


Database system implementation

1. Introduction

http://sok.susu.ru/e_index.html


[\[Russian Homepage\]](#) [\[Research Interests\]](#) [\[Bio Sketch\]](#) [\[Selected Publications\]](#) [\[Recent Talks\]](#)

 **Doctor of Science (Phys.&Math.), Prof.**
Leonid B. Sokolinsky
[South Ural State University](#) Faculty of Computational Mathematics & Informatics
*Lenin prospekt 76,
Chelyabinsk 454080, Russia*
Phone: +7(351) 272-35-00.
✉ Leonid.Sokolinsky@susu.ru
🌐 <http://sok.susu.ru/>
ACM V-Card: <http://member.acm.org/~sokolinsky>

⚠ **Attention! The contents of English and Russian pages are different.**

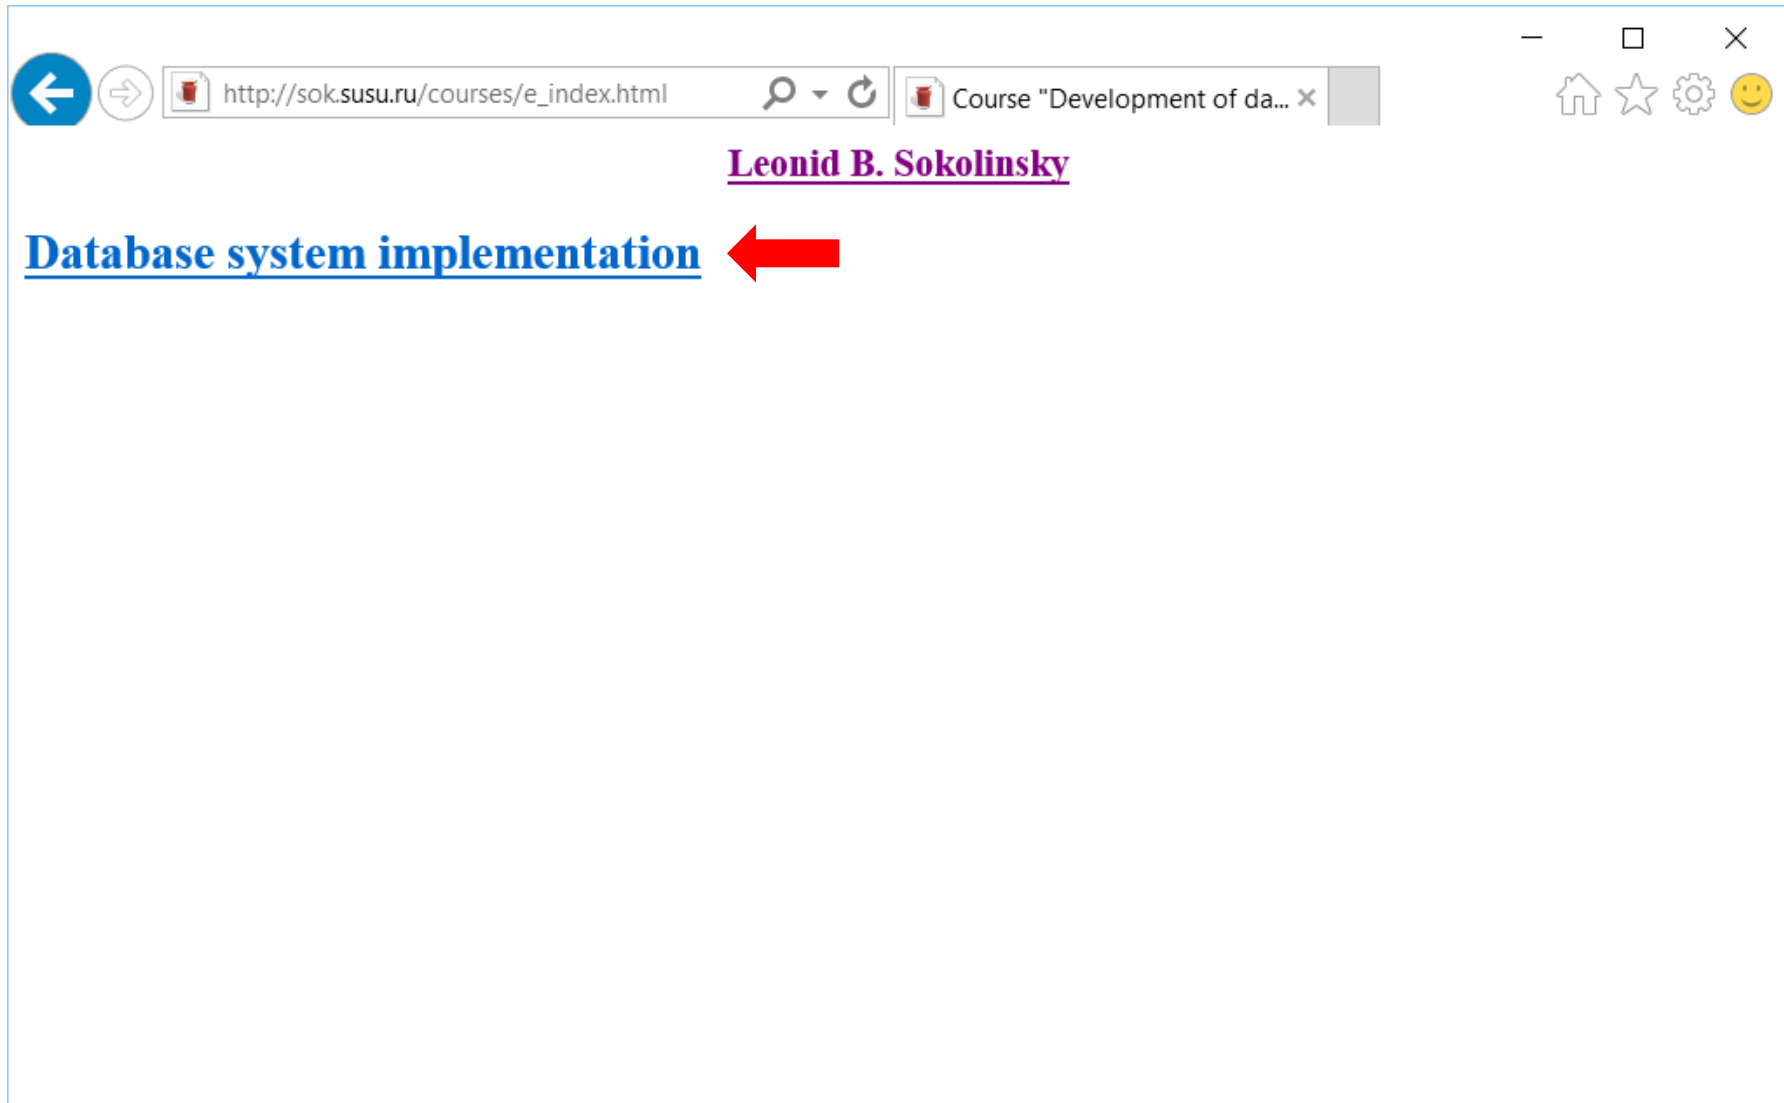
Leonid Sokolinsky is the Vice-Rector for Informatization of [South Ural State University \(National Research University\)](#), doctor of science (phys.&math.), professor, honorary teacher of higher professional education of Russian Federation. Leonid Sokolinsky published more than 100 scientific papers. Leonid Sokolinsky is a member of the [ACM](#) since 1997. Research interests: parallel computing, database management systems, software for supercomputers, cloud computing, e-learning, mathematical programming. Leonid Sokolinsky maintains:

- [Database Systems & Systems Programming Bibliographic Catalogue](#) (in Russian).

● [Research interests](#)
● [Bio sketch](#)
● [Selected publications](#)
● [Teaching courses](#) 
● [Favorite WWW links](#)

