Stochastic process creation

T. Brázdil J. Esparza S. Kiefer M. Luttenberger

Masaryk University, Brno

Technische Universität München

MFCS 2009

Back in victorian Britain . . .

There was concern amongst the Victorians that aristocratic families were becoming extinct.

Back in victorian Britain . . .

There was concern amongst the Victorians that aristocratic families were becoming extinct.

Francis Galton (1822-1911), anthropologist and polymath:
Are families of English peers more likely to die out than the families of ordinary men?

> *Let $p_0, p_1, \ldots, p_n$ be the respective probabilities that a man has 0, 1, 2, . . . n sons, let each son have the same probability for sons of his own, and so on. What is the probability that the male line goes extinct?*

Back in victorian Britain . . .

There was concern amongst the Victorians that aristocratic families were becoming extinct.

Francis Galton (1822-1911), anthropologist and polymath:
Are families of English peers more likely to die out than the families of ordinary men?

> *Let $p_0, p_1, \ldots, p_n$ be the respective probabilities that a man has 0, 1, 2, . . . n sons, let each son have the same probability for sons of his own, and so on. What is the probability that the male line goes extinct?*

Henry William Watson (1827-1903), vicar and mathematician:
The probability that the line goes extinct is the least solution of

$$x = p_0 + p_1 x + p_2 x^2 + \ldots + p_n x^n$$

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Our work in the last months (and ongoing)

Investigate SBPs as models for the stochastic analysis of CS systems with process creation:

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Our work in the last months (and ongoing)

Investigate SBPs as models for the stochastic analysis of CS systems with process creation: multi-threaded programs,

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Our work in the last months (and ongoing)

Investigate SBPs as models for the stochastic analysis of CS systems with process creation: multi-threaded programs, OS tasks,

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Our work in the last months (and ongoing)

Investigate SBPs as models for the stochastic analysis of CS systems with process creation: multi-threaded programs, OS tasks, computer viruses,

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Our work in the last months (and ongoing)

Investigate SBPs as models for the stochastic analysis of CS systems with process creation: multi-threaded programs, OS tasks, computer viruses, dynamic data structures,

Stochastic branching theory

Stochastic branching processes (SBPs)

Stochastic processes for the behaviour of populations whose individuals die and reproduce.

Used as models of reproduction of biological species, evolution of gene pools, chemical and nuclear reactions.

Very well studied by mathematicians (several standard textbooks).

Our work in the last months (and ongoing)

Investigate SBPs as models for the stochastic analysis of CS systems with process creation: multi-threaded programs, OS tasks, computer viruses, dynamic data structures, divide-and-conquer algorithms . . .

A classification of SBPs

Two classical dimensions

Single-type/Multi-type
(one/several "subspecies" with different offspring probabilities).

Untimed/Timed

A classification of SBPs

Two classical dimensions

Single-type/Multi-type
(one/several “subspecies” with different offspring probabilities).

Untimed/Timed

A new dimension for CS systems

Distinction between processes and processors

A classification of SBPs

Two classical dimensions

Single-type/Multi-type
(one/several “subspecies” with different offspring probabilities).

Untimed/Timed

A new dimension for CS systems

Distinction between processes and processors

Unboundedly many processors
(new processes immediately allocated to fresh processors)
All research on SBPs has only considered this model.

A classification of SBPs

Two classical dimensions

Single-type/Multi-type
(one/several “subspecies” with different offspring probabilities).

Untimed/Timed

A new dimension for CS systems

Distinction between processes and processors

Unboundedly many processors
(new processes immediately allocated to fresh processors)
All research on SBPs has only considered this model.

Single processor

A classification of SBPs

Two classical dimensions

Single-type/Multi-type
(one/several "subspecies" with different offspring probabilities).

Untimed/Timed

A new dimension for CS systems

Distinction between processes and processors

Unboundedly many processors
(new processes immediately allocated to fresh processors)
All research on SBPs has only considered this model.

Single processor

K-processors, variable number of processors ...

In this talk . . .

. . . single-typed, untimed systems, with either unboundedly many or a single processor.

In this talk . . .

. . . single-typed, untimed systems, with either unboundedly many or a single processor.

- restriction to single-type only for expository reasons

In this talk . . .

. . . single-typed, untimed systems, with either unboundedly many or a single processor.

- restriction to single-type only for expository reasons
- deterministic lifetimes can be modelled within the untimed model

In this talk . . .

. . . single-typed, untimed systems, with either unboundedly many or a single processor.

- restriction to single-type only for expository reasons
- deterministic lifetimes can be modelled within the untimed model
- extensions to stochastic lifetimes and k-processors are future work.

In this talk . . .

. . . single-typed, untimed systems, with either unboundedly many or a single processor.

- restriction to single-type only for expository reasons
- deterministic lifetimes can be modelled within the untimed model
- extensions to stochastic lifetimes and k-processors are future work.

Mix of survey and new results

Describing systems

- A process “dies” when it generates its children.
- Thread creation is simulated by assuming that one child is the continuation of the parent.

Describing systems

- A process "dies" when it generates its children.
- Thread creation is simulated by assuming that one child is the continuation of the parent.
- Our running example:

$$X \xrightarrow{0.1} \langle X, X, X \rangle \quad X \xrightarrow{0.2} \langle X, X \rangle \quad X \xrightarrow{0.1} X \quad X \xrightarrow{0.6} \epsilon$$

Describing systems

- A process "dies" when it generates its children.
- Thread creation is simulated by assuming that one child is the continuation of the parent.
- Our running example:

$$X \xrightarrow{0.1} \langle X, X, X \rangle \quad X \xrightarrow{0.2} \langle X, X \rangle \quad X \xrightarrow{0.1} X \quad X \xrightarrow{0.6} \epsilon$$

Probability generating function $f(x)$

$$f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$$

Describing executions: family trees

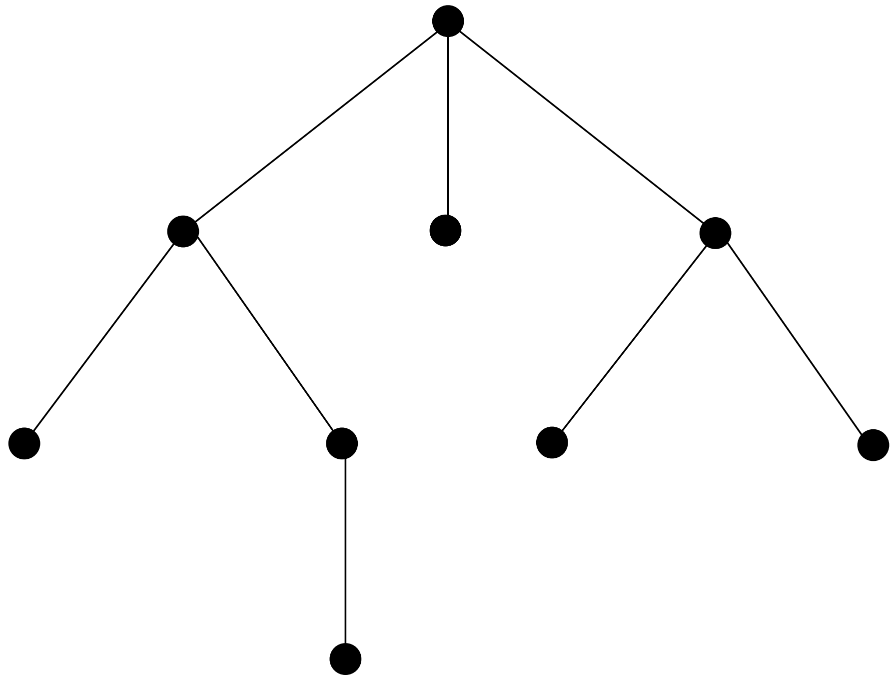

Describing executions: family trees

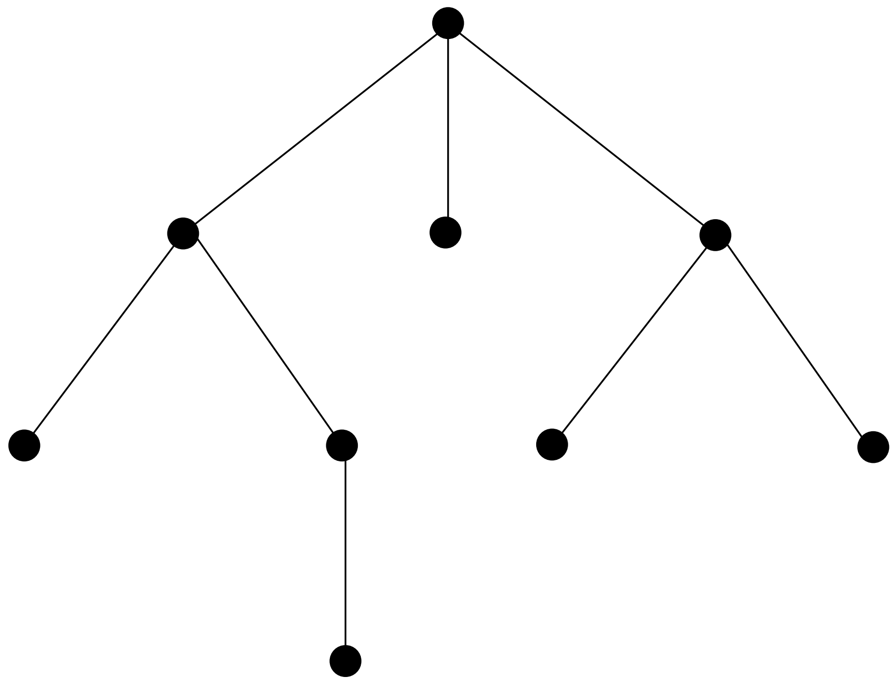

Executing a family tree

∞-processors: generation-wise

1-processor: scheduler (system det. by pgf and scheduler)

Probability of termination (extinction)

Observe

The probability of extinction is independent of the number of processors. (More processors accelerate a computation, but don't change it.)

Probability of termination (extinction)

Observe

The probability of extinction is independent of the number of processors. (More processors accelerate a computation, but don't change it.)

Theorem (well known)

The probability of extinction of the process types is the least nonnegative fixed point of the pgf, i.e., the smallest nonnegative solution of $x = f(x)$.

Probability of termination (extinction)

Observe

The probability of extinction is independent of the number of processors. (More processors accelerate a computation, but don't change it.)

Theorem (well known)

The probability of extinction of the process types is the least nonnegative fixed point of the pgf, i.e., the smallest nonnegative solution of $x = f(x)$.

The least solution for $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$ is 1.

Probability of termination (extinction)

Observe

The probability of extinction is independent of the number of processors. (More processors accelerate a computation, but don't change it.)

Theorem (well known)

The probability of extinction of the process types is the least nonnegative fixed point of the pgf, i.e., the smallest nonnegative solution of $x = f(x)$.

The least solution for $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$ is 1.

The least solution for $f(x) = 2/3x^2 + 1/3$ is 1/2.

Critical and subcritical systems

We consider systems that terminate with probability 1.

Further classified into:

- Critical: expected number of children is 1.
- Subcritical: expected number of children smaller than 1.

Probability space

- Elementary events: family trees.
- Probability of a family tree: product of the probabilities of its nodes.

Probability space

- Elementary events: family trees.
- Probability of a family tree: product of the probabilities of its nodes.

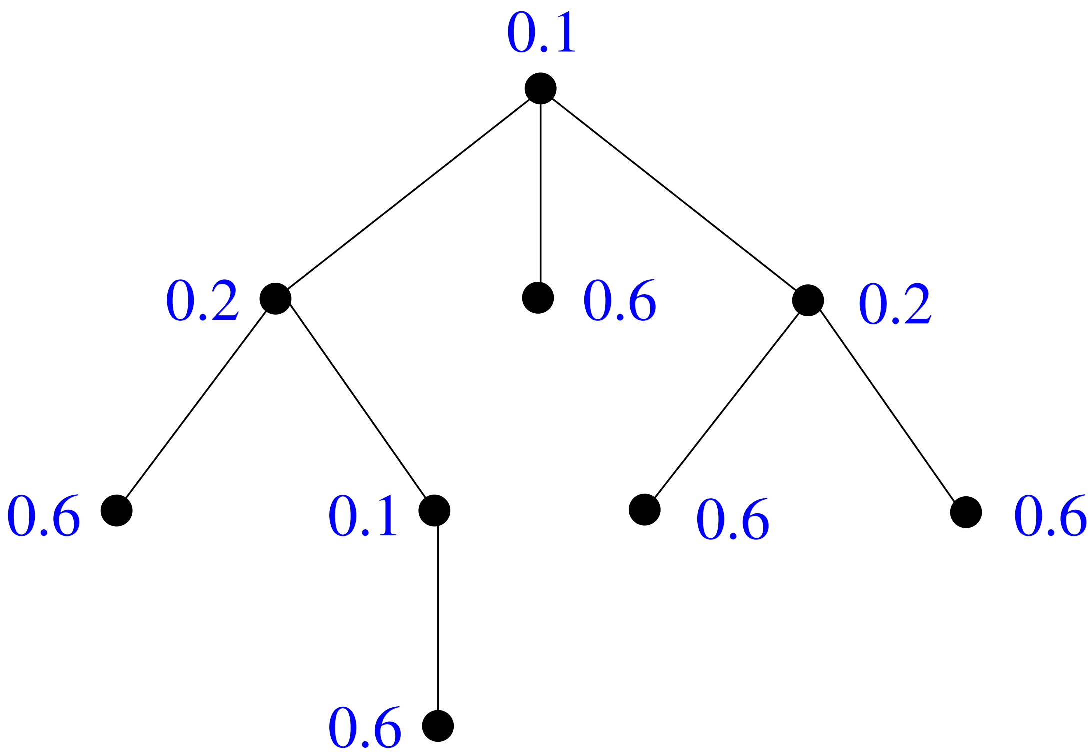

The ∞-processor case: random variables

Completion time (time to extinction)

Random variable T that assigns to a family tree its number of generations.

