Chapter 9

•**Multi-valued dimensions**
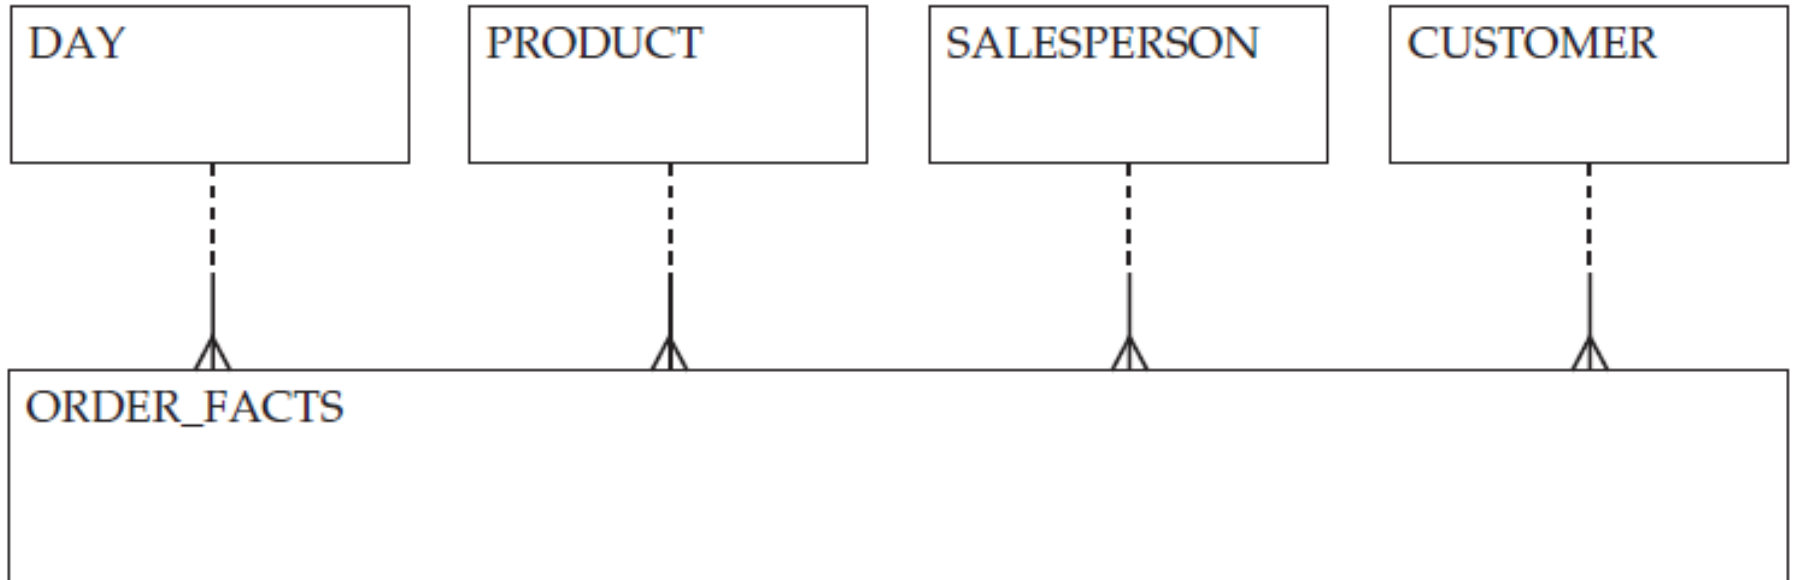
•Multi-valued attributes

Multi-valued Dimension

All schemas up to this point had 1-n relationships:
See fig 9-1

(Aside: What does fig 9-1 correspond to in the Chen notation?)

Multi-valued Dimension

Typical query

SELECT product.brand_name, sum(order_facts.order_dollars)
FROM product inner join order_facts
ON (product.product_key = order_facts.product_key)
GROUP BY
        product.brand_name

Now suppose a fact row must be related to more than one row of a dimension.

