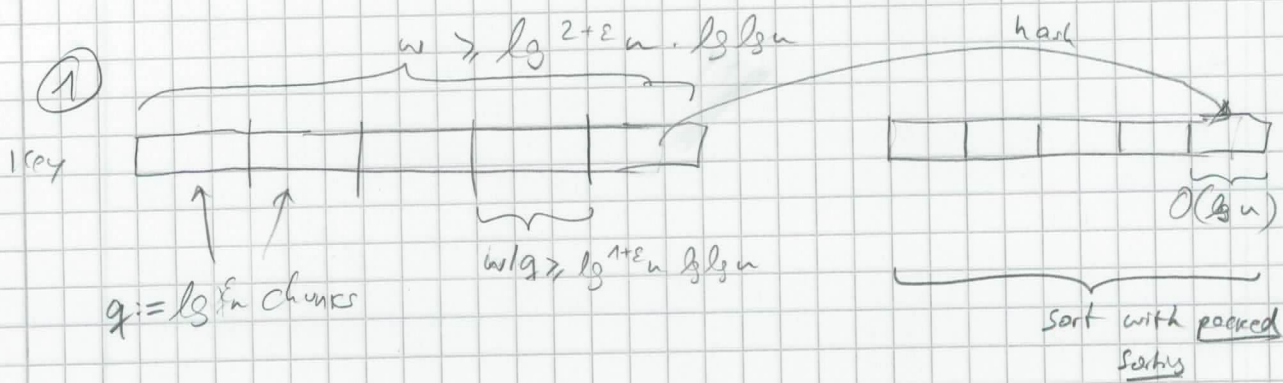


# Zusammenfassung Teil 1 von Radix Sort:



$\Rightarrow$  hash  $nq$  elements to table of size  $2^{O(\lg n)} = n^{O(1)}$  with univ. h.f.  $h$

$$IP[h(x) = h(y)] \leq \frac{1}{n}$$

$$\Rightarrow IP[\text{single collision}] \leq \binom{nq}{2} \cdot \frac{1}{n} \leq \frac{n^2 q^2}{n}$$

$$\text{set } n \leftarrow n^4 \Rightarrow IP[\text{single collision}] \leq \frac{n^2 \cdot \lg^{2\epsilon} n}{n^4} = \frac{\lg^{2\epsilon} n}{n^2}$$

$$\Rightarrow IP[h \text{ injective}] \leq 1 - IP[\text{single collision}] \approx 1 - \frac{1}{n^2} \quad (\text{"high"})$$

DO NOT CHECK INJECTIVITY IMMEDIATELY

② Use trie + recursion to account for the fact that hashes might not preserve order


recursion stops after  $O(1 + \frac{1}{\epsilon}) = O(1)$  calls with key lengths  $O(\lg n + \frac{w}{\lg^{1+\epsilon} n})$ .

$\underbrace{\hspace{10em}}_{O(w)} \quad \underbrace{\hspace{10em}}_{\frac{\leq w}{\lg n \cdot \lg n}}$


$\Rightarrow$  use Packed Radix for base case

③ check if output sorted and restart if not (only  $O(1)$  trials).

⑥



⇒ rebuilt tree is



⇒ final result  $x_2 \leq x_4 \leq x_3 \leq x_1$


So why didn't we use recursion in the first place (on the chunks of the keys)? The reason is that there are too many such chunks:  $n \cdot q$  many, whereas the "trie-technique" keeps their number linear in  $n$ . With the original chunks the running time would follow the recursion

$$T(n, w) = T(nq, \frac{w}{q}) = O(n \cdot w),$$

even worse than radix sort!

### 3.3 Packed Sorting

We now show how to sort  $n$  <sup>b-bit</sup> keys in  $O(n)$  time, provided that the word size of the computer is at least  $2(b+1) \lg n \lg \lg n$ . The idea is to pack  $k := \lg n \lg \lg n$  keys into one computer word, sort them with bit-tricks in  $O(k)$  time, and then merge the  $\frac{n}{k}$  sorted lists in additional  $O(n)$  time.