Processor number

Random variable N that assigns to a family tree the maximal size of a generation.

An example

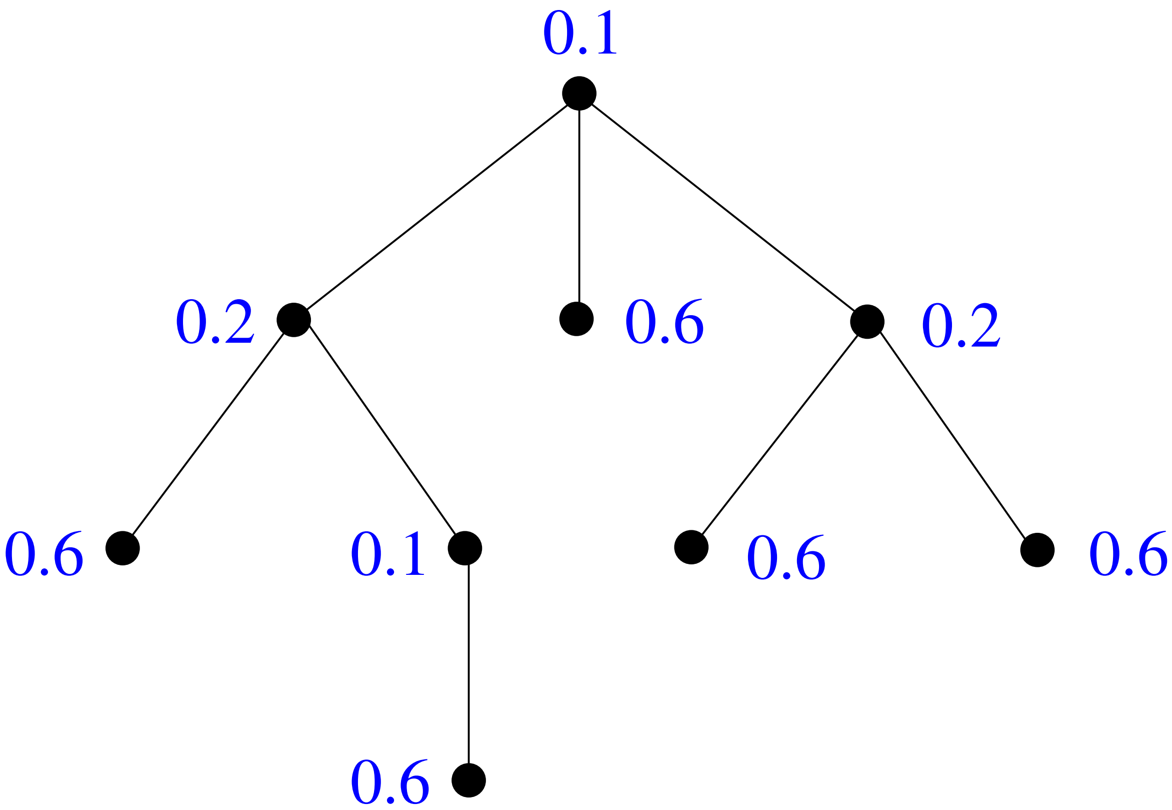

Completion time $= 4$ (four generations)

Processor number $= 4$ (size of the 3rd generation)

Analyzing the completion time

Proposition

The probabilities $\Pr(T \leq 1)$, $\Pr(T \leq 2)$, $\Pr(T \leq 3)$, $\ldots$ *of termination in at most* $1, 2, 3, \ldots$ *generations are equal to*

$$f(0)\ ,\ f(f(0)) = f^2(0)\ ,\ f(f(f(0))) = f^3(0)\ ,\ \ldots$$

Analyzing the completion time

Proposition

The probabilities $\Pr(T \leq 1)$, $\Pr(T \leq 2)$, $\Pr(T \leq 3)$, $\ldots$ *of termination in at most* $1, 2, 3, \ldots$ *generations are equal to*

$$f(0)\ ,\ f(f(0)) = f^2(0)\ ,\ f(f(f(0))) = f^3(0)\ ,\ \ldots$$

Proof by example.

Let $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

□

Analyzing the completion time

Proposition

The probabilities $\Pr(T \le 1)$, $\Pr(T \le 2)$, $\Pr(T \le 3)$, $\ldots$ *of termination in at most* $1, 2, 3, \ldots$ *generations are equal to*

$$f(0) \;,\; f(f(0)) = f^2(0) \;,\; f(f(f(0))) = f^3(0) \;,\; \ldots$$

Proof by example.

Let $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.
Let p_{k+1} be the probability of termination in at most $k+1$-generations. We have

□

Analyzing the completion time

Proposition

The probabilities $\Pr(T \leq 1)$, $\Pr(T \leq 2)$, $\Pr(T \leq 3)$, $\ldots$ *of termination in at most* $1, 2, 3, \ldots$ *generations are equal to*

$$f(0) \; , \; f(f(0)) = f^2(0) \; , \; f(f(f(0))) = f^3(0) \; , \; \ldots$$

Proof by example.

Let $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.
Let p_{k+1} be the probability of termination in at most $k+1$-generations. We have

$$p_{k+1} = 0.1 \cdot p_k^3 + 0.2 \cdot p_k^2 + 0.1 \cdot p_k + 0.6$$

□

Analyzing the completion time

Proposition

The probabilities $\Pr(T \leq 1)$, $\Pr(T \leq 2)$, $\Pr(T \leq 3)$, $\ldots$ *of termination in at most* $1, 2, 3, \ldots$ *generations are equal to*

$$f(0)\ ,\ f(f(0)) = f^2(0)\ ,\ f(f(f(0))) = f^3(0)\ ,\ \ldots$$

Proof by example.

Let $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.
Let p_{k+1} be the probability of termination in at most $k+1$-generations. We have

$$p_{k+1} = 0.1 \cdot p_k^3 + 0.2 \cdot p_k^2 + 0.1 \cdot p_k + 0.6$$

$= f(p_k)$ □

Analyzing the completion time

By Kleene's theorem, the least fixed point of $f(x)$ is the limit of $f(0), f^2(0), f^3(0) \ldots$

Analyzing the completion time

By Kleene's theorem, the least fixed point of $f(x)$ is the limit of $f(0), f^2(0), f^3(0) \ldots$

Least fixed point of $f(x)$ = probability of termination.

Analyzing the completion time

By Kleene's theorem, the least fixed point of $f(x)$ is the limit of $f(0)$, $f^2(0)$, $f^3(0)$ ...

Least fixed point of $f(x)$ = probability of termination.

k-th Kleene approximant to the least fixed point

=

probability of termination after at most k generations.

Analyzing the process number

Much harder problem, studied in the 70s by Lindvall and Nerman for one-type systems.

Analyzing the process number

Much harder problem, studied in the 70s by Lindvall and Nerman for one-type systems.

Fact

The pgf of a subcritical system has exactly two fixed points.

Analyzing the process number

Much harder problem, studied in the 70s by Lindvall and Nerman for one-type systems.

Fact

The pgf of a subcritical system has exactly two fixed points.

Theorem (Lindvall 76,Nerman 77)

Let $a > 1$ be the greatest fixed point of the pgf. For all $n \geq 1$

$$\Pr[N > n] < \frac{a-1}{a^n - 1} \quad \text{and} \quad \Pr[N > n] \in \Theta\left(\frac{1}{na^n}\right).$$

Analyzing the process number

Much harder problem, studied in the 70s by Lindvall and Nerman for one-type systems.

Fact

The pgf of a subcritical system has exactly two fixed points.

Theorem (Lindvall 76,Nerman 77)

Let $a > 1$ be the greatest fixed point of the pgf. For all $n \geq 1$

$$\Pr[N > n] < \frac{a-1}{a^n - 1} \quad \textit{and} \quad \Pr[N > n] \in \Theta\left(\frac{1}{na^n}\right) .$$

For $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$ we have $a \approx 1.3722$.
For instance, $\Pr[N > n] \leq 0.01$ for $n \geq 12$.

The single processor case: random variables

Recall: a scheduler repeatedly chooses a process from from the pool of current processes awaiting execution.

The single processor case: random variables

Recall: a scheduler repeatedly chooses a process from from the pool of current processes awaiting execution.

Time to termination (time to extinction)

Random variable T that assigns to a family tree its size. Independent of the scheduler.

The single processor case: random variables

Recall: a scheduler repeatedly chooses a process from from the pool of current processes awaiting execution.

Time to termination (time to extinction)

Random variable T that assigns to a family tree its size. Independent of the scheduler.

Completion space

Random variable S^σ that assigns to a family tree the maximal size reached by the pool during the execution of the tree by the scheduler σ.

An example

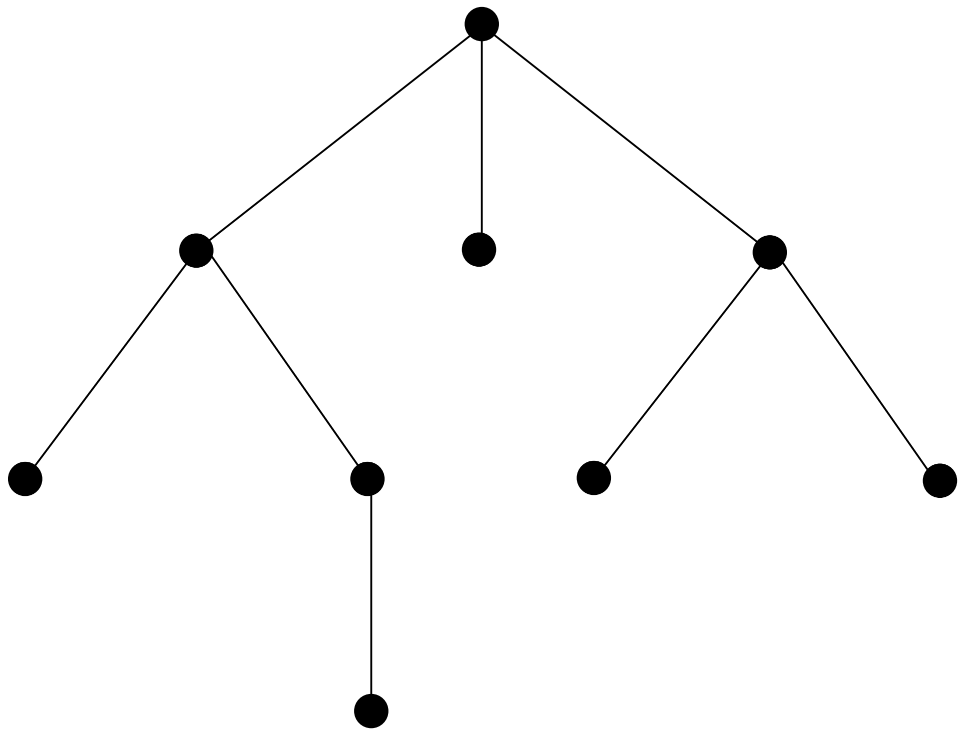

Completion time $= 9$, completion space between 3 and 5

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

□

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

$$E[T] = 0.6 \cdot 1$$

□

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

$$\begin{aligned} E[T] &= \quad 0.6 \cdot 1 \\ &\quad + \; 0.1 \cdot (\, 1 + E[T] \,) \end{aligned}$$

□

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

$$\begin{aligned} E[T] &= \quad 0.6 \cdot 1 \\ &\quad + \; 0.1 \cdot (\, 1 + E[T] \,) \\ &\quad + \; 0.2 \cdot (\, 1 + 2E[T] \,) \end{aligned}$$

□

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

$$\begin{aligned} E[T] &= && 0.6 \cdot 1 \\ & &+& \; 0.1 \cdot (\, 1 + E[T] \,) \\ & &+& \; 0.2 \cdot (\, 1 + 2E[T] \,) \\ & &+& \; 0.1 \cdot (\, 1 + 3E[T] \,) \end{aligned}$$

□

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

$$\begin{aligned} E[T] &= \quad 0.6 \cdot 1 \\ &\quad + \; 0.1 \cdot (\, 1 + E[T] \,) \\ &\quad + \; 0.2 \cdot (\, 1 + 2E[T] \,) \\ &\quad + \; 0.1 \cdot (\, 1 + 3E[T] \,) \\ &= \quad 1 + 0.8 \cdot E[T] \end{aligned}$$

□

Analyzing the completion time

Proposition

The expected value of T is the solution of a linear equation.

Proof by example.

Consider $f(x) = 0.1x^3 + 0.2x^2 + 0.1x + 0.6$.

$$
\begin{aligned}
E[T] &= \quad 0.6 \cdot 1 \\
&\quad + \; 0.1 \cdot (\, 1 + E[T] \,) \\
&\quad + \; 0.2 \cdot (\, 1 + 2E[T] \,) \\
&\quad + \; 0.1 \cdot (\, 1 + 3E[T] \,) \\
&= \quad 1 + 0.8 \cdot E[T]
\end{aligned}
$$

and so $E[T] = 5$. □

A theorem by Dwass

Theorem (Dwass69)

If $p_0 > 0$ *then*

$$\Pr[T = j] = \frac{1}{j}\, p_{j,j-1}$$

for every $j \geq 0$*, where* $p_{j,j-1}$ *denotes the probability that a generation has* $j - 1$ *processes under the condition that the parent generation has* j *processes.*

Analyzing the completion space

Scheduler

Function that assigns to a family tree one of its executions.

Analyzing the completion space

Scheduler

Function that assigns to a family tree one of its executions.

Offline schedulers

Know the complete family tree in advance.

Analyzing the completion space

Scheduler

Function that assigns to a family tree one of its executions.

Offline schedulers

Know the complete family tree in advance.

Online schedulers

Only know the part of the family tree executed so far.