E.g. Consider that more than one sales rep is associated with an order.
See figure 9-2

Fig 9-3 is one solution – but too limited

We are primarily interested in fig 9-4
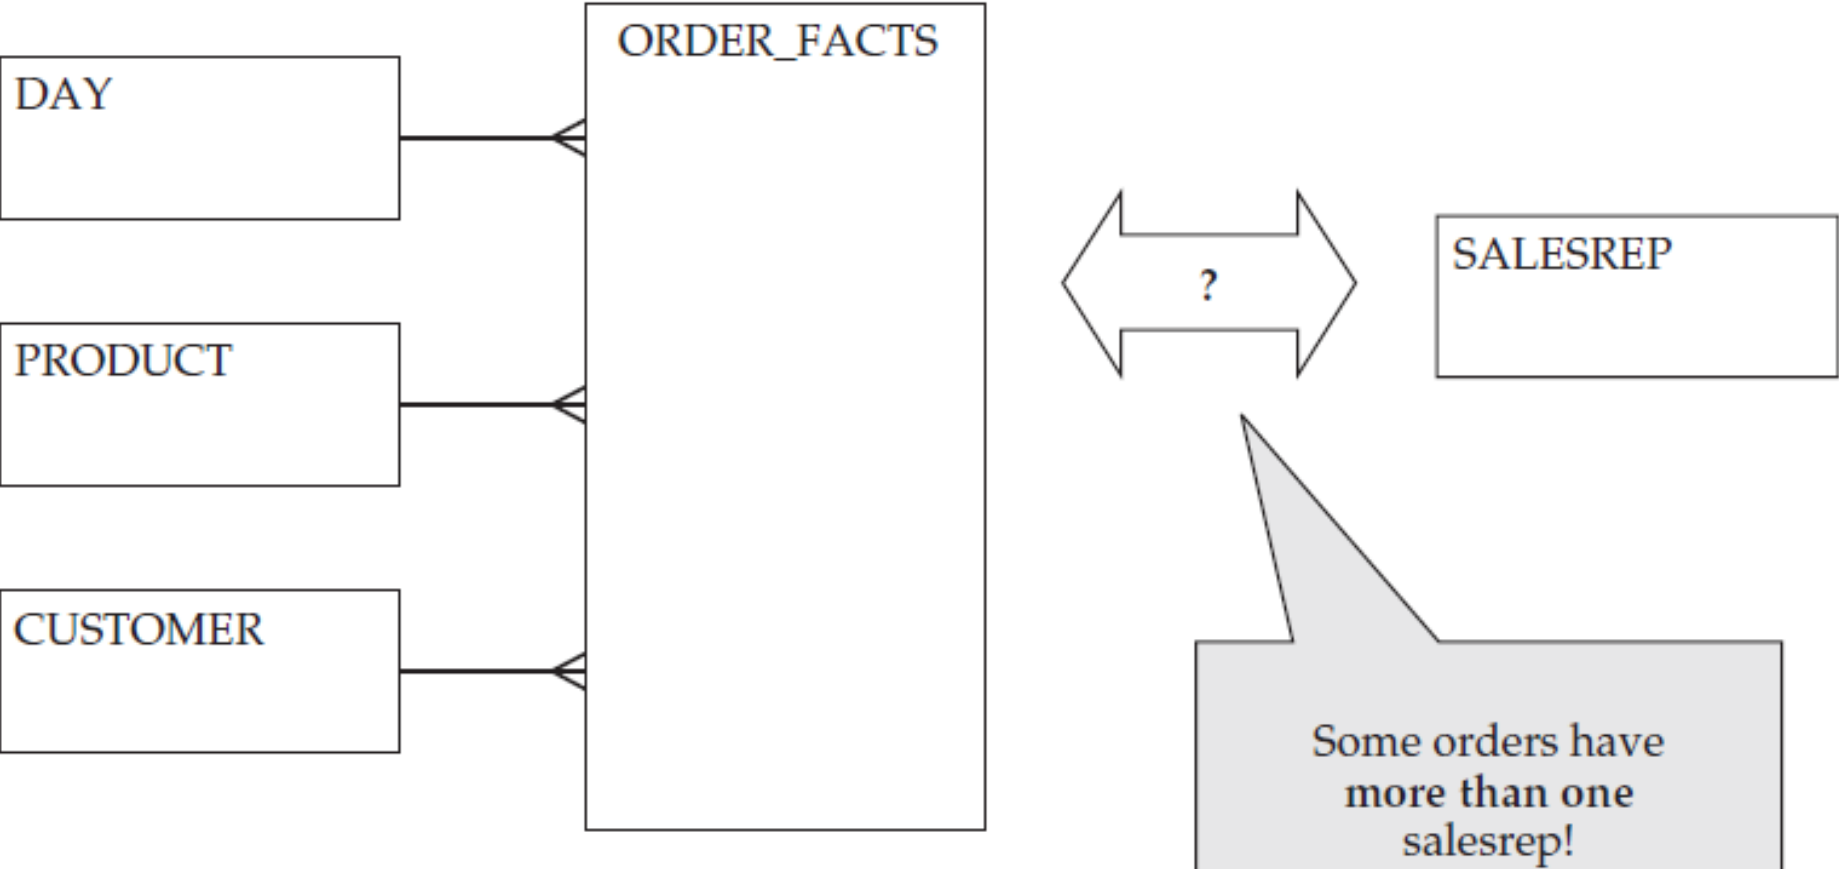
Multi-valued Dimension

