8. Universal algorithms

8.1 Sorting algorithms

Bubble sort is a simple sorting algorithm. It is a comparison sort in which the sorted sequence is created one element at a time by finding the minimum value in the unsorted portion and moving it to the sorted portion of the sequence. The algorithm gets its name from the way smaller elements "bubble" to the top of the list.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function bubbleSort(A : list of sortable items)
    n = length(A)
    repeat
        swapped = false
        for i = 1 to n - 1 inclusive do
            if A[i - 1] > A[i] then
                swap(A[i - 1], A[i])
                swapped = true
            end if
        end for
    until not swapped
end function
```

Selection sort is a sorting algorithm, specifically an in-place comparison sort. It has $O(n^2)$ time complexity, making it inefficient on large lists, and generally performs worse than the similar insertion sort. Selection sort is noted for its simplicity, and also has performance advantages over more complicated algorithms in certain situations, particularly where auxiliary memory is limited.

Time complexity: $O(n^2)$ as there are two nested loops.

Space complexity: $O(1)$ (in-place)

```
function selectionSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 inclusive do
        min = i
        for j = i + 1 to n - 1 inclusive do
            if A[j] < A[min] then
                min = j
            end if
        end for
        if min ≠ i then
            swap(A[i], A[min])
        end if
    end for
end function
```

Insertion sort is a simple sorting algorithm that builds the final sorted array (or list) one item at a time. It is much less efficient on large lists than more advanced algorithms such as quicksort, heapsort, or merge sort. However, insertion sort provides several advantages: simple implementation, efficient for (quite) small data sets, much like other quadratic sorting algorithms, more efficient in practice than most other simple quadratic algorithms such as selection sort or bubble sort, adaptive, i.e., efficient for data sets that are already substantially sorted: the time complexity is $O(kn)$ when each element in the input is no more than k places away from its sorted position, stable; i.e., does not change the relative order of elements with equal keys, in-place; i.e., only requires a constant amount $O(1)$ of additional memory space, online; i.e., can sort a list as it receives it.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function insertionSort(A : list of sortable items)
    n = length(A)
    for i = 1 to n - 1 inclusive do
        j = i
        while j > 0 and A[j - 1] > A[j] do
            swap(A[j], A[j - 1])
            j = j - 1
        end while
    end for
end function
```

Merge sort is an efficient, general-purpose, comparison-based sorting algorithm. Most implementations produce a stable sort, which means that the implementation preserves the input order of equal elements in the sorted output. Merge sort is a divide and conquer algorithm that was invented by John von Neumann in 1945. A detailed description and analysis of bottom-up merge sort appeared in a report by Goldstine and von Neumann as early as 1948.

Time complexity: $O(n \log n)$ in worst case

Space complexity: $O(n)$

```
function mergeSort(m)
    var list left, right, result
    if length(m) ≤ 1
        return m
    else
        var middle = length(m) / 2
        for each x in m up to middle - 1
            add x to left
        for each x in m at middle to last
            add x to right
        left = mergeSort(left)
        right = mergeSort(right)
        result = merge(left, right)
        return result
```

There is a simple implementation of merge function, allocating a new list ("result") for the return value, and destroying the two input lists in the process:

```
function merge(left, right)
    var list result
    while length(left) > 0 and length(right) > 0
        if first(left) ≤ first(right)
            append first(left) to result
            left = rest(left)
        else
            append first(right) to result
            right = rest(right)
    if length(left) > 0
        append left to result
    if length(right) > 0
        append right to result
    return result
```

Quicksort is a sorting algorithm developed by Tony Hoare that, on average, makes $O(n \log n)$ comparisons to sort n items. In the worst case, it makes $O(n^2)$ comparisons, though this behavior is rare. Quicksort is often faster in practice than other $O(n \log n)$ algorithms. Additionally, quicksort's sequential and localized memory references work well with a cache. Quicksort is a comparison sort and, in efficient implementations, is not a stable sort. Quicksort can be implemented with an in-place partitioning algorithm, so the entire sort can be done with only $O(\log n)$ additional space used by the stack during the recursion.

Time complexity: $O(n^2)$ in worst case, $O(n \log n)$ in best case

Space complexity: $O(n)$

```
function quicksort(q)
    var list less, pivotList, greater
    if length(q) ≤ 1
        return q
    else
        select a pivot value pivot from q
        for each x in q
            if x < pivot then add x to less
            if x = pivot then add x to pivotList
            if x > pivot then add x to greater
        return concatenate(quicksort(less), pivotList, quicksort(greater))
```

Heapsort is a comparison-based sorting algorithm to create a sorted array (or list), and is part of the selection sort family. Although somewhat slower in practice on most machines than a well-implemented quicksort, it has the advantage of a more favorable worst-case $O(n \log n)$ runtime. Heapsort is an in-place algorithm, but is not a stable sort.