An example

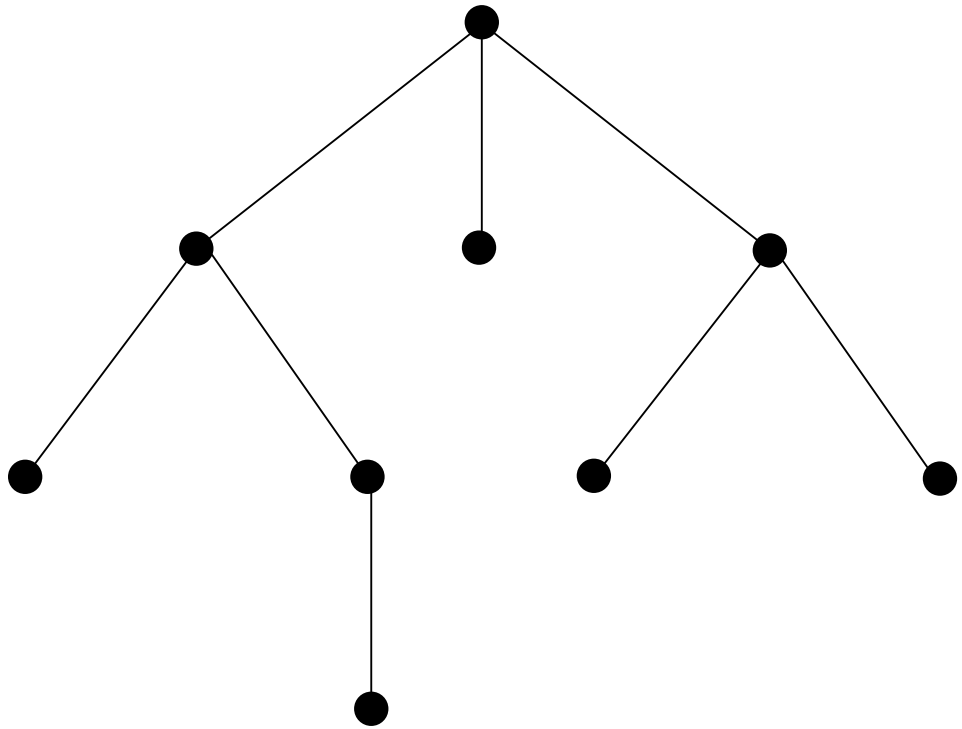

An example

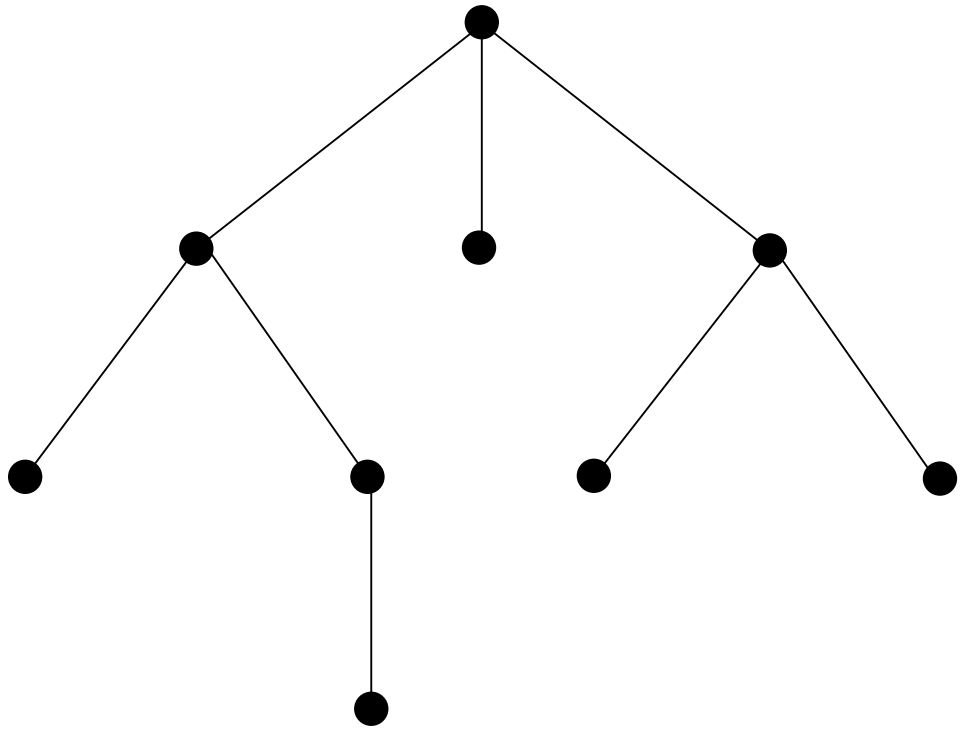

Goal: obtain bounds valid for all online schedulers, and compare them with the optimal offline scheduler

