Stochastic Networks with Resource Sharing

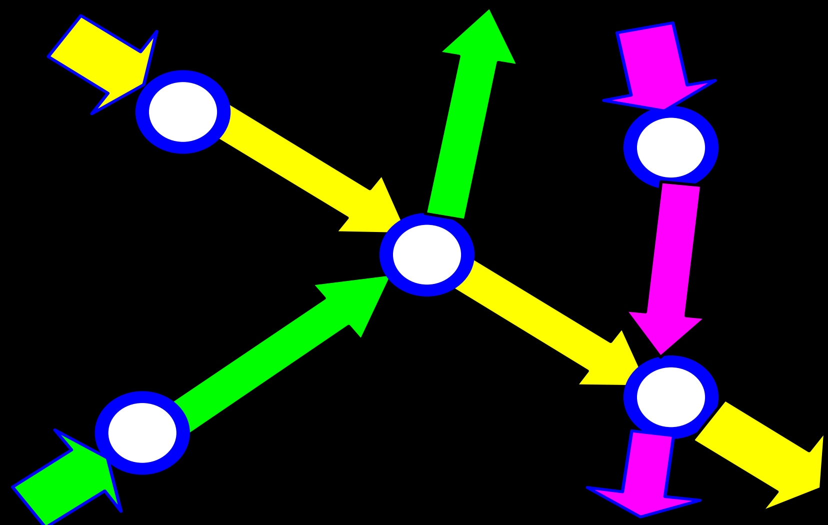

Ruth J. Williams

University of California, San Diego

http://math.ucsd.edu/~williams

Joint work with Weining Kang, Frank Kelly, Nam Lee

2007 MARKOV LECTURE ABSTRACT

Stochastic networks are used as models for complex systems involving dynamic interactions subject to uncertainty. Application domains include manufacturing, the service industry, telecommunications, and computer systems. Networks arising in modern applications are often highly complex and heterogeneous, with network features that transcend those of conventional queueing models. The control and analysis of such networks present challenging mathematical problems. In this talk, a concrete application will be used to illustrate a general approach to the study of stochastic networks using more tractable approximate models. Specifically, we consider a connection-level model of Internet congestion control that represents the randomly varying number of flows present in a network where bandwidth is shared fairly amongst elastic documents. This model, introduced by Massoulie and Roberts, can be viewed as a stochastic network with simultaneous resource possession. Elegant fluid and diffusion approximations will be used to study the stability and performance of this model. The talk will conclude with a summary of the current status and description of open problems associated with the further development of approximate models for general stochastic networks. This talk is based in part on joint work with W. Kang, F. P. Kelly, and N. H. Lee. Discussants: Kavita Ramanan and Mark Squillante.

Outline

- Stochastic processing networks
- Flow level model of congestion control
- Questions: stability and performance
- Approximations: fluid and diffusion
- Perspective and open problems

STOCHASTIC PROCESSING NETWORKS

Stochastic Processing Networks (cf. Harrison '00)

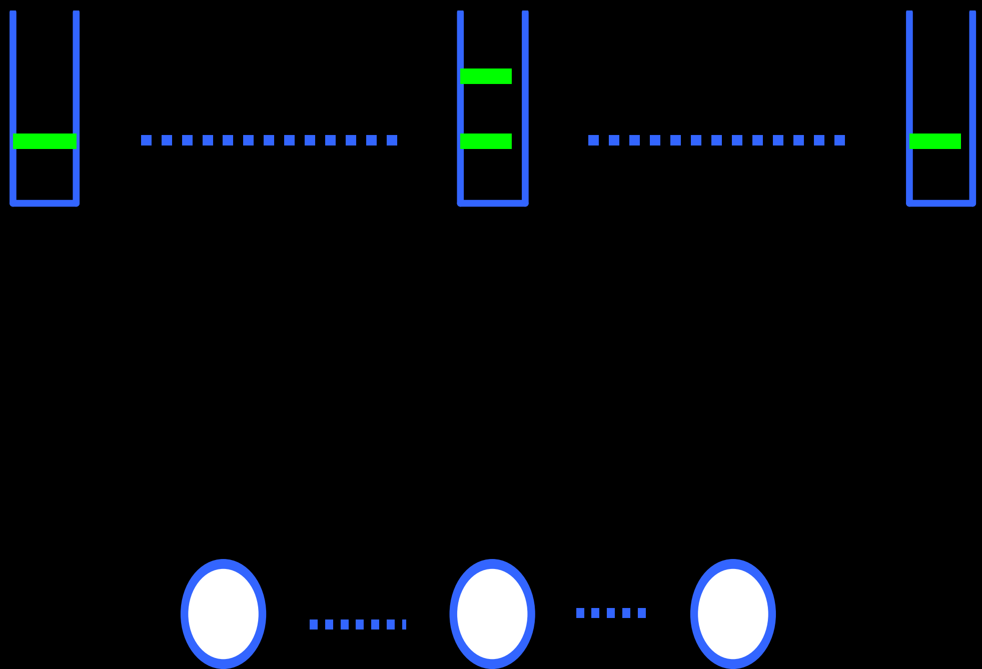

buffers (classes)

activities

servers (resources)

An activity consumes from certain classes,
produces for certain (possibly different) classes,
and uses certain servers.