Fig 9-2

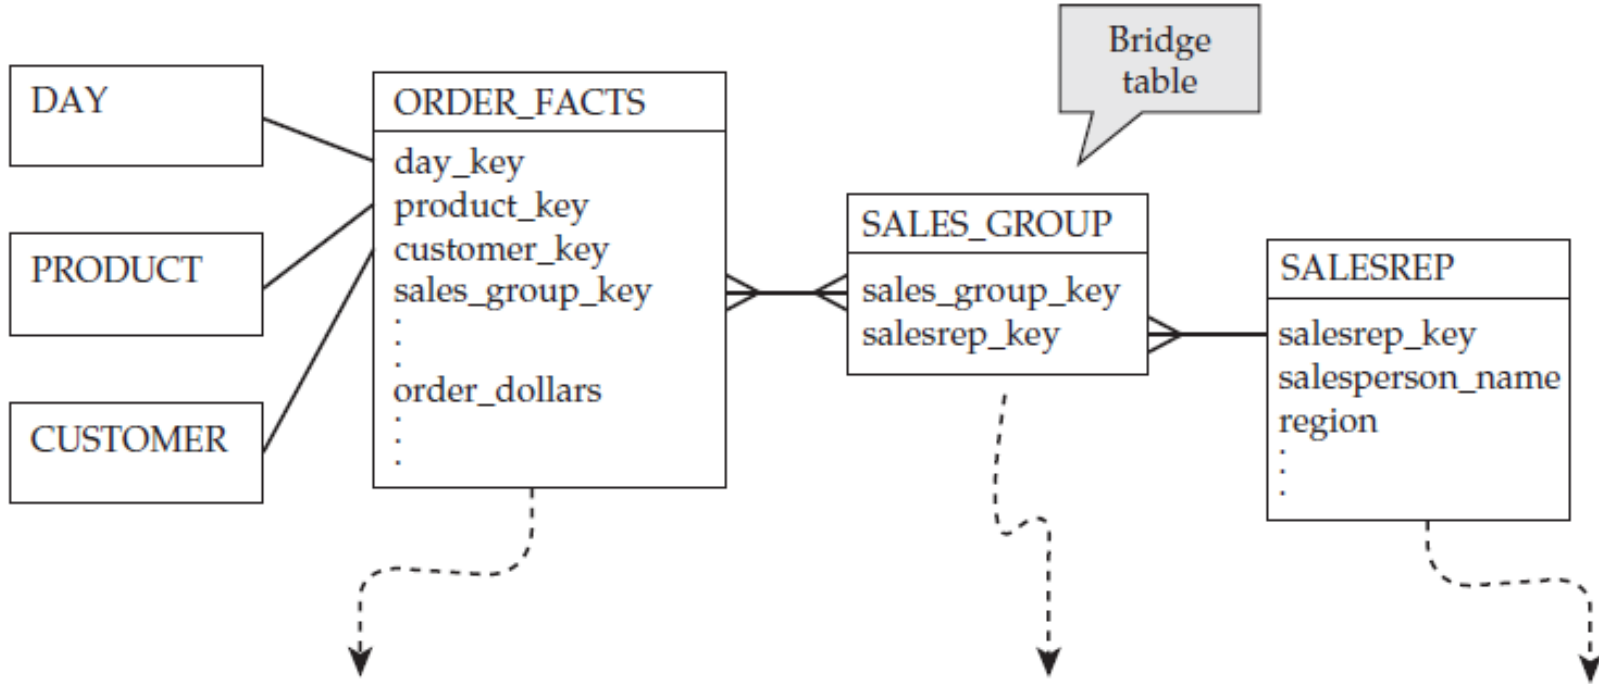The standard DM solution is to use a special type of bridge table
(Aside: not the traditional intersection table
What would that look like?)