Time complexity: $O(n \log n)$ in worst case

Space complexity: $O(1)$ (in-place)

```
function heapSort(a, count) is
    input: an unordered array a of length count
    (Build the heap in array a so that largest value is at the root)
    heapify(a, count)
    (The following loop maintains the invariants that a[0:end] is a heap and every element
     beyond end is greater than everything before it (so a[end:count] is in sorted order))
    end ← count - 1
    while end > 0 do
        (a[0] is the root and largest value. The swap moves it in front of the sorted elements.)
        swap(a[end], a[0])
        (the heap size is reduced by one)
        end ← end - 1
        (the swap ruined the heap property, so restore it)
        siftDown(a, 0, end)
```

Bucket sort is a sorting algorithm that works by partitioning an array into a number of buckets. Each bucket is then sorted individually, either using a different sorting algorithm, or by recursively applying the bucket sorting algorithm. It is a distribution sort, and is a cousin of radix sort in the most to least significant digit flavour. Bucket sort is a generalization of pigeonhole sort. Bucket sort can be implemented with comparisons and therefore can also be considered a comparison sort algorithm. The computational complexity estimates involve the number of buckets.

Time complexity: $O(n + k)$ in worst case, where k is the number of buckets

Space complexity: $O(n)$

```
function bucketSort(A, n) is
    buckets ← new array of n empty lists
    for i = 0 to (length(A) - 1) do
        insert A[i] into buckets[floor(n × A[i])]
    for i = 0 to (n - 1) do
        nextSort(buckets[i])
    return the concatenation of buckets[0], ..., buckets[n - 1]
```

Radix sort is a non-comparative integer sorting algorithm that sorts data with integer keys by grouping keys by the individual digits which share the same significant position and value. A positional notation is required, but because integers can represent strings of characters (e.g., names or dates) and specially formatted floating point numbers, radix sort is not limited to integers. Radix sort dates back as far as 1887 to the punch card sorting machines.

Time complexity: $O(nk)$ in worst case, where k is the number of passes over the keys

Space complexity: $O(n + k)$

```
function radixSort(A, d) is
    for i = 1 to d do
        allocate buckets[0 ... (n - 1)]
        for j = 0 to (n - 1) do
            place A[j] at buckets[(A[j] ÷ (n^(i - 1))) mod n]
        A ← concatenate(buckets[0], ..., buckets[n - 1])
```

Counting sort is a sorting algorithm which assumes that each of the n input elements is an integer with value between 0 and some integer k. When $k = O(n)$, the sort runs in $O(n)$ time. Counting sort is not a comparison sort, and it is not limited by the $O(n \log n)$ lower bound for comparison sorts. Counting sort is a stable sort.

Time complexity: $O(n + k)$ in worst case, where k is the number of passes over the keys

Space complexity: $O(n + k)$

```
function countingSort(A, k) is
    count ← array of k + 1 zeros
    for i = 0 to (length(A) - 1) do
        count[A[i]] ← count[A[i]] + 1
    sum ← 0
    for i = 0 to k do
        c ← count[i]
        count[i] ← sum
        sum ← sum + c
    B ← array of same length as A
    for i = 0 to (length(A) - 1) do
        B[count[A[i]]] ← A[i]
        count[A[i]] ← count[A[i]] + 1
    return B
```

Topological sorting is a sorting algorithm for directed graphs. For example, a partial ordering on the set of all points in the plane is the containment order of disks: given two disks in the plane, they are ordered if one of them contains the other, and unordered otherwise. A topological sort of this ordering would be a sequence of the disks such that no disk contains any of the previous disks. Topological sorting is not possible if the graph is not a directed acyclic graph (DAG). Any DAG has at least one topological ordering, and there are linear time algorithms for finding such an ordering.