Stochastic Processing Networks

Activities are Very General

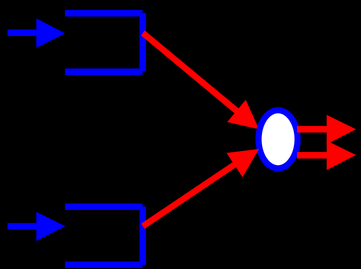

Multiclass Queueing Network

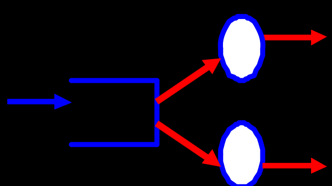

Alternate Routing

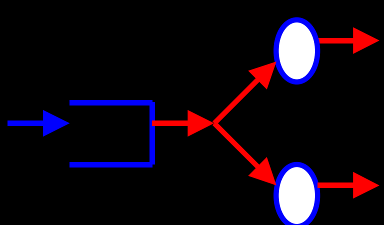

Simultaneous actions

FLOW LEVEL MODEL OF CONGESTION CONTROL

Flow Level Model of Congestion Control

(Massoulie-Roberts ‘00)

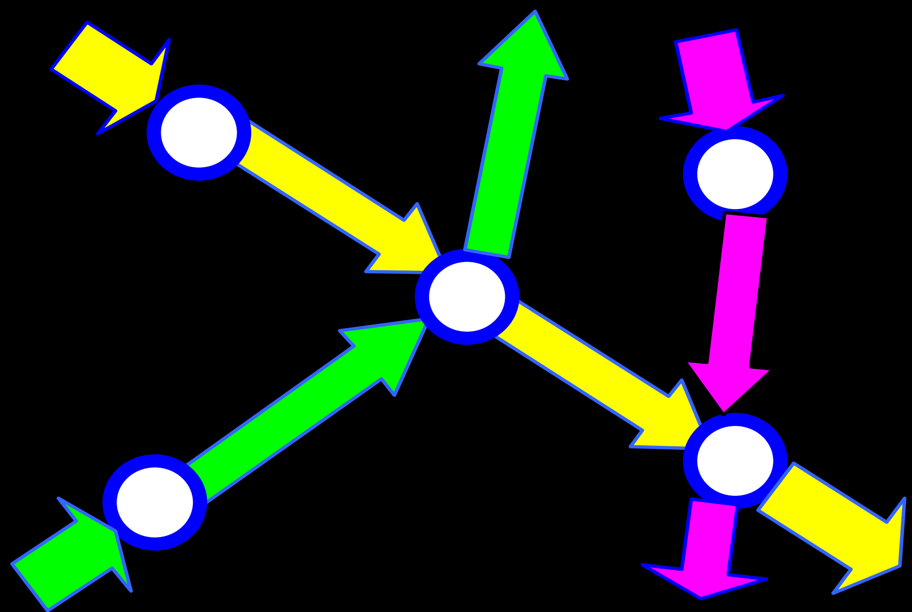

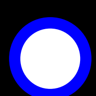 Link

 Route

Flow Level Model

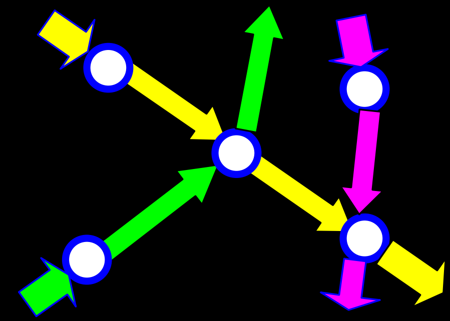

- Documents arrive to routes
- A flow corresponds to the continuous transfer of a document over a route
- Assume a “separation of time scales”

  (zero transmission time through the network)
- Bandwidth is allocated dynamically to the routes and is shared equally amongst all active flows on a route

Stochastic Processing Network with Simultaneous Resource Possession

- *Flow level model*

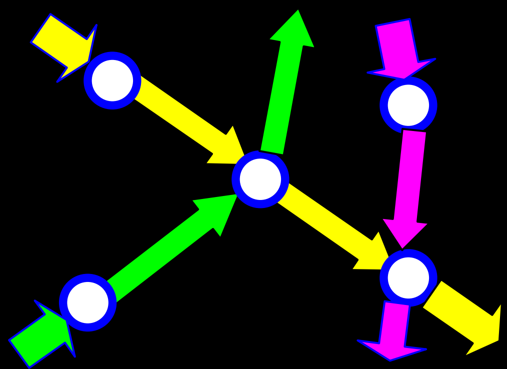

- *Stochastic processing network*

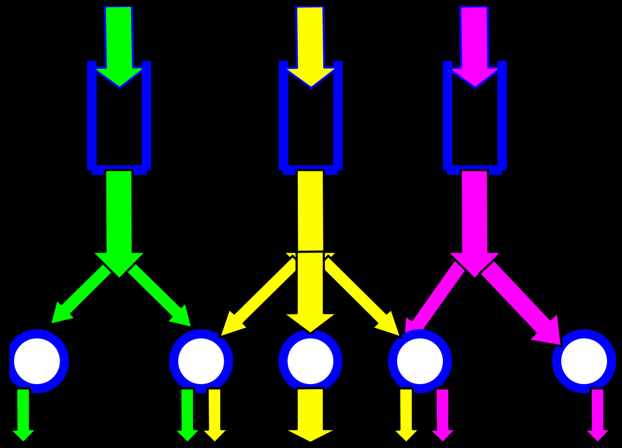

Network Structure for Flow Level Model

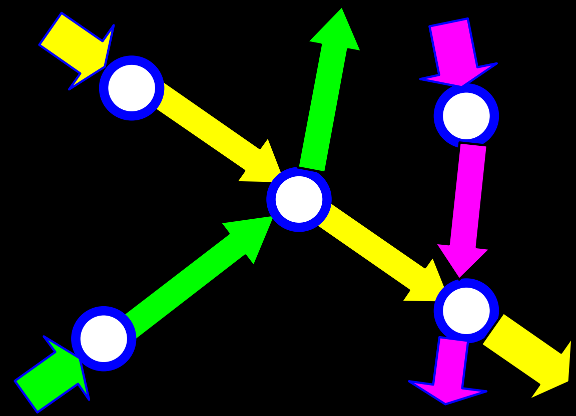

$\mathbb{I}$ routes → $\mathbb{J}$ **links**

Bandwidth (capacity) for link *j:* C_j

Incidence matrix *(full row rank):*

$$A_{ji} = \begin{cases} 1 & \text{if route } i \text{ uses link } j \\ 0 & \text{otherwise} \end{cases}$$

