langle v_{1i}, v_{3j} \rangle \mid j:1..m] \mid i:1..n] \quad // \text{ cross product}$$

$$\text{and including } v_2: [[\langle v_{1i}, v_2, v_{3j} \rangle \mid j:1..m] \mid i:1..n]$$

Binary product, $\delta = 1$:

$$a \times b = [[\langle a_i, b_j \rangle] | b_j \leftarrow b] | a_i \leftarrow a]$$

$$(\text{eval}_2 P \langle a, b \rangle) = (\text{map} (\text{eval}_1 P) a \times b)$$

Binary product, $\delta = 1$:

$$a \times b = [[\langle a_i, b_j \rangle] | b_j \leftarrow b] | a_i \leftarrow a]$$

$$(\text{eval}_2 P \langle a, b \rangle) = (\text{map} (\text{eval}_1 P) a \times b)$$

Generalized to arbitrary depths:

$$(v, d_1) \otimes (w, d_2) = \begin{cases} [[(v_i, w_j) | w_j \leftarrow w] | v_i \leftarrow v] & \text{if } d_1 > 0, d_2 > 0 \\ [(v_i, w) | v_i \leftarrow v] & \text{if } d_1 > 0, d_2 = 0 \\ [(v, w_j) | w_j \leftarrow w] & \text{if } d_1 = 0, d_2 > 0 \\ (v, w) & \text{if } d_1 = 0, d_2 = 0 \end{cases}$$

...and to n operands: $\otimes_{i:1\dots n} (v_i, d_i)$

Binary product, $\delta = 1$:

$$a \times b = [[\langle a_i, b_j \rangle] | b_j \leftarrow b] | a_i \leftarrow a]$$

$$(\text{eval}_2 P \langle a, b \rangle) = (\text{map} (\text{eval}_1 P) a \times b)$$

Generalized to arbitrary depths:

$$(v, d_1) \otimes (w, d_2) = \begin{cases} [[(v_i, w_j) | w_j \leftarrow w] | v_i \leftarrow v] & \text{if } d_1 > 0, d_2 > 0 \\ [(v_i, w) | v_i \leftarrow v] & \text{if } d_1 > 0, d_2 = 0 \\ [(v, w_j) | w_j \leftarrow w] & \text{if } d_1 = 0, d_2 > 0 \\ (v, w) & \text{if } d_1 = 0, d_2 = 0 \end{cases}$$

...and to n operands: $\otimes_{i:1\dots n} (v_i, d_i)$

Finally: general functional semantics for collection-based processing

$$\begin{aligned} & (\text{eval}_l P \langle (v_1, d_1), \dots, (v_n, d_n) \rangle) \\ &= \begin{cases} (P \langle v_1, \dots, v_n \rangle) & \text{if } l = 0 \\ (\text{map} (\text{eval}_{l-1} P) \otimes_{i:1\dots n} \langle v_i, d_i \rangle) & \text{if } l > 0 \end{cases} \end{aligned}$$

The iteration structure can be determined **statically**

X_1	X_2	X_3
P		
Y		

$$dd(X_1) = 0, \text{ ad}(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \text{ ad}(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \text{ ad}(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$$[i_1 . i_2 . \dots . i_k] = \underline{\hspace{15em}}$$

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

$$dd(X_1) = 0, \quad ad(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$$[i_1 . i_2 . \dots . i_k] = \underline{\hspace{15em}}$$

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

$$dd(X_1) = 0, \quad ad(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$$[i_1 . i_2 . \dots . i_k] = \frac{\delta_1}{\dots}$$

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

$$dd(X_1) = 0, \quad ad(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$[i_1 . i_2 . \dots . i_k] =$

X_1

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

$$dd(X_1) = 0, \quad ad(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$$[i_1 \ . \ i_2 \ . \ \dots \ . \ i_k] = \frac{\delta_2}{X_1}$$

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

$$dd(X_1) = 0, \quad ad(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$$[i_1 . i_2 . \dots . i_k] =$$

$$\begin{array}{c} \text{-----} \\ X_1 \end{array} \quad \begin{array}{c} \text{-----} \\ X_2 \end{array}$$

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}

$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$[i_1 . i_2 . \dots . i_k] =$

$\frac{\quad}{X_1} \quad \frac{\quad}{X_2}$

$\frac{\delta_k}{\quad}$

$$dd(X_1) = 0, \quad ad(v_1) = 1 \quad \Rightarrow \quad \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \quad \Rightarrow \quad \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \quad \Rightarrow \quad \delta_3 = 1$$

The iteration structure can be determined **statically**

(0,1) (1,1) (0,1)

X_1	X_2	X_3
P		
Y		

(0,2)

$$dd(X_1) = 0, \quad ad(v_1) = 1 \implies \delta_1 = 1$$

$$dd(X_2) = 1, \quad ad(v_2) = 1 \implies \delta_2 = 0$$

$$dd(X_3) = 0, \quad ad(v_3) = 1 \implies \delta_3 = 1$$

this leads to a simple mapping rule:

index of an output list value \rightarrow {index of input values}


$$Y[i..j] \rightarrow X_1[i], X_2[], X_3[j]$$

$$[i_1 \ . \ i_2 \ . \ \dots \ . \ i_k] =$$




1. traverse the workflow graph (small) rather than the provenance trace (large)
2. use static prediction of iterations to trace through collection elements
3. “parachute” into the actual trace only at the end

Workflow graph




$$y = \begin{bmatrix} [y_{11} \dots y_{1n}], \\ \dots \\ [y_{m1} \dots y_{mn}] \end{bmatrix}$$

Provenance graph




1. traverse the workflow graph (small) rather than the provenance trace (large)
2. use static prediction of iterations to trace through collection elements
3. “parachute” into the actual trace only at the end

