Stochastic Scheduling with two Processors and Arbitrary Precedence Relations

Carlos Camino

August 4, 2015

Overview

I. Introduction

II. Theoretical Foundations

III. Methods

IV. Results and Discussion

V. Conclusion

Overview

I. Introduction

II. Theoretical Foundations

III. Methods

IV. Results and Discussion

V. Conclusion

What is scheduling?

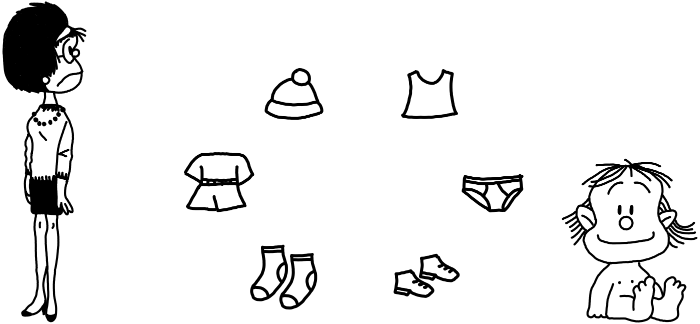

What is scheduling?

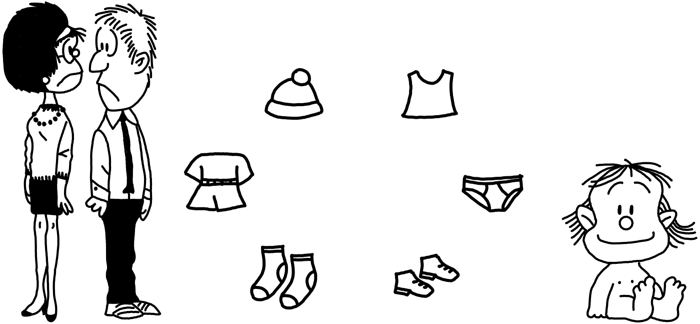

What is scheduling?

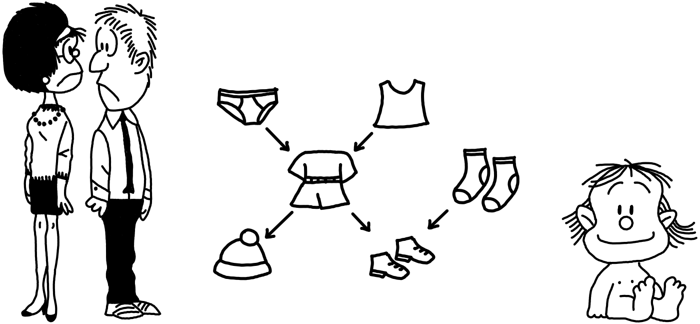

What is scheduling?