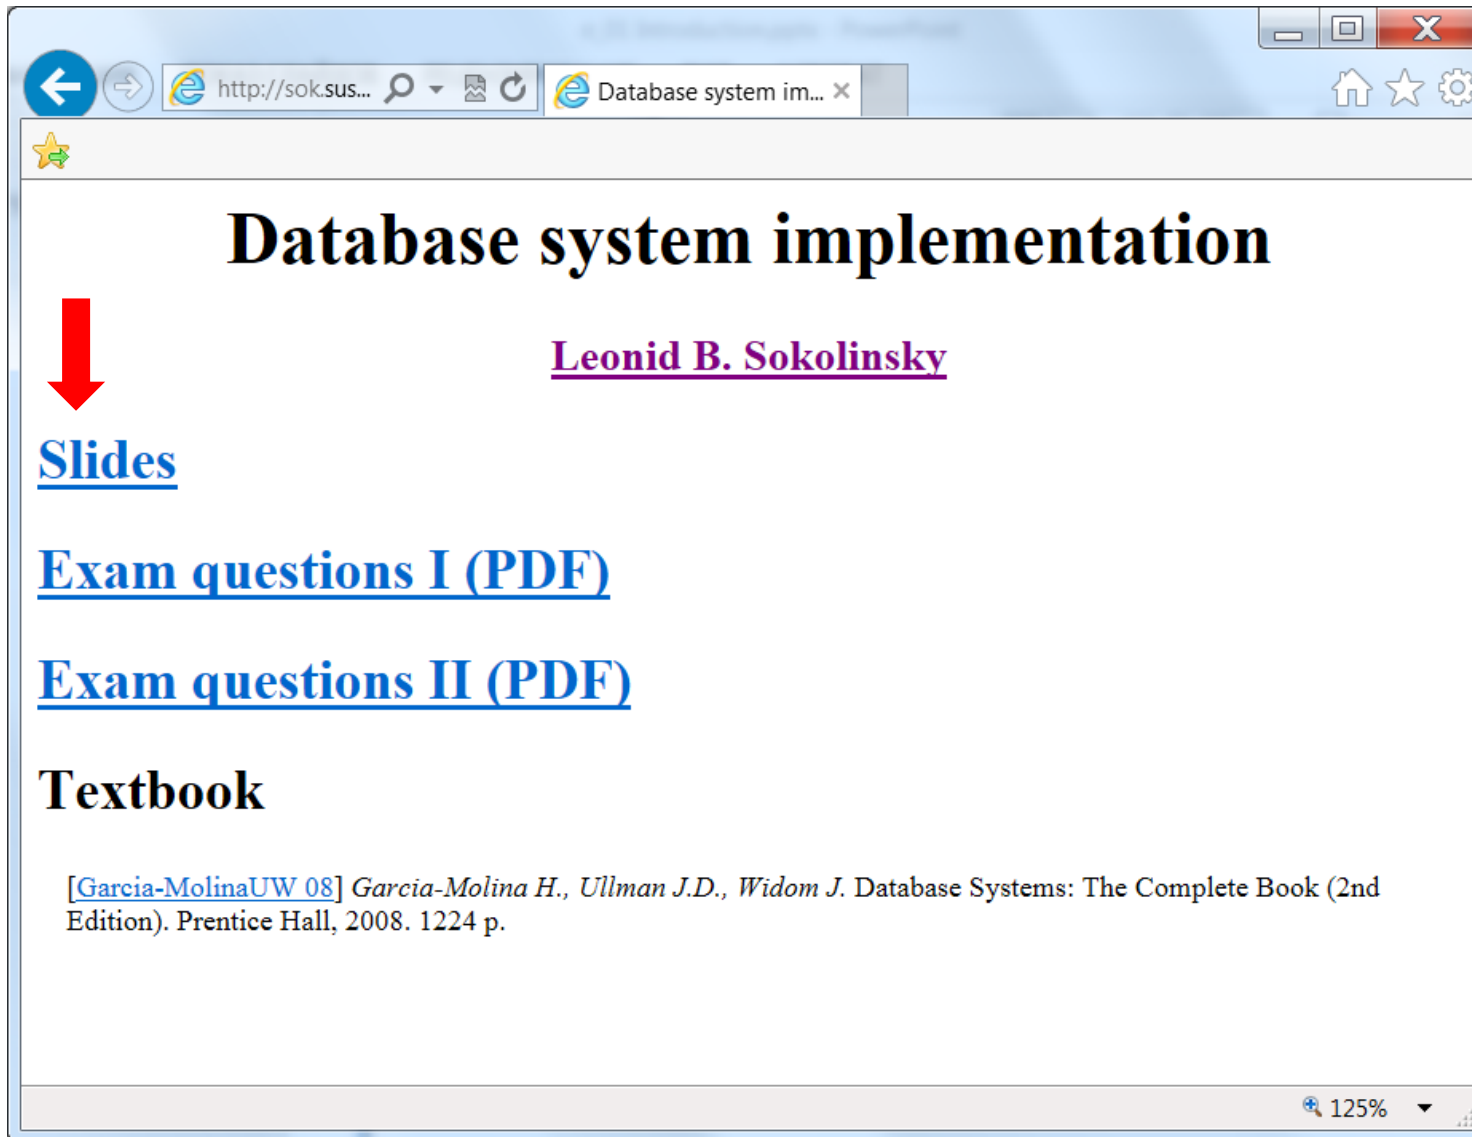
Course «Development of database management systems»

http://sok.susu.ru/courses/e_index.html




Part 1. Database system implementation

http://sok.susu.ru/courses/QueryProc/e_index.html



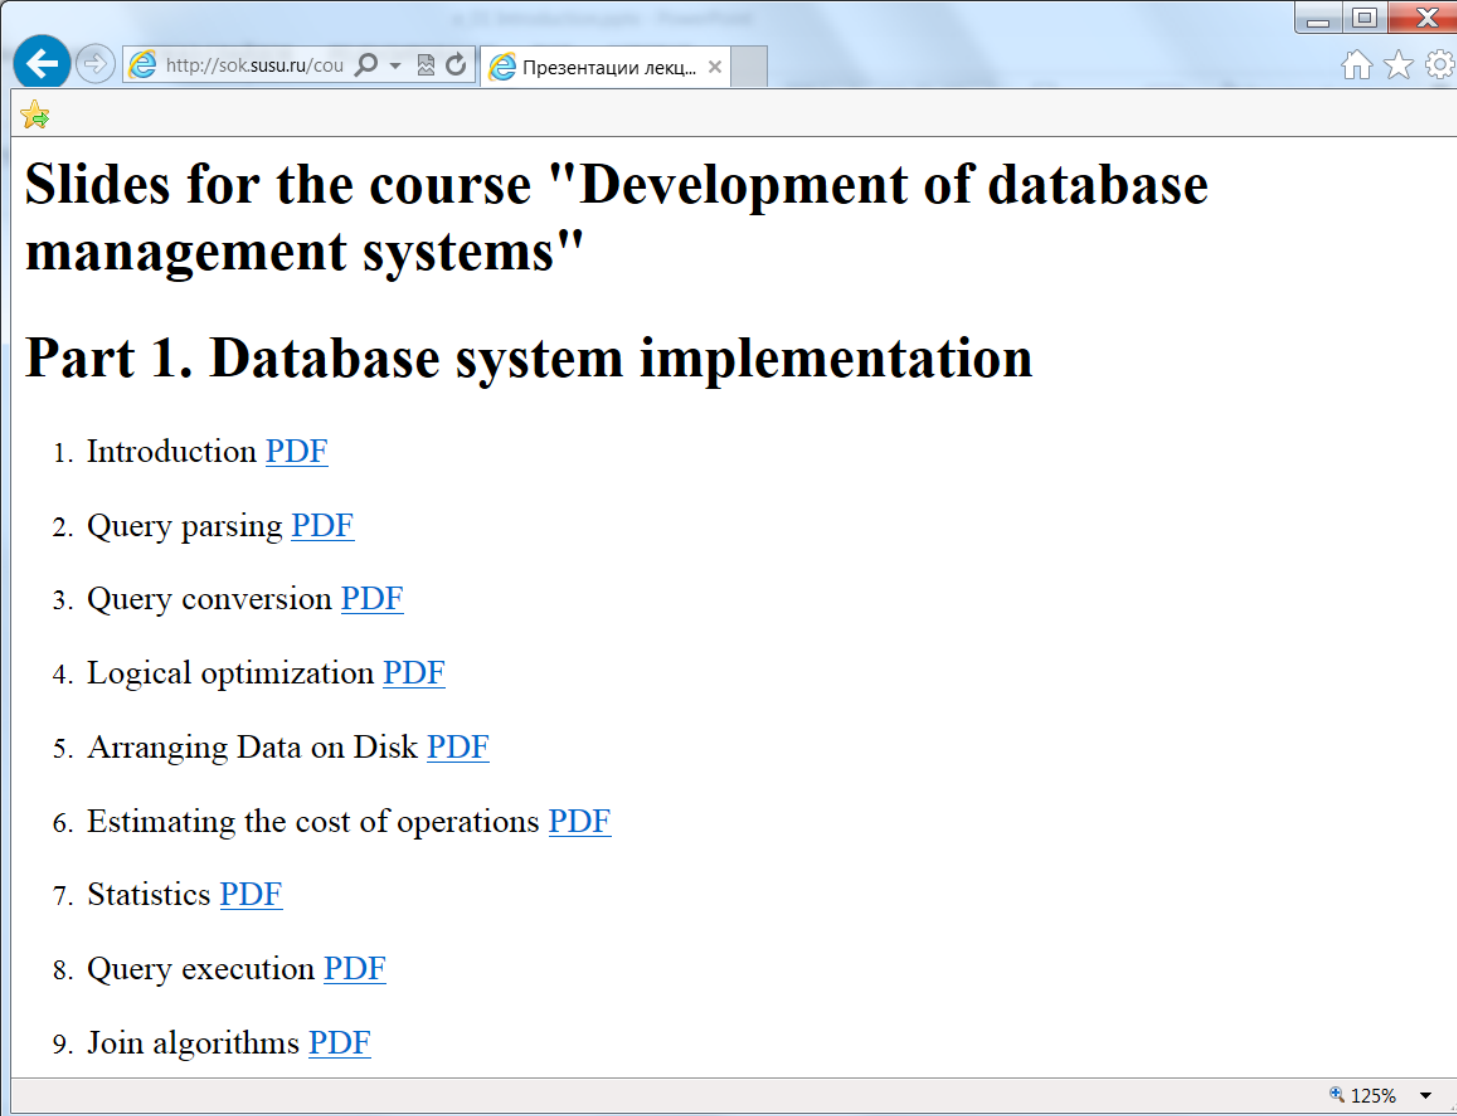
The screenshot shows a web browser window with the following content:

- Address bar: http://sok.susu.ru/courses/QueryProc/e_index.html
- Page title: Database system im...
- Page content:
 - ## Database system implementation
 - 
 - [Leonid B. Sokolinsky](#)
 - [Slides](#)
 - [Exam questions I \(PDF\)](#)
 - [Exam questions II \(PDF\)](#)
 - ### Textbook
 - [[Garcia-MolinaUW 08](#)] *Garcia-Molina H., Ullman J.D., Widom J.* Database Systems: The Complete Book (2nd Edition). Prentice Hall, 2008. 1224 p.