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

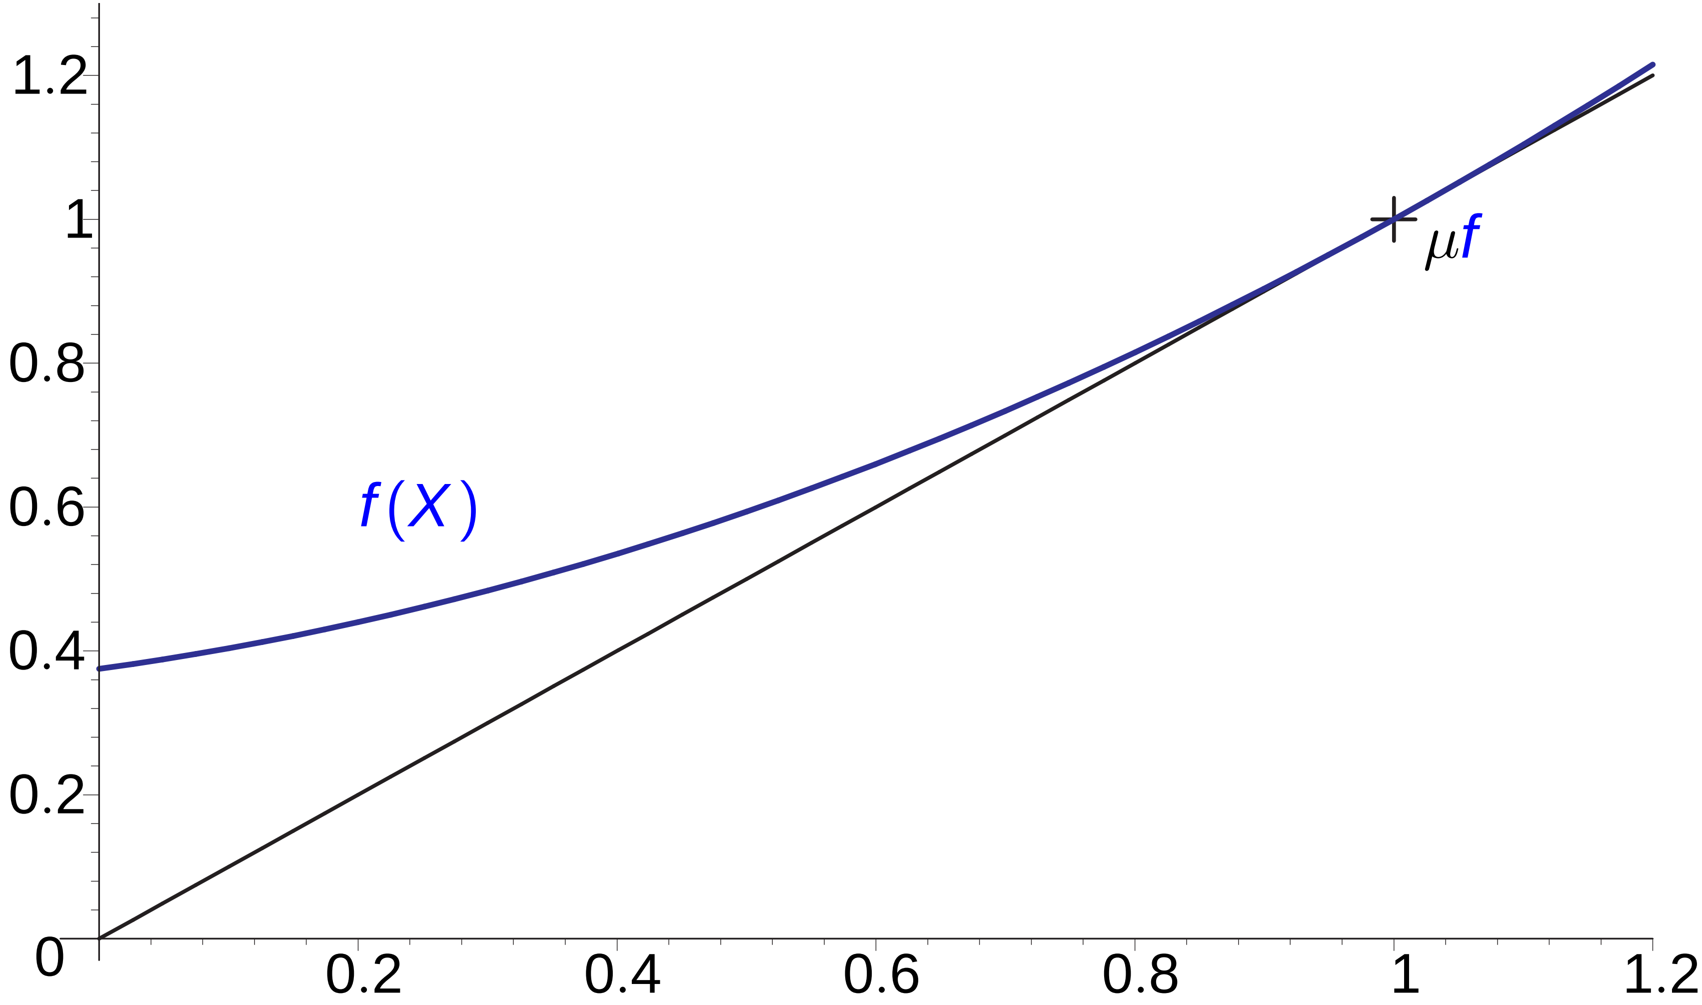

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

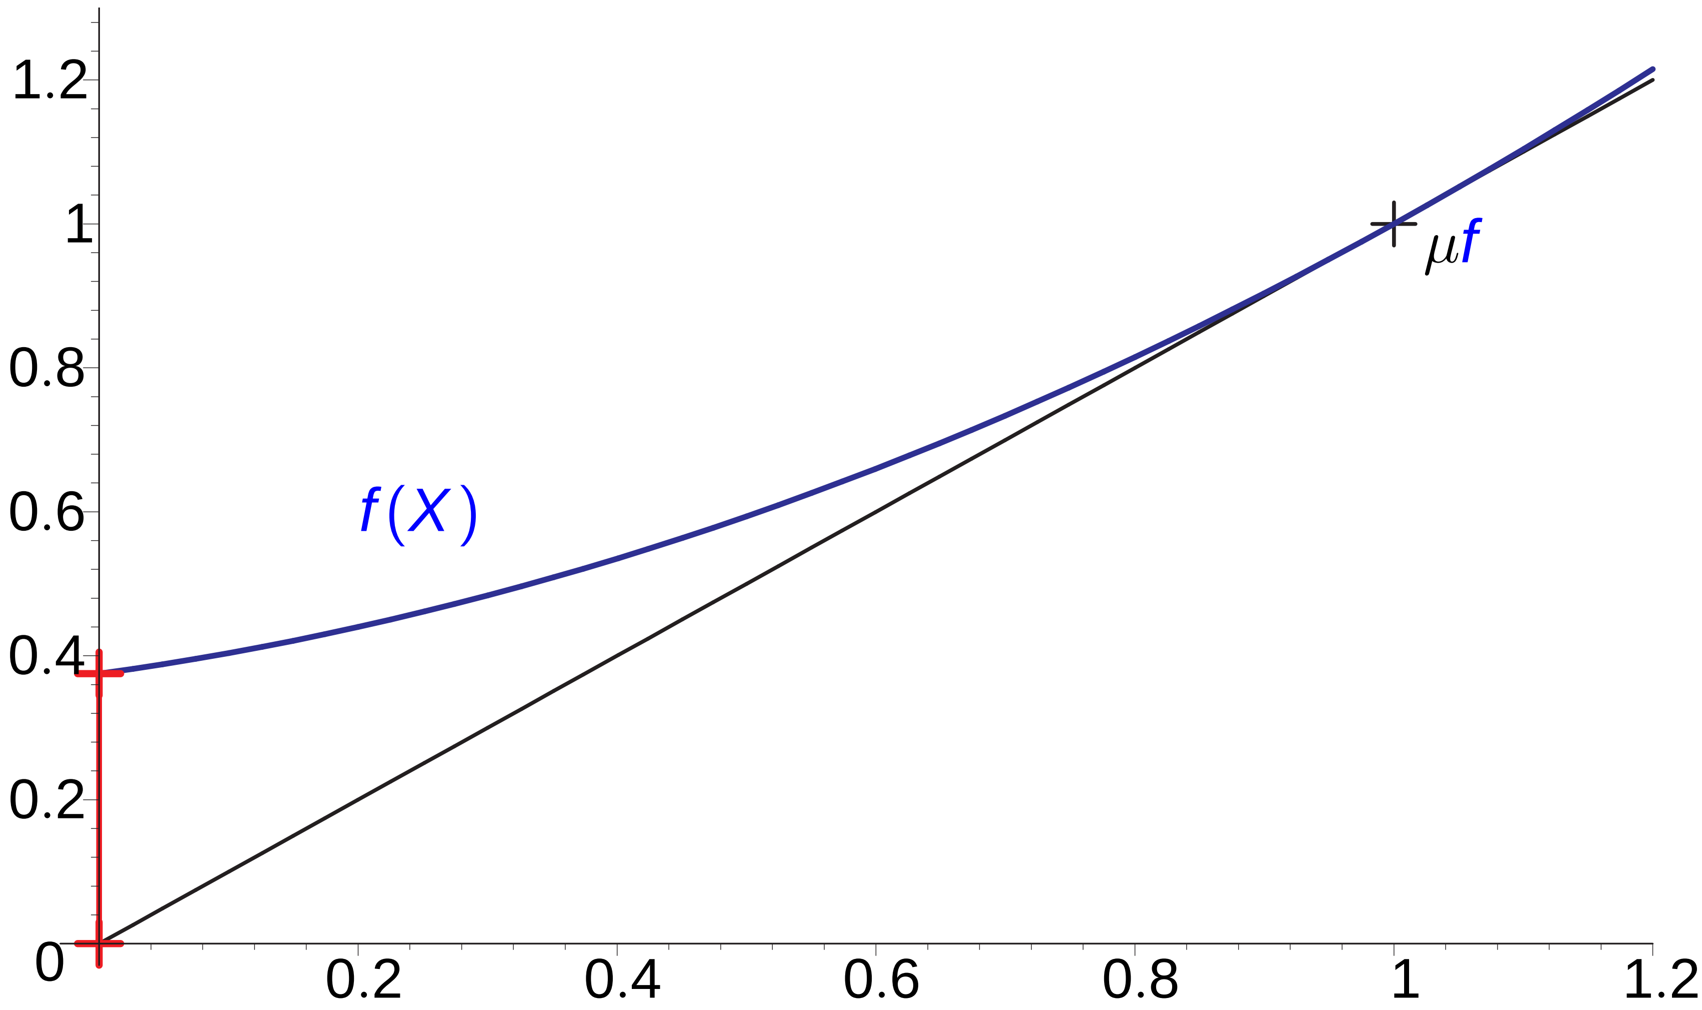

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

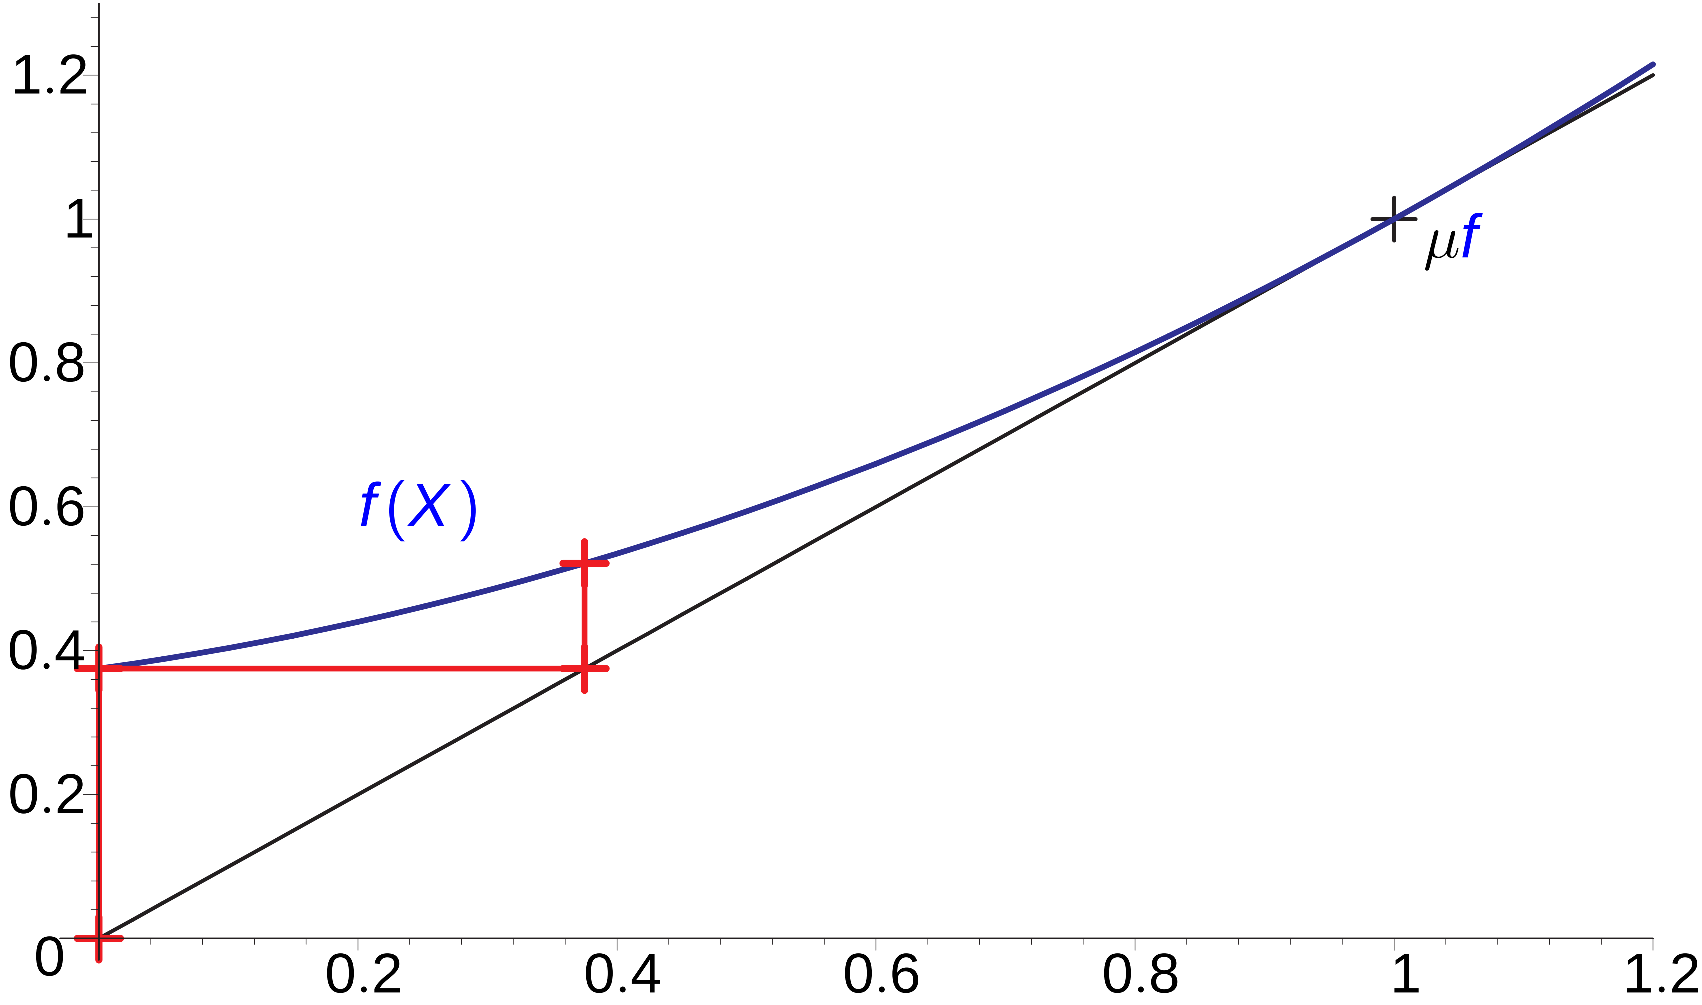

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

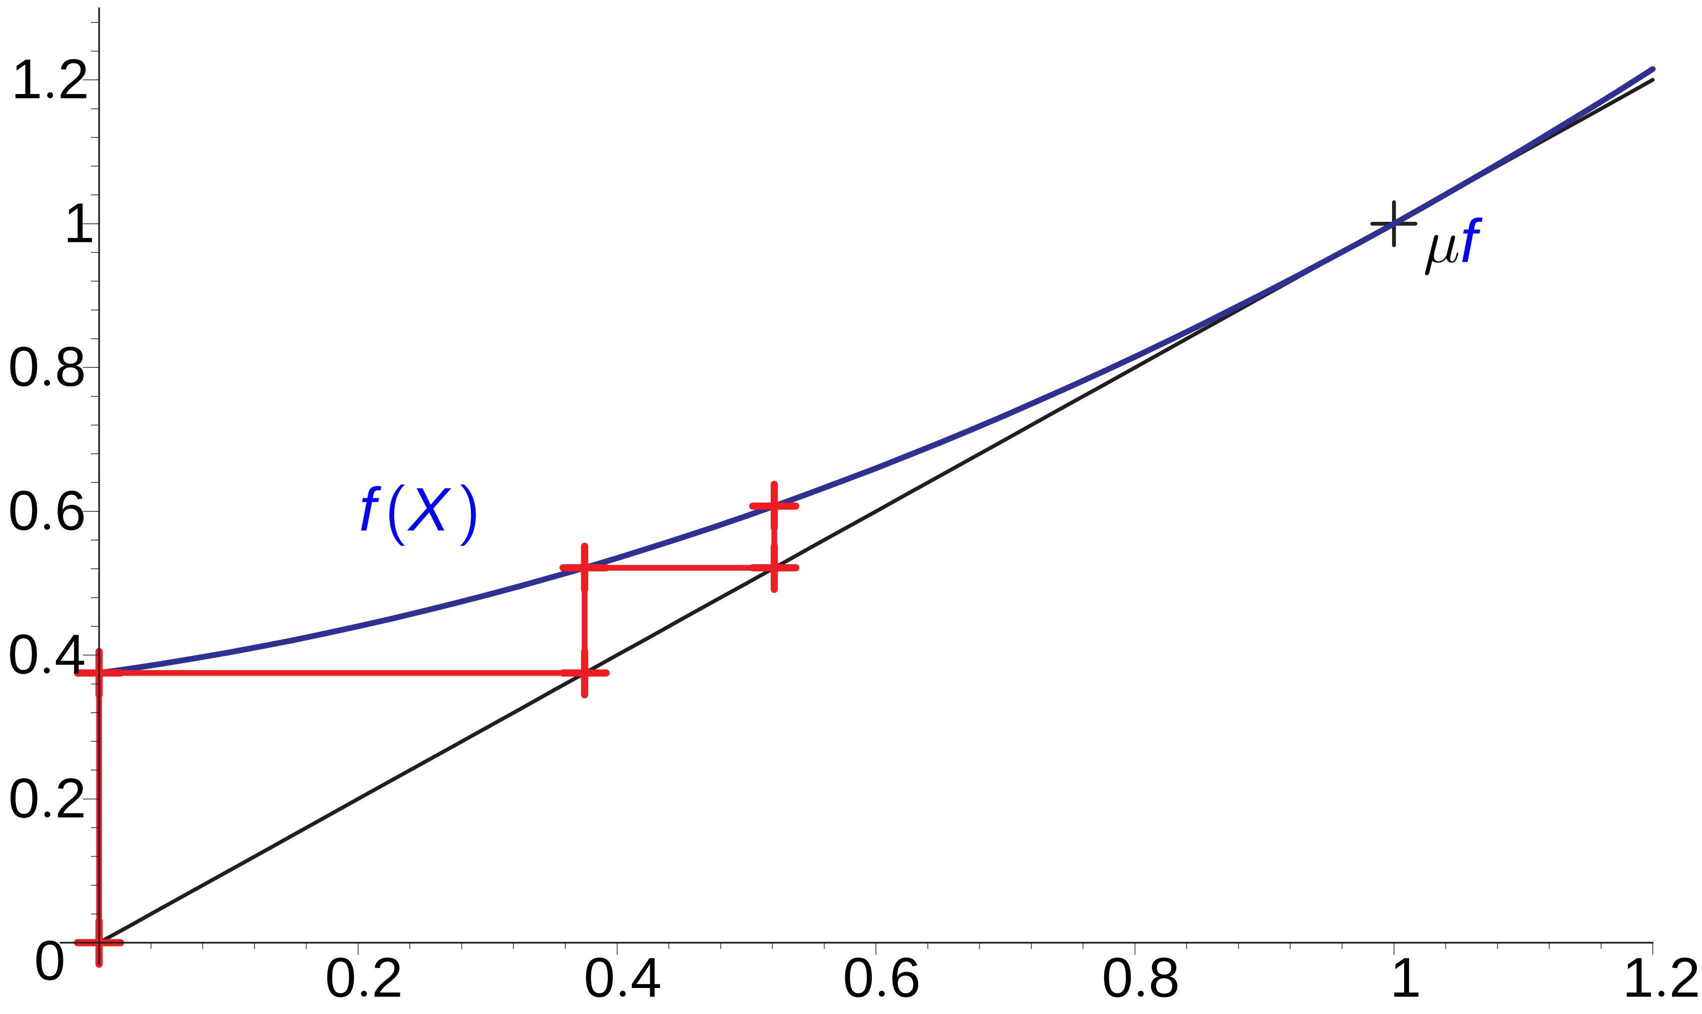

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

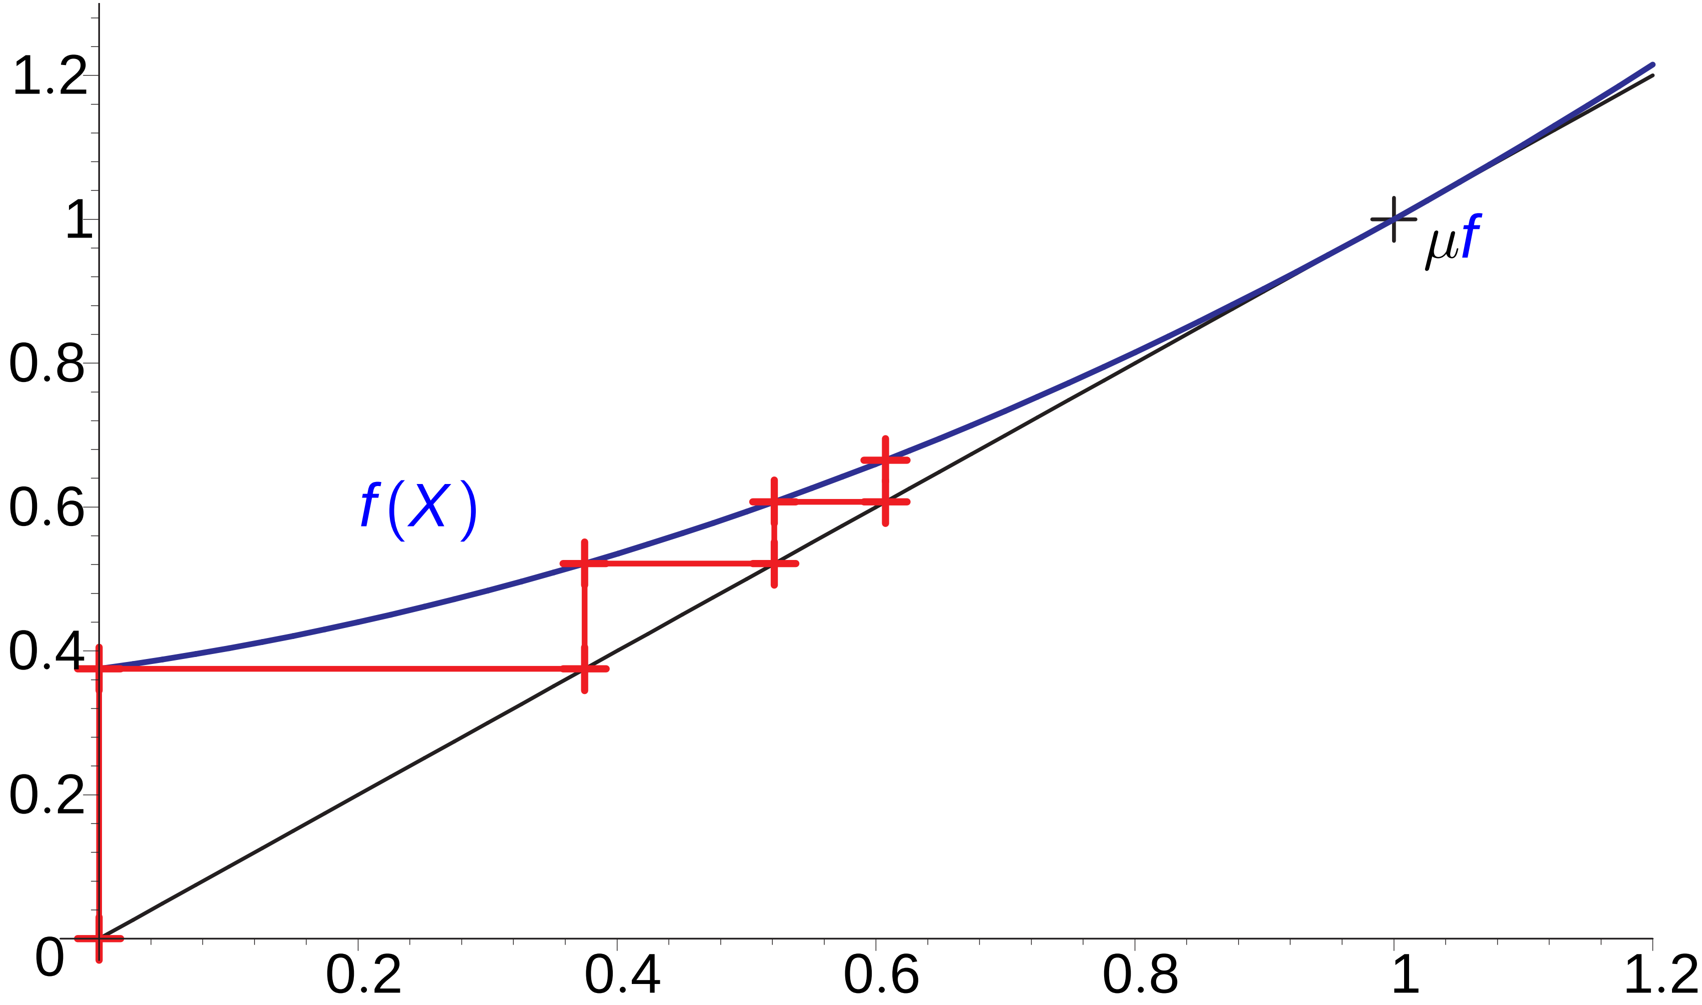

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

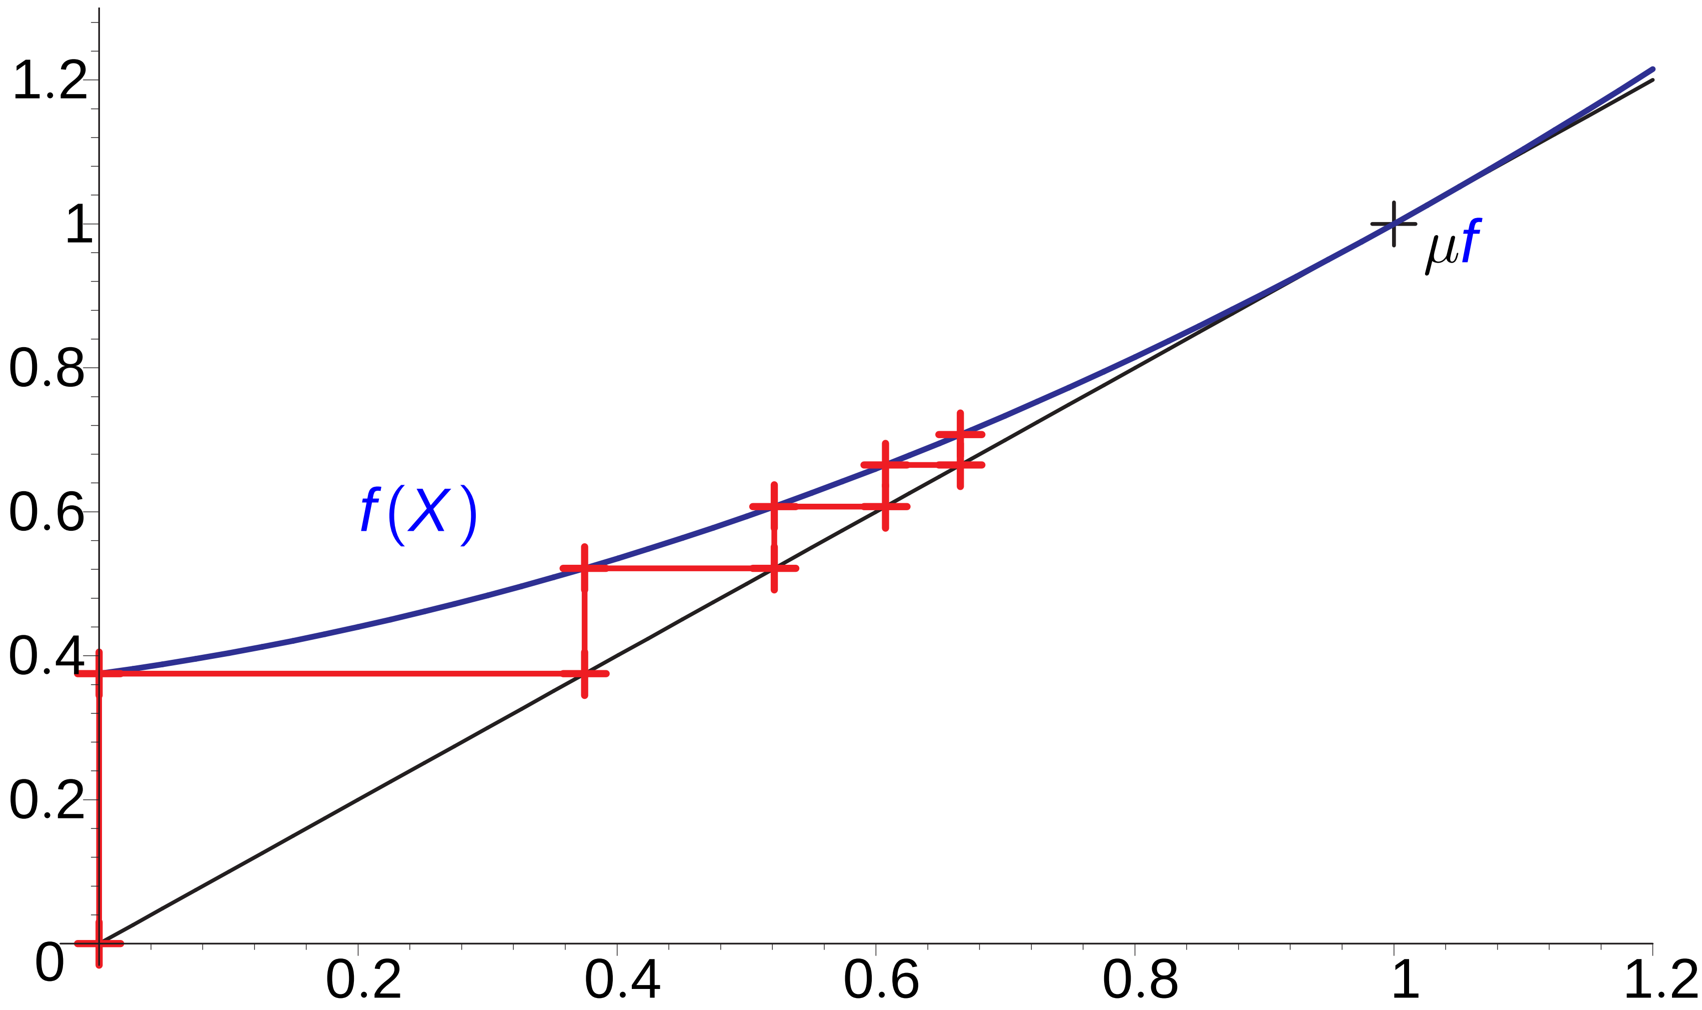

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

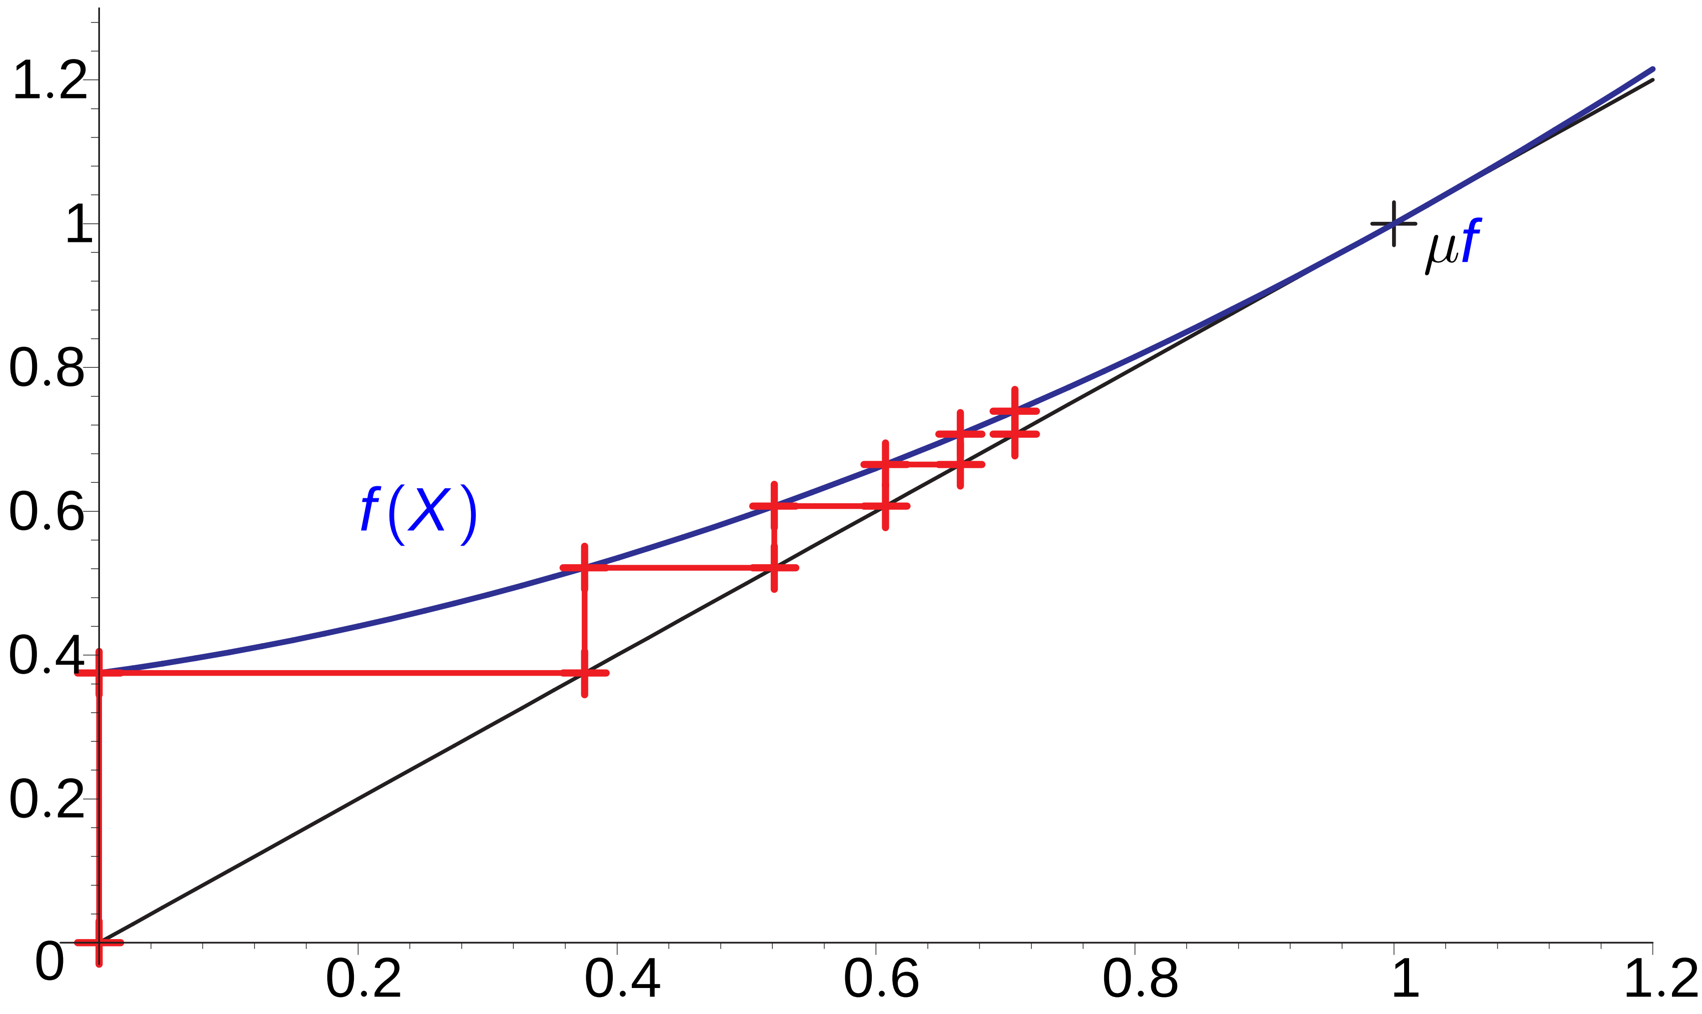

Kleene Iteration

Consider $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

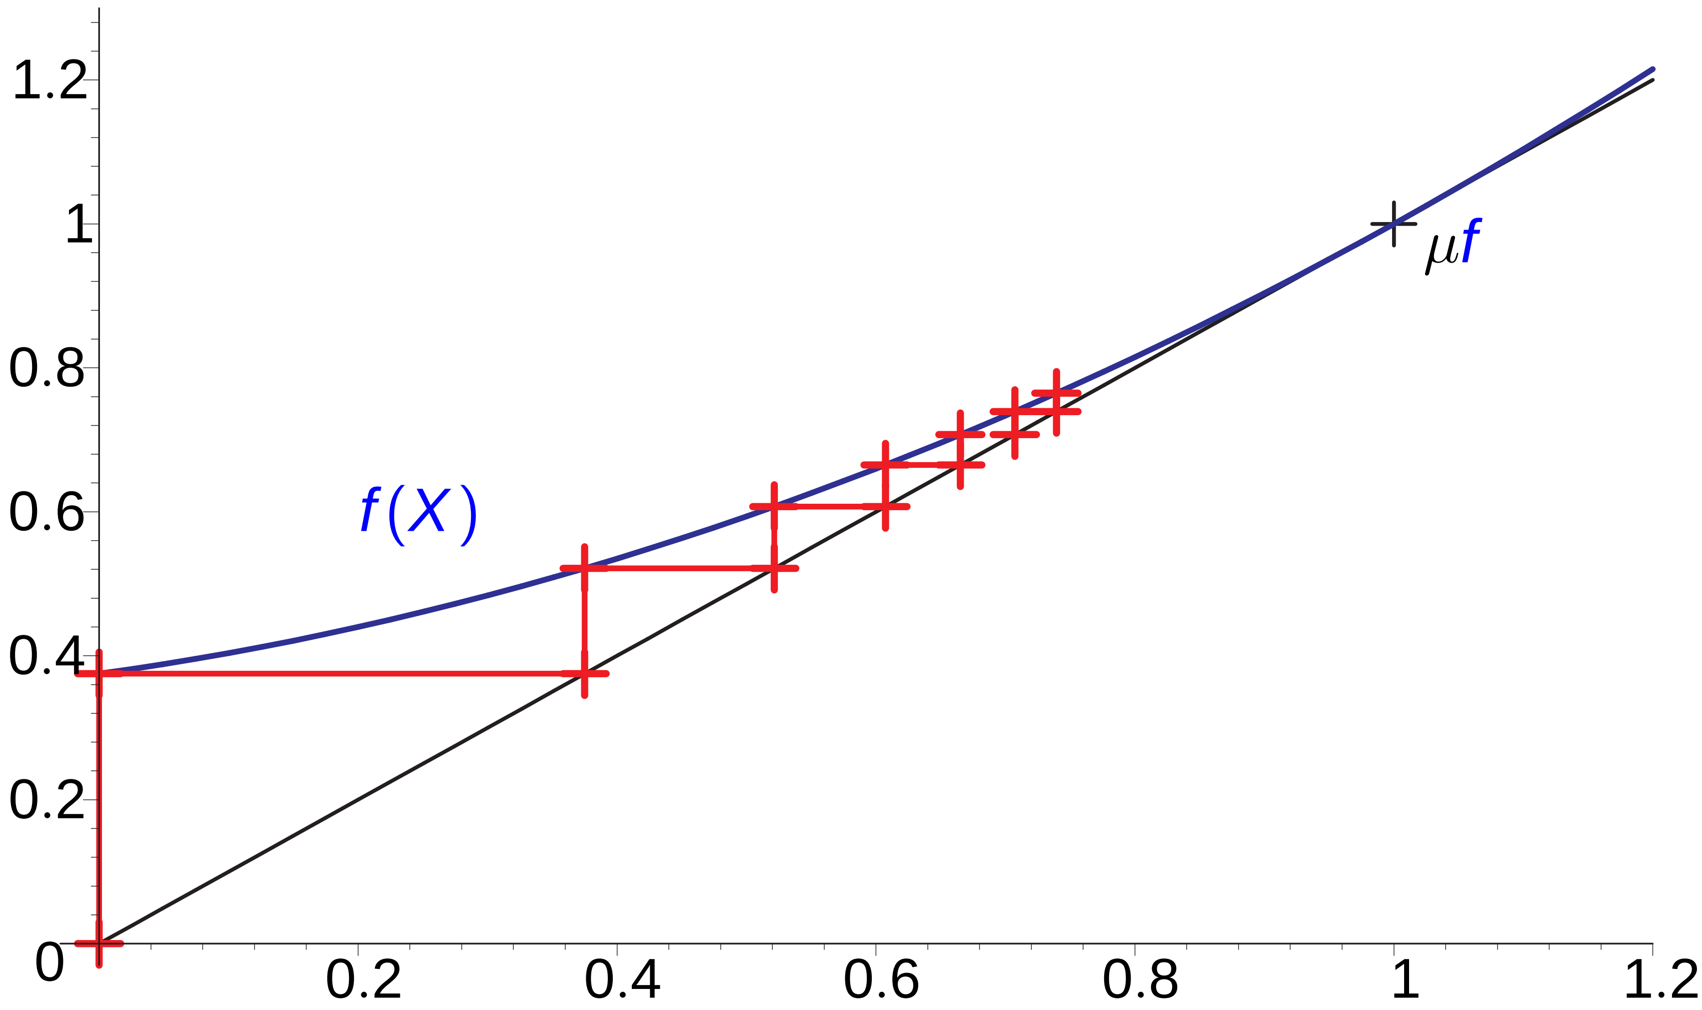

Newton's method

Consider $x = f(x)$ with $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

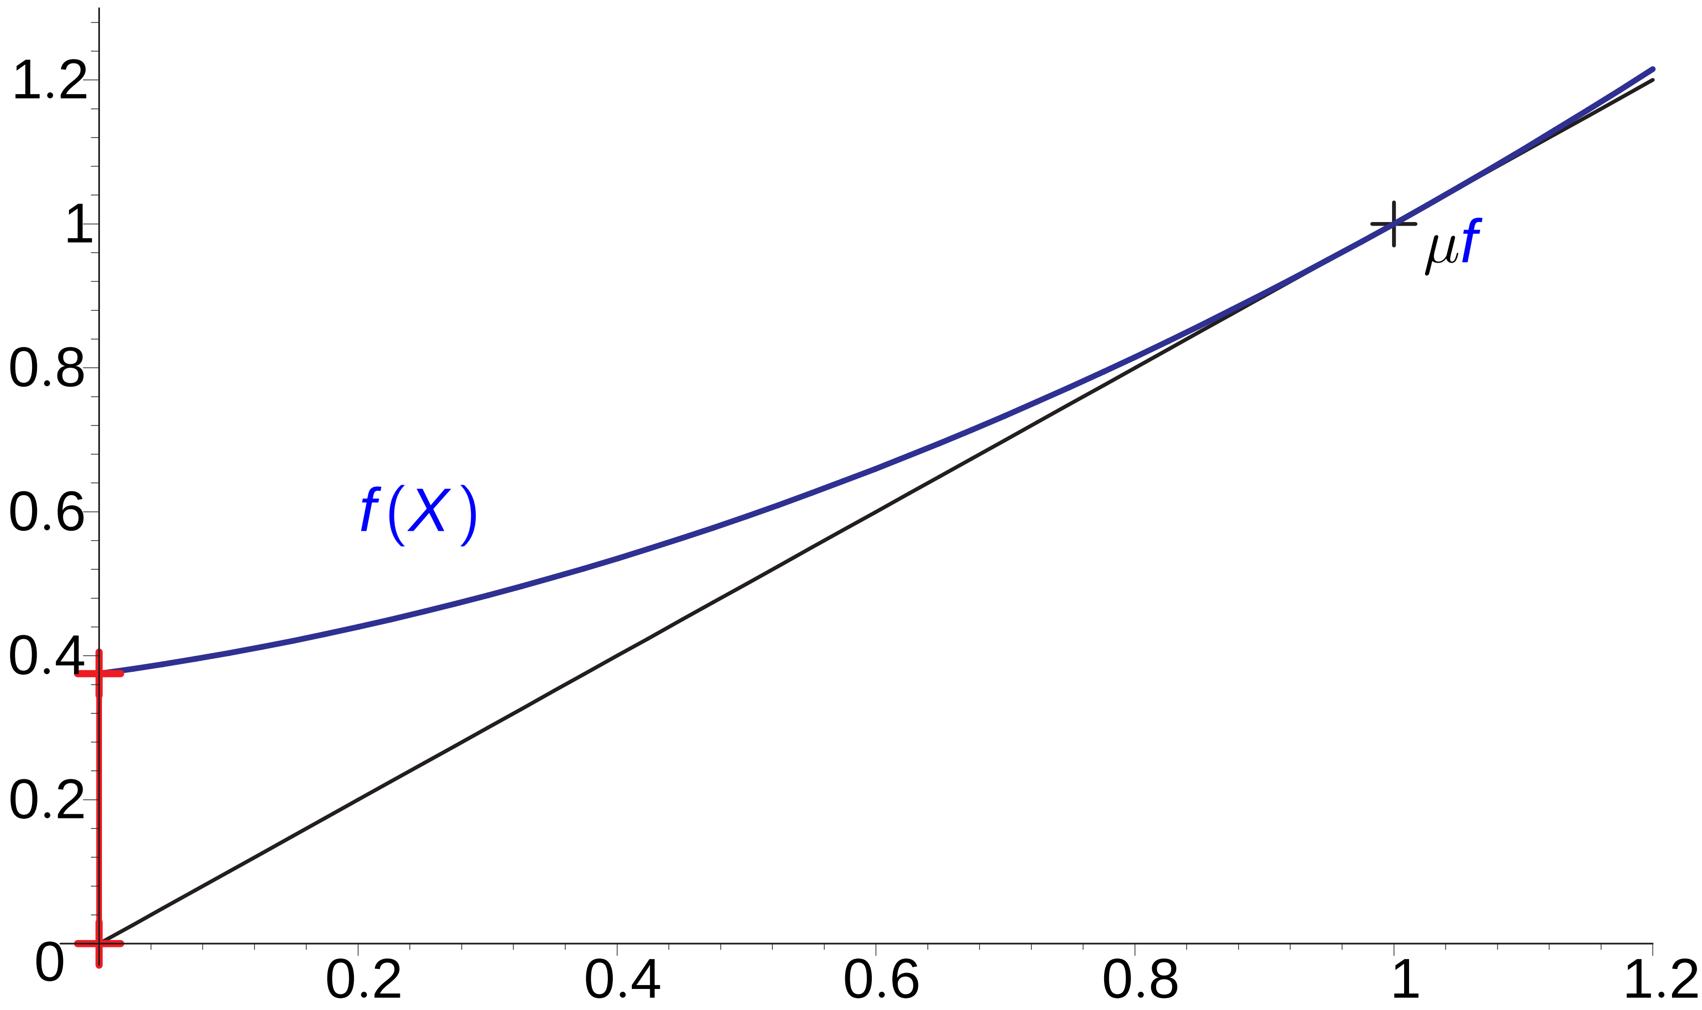

Newton's method

Consider $x = f(x)$ with $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

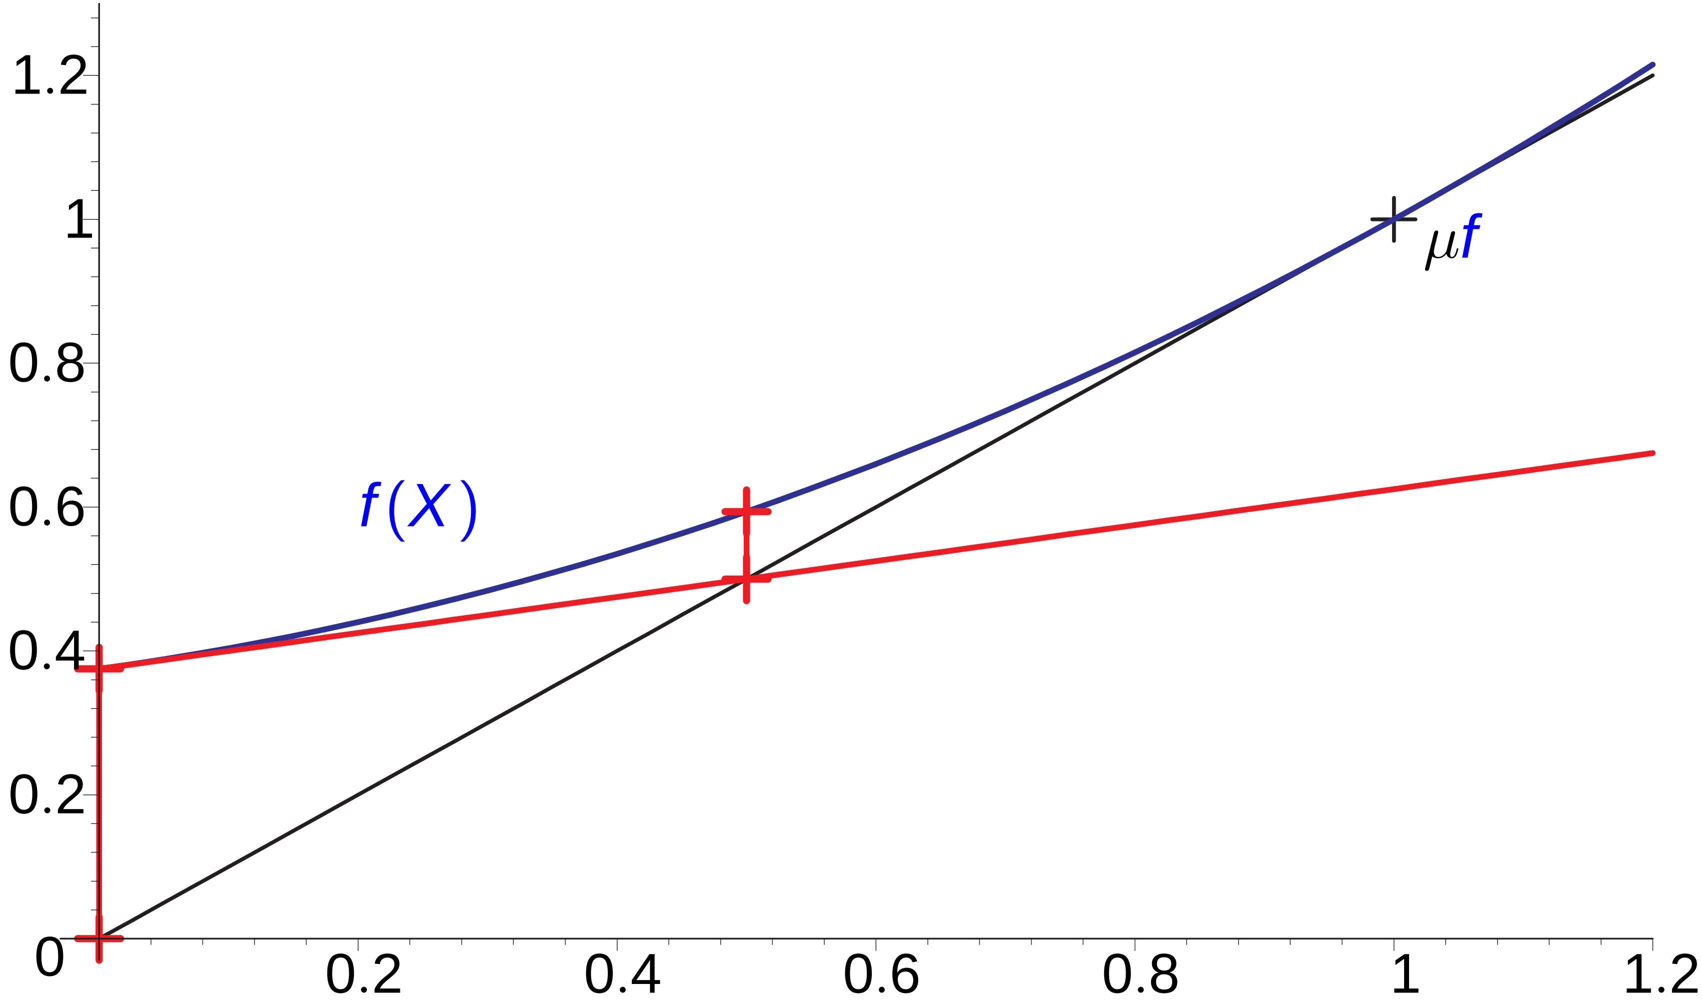

Newton's method

Consider $x = f(x)$ with $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

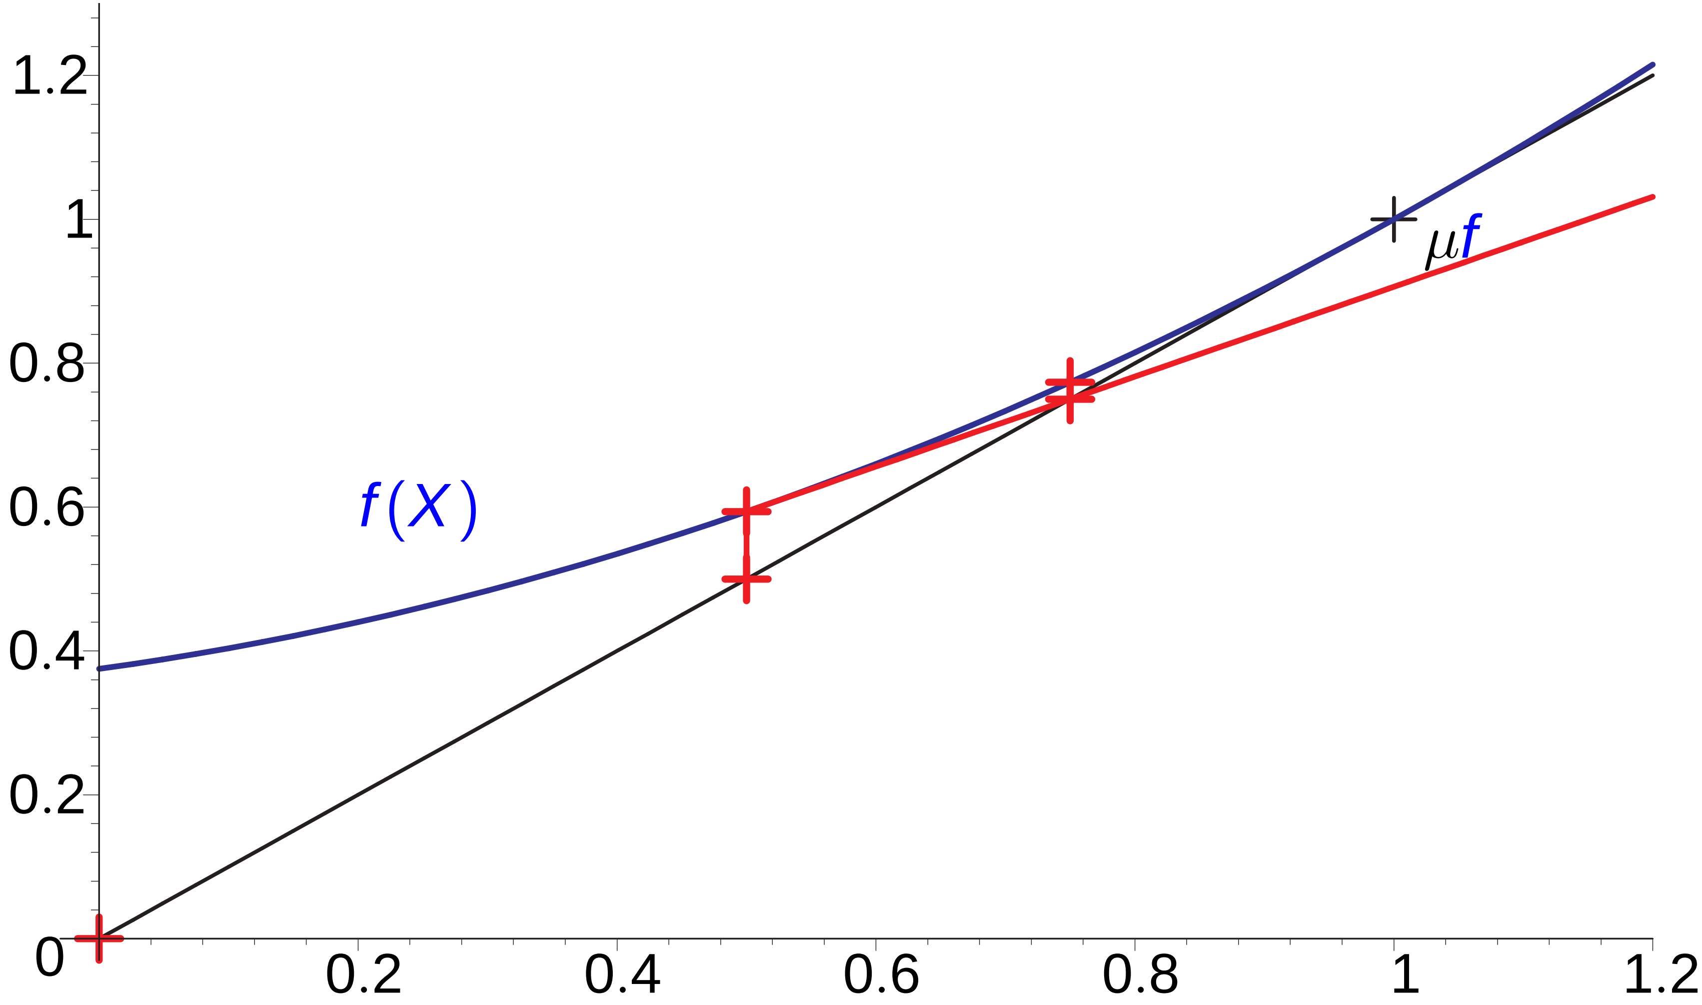

Newton's method

Consider $x = f(x)$ with $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

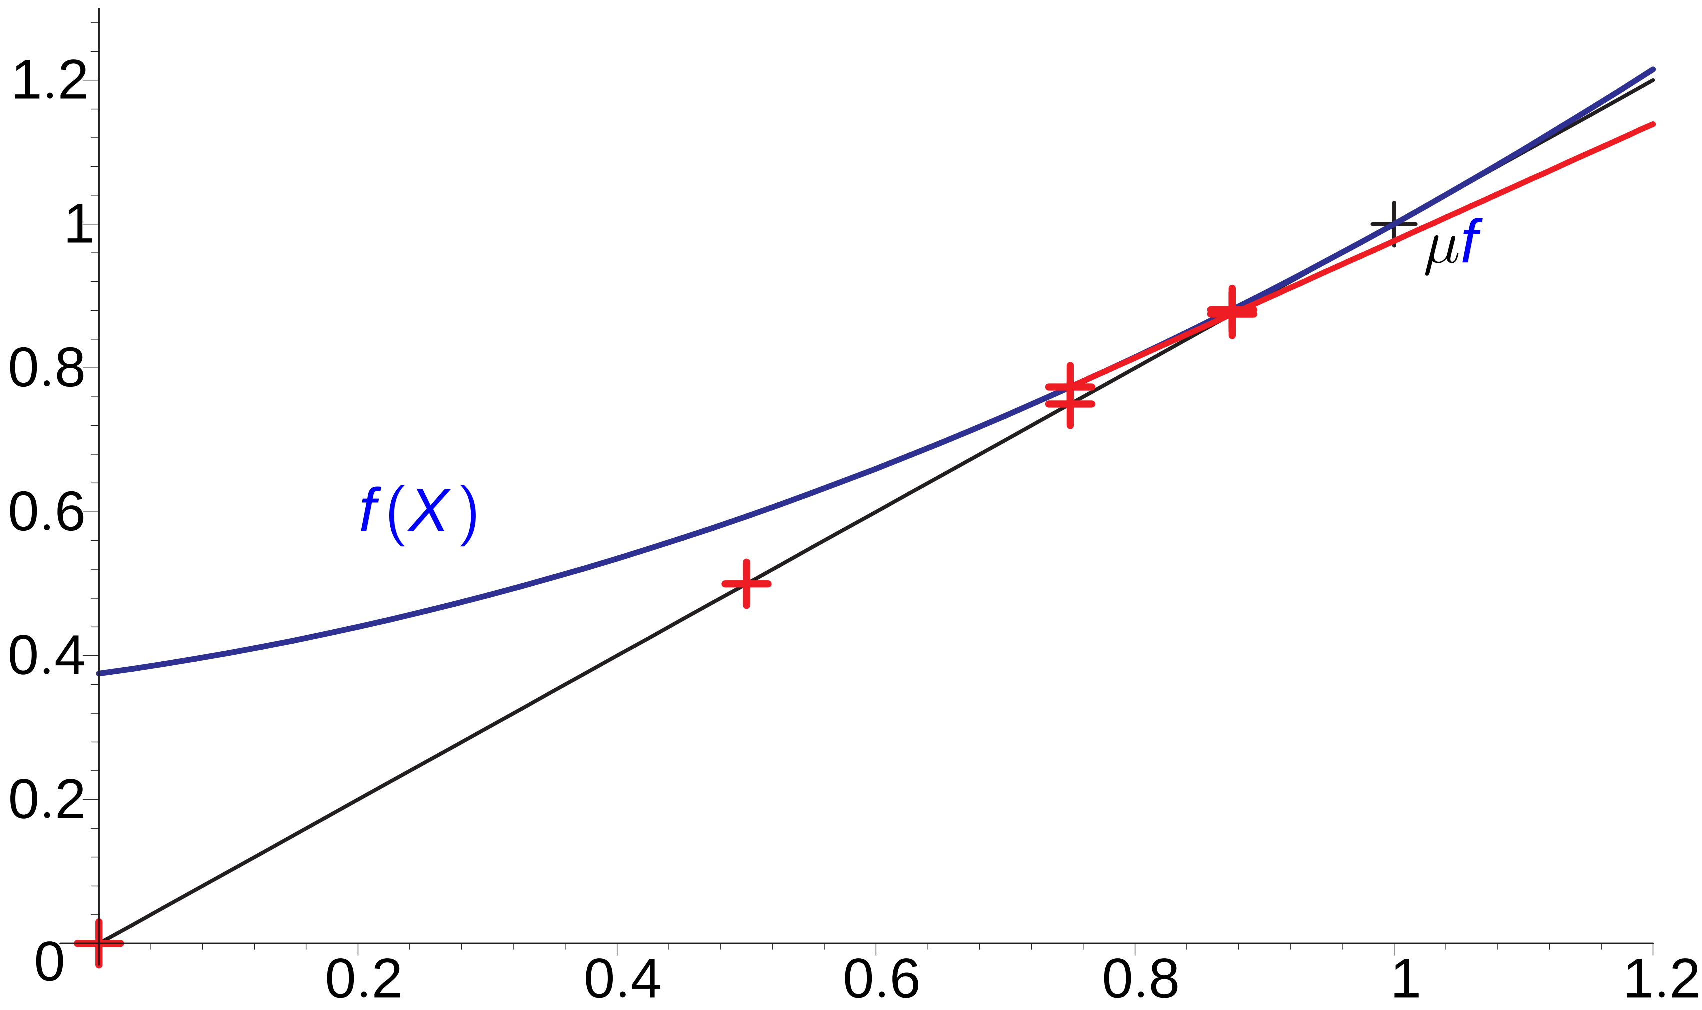

Newton's method

Consider $x = f(x)$ with $f(x) = \frac{3}{8}x^2 + \frac{1}{4}x + \frac{3}{8}$

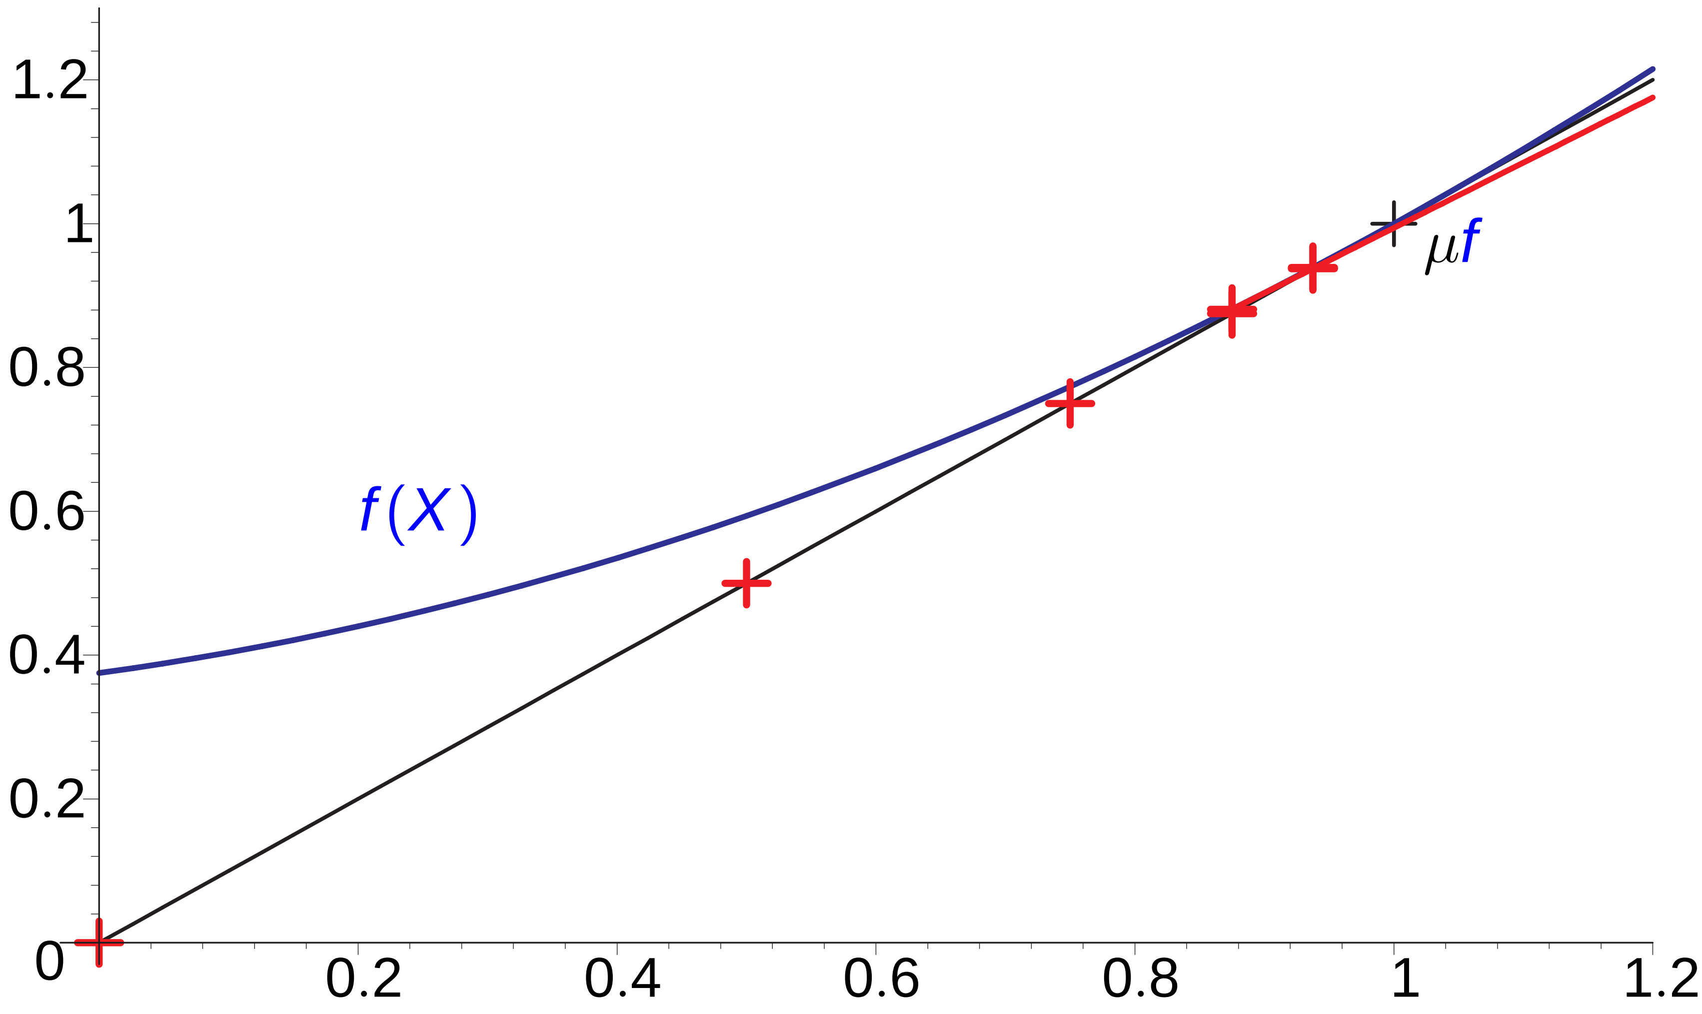

Mathematical formulation of Newton's method

The Newton approximants to the least fixed point of $f(x)$ are given by:

$$\begin{aligned} \nu^{(0)} &= 0 \\ \nu^{(i+1)} &= \nu^{(i)} + \frac{f\left(\nu^{(i)}\right) - \nu^{(i)}}{1 - f'\left(\nu^{(i)}\right)} \end{aligned}$$

Completion space of the optimal scheduler

Proposition