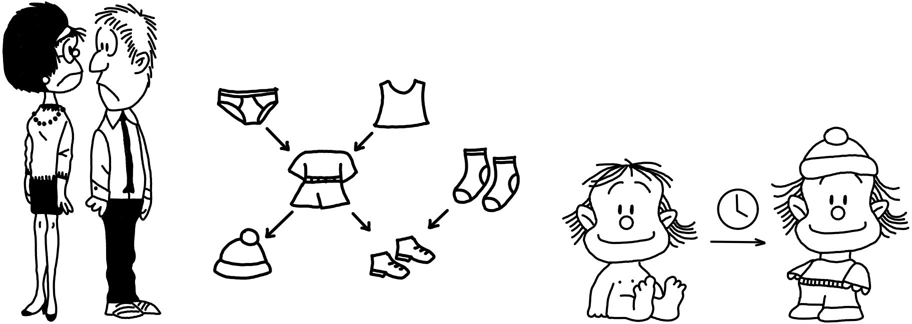

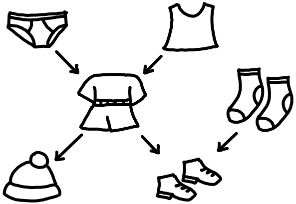

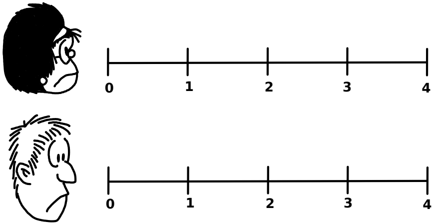
0
1
2
3
4
0
1
2
3
4

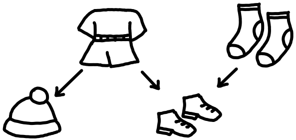

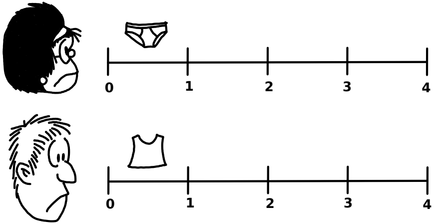
0
1
2
3
4
0
1
2
3
4

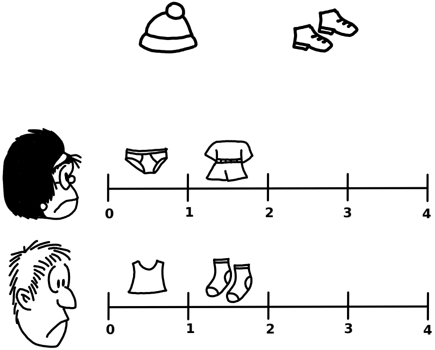
0
1
2
3
4
0
1
2
3
4

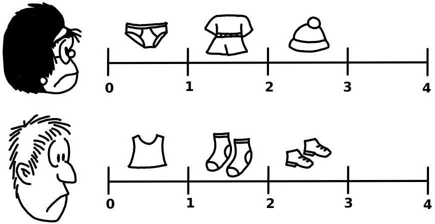
0
1
2
3
4
0
1
2
3
4

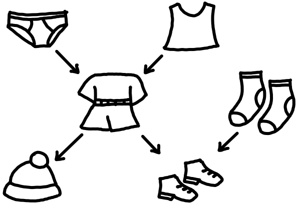

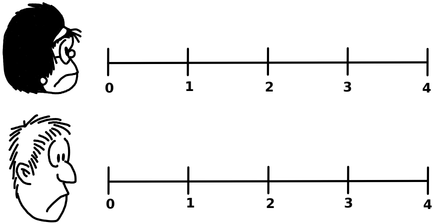
0
1
2
3
4
0
1
2
3
4

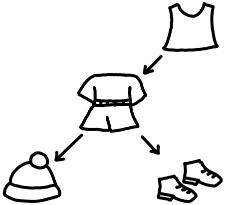

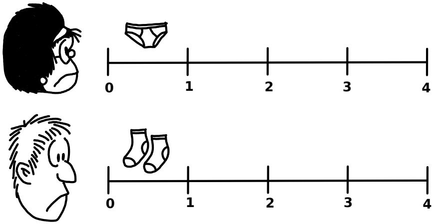
0 1 2 3 4
0 1 2 3 4

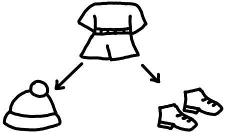

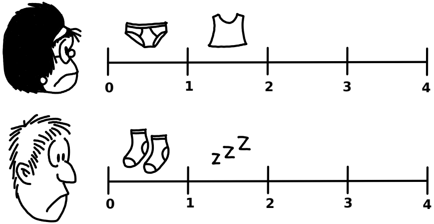
0
1
2
3
4
zzZ
0
1
2
3
4

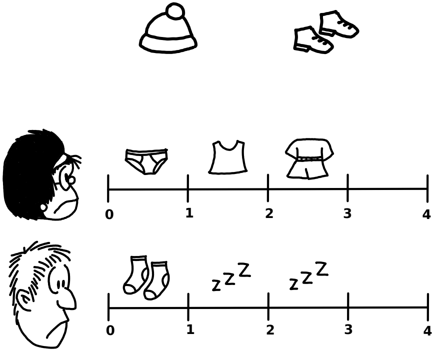
0
1
2
3
4
0
1
2
3
4

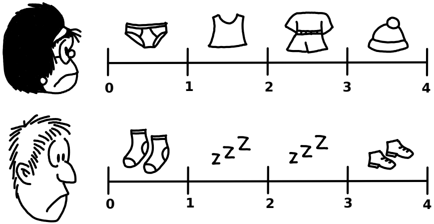

0
1
2
3
4
zzZ
zzZ
0
1
2
3
4

Literature Review

Problem	Reference
$P2 \mid prec \mid C_{\max}$	Solved by Coffman and Graham in 1972.
$P2 \mid intree, exp \mid \mathbb{E}[C_{\max}]$	Solved by Chandy and Reynolds in 1975.
$P2 \mid prec, exp \mid \mathbb{E}[C_{\max}]$	Subject of this thesis.

Overview

I. Introduction

II. Theoretical Foundations

III. Methods

IV. Results and Discussion

V. Conclusion

Probability Theory

Why consider exponentially distributed random variables $X_i \sim \text{Exp}(\lambda_i)$?

- Expected value:
$$\mathbb{E}[X_i] = \frac{1}{\lambda_i}.$$
- Minimum:
$$\min\{X_1, \ldots, X_n\} \sim \text{Exp}(\lambda_1 + \ldots + \lambda_n).$$
- Memorylessness:
$$\Pr[X_i > x + y | X_i > y] = \Pr[X_i > x] \qquad (x, y > 0).$$

Snapshots

Let m be the number of available machines. A snapshot $S = (T, P, A)$ consists of

- a finite set T,