The browser interface includes navigation buttons (back, forward, search, refresh), a star icon for bookmarks, and a zoom level of 125% at the bottom right.

Slides

http://sok.susu.ru/courses/QueryProc/lectures/e_index.html

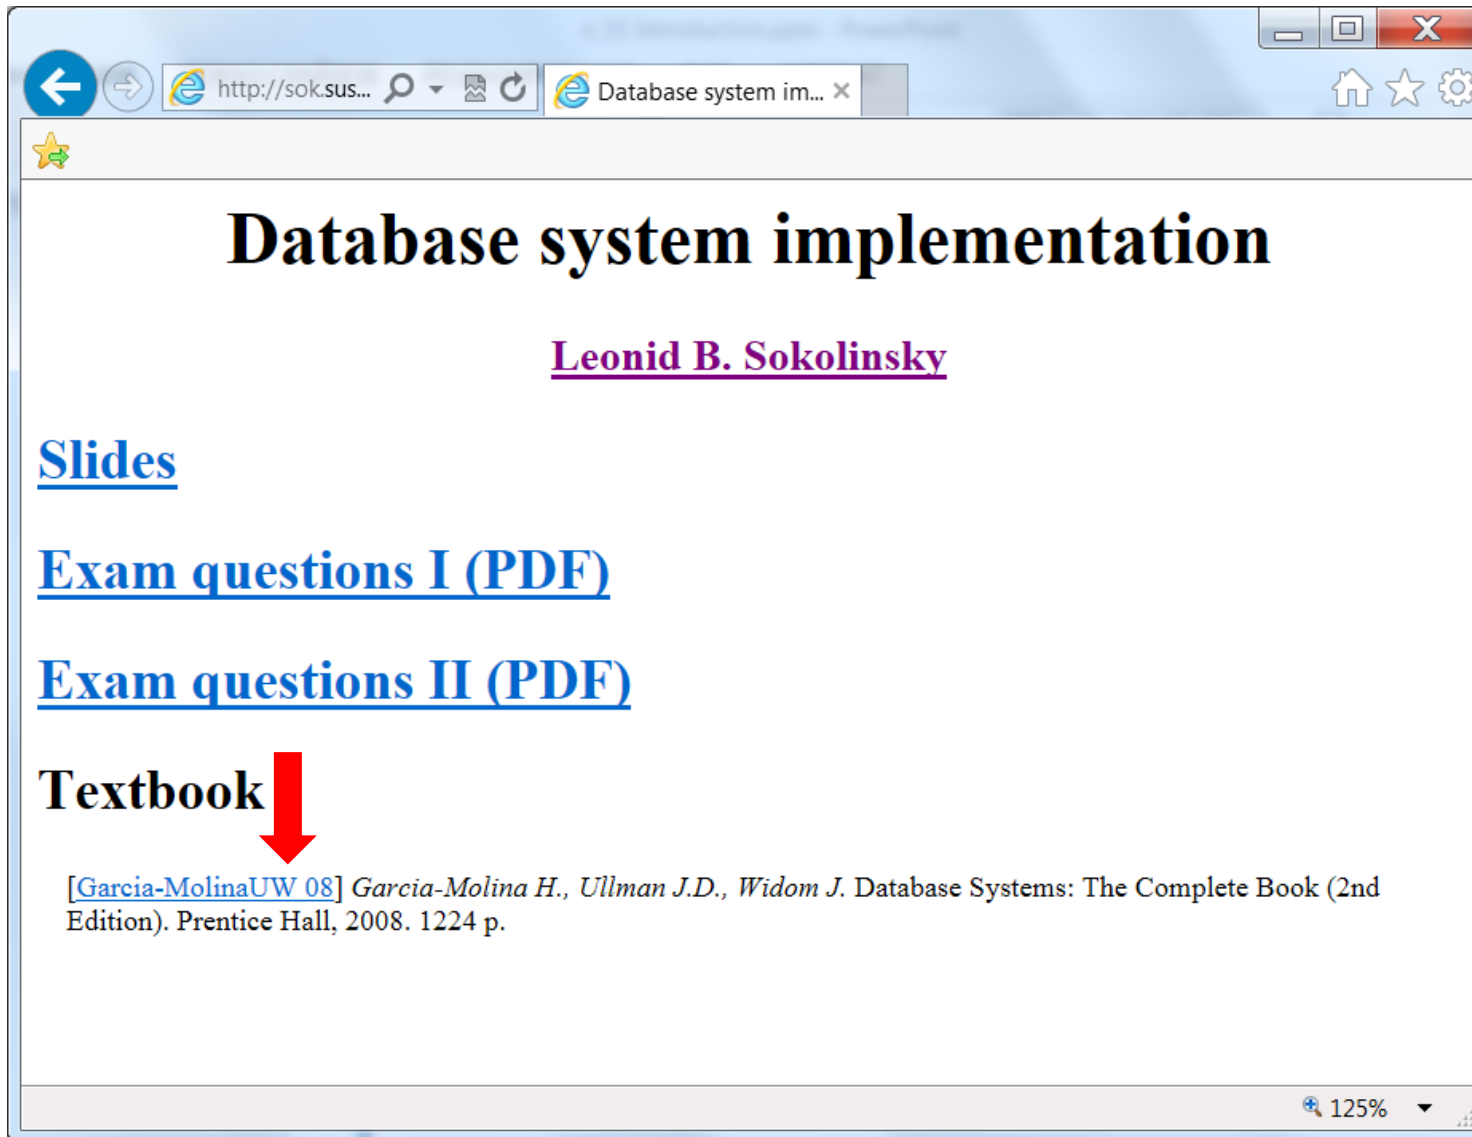


The screenshot shows a web browser window with the following content:

- Address bar: <http://sok.susu.ru/cou>
- Page title: Презентации лекц...
- Star icon
- Section header: **Slides for the course "Development of database management systems"**
- Section header: **Part 1. Database system implementation**
- List of slides:
 1. Introduction [PDF](#)
 2. Query parsing [PDF](#)
 3. Query conversion [PDF](#)
 4. Logical optimization [PDF](#)
 5. Arranging Data on Disk [PDF](#)
 6. Estimating the cost of operations [PDF](#)
 7. Statistics [PDF](#)
 8. Query execution [PDF](#)
 9. Join algorithms [PDF](#)
- Zoom: 125%

Part 1. Database system implementation

http://sok.susu.ru/courses/QueryProc/e_index.html



The screenshot shows a web browser window with the following content:

- Address bar: [http://sok.susu...](http://sok.susu.ru/courses/QueryProc/e_index.html)
- Page Title: Database system im...
- Page Content:
 - Database system implementation**
 - Leonid B. Sokolinsky
 - [Slides](#)
 - [Exam questions I \(PDF\)](#)
 - [Exam questions II \(PDF\)](#)
 - Textbook** (with a red arrow pointing down)
 - [[Garcia-MolinaUW 08](#)] *Garcia-Molina H., Ullman J.D., Widom J.* Database Systems: The Complete Book (2nd Edition). Prentice Hall, 2008. 1224 p.
- Bottom right: 125% zoom, play button icon.

Bibliography on Information Technology

<http://bit.sp.susu.ac.ru/id/9645/view>

(Login)

The screenshot shows a web browser window with the URL <http://bit.sp.susu.ac.ru/id/9645/view>. The page title is "Информационные технологии". The main content area displays a search result for the book "Database Systems: The Complete Book (2nd Edition)" by Garcia-Molina H., Ullman J.D., and Widom J. The authors are listed as Garcia-Molina H., Ullman J.D., and Widom J. The bibliographic reference is: Garcia-Molina H., Ullman J.D., Widom J. Database Systems: The Complete Book (2nd Edition). Prentice Hall, 2008. 1224 p. The annotation states: "Database Systems: The Complete Book is ideal for Database Systems and Database Design and Application courses offered at the junior, senior and graduate levels in Computer Science departments. A basic understanding of algebraic expressions and laws, logic, basic data structure, OOP concepts, and programming environments is implied. Written by well-known computer scientists, this introduction to database systems offers a comprehensive approach, focusing on database design, database use, and implementation of database applications and database management systems. The first half of the book provides in-depth coverage of databases from the point of view of the database designer, user, and application programmer. It covers the..."

Login

<http://bit.sp.susu.ac.ru/accounts/login/>

The screenshot shows a web browser window displaying the login page of the bit.sp.susu.ac.ru website. The browser's address bar shows the URL <http://bit.sp.susu.ac.ru/accounts/login/>. The page has a green header with a search bar and a 'Поиск' button. Below the header, there is a navigation menu with links to 'Главная', 'Электронная библиотека', 'Предметный указатель', 'Новости', and 'Конференции'. The main content area is titled 'Аутентификация' and contains a login form with the following fields and buttons:

- Имя** (Name): A text input field.
- Пароль** (Password): A text input field.
- Войти** (Enter): A button to submit the login form.