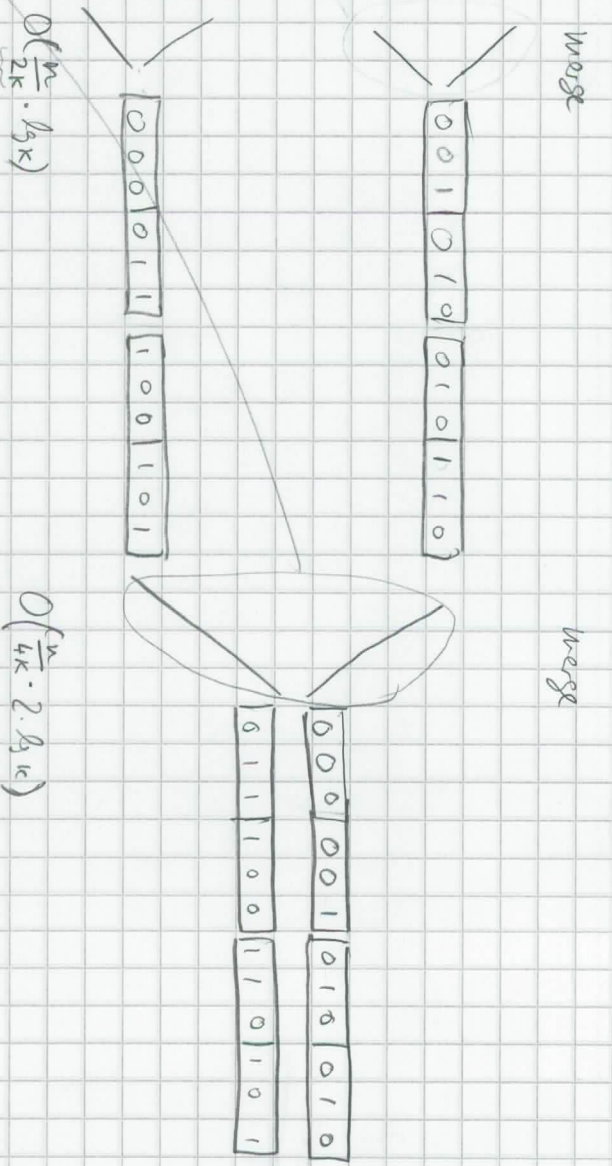


Example keys = 010, 010, 010, 001, 010, 101, 0,

b = 2, n = 8, k = 2




$O(\frac{n}{k} \cdot k) = O(n)$



$O(\frac{n}{2k} \cdot 2k)$

$O(\frac{n}{4k} \cdot 2 \cdot 2k)$

"normal" merge sort with  $\frac{n}{k}$  leaves as base case in  $O(k)$  time and merging of  $r$  sorted words with  $k$  keys each in  $O(r \lg k)$  time.