The probability $\Pr(S^{op} \leq k)$ *of completing execution within space at most* k *is equal to the* k*-th Newton approximant* $\nu^{(k)}$ *of the least fixed point of* $f(x)$.

Completion space of the optimal scheduler

Proposition

The probability $\Pr(S^{op} \leq k)$ of completing execution within space at most k is equal to the k-th Newton approximant $\nu^{(k)}$ of the least fixed point of $f(x)$.

Proof idea.

Show that $\{\Pr(S^{op} \leq k)\}_{k \geq 0}$ and $\{\nu^k\}_{k \geq 0}$ satisfy the same recurrence equation. □

Completion space of the optimal scheduler

Proposition

The probability $\Pr(S^{op} \leq k)$ *of completing execution within space at most* k *is equal to the* k*-th Newton approximant* $\nu^{(k)}$ *of the least fixed point of* $f(x)$*.*

Proof idea.

Show that $\{\Pr(S^{op} \leq k)\}_{k \geq 0}$ and $\{\nu^{k}\}_{k \geq 0}$ satisfy the same recurrence equation. □

Least fixed point of $f(x)$ = probability of termination

Completion space of the optimal scheduler

Proposition

The probability $\Pr(S^{op} \leq k)$ *of completing execution within space at most* k *is equal to the* k*-th Newton approximant* $\nu^{(k)}$ *of the least fixed point of* $f(x)$.

Proof idea.

Show that $\{\Pr(S^{op} \leq k)\}_{k \geq 0}$ and $\{\nu^k\}_{k \geq 0}$ satisfy the same recurrence equation. □

Least fixed point of $f(x)$ = probability of termination

k-th Newton approximant to the least fixed point

=

probability of termination within space at most k

Exploiting the result

Applying our recent results on the convergence speed of Newton's method [STOC'07 and STACS'08EKL08]:

Theorem

For a subcritical system there are $c > 0$ and $0 < d < 1$ such that $\Pr[S^{op} \geq k] \leq c \cdot d^{2^k}$ for every $k \in \mathbb{N}$.

Consequence: the optimal scheduler always has finite expected completion space

Theorem

For a critical system there are $c > 0$ and $0 < d < 1$ such that $\Pr[S^{op} \geq k] \leq c \cdot d^{k}$ for every $k \in \mathbb{N}$.