A red arrow points to the 'Войти' button. Below the form, there is a link to [зарегистрироваться](#) (register) for users who do not have an account. The footer of the page contains the following text:

Статистика: рубрик 34/170/482, записей 5971, публикаций 5634, новостей 157, конференций 99, интернет-ресурсов 70.
Техническая поддержка
© ФГБОУ ВПО «Южно-Уральский государственный университет» (НИУ)

Find

blf http://bit.sp.susu.ac.ru/ Информационные технол...

Файл Правка Вид Избранное Сервис Справка

Главная Электронная библиотека Предметный указатель Новости Конференции Студент каф. С. П. Выйти

blf

Database Systems: The Complete Book Поиск (Find)

Расширенный поиск

Новости

30/09 9-я конференция .NET-разработчиков

12 октября 2014 г. в г. Челябинск состоится 9-я конференция .NET-разработчиков. Основная цель – обсуждение передовых практик и инструментов .NET-разработчиков, обмен знаниями в сфере создания новых программ. Ежегодно в конференции принимают участие более 1500 разработчиков из Екатеринбурга, Уфы, Тюмени, Санкт-Петербурга, Москвы, Озерска, Снежинска, Омска, Новосибирска и других городов России и стран СНГ. Конференция предназначена как для профессионалов, так и для начинающих программистов. Вход на конференцию бесплатный после предварительной регистрации.

[Подробнее...](#)

19/09 Открытая лекция проф. Мориса Доусона "Использование интеллектуального анализа данных в конкурентной борьбе"

Access to full text

The screenshot shows a web browser window with the URL <http://bit.sp.susu.ac.ru/sear>. The page title is "Информационные технологии". A search bar contains the text "Database Systems: The Complete Book" and a "Поиск" button. Below the search bar, the results are listed under the heading "Поиск".

1. [Database Systems: The Complete Book \(2nd Edition\)](#) ←
Garcia-Molina H., Ullman J.D., Widom J. Database Systems: The Complete Book (2nd Edition). Prentice Hall, 2008. 1224 p.
Ключевые слова: Системы баз данных, Общие руководства, Учебники, Введение в системы баз данных, Реализация СУБД
Составитель: Соколинский Леонид Борисович (30.09.2012 10:47:40) [Подробнее...](#)
Модератор: Соколинский Леонид Борисович (30.09.2012 10:51:09)

2. [In-Memory Data Management: An Inflection Point for Enterprise Applications](#)
Plattner H., Zeier A. In-Memory Data Management: An Inflection Point for Enterprise Applications. Springer, 2011. 254 p.
Ключевые слова: Системы баз данных, Системы баз данных в оперативной памяти, Обзоры, Параллельная обработка данных, Хранение данных по столбцам, Колоночные СУБД, Бизнес-аналитика, OLAP
Составитель: Аксенова Елена Владимировна (06.07.2012 14:29:02) [Подробнее...](#)
Модератор: Соколинский Леонид Борисович (19.07.2012 09:56:03)


3. [Intelligent Database Systems](#)
Bertino E., Catania B., Zarrì G.P. Intelligent Database Systems // Addison-Wesley, 2001. No. 0. P. 400-0.
Ключевые слова: Системы баз данных, Базы знаний, Общие руководства
Составитель: Соколинский Леонид Борисович (29.11.2010 13:01:16) [Подробнее...](#)
Модератор: Соколинский Леонид Борисович (15.11.2010 10:00:00)

Full text (PDF 30 MB)

The screenshot shows a web browser window with the address bar displaying 'http://bit.sp.susu.ac.ru/id/9'. The page title is 'Электронная библиотека'. The main content area is titled 'Информационные технологии' and features a search bar with the text 'Расширенный поиск'. Below the search bar, the breadcrumb trail reads 'Электронная библиотека > Книги, учебники, монографии'. The main entry is for the book 'Database Systems: The Complete Book (2nd Edition)'. The authors listed are Garcia-Molina H., Ullman J.D., and Widom J. The bibliographic reference is: Garcia-Molina H., Ullman J.D., Widom J. Database Systems: The Complete Book (2nd Edition). Prentice Hall, 2008. 1224 p. The annotation states: 'Database Systems: The Complete Book is ideal for Database Systems and Database Design and Application courses offered at the junior, senior and graduate levels in Computer Science departments. A basic understanding of algebraic expressions and laws, logic, basic data structure, OOP concepts, and programming environments is implied. Written by well-known computer scientists, this introduction to database systems offers a comprehensive approach, focusing on database design, database use, and implementation of database applications and database management systems. The first half of the book provides in-depth coverage of databases from the point of view of the database designer, user, and application programmer. It covers the

Database Systems: The Complete Book

Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer Widom



Hector Garcia-Molina



Jeffrey D. Ullman



Jennifer Widom



Database schema "Supplies"

Relation	Semantics
S (ID_S, Name_S, City_S, Rating, Deposit)	Information about suppliers
P (ID_P, Name_P, City_P, Size_P, Color)	Information about parts
H (ID_H, Volume_H)	Hangars for storage in the assembly plant
SP (ID_S, ID_P, Quantity, Price, ID_H)	Supplies of parts



Relation S (Suppliers)

Attribute	Semantics
ID_S*	Supplier ID
Name_S	Supplier name
City_S	Supplier city
Rating	Supplier rating (reliability)
Deposit	Security deposit



Relation P (Parts)

Attribute	Semantics
ID_P*	Part ID
Name_P	Part name
City_P	Part city
Size_P	Part size
Color	Part color



Nut



Bolt

Relation H (Hangars)

Attribute	Semantics
ID_H*	Hangar ID
Volume_H	Hangar free volume

Relation SP (Supplies of parts)

Attribute	Semantics
ID_S*	Supplier ID
ID_P*	Part ID
Quantity	Number of pieces in the supply
Price	Price per piece
ID_H	Hangar ID for storage

Example of Database

S (Suppliers)

ID_S	Name_S	City_S	Rating	Deposit
1	Peterson	Moscow	100	1000
2	Johnson	Tokyo	60	2000
3	Simpson	Moscow	80	1000
4	Abrahams	Sydney	50	5000

SP (supplies of parts)

ID_S	ID_P	Quantity	Price	ID_H
1	1	2000	100	1
1	2	3000	100	2
2	3	2500	200	3
4	4	500	150	3

P (parts)

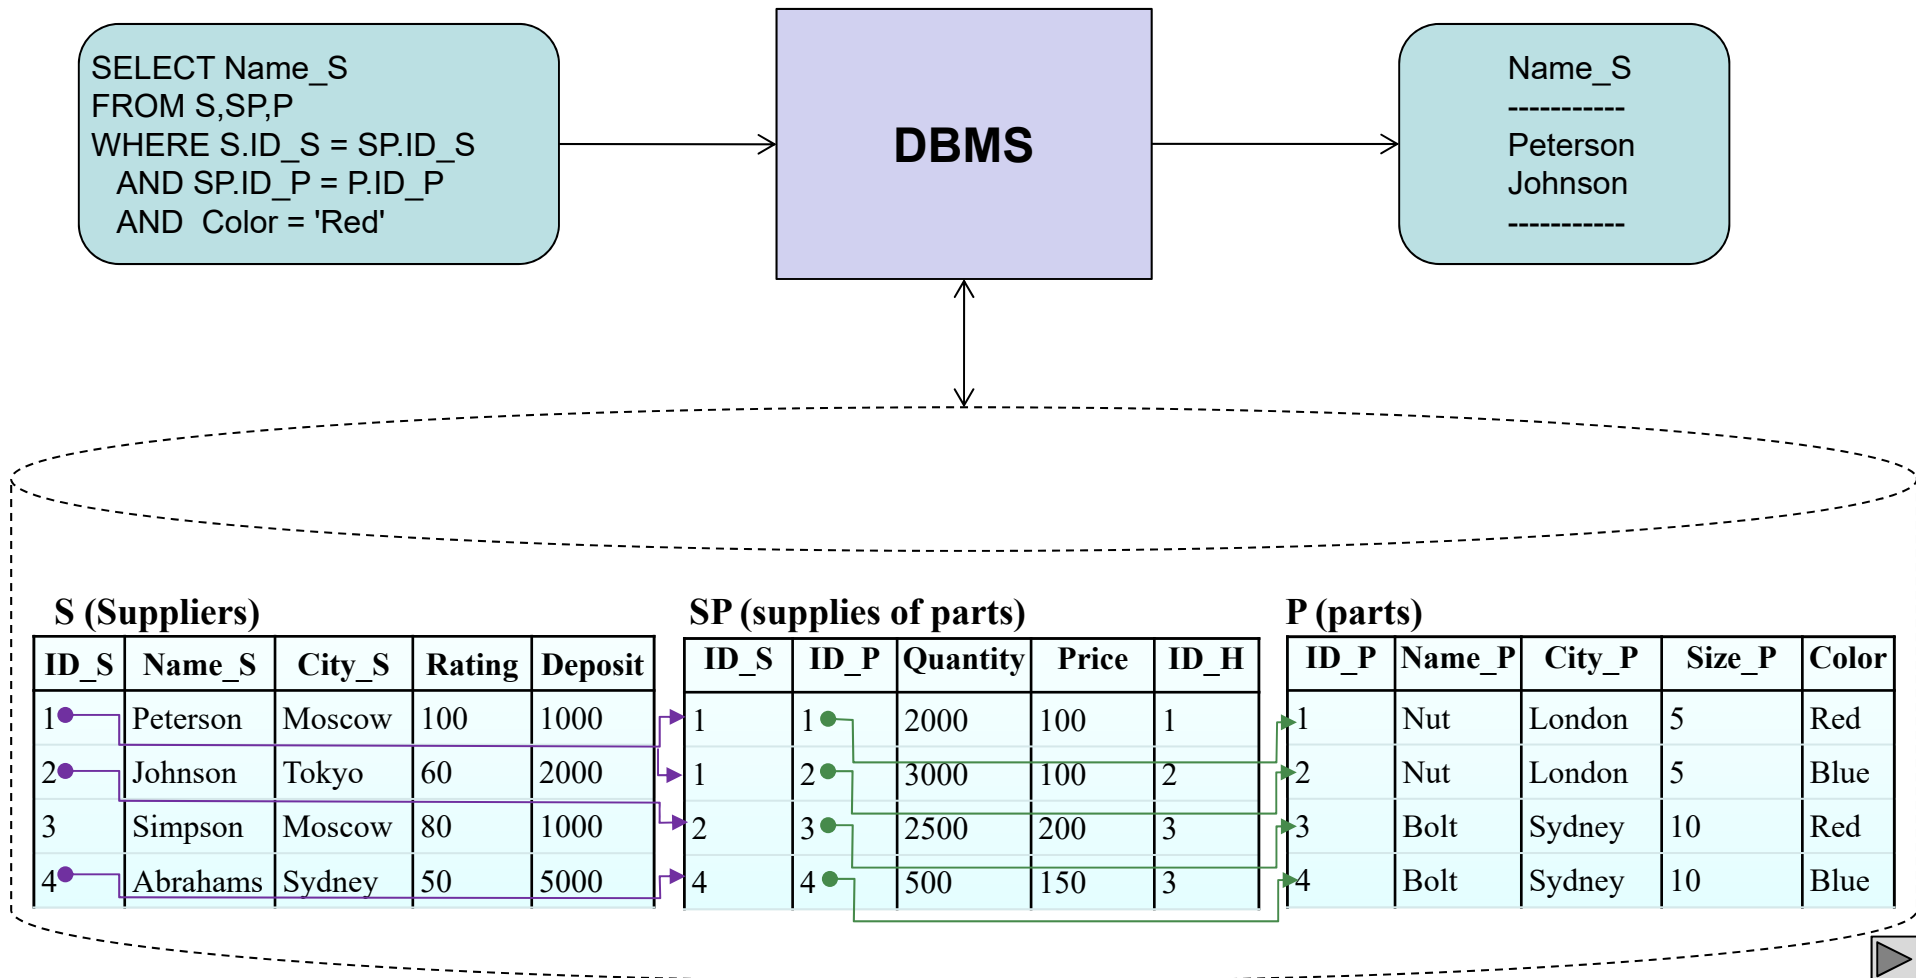
ID_P	Name_P	City_P	Size_P	Color
1	Nut	London	5	Red
2	Nut	London	5	Blue
3	Bolt	Sydney	10	Red
4	Bolt	Sydney	10	Blue