Stochastic Assumptions for Flow Level Model

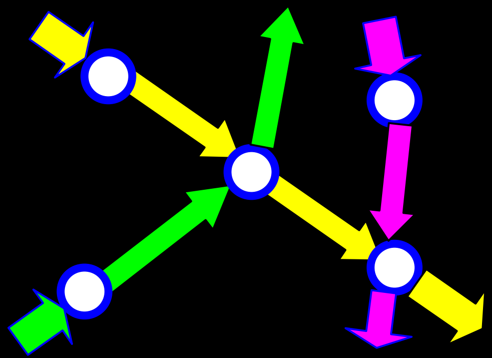

- Poisson arrivals of documents at rate ν_i to route *i*
- Document sizes: exponentially distributed with mean μ_i^{-1} for route *i*
- Interarrival times and document sizes are all mutually independent
- Traffic intensity $\rho_i = \nu_i / \mu_i$

Proportional Fair Bandwidth Sharing Policy (Kelly '97)

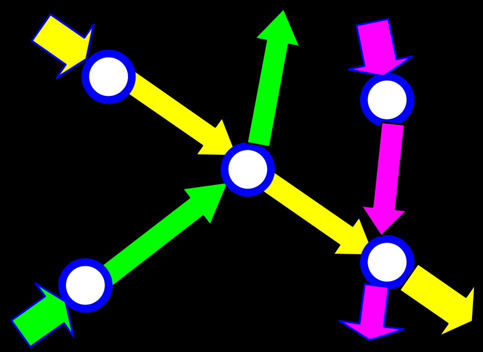

N_i = # of documents on route *i*

$\Lambda_i(N)$ = bandwidth allocation for route *i*

$$\Lambda(N) = \operatorname{argmax}\left\{\sum_i N_i \log(\Lambda_i) : A\Lambda \le C, \Lambda \ge 0\right\}$$

Bandwidth Allocations

$$\Lambda_i(N) = \frac{N_i}{\sum_j A_{ji} q_j(N)}$$

$q_j(N) =$ Lagrange multiplier for the link j capacity constraint

Stochastic Network Model

Number of documents on route i at time t:

$$N_i(t) = N_i(0) + E_i(t) - S_i(T_i(t))$$

Cumulative unused capacity for link j up to time t:

$$U_j(t) = C_j t - (AT(t))_j$$

Cumulative bandwidth allocated to route i up to time t:

$$T_i(t) = \int_0^t \Lambda_i(N(s))ds$$

$E_i(\cdot) =$ Poisson process of rate ν_i

$S_i(\cdot) =$ Poisson process of rate μ_i

Outline of rest of talk

- Stability
- Performance in heavy traffic
  - balanced fluid model and invariant manifold
  - (multiplicative) state space collapse
  - diffusion approximation
  - example: a linear network (entrainment)
- Further problems and perspective

STABILITY

Fluid model

(formal functional law of large numbers limit)

$$\bar{N}_i(t) = \bar{N}_i(0) + \nu_i t - \mu_i \bar{T}_i(t) \geq 0$$

$$\bar{U}_j(t) = C_j t - (A\bar{T}(t))_j, \quad \bar{U}_j \uparrow$$

$\bar{T}$ **is uniformly Lipschitz continuous and** $\bar{T}(0) = 0.$

At a.e. *t*, for each *i*,

$$\frac{d}{dt}\bar{T}_i(t) = \begin{cases} \Lambda_i(\bar{N}(t)) & if \quad \bar{N}_i(t) > 0 \\ \rho_i & if \quad \bar{N}_i(t) = 0 \end{cases}$$

where $\rho_i = \nu_i / \mu_i.$

Stability

Lyapunov function

$$F(N) = \sum_i N_i^2 / \nu_i$$

Theorem (De Veciana et al. '01, Bonald & Massoulie '01, Kelly-W '04)

The Markov chain $N(\cdot)$ is positive recurrent if and only if $A\rho < C$

PERFORMANCE IN HEAVY TRAFFIC

$(A\rho = C)$

Balanced Fluid Model $(A\rho = C)$

Defn: $n \in \mathbb{R}_+^{\mathbb{I}}$ is an invariant state if there is a fluid model solution $\bar{N}(\cdot)$ such that $\bar{N}(t) = n$ for all $t \geq 0$.

Balanced Fluid Model $(A\rho = C)$

Theorem (Kelly-W '04) The following are equivalent:

(a) *n* is an invariant state

(b) $\Lambda_i(n) = \rho_i$ for all $i : n_i > 0$

(c) $\exists\, q \in \mathbb{R}_+^{\mathbb{J}} :\; n_i = \rho_i (A'q)_i$ for all i

(d) $n = \Delta w(n)$ where $w(n) = A\, diag(\mu)^{-1} n$

$$\Delta = diag(\rho) A' [A diag(\mu)^{-1} diag(\rho) A']^{-1}$$

Balanced Fluid Model $(A\rho = C)$

Theorem (Kelly-W '04) The following are equivalent:

(a) *n* is an invariant state

(b) $\Lambda_i(n) = \rho_i$ for all $i : n_i > 0$

(c) $\exists\, q \in \mathbb{R}_+^{\mathbb{J}} :\ n_i = \rho_i (A'q)_i$ for all i

(d) $n = \Delta w(n)$ where $w(n) = A\, diag(\mu)^{-1} n$

$$\Delta = diag(\rho) A' [A diag(\mu)^{-1} diag(\rho) A']^{-1}$$

Furthermore, fluid model solutions converge uniformly to the invariant manifold starting in a compact set.