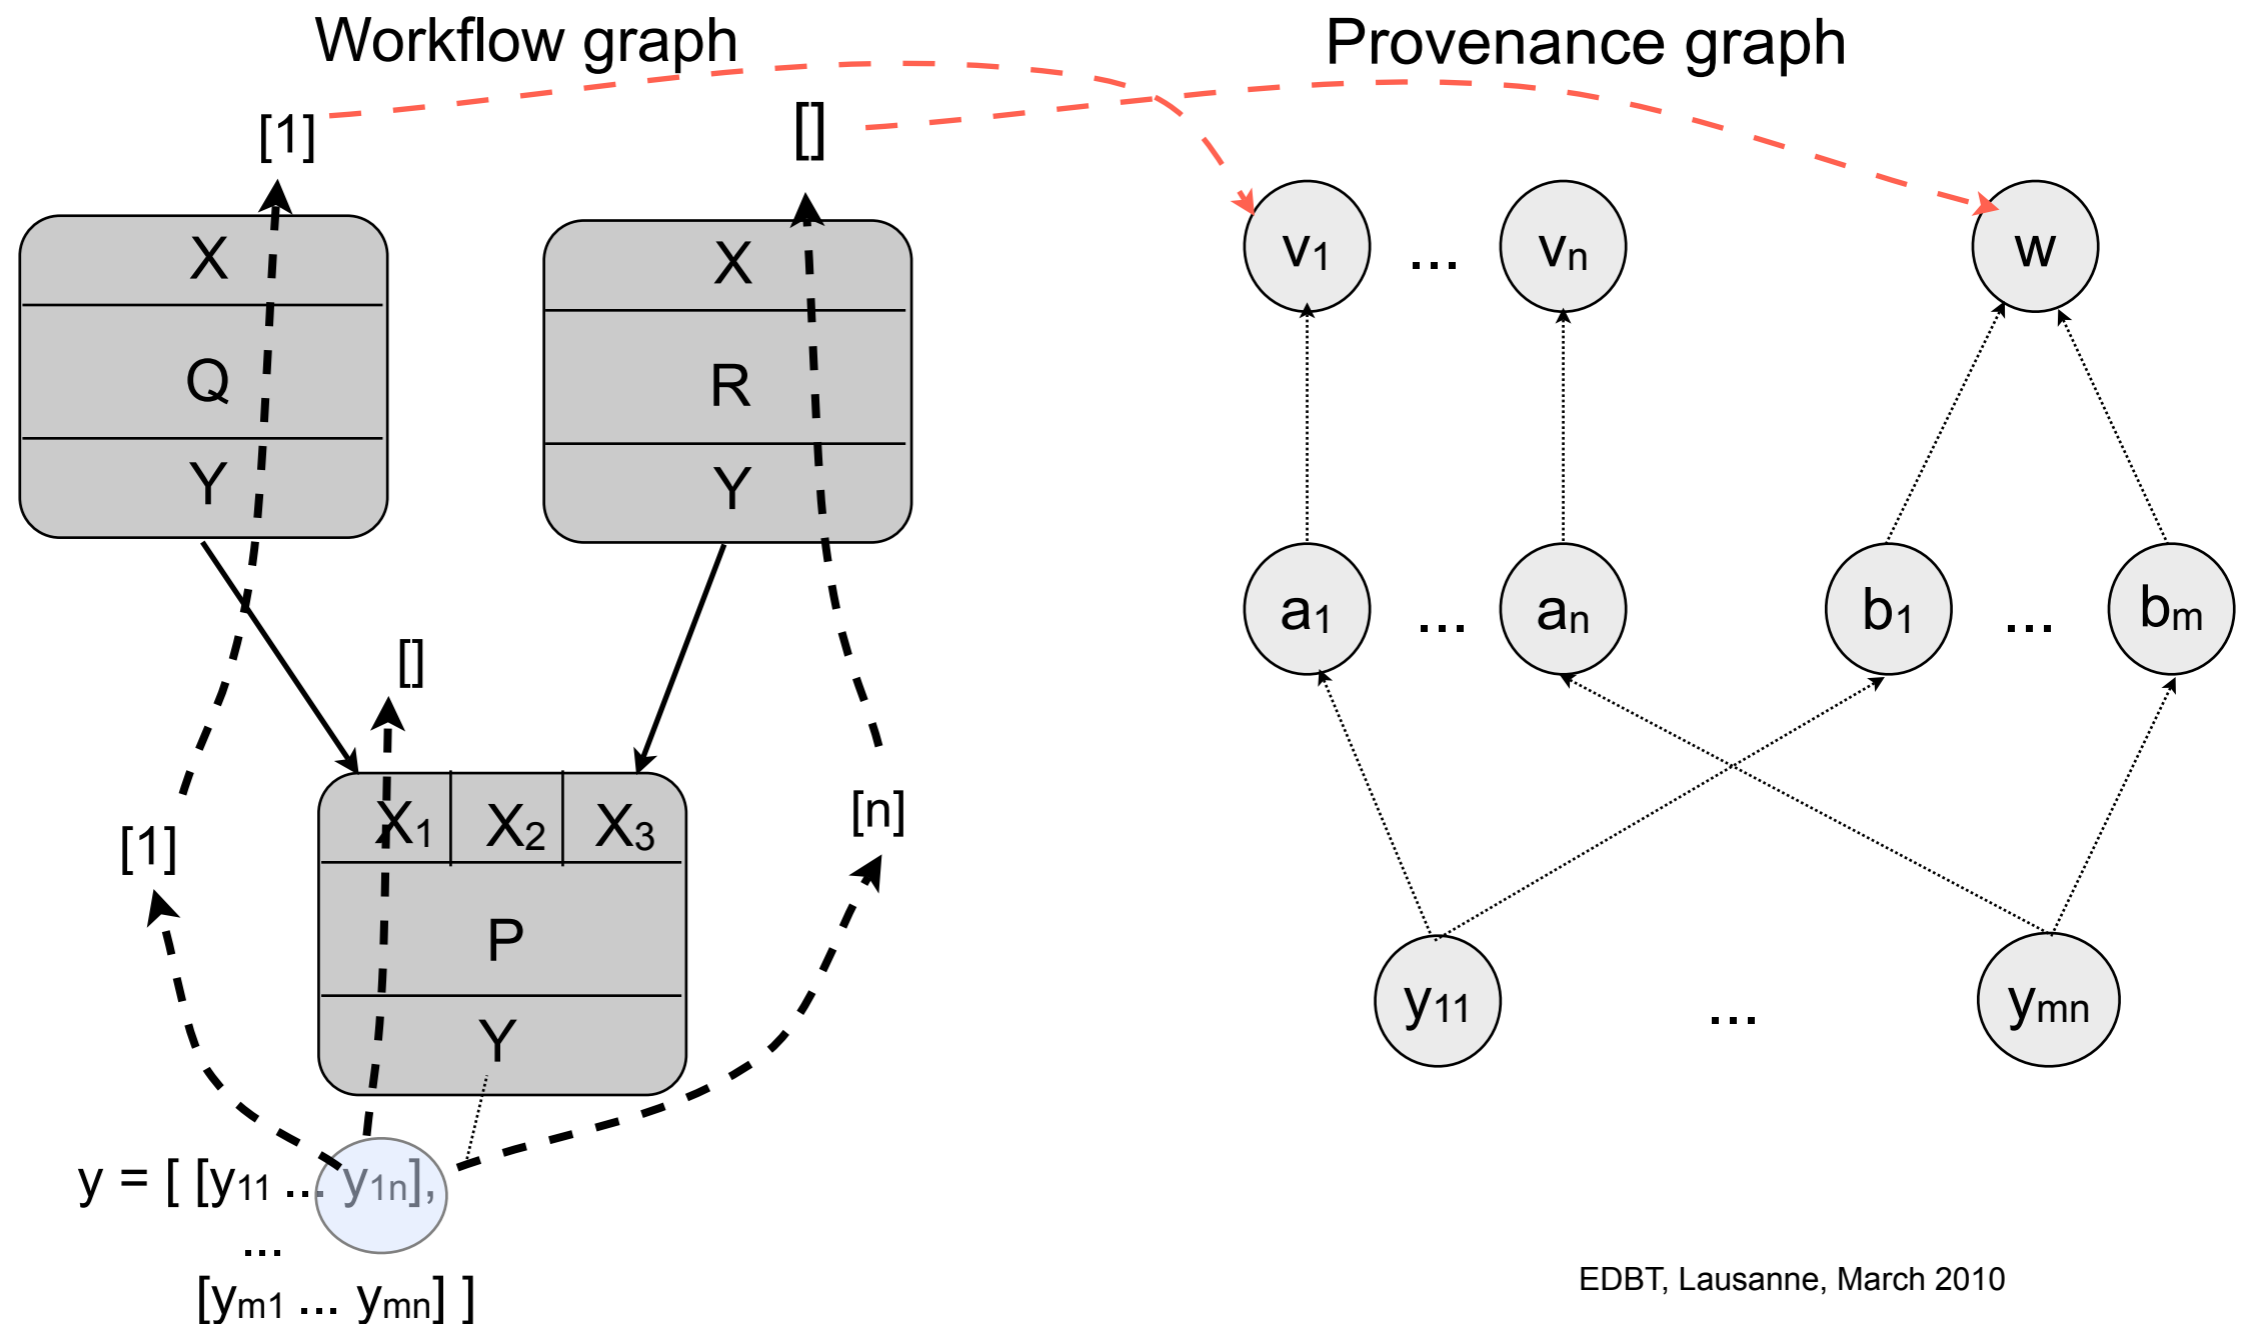
Workflow graph



Provenance graph



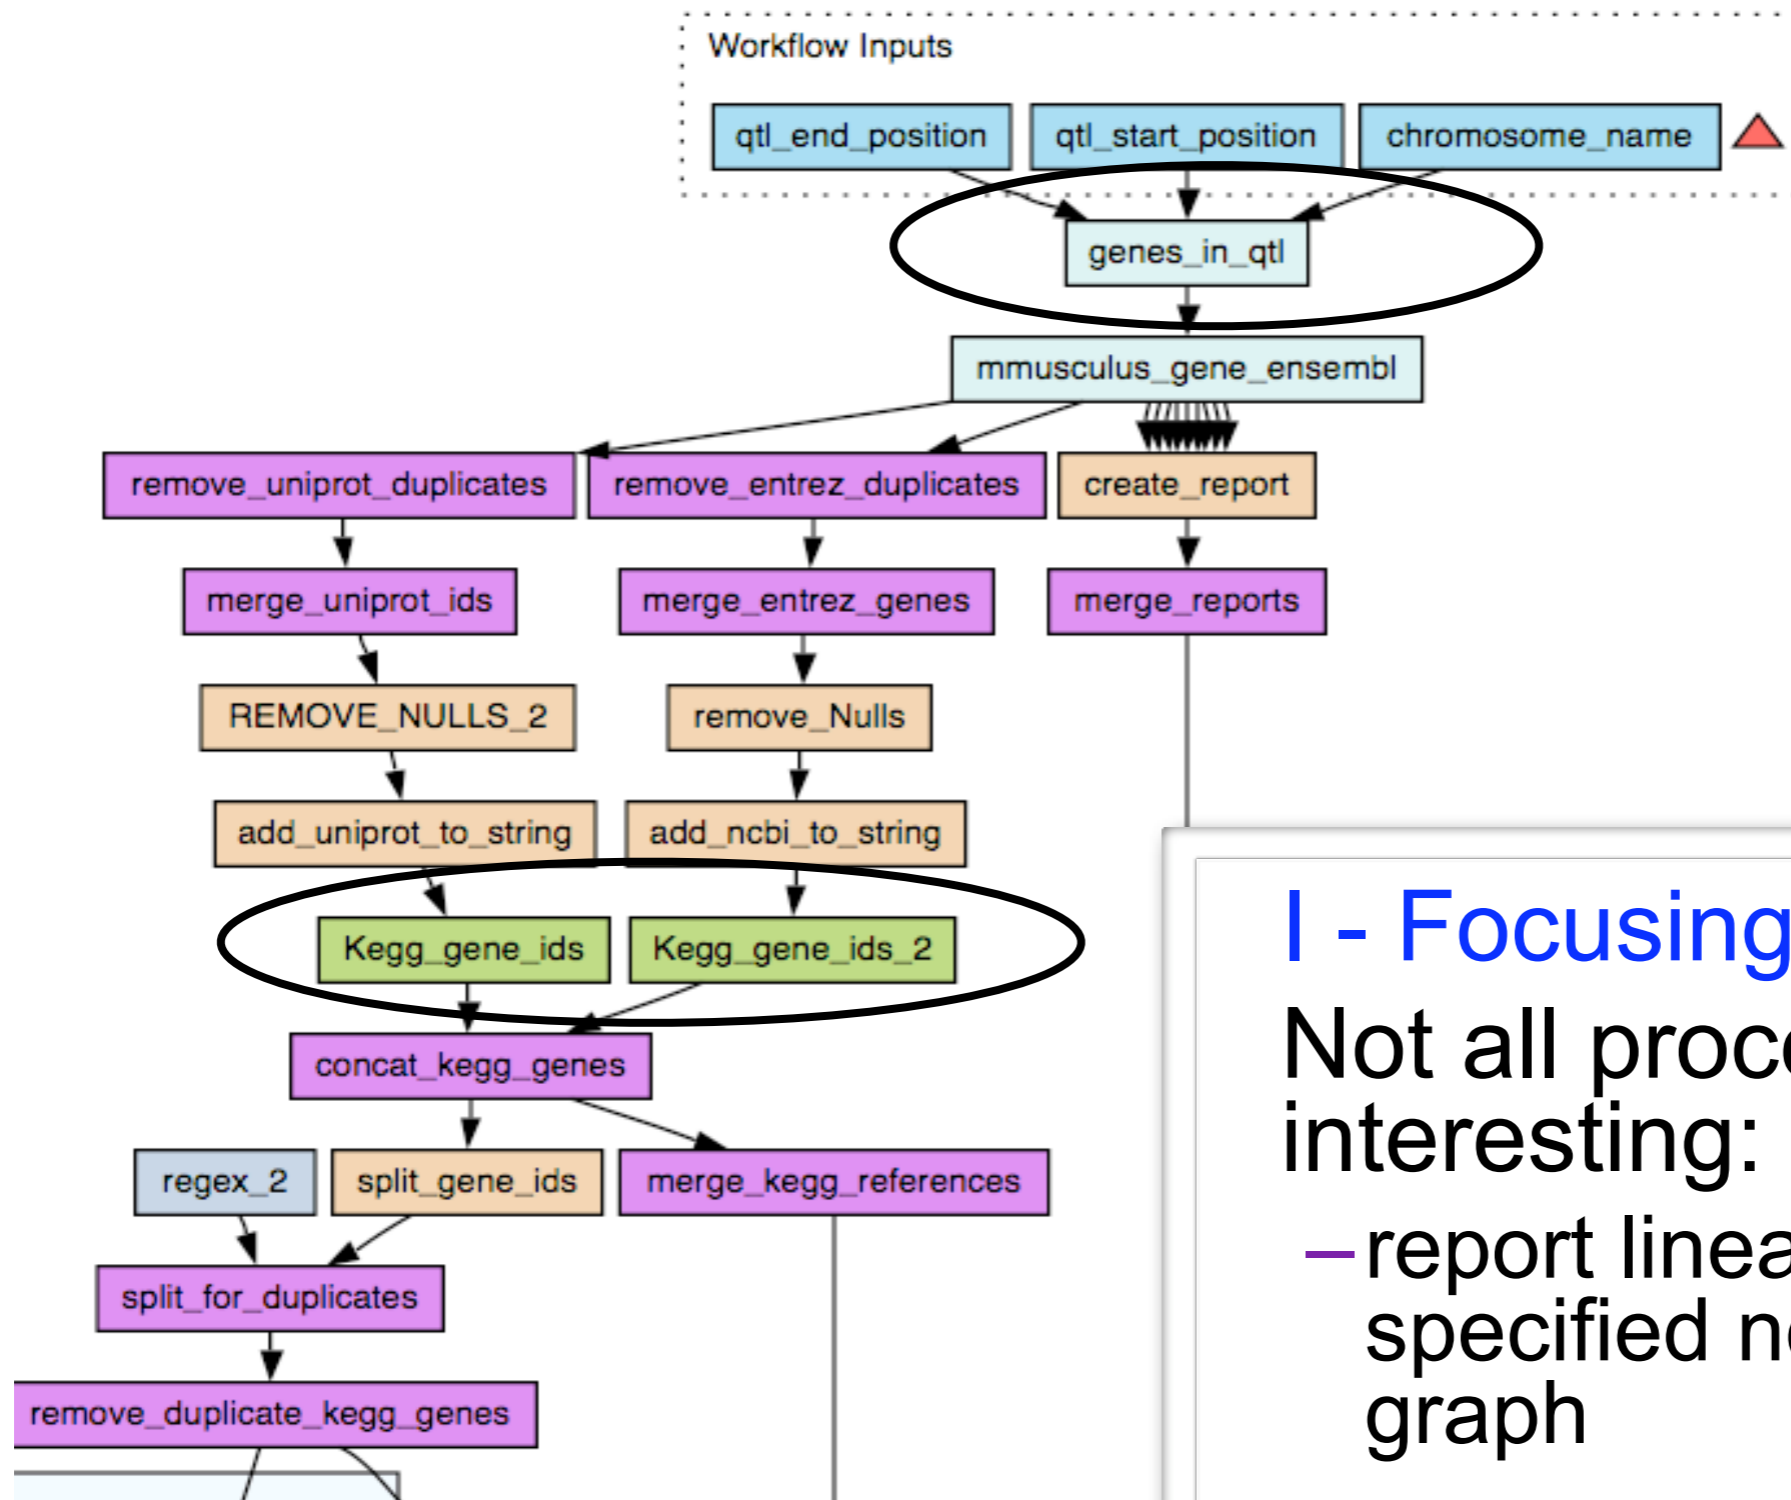
1. traverse the workflow graph (small) rather than the provenance trace (large)
2. use static prediction of iterations to trace through collection elements
3. “parachute” into the actual trace only at the end