Time complexity: $O(n + m)$ in worst case, where m is the number of edges

Space complexity: $O(n)$

```
L ← Empty list that will contain the sorted nodes
while there are unmarked nodes do
    select an unmarked node n
    visit(n)
function visit(node n)
    if n has a permanent mark then return
    if n has a temporary mark then stop (not a DAG)
    mark n with a temporary mark
    for each node m with an edge from n to m do
        visit(m)
    remove temporary mark from n
    mark n with a permanent mark
    add n to head of L
```

Pancake sorting is the mathematical problem of sorting a disordered stack of pancakes in order of size when a spatula can be inserted at any pancake and all pancakes above it flipped. The pancake problem was first discussed by Jacob E. Goodman. A variant of the problem is concerned with the minimum number of flips required to sort a stack, where each pancake is marked with a burnt side. This variant is of theoretical interest only; it is not edible.

Time complexity: $O(n)$ in worst case

Space complexity: $O(1)$ (in-place)

```
function pancakeSort(A : list of sortable items)
    n = length(A)
    for i = n - 1 downto 1 inclusive do
        m = index of largest element in A[1..i]
        if m ≠ i then
            flip(A, m)
            flip(A, i)
    end for
end function
```

Spaghetti sort is a sorting algorithm named after spaghetti because it is reminiscent of placing a bundle of spaghetti in front of you and then lowering it to the table. It is a linear time, analog algorithm for sorting a sequence of items, requiring $O(n)$ stack space, where n is the number of items to be sorted. Spaghetti sort is simply a variant of bucket sort with unit-length buckets.

Time complexity: $O(n)$ in worst case

Space complexity: $O(n)$

```
function spaghettiSort(A : list of sortable items)
    n = length(A)
    for i = 1 to n do
        spaghetti[i] = new Thread(sleep(A[i]))
    end for
    for i = 1 to n do
        wait for next thread to terminate
        print value that thread slept
    end for
end function
```

Sorting network is an abstract mathematical model of a network of wires and comparator modules that is used to sort a sequence of numbers. Each comparator connects two wires and sorts the values by outputting the smaller value to one wire and the larger value to the other. The design of sorting networks differs from that of general algorithms in that the sequence of which pairs of inputs are to be compared must be specified in advance. Sorting networks can be implemented either in hardware or in software. Sorting networks are typically evaluated according to two parameters: their size and their depth. Size is the number of comparators in the network, and depth is the maximum number of comparators that any input value can pass through. Sorting networks are also known as Batcher networks, after Ken Batcher who first described the odd-even mergesort and bitonic sorter in 1968.

Time complexity: $O(n)$ in worst case

Space complexity: $O(n)$

```
for each i = 1 to n do
    spaghetti[i] = new Thread(sleep(A[i]))
for each i = 1 to n do
    wait for next thread to terminate
    print value that thread slept
```

Bitonic sorter is a sorting network that is particularly efficient on sequences with lengths that are powers of two. It was originally described by Batcher, who called it a bitonic sorting network. A sorting network is a sorting algorithm that is suitable for implementation in hardware or parallel processors. Bitonic sorter is a parallel algorithm, but not all parallel sorting algorithms are sorting networks. Bitonic sorter is also notable for being relatively easy to implement in hardware, because the number of comparisons and the pattern of which elements are compared does not depend on the data at all.

Time complexity: $O(n)$ in worst case

Space complexity: $O(n)$

```
function bitonicSort(A : list of sortable items)
    n = length(A)
    for i = 1 to n do
        for j = 1 to n do
            if (i & 1) = (j & 1) then
                swap(A[i], A[j])
            end if
        end for
    end for
end function
```

Bogosort is a sorting algorithm based on the generate and test paradigm. The function successively generates permutations of its input until it finds one that is sorted. It is not useful for sorting, but may be used for educational purposes, to contrast it with more efficient algorithms. Two versions of this algorithm exist: a deterministic version that enumerates all permutations until it hits a sorted one, and a randomized version that randomly permutes its input. An analogy for the randomized version is that it is the way one would sort cards by throwing them in the air, picking them up randomly, and then repeating the procedure until the cards are sorted.

Time complexity: $O(n!)$ in worst case

Space complexity: $O(1)$ (in-place)

```
while not isSorted(deck) do
    deck := random permutation of deck
```

Stooge sort is a recursive sorting algorithm with poor time complexity. It is defined as follows:

```
function stoogesort(A, i, j) is
    if A[j] < A[i] then
        swap A[i] and A[j]
    if (j - i + 1) > 2 then
        t := (j - i + 1) / 3
        stoogesort(A, i, j - t)
        stoogesort(A, i + t, j)
        stoogesort(A, i, j - t)
    return A
```