Fluid and Diffusion Scaling

Fluid scaling

$$\overline{N}^r(t) = N^r(rt)/r$$

Fluid and Diffusion Scaling

Fluid scaling

$$\overline{N}^r(t) = N^r(rt)/r$$

Diffusion scaling

$$\hat{N}^r(t) = N^r(r^2 t)/r$$

Fluid and Diffusion Scaling

Fluid scaling

$$\overline{N}^r(t) = N^r(rt)/r$$

Diffusion scaling

$$\hat{N}^r(t) = N^r(r^2 t)/r$$

$$\hat{W}^r(t) = A\,diag(\mu)^{-1}\hat{N}^r(t)$$

(Multiplicative) State Space Collapse

Theorem (Kang-Kelly-Lee-W) Suppose the system starts empty.

For each $t \geq 0$,

$$\frac{\sup_{0 \leq s \leq t} \left\| \hat{N}^r(s) - \Delta \hat{W}^r(s) \right\|}{\sup_{0 \leq s \leq t} \left(\left\| \hat{N}^r(s) \right\| \vee 1 \right)} \Rightarrow 0$$

as $r \to \infty$.

Proof: Use asymptotic behavior of balanced fluid model and adapt Bramson '98.

Diffusion Scaled Workload

$$\hat{W}^r(t) \equiv A\, diag(\mu)^{-1} \hat{N}^r(t) = \hat{X}^r(t) + \hat{U}^r(t)$$

where

$$\hat{X}^r(t) = A\, diag(\mu)^{-1} \left(\hat{E}^r(t) - \hat{S}^r\left(\bar{\bar{T}}^r(t) \right) \right)$$

$$\approx \textit{Brownian motion}$$

Conjectured Diffusion Approximation in Heavy Traffic $(A\rho = C)$

Conjecture: $\hat{W}^r \Rightarrow \tilde{W}$ as $r \to \infty$ where $\tilde{W} = \tilde{X} + \tilde{U}$ is a semimartingale reflecting Brownian motion in the polyhedral cone

$$\mathcal{W} = \{A\,diag(\mu)^{-1} diag(\rho) A' q : q \in \mathbb{R}_+^{\mathbb{J}}\}$$