8

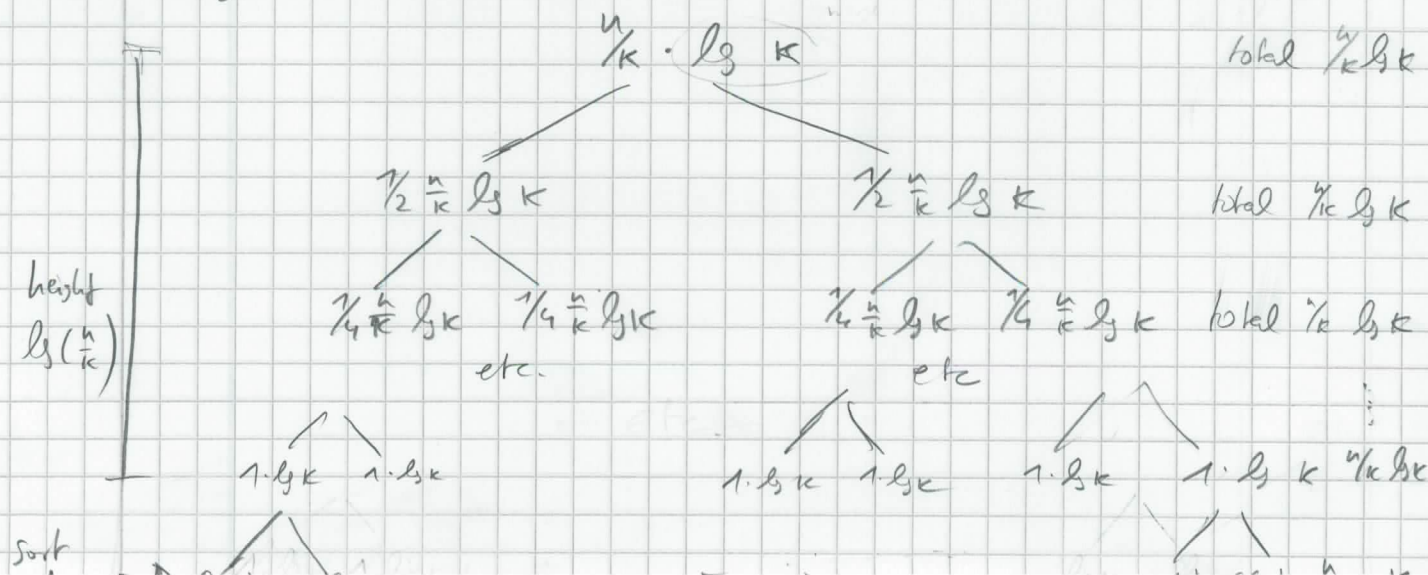
Suppose we have a black box for merging a pair of words consisting of  $k$  sorted elements each into one word consisting of  $2k$  sorted elements in  $O(\lg k)$  time (instead of  $O(k)$  time with "normal" merging), and another black box for sorting 1 word of  $k$  elements in  $O(k)$  time (instead of  $k \lg k$  with "normal" merge sort).

Then we can merge 2 sorted lists of  $r$  sorted words into one sorted list of  $2r$  sorted words in  $O(r \lg k)$  time by

1. merging the first words of each list in  $O(\lg k)$
2. outputting the first half of the result
3. replacing the <sup>merged</sup> word containing the maximum key with the second half of the result
4. removing the other merged word (not containing the max)
5. if both lists <sup>are</sup> nonempty, goes back to 1.
6. outputting all remaining words in order.

This algorithm spends  $O(\lg k)$  time per word in the output, so the running time is  $O(r \lg k)$ .

Then we can also merge all  $\frac{n}{k}$  <sup>sorted</sup> words in  $O(n)$  time by a standard merge-sort with the  $r \lg k$  sub-routine as a merger.






9) Hence, the total time is  $O(\underbrace{\lg \frac{n}{k}}_{\text{height}} \cdot \underbrace{\frac{n}{k} \lg k}_{\text{merge per level}} + \underbrace{\frac{n}{k} \cdot k}_{\text{sort leaves}})$   
 $= O(\frac{n}{k} \lg k \lg n + n)$   
 $= O(n)$  with  $k = \lg n \lg \lg n$ .

### 3.4. Bitonic Sorting

The base case (sorting a <sup>packed</sup> word consisting of  $k$  keys) is handled with a bit-parallel version of bitonic sorting.

(Def) A sequence  $x_1, \dots, x_n$  is called bitonic if it is a cyclic shift of the concatenation of an increasing and a decreasing sequence.

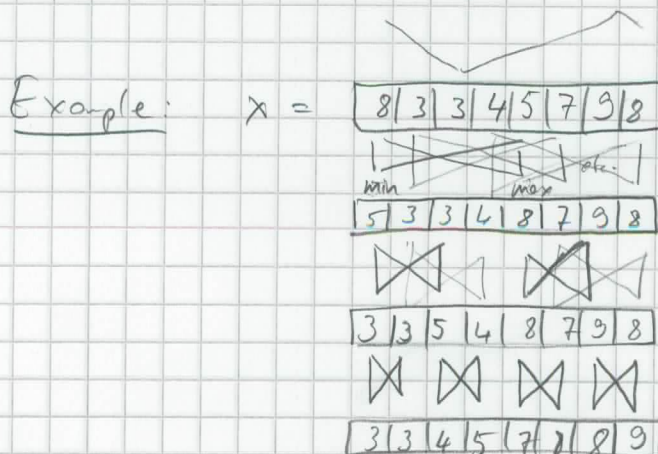


Bitonic sequences are easy to sort in parallel:

procedure sort ( $x_1, \dots, x_n$ ):  
 for  $i = 1$  to  $n/2$

if  $(x_i > x_{i+n/2})$  swap  $x_i$  with  $x_{i+n/2}$

sort ( $x_1, \dots, x_{n/2}$ ) and  $(x_{n/2+1}, \dots, x_n)$  recursively in parallel




10

The correctness of this algorithm follows from the following 2 facts:

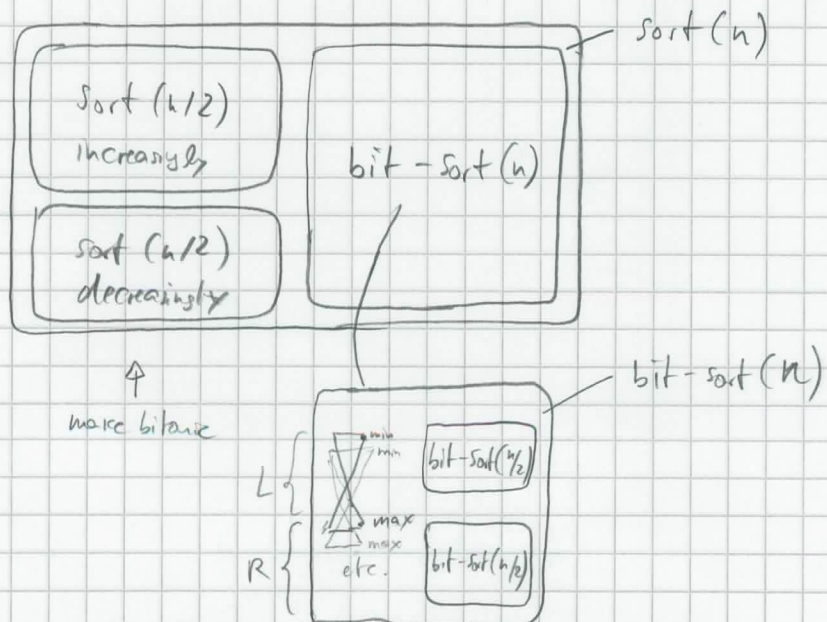
1) both  $(\min(x_1, x_{n/2+1}), \dots, \min(x_{n/2}, x_n)) =: L$   
and  $(\max(x_1, x_{n/2+1}), \dots, \max(x_{n/2}, x_n)) =: R$   
are bitonic

2) all elements in  $L$  are smaller than those in  $R$

This can be proved by looking at all possibilities of where the half falls into, e.g.



The idea for sorting any sequences (not necessarily bitonic) is now to first make them bitonic by sorting the first half increasingly and the second decreasingly, and then use the bitonic-sorter to finish.



[Note the running time of  $O(n \log^2 n)$  in the sequential case.]




11


### 3.5. Black Box 1: Merging Two Sorted Words.

We need to show how to merge two sorted words of  $k = \lg n$  by  $n$  elements into one sorted word consisting of  $2k$  elements. To adopt the ideas from bitonic sorting, we first reverse the first word ( $n$  bit-parallel way) and then concatenate the two words to get a bitonic sequence. This sequence can be sorted with a bit-parallel version of bitonic sorting — the result is the merged word!


To reverse the keys in a word, observe that  $\text{rev}(LR) = \text{rev}(R) \cdot \text{rev}(L)$ . Hence we can do the following.



and likewise with R to get



But this can be done in parallel by




And so on until all elements are reversed; this loops after  $\lg(k)$  iterations.

We now come to the last step of the merging algorithm: the bit-parallel version of bitonic sorting.