- a binary relation $P \subseteq T \times T$ and
- a subset $A \subseteq T$,

such that (T, P) is a DAG, $|A| \leq m$ and every $t \in A$ is a source in (T, P).

$T(S)$, $P(S)$ and $A(S)$ can be defined as T, P and A, respectively.

Example

Let $m = 2$. The snapshot S with

- $T(S) = \{1, 2, 3, 4, 5\}$,
- $P(S) = \{(1, 3), (1, 4), (2, 4)\}$ and
- $A(S) = \{1\}$

can be represented as

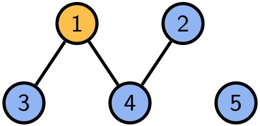

S

Strategies and Direct Snapshots

- A strategy σ is a function that maps a snapshot $S = (T, P, A)$ to a snapshot $\sigma(S) = (T, P, A')$ with $A \subseteq A'$.
- The resulting snapshot S' after removing a task $t \in A$ from a snapshot $S = (T, P, A)$ is called direct subsnapshot of S.

Example

Snapshots S, S', S'' with

- $\sigma(S) = S'$ for some strategy σ and
- S'' direct subsnapshot of S':

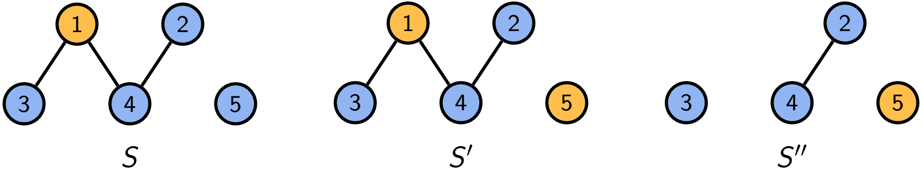

A formula for the expected makespan $\mathbb{E}[T_S^\sigma]$

Assuming independent finishing times $X_1, \ldots, X_n \sim \mathrm{Exp}(1)$ we can derive for a given strategy σ:

$$
\begin{aligned}
\mathbb{E}[T_S^\sigma] &= \mathbb{E}\left[T_{\sigma(S)}^\sigma\right] \\
&= \mathbb{E}\left[{T_{\sigma(S)}^\sigma}' + {T_{\sigma(S)}^\sigma}''\right] \\
&= \mathbb{E}\left[{T_{\sigma(S)}^\sigma}'\right] + \mathbb{E}\left[{T_{\sigma(S)}^\sigma}''\right] \\
&= \mathbb{E}[\min\{X_t \mid t \in A(\sigma(S))\}] + \sum_{S' \in D(\sigma(S))} \mathbb{E}[T_{S'}^\sigma] \cdot \frac{1}{|D(\sigma(S))|} \\
&= \frac{1}{|D(\sigma(S))|} \left(1 + \sum_{S' \in D(\sigma(S))} \mathbb{E}[T_{S'}^\sigma]\right),
\end{aligned}
$$

where $D(S)$ is the set of all direct subsnapshots of a snapshot S.