Here $\tilde{U}_j$ increases on the boundary corresponding to $q_j = 0$

SRBM IN A 3-D CONE

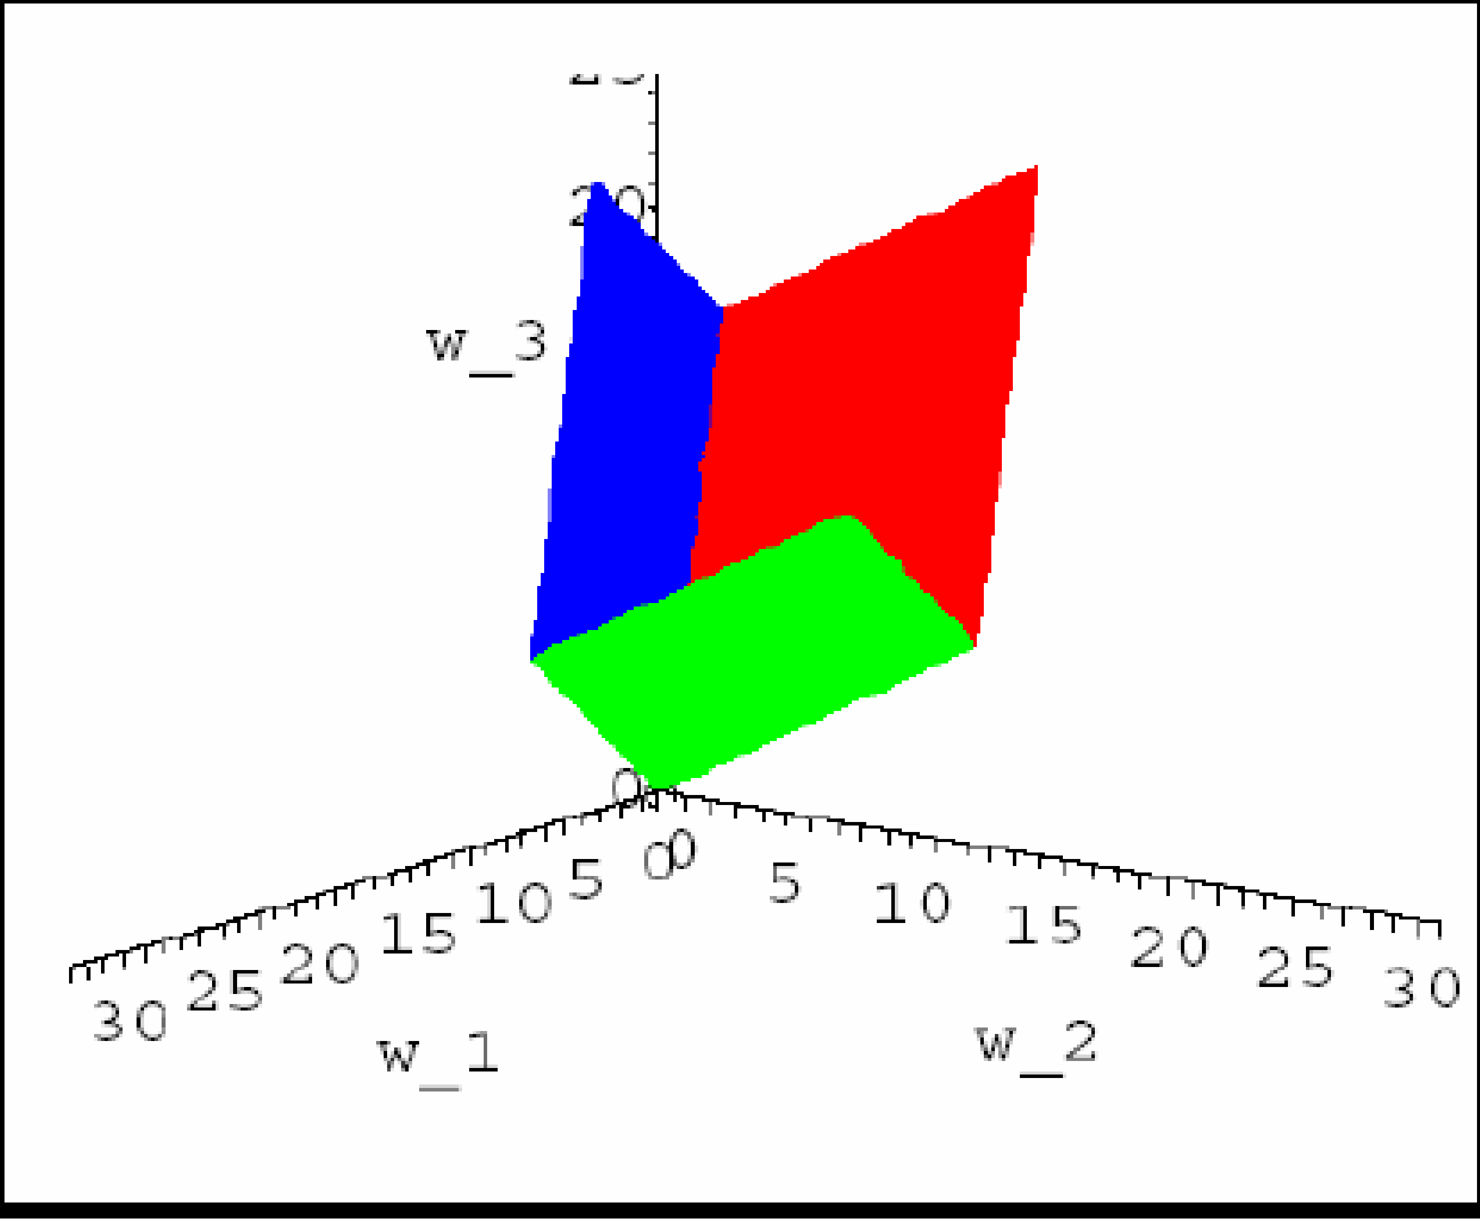

SRBM IN A 3-D CONE

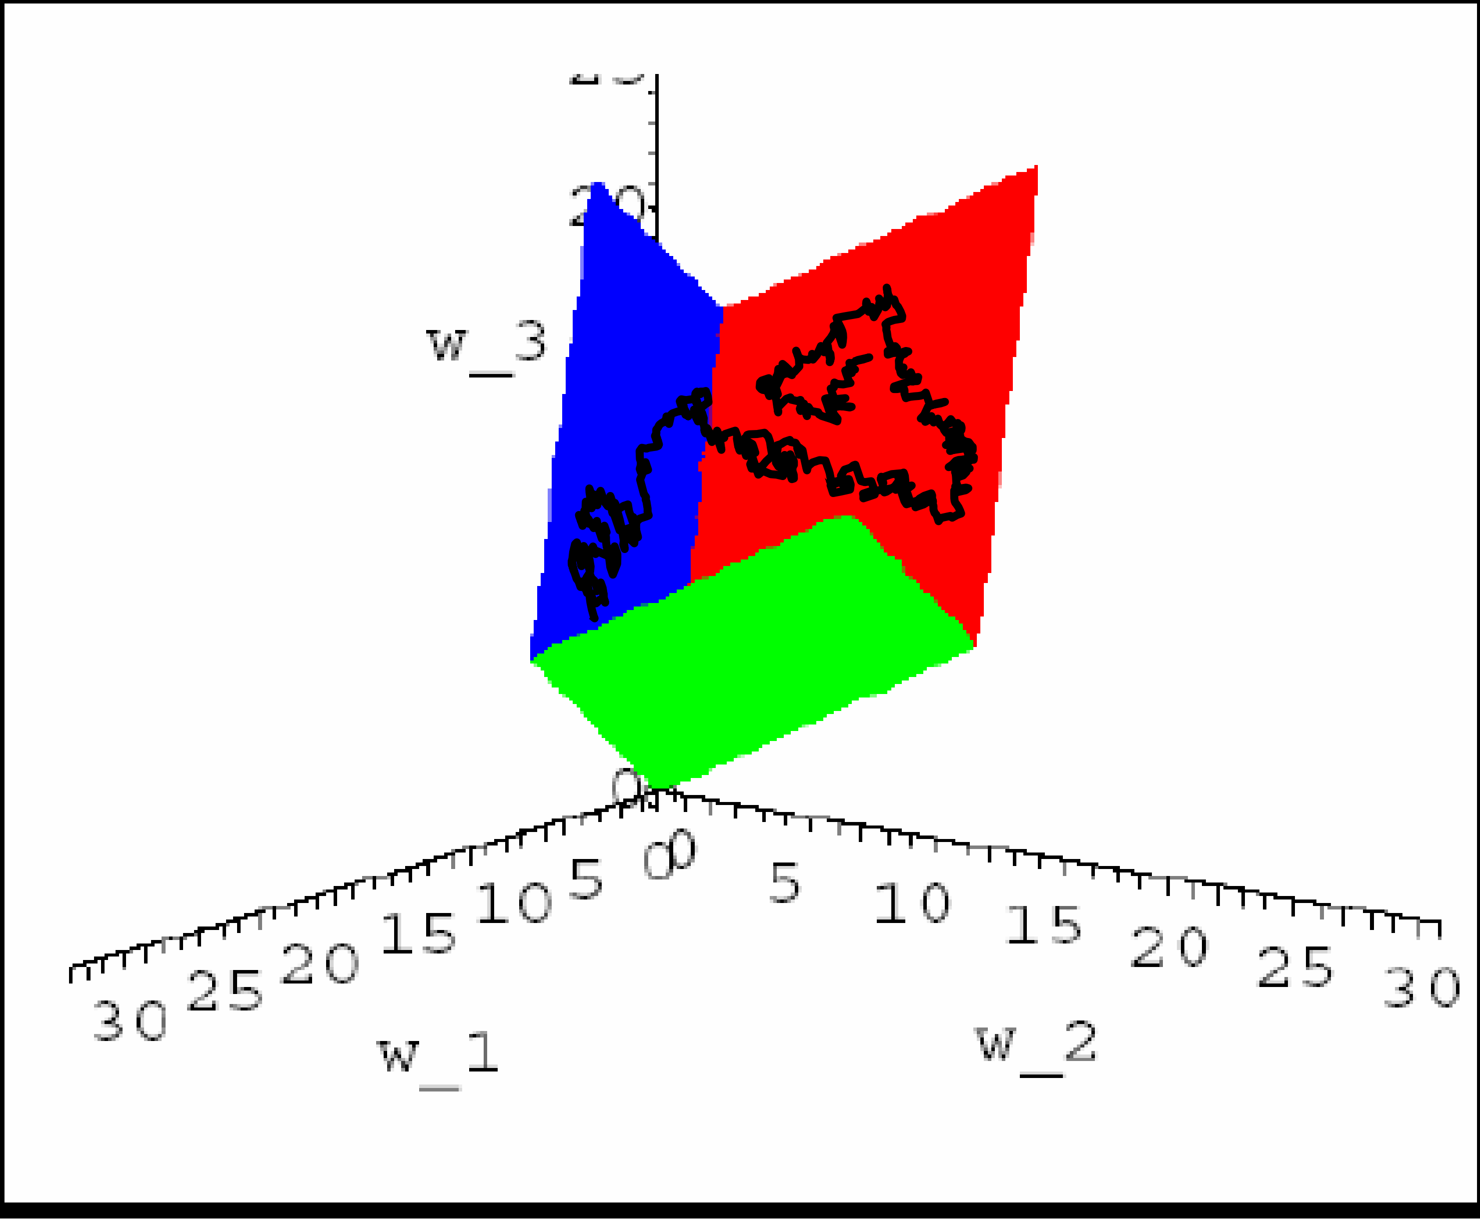

SRBM IN A 3-D CONE

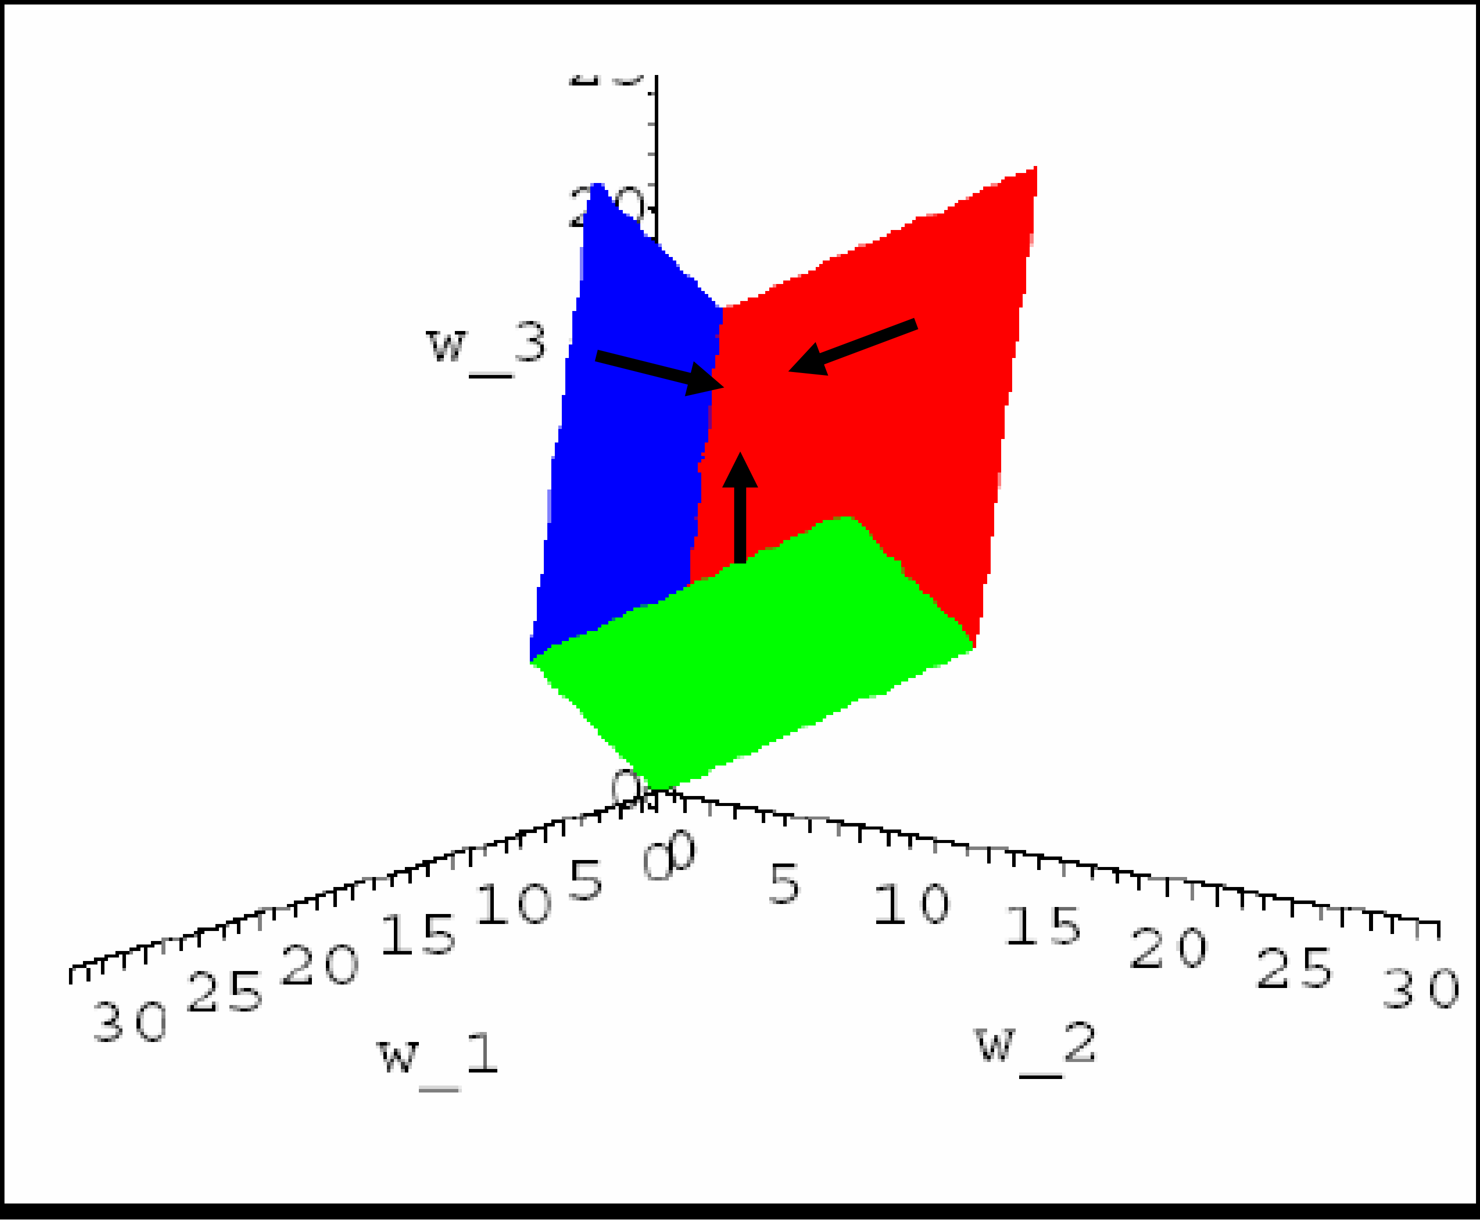

DIFFUSION APPROXIMATION

Diffusion Approximation

Assumption (local traffic): For each link j there is a route i that goes only through j

Diffusion Approximation

Assumption (local traffic): For each link j

there is a route i that goes only through j

Theorem (Kang-Kelly-Lee-W) Suppose that the local traffic assumption holds. Then,

$$(\hat{W}^r, \hat{N}^r) \Rightarrow (\tilde{W}, \tilde{N}) \text{ as } r \to \infty$$

where $\tilde{N} = \Delta \tilde{W}$ and $\tilde{W} = \tilde{X} + \tilde{U}$ is an SRBM in the polyhedral cone $\mathcal{W}$

Diffusion Approximation

Assumption (local traffic): For each link j there is a route i that goes only through j

Theorem (Kang-Kelly-Lee-W) Suppose that the local traffic assumption holds. Then,

$$(\hat{W}^r, \hat{N}^r) \Rightarrow (\tilde{W}, \tilde{N}) \text{ as } r \to \infty$$

where $\tilde{N} = \Delta \tilde{W}$ and $\tilde{W} = \tilde{X} + \tilde{U}$ is an SRBM in the polyhedral cone $\mathcal{W}$