Summary so far:

- whenever iterations are involved, we can trace the provenance of individual elements of a processor's output
- iterations are explained in terms of a functional model and based on list depth discrepancies
- The relationships between output and input indexes are derived using the workflow specification graph (statically)

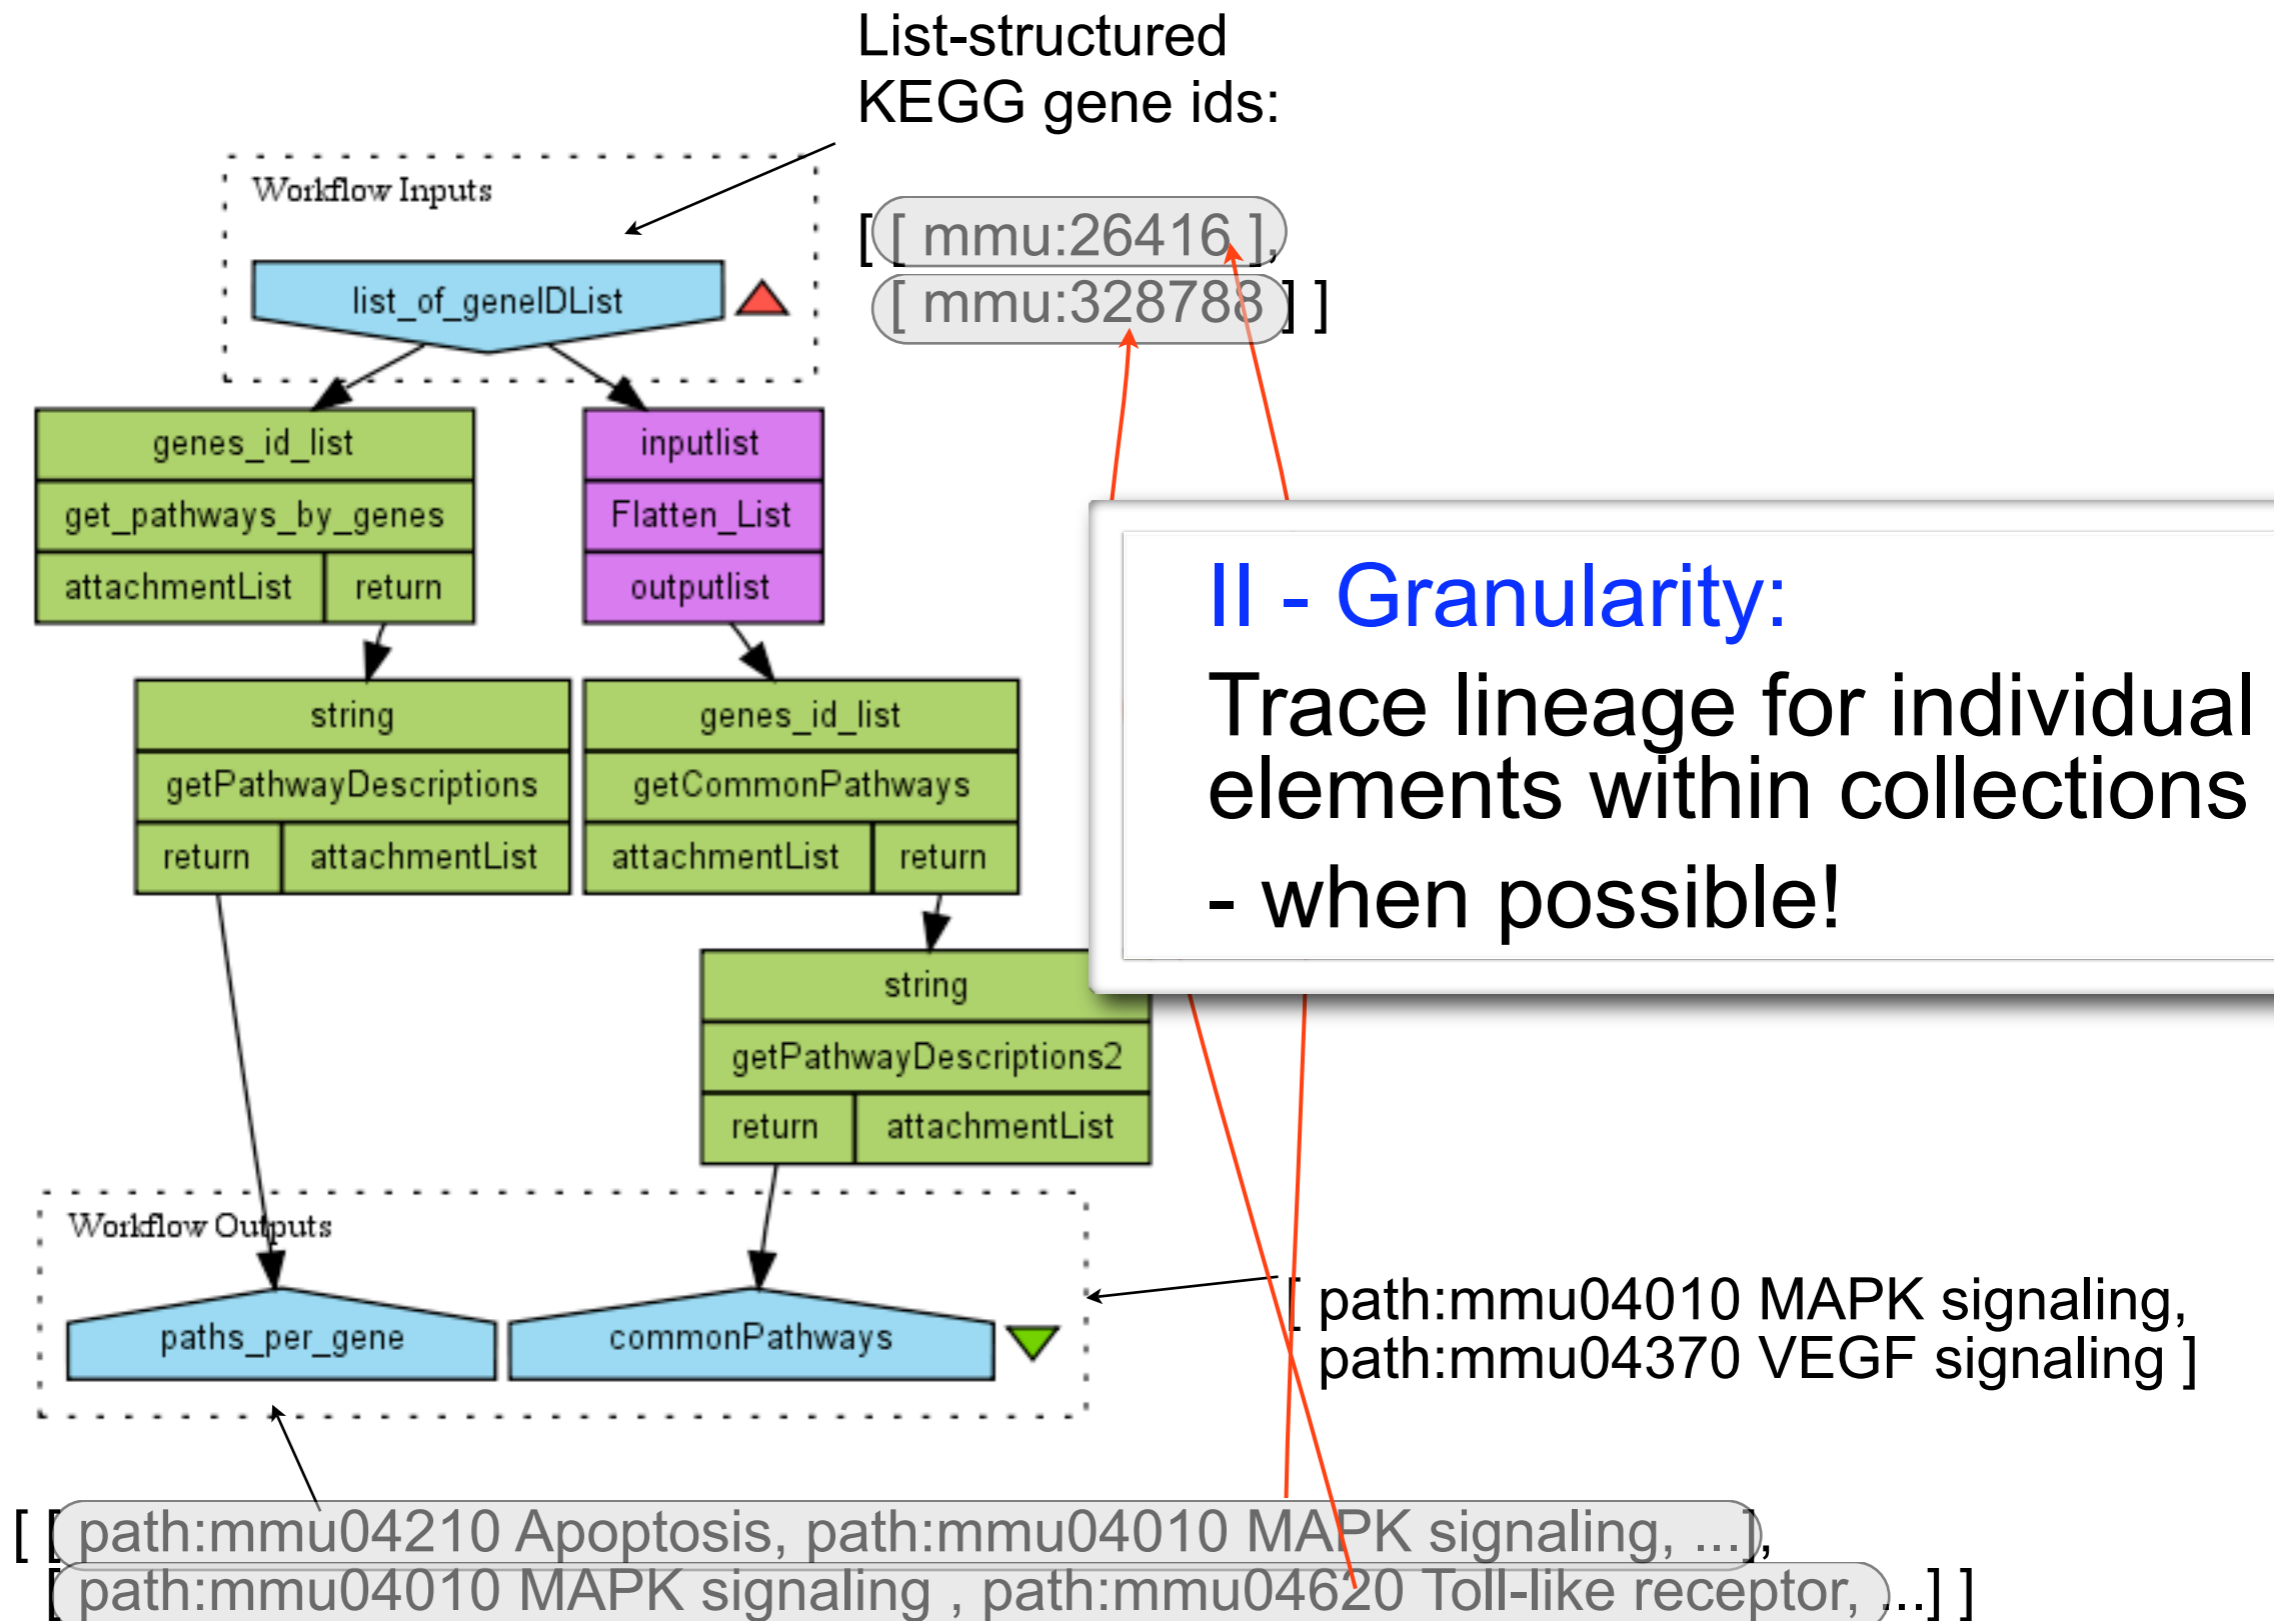
How about expressivity and efficient processing of lineage queries?