12

To this end, separate the reversed keys from the first word into one word, separated by 1-bits, and the keys from the second word B by 0-bits. Then a subtraction yields a parallel comparison (like with fusion trees):



$$\begin{array}{r}
 A = \overbrace{1 \mid A\text{-key}_1 \mid 1 \mid A\text{-key}_2 \mid 1 \mid A\text{-key}_3 \mid 1 \mid \dots \mid 1 \mid A\text{-key}_k}^{w/2} \mid \overbrace{0 \dots 0}^{w/2} \\
 B = \overbrace{0 \mid B\text{-key}_1 \mid 0 \mid B\text{-key}_2 \mid 0 \mid B\text{-key}_3 \mid 0 \mid \dots \mid 0 \mid B\text{-key}_k}^{w/2} \mid \overbrace{0 \dots 0}^{w/2} \\
 \hline
 A-B = \overbrace{1 \mid ? \dots ? \mid 1 \mid ? \dots ? \mid 1 \mid ? \dots ? \mid 1 \mid \dots \mid 1 \mid ? \dots ?}^{w/2} \mid \dots \mid \overbrace{1 \mid ? \dots ?}^{w/2}
 \end{array}$$

0 iff b-key > a-key

Now the aim is to construct a word containing exactly the smaller keys, (and another one containing exactly the larger ones). That is, we need a mask M with 1's at exactly those positions where A-B has 1-separators, and AND M with B. This gives those keys from B that are smaller than their counterparts in A:

$$\begin{array}{r}
 A-B = \overbrace{0 \dots 1 \dots 1 \dots 1 \dots 0} \\
 (A-B) \cdot [ (A-B) \gg 1 ] = M = \overbrace{0 \dots 0 \dots 1 \dots 1 \dots 1 \dots 0} = A \\
 B \cdot \text{AND } M = \overbrace{0 \dots 0 \mid B\text{-key}_2 \mid 0 \mid B\text{-key}_3 \mid 0 \dots 0} =: \text{SMALL-B}
 \end{array}$$

For the smaller keys from A, we do the following:

$$\begin{array}{r}
 A-B = \overbrace{0 \dots 1 \dots 1 \dots 1 \dots 0} \\
 \neg M = \overbrace{1 \dots 1 \dots 1 \dots 0 \dots 0 \dots 1 \dots 0 \dots 1 \dots 1 \dots 1} \\
 M' = \overbrace{0 \dots 1 \dots 1 \dots 0 \dots 1 \dots 1 \dots 0 \dots 1 \dots 1 \dots 0 \dots 0 \dots 1 \dots 1} \\
 M'' = \neg M \text{ AND } M' = \overbrace{0 \dots 1 \dots 1 \dots 0 \dots 0 \dots 0 \dots 0 \dots 0 \dots 0 \dots 0 \dots 1 \dots 1} \\
 A \text{ AND } M'' = \overbrace{0 \mid A\text{-key}_1 \mid 0 \dots 0 \dots 0 \dots 0 \mid A\text{-key}_k} =: \text{SMALL-A}
 \end{array}$$

$$\text{SMALL-A OR SMALL-B} = \overbrace{0 \mid A\text{-key}_1 \mid 0 \mid B\text{-key}_2 \mid 0 \mid B\text{-key}_3 \mid 0 \dots 0 \mid A\text{-key}_k} =: \text{SMALLS}$$


We perform a symmetric calculation to obtain a word LARGES




13

containing those keys that are larger than their counterpart.

For the next round, the goal is to compare the elements from SMALLS and LARGES separately:




Like with swapping, this can be done in parallel:



And so on. In total, we need  $\lg k$  rounds, each taking  $O(1)$  time. Hence we can merge in  $O(\lg k)$  time.

### 3.6 Black Box 2: Sorting a Packed Word

We can simulate merge sort within a word:



Merging on the last level works like black box 1, but only on two words of half the size  $\Rightarrow O(\lg(k))$  time

(14)

Likewise, merging on any level can be implemented like block box 1, with the appropriate bits masked. Hence, the running time follows the recursion

$$\begin{aligned} T(k) &= O(\lg k) + 2 T(k/2) \\ &= O(k). \end{aligned}$$