Online schedulers

Theorem

Let $a > 1$ be the greatest fixed point of the pgf of a subcritical system (in a certain normal form). Then

$$\Pr[S^\sigma \geq n] = \frac{a-1}{a^n - 1}$$

for every online scheduler σ and for every $n \geq 1$.

Online schedulers

Theorem

Let $a > 1$ be the greatest fixed point of the pgf of a subcritical system (in a certain normal form). Then

$$\Pr[S^\sigma \geq n] = \frac{a-1}{a^n - 1}$$

for every online scheduler σ and for every $n \geq 1$.

- All online schedulers have the same distribution.
  (No longer true for multitype systems!!)

Online schedulers

Theorem

Let $a > 1$ be the greatest fixed point of the pgf of a subcritical system (in a certain normal form). Then

$$\Pr[S^\sigma \geq n] = \frac{a-1}{a^n - 1}$$

for every online scheduler σ and for every $n \geq 1$.

- All online schedulers have the same distribution. (No longer true for multitype systems!!)
- Gap between online and offline schedulers:
  - $\Pr[S^{op} \geq k] \leq c \cdot d^{2^k}$ for the optimal scheduler.
  - $\Pr[S^\sigma \geq n] = \frac{a-1}{a^n-1}$ for any online scheduler σ.

Online schedulers

Theorem

Let $a > 1$ be the greatest fixed point of the pgf of a subcritical system (in a certain normal form). Then

$$\Pr[S^\sigma \geq n] = \frac{a-1}{a^n - 1}$$

for every online scheduler σ and for every $n \geq 1$.

- All online schedulers have the same distribution. (No longer true for multitype systems!!)
- Gap between online and offline schedulers:
  - $\Pr[S^{op} \geq k] \leq c \cdot d^{2^k}$ for the optimal scheduler.
  - $\Pr[S^\sigma \geq n] = \frac{a-1}{a^n-1}$ for any online scheduler σ.
- The optimal scheduler always has finite expected space, online schedulers may not.

Conclusions

- Stochastic branching processes are important for computer science.