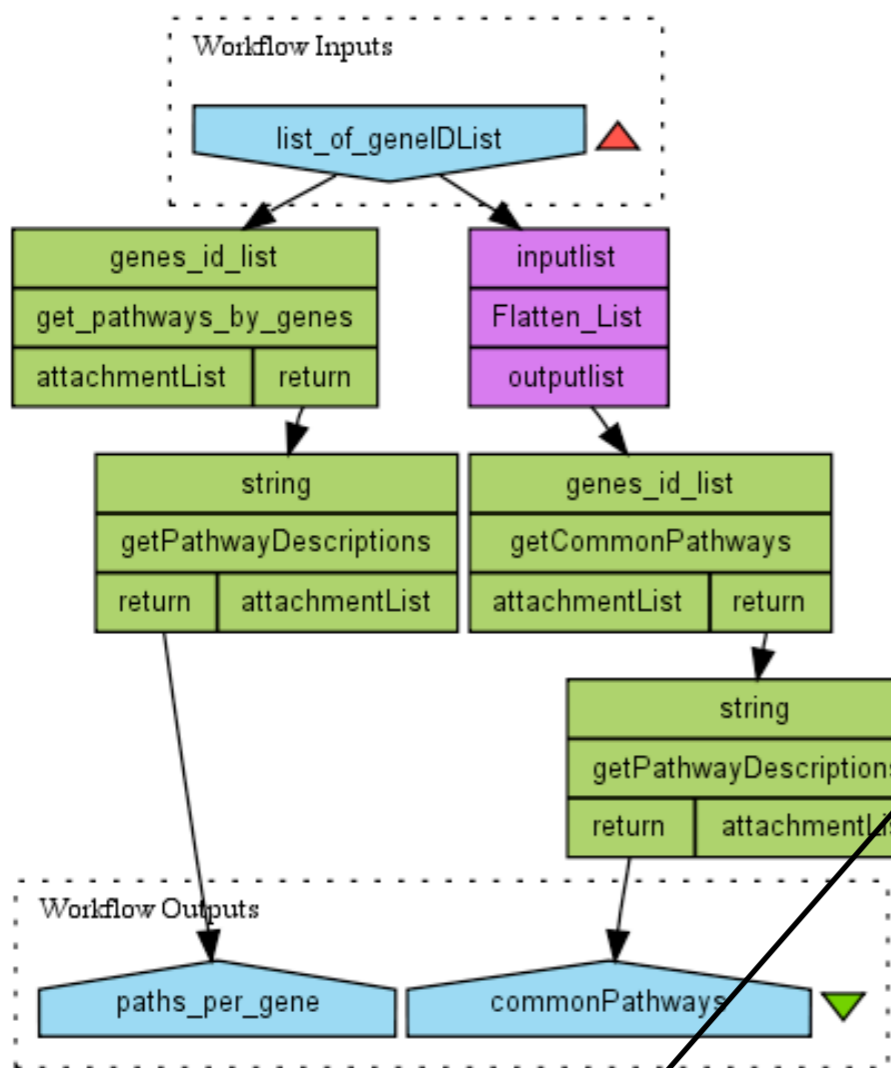


I - Focusing:

Not all processors are interesting:

- report lineage only at specified nodes in the graph





workflow scope defaults to latest run

optionally specifies one or more runs for the target workflow

```

<pquery>
  <scope workflow="keggPathways">
    <run id="ae1e2b6b-3bc5-4c93-a250-c4dd0210c3b3" />
  </scope>
  <select>
    <outputPort name="paths_per_gene" index="[1,2]" />
    <outputPort name="paths_per_gene" index="[3,4]" />
    <outputPort name="commonPathways" index="[1]" />
    <processor name="getPathwayDescriptions">
      <outputPort name="return" />
    </processor>
  </select>
  <focus>
    <processor name="get_pathway_by_genes" />
  </focus>
</pquery>

```

port values for which lineage is sought:
global outputs or processor-qualified

processors where lineage is to be reported
- possibly workflow-qualified

- **Scalability:**
 - query time depends on size of workflow graph, not size of provenance graph
 - workflow graphs are small, fit in memory, can be indexed easily
- **Graceful degradation:**
 - worst case is a completely unfocused query
 - no worse than other approaches
- Fine-grain answers provided at the same time