Proof : Uses invariance principle of Kang-W '07

Example: Linear Network

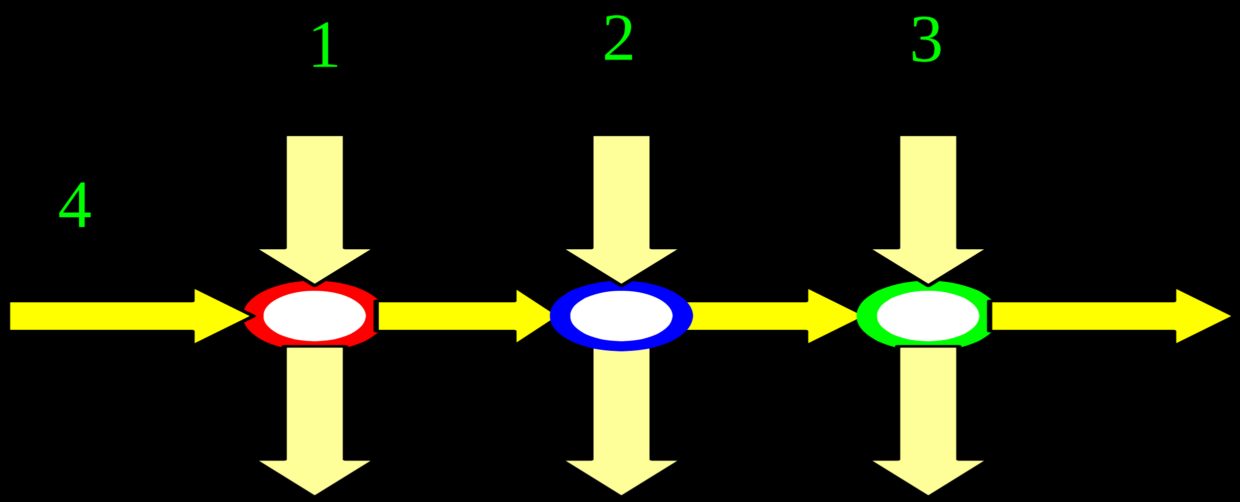

$$\mu_i = 1,\ i = 1, 2, 3, 4$$

$$\rho_1 + \rho_4 = 1,\ \rho_2 + \rho_4 = 1, \rho_3 + \rho_4 = 1,$$

$$C_j = 1,\ j = 1, 2, 3$$

Linear Network: Workload State Space

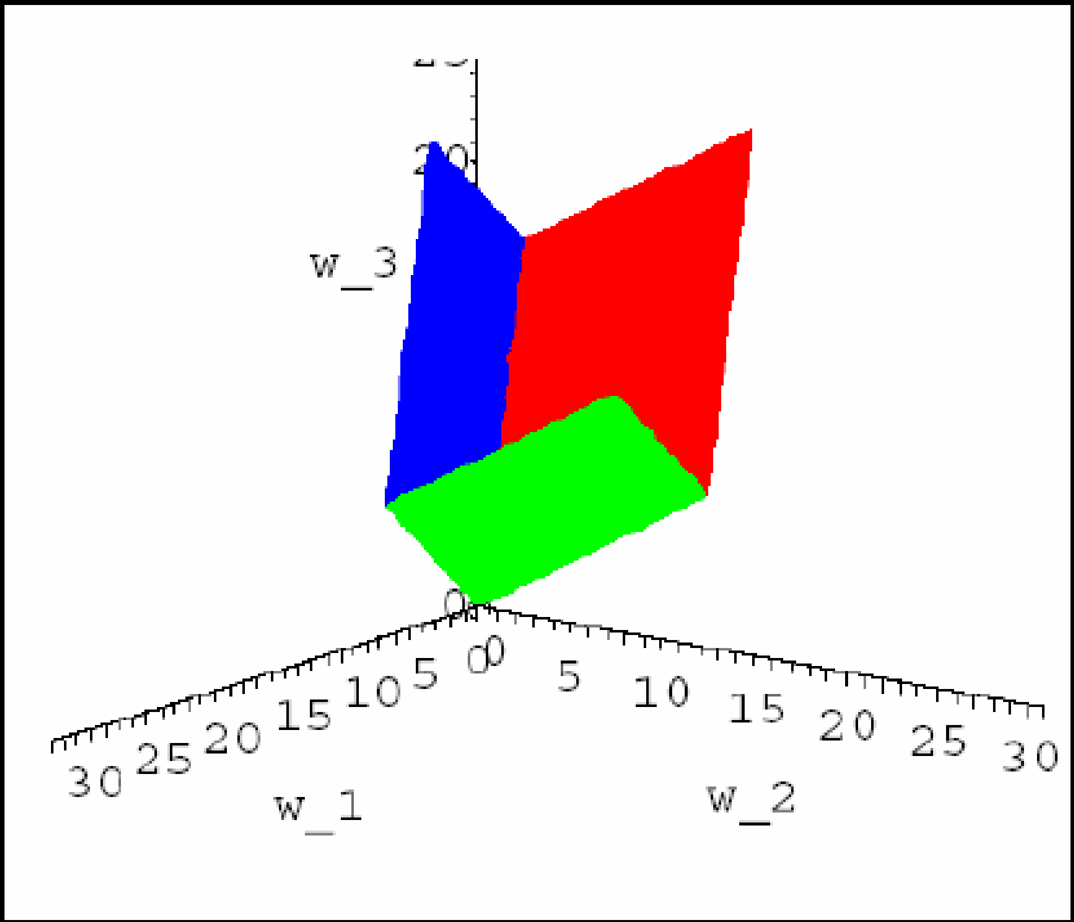

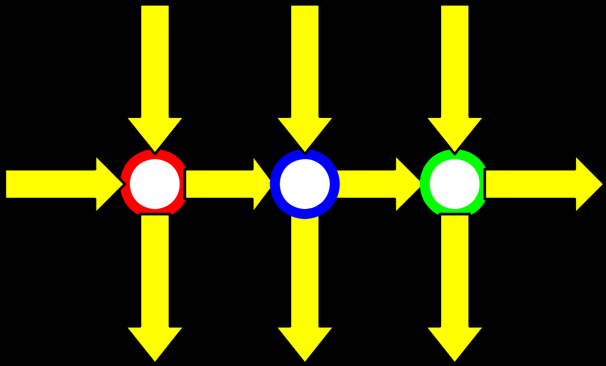

STATIONARY DISTRIBUTION

Product Form Stationary Distribution

Theorem (Kang-Kelly-Lee-W) Suppose that the SRBM $\tilde{W}$ has a drift $\theta < 0$ and covariance matrix Γ. Then, $\tilde{W}$ has a product form stationary distribution with density

$$p(w) = c \exp(2\Gamma^{-1}\theta \cdot w), \; w \in \mathcal{W}$$

Product Form Stationary Distribution

Theorem (Kang-Kelly-Lee-W) Suppose that the SRBM $\tilde{W}$ has a drift $\theta < 0$ and covariance matrix Γ.
Then, $\tilde{W}$ has a product form stationary distribution with density

$$p(w) = c \exp(2\Gamma^{-1}\theta \cdot w), \ w \in \mathcal{W}$$

Hence, $\widetilde{N} = \Delta \tilde{W}$ has a stationary distribution expressed as a linear combination of independent exponential random variables.

Extensions

- Theorems extend to document sizes distributed as finite mixtures of exponentials (insensitivity).
  [cf. Massoulie & Roberts '00, Bonald & Massoulie '01 exact results for single link, linear network, grid network with equal capacities.]

- Some extension to models with multipath routing

FURTHER PROBLEMS

Further Problems for Flow Level Model

- Prove diffusion approximation for more general utility based bandwidth sharing policies

  (e.g., alpha fair policies of Mo and Walrand '00 ---- diffusion workload state space can be a non-polyhedral cone)

Further Problems for Flow Level Model

- General document size distributions (non-HL)

  (Massoulie '07, Gromoll-W '07)

Further Problems for SPN

- Stability and performance via fluid and diffusion approximations for other stochastic processing networks

PERSPECTIVE

	MQN	SPN
HL	Sufficient conditions for stability and diffusion approximations	e.g., parallel server system, packet switch
Non-HL	e.g., LIFO, EDF, Processor Sharing	e.g., congestion control model with general document distributions

THE END