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet
  - No distinction between processes and processors.

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet
  - No distinction between processes and processors.
  - No study of "CS random variables" like space consumption.

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet
  - No distinction between processes and processors.
  - No study of "CS random variables" like space consumption.
- Beautiful theory! Surprising connectiosn between approximants to fixed points and random variables of interest.

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet
  - No distinction between processes and processors.
  - No study of "CS random variables" like space consumption.
- Beautiful theory! Surprising connectiosn between approximants to fixed points and random variables of interest.
- Much to do:

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet
  - No distinction between processes and processors.
  - No study of "CS random variables" like space consumption.
- Beautiful theory! Surprising connectiosn between approximants to fixed points and random variables of interest.
- Much to do: *k*-processors,

Conclusions

- Stochastic branching processes are important for computer science.
- Mathematicians haven't studied SBPs for computer science yet
  - No distinction between processes and processors.
  - No study of "CS random variables" like space consumption.
- Beautiful theory! Surprising connectiosn between approximants to fixed points and random variables of interest.
- Much to do: *k*-processors, non-terminating systems, light-first schedulers . . .

Back to victorian Britain . . .

There was concern amongst the Victorians that aristocratic families were becoming extinct.

Francis Galton (1822-1911), anthropologist and polymath:
Are families of English peers more likely to die out than the families of ordinary men?