The rows of the bridge table are organized into groups and the fact table identifies a single group.
Consider figure 9-4.

Multi-valued Dimension

Fig 9-4

Need to be careful when querying:

SELECT salesrep.region, sum (order_facts.order_dollars)
FROM   salerep, sales_group, order_facts
WHERE order_facts.sales_group_key = sales_group.sales_group_key
AND   sales_group.salesrep_key = salesrep.salesrep_key
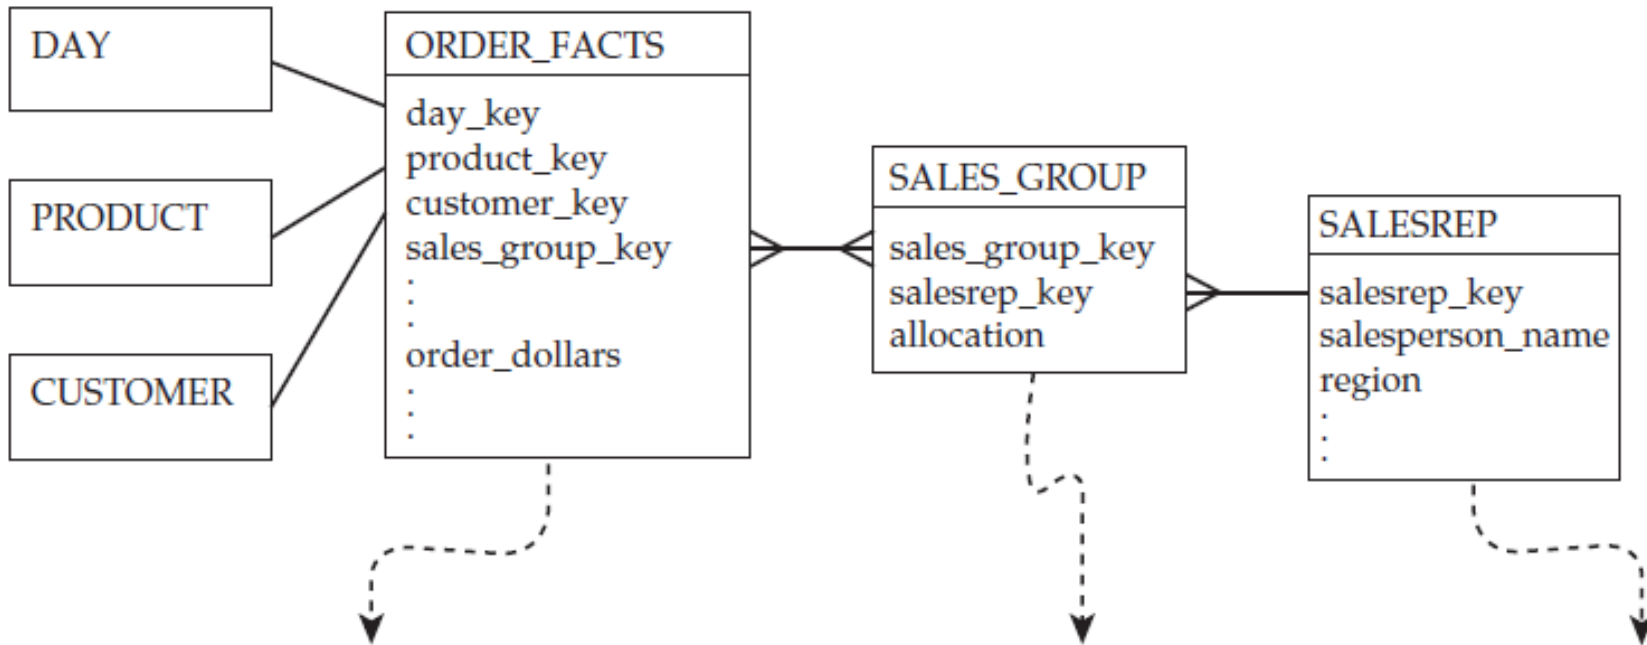GROUP BY salesrep.region

When executed, the DBMS finds two rows in the salesrep for the single order booked by Ann and Henry. In the join the order is included twice: once for Ann and once for Henry.

Due to the aggregation Ann and Henry's joint order is counted twice for region "East."

Multi-valued Dimension

The bridge table can be augmented with an allocation factor
– see fig 9-5

Multi-valued Dimension

Fig 9-5