Time complexity: $O(n^{\log 3 / \log 1.5}) = O(n^{2.7095...})$ in worst case

Space complexity: $O(n)$

Slowsort is a sorting algorithm. It is of humorous nature and not useful. It is a multiply and surrender algorithm, following the principle of divide and conquer. It was published in 1986 by Andrei Broder and Jorge Stolfi in the paper *Pessimal Algorithms and Simplexity Analysis* (a parody of optimal algorithms and complexity analysis). The algorithm is defined as follows:

```
function slowsort(A, i, j) is
    if i ≥ j then
        return
    m := floor((i + j) / 2)
    slowsort(A, i, m)
    slowsort(A, m + 1, j)
    if A[j] < A[m] then
        swap A[j] and A[m]
    slowsort(A, i, j - 1)
```

Time complexity: $O(n^{\log n})$ in worst case

Space complexity: $O(n)$

Gnome sort is a sorting algorithm which is similar to insertion sort, except that moving an element to its proper place is accomplished by a series of swaps, as in bubble sort. It is conceptually simple, requiring no nested loops. The time required by gnome sort to sort an array is proportional to $\sum_{i=0}^{n-1} (1 + |g_i|)$ where g_i is the distance a gnome i must move to the left after placing stone i in the correct position. The average value of $|g_i|$ is $n/3$, so gnome sort is an $O(n^2)$ algorithm. The best case is when the array is already sorted, in which case gnome sort is $O(n)$.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function gnomeSort(a[0..size - 1])
    i := 1
    while i < size do
        if a[i - 1] ≤ a[i] then
            i := i + 1
        else
            swap a[i - 1] and a[i]
            i := i - 1
            if i = 0 then
                i := 1
```

Comb sort is a relatively simple sorting algorithm originally designed by Włodzimierz Dobosiewicz in 1980. Later it was rediscovered by Stephen Lacey and Richard Box in 1991. Comb sort improves on bubble sort. The basic idea is to eliminate turtles, or small values near the end of the list, since in a bubble sort these slow the sorting down tremendously. Rabbits, large values around the beginning of the list, do not pose a problem in bubble sort.

Time complexity: $O(n^2)$ in worst case, $O(n \log n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function combSort(A : list of sortable items)
    gap = length(A)
    shrink = 1.3
    sorted = false
    while sorted = false do
        gap = floor(gap / shrink)
        if gap ≤ 1 then
            gap = 1
            sorted = true
        end if
        i = 0
        while i + gap < length(A) do
            if A[i] > A[i + gap] then
                swap(A[i], A[i + gap])
                sorted = false
            end if
            i = i + 1
        end while
    end while
end function
```

Odd-even sort is a relatively simple sorting algorithm, developed originally for use on parallel processors with local interconnections. It works by comparing all (odd, even) indexed pairs of adjacent elements in the list and, if a pair is in the wrong order, swapping the elements. The next step repeats this for (even, odd) indexed pairs. It then alternates between (odd, even) and (even, odd) steps until the list is sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function oddEvenSort(A : list of sortable items)
    n = length(A)
    sorted = false
    while sorted = false do
        sorted = true
        for i = 1 to n - 2 inclusive step 2 do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                sorted = false
            end if
        end for
        for i = 0 to n - 2 inclusive step 2 do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                sorted = false
            end if
        end for
    end while
end function
```

Cocktail sort is a variation of bubble sort that is both a stable sorting algorithm and a comparison sort. It differs from a bubble sort in that it sorts in both directions on each pass through the list. This sorting algorithm is only marginally more difficult to implement than a bubble sort, and solves the problem of turtles in bubble sort. It provides only a slight improvement in efficiency over bubble sort, however; while turtles are no longer a problem, rabbits are still very slow.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function cocktailSort(A : list of sortable items)
    n = length(A)
    repeat
        swapped = false
        for i = 1 to n - 2 inclusive do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                swapped = true
            end if
        end for
        if swapped = false then
            break do-while loop
        end if
        swapped = false
        for i = n - 2 downto 1 inclusive do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                swapped = true
            end if
        end for
    until swapped = false
end function
```

Binary tree sort is a sorting algorithm that builds a binary search tree from the elements to be sorted, and then traverses the tree in order to generate the sorted list. It is the typical way of sorting elements in augmented binary search tree data structure. Binary tree sort is a form of selection sort. Adding one item to a balanced binary search tree is $O(\log n)$, so adding n items is $O(n \log n)$. Binary tree sort can be used as an online algorithm, since it is possible to add any unsorted element to a binary search tree at any time.