> *Let $p_0, p_1, \ldots, p_n$ be the respective probabilities that a man has 0, 1, 2, . . . n sons, let each son have the same probability for sons of his own, and so on. What is the probability that the male line goes extinct?*

Henry William Watson (1827-1903), vicar and mathematician:
The probability is the least solution of

$$X = p_0 + p_1 X + p_2 X^2 + \ldots + p_n X^n$$

English peers again . . .

Due to an algebraic error, Watson concluded wrongly that all families eventually die out.

English peers again . . .

Due to an algebraic error, Watson concluded wrongly that all families eventually die out.

But Galton found a fact, that, with hindsight, provides a possible explanation for the observed data:

English peers again . . .

Due to an algebraic error, Watson concluded wrongly that all families eventually die out.

But Galton found a fact, that, with hindsight, provides a possible explanation for the observed data:

- English peers tended to marry heiresses (daughters without brothers)

English peers again . . .

Due to an algebraic error, Watson concluded wrongly that all families eventually die out.

But Galton found a fact, that, with hindsight, provides a possible explanation for the observed data:

- English peers tended to marry heiresses (daughters without brothers)
- Heiresses come from families with lower fertility rates (lower probabilities $p_1, p_2, p_3, \ldots$).

English peers again . . .

Due to an algebraic error, Watson concluded wrongly that all families eventually die out.

But Galton found a fact, that, with hindsight, provides a possible explanation for the observed data:

- English peers tended to marry heiresses (daughters without brothers)
- Heiresses come from families with lower fertility rates (lower probabilities $p_1, p_2, p_3, \ldots$).
- . . . which increases the probability of the family dying out.