Query example

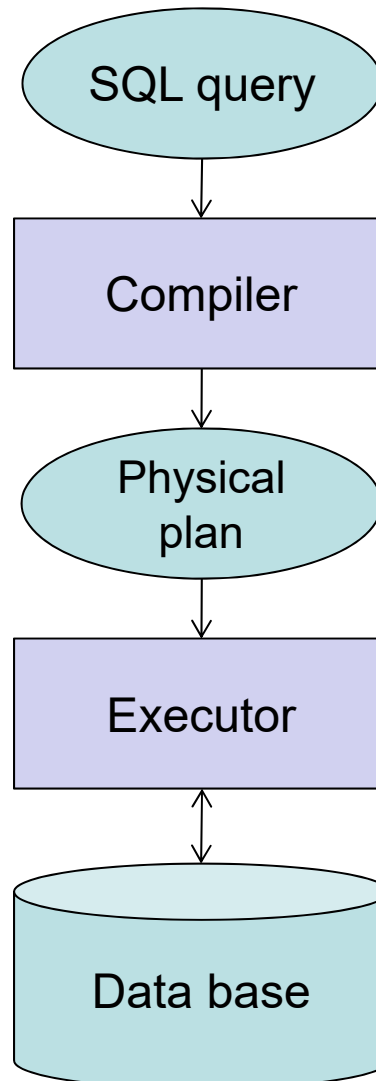
The names of suppliers who supply at least one red part.

```
SELECT Name_S
FROM S, SP, P
WHERE S.ID_S = SP.ID_S
      AND SP.ID_P = P.ID_P
      AND Color = 'Red';
```


How DBMS gets the answer?



Outline of query processing



On the way to the physical plane



The main operations of relational algebra

Operation	Name	Semantics
$\pi_{A,C}(R)$	Projection	Projection removes some columns from the relation.
$\sigma_{\theta}(R)$	Selection	The selection produces a result consisting of all tuples of the relation that satisfy the selection condition θ .
$R \times S$	Cartesian product	Pairing the tuples of two relations in all possible ways.
$R \bowtie S$	Natural join	Pairing the tuples of two relations that agree on all common attributes.
$R \bowtie_{\theta} S$	Theta-join	Pairing the tuples of two relations that meet a condition θ .
$\delta(R)$	Duplicate elimination	Turns a bag into a set by eliminating all but one copy of each tuple.
$\gamma_{B, \text{AVG}(C) \rightarrow X}(R)$	Grouping	Operator that combines the effect of grouping and aggregation.



Projection

R

A*	B	C
1	20	100
2	40	300
3	20	100
4	10	300

 $\pi_{A,C}(R)$

A*	C
1	100
2	300
3	100
4	300

Selection

R

A*	B	C
1	20	100
2	40	300
3	20	100
4	10	300

 $\sigma_{C < 150}(\mathbf{R})$

A*	B	C
1	20	100
3	20	100

Cartesian product

R

A*	B	C
1	20	100
2	40	300
3	20	100
4	10	300

S

D*	A#	E
1	3	0.2
2	1	0.5
3	1	0.5

Cartesian product pairs the tuples of two relations in all possible ways.

R×S

R.A*	B	C	D*	S.A	E
1	20	100	1	3	0.2
1	20	100	2	1	0.5
1	20	100	3	1	0.5
2	40	300	1	3	0.2
2	40	300	2	1	0.5
2	40	300	3	1	0.5
3	20	100	1	3	0.2
3	20	100	2	1	0.5
3	20	100	3	1	0.5
4	10	300	1	3	0.2
4	10	300	2	1	0.5
4	10	300	3	1	0.5



Theta-join

R

A*	B	C
1	20	600
2	40	300
3	20	250
4	10	300

S

D*	A#	E
1	3	0.2
2	1	0.7
3	1	0.5

$R \bowtie S$
 $R.A = S.A \wedge C > E * 1000$

R.A*	B	C	D*	S.A	E
1	20	600	3	1	0.5
3	20	250	1	3	0.2

Theta-join pairs the tuples of two relations that meet a condition theta.