For $m = 2$ machines we obtain:

$$\mathbb{E}[T_S^\sigma] = \begin{cases} 0 & \text{if } D(\sigma(S)) = \emptyset \\ 1 + \mathbb{E}[T_{S'}^\sigma] & \text{if } D(\sigma(S)) = \{S'\} \\ \frac{1}{2} + \frac{1}{2}\mathbb{E}[T_{S'}^\sigma] + \frac{1}{2}\mathbb{E}[T_{S''}^\sigma] & \text{if } D(\sigma(S)) = \{S', S''\} \text{ for } S' \neq S''. \end{cases}$$

Overview

I. Introduction

II. Theoretical Foundations

III. Methods

IV. Results and Discussion

V. Conclusion

Schedule Visualization

PROSIT (**Pro**babilistic **S**cheduling **I**nteractive **T**ool) allows to

- draw a snapshot and
- visualize its schedule.

It can be downloaded from

`www.carlos-camino.de/prosit`

Demo: Schedule example 1

Computing Expectancies

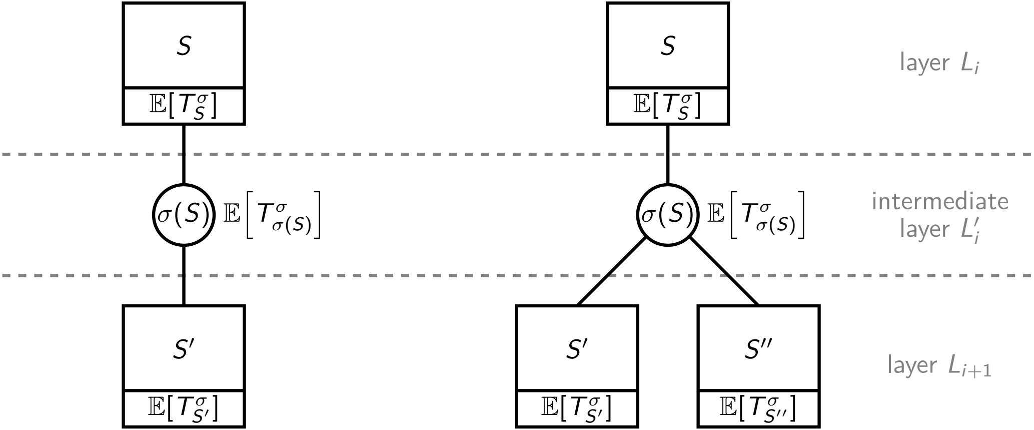

$$\mathbb{E}[T_S^\sigma] = \mathbb{E}\left[T_{\sigma(S)}^\sigma\right] = 1 + \mathbb{E}[T_{S'}^\sigma] \qquad \mathbb{E}[T_S^\sigma] = \mathbb{E}\left[T_{\sigma(S)}^\sigma\right] = \tfrac{1}{2} + \tfrac{1}{2}\mathbb{E}[T_{S'}^\sigma] + \tfrac{1}{2}\mathbb{E}[T_{S''}^\sigma]$$

The Coffman-Graham Algorithm

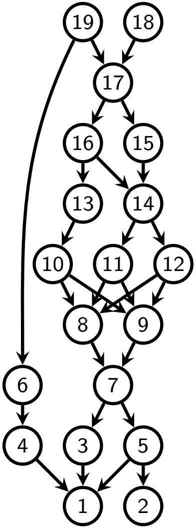

The Coffman-Graham Algorithm

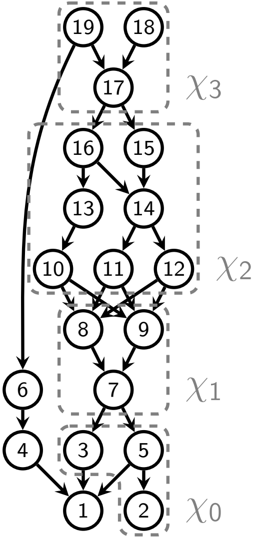

Demo: Schedule example 2

Overview

I. Introduction

II. Theoretical Foundations

III. Methods

IV. Results and Discussion

V. Conclusion

Case Studies on the HLF Policy

The snapshot

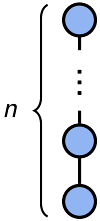

has optimal makespan n in the deterministic setting and optimal expected makespan n in the probabilistic setting.