Time complexity: $O(n \log n)$ in worst case

Space complexity: $O(n)$

```
function binaryTreeSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 inclusive do
        insert(A[i])
    end for
    for i = 0 to n - 1 inclusive do
        A[i] = findMin()
        deleteMin()
    end for
end function
```

Pigeonhole sort is a sorting algorithm that is suitable for sorting lists of elements where the number of elements and the number of possible key values are approximately the same. It requires $O(n + k)$ time where n is the number of elements and k is the number of possible values. Pigeonhole sort is not a comparison sort, and it is not limited by the $O(n \log n)$ lower bound for comparison sorts.

Time complexity: $O(n + k)$ in worst case, where k is the number of passes over the keys

Space complexity: $O(n + k)$

```
function pigeonholeSort(A, k) is
    holes ← array of k + 1 empty lists
    for i = 0 to (length(A) - 1) do
        insert A[i] into holes[A[i]]
    for i = 0 to k do
        nextSort(holes[i])
    return the concatenation of holes[0], ..., holes[k]
```

Cycle sort is an in-place, unstable sorting algorithm, a comparison sort that is theoretically optimal in terms of the total number of writes to the original array, unlike any other in-place sorting algorithm. It is based on the idea that the permutation to be sorted can be factored into cycles, which can individually be rotated to give a sorted result. It is theoretically optimal in the sense that it uses the minimal number of writes (overwrites) to the original array, one fewer than the number of elements, regardless of the initial ordering. Minimizing the number of writes is useful when writing to memory or to flash storage is expensive or limited, such as with EEPROMs or flash memory.

Time complexity: $O(n^2)$ in worst case

Space complexity: $O(1)$ (in-place)

```
function cycleSort(A : list of sortable items)
    n = length(A)
    for cycleStart = 0 to n - 2 inclusive do
        item = A[cycleStart]
        pos = cycleStart
        for i = cycleStart + 1 to n - 1 inclusive do
            if A[i] < item then
                pos = pos + 1
            end if
        end for
        if pos = cycleStart then
            continue
        end if
        while item = A[pos] do
            pos = pos + 1
        end while
        swap(item, A[pos])
        while pos ≠ cycleStart do
            pos = cycleStart
            for i = cycleStart + 1 to n - 1 inclusive do
                if A[i] < item then
                    pos = pos + 1
                end if
            end for
            while item = A[pos] do
                pos = pos + 1
            end while
            swap(item, A[pos])
        end while
    end for
end function
```

Strand sort is a sorting algorithm that works by repeatedly pulling sorted subsequences out of the list to be sorted and merging them with a result array. It is a comparison sort with $O(n^2)$ worst case behavior, but it is adaptive and can perform reasonably well when the input list is nearly sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(n)$

```
function strandSort(A : list of sortable items)
    n = length(A)
    result = new empty list
    while length(A) > 0 do
        sublist = new list with A[0]
        remove A[0] from A
        for i = 0 to length(A) - 1 do
            if A[i] > sublist[length(sublist) - 1] then
                append A[i] to sublist
                remove A[i] from A
            end if
        end for
        merge sublist into result
    end while
    return result
end function
```

Bead sort is a sorting algorithm that is inspired by gravity, or alternatively by bubbles in water. It is also known as gravity sort. The algorithm is as follows:

```
function beadSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 do
        for j = 0 to A[i] - 1 do
            add a bead to column j
        end for
    end for
    for i = 0 to n - 1 do
        for j = 0 to A[i] - 1 do
            remove a bead from column j
        end for
    end for
end function
```

Time complexity: $O(n)$ in worst case

Space complexity: $O(n)$

Simple pancake sort is a sorting algorithm that is inspired by the manner in which one sometimes sorts a stack of pancakes. It is a comparison sort with $O(n^2)$ worst case behavior, but it is adaptive and can perform reasonably well when the input list is nearly sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function pancakeSort(A : list of sortable items)
    n = length(A)
    for i = n - 1 downto 1 inclusive do
        m = index of largest element in A[1..i]
        if m ≠ i then
            flip(A, m)
            flip(A, i)
    end for
end function
```

Stupid sort is a sorting algorithm that is inspired by the manner in which one sometimes sorts a stack of pancakes. It is a comparison sort with $O(n^2)$ worst case behavior, but it is adaptive and can perform reasonably well when the input list is nearly sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function stupidSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 do
        for j = i + 1 to n - 1 do
            if A[i] > A[j] then
                swap(A[i], A[j])
            end if
        end for
    end for
end function
```

Bogosort is a sorting algorithm based on the generate and test paradigm. The function successively generates permutations of its input until it finds one that is sorted. It is not useful for sorting, but may be used for educational purposes, to contrast it with more efficient algorithms. Two versions of this algorithm exist: a deterministic version that enumerates all permutations until it hits a sorted one, and a randomized version that randomly permutes its input. An analogy for the randomized version is that it is the way one would sort cards by throwing them in the air, picking them up randomly, and then repeating the procedure until the cards are sorted.