Natural join

R			S		
A*	B	C	D*	A#	E
1	20	100	1	3	0.2
2	40	300	2	1	0.5
3	20	100	3	1	0.5
4	10	300			

$R \bowtie S$
 $R.A=S.A$

R.A*	B	C	D*	S.A	E
1	20	100	2	1	0.5
1	20	100	3	1	0.5
3	20	100	1	3	0.2

$$R \bowtie S = \pi_{* \setminus S.A} \left(R \bowtie S \right)_{R.A=S.A}$$

A*	B	C	D*	E
1	20	100	2	0.5
1	20	100	3	0.5
3	20	100	1	0.2

The appearance of duplicates

R

A*	B	C
1	20	100
2	40	300
3	20	100
4	10	300

 $\pi_{B,C}(R)$

B	C
20	100
40	300
20	100
10	300

Duplicate elimination

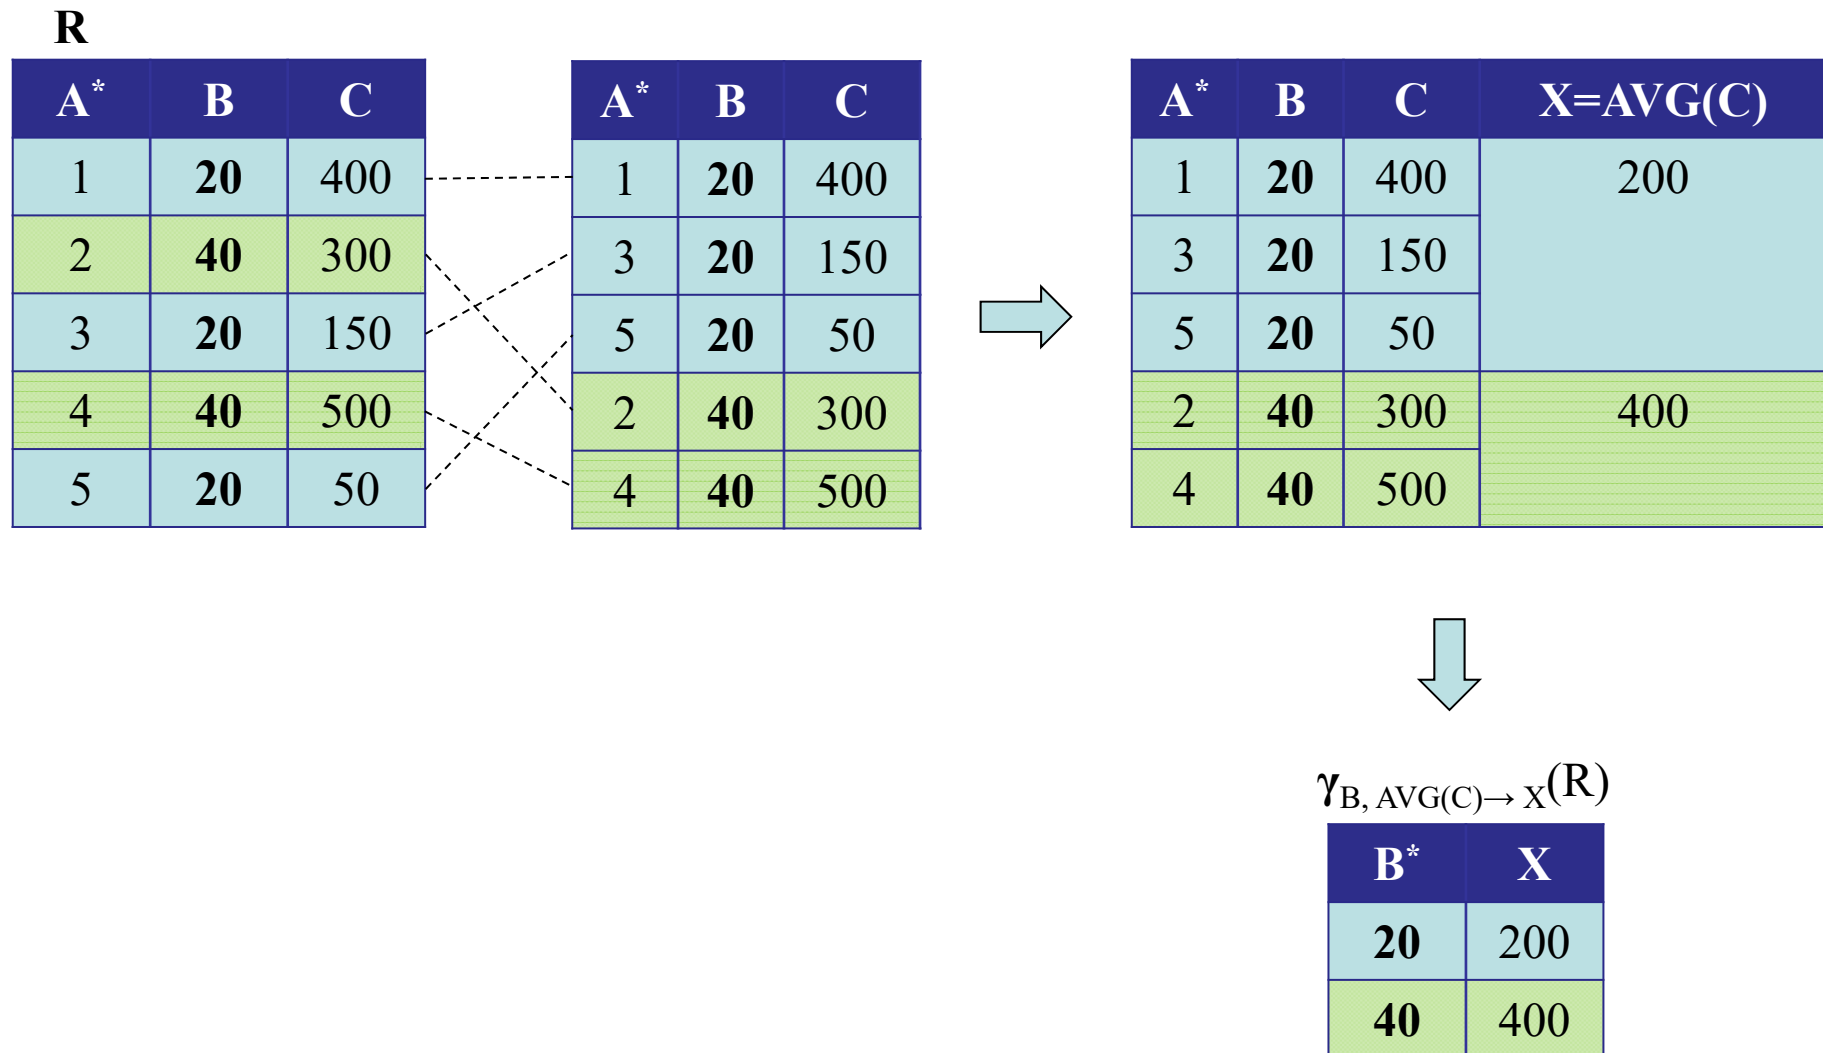
R

B	C
20	100
40	300
20	100
10	300

 $\delta(\mathbf{R})$

B*	C
20	100
40	300
10	300

Grouping



Aggregation

R

A*	B	C
1	2	400
2	4	300
3	2	150
4	4	500
5	2	50

→

$\gamma_{\text{SUM}(B*C)}(R)$

4400

Aggregate functions

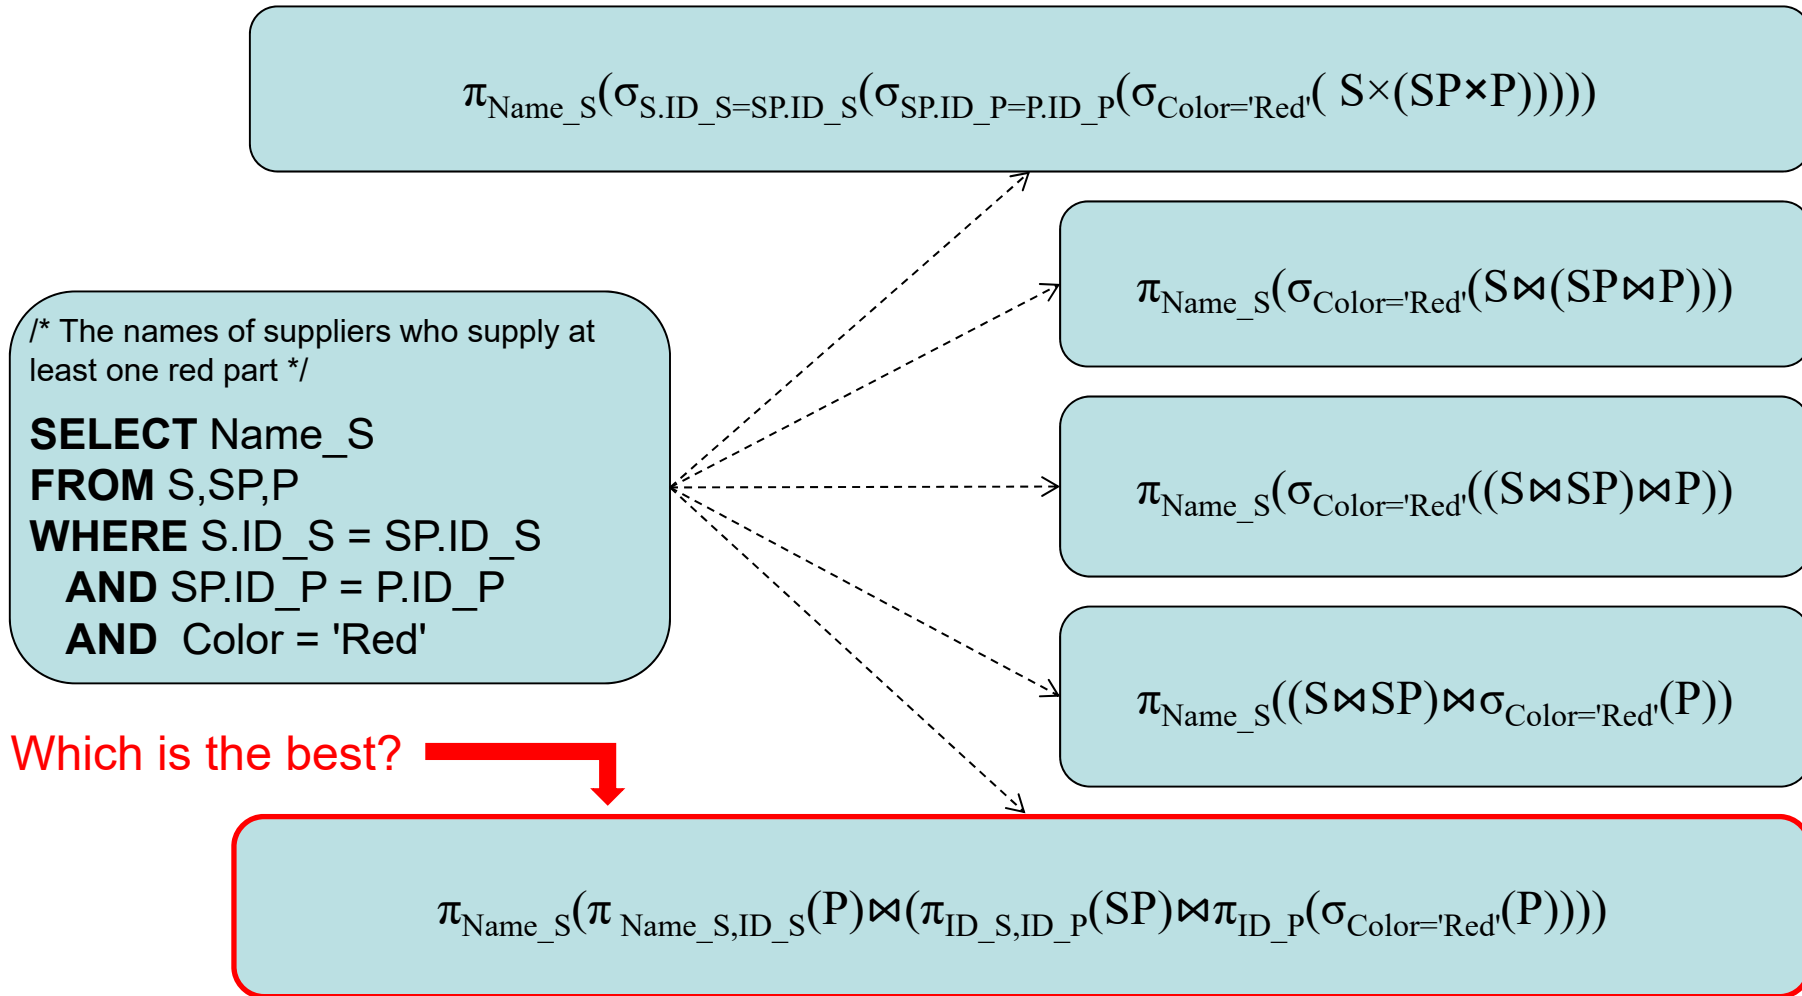
Function	Semanics
SUM	The sum of values
AVG	The average value
MIN	The least value
MAX	The greatest value
COUNT	The number of tuples