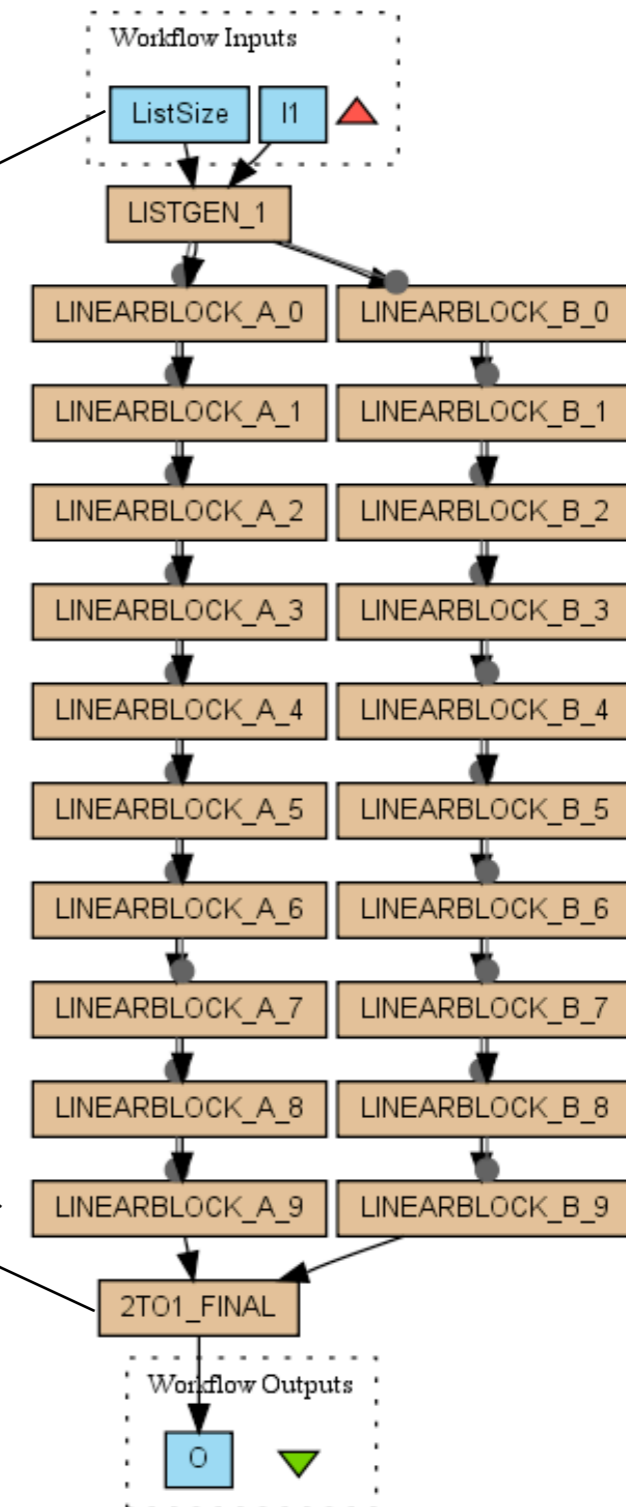
- Assumption:
 - *Black box* provenance of workflow data products
 - ✓ Fine-grained provenance:
 - tracking provenance through collection elements
 - motivation, **functional model** of collection-oriented workflow processing
 - ✓ Efficient query processing:
 - ✓ leveraging the functional model to achieve efficient processing for a simple query model
- ➔ **Experimental evaluation**


- Performance evaluation performed on programmatically generated dataflows

– the “T-towers”

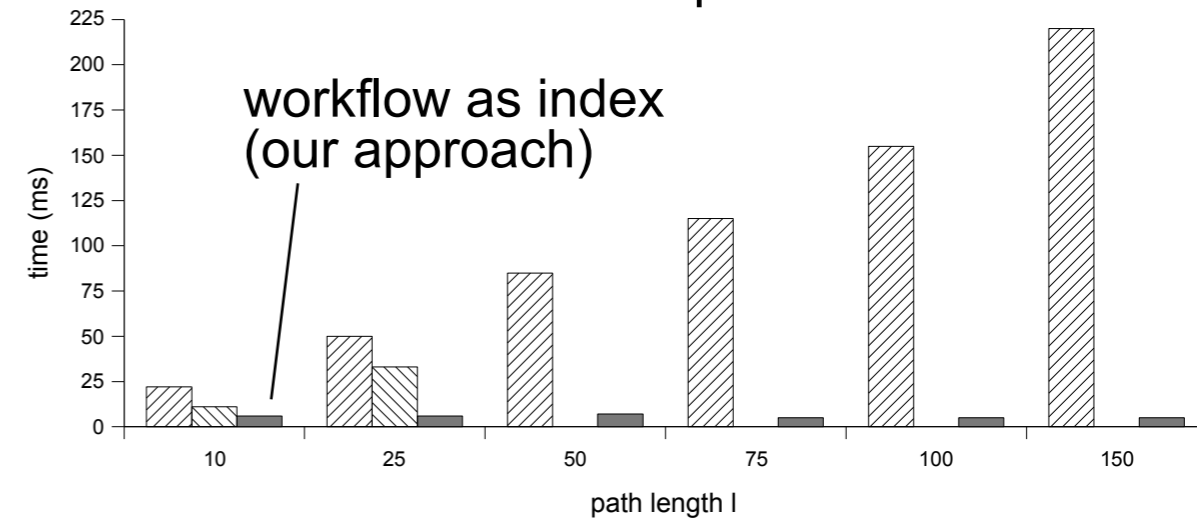
parameters:

- size of the lists involved
- length of the paths
- includes one cross product







Naive traversal of provenance graph



workflow pre-processing time by graph size



performance degradation on fully unfocused queries



- A simple lineage query model for Taverna
 - grounded in the semantics of collection-oriented processing
 - combines fine-grain answers with efficient query processing
- Ongoing work:
 - space compression, indexing
 - QLP?
 - semantic provenance (initial paper submitted)
- Currently part of the Taverna 2.1 release