Time complexity: $O(n!)$ in worst case

Space complexity: $O(1)$ (in-place)

```
while not isSorted(deck) do
    deck := random permutation of deck
```

Stooge sort is a recursive sorting algorithm with poor time complexity. It is defined as follows:

```
function stoogesort(A, i, j) is
    if A[j] < A[i] then
        swap A[i] and A[j]
    if (j - i + 1) > 2 then
        t := (j - i + 1) / 3
        stoogesort(A, i, j - t)
        stoogesort(A, i + t, j)
        stoogesort(A, i, j - t)
    return A
```

Time complexity: $O(n^{\log 3 / \log 1.5}) = O(n^{2.7095...})$ in worst case

Space complexity: $O(n)$

Slowsort is a sorting algorithm. It is of humorous nature and not useful. It is a multiply and surrender algorithm, following the principle of divide and conquer. It was published in 1986 by Andrei Broder and Jorge Stolfi in the paper *Pessimal Algorithms and Simplexity Analysis* (a parody of optimal algorithms and complexity analysis). The algorithm is defined as follows:

```
function slowsort(A, i, j) is
    if i ≥ j then
        return
    m := floor((i + j) / 2)
    slowsort(A, i, m)
    slowsort(A, m + 1, j)
    if A[j] < A[m] then
        swap A[j] and A[m]
    slowsort(A, i, j - 1)
```

Time complexity: $O(n^{\log n})$ in worst case

Space complexity: $O(n)$

Gnome sort is a sorting algorithm which is similar to insertion sort, except that moving an element to its proper place is accomplished by a series of swaps, as in bubble sort. It is conceptually simple, requiring no nested loops. The time required by gnome sort to sort an array is proportional to $\sum_{i=0}^{n-1} (1 + |g_i|)$ where g_i is the distance a gnome i must move to the left after placing stone i in the correct position. The average value of $|g_i|$ is $n/3$, so gnome sort is an $O(n^2)$ algorithm. The best case is when the array is already sorted, in which case gnome sort is $O(n)$.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function gnomeSort(a[0..size - 1])
    i := 1
    while i < size do
        if a[i - 1] ≤ a[i] then
            i := i + 1
        else
            swap a[i - 1] and a[i]
            i := i - 1
            if i = 0 then
                i := 1
```

Comb sort is a relatively simple sorting algorithm originally designed by Włodzimierz Dobosiewicz in 1980. Later it was rediscovered by Stephen Lacey and Richard Box in 1991. Comb sort improves on bubble sort. The basic idea is to eliminate turtles, or small values near the end of the list, since in a bubble sort these slow the sorting down tremendously. Rabbits, large values around the beginning of the list, do not pose a problem in bubble sort.

Time complexity: $O(n^2)$ in worst case, $O(n \log n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function combSort(A : list of sortable items)
    gap = length(A)
    shrink = 1.3
    sorted = false
    while sorted = false do
        gap = floor(gap / shrink)
        if gap ≤ 1 then
            gap = 1
            sorted = true
        end if
        i = 0
        while i + gap < length(A) do
            if A[i] > A[i + gap] then
                swap(A[i], A[i + gap])
                sorted = false
            end if
            i = i + 1
        end while
    end while
end function
```

Odd-even sort is a relatively simple sorting algorithm, developed originally for use on parallel processors with local interconnections. It works by comparing all (odd, even) indexed pairs of adjacent elements in the list and, if a pair is in the wrong order, swapping the elements. The next step repeats this for (even, odd) indexed pairs. It then alternates between (odd, even) and (even, odd) steps until the list is sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function oddEvenSort(A : list of sortable items)
    n = length(A)
    sorted = false
    while sorted = false do
        sorted = true
        for i = 1 to n - 2 inclusive step 2 do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                sorted = false
            end if
        end for
        for i = 0 to n - 2 inclusive step 2 do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                sorted = false
            end if
        end for
    end while
end function
```

Cocktail sort is a variation of bubble sort that is both a stable sorting algorithm and a comparison sort. It differs from a bubble sort in that it sorts in both directions on each pass through the list. This sorting algorithm is only marginally more difficult to implement than a bubble sort, and solves the problem of turtles in bubble sort. It provides only a slight improvement in efficiency over bubble sort, however; while turtles are no longer a problem, rabbits are still very slow.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function cocktailSort(A : list of sortable items)
    n = length(A)
    repeat
        swapped = false
        for i = 1 to n - 2 inclusive do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                swapped = true
            end if
        end for
        if swapped = false then
            break do-while loop
        end if
        swapped = false
        for i = n - 2 downto 1 inclusive do
            if A[i] > A[i + 1] then
                swap(A[i], A[i + 1])
                swapped = true
            end if
        end for
    until swapped = false
end function
```