Relational algebra operations in SQL

Operation	SQL
$\pi_{A,C}(R)$	SELECT A,C FROM R
$\sigma_{\theta}(R)$	SELECT * FROM R WHERE θ
$R \times S$	R CROSS JOIN S
$R \bowtie S$	R NATURAL JOIN S
$R \bowtie_{\theta} S$	R JOIN S ON θ
$\delta(R)$	SELECT DISTINCT * FROM R
$\gamma_{B,AVG(C) \rightarrow X}(R)$	SELECT B, AVG(C) AS X FROM R GROUP BY B



Constructing an expression of relational algebra



Why?

How to construct?

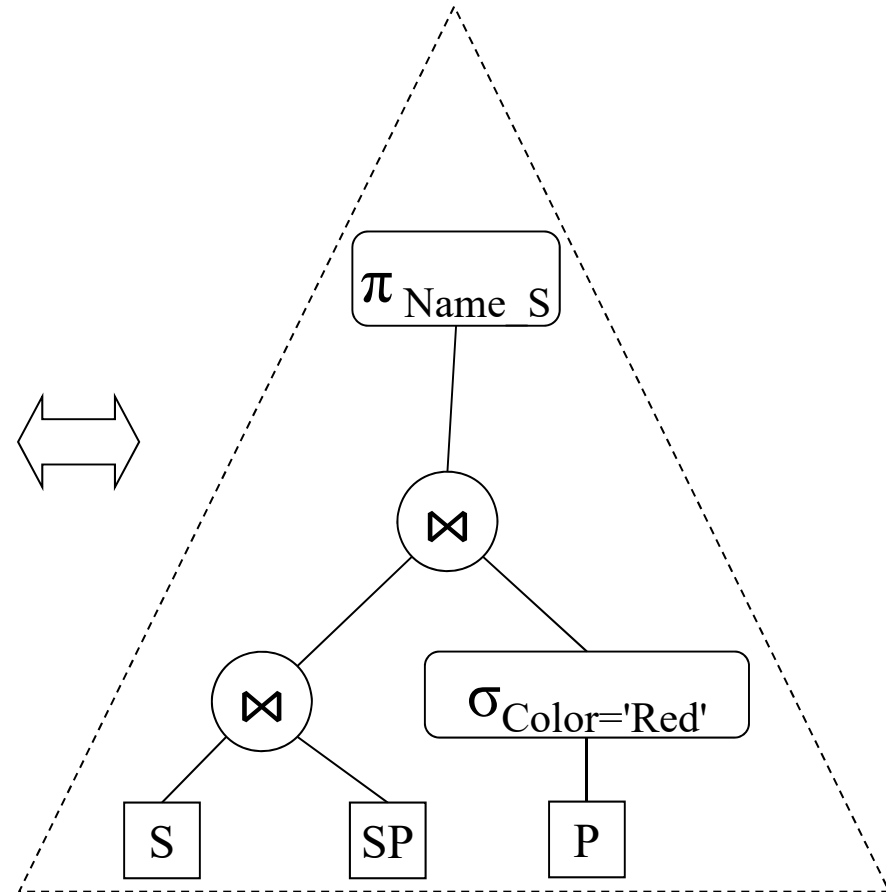

How to choose?



Representing relational algebra expression as a logical plan

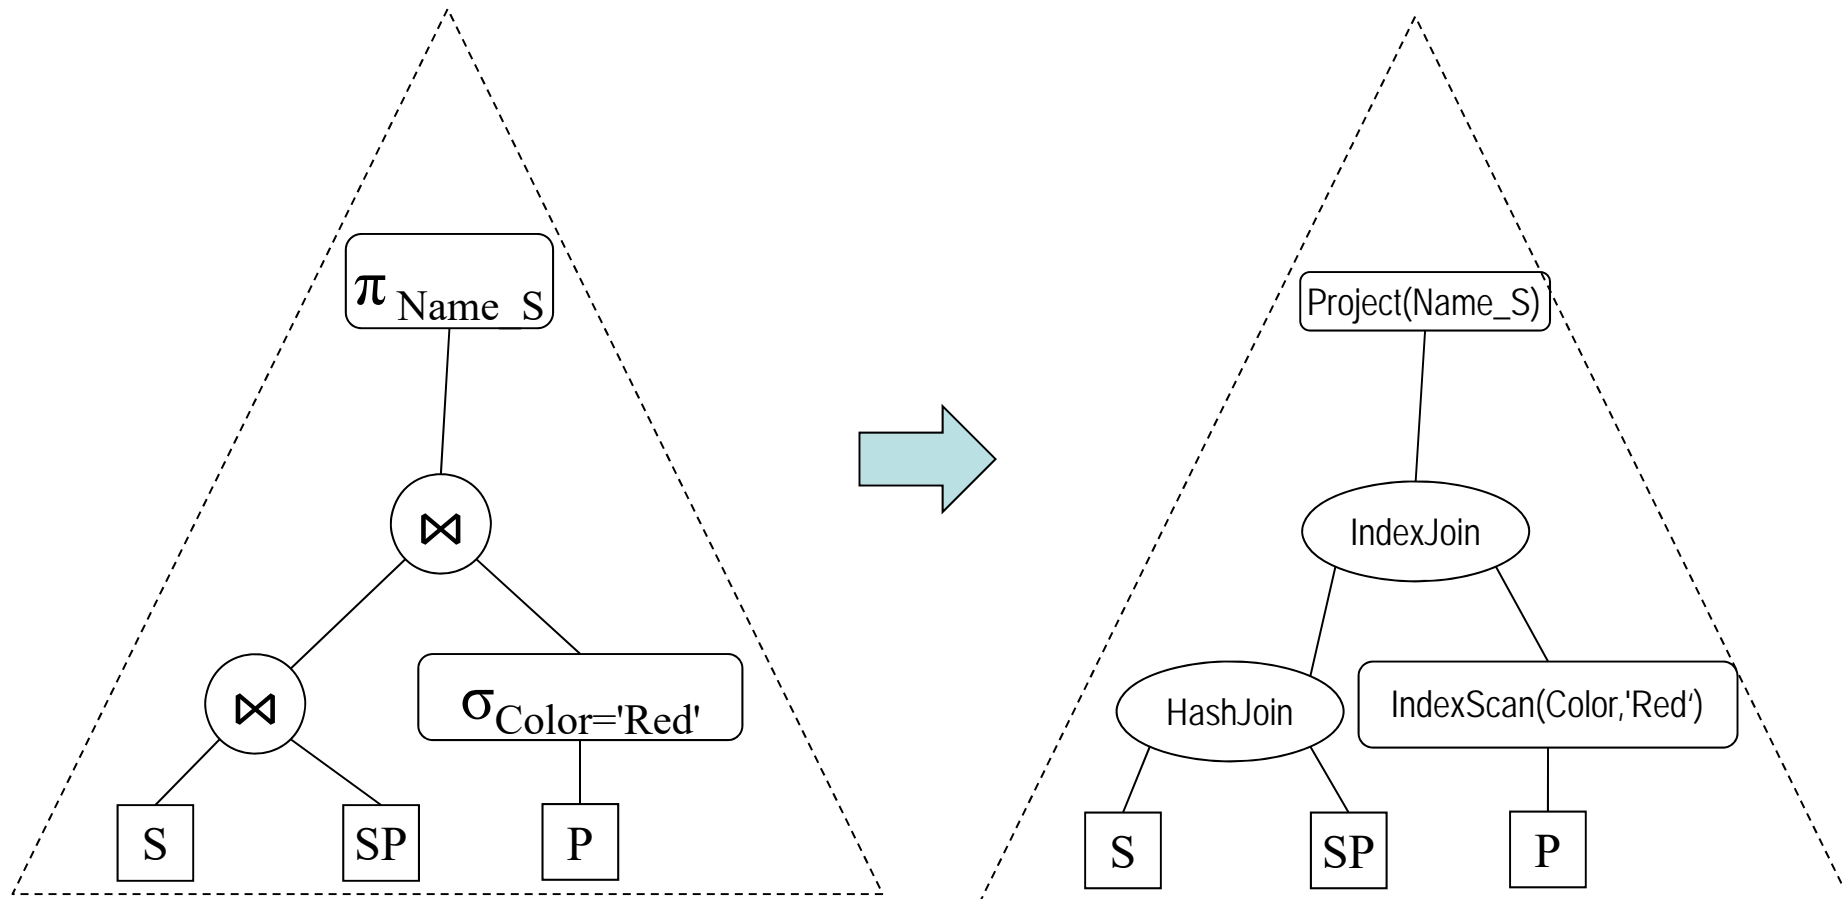
$$\pi_{\text{Name}_S}((S \bowtie SP) \bowtie \sigma_{\text{Color}='Red'}(P))$$