The snapshot

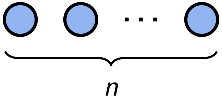

has optimal makespan $\lceil \frac{n}{2} \rceil$ in the deterministic setting and optimal expected makespan $\frac{n+1}{2}$ in the probabilistic setting.

Demos: Two parallel chains and a small counterexample for HLF

Low Level Tasks with High Priority

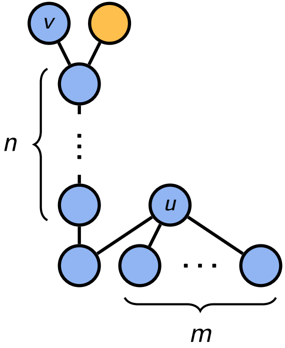

For which values of m, n would an optimal strategy ...

... assign u?

... assign v?

The following snapshot $L_{m,n}$

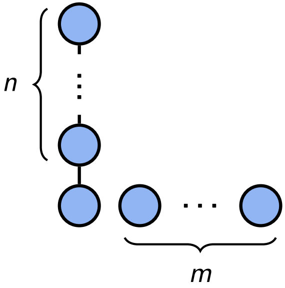

with optimal expected makespan $l_{m,n}$ satisfies $l_{m,0} = \frac{m}{2} + 1$, $l_{0,n} = n + 1$ and

$$l_{m,n} = \frac{1}{2} l_{m-1,n} + \frac{1}{2} l_{m,n-1} + \frac{1}{2}$$

for $m, n \geq 1$.

The following snapshot $H_{m,n}$

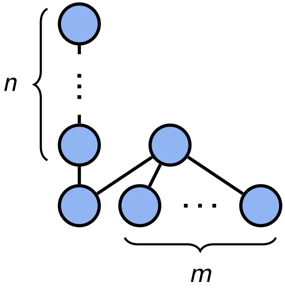

with optimal expected makespan $h_{m,n}$ satisfies $h_{m,0} = \frac{m}{2} + 2$ and

$$h_{m,n} = \frac{1}{2} h_{m,n-1} + \frac{1}{2} l_{m,n} + \frac{1}{2}$$

for $m, n \geq 1$.

The following snapshots $A^u_{m,n}$ and $A^v_{m,n}$

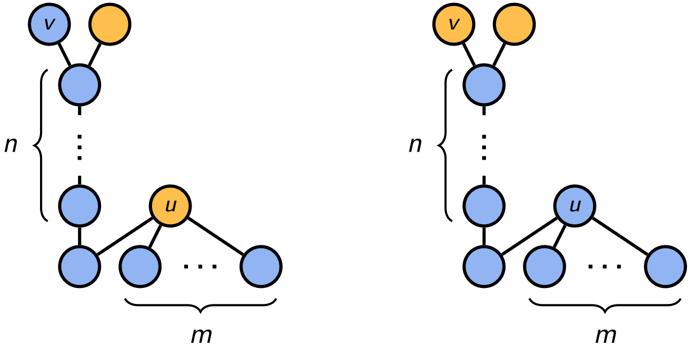

with optimal expected makespans $a^u_{m,n}$ and $a^v_{m,n}$ satisfy

$$a^u_{m,n} = \frac{1}{2} l_{m,n+1} + \frac{1}{2} h_{m,n+1} + \frac{3}{4} \qquad \text{and} \qquad a^v_{m,n} = h_{m,n+1} + \frac{1}{2}.$$

For which m, n it holds $a^u_{m,n} < a^v_{m,n}$? For which $a^u_{m,n} > a^v_{m,n}$?

