


Physical Storage Hierarchy

- Main memory
- Disk storage
 - provides *direct access*
 - is unavoidable because of data *volume* and need for *persistence*
 - is critical for *performance*
- Tapes (*sequential access*) used for *logging* and *archiving*

The Need for Clustering

- *Access cost* depends on how many and which blocks where fetched
- *Seek time* is much larger than *latency* time, especially if the arm must move a lot
- *Consequence*: We must cluster related records in the *same or consecutive blocks*



Mixed Record Types Facilitate Clustering

- Records of different relations stored in the same block
- Facilitate frequent joins

Find theaters playing movies of "Berto"


M: Wild; Lynch; Winger	T: Wild; Odeon
M: Sky; Berto; Winger	T: Sky; Forum

The above clustering facilitates the join on Title

Sequencing

- Ordering records in file by some attribute (usually key)
- efficient sorted access
- implementation options:
 - next record physically consecutive
 - linked list
 - overflow area and eventual reorganization


Reds	Beatty	Beatty
Sky	Berto	Winger
Tango	Berto	Brando
Tango	Berto	Winger
Tango	Berto	Snyder
Wild	Lynch	Winger



Sequencing (cont'd)

Insertion of a new record: find where it should go and if there is a free record on the block insert

Otherwise put into overflow space



Spanned vs Unspanned Records

- Unspanned: Records must be in the same block
- Spanned: A record may reside in two blocks