Relational algebra expression



Logical plan

Converting logical plan to physical plan

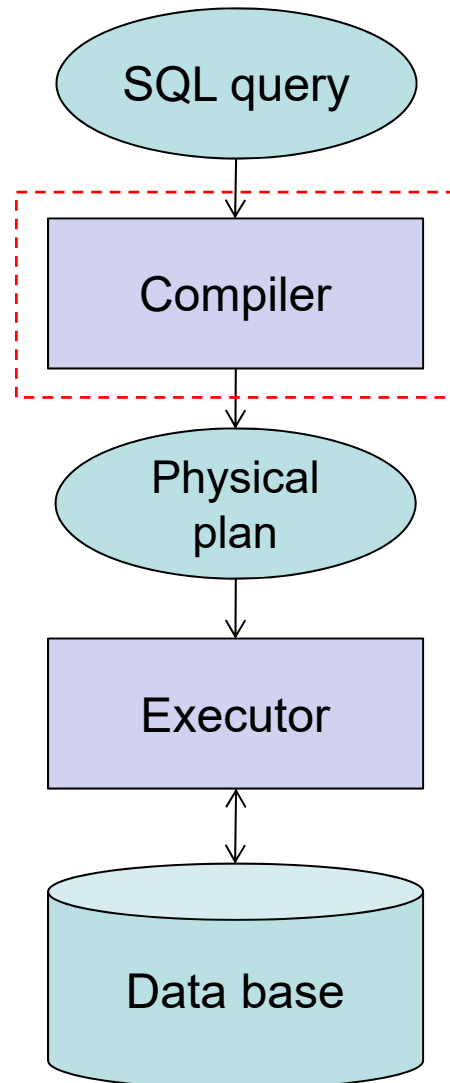


Logical plan

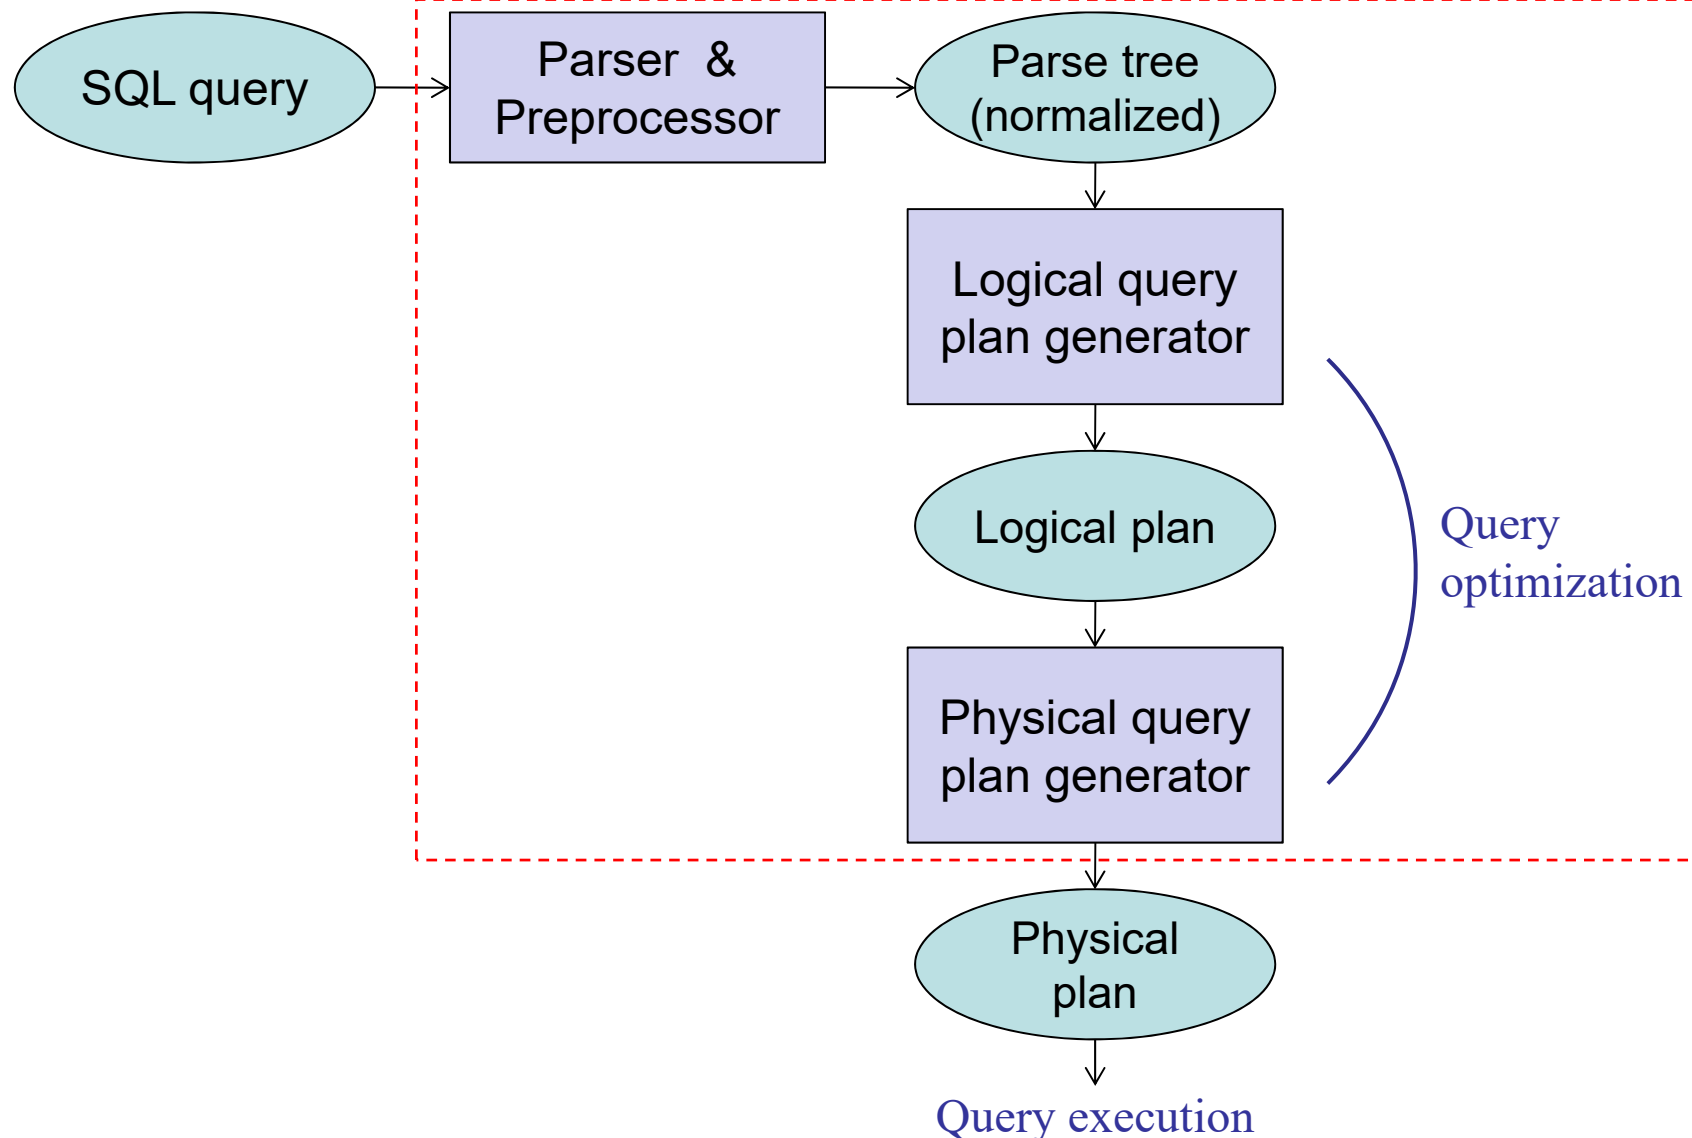
Physical plan

DBMS converts the logical plan into a physical plan replacing the relational algebra operations by the physical algebra operations.

Outline of query processing



Outline of query compilation



Query optimization

- Which of the algebraically equivalent forms of a query leads to the most efficient algorithm for answering the query?
- For each operation of the selected form, what algorithm should we use to implement that operation?
- How should the operations pass data from one to the other (in a pipelined fashion, in main-memory buffers, or via the disk)?