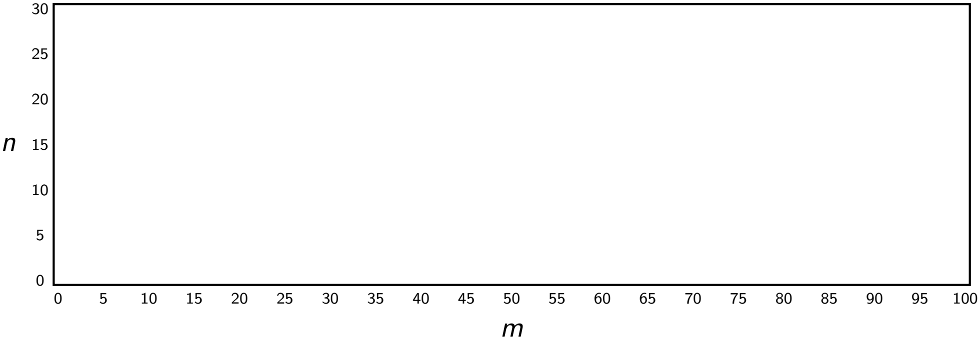

For which m, n it holds $a^u_{m,n} < a^v_{m,n}$? For which $a^u_{m,n} > a^v_{m,n}$?

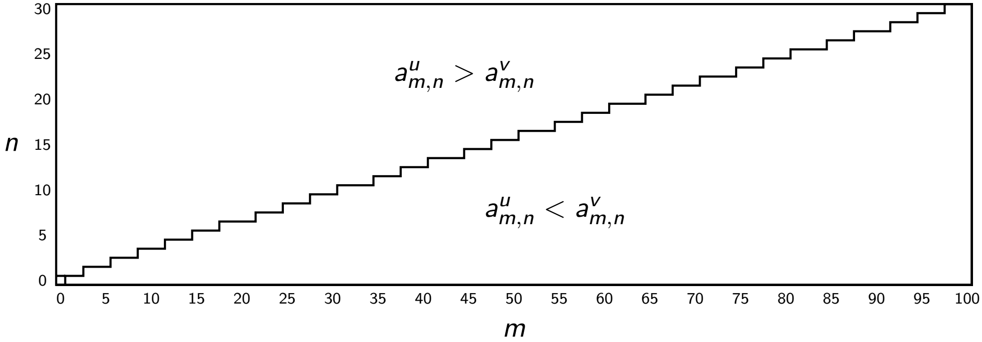

How big is the difference $a^u_{m,n} - a^v_{m,n}$?

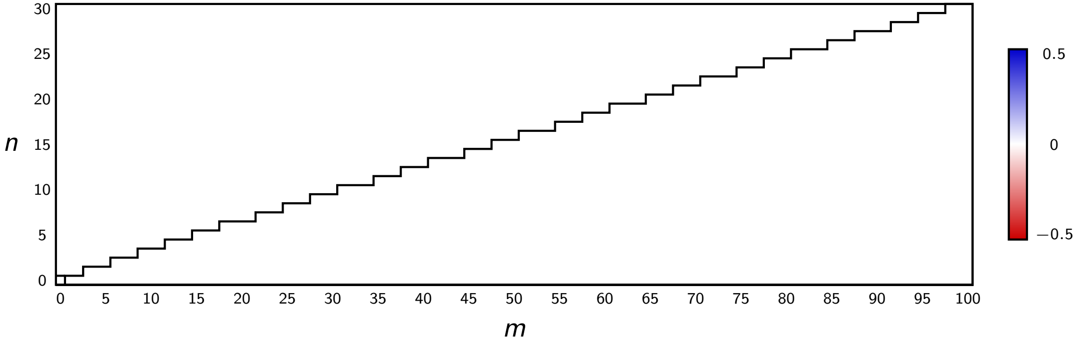

How big is the difference $a^u_{m,n} - a^v_{m,n}$?