Binary tree sort is a sorting algorithm that builds a binary search tree from the elements to be sorted, and then traverses the tree in order to generate the sorted list. It is the typical way of sorting elements in augmented binary search tree data structure. Binary tree sort is a form of selection sort. Adding one item to a balanced binary search tree is $O(\log n)$, so adding n items is $O(n \log n)$. Binary tree sort can be used as an online algorithm, since it is possible to add any unsorted element to a binary search tree at any time.

Time complexity: $O(n \log n)$ in worst case

Space complexity: $O(n)$

```
function binaryTreeSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 inclusive do
        insert(A[i])
    end for
    for i = 0 to n - 1 inclusive do
        A[i] = findMin()
        deleteMin()
    end for
end function
```

Pigeonhole sort is a sorting algorithm that is suitable for sorting lists of elements where the number of elements and the number of possible key values are approximately the same. It requires $O(n + k)$ time where n is the number of elements and k is the number of possible values. Pigeonhole sort is not a comparison sort, and it is not limited by the $O(n \log n)$ lower bound for comparison sorts.

Time complexity: $O(n + k)$ in worst case, where k is the number of passes over the keys

Space complexity: $O(n + k)$

```
function pigeonholeSort(A, k) is
    holes ← array of k + 1 empty lists
    for i = 0 to (length(A) - 1) do
        insert A[i] into holes[A[i]]
    for i = 0 to k do
        nextSort(holes[i])
    return the concatenation of holes[0], ..., holes[k]
```

Cycle sort is an in-place, unstable sorting algorithm, a comparison sort that is theoretically optimal in terms of the total number of writes to the original array, unlike any other in-place sorting algorithm. It is based on the idea that the permutation to be sorted can be factored into cycles, which can individually be rotated to give a sorted result. It is theoretically optimal in the sense that it uses the minimal number of writes (overwrites) to the original array, one fewer than the number of elements, regardless of the initial ordering. Minimizing the number of writes is useful when writing to memory or to flash storage is expensive or limited, such as with EEPROMs or flash memory.

Time complexity: $O(n^2)$ in worst case

Space complexity: $O(1)$ (in-place)

```
function cycleSort(A : list of sortable items)
    n = length(A)
    for cycleStart = 0 to n - 2 inclusive do
        item = A[cycleStart]
        pos = cycleStart
        for i = cycleStart + 1 to n - 1 inclusive do
            if A[i] < item then
                pos = pos + 1
            end if
        end for
        if pos = cycleStart then
            continue
        end if
        while item = A[pos] do
            pos = pos + 1
        end while
        swap(item, A[pos])
        while pos ≠ cycleStart do
            pos = cycleStart
            for i = cycleStart + 1 to n - 1 inclusive do
                if A[i] < item then
                    pos = pos + 1
                end if
            end for
            while item = A[pos] do
                pos = pos + 1
            end while
            swap(item, A[pos])
        end while
    end for
end function
```

Strand sort is a sorting algorithm that works by repeatedly pulling sorted subsequences out of the list to be sorted and merging them with a result array. It is a comparison sort with $O(n^2)$ worst case behavior, but it is adaptive and can perform reasonably well when the input list is nearly sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(n)$

```
function strandSort(A : list of sortable items)
    n = length(A)
    result = new empty list
    while length(A) > 0 do
        sublist = new list with A[0]
        remove A[0] from A
        for i = 0 to length(A) - 1 do
            if A[i] > sublist[length(sublist) - 1] then
                append A[i] to sublist
                remove A[i] from A
            end if
        end for
        merge sublist into result
    end while
    return result
end function
```

Bead sort is a sorting algorithm that is inspired by gravity, or alternatively by bubbles in water. It is also known as gravity sort. The algorithm is as follows:

```
function beadSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 do
        for j = 0 to A[i] - 1 do
            add a bead to column j
        end for
    end for
    for i = 0 to n - 1 do
        for j = 0 to A[i] - 1 do
            remove a bead from column j
        end for
    end for
end function
```