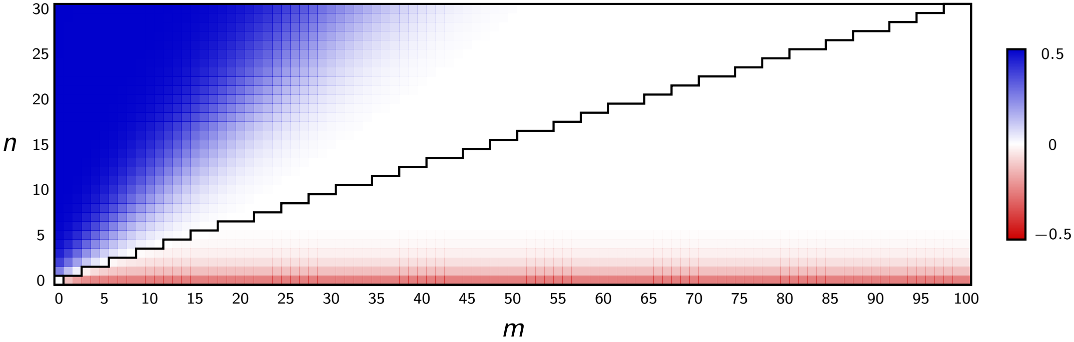

A formula for $I_{m,n}$

The recurrence relation

$$I_{m,n} = \frac{1}{2} I_{m-1,n} + \frac{1}{2} I_{m,n-1} + \frac{1}{2}$$

can be visualized as

$$\begin{tikzcd}
 & \vdots & \vdots & \vdots & \vdots & \vdots & \\
I_{0,2} \arrow[r, "\frac{1}{2}"] & I_{1,2} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{2,2} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{3,2} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{4,2} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{5,2} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & \cdots \\
I_{0,1} \arrow[r, "\frac{1}{2}"] & I_{1,1} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{2,1} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{3,1} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{4,1} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & I_{5,1} \arrow[r, "\frac{1}{2}"] \arrow[u, "\frac{1}{2}"'] & \cdots \\
I_{0,0} & I_{1,0} \arrow[u, "\frac{1}{2}"'] & I_{2,0} \arrow[u, "\frac{1}{2}"'] & I_{3,0} \arrow[u, "\frac{1}{2}"'] & I_{4,0} \arrow[u, "\frac{1}{2}"'] & I_{5,0} \arrow[u, "\frac{1}{2}"'] &
\end{tikzcd}$$

Using combinatorical arguments we can derive

$$
\begin{aligned}
l_{m,n} = & \sum_{i=0}^{m-1} \binom{n+i-1}{i} \left(\frac{1}{2}\right)^{n+i} l_{m-i,0} \\
& + \sum_{j=0}^{n-1} \binom{m+j-1}{j} \left(\frac{1}{2}\right)^{m+j} l_{0,n-j} \\
& + \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} \binom{i+j}{i} \left(\frac{1}{2}\right)^{i+j+1}
\end{aligned}
$$

with $l_{m-i,0} = \frac{m-i}{2} + 1$ and $l_{0,n-j} = n - j + 1$ for all $(m, n) \in \mathbb{N}_0 \times \mathbb{N}_0 \setminus \{(0,0)\}$.

Using generating functions we can derive the more complete formula

$$
\begin{aligned}
l_{m,n} = & \sum_{i=0}^{m} \binom{n+i-1}{i} \left(\frac{1}{2}\right)^{n+i} l_{m-i,0} \\
& + \sum_{j=0}^{n} \binom{m+j-1}{j} \left(\frac{1}{2}\right)^{m+j} l_{0,n-j} \\
& + \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} \binom{i+j}{i} \left(\frac{1}{2}\right)^{i+j+1} \\
& - \left(\frac{1}{2}\right)^{m+n} \binom{m+n}{m} l_{0,0}
\end{aligned}
$$

with $l_{m-i,0} = \frac{m-i}{2} + 1$ and $l_{0,n-j} = n - j + 1$ for all $(m, n) \in \mathbb{N}_0 \times \mathbb{N}_0$.

Overview

I. Introduction

II. Theoretical Foundations

III. Methods

IV. Results and Discussion

V. Conclusion

Adaptations of PROSIT

- other distributions
- more sophisticated strategies
- runtime optimizations

Further Analyses

- consider other task parameters
- find closed-form expressions for $l_{m,n}$, $h_{m,n}$, ...