Time complexity: $O(n)$ in worst case

Space complexity: $O(n)$

Simple pancake sort is a sorting algorithm that is inspired by the manner in which one sometimes sorts a stack of pancakes. It is a comparison sort with $O(n^2)$ worst case behavior, but it is adaptive and can perform reasonably well when the input list is nearly sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function pancakeSort(A : list of sortable items)
    n = length(A)
    for i = n - 1 downto 1 inclusive do
        m = index of largest element in A[1..i]
        if m ≠ i then
            flip(A, m)
            flip(A, i)
    end for
end function
```

Stupid sort is a sorting algorithm that is inspired by the manner in which one sometimes sorts a stack of pancakes. It is a comparison sort with $O(n^2)$ worst case behavior, but it is adaptive and can perform reasonably well when the input list is nearly sorted.

Time complexity: $O(n^2)$ in worst case, $O(n)$ in best case (if the list is already sorted)

Space complexity: $O(1)$ (in-place)

```
function stupidSort(A : list of sortable items)
    n = length(A)
    for i = 0 to n - 1 do
        for j = i + 1 to n - 1 do
            if A[i] > A[j] then
                swap(A[i], A[j])
            end if
        end for
    end for
end function
```

8.2 Searching algorithms

Linear search is a search algorithm that finds a value in a list by checking each element in turn until the desired value is found. It is the simplest search algorithm; it is a special case of brute force search. Its worst case cost is proportional to the number of elements in the list; and so is its expected cost, if all list elements are equally likely to be searched for. Therefore, linear search has worst case complexity $O(n)$ and average case complexity $O(n)$. If the list has more than a few elements, and if there is enough memory for a larger list, then it is usually better to sort the list first, and then use a binary search or a hash table. Linear search is rarely practical because other search algorithms and schemes, such as the binary search algorithm and hash tables, allow significantly faster searching for all but short lists.

Time complexity: $O(n)$ in worst case

Space complexity: $O(1)$ (in-place)

```
function linearSearch(A : list of sortable items, x : item to find)
    n = length(A)
    for i = 0 to n - 1 do
        if A[i] = x then
            return i
        end if
    end for
    return -1
end function
```

Binary search is a search algorithm that finds the position of a target value within a sorted array. Binary search compares the target value to the middle element of the array. If they are not equal, the half in which the target cannot lie is eliminated and the search continues on the remaining half, again taking the middle element to compare to the target value, and repeating this until the target value is found. If the search ends with the remaining half being empty, the target is not in the array. Binary search runs in logarithmic time in the worst case, making $O(\log n)$ comparisons, where n is the number of elements in the array. Binary search is a special case of three-way search. It is also known as half-interval search, logarithmic search, or binary chop.

Time complexity: $O(\log n)$ in worst case

Space complexity: $O(1)$ (in-place)

```
function binarySearch(A : list of sortable items, x : item to find)
    n = length(A)
    left = 0
    right = n - 1
    while left ≤ right do
        mid = floor((left + right) / 2)
        if A[mid] = x then
            return mid
        else if A[mid] < x then
            left = mid + 1
        else
            right = mid - 1
        end if
    end while
    return -1
end function
```

Jump search is a search algorithm for ordered lists. It works by first checking all items $A[k \cdot m]$, where $k \in \mathbb{N}$ and m is the block size, until an item is found that is larger than the search key. To find the exact position of the search key, the algorithm performs a linear search on the m elements $A[(k-1) \cdot m]$ to $A[k \cdot m]$. The optimal value of m is $\sqrt{n}$, where n is the length of the list A. Because both the linear search and the jump search examine $O(\sqrt{n})$ items, the algorithm runs in $O(\sqrt{n})$ time. Jump search is also known as block search.

Time complexity: $O(\sqrt{n})$ in worst case

Space complexity: $O(1)$ (in-place)

```
function jumpSearch(A : list of sortable items, x : item to find)
    n = length(A)
    m = floor(sqrt(n))
    for i = 0 to n - 1 step m do
        if A[i] ≥ x then
            for j = i - m + 1 to i do
                if A[j] = x then
                    return j
                end if
            end for
            return -1
        end if
    end for
    return -1
end function
```

Interpolation search is a search algorithm for searching in ordered arrays. It is similar to binary search, but it estimates the position of the search key in the array by interpolating between the values at the ends of the array. It works on the principle of dividing the array into two parts, and then estimating the position of the search key in the array by interpolating between the values at the ends of the array. It is